

An Improved Lower Bound for the Complementation of Rabin Automata

Yang Cai
MIT CSAIL
The Stata Center, 32-G580
Cambridge, MA 02139 USA
ycai@csail.mit.edu

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

Haifeng Luo
Princeton University
201 Sherrerd Hall
Princeton, NJ 08544 USA
haifengl@princeton.edu

Abstract

Automata on infinite words (ω -automata) have wide applications in formal language theory as well as in modeling and verifying reactive systems. Complementation of ω -automata is a crucial instrument in many these applications, and hence there have been great interests in determining the state complexity of the complementation problem. However, obtaining nontrivial lower bounds has been difficult. For the complementation of Rabin automata, a significant gap exists between the state-of-the-art lower bound $2^{\Omega(N \lg N)}$ and upper bound $2^{O(kN \lg N)}$, where k , the number of Rabin pairs, can be as large as 2^N . In this paper we introduce multidimensional rankings to the full automata technique. Using the improved technique we establish an almost tight lower bound for the complementation of Rabin automata. We also show that the same lower bound holds for the determinization of Rabin automata.

1 Introduction

Automata on infinite words (ω -automata) have wide applications in formal language theory as well as in modeling and verifying reactive systems. In many these applications a crucial instrument is *complementation*, which is to construct an automaton $C\mathcal{A}$ from a given an automaton \mathcal{A} , such that $C\mathcal{A}$ accepts an infinite word if and only if \mathcal{A} does not accept. For instance, in automata-theoretic model checking, to find out whether a system represented by an automaton \mathcal{B} satisfies a specification represented by another automaton \mathcal{A} , one checks if $\mathcal{L}(\mathcal{B}) \subseteq \mathcal{L}(\mathcal{A})$, which reduces to $\mathcal{L}(\mathcal{B}) \cap \mathcal{L}(C\mathcal{A}) = \emptyset$ [Kur94, VW94]. Thus, determining the state complexity of the complementation problem of ω -automata has a significant value in practice and it has attracted great interests in the last four decades [Var07].

Büchi first invented a kind of ω -automata (now

called Büchi automata) as a tool to study decision problems of second-order arithmetic [Büc62]. Over the years many variants of ω -automata have been proposed, including Rabin automata and Streett automata. These common variants only differ at the definition of acceptance conditions and they all recognize ω -regular languages. Although equivalent in expressiveness, ω -automata with rich acceptance conditions, such as Streett automata and Rabin automata can express properties more easily and succinctly than Büchi automata. For example, the strong fairness condition [FK84, Fra86] that every infinitely enabled transition in a run is also taken infinitely often, can be expressed straightforwardly by a Streett acceptance condition. On the other hand, Rabin condition, defined as the dual of Streett condition, can directly express unfairness. The fair termination problem [Fra86, KK91], that is, whether all fair computations terminate, can be easily encoded into a Rabin condition that requires all infinite computations be either unfair or eventually forever idling. For these reasons, tightening the bounds for the complementation problem of other types of ω -automata also has great practical value.

The complementation problem of Büchi automata have been investigated for over 40 years. The current best algorithm has $O(N^2((0.76 + c_0)N)^N)$ state blow-up (for a fixed $c_0 \in (0, 1)$) [Sch09], which tightly matches the best lower bound $\Omega(((0.76 + c_0)N)^N)$ (for the same $c_0 \in (0, 1)$) [Yan06]. However, for the complementation of Rabin automata, a huge gap still exists between the state-of-the-art lower bound $2^{\Omega(N \lg N)}$ [Yan06] and upper bound $2^{O(kN \lg N)}$ [KV05a], where k , the number of Rabin pairs, can be as large as 2^N . A similar huge gap exists for the complementation of Streett automata [KV05a, Yan06].

In this paper we generalize the *full automata* technique [Yan06] to incorporate multiple dimensional ranking functions. Using the generalized method we show that for any $\epsilon > 0$, the lower bound for the com-

plementation of Rabin automata with N states, k Rabin pairs, and an alphabet of size N^2 is $2^{\Omega(kN \lg N)}$ if $k \leq 2^{N(1-\epsilon)}$, and is $2^{\Omega(2^{N(1-\epsilon)}N \lg N)}$ if $k > 2^{N(1-\epsilon)}$. For a Rabin automaton with N states, the number of effective Rabin pairs can be at most 2^N because we can merge two Rabin pairs (G_0, B) and (G_1, B) into one pair $(G_0 \cup G_1, B)$ without changing the recognized language. In this sense, our lower bound is almost tight in terms of big- O notation.

The condition of using alphabets of unbounded cardinality can be removed via an encoding trick. We show that for any $\epsilon > 0$, a constant d exists such that the above lower bound holds for any alphabet of size d . We can further reduce the alphabet size to a fixed small constant (independent of ϵ) with the lower bound adjusted to $2^{\Omega(kN \lg N)}$ for $k \leq 2^{\frac{N}{2}(1-\epsilon)}$ and to $2^{\Omega(2^{\frac{N}{2}(1-\epsilon)}N \lg N)}$ for $k > 2^{\frac{N}{2}(1-\epsilon)}$. We also show that these lower bound results apply to the determinization of Rabin automata. Note that all lower bounds in this paper apply to any complementation or determinization algorithm which outputs ω -automata of common types (see Acceptance Conditions in Section 2).

Related Work and Comparison. The first complementation construction for Büchi automata was given by Büchi and that construction requires $2^{2^{O(N)}}$ states [Büc62]. The construction was improved to $2^{O(N^2)}$ states by Sistla, Vardi and Wolper [SVW87]. Safra gave a determinization construction for Büchi automata, from which a complementation construction with $2^{O(N \lg N)}$ states was obtained [Saf88]. This upper bound matches well with the lower bound $N! \approx (0.36N)^N = 2^{\Omega(N \lg N)}$ proved by Michel [Mic88, Löd99]. Klarlund later gave a construction with $2^{O(N \lg N)}$ states without using determinization [Kla91]. Klarlund's construction relies on *quasi co-Büchi measure*, which is a ranking function on states in a run graph, measuring the progress of a run toward being accepted. Kupferman and Vardi proposed a complementation construction that uses *co-Büchi ranking* (similar to quasi co-Büchi measure). The construction is essentially the same as Klarlund's, but provides a better lower bound $O((6N)^N) = 2^{O(N \lg N)}$ [KV01]. With refined constructions, the upper bound was further improved to $O((0.9624N)^N)$ by Friedgut, Kupferman and Vardi [FKV06], and then most recently to $O(N^2(((0.76 + c_0)N)^N))$ (for a fixed $c_0 \in (0, 1)$) by Schewe [Sch09]. In 2006, Yan introduced rankings into *full automata* technique and obtained a sequence of sharper lower bounds for complementation and determinization of ω -automata [Yan06]. In particular, Yan sharpened the lower bound for the complementation of Büchi au-

