

Can Nondeterminism Help Complementation?*

Yang Cai

MIT CSAIL
The Stata Center, 32-G696
Cambridge, MA 02139 USA
ycai@csail.mit.edu

Ting Zhang

Iowa State University
226 Atanasoff Hall
Ames, IA 50011 USA
tingz@cs.iastate.edu

Complementation and determinization are two fundamental notions in automata theory. The close relationship between the two has been well observed in the literature. In the case of nondeterministic finite automata on finite words (NFA), complementation and determinization have the same state complexity, namely $\Theta(2^n)$ where n is the state size. The same similarity between determinization and complementation was found for Büchi automata, where both operations were shown to have $2^{\Theta(n \lg n)}$ state complexity. An intriguing question is whether there exists a type of ω -automata whose determinization is considerably harder than its complementation. In this paper, we show that for all common types of ω -automata, the determinization problem has the same state complexity as the corresponding complementation problem at the granularity of $2^{\Theta(\cdot)}$. In particular, we show a determinization construction for Streett automata with state complexity bounded by $2^{12n}(0.37n)^{n^2+8n}$.

1 Introduction

Automata on infinite words (ω -automata) have wide applications in synthesis and verification of reactive concurrent systems. Complementation and determinization are two fundamental notions in automata theory. The complementation problem is to construct, given an ω -automaton \mathcal{A} , an ω -automaton \mathcal{B} that recognizes the complementary language of \mathcal{A} . Provided that a given \mathcal{A} is nondeterministic, the determinization problem is to construct a deterministic ω -automaton \mathcal{B} that recognizes the same language as \mathcal{A} does.

The close relationship between determinization and complementation has been well observed in the literature (see [5, 23] for discussions). A deterministic ω -automaton can be trivially complemented by dualizing its acceptance condition. As a result, a lower bound on complementation applies to determinization while an upper bound on determinization applies to complementation. Besides easily rendering complementation, determinization is crucial in decision problems for tree temporal logics, logic games and system synthesis. For example, using game-theoretical semantics [9], complementation of ω -tree automata (ω -tree complementation) not only requires complementation of ω -automata, but also requires the complementary ω -automata be deterministic (called co-determinization). Therefore, lowering the cost of ω -determinization also improves the performance of ω -tree complementation.

In the case of nondeterministic finite automata on finite words (NFA), complementation and determinization have the same state complexity, namely $\Theta(2^n)$ where n is the state size. The same similarity between determinization and complementation was found for Büchi automata, where both operations were shown to have $2^{\Theta(n \lg n)}$ state complexity. An intriguing question is whether there exists a type of ω -automata whose determinization is considerably harder than its complementation. In this paper, we show that for all common types of ω -automata, the determinization problem has the same state complexity as the corresponding complementation problem at the granularity of $2^{\Theta(\cdot)}$.

*The research was supported by NSF grants 0954132 and 1143830.

Related Work. Büchi started the theory of ω -automata. The original ω -automata used to establish the decidability of S1S [1] are now referred to as Büchi automata. Shortly after Büchi’s work, McNaughton proved a fundamental theorem in the theory of ω -automata, that is, ω -regular languages (the languages that are recognized nondeterministic Büchi automata) are exactly those recognizable by the deterministic version of a type of ω -automata, now referred to as Rabin automata¹ [12].

The complexity of McNaughton’s Büchi determinization is double exponential. In 1988, Safra proposed a type of tree structures (now referred to as Safra trees) to obtain a Büchi determinization that converts a nondeterministic Büchi automaton of state size n to an equivalent deterministic Rabin automaton of state size $2^{O(n \lg n)}$ and index size $O(n)$ [17]. Safra’s construction is essentially optimal as the lower bound on state size for Büchi complementation is $2^{\Omega(n \lg n)}$ [13, 11]. Later, Safra generalized the Büchi determinization to a Streett determinization which, given a nondeterministic Streett automaton of state size n and index size k , produces an equivalent deterministic Rabin automaton of state size $2^{O(nk \lg nk)}$ and index size $O(nk)$ [19]. A variant Büchi determinization using a similar tree structure was proposed by Muller and Schupp [14]. Recently, Piterman improved both of Safra’s determinization procedures with a node renaming scheme [15]. Piterman’s constructions are more efficient than Safra’s, though the asymptotical bounds in terms of $2^{O(\cdot)}$ are the same. A big advantage of Piterman’s constructions, however, is to output deterministic parity automata, which is easier to manipulate than deterministic Rabin automata. For example, there exists no efficient procedure to complement deterministic Rabin automata to Büchi automata [18], while such complementation is straightforward and efficient for deterministic parity automata.

In [4, 2, 3] we established tight bounds for the complementation of ω -automata with rich acceptance conditions, namely Rabin, Streett and parity automata. The state complexities of the corresponding determinization problems, however, are yet to be settled. In particular, a large gap exists between the lower and upper bounds for Streett determinization. A Streett automaton can be viewed as a Büchi automaton $\langle Q, \Sigma, Q_0, \Delta, \mathcal{F} \rangle$ except that the acceptance condition $\mathcal{F} = \langle G, B \rangle_I$, where $I = [1..k]$ for some k and $G, B : I \rightarrow 2^Q$, comprises k pairs of *enabling sets* $G(i)$ and *fulfilling sets* $B(i)$. A run is accepting if for every $i \in I$, if the run visits $G(i)$ infinitely often, then so does it to $B(i)$. Therefore, Streett automata naturally express *strong fairness* conditions that characterize meaningful computations [8, 7]. Moreover, Streett automata can be exponentially more succinct than Büchi automata in encoding infinite behaviors of systems [18]. As a result, Streett automata have an advantage in modeling the behaviors of concurrent and reactive systems. For Streett determinization, the gap between current lower and upper bounds is huge: the lower bound is $2^{\Omega(n^2 \lg n)}$ [2] and the upper bound is $2^{O(nk \lg nk)}$ [19, 15] when k is large (say $k = \omega(n)$).

In this paper we focus on Streett determinization. We show a Streett determinization whose state complexity matches the lower bounds established in [2]. More precisely, our construction has state complexity $2^{O(n \lg n + nk \lg k)}$ for $k = O(n)$ and $2^{O(n^2 \lg n)}$ for $k = \omega(n)$. We note that this improvement is not only meant for large k . When $k = O(\log n)$, the difference between $2^{O(n \lg n + nk \lg k)}$ and $2^{O(nk \lg nk)}$ is already substantial. In fact, no matter how large k is, our state complexity is always bounded by $2^{12n} (0.37n)^{n^2 + 8n}$ while the current best bound for $k = n - 1$ is $(12)^{n^2} n^{3n^2 + n}$ [15].

The phenomenon that determinization and complementation have the same state complexity does not stop at Streett automata; we also show that this phenomenon holds for generalized Büchi automata, parity automata and Rabin automata. This raises a very interesting question: do determinization and complementation always “walk hand in hand”? Although the exact complexities of complementation and determinization for some or all types of ω -automata could be different, the “coincidence” at the

¹[12] used Muller condition which, however, was converted from Rabin condition.

granularity of $2^{\Theta(\cdot)}$ is already intriguing².