tomata to $\Omega(((0.76 + c_0)N)^N)$ (for the same $c_0 \in (0, 1)$), which now is only quadratically smaller than the best upper bound obtained by Schewe. In [Yan06], Yan also showed that the lower bound holds for any complementation construction whose output automata are of common types. This immediately gives a $2^{\Omega(N \lg N)}$ lower bound for the complementation of Rabin automata because a Büchi automaton can be viewed as a Rabin automaton by simply reinterpreting the Büchi condition as a Rabin condition. In the contrast, the state-of-the-art complementation construction for Rabin automata, introduced by Kupferman and Vardi in [KV05a], requires $2^{O(kN \lg N)}$ states.

Sakoda and Sipser introduced the *full automata*¹ technique and used it to obtain several completeness and lower bound results on transformations involving 2-way finite automata [SS78]. Full automata operate on unconventional and large alphabets; in a full automaton of N states, every possible unit transition graph (bipartite graph with $2N$ vertices) is identified with a letter, and words are nothing but potential runs of the automaton. Using this technique, in [SS78] Sakoda and Sipser also proved the classic result that the lower bound of complementing finite automata (on finite words) is 2^N . In the proof, the difficult word serving as the witness of this lower bound is just a difficult run that cannot be accepted by any automaton with less than 2^N states. The power of full automata technique, however, does not rely on using alphabets of unbounded cardinality; by an encoding trick, a large alphabet can be mapped to a small alphabet containing only a few letters, with little compromise to the lower bound results [Sip79, Yan06].

To the best of our knowledge, although the full automata technique offers a systematic way for constructing difficult witnesses, it has never been applied to obtain lower bounds for transformations of ω -automata until Yan extended the technique with rankings [Yan06]. Rankings used in [Yan06] bear certain similarities to those used in the work of Klarlund, Friedgut, Kupferman and Vardi [Kla91, KV01, FKV06, KV05a]. However, these two kinds of rankings are designed for different purposes. In [Kla91, KV01, FKV06, KV05a], a word is in the complementary language of a Büchi automaton if and only if there exists an odd co-Büchi ranking (or quasi co-Büchi measure) on the run graph of the word. A complementary automaton is so constructed to recognize run graphs with odd co-Büchi rankings (or quasi co-Büchi measures). In [Yan06], the rankings are designed to show that for a family of full automata $\{\mathcal{F} \mathcal{A}_n\}$, a family of difficult

¹The term was first coined in [Yan06] where the definition is slightly different from that used in [SS78].

words (difficult runs) $\{\alpha_n\}$ exists such that for each n , α_n is not recognized by \mathcal{FA}_n nor by any “small” complementary automaton of \mathcal{FA}_n . Our lower bound proof relies on rankings in the same vein as Yan’s. To obtain tighter lower bounds, however, we use a type of multi-dimensional rankings, which result in a construction considerably more sophisticated than that used in [Yan06]. Our rankings are also different from the old Streett rankings that were used in [KV05a] for the complementation of Rabin automata.

Our generalization relies on two key notions: (1) Q_k -rankings and (2) Υ -graphs. A Q_k -ranking is k -dimensional function mapping states to k -tuples of integers. In fact a Q_k -ranking can be viewed as k independent bijective functions on states. A transition graph is called Q_k -ranked if every level of the graph is associated with a Q_k ranking. A Υ -graph is a special Q_k -ranked transition graph that satisfies four properties designed for constructing difficult words. These properties are parameterized with a pair of states and an index value in between 1 and k . It is not hard to construct a Υ -graph that satisfies the four properties for a specific instantiation of the parameters. The technical difficulty, however, lies in how to accommodate the four properties for each pair of states and for each index simultaneously. Our solution is to use “bypasses” (called *Refuge* and *Tunnel*) to concatenate a sequence of Υ -graphs each of which satisfies the four properties for a specific pair of states and for a specific index. The bypasses make the concatenation behaves like a parallel composition so that properties satisfied by each fragment are all preserved in the final concatenation.

Paper Organization. Section 2 presents notations used in this paper and basic terminology in automata theory. Section 3 generalizes the full automata technique and proves the lower bound. Section 4 presents the detailed construction with examples. Section 5 concludes with a discussion of future work. Due to space limitation, figures and the proofs for technical lemmas and theorems are omitted. They are available in the full version of this paper, which is available on authors’ webpages.

2 Preliminaries

Basic Notations. \mathbb{N} denotes natural numbers. We write $[i..j]$ for $\{k \in \mathbb{N} \mid i \leq k \leq j\}$ and $[n]$ to mean $[0..n-1]$. If u be a sequence, we use $|u|$ to denote the length of u , $u(i)$ ($i \in [|u|]$) to denote the object at the i -th position, and $u[i..j]$ ($i, j \in [|u|]$) to denote the subsequence of u from position i to position j . We use $u \circ v$ to denote the concatenation of u and v and

when little confusion is present, we simply use the juxtaposition uv .

ω -automata. A finite automaton on infinite words is a tuple $\mathcal{A} = (\Sigma, S, I, \Delta, \mathcal{F})$ where Σ is an alphabet, S is a finite set of states, $I \subseteq S$ is a set of initial states, $\Delta \subseteq S \times \Sigma \times S$ is a set of transition relations, and \mathcal{F} is an acceptance condition. We say that \mathcal{A} is *deterministic* when $|I| = 1$ and for all $p \in S, \sigma \in \Sigma, \{q \in S \mid \langle p, \sigma, q \rangle \in \Delta\} \leq 1$. Unless stated otherwise, all automata we consider in this paper are nondeterministic ω -automata.

Run of Automata. An infinite word (ω -words) over Σ is an infinite sequence of letters in Σ . A *run* of \mathcal{A} on an ω -word w is an infinite sequence of states in S such that $\rho(0) \in I$ and, $\langle \rho(i), w(i), \rho(i+1) \rangle \in \Delta$ for $i \in \mathbb{N}$. We use $\rho[v_1, v_2]$ to denote the subsequence $\rho(v_1)\rho(v_1+1)\cdots\rho(v_2)$. Let $Occ(\rho)$ be the set of states occurring in ρ and $Inf(\rho)$ the set of states that occur infinitely many times in ρ . A *finite run* of \mathcal{A} from state p to state q over a finite word w is a finite sequence of states $\rho = \rho(0)\rho(1)\cdots\rho(|w|)$ such that $\rho(0) = p, \rho(|w|) = q$ and $\langle \rho(i), w(i), \rho(i+1) \rangle \in \Delta$ for all $i \in [|w|]$.