Our Approaches. Our improved construction for Streett determinization bases on two ideas. The first one is what we have exploited in obtaining tight upper bounds for Streett complementation [3], namely, the larger the size of Streett index size, the higher correlations in the runs of Streett automata. We used two tree structures: *ITS* (*Increasing Tree of Sets*) and *TOP* (*Tree of Ordered Partitions*) to characterize those correlations. We observed that there is a similarity between *TOP* and Safra trees for Büchi determinization [17]. As Safra trees for Streett determinization are generalization of those for Büchi determinization, we conjectured that *ITS* should have a role in improving Streett determinization. Our study confirmed this expectation; using *ITS* we can significantly reduce the size of Safra trees for Streett determinization.

The second idea is a new naming scheme. Bounding the size of Safra trees alone cannot bring down the state complexity when the Streett index is small (i.e., $k = O(n)$), because the naming cost becomes a dominating factor in this case. Naming is an integral part of Safra trees. Every node in a Safra tree is associated with a name, which is used to track changes of the node between the tree (state) and its successors. The current name allocation is a retail-style strategy; when a new node is created, a name from the pool of unused names is selected arbitrarily and assigned, and when a node is removed, its name is recycled to the pool. In contrast, our naming scheme is more like wholesale; the name space is divided into even blocks and every block is allocated at the same time. When a branch is created, an unused block is assigned to it, and when a branch is changed, the corresponding block is recycled.

Paper Organization. Section 2 presents notations and basic terminology in automata theory. Section 3 introduces Safra’s construction for Büchi and Streett determinization. Section 4 presents our construction for Streett determinization. Section 5 establishes tight upper bounds for the determinization of Streett, generalized Büchi, parity, and Rabin automata. Section 6 concludes with some discussion on future work. Due to space limit, most of the proofs are not included, but they can be found in the full version of this paper on authors’ websites.

2 Preliminaries

Basic Notations. Let \mathbb{N} denote the set of natural numbers. We write $[i..j]$ for $\{k \in \mathbb{N} \mid i \leq k \leq j\}$, $[i..j]$ for $[i..j-1]$, and $[n]$ for $[0..n]$. For an infinite sequence ρ , we use $\rho(i)$ to denote the i -th component for $i \in \mathbb{N}$. For a finite sequence α , we use $|\alpha|$ to denote the length of α , $\alpha[i]$ ($i \in [1..|\alpha|]$) to denote the object at the i -th position, and $\alpha[i..j]$ (resp. $\alpha[1..j]$) to denote the subsequence of α from position i to position j (resp. $j-1$). We reserve n and k as parameters of a determinization instance (n for state size and k for index size). Define $I = [1..k]$.

Automata and Runs. An ω -automaton is a tuple $\mathcal{A} = (\Sigma, Q, Q_0, \Delta, \mathcal{F})$ where Σ is an alphabet, Q is a finite set of states, $Q_0 \subseteq Q$ is a set of initial states, $\Delta \subseteq Q \times \Sigma \times Q$ is a set of transition relations, and \mathcal{F}

²Recent work by Colcombet and Zdanowski, and by Schewe showed that the state complexity of determinization of Büchi automata on alphabets of *unbounded* size is between $\Omega((1.64n)^n)$ [6] and $O((1.65n)^n)$ [21], which is strictly higher than the state complexity of Büchi complementation which is between $\Omega(L(n))$ [24] and $O(n^2L(n))$ [20] (where $L(n) \approx (0.76n)^n$). However, Schewe’s determinization construction produces Rabin automata with exponentially large index size, and it is not known whether Colcombet and Zdanowski’s lower bound result can be generalized to Büchi automata on conventional alphabets of fixed size.

is an acceptance condition. A *finite run* of \mathcal{A} from state q to state q' over a finite word w is a sequence of states $\rho = \rho(0)\rho(1)\cdots\rho(|w|)$ such that $\rho(0) = q, \rho(|w|) = q'$ and $\langle \rho(i), w(i), \rho(i+1) \rangle \in \Delta$ for all $i \in [|w|]$. An infinite word (ω -words) over Σ is an infinite sequence of letters in Σ . A *run* ρ of \mathcal{A} over an ω -word w is an infinite sequence of states in Q such that $\rho(0) \in Q_0$ and $\langle \rho(i), w(i), \rho(i+1) \rangle \in \Delta$ for $i \in \mathbb{N}$. Let $\text{Inf}(\rho)$ be the set of states that occur infinitely many times in ρ . An automaton accepts w if ρ satisfies \mathcal{F} , which usually is defined as a predicate on $\text{Inf}(\rho)$. By $\mathcal{L}(\mathcal{A})$ we denote the set of ω -words accepted by \mathcal{A} .

Acceptance Conditions and Types. ω -automata are classified according their acceptance conditions. Below we list automata of common types. Let G, B be two functions $I \rightarrow 2^Q$.

- *Generalized Büchi*: $\langle B \rangle_I: \forall i \in I, \text{Inf}(\rho) \cap B(i) \neq \emptyset$.
- *Büchi*: $\langle B \rangle_I$ with $I = \{1\}$ (i.e., $k = 1$).
- *Streett*, $\langle G, B \rangle_I: \forall i \in I, \text{Inf}(\rho) \cap G(i) \neq \emptyset \rightarrow \text{Inf}(\rho) \cap B(i) \neq \emptyset$.
- *Parity*, $\langle G, B \rangle_I$ with $B(1) \subset G(1) \subset \cdots \subset B(k) \subset G(k)$.
- *Rabin*, $[G, B]_I: \exists i \in I, \text{Inf}(\rho) \cap G(i) \neq \emptyset \wedge \text{Inf}(\rho) \cap B(i) = \emptyset$.

For simplicity, we denote a Büchi condition by F (i.e., $F = B(1)$) and call it the final set of \mathcal{A} . Note that Büchi, generalized Büchi and parity automata are all subclasses of Streett automata. For a Streett condition $\langle G, B \rangle_I$, if there exist $i, i' \in I, B(i) = B(i')$, then we can simplify the condition by replacing both $\langle G(i), B(i) \rangle$ and $\langle G(i'), B(i') \rangle$ by $\langle G(i) \cup G(i'), B(i) \rangle$. For this reason, for every Streett condition $\langle G, B \rangle_I$, we assume that B is injective and hence $k = |I| \leq 2^n$.

Trees. A tree is a set $V \subseteq \mathbb{N}^*$ such that if $v \cdot i \in V$ then $v \in V$ and $v \cdot j \in V$ for $j \leq i$. In this paper we only consider finite trees. Elements in V are called *nodes* and ε is called the *root*. Nodes $v \cdot i$ are *children* of v and they are *siblings* to one another. The set of v 's children is denoted by $ch(v)$. A node is a leaf if it has no children. Given an alphabet Σ , a Σ -*labeled tree* is a pair $\langle V, L \rangle$, where V is a tree and $L: V \rightarrow \Sigma$ assigns a letter to each node in V . We refer to $v \in V$ as V -*value* (*structural value*) and $L(v)$ as L -*value* (*label value*) of v .

3 Safra's Streett Determinization

In this section we introduce Safra's constructions for Streett determinization.

We start with the idea behind Büchi determinization [17]. Let $\mathcal{A} = \langle Q, Q_0, \Sigma, \Delta, F \rangle$ be a nondeterministic Büchi automaton. By the standard subset construction [16], it is not hard to build a deterministic ω -automaton \mathcal{B} such that if a run $\rho = q_0, q_1, \dots \in Q^\omega$ of \mathcal{A} over an infinite word w is accepting, then so is the run $\tilde{\rho} = \tilde{q}_0, \tilde{q}_1, \dots \in (2^Q)^\omega$ of \mathcal{B} over w , and $q_i \in \tilde{q}_i$ for every $i \geq 0$. The hard part is to also guarantee that an accepting run $\tilde{\rho}$ of \mathcal{B} over w induces an accepting run ρ of \mathcal{A} over w . Let $S \xrightarrow{F} S'$ denote that for every state $q' \in S'$, there exists a state $q \in S$ and a finite run ρ of \mathcal{A} such that ρ goes from q to q' and visits F . The key idea in many determinization constructions [12, 17, 19, 14, 15] relies on the following lemma, which itself is a consequence of König's lemma.