Acceptance Conditions. Automata on infinite words are classified according to acceptance conditions. We say that a type of ω -automata is *common* if the acceptance condition of this type is defined solely with respect to $Inf(\rho)$. Below we list some common types of ω -automata relevant to this paper.

- *Büchi automata*, where $\mathcal{F} \subseteq S$, and ρ is accepting iff $Inf(\rho) \cap \mathcal{F} \neq \emptyset$.
- *Rabin automata*, where $\mathcal{F} = \{\langle G_1, B_1 \rangle, \dots, \langle G_k, B_k \rangle\}$, and ρ is accepting iff for some $i \in [1..k]$, we have that $Inf(\rho) \cap G_i \neq \emptyset$ and $Inf(\rho) \cap B_i = \emptyset$. States in G_i and B_i are called, respectively, *reconfirming* states and *invalidating* states.
- *Streett automata*, where $\mathcal{F} = \{\langle G_1, B_1 \rangle, \dots, \langle G_k, B_k \rangle\}$, and ρ is accepting iff for all $i \in [1..k]$, if $Inf(\rho) \cap G_i \neq \emptyset$, then $Inf(\rho) \cap B_i \neq \emptyset$.

We say that an ω -word w is accepted by \mathcal{A} if there exists an accepting run of \mathcal{A} over w . By $\mathcal{L}(\mathcal{A})$ we denote the set of ω -words accepted by \mathcal{A} .

Δ -Graph. Let $\mathcal{A} = (\Sigma, S, I, \Delta, \mathcal{F})$ be an automaton. A Δ -graph² of an ω -word w under \mathcal{A} is a directed graph $\mathcal{G}_w = (V, E)$ where $V = S \times \mathbb{N}$ and $E = \{\langle \langle p, i \rangle, \langle q, i+1 \rangle \rangle \in V \times V \mid p, q \in S, i \in \mathbb{N}, \langle p, w(i), q \rangle \in \Delta\}$. By the i -th level, we mean the vertex set $S \times \{i\}$. The Δ -graph of

a finite word is defined similarly. Let w be a finite word. By $|\mathcal{G}_w|$ we denote the length of \mathcal{G}_w , which is the same as $|w|$. We call a path in \mathcal{G}_w a *full path* if the path goes from level 0 to level $|\mathcal{G}_w|$. By $\mathcal{G}_w \circ \mathcal{G}_{w'}$, we mean the concatenation of \mathcal{G}_w and $\mathcal{G}_{w'}$, which is the graph obtained by merging the last level of \mathcal{G}_w with the first level of $\mathcal{G}_{w'}$. Note that $\mathcal{G}_w \circ \mathcal{G}_{w'} = \mathcal{G}_{ww'}$.

\mathcal{G}_w is a visualization of the *complete* behavior of \mathcal{A} over w . It is easily seen that Δ can be identified with a function $\Delta' : \Sigma \rightarrow 2^{S \times S}$ such that $\langle p, q \rangle \in \Delta'(\sigma)$ iff $\langle p, \sigma, q \rangle \in \Delta$ for any $\sigma \in \Sigma$. With indices dropped, \mathcal{G}_σ , the Δ -graph of a letter σ , is a just the graph of the relation $\Delta'(\sigma)$. By abusing notation, we identify $\Delta'(\sigma)$ with \mathcal{G}_σ and \mathcal{G}_w (where $w = \sigma_0\sigma_1\dots$) with $\Delta'(w) = \Delta'(\sigma_0) \circ \Delta'(\sigma_1) \circ \dots$.

Let w be a finite word. For $v_0, v_1 \in \mathbb{N}$, $p_0, p_1 \in S$ we write $(p_0, v_0) \xrightarrow{w} (p_1, v_1)$ to mean that there exists a run ρ of \mathcal{A} such that $\rho[v_0, v_1]$ is a finite run of \mathcal{A} from p_0 to p_1 over w . For $S_0, S_1 \subseteq S$, we use $(p_0, v_0) \xrightarrow[S_0, S_1]{w} (p_1, v_1)$ to mean, in addition, that $\rho[v_0, v_1]$ contains an S_0 -state but no S_1 -state. We write $(p_0, v_0) \xrightarrow[\sim_{S_1}]{w} (p_1, v_1)$ to mean that $\rho[v_0, v_1]$ does not contain an S_1 -state. In case the indices of a run are of no importance, we simply drop them and write $p_0 \xrightarrow{w} p_1$, $p_0 \xrightarrow[S_0, S_1]{w} p_1$, and $p_0 \xrightarrow[\sim_{S_1}]{w} p_1$.

Full Automata. A full automaton $\mathcal{A} = (\Sigma, S, I, \Delta, \mathcal{F})$ is an automaton such that $\Sigma = 2^{(S \times S)}$, $\Delta \subseteq S \times 2^{(S \times S)} \times S$, and for all $p, q \in S$, $\sigma \in \Sigma$, $\langle p, \sigma, q \rangle \in \Delta$ if and only if $\langle p, q \rangle \in \sigma$ [SS78, Yan06].

For a full automaton, Σ and Δ are completely determined by S . Now Δ' is just the identity function on $2^{(S \times S)}$ and hence there is no difference between words and their corresponding Δ -graphs. From now on we use two terms interchangeably.

3 Lower Bound

In this section we extend the full automata technique with multidimensional ranking functions and we use the generalized technique to obtain an almost tight lower bound for the complementation of Rabin automata.