Lemma 1 ([12, 17]). *Let $\tilde{\rho} = \tilde{q}_0, \tilde{q}_1, \dots \in (2^Q)^\omega$ and $(S_i)_\omega = S_0, S_1, \dots$ an infinite subsequence of $\tilde{\rho}$ such that for every $i \geq 0, S_i \xrightarrow{F} S_{i+1}$, then there is an accepting run $\rho = q_0, q_1, \dots \in Q^\omega$ of \mathcal{A} such that $q_i \in \tilde{q}_i$ for every $i \geq 0$.*

The idea was generalized to Streett determinization [19]. Let $\mathcal{A} = \langle Q, Q_0, \Sigma, \Delta, \langle G, B \rangle_I \rangle$ be a Streett automaton. An index set $J \subseteq I$ serves as a *witness set* for an accepting run ρ of \mathcal{A} in the following sense: for every j in J , ρ visits $B(j)$ infinitely often while for every $j \in I \setminus J$, ρ visits $G(j)$ only finitely many times. We call indices in $I \setminus J$ *negative obligations* and indices in J *positive obligations*. It is easily seen that a run ρ is accepting if and only if ρ admits a witness set $J \subseteq I$. We say that a finite run ρ *fulfills* J if for every j in J , ρ visits $B(j)$, but for every $j \in I \setminus J$, ρ does not visit $G(j)$. By $S \xrightarrow{J} S'$ we mean that every state in S' is reachable from a state in S via a finite run ρ that fulfills J . We have the following lemma analogous to Lemma 1.

Lemma 2 ([19]). *Let $\tilde{\rho} = \tilde{q}_0, \tilde{q}_1, \dots \in (2^Q)^\omega$ and $(S_i)_\omega = S_0, S_1, \dots$ an infinite subsequence of $\tilde{\rho}$ such that for every $i \geq 0$, $S_i \xrightarrow{J} S_{i+1}$, then there is an accepting run $\rho = q_0, q_1, \dots \in Q^\omega$ of \mathcal{A} such that $q_i \in \tilde{q}_i$ for every $i \geq 0$.*

The ingenuity of Safra construction is to efficiently organize information in a type of tree structures, now referred to as Safra trees. In a Safra tree for Streett determinization, each node is labeled by a witness set (index label) and a state set (state label). We simply say that a node *contains* the states in its state label. The standard subset construction is carried out on all nodes in parallel. Each node v with witness set J tracks runs that fulfill J . We say that v turns *green* when J is fulfilled, and we have $S \xrightarrow{J} S'$ where S and S' are the state labels of v at two consecutive moments of turning green. If v turns green infinitely often, then we have a desired $(S_i)_\omega$ as stated in Lemma 2. We illustrate this idea by an example. Let $k = 3$ and $I = [1..3]$.

The initial state in the deterministic automaton \mathcal{B} is a root v labeled with I and Q_0 . The obligation of v is to detect runs that fulfill I . Once a run visits $B(3)$ (fulfilling a positive obligation), the run moves to a new child, waiting to visit $B(2)$, and once the run visits $B(2)$, the run moves to another new child, waiting to visit $B(1)$, and so on. Technically speaking, states are moved around, which induces moving of runs.

This *sequential sweeping* can be stalled because some runs in v , from some point on, could never visit $B(3)$. So v should spawn a child v_3 with the witness set $I \setminus \{3\} = \{1, 2\}$, to detect runs that fulfill $\{1, 2\}$ (see in Figure 1a). A run in v_3 should never visit $G(3)$, its negative obligation. Or the run will be reset (the exact meaning of reset is shown in Step (1.3.2.3.2)). If a run fulfills I , then the run moves into a special child that also has I as its witness set. If all children of v are special, then we say v turns *green*. It is not hard to verify that if S and S' are the state sets of v at any two consecutive moments when v turns green, then we have $S \xrightarrow{J} S'$. But v may never turn green because not all runs fulfill I . Thus the special children of v also have I as their witness sets, for they also spawn children, behaving just as v . This spawning should be recursively applied until we arrive at leaves whose witness sets are singletons. We refer the reader to [19, 22, 15] for a detailed exposition of this spawning and sweeping process. Figure 1 shows part of a Safra tree for Streett determinization in two equivalent representations (with respect to witness sets). The representation shown in Figure 1b is used in Definition 1.

A key ingredient in Safra's construction is to assign names and colors to nodes, in order to track changes on nodes to identify those that turn green infinitely many times. In this paper, use three colors: *green, red* and *yellow*.

Definition 1 (Safra Trees for Streett Determinization (STS)). *A Safra tree for Streett determinization (STS) is a labeled tree $\langle V, L \rangle$ with $L = \langle L_n, L_s, L_c, L_h \rangle$ where*

1.1 $L_n : V \rightarrow [1..nk]$ *assigns each node a unique name.*

1.2 $L_s : V \rightarrow 2^Q$ *assigns each node a subset of Q such that for every node v , $L_s(v) = \cup_{v' \in ch(v)} L_s(v')$ and $L_s(v') \cap L_s(v'') = \emptyset$ for every two distinct $v', v'' \in ch(v)$.*

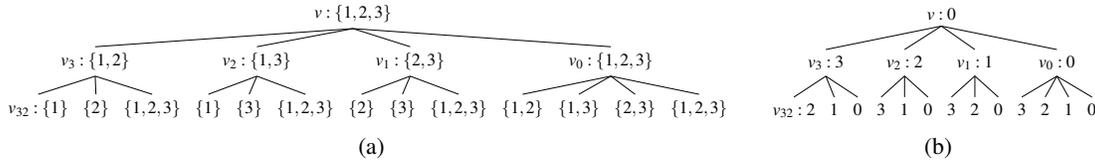


Figure 1: Two equivalent representations of witness sets. In (a), nodes are explicitly labeled with witness sets. In (b), witness sets are implicitly represented by index labels. In this representation, the witness set of a node is the set of indices that do not appear on the path from the root to the node.

1.3 $L_c : V \rightarrow \{\text{green}, \text{red}, \text{yellow}\}$ assigns each node a color.

1.4 $L_h : V \rightarrow I \cup \{0\}$ assigns each node an index in $I \cup \{0\}$. For a node v , let $L_h^{\rightarrow}(v)$ denote the sequence of I -elements from the root to v with 0 excluded and $L_h^{\text{set}}(v)$ the set of I -elements occurring in $L_h^{\rightarrow}(v)$. We require that for every node v , there is no repeated index in the sequence $L_h^{\rightarrow}(v)$.

More precisely, an STS is a labeled and ordered tree; nodes are partially ordered by *older-than* relation. In the tree transformation (see Procedure 1), a newly added node is considered *younger* than its existing siblings. We choose not to define the ordering formally as the meaning is clear from the context. Also, for the sake of presentation clarity, we separate the following naming and coloring convention from the core construction (Step 1.3 of Procedure 1).

Rule 1 (Naming and Coloring Convention). (1) newly created nodes are marked red, (2) nodes whose all descendants are removed are marked green, (3) nodes are marked yellow unless they have been marked red or green, (4) nodes with the empty state label are removed, and (5) when a node is created, a name from the pool of unused name is selected and assigned to the node; when a node is removed, its name is recycled back to the pool.

Procedure 1 (Streett Determinization [19]). Let $\mathcal{A} = \langle Q, Q_0, \Sigma, \Delta, \langle G, B \rangle_I \rangle$ ($I = [1..k]$) be a nondeterministic Streett automaton. The following procedure outputs a deterministic Rabin automaton $\mathcal{B} = \langle \tilde{Q}, \tilde{q}_0, \Sigma, \tilde{\Delta}, [\tilde{G}, \tilde{B}]_{\tilde{I}} \rangle$ ($\tilde{I} = [1..nk]$) such that

1.1 \tilde{Q} is the set of STS.

1.2 $\tilde{q}_0 \in \tilde{Q}$ is the tree with just the root v such that $L_n(v) = 1$, $L_s(v) = Q_0$, $L_c(v) = \text{red}$ and $L_h(v) = 0$.

1.3 $\tilde{q}' = \tilde{\Delta}(\tilde{q})$ is defined the STS obtained by applying the following transformation rule to \tilde{q} .

1.3.1 Subset Construction: for each node v in \tilde{Q} , update state label $L_s(v)$ to $\Delta(L_s(v))$.

1.3.2 Expansion. Apply the following transformations downwards from the root:

1.3.2.1 If v is a leaf with $L_h^{\text{set}}(v) = I$, then stop.

1.3.2.2 If v is a leaf with $L_h^{\text{set}}(v) \neq I$, then add a child v' to v such that $L_s(v') = L_s(v)$ and $L_h(v') = \max(I \setminus L_h^{\text{set}}(v))$.

1.3.2.3 If v is a node with j children v_1, \dots, v_j , then let i_1, \dots, i_j be the corresponding index labels, and consider the following two cases for each $j' \in [1..j]$.

1.3.2.3.1 If $L_s(v_{j'}) \cap B(i_{j'}) \neq \emptyset$, then add a child v' to v with $L_s(v') = L_s(v_{j'}) \cap B(i_{j'})$ and $L_h(v') = \max([0..i_{j'}] \cap ((I \cup \{0\}) \setminus L_h^{\text{set}}(v)))$, and remove the states in $L_s(v_{j'}) \cap B(i_{j'})$ from $v_{j'}$ and all the descendants of $v_{j'}$.

1.3.2.3.2 If $L_s(v_{j'}) \cap B(i_{j'}) = \emptyset$ and $L_s(v_{j'}) \cap G(i_{j'}) \neq \emptyset$, then add a child v' to v with $L_s(v') = L_s(v_{j'}) \cap G(i_{j'})$ and $L_h(v') = L_h(v)$, and remove the states in $L_s(v_{j'}) \cap G(i_{j'})$ from $v_{j'}$ and all the descendants of $v_{j'}$.

1.3.3 Horizontal Merge. For any state q and any two siblings v and v' such that $q \in L_s(v) \cap L_s(v')$, if $L_h(v) < L_h(v')$, or $L_h(v) = L_h(v')$ and v is older than v' , then remove q from v' and all its descendants.

1.3.4 Vertical Merge. For each v , if all children of v have index label 0, then remove all descendants of v .

1.4 $[\tilde{G}, \tilde{B}]_I$ is such that for every $i \in \tilde{I}$

$$\tilde{G}(i) = \{\tilde{q} \in \tilde{Q} \mid \tilde{q} \text{ contains a red node with name } i \text{ or does not contains a node with name } i\}$$

$$\tilde{B}(i) = \{\tilde{q} \in \tilde{Q} \mid \tilde{q} \text{ contains a green node with name } i\}$$

Step (1.3.2.3.1) says that if a run in v_j fulfills the positive obligation i_j by visiting $B(i_j)$ ($L_s(v_j) \cap B(i_j) \neq \emptyset$), then the run moves into a new node v' ($L_s(v') = L_s(v_j) \cap B(i_j)$), and v' continues to monitor if the run hits the next largest positive obligation in the witness set of its parent ($L_h(v') = \max([0..i_j] \cap ((I \cup \{0\}) \setminus L_h^{set}(v)))$). Step (1.3.2.3.2) says that if this is not the case and the run in v_j also violates the negative obligation i_j by visiting $G(i_j)$ ($L_s(v_j) \cap G(i_j) \neq \emptyset$), then this run is *reset*, in the sense that the states in $L_s(v_j) \cap G(i_j)$ moved into a new child of v .

4 Improved Streett Determinization

In this section, we show our improved construction for Streett determinization. Define $\mu = \min(n, k)$.

4.1 Improvement I

The first idea is what we have applied to Streett complementation [3], namely, the larger the k , the more overlaps between $B(i)$'s and between $G(i)$'s ($i \in I$). Let us revisit the previous example, illustrated in Figure 1. Assume that $G(2) \subseteq G(3)$ (we say that $G(3)$ covers $G(2)$). If a run stays at v_3 , then the run is not supposed to visit $G(3)$, or otherwise the run should have been reset by Step (1.3.2.3.2). Since the run cannot visit $G(2)$ either, there is no point to check if it is to visit $B(2)$, and hence v_3 does not need to have a child with index label 2 (in this case the node v_{32}). This simple idea already puts a cap on the size of *STS*. But it turns out that we can save the most if we exploit the redundancy on B instead of on G .

Reduction of Tree Size. Step (1.3.2.3.1) at a non-leaf node v is to check, for every child v' of v , if a run visits $B(L_h(v'))$, and in the positive case, move the run into a new node. Let v' be a non-root node and v the parent of v' . Let $I_v = L_h^{set}(v)$. As Step (1.3.2) is executed recursively from top to bottom, it can be assured that at the moment of its arriving at v , for every $i \in I_v$, we have $L_s(v') \cap B(i) = \emptyset$. By an abuse of notation, we write $B(v)$ for $\cup_{j \in I_v} B(j)$, and hence we have $L_s(v') \cap B(v) = \emptyset$. Thus, there is no chance of missing a positive obligation even if we restrict L_h to be such that $B(L_h(v')) \not\subseteq B(v)$. It follows that each node “watches” at least one more state that has not been watched by its ancestors, and therefore along any path of an *STS*, there are at most μ nodes with non-zero index labels (recall that $\mu = \min(n, k)$ and note that the root is excluded as its index label is 0). Also, it can be shown by induction on tree height that an *STS* contains at most n nodes with index label 0, using the fact that state labels of sibling nodes are pairwise disjoint and the fact that if v is the parent of v' and $L_h(v') = 0$, then $L_s(v') \subset L_s(v)$. Therefore, the number of nodes in an *STS* is bounded by $n(\mu + 1)$.

Reduction of Index Labels. Let $I'_v = \{i \in I \mid B(i) \subseteq B(v)\}$. The above analysis tells us that $L_h(v') \in (I \setminus I'_v)$. However, we can further improve L_h such that there is no $j \in (I \setminus I'_v)$, $(B(j) \setminus B(v)) \subset (B(L_h(v')) \setminus B(v))$ and for any $j \in (I \setminus I'_v)$, $(B(j) \setminus B(v)) = (B(L_h(v')) \setminus B(v))$ implies $L_h(v') < j$. We say that $L_h(v')$ minimally extends $L_h^\rightarrow(v)$ if this condition holds.