Proof Plan. The key of this lower bound proof is to construct a family of full Rabin automata $\{\mathcal{F}\mathcal{R}_n\}$ and a corresponding family of difficult words $\{\alpha_n\}$ such that

²The Δ -Graphs defined in this paper are slightly different from the *run graphs* (or *run dags*) used in [Kla91, KV01, KV05a, FKV06]. We do not require that any vertex in a Δ -graph be reachable from an initial state at the first level.

for each n , $\alpha_n \notin \mathcal{L}(\mathcal{F}\mathcal{R}_n)$, yet no ω -regular language that separates α_n from $\mathcal{L}(\mathcal{F}\mathcal{R}_n)$ can be recognized by a “small” ω -automaton of any common type. We first construct $\mathcal{F}\mathcal{R}_n$ (Definition 1) with respect to which we define Q_k -rankings (Definition 2) and Υ -graphs (Definition 3). A Q_k -ranking is a function on the state set of $\mathcal{F}\mathcal{R}_n$ into a set of k -tuples of integers. A Δ -graph is said to be Q_k -ranked if its every level is associated with a Q_k -ranking. Υ -graphs are special Q_k -ranked Δ -graphs that satisfy several properties that are consequences of Properties (3.1)-(3.4) (Definition 3). First and most importantly, for any pair of Q_k -rankings $\langle f, g \rangle$, there exists a Υ -graph $\mathcal{G}_{\langle f, g \rangle}$ whose first level ranking is f and last level ranking is g (Theorem 1). Such a $\mathcal{G}_{\langle f, g \rangle}$ is said to be R -compatible with $\langle f, g \rangle$ (Definition 4). Theorem 1 is the most difficult technical part of this paper and we leave its proof to Section 4. Second, for any Q_k -rankings f, g, h , if $\mathcal{G}_{\langle f, g \rangle}$ is R -compatible with $\langle f, g \rangle$ and $\mathcal{G}_{\langle g, h \rangle}$ is R -compatible with $\langle g, h \rangle$, then $\mathcal{G}_{\langle f, g \rangle} \circ \mathcal{G}_{\langle g, h \rangle}$ is R -compatible with $\langle f, h \rangle$ (Lemma 2). By this transitivity property, we can construct α_n by repeatedly enumerating Q_k -rankings. Property (3.1) guarantees that $\alpha_n \notin \mathcal{L}(\mathcal{F}\mathcal{R}_n)$, and Properties (3.2) and (3.3) imply that for any ω -automata \mathcal{A} that accepts α_n , if $\mathcal{L}(\mathcal{A}) \cap \mathcal{L}(\mathcal{F}\mathcal{R}_n) = \emptyset$, then the number of states of \mathcal{A} must exceed the number of distinct Q_k -rankings. Our lower bound is so obtained.

Definition 1 (Full Rabin Automata). We define a family of full Rabin automata $\mathcal{F}\mathcal{R}_n = (\Sigma, S, I, \Delta, \mathcal{F})$ such that

1.1 $S = I \cup R \cup T \cup G \cup B$ where I, R, T, G and B are pairwise disjoint and assume the following forms:

$$I = \{s_0, \dots, s_{n-1}\}, \quad T = \{t_0, \dots, t_{n-1}\}, \\ R = \{r_0, \dots, r_{n-1}\}, \quad B = \{b_1, \dots, b_{2\gamma}\}, \quad G = \{\ddot{g}\};$$

1.2 $\mathcal{F} = \{\langle G, B_1 \rangle, \dots, \langle G, B_k \rangle\}$ where $B_i \subseteq B$, $|B_i| = \gamma$ and $B_i \neq B_j$ if $i \neq j$.

Note that $\mathcal{F}\mathcal{R}_n$ is defined with three parameters: n , γ and k , all required to be positive integers. Thus all notions derived from $\mathcal{F}\mathcal{R}_n$ should also be parameterized with the three. But for notation simplicity, we selectively omit all or some of them unless there is a chance of confusion. Now let N denote the number of states of $\mathcal{F}\mathcal{R}_n$.

$\mathcal{F}\mathcal{R}_n$ is designed to accommodate sophisticated properties of Υ -graphs (Definition 3) to be introduced below. To reduce the construction complexity, we divides the state set S to five pairwise disjoint subsets I, R, T, G and B , each designated for a different task. I is intended to be the domain of Q_k -rankings. R and T , called *Refuge* and *Tunnel*, respectively, are solely for

building bypasses. G and B are pools from which re-confirming states and invalidating states are chosen to form G_i 's and B_i 's, respectively. The exact lower bound actually is the number of total Q_k -rankings which will be shown to be $(n!)^k = 2^{\Omega(nk \lg n)}$ (Lemma 1). To get the desired lower bound $2^{\Omega(kN \lg N)}$, we want n to be as close to N as possible and k to be as close to 2^N as possible. It turns out that it suffices to let G be a singleton and then let all G_i 's be the same as G . To make k close to 2^N , we require that B_i 's be pairwise unequal subsets of B , all with the cardinality $|B|/2$. In this way, k can be as large as $\binom{2^N}{\gamma}$.

Now we introduce Q_k -ranking, a function that associates a k -tuple of integers with each state in I .

Definition 2 (Q_k -Ranking). A Q_k -ranking for \mathcal{FR}_n is a function $f : I \rightarrow [n]^k$ such that for each $i \in [1..k]$, $f^{(i)} : I \rightarrow [n]$, the projection of f on the i -th coordinate (i.e., $f^{(i)}(p) = f(p)(i)$ for $p \in I$), is bijective.

A Q_k -ranking can be viewed k independent bijective functions from I into $[n]$. Let \mathcal{N} denote the number of Q_k -rankings (again, for simplicity we omit the parameters n and k). We have

Lemma 1. $\mathcal{N} = 2^{\Omega(nk \lg n)}$.

Proof. By \mathcal{N}_1 we denote the number of Q_1 -rankings. By definition, a Q_1 -ranking is a bijective function from I to $[n]$, and therefore $\mathcal{N}_1 = n! = 2^{\Omega(n \lg n)}$. Note that a Q_k -ranking consists of k independent Q_1 -rankings. Therefore $\mathcal{N} = (\mathcal{N}_1)^k = 2^{\Omega(nk \lg n)}$. \square

A Δ -graph is called Q_k -ranked if its every level is associated with a Q_k -ranking. We write $rank_j$ to denote the Q_k -ranking at level j and $rank_j^{(i)}$ to denote the i -th projection of $rank_j$ at level j . Let X be a subset of S . We call a vertex (p, v) in a Δ -graph of \mathcal{FR}_n X -vertex if $p \in X$. When there is no confusion, we just write p for (p, v) . We are interested in a special kind of Q_k -ranked Δ -graphs.

Definition 3 (Υ -Graph). We say that a Q_k -ranked Δ -graph \mathcal{G} is an Υ -graph if the following conditions hold.

3.1 For any $p, q \in I$, if in \mathcal{G} there exists a path l from (p, v_0) to (q, v_1) (for some v_0, v_1) such that l contains no other I -vertex, then for each $i \in [1..k]$,

3.1a if l contains a G -vertex, then $rank_{v_0}^{(i)}(p) > rank_{v_1}^{(i)}(q)$;

3.1b if $rank_{v_0}^{(i)}(p) < rank_{v_1}^{(i)}(q)$, then l contains a B_i -vertex.