To formalize the intuition of *minimal extension*, we introduce two functions $Cover : I^* \rightarrow 2^I$ and $Mini : I^* \rightarrow 2^I$ as in [3]. $Cover$ maps finite sequences of I -elements to subsets of I such that

$$Cover(\alpha) = \{j \in I \mid B(j) \subseteq \bigcup_{i=1}^{|\alpha|} B(\alpha[i])\}.$$

Note that $Cover(\varepsilon) = \emptyset$. $Mini$ also maps finite sequences of I -elements to subsets of I such that $j \in Mini(\alpha)$ if and only if $j \in I \setminus Cover(\alpha)$ and

$$\forall j' \in I \setminus Cover(\alpha) \left[j' \neq j \rightarrow \neg \left(B(j') \cup \bigcup_{i=1}^{|\alpha|} B(\alpha[i]) \subseteq B(j) \cup \bigcup_{i=1}^{|\alpha|} B(\alpha[i]) \right) \right], \quad (1)$$

$$\forall j' \in I \setminus Cover(\alpha) \left[j' < j \rightarrow \left(B(j') \cup \bigcup_{i=1}^{|\alpha|} B(\alpha[i]) \neq B(j) \cup \bigcup_{i=1}^{|\alpha|} B(\alpha[i]) \right) \right]. \quad (2)$$

$Mini(\alpha)$ consists of index candidates to *minimally* enlarge $Cover(\alpha)$; ties (with respect to set inclusion) are broken by numeric minimality (Condition (2)).

$Mini$ plays a crucial role in reducing the combination of index labels. In the reduced Safra's trees (Definition 3), we identify a collection of paths such that each node appears on exactly one of those paths. The sequence of index labels on each of those paths corresponds to a path in a specific tree structure, which we refer to as *increasing tree of sets (ITS)* [3]. Counting the number of paths in ITS give us a better upper bound on the combination of index labels. Here we switch to an informal notation of labeled trees and we identify a node with the sequence of labels from the root to the node.

Definition 2 (Increasing Tree of Sets (ITS) [3]). *An ITS $\mathcal{T}(n, k, B)$ is an unordered I -labeled tree such that a node α exists in $\mathcal{T}(n, k, B)$ if and only if $\forall i \in [1..|\alpha|]$, $\alpha[i] \in Mini(\alpha[1..i])$.*

Several properties are easily seen from the definition. First, an ITS is uniquely determined by parameter n, k and B . Second, the length of the longest path in $\mathcal{T}(n, k, B)$ is bound by μ . Third, if β is a direct child of α , then β must contribute at least one new element that has not been seen from the root to α . Forth, the new contributions made by β cannot be covered by contributions made by another sibling β' , with ties broken by selecting the one with smallest index. As $B : I \rightarrow 2^Q$ is one to one, we also view ITS as 2^Q -labeled trees.

Example 1 (ITS). Consider $n = 3, k = 4, Q = \{q_0, q_1, q_2\}$, and $B : [1..4] \rightarrow 2^Q$ such that

$$B(1) = \{q_0, q_1\}, \quad B(2) = \{q_0\}, \quad B(3) = \{q_1, q_2\}, \quad B(4) = \{q_2\}.$$

Figure 2 shows the corresponding $\mathcal{T}(n, k, B)$.

4.2 Improvement II

The second idea is a batch-mode naming scheme to reduce name combinations. As shown before, an STS can have $n(\mu + 1)$ nodes, which translates to $(n(\mu + 1))!$ name combinations according to the current “first-come-first-serve” naming scheme, that is, picking an unused name when a new node is created and recycling the name when a node is removed. When $k = O(n)$, the naming cost is higher than all other complexity factors combined, as $(n(\mu + 1))! = 2^{O(nk \lg nk)}$. However, this can be overcome by dividing the name space into even buckets and “wholesaling” buckets to specific paths in an STS .

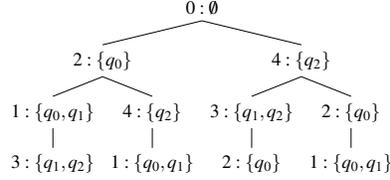


Figure 2: The *ITS* $\mathcal{T}(n, k, B)$ in Example 1. Note that $\{q_0, q_1\}$ and $\{q_1, q_2\}$ cannot appear in the first level because $\{q_0, q_1\}$ covers $\{q_0\}$ and $\{q_1, q_2\}$ covers $\{q_2\}$. The leftmost node at the bottom level is labeled by $\{q_1, q_2\}$ instead of by $\{q_2\}$ due to the index minimality requirement.

Reduction of Names. Let t be an *STS*. A *left spine* (*LS*) is a maximal path $l = v_1 \cdots v_m$ such that v_m is a leaf, for any $i \in [2..m]$, v_i is the left-most child of v_{i-1} , and v_1 is not a left-most child of its parent. We call v_1 the head of l . Let $head : V \rightarrow V$ be such that $head(v) = v'$ if v' is the head of the *LS* where v belongs to. We say that l is the i -th *LS* of t if v_m is the i -th leaf of t , counting from left to right. It is clear that a tree with m leaves has m *LS*, and every node is on exactly one *LS*. Thus, an *STS* has at most n *LS*. For this renaming scheme to work, we require that a newly created node be added to the right of all existing siblings of the same L_h -value. If a non-head node on an *LS* has index label 0, then so should all of its siblings. But then all of them should have been removed. Therefore, only the head of an *LS* can have index label 0, which means that an *LS* can have at most $\mu + 1$ nodes.

We use $n(\mu + 1)$ names and divide them evenly into n buckets. Let b_i ($i \in [1..n]$) denote the i -th bucket, $b_{ij} = (\mu + 1)(i - 1) + j$ ($i \in [1..n], j \in [1..\mu + 1]$) the j -th name in the i -th bucket. We say that $b_{i1} = (\mu + 1)(i - 1) + 1$ is the *initial value* of b_i . Our naming strategy is as follows. Every *LS* $l = v_1 \cdots v_m$ in an *STS* is associated with a name bucket b and v_1, \dots, v_m are assigned names continuously from the initial value of b . For example, if l is associated with bucket b_t , then $L_n(v_i) = (\mu + 1)(t - 1) + i$ for $i \in [1..m]$. The bucket association for each *LS* in an *STS* t can be viewed as selection function $bucket : V \rightarrow [1..n]$ such that $bucket(v) = i$ if node v is assigned a name in the i -th bucket.

This naming strategy, however, comes with a complication; what if a leftmost sibling v is removed (due to Step (1.3.2.3)), and the second leftmost sibling v' (if exists) and all nodes belonging to the *LS* of which v' is the head, “graft into” the *LS* that v belongs to? The answer is that at the end of tree transformation, we need to rename those nodes that have moved into another *LS*. If a tree transformation turns t into t' , and during the process a node v joins another *LS* l in t' (which is also in t before the transformation), then we rename v to a name in the bucket that l uses in t , and recycle the bucket with which v was associated in t .

Example 2 (New Naming Scheme). *Figure 3 illustrates the changes of names in a sequence of tree transformations. We assume that there are 10 buckets $b_1 - b_{10}$, each of which is of size 4. Nodes in the graphs are denoted in the form $v : L_n(v)$; all other types of labels are omitted for simplicity.*

Definition 3 (Reduced Safra Trees for Streett Determinization (μ STS)). A reduced Safra tree for Streett determinization (μ STS) is an *STS* $\langle V, L \rangle$ with $L = \langle L_n, L_s, L_c, L_h \rangle$ that satisfies the following additional conditions:

- 3.1 Condition on L_h . For each node v , if $Mini(L_h^{\rightarrow}(v)) \neq \emptyset$, then v is not a leaf node and for any child v' of v , $L_h(v') \in Mini(L_h^{\rightarrow}(v))$.
- 3.2 Condition on L_n . There exists a function $bucket : V \rightarrow [1..n]$ such that for every *LS* $v_1 \cdots v_m$, we have $bucket(v_i) = bucket(v_j)$ for $i, j \in [1..m]$ and $L_n(v_i) = (\mu + 1)(bucket(v_i) - 1) + i$ for $i \in [1..m]$.

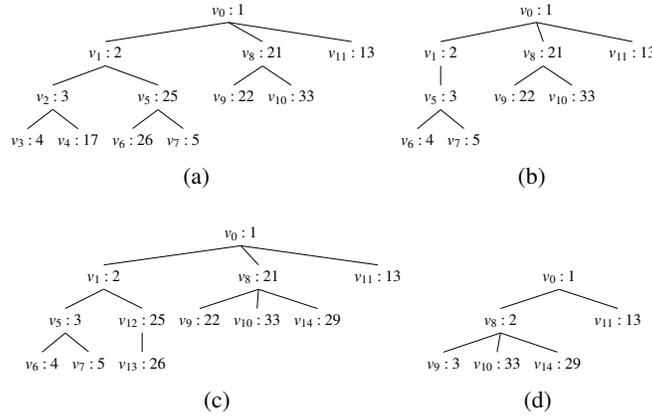


Figure 3: (a) shows an *STS* with 7 *LS*: $l_1 : v_0v_1v_2v_3$, $l_2 : v_4$, $l_3 : v_5v_6$, $l_4 : v_7$, $l_5 : v_8v_9$, $l_6 : v_{10}$ and $l_7 : v_{11}$, associated with buckets $b_1, b_5, b_7, b_2, b_6, b_9$ and b_4 , respectively. (b) shows the resulting *STS* after removing v_2 and its descendants. Nodes v_5 and v_6 migrate into l_1 , and accordingly their name bucket b_7 is recycled. Also recycled is bucket b_5 due to the deletion of v_4 . (c) shows the resulting *STS* after adding nodes v_{12}, v_{13} and v_{14} and forming two new *LS*: $v_{12}v_{13}$ and v_{14} . Here v_{12} and v_{13} reuse the previously recycled bucket b_7 , while node v_{14} takes an unused bucket b_8 . (d) shows the resulting *STS* after removing v_1 and its descendants. Nodes v_8 and v_9 migrate into l_1 and take names 2 and 3, respectively. Buckets b_2, b_7 and b_6 are recycled accordingly.

As before, nodes in a μ *STS* are partially ordered by *older-than* relation. But we impose an additional *structural ordering* on nodes, that, for any two sibling v and v' , v' is placed to the right of v if and only if $L_h(v) > L_h(v')$, or $L_h(v) = L_h(v')$ and v is *older than* v' . This structural ordering is needed for our renaming scheme (see the proof of Theorem 2).

Condition (3.2), together with the requirement that L_n is injective (Condition (1.1)), guarantees that no two distinct *LS* in a tree share a bucket. Condition (3.1) says that every μ *STS* has fully grown left spines, that is, no leaf v can be further extended, as $\text{Mini}(L_h^{\rightarrow}(v)) = \emptyset$. To achieve this, we need the following procedure applied as the last step of each tree transformation.

Procedure 2 (Grow Left Spine). *Repeat the following procedure until no new nodes can be added: if v is a leaf and $\text{Mini}(L_h^{\rightarrow}(v)) \neq \emptyset$, add a new child v' to v with $L_s(v') = L_s(v)$, $L_h(v') = \max(\text{Mini}(L_h^{\rightarrow}(v)))$, $L_c(v') = \text{red}$, and $L_n(v') = L_n(v) + 1$.*

We note that the requirement that a μ *STS* has fully grown left spines is not essential; we can “grow” a μ *STS* “on-the-fly” as in [19, 15]. But this requirement simplifies the analysis on the number of combinations of index labels (see the proof of Theorem 2).

Rule 2 (Naming and Coloring on μ *STS*). *Naming and Coloring convention for μ STS is the one for STS plus the following: (1) nodes in an *LS*, from the head downwards, are assigned continuously increasing names, starting from the initial value of a bucket; (2) when an *LS* is created, nodes in the *LS* are assigned names from an unused name bucket; when an *LS* is removed (which only happens when its head is removed), the name bucket of the *LS* is recycled; (3) when an *LS* l is grafted into another *LS* l' , the name bucket of l is recycled and nodes on l are renamed according to (1), as if they were on l' originally; (4) renamed nodes are marked red.*

Procedure 3 (Improved Streett Determinization). *Let $\mathcal{A} = \langle Q, Q_0, \Sigma, \Delta, \langle G, B \rangle_I \rangle$ be a nondeterministic Streett automaton. This procedure outputs a deterministic Rabin automaton $\mathcal{B} = \langle \tilde{Q}, \tilde{q}_0, \Sigma, \tilde{\Delta}, [\tilde{G}, \tilde{B}]_{\tilde{I}} \rangle$,*

where \tilde{Q} is the set of μ STS, $\tilde{I} = [1..n(\mu + 1)]$, $[\tilde{G}, \tilde{B}]_{\tilde{I}}$ is as defined in Procedure 1, \tilde{q}_0 is a tree that is just a fully grown left spine obtained by growing the single root tree (as defined in Procedure 1) according to Procedure (2), and $\tilde{\Delta}$ is defined such that $\tilde{q}' = \tilde{\Delta}(\tilde{q})$ if and only if \tilde{q}' is the μ STS obtained by applying the following transformation rule to \tilde{q} .

3.1 Subset Construction. For each node v in \tilde{Q} , update state label $L_s(v)$ to $\Delta(L_s(v))$.

3.2 Expansion. Apply the following transformations to non-leaf nodes recursively from the root. Let v be a node with j children v_1, \dots, v_j with i_1, \dots, i_j as the corresponding index labels. Consider the following cases for each $j' \in [1..j]$.

3.2.1 If $L_s(v_{j'}) \cap B(i_{j'}) \neq \emptyset$, then add a child v' to v with $L_s(v') = L_s(v_{j'}) \cap B(i_{j'})$ and $L_h(v') = \max([0..i_{j'}] \cap (\{0\} \cup \text{Mini}(L_h^{\rightarrow}(v))))$, and remove the states in $L_s(v_{j'}) \cap B(i_{j'})$ from $v_{j'}$ and all the descendants of $v_{j'}$.

3.2.2 If $L_s(v_{j'}) \cap B(i_{j'}) = \emptyset$ and $L_s(v_{j'}) \cap G(i_{j'}) \neq \emptyset$, then add a child v' to v with $L_s(v') = L_s(v_{j'}) \cap G(i_{j'})$ and $L_h(v') = L_h(v)$, and remove the states in $L_s(v_{j'}) \cap G(i_{j'})$ from $v_{j'}$ and all the descendants of $v_{j'}$.

3.3 Horizontal Merge. For any state q and any two siblings v and v' such that $q \in L_s(v) \cap L_s(v')$, if $L_h(v) < L_h(v')$, or $L_h(v) = L_h(v')$ and v is older than v' , then remove q from v' and all its descendants.