3.2 For any $p, q \in I$, $i \in [1..k]$, if $rank_0^{(i)}(p) > rank_{|\mathcal{G}|}^{(i)}(q)$, then $p \xrightarrow[\langle G, B_i \rangle]{\mathcal{G}} q$.

3.3 For any $p, q \in I$, $i \in [1..k]$, if $rank_0^{(i)}(p) = rank_{|\mathcal{G}|}^{(i)}(q)$, then $p \xrightarrow[\sim B_i]{\mathcal{G}} q$.

3.4 In \mathcal{G} only I -vertices have outgoing edges at the first level and incoming edges at the last level, and for every I -vertex at the first level there exists an outgoing path from that vertex to an I -vertex at the last level. In particular, for any $p, q \in S$, if $p \xrightarrow{\mathcal{G}} q$, then $p, q \in I$.

Property (3.1) is of local and *universal* nature; it requires that all paths in between I -vertices satisfy certain conditions, which induces a Streett condition on any infinite path in α_n so that $\alpha_n \notin \mathcal{L}(\mathcal{FR}_n)$ (Lemma 3). In the contrast, Property (3.2) is of global and *existential* nature; it ensures that in α_n , from a vertex with higher rank to a vertex with lower rank (with respect to some index), there exists a "bad" finite path which, if repeated forever, generates an infinite path that satisfies the dual Rabin condition. This property is intended to show that it is hard to separate α_n from $\mathcal{L}(\mathcal{FR}_n)$, in the sense that we can construct a word α'_n from α_n such that α'_n is accepted by \mathcal{FR}_n and no "small" automaton can distinguish α'_n from α_n . Property (3.3) is also of global and *existential* nature; it ensures that in α_n , in between two vertices of the same rank (with respect to certain index i), there is a path that does not visit B_i . This makes sure that Property (3.2) can be maintained during concatenation. Property (3.4) is mainly technical; it guarantees graph connectivity under concatenation.

Definition 4 (R -Compatibility). We say that a word $\mathcal{G} \in (\Sigma)^*$ is R -compatible with an ordered pair of Q_k -rankings $\langle f, g \rangle$ for \mathcal{FR}_n if there exists a Υ -graph of w in which the first level and the last level are ranked by f and g , respectively.

Theorem 1. For any pair of Q_k -rankings $\langle f, g \rangle$, there exists a Υ -graph \mathcal{G} that is R -compatible with $\langle f, g \rangle$.

This is the key theorem of this paper. We leave its proof to Section 4.

Lemma 2 (Transitivity). Let f, g, h be Q_k -rankings for \mathcal{FR}_n , and $\mathcal{G}_{\langle f, g \rangle}$ and $\mathcal{G}_{\langle g, h \rangle}$ be Υ -graphs that are R -compatible with $\langle f, g \rangle$ and $\langle g, h \rangle$, respectively. Then $\mathcal{G}_{\langle f, g \rangle} \circ \mathcal{G}_{\langle g, h \rangle}$ is R -compatible with $\langle f, h \rangle$.

Theorem 1 and Lemma 2 allow us to construct an ω -word which does not belong to $\mathcal{L}(\mathcal{FR}_n)$. Let f_0, f_1, \dots be a repeated enumeration of Q_k -rankings such that (1) for any $i, j < \mathcal{N}$, $f_i \neq f_j$ and for any $i, j \in \mathbb{N}$, $f_i = f_{\mathcal{N}j+i}$. Now define α_n to be the ω -word $\mathcal{G}_0 \mathcal{G}_1 \mathcal{G}_2 \dots$,

such that for $i \geq 0$, $\mathcal{G}_i = \mathcal{G}_{\langle f_i, f_{i+1} \rangle}$, the word compatible with $\langle f_i, f_{i+1} \rangle$.

Lemma 3. *We have $\alpha_n \notin \mathcal{L}(\mathcal{FR}_n)$.*

Proof. Suppose that $\alpha_n \in \mathcal{L}(\mathcal{FR}_n)$. Let ρ be a successful run of \mathcal{FR}_n over α_n . By Property (3.4), infinitely many I -vertices occurs in ρ . So we can assume ρ to be of the form $\rho(k_0) \cdots \rho(k_1) \cdots \rho(k_2) \cdots$ where $k_0 = 0$, and for any $i \in \mathbb{N}$, $\rho(k_i) \in I$ and no intermediate states from $\rho(k_i)$ to $\rho(k_{i+1})$ belong to I . Since ρ is accepted by \mathcal{FR}_n , there is an $i' \in [1..k]$ and a sufficiently large j , such that for any $j' \geq j$, $\rho(j') \notin B_{i'}$. According to Property (3.1b), we have $\text{rank}_{k_j}^{(i')}(\rho(k_j)) \geq \text{rank}_{k_{j+1}}^{(i')}(\rho(k_{j+1})) \geq \text{rank}_{k_{j+2}}^{(i')}(\rho(k_{j+2})) \geq \cdots$, in which strict inequality can only occur finitely many times. As a consequence, by Property (3.1a), the G -vertex can only appear finitely often in ρ , a contradiction. \square

Although α_n is not recognized by \mathcal{FR}_n , it closely “resembles” a word in $\mathcal{L}(\mathcal{FR}_n)$ in the sense that any ω -regular language that separates α_n from $\mathcal{L}(\mathcal{FR}_n)$ cannot be recognized by an automaton with less than \mathcal{N} states. This is established by the following lemma which is tailored for Rabin automata from Lemma 5 in [Yan06]. The proof relies on a counting argument similar to that of Pumping Lemma.

Lemma 4. *Let \mathcal{A} be an ω -automaton of any common type defined on the same alphabet as \mathcal{FR}_n such that $\alpha_n \in \mathcal{L}(\mathcal{A})$ and $\mathcal{L}(\mathcal{A}) \cap \mathcal{L}(\mathcal{FR}_n) = \emptyset$. Then \mathcal{A} must have more than \mathcal{N} states.*

Proof. Suppose the contrary. Let $\mathcal{A} = (\Sigma, S', I', \Delta', \mathcal{F}')$ be an ω -automaton such that $|S'| < \mathcal{N}$, and $\rho = \rho(0)\rho(1) \cdots \in (S')^\omega$ be an accepting run of \mathcal{A} over α_n . We shall show that $\mathcal{L}(\mathcal{A}) \cap \mathcal{L}(\mathcal{FR}_n) \neq \emptyset$.