3.4 Vertical Merge. For each v , if all children of v have index label 0, then remove all descendants of v .

3.5 Grow the tree fully according to Procedure (2).

Note that nodes are assigned or reassigned names in batch only after the whole expansion phase is carried out and the tree is fully grown. Besides naming and renaming, the only major difference between Procedures 3 and 1 is the way of selecting the next positive obligation. In Step (3.2.1), $\text{Mini}(L_h^{\rightarrow}(v))$ is used in the calculation, while in Step (1.3.2.3.1), $L_h^{\text{set}}(v)$ is used.

Theorem 1 (Streett Determinization: Correctness). *Let $\mathcal{A} = \langle Q, Q_0, \Sigma, \Delta, \langle G, B \rangle_I \rangle$ be a Streett automaton with $|Q| = n$ and $I = [1..k]$, and $\mathcal{B} = \langle \tilde{Q}, \tilde{q}_0, \Sigma, \tilde{\Delta}, [\tilde{G}, \tilde{B}]_{\tilde{I}} \rangle$ the deterministic Rabin automaton obtained by Procedure 3. We have $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{B})$.*

Proof. ($\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{B})$). This part of proof is almost identical to the one in [19]. We ought to show that if $\rho = \tilde{q}_0 \tilde{q}_1 \dots$ is an run of \mathcal{B} over an infinite word $w = w_0 w_1 \dots \in \mathcal{L}(\mathcal{A})$, then (1) a node v exists in every state in ρ from some point on, (2) v turns green infinitely often, and (3) v has a fixed name $i \in \tilde{I}$. The argument in [19] guarantees the existence of such a node v with the first two properties. The only complication comes from renaming. We have the situation that v with name i exists in \tilde{q}_j , but it is renamed to i' in the following state \tilde{q}_{j+1} . This happens when v is on an LS l whose head is a second left-most sibling in \tilde{q}_j and the corresponding leftmost sibling is removed in \tilde{q}_{j+1} , resulting in nodes on l (including v) joining another LS l' in \tilde{q}_{j+1} . However, such ‘‘grafting’’ can only happen to v finitely many times, as each time the height of the head of l' is strictly smaller than the height of the head of l . Therefore, v is eventually assigned a fixed name i , which gives us the third property.

($\mathcal{L}(\mathcal{B}) \subseteq \mathcal{L}(\mathcal{A})$). We ought to show that if $w = w_0 w_1 \dots \in \mathcal{L}(\mathcal{B})$, then the run $\rho = \tilde{q}_0 \tilde{q}_1 \dots$ of \mathcal{B} over w induces an accepting run of \mathcal{A} over w . The assumption that ρ is accepting means that there exists an $i \in \tilde{I}$ such that ρ eventually never visits $\tilde{G}(i)$, but visits $\tilde{B}(i)$ infinitely often, or equivalently, i names a green or yellow node in every state in a suffix of ρ and there are infinitely many occurrences when the nodes named by i are green. Since renamed nodes are marked red, all nodes named by i in the suffix have

Cost	Safra's	Piterman's	Ours
Ordered Tree	$2^{O(nk)}$	$2^{O(nk)}$	$2^{O(n \lg n)}$
Name	$2^{O(nk \lg nk)}$	$2^{O(nk \lg nk)}^\dagger$	$2^{O(n \lg n)}$
Color	$2^{O(nk)}$	$O((nk)^2)^\dagger$	$2^{O(n \lg k)}$ $k = O(n)$
			$2^{O(n \lg n)}$ $k = \omega(n)$
Set Label	$2^{O(n \lg n)}$	$2^{O(n \lg n)}$	$2^{O(n \lg n)}$
Index Label	$2^{O(nk \lg nk)}$	$2^{O(nk \lg nk)}$	$2^{O(n \lg n + nk \lg k)}$ $k = O(n)$
			$2^{O(n^2 \lg n)}$ $k = \omega(n)$
Total	$2^{O(nk \lg nk)}$	$2^{O(nk \lg nk)}$	$2^{O(n \lg n + nk \lg k)}$ $k = O(n)$
			$2^{O(n^2 \lg n)}$ $k = \omega(n)$

Figure 4: Cost breakdown of Streett determinization.

to be the same one. It follows that a node v eventually stays in every state in a suffix of ρ and v turns green infinitely often. The rest of the proof is the same as the one in [19], with the help of Lemma 2 and the fact that using *Mini* to select the index labels of the children of v is sound and complete. \square

5 Complexity

In this section we state the complexity results for the determinization of Streett, generalized Büchi, parity and Rabin automata.

Theorem 2 (Streett Determinization: Complexity). *Let $\mathcal{A} = \langle Q, Q_0, \Sigma, \Delta, \langle G, B \rangle_I \rangle$ be a Streett automaton with $|Q| = n$ and $I = [1..k]$, and $\mathcal{B} = \langle \tilde{Q}, \tilde{q}_0, \Sigma, \tilde{\Delta}, [\tilde{G}, \tilde{B}]_{\tilde{I}} \rangle$ the deterministic Rabin automaton obtained by Procedure 3. For any k , $|\tilde{Q}| \leq 2^{12n} (0.37n)^{n^2+8n} = 2^{O(n^2 \lg n)}$ and $|\tilde{I}| = O(n^2)$. If $k = O(n)$, we have $|\tilde{Q}| = 2^{O(n \lg n + nk \lg k)}$ and $|\tilde{I}| = O(nk)$.*

Figure 4 breaks down the total cost into five categories and compares our construction with previous ones in each category. It turns out that the complexity analysis can be easily adapted for generalized Büchi and parity automata as they are subclasses of Streett automata.

Corollary 1 (Generalized Büchi Determinization: Complexity). *Let \mathcal{A} be a generalized Büchi automaton with state size n and index size k . There is an equivalent deterministic Rabin automaton \mathcal{B} with state size $2^{O(n \lg nk)}$. The index size is $O(nk)$ if $k = O(n)$ and $O(n^2)$ if $k = \omega(n)$.*

Corollary 2 (Parity Determinization: Complexity). *Let \mathcal{A} be a parity automaton with state size n and index size k . There is an equivalent deterministic Rabin automaton \mathcal{B} with state size $2^{O(n \lg n)}$ and index size $O(nk)$.*

The determinization of a Rabin automaton with acceptance condition $[G, B]_I$ (the dual of $\langle G, B \rangle_I$ for $I = [1..k]$) can be straightforwardly obtained by running, in parallel, k modified Safra trees, each of which monitors runs for an individual Rabin condition $[G(i), B(i)]$ ($i \in I$).

Theorem 3 (Rabin Determinization: Complexity). *Let \mathcal{A} be a Rabin automaton with state size n and index size k . There is an equivalent deterministic Rabin automaton \mathcal{B} with state size $2^{O(nk \lg n)}$ and index size $O(nk)$.*

† In Piterman's construction, the number of ordered trees times the number of name combinations is bounded by $(nk)^{nk}$. However, the second factor is still the dominating one, costing $2^{\Omega(nk \lg nk)}$. Also, there is no notion of color in Piterman's construction. Instead, each tree is associated with two special names both ranging from 1 to nk , resulting in the cost $O((nk)^2)$.