By the construction of α_n , we know there exists an infinite sequence of indices k_0, k_1, \dots such that $k_0 = 0$, and for any $i \geq 0$, $k_{i+1} - k_i = |\mathcal{G}_i|$ and $\rho(k_i) \xrightarrow{\mathcal{G}_i} \rho(k_{i+1})$. For each $i < \mathcal{N}$, let

$$S_i = \{s^* \in S' \mid \rho(k_{\mathcal{N}+i}) = s^* \text{ for infinitely many } j \in \mathbb{N}\}.$$

Obviously, for each $i < \mathcal{N}$, $S_i \neq \emptyset$. Since \mathcal{A} has less than \mathcal{N} states, there exist two different indices $\mu, \nu < \mathcal{N}$ such that $S_\mu \cap S_\nu \neq \emptyset$.

Since $f_\mu \neq f_\nu$, without loss of generality, we assume that for some state $q \in I$ and $t \in [1..k]$, we have $f_\mu^{(t)}(q) > f_\nu^{(t)}(q)$. Let $s^* \in S_\mu \cap S_\nu$. Because there are infinitely many l satisfying $\rho(k_{\mathcal{N}+l}) = s^*$, there exists a sufficiently large $\zeta \in \mathbb{N}$, such that (1) $\rho(k_{\mathcal{N}+\zeta+\mu}) = s^*$, (2) for every $l > \mathcal{N}\zeta + \mu$, $\rho(l) \in \text{Inf}(\rho)$. For the same reason we can find a sufficiently large η such that (1)

$k_{\mathcal{N}+\eta+\nu} > k_{\mathcal{N}\zeta+\mu}$, (2) $\rho(k_{\mathcal{N}+\eta+\nu}) = s^*$ and (3) $\text{Inf}(\rho) = \text{Occ}(\rho[k_{\mathcal{N}\zeta+\mu}, k_{\mathcal{N}+\eta+\nu-1}])$. Let $\mathcal{G}_u = \mathcal{G}_0\mathcal{G}_1 \cdots \mathcal{G}_{\mathcal{N}\zeta+\mu-1}$, $\mathcal{G}_v = \mathcal{G}_{\mathcal{N}\zeta+\mu}\mathcal{G}_{\mathcal{N}\zeta+\mu+1} \cdots \mathcal{G}_{\mathcal{N}+\eta+\nu-1}$, and $\alpha'_n = \mathcal{G}_u(\mathcal{G}_v)^\omega$. It is easily seen that $\text{Inf}(\rho) = \text{Inf}((\rho[k_{\mathcal{N}\zeta+\mu}, k_{\mathcal{N}+\eta+\nu-1}])^\omega)$. Therefore, $\rho' = \rho[0..k_{\mathcal{N}\zeta+\mu-1}] \circ (\rho[k_{\mathcal{N}\zeta+\mu}, k_{\mathcal{N}+\eta+\nu-1}])^\omega$ is a run of \mathcal{A} over α'_n such that $\text{Inf}(\rho') = \text{Inf}(\rho)$. So $\alpha'_n \in \mathcal{L}(\mathcal{A})$.

Next we show that $\alpha'_n \in \mathcal{L}(\mathcal{FR}_n)$. Recall that there exists $t \in [1..k]$ such that $f_\mu^{(t)}(q) > f_\nu^{(t)}(q)$. By Lemma 2, \mathcal{G}_u is R -compatible with $\langle f_0, f_{\mathcal{N}\zeta+\mu} \rangle = \langle f_0, f_\mu \rangle$, and \mathcal{G}_v is R -compatible with $\langle f_{\mathcal{N}\zeta+\mu}, f_{\mathcal{N}+\eta+\nu} \rangle = \langle f_\mu, f_\nu \rangle$. Since $f_0^{(t)}$ is bijective, there is a state $q_0 \in I$ such that $f_0^{(t)}(q_0) = f_\mu^{(t)}(q)$. By Property (3.3), we have $q_0 \xrightarrow{\mathcal{G}_u} q$. Also, because $f_\mu^{(t)}(q) > f_\nu^{(t)}(q)$, Property (3.2) gives us $q \xrightarrow{\mathcal{G}_v} q$. Now $q_0 \xrightarrow{\mathcal{G}_u} q \xrightarrow{\mathcal{G}_v} q \xrightarrow{\mathcal{G}_v} q \cdots$ gives us an accepting run of \mathcal{FR}_n over α'_n . Thus $\alpha'_n \in \mathcal{L}(\mathcal{FR}_n)$. \square

Note that in Lemma 4, \mathcal{A} only needs to be an automaton of a common type.

Now we are ready for the lower bound proof.

Theorem 2. *For any $\epsilon > 0$, the lower bound for the complementation problem of Rabin automata with N states and k Rabin pairs is $2^{\Omega(kN \lg N)}$ if $k \leq 2^{N(1-\epsilon)}$, and is $2^{\Omega(2^{N(1-\epsilon)}N \lg N)}$ if $k > 2^{N(1-\epsilon)}$.*

Proof. By Lemmas 1 and 4, the lower bound of complementing \mathcal{FR}_n is $2^{\Omega(kn \lg n)}$ if $k \leq \binom{2\gamma}{\gamma}$. Let c be a constant and $\gamma = cn$. So $N = 3n + 2cn + 1 = O(n)$ and we are left to show that if $k \leq 2^{N(1-\epsilon)}$, then $k \leq \binom{2\gamma}{\gamma}$. It follows from Stirling’s inequality ([Rob55]) that for any $m > 0$, $\binom{2m}{m} > \frac{2^{2m}}{2\pi m}$. So for a sufficiently large c we indeed have

$$k \leq 2^{N(1-\epsilon)} < \frac{2^{\frac{(N-1)2c}{3+2c}}}{2\pi \frac{(N-1)c}{3+2c}} < \binom{\frac{(N-1)2c}{3+2c}}{\frac{(N-1)c}{3+2c}} = \binom{2\gamma}{\gamma}. \quad \square$$

As stated in the introduction, no generality is lost by using alphabets of unbounded cardinalities. We can map large alphabets to an alphabet of constant size, with little compromise to our lower bounds.