Type	Bound	Lower	Upper
Büchi	$2^{\Theta(n \lg n)}$	[13]	[17]
Generalized Büchi	$2^{\Theta(n \lg nk)}$	[24]	[10]
Streett	$2^{\Theta(n \lg n + nk \lg k)}$ $k = O(n)$ $2^{\Theta(n^2 \lg n)}$ $k = \omega(n)$	[2]	[3]
Rabin	$2^{\Theta(nk \lg n)}$	[4]	[10]
Parity	$2^{\Theta(n \lg n)}$	[13]	[3]

Figure 5: Complementation and determinization complexities for ω -automata of common types. The listed citations are meant for complementation only.

We note that it is unlikely that there exists a Safra-tree style determinization for Rabin automata, because an analogue of Lemma 2 fails due to the existential nature of Rabin acceptance conditions.

6 Concluding Remarks

In this paper we improved Safra’s construction and obtained tight upper bounds on the determinization complexities of Streett, generalized Büchi, parity and Rabin automata. Figure 5 summarizes these complexity results.

Our results show an interesting phenomenon that in the asymptotic notation $2^{\Theta(\cdot)}$, complementation complexity is identical to determinization complexity. The same phenomenon happens to finite automata on finite words. We believe it is worth investigating the reason behind this phenomenon.

As mentioned earlier, determinization procedures that output parity automata, like Piterman’s constructions, are preferable to the classic ones that output Rabin automata. We plan to investigate how to combine Piterman’s node renaming scheme with ours to obtain determinization procedures that output parity automata with optimal state complexity.

References

- [1] Richard Büchi (1966): *Symposium on Decision Problems: On a Decision Method in Restricted Second Order Arithmetic*. In: *Logic, Methodology and Philosophy of Science, Proceeding of the 1960 International Congress, Studies in Logic and the Foundations of Mathematics* 44, Elsevier, pp. 1–11, doi:10.1016/S0049-237X(09)70564-6.
- [2] Yang Cai & Ting Zhang (2011): *A Tight Lower Bound for Streett Complementation*. In: *Proceedings of the IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2011)*, pp. 339–350, doi:10.4230/LIPIcs.FSTTCS.2011.339.
- [3] Yang Cai & Ting Zhang (2011): *Tight Upper Bounds for Streett and Parity Complementation*. In: *Proceedings of the 20th Conference on Computer Science Logic (CSL 2011)*, Dagstuhl Publishing, pp. 112–128, doi:10.4230/LIPIcs.CSL.2011.112.
- [4] Yang Cai, Ting Zhang & Haifeng Luo (2009): *An Improved Lower Bound for the Complementation of Rabin Automata*. In: *Proceedings of the 24th Annual IEEE Symposium on Logic In Computer Science*, IEEE Computer Society, pp. 167–176, doi:10.1109/LICS.2009.13.
- [5] Yaacov Choueka (1974): *Theories of automata on ω -types: a simplified approach*. *Journal of Computer and System Science* 8(2), pp. 117–141, doi:10.1016/S0022-0000(74)80051-6.

- [6] Thomas Colcombet & Konrad Zdanowski (2009): *A Tight Lower Bound for Determinization of Transition Labeled Büchi Automata*. In: *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP 2009)*, Springer-Verlag, pp. 151–162, doi:10.1007/978-3-642-02930-1_13.
- [7] Nissim Francez (1986): *Fairness*. Springer-Verlag New York, Inc.
- [8] Nissim Francez & Dexter Kozen (1984): *Generalized fair termination*. In: *Proceedings of the 11th ACM SIGACT-SIGPLAN symposium on Principles of programming languages (POPL'84)*, ACM, pp. 46–53, doi:10.1145/800017.800515.
- [9] Yuri Gurevich & Leo Harrington (1982): *Trees, Automata, and Games*. In: *Proceedings of the 14th annual ACM symposium on Theory of computing (STOC'82)*, pp. 60–65, doi:10.1145/800070.802177.
- [10] Orna Kupferman & Moshe Y. Vardi (2005): *Complementation Constructions for Nondeterministic Automata on Infinite Words*. In: *Tools and Algorithms for the Construction and Analysis of Systems, Lecture Notes in Computer Science 3440*, Springer Berlin / Heidelberg, pp. 206–221, doi:10.1007/978-3-540-31980-1_14.
- [11] Christof Löding (1999): *Optimal Bounds for Transformations of omega-Automata*. In: *Proceedings of the 19th Conference on Foundations of Software Technology and Theoretical Computer Science*, Springer-Verlag, pp. 97–109, doi:10.1007/3-540-46691-6_8.
- [12] Robert McNaughton (1966): *Testing and Generating Infinite Sequences by a Finite Automaton*. *Information and Control* 9(5), pp. 521–530, doi:10.1016/S0019-9958(66)80013-X.
- [13] M. Michel (1988): *Complementation is more difficult with automata on infinite words*. CNET.
- [14] David E. Muller & Paul E. Schupp (1995): *Simulating alternating tree automata by nondeterministic automata: New results and new proofs of the theorems of Rabin, McNaughton and Safra*. *Theoretical Computer Science* 141(1-2), pp. 69–107, doi:10.1016/0304-3975(94)00214-4.
- [15] N. Piterman (2006): *From Nondeterministic Buchi and Streett Automata to Deterministic Parity Automata*. In: *21st Annual IEEE Symposium on Logic in Computer Science*, pp. 255–264, doi:10.1109/LICS.2006.28.
- [16] M. O. Rabin & D. Scott (1959): *Finite automata and their decision problems*. *IBM Journal of Research and Development* 3, pp. 114–125, doi:10.1147/rd.32.0114.
- [17] S. Safra (1988): *On the complexity of omega -automata*. In: *Proceedings of the 29th Annual Symposium on Foundations of Computer Science*, IEEE Computer Society, pp. 319–327, doi:10.1109/SFCS.1988.21948.
- [18] S. Safra & M. Y. Vardi (1989): *On ω -automata and temporal logic*. In: *Proceedings of the 21st annual ACM symposium on Theory of computing (STOC'89)*, ACM, pp. 127–137, doi:10.1145/73007.73019.
- [19] Shmuel Safra (1992): *Exponential Determinization for omega-Automata with Strong-Fairness Acceptance Condition (Extended Abstract)*. In: *Proceedings of the 24th annual ACM symposium on Theory of computing (STOC'92)*, pp. 275–282, doi:10.1145/129712.129739.
- [20] Sven Schewe (2009): *Büchi Complementation Made Tight*. In: *Proceedings of the 26th International Symposium on Theoretical Aspects of Computer Science (STACS 2009)*, pp. 661–672, doi:10.4230/LIPIcs.STACS.2009.1854.
- [21] Sven Schewe (2009): *Tighter Bounds for the Determinization of Büchi Automata*. In: *Proceedings of the 12th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS 2009)*, pp. 167–181.
- [22] Stefan Schwoon (2002): *Determinization and Complementation of Streett Automata*. In: *Automata Logics, and Infinite Games, Lecture Notes in Computer Science 2500*, Springer Berlin / Heidelberg, pp. 257–264, doi:10.1007/3-540-36387-4_5.
- [23] Moshe Y. Vardi (2007): *The Büchi complementation saga*. In: *Proceedings of the 24th International Symposium on Theoretical Aspects of Computer Science (STACS 2007)*, Springer-Verlag, pp. 12–22, doi:10.1007/978-3-540-70918-3_2.
- [24] Qiqi Yan (2006): *Lower Bounds for Complementation of ω -Automata Via the Full Automata Technique*. In: *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP 2006)*, pp. 589–600, doi:10.1007/11787006_50.