Theorem 3. *For any $\epsilon > 0$, there exists a constant $d > 0$ such that the lower bound for the complementation problem of Rabin automata over an alphabet of size d and with N states and k Rabin pairs is $2^{\Omega(kN \lg N)}$ if $k \leq 2^{N(1-\epsilon)}$, and is $2^{\Omega(2^{N(1-\epsilon)}N \lg N)}$ if $k > 2^{N(1-\epsilon)}$. For an alphabet of small constant size, the lower bound is $2^{\Omega(kN \lg N)}$ if $k \leq 2^{\frac{N}{2}(1-\epsilon)}$, and is $2^{\Omega(2^{\frac{N}{2}(1-\epsilon)}N \lg N)}$ if $k > 2^{\frac{N}{2}(1-\epsilon)}$.*

As Lemma 4, Theorems 2 and 3 apply to any complementation algorithm that outputs automata of common types. The same lower bounds hold for the determination of Rabin automata.

Corollary 1. *The lower bounds in Theorems 2 and 3 apply to the determinization of Rabin automata.*

Proof. Suppose the contrary. Let \mathcal{A} be a Rabin automaton and \mathcal{B} a deterministic ω -automaton of a common type T such that $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{B})$ and the state size of \mathcal{B} is below our lower bounds. Let T^c denote the dual type of T such that the acceptance condition of T^c is just the negation of the acceptance condition of T . Obviously, T^c is also a common type. Since \mathcal{B} is deterministic, we can obtain an automaton \mathcal{C} of type T^c that complements \mathcal{B} by simply negating the acceptance condition of \mathcal{B} . Now \mathcal{C} complements \mathcal{A} with state size below our lower bounds, a contradiction. \square

4 Construction of Difficult Words

In this section we prove Theorem 1. We shall show that for any pair of Q_k -rankings $\langle f, g \rangle$, there exists a Υ -graph $\mathcal{G}_{\langle f, g \rangle}$ that is R -compatible with $\langle f, g \rangle$.

Proof Plan. We need a construction to simultaneously satisfy all properties in Definition 3. The key challenge lies in making Property (3.2) in harmony with other properties. Our solution is to divide $\mathcal{G}_{\langle f, g \rangle}$ into two sequential subgraphs \mathcal{G}_f and $\mathcal{G}_{f, g}$. \mathcal{G}_f is a Υ -graph with respect to $\langle f, f \rangle$ while $\mathcal{G}_{f, g}$ is almost Υ -graph with respect to $\langle f, g \rangle$ except that it does not satisfy Property (3.2). However, it turns out that Property (3.2) will hold in $\mathcal{G}_f \circ \mathcal{G}_{f, g}$ as follows.

For any $p, q \in I$, $i \in [1..k]$, suppose that $\text{rank}_0^{(i)}(p) > \text{rank}_{|\mathcal{G}_{\langle f, g \rangle}|}^{(i)}(q)$. Since \mathcal{G}_f is a Υ -graph for $\langle f, f \rangle$, we have $\text{rank}_0^{(i)} = \text{rank}_{|\mathcal{G}_f|}^{(i)}$. So we can find a vertex $r \in I$ in the boundary of \mathcal{G}_f and $\mathcal{G}_{f, g}$ such that $\text{rank}_0^{(i)}(p) > \text{rank}_{|\mathcal{G}_f|}^{(i)}(r) = \text{rank}_{|\mathcal{G}_{\langle f, g \rangle}|}^{(i)}(q)$. By Property (3.2), there exists a path l_f from $(p, 0)$ to $(r, |\mathcal{G}_f|)$ such that l_f visits G but no B_i . By Property (3.3), there exists a path $l_{f, g}$ from $(r, |\mathcal{G}_f|)$ to $(q, |\mathcal{G}_{\langle f, g \rangle}|)$ such that $l_{f, g}$ visits no B_i -vertices either. Therefore, $l_f[0..|l_f|] \circ l_{f, g}$ is the desired path from $(p, 0)$ to $(q, |\mathcal{G}_{\langle f, g \rangle}|)$.

\mathcal{G}_f and $\mathcal{G}_{f, g}$ are constructed in a similar manner. First, we construct a sequence of graph fragments, each of which has a portion for building bypasses, and satisfies the corresponding properties with respect to a specific combination of $i \in [1..k]$, $p, q \in I$. Second, we concatenate these fragments in such a way that bypasses are in place to guarantee properties of these fragments are all preserved under concatenation. Let us take a look at an example before going into the details. (The reader should refer to the full version of this paper for the figures used in this example and the examples to follow.)

Example 1 (Υ -Graph $\mathcal{G}_{\langle f, g \rangle}$). Consider \mathcal{FR}_4 where $\gamma = 1$, $k = 2$, and

$$\begin{aligned} I &= \{s_0, s_1, s_2, s_3\}, & R &= \{r_0, r_1, r_2, r_3\}, & B_1 &= \{b_1\}, \\ T &= \{t_0, t_1, t_2, t_3\}, & G &= \{\check{g}\}, & B_2 &= \{b_2\}. \end{aligned}$$

Consider the following pair of Q_2 -rankings f and g .

$$\begin{aligned} f(s_0) &= (3, 3) & f(s_1) &= (1, 2) & f(s_2) &= (2, 0) & f(s_3) &= (0, 1) \\ g(s_0) &= (0, 3) & g(s_1) &= (3, 2) & g(s_2) &= (1, 1) & g(s_3) &= (2, 0) \end{aligned}$$

The first subgraph \mathcal{G}_f is shown in Figure 1 where we omit state sets R and T as they have no use. The second subgraph $\mathcal{G}_{f, g}$ has two parts, which are shown, respectively, as \mathcal{G}_1 in Figure 2 and \mathcal{G}_2 in Figure 3. The complete graph $\mathcal{G}_{\langle f, g \rangle}$ has 45 levels (0..44) in total. Since $f((s_0, 0)) = (3, 3)$

and $g((s_2, 44)) = (1, 1)$, we have $(s_0, 0) \xrightarrow[\langle G, B_1 \rangle]{\mathcal{G}_{\langle f, g \rangle}} (s_2, 44)$ and

$(s_0, 0) \xrightarrow[\langle G, B_2 \rangle]{\mathcal{G}_{\langle f, g \rangle}} (s_2, 44)$. We mark a red (double dotted) path that satisfies the former and a green (double lined) path that satisfies the latter.

The existence of the first subgraph and the second subgraph is established, respectively, by Lemma 5 and Lemma 6. Due to space limitation, for each lemma, we only present a construction for \mathcal{FR}_n with $\gamma = 1$ (denoted by $\mathcal{FR}_n^{(1)}$) and leave the generalization and detailed proof to the appendix.

Lemma 5. *For any Q_k -rankings f , there exists a Υ -graph \mathcal{G}_f that is R -compatible with $\langle f, f \rangle$.*

Proof Sketch. We show the construction of \mathcal{G}_f for $\mathcal{FR}_n^{(1)}$. Since $\gamma = 1$, we have $k \leq 2$, but we keep k as a parameter for the sake of later generalization. Also in this setting B_i 's ($i \in [1..k]$) are singletons, and hence with no loss of generality we assume that $B_i = \{b_i\}$.

As mentioned before, \mathcal{G}_f is obtained by concatenating a sequence of graph fragments, each of which satisfies Properties (3.1)-(3.4) with respect to a specific combination of $i \in [1..k]$, $p, q \in I$. It turns out that I suffices for building bypasses and so there is no need of R or T in \mathcal{G}_f .

The construction uses the following letters:

$$\begin{aligned} Id(I) &= \{\langle p, p \rangle \mid p \in I\} & ToG(p) &= Id(I) \cup \{\langle p, \check{g} \rangle\} \\ GtoB_j &= Id(I) \cup \{\langle \check{g}, b_j \rangle\} & B_i to B_j &= Id(I) \cup \{\langle b_i, b_j \rangle\} \\ FrG(p) &= Id(I) \cup \{\langle \check{g}, p \rangle\} & FrB_i(p) &= Id(I) \cup \{\langle b_i, p \rangle\} \end{aligned}$$

where $p \in I$ and $i, j \in [1..k]$. For $i \in [1..k]$, $p, q \in I$ we define $d^{(i)}(p, q)$ to be

$$ToG(p) \circ GtoB_k \circ \dots \circ B_{i+1} to B_{i-1} \circ \dots \circ B_2 to B_1 \circ FrB_1(q)$$

Note that $d^{(k)}(p, q) = \text{ToG}(p) \circ \text{GtoB}_{k-1} \circ \dots \circ \text{B}_2 \text{toB}_1 \circ \text{FrB}_1(q)$ and similar adjustment for $d^{(1)}(p, q)$. It is not hard to verify that for fixed i, p, q , if $f^{(i)}(p) > f^{(i)}(q)$, then $d^{(i)}(p, q)$ satisfies Properties (3.1)-(3.4). Indeed, since G is visited in $d^{(i)}(p, q)$, Property (3.2) is satisfied. Property (3.3) holds trivially because B_i is not visited in $d^{(i)}(p, q)$. Property (3.1) also holds trivially because of the assumption $f^{(i)}(p) > f^{(i)}(q)$. Property (3.4) obviously holds because all letters contain $\text{Id}(I)$.

Formally, \mathcal{G}_f is a concatenation of subgraphs $\mathcal{G}_1, \dots, \mathcal{G}_k$ where each \mathcal{G}_i in turn is a concatenation of graphs of the form $d^{(i)}(p, q)$ where p, q satisfy $f^{(i)}(p) > f^{(i)}(q)$. Every level of \mathcal{G}_f is ranked by f . By the definition of Q_k -rankings, for each $i \in [1..k], j \in [n]$, there is one and only one state $p_j^{(i)} \in I$, such that $f^{(i)}(p_j^{(i)}) = j$. So the set $\{(p_{j_1}^{(i)}, p_{j_2}^{(i)}) \mid j_1, j_2 \in [n], j_1 > j_2\}$ contains all the combinations of (p, q) that we need. It is not hard to see that the set $\{(p_{j+1}^{(i)}, p_j^{(i)}) \mid j \in [n-1]\}$ just serves our purpose as long as we let a path go down through ranks step by step. Put all together, we have $\mathcal{G} = \mathcal{G}_1 \circ \dots \circ \mathcal{G}_k$ where for $i \in [1..k], \mathcal{G}_i$ is

$$d^{(i)}(p_n^{(i)}, p_{n-1}^{(i)}) \circ d^{(i)}(p_{n-1}^{(i)}, p_{n-2}^{(i)}) \circ \dots \circ d^{(i)}(p_2^{(i)}, p_1^{(i)}). \quad \square$$

We show how to select a path l_f from $(p, 0)$ to $(q, |\mathcal{G}_f|)$ that satisfies Property (3.2) for a fixed index i . Beginning from $(p, 0)$, l_f takes horizontal edges passing through $\mathcal{G}_0, \dots, \mathcal{G}_{i-1}$ until it enters \mathcal{G}_i and reaches $d^{(i)}(p, p')$ for some $p' \in I$ such that $f^{(i)}(p) = f^{(i)}(p') + 1$. Then l_f takes the only path that leads to p' in another horizontal track, decreasing the i -th rank by 1. Repeating this process, each step going to a vertex whose i -th rank is one less, eventually l_f reaches q at some level and from there it only takes horizontal edges till reaching $(q, |\mathcal{G}_f|)$.

Example 2 (Υ -Graph \mathcal{G}_f). Let us revisit Example 1 and take a close look at \mathcal{G}_f which is shown in Figure 1. Every level of \mathcal{G}_f is ranked by f and \mathcal{G}_f has two parts: \mathcal{G}_1 and \mathcal{G}_2 . In \mathcal{G}_1 , after a path leaves an I -vertices, it visits G , then B_2 and comes back to an I -vertex whose 1-st rank is 1 less than that of the last visited I -vertex. Similarly, in \mathcal{G}_2 , after a path leaves an I -vertices, it visits G , then B_1 and comes back to an I -vertex whose 2-nd rank is 1 less than that of the last visited I -vertex. Property (3.1) therefore holds for both indices 1 and 2. The red (double dotted) path is a witness for Property (3.2) with respect to $(s_0, 0)$, $(s_1, 18)$ and index 1, and the green (double lined) path is a witness for Property (3.2) with respect to $(s_0, 0)$, $(s_3, 18)$ and index 2. Property (3.3) is satisfied by any horizontal path from level 0 to level 18. Property (3.4) is obvious.

Lemma 6. For any pair of Q_k -rankings $\langle f, g \rangle$, there exists a Q_k -ranked Δ -graph $\mathcal{G}_{f,g}$ such that the first level and the

last level of G is ranked by f and g respectively, and $\mathcal{G}_{f,g}$ satisfies Properties (3.1), (3.3), and (3.4).

Proof Sketch. We show the construction of $\mathcal{G}_{f,g}$ for $\mathcal{FR}_n^{(1)}$. As before we keep k as a parameter for the later generalization and we assume that $B_i = \{b_i\}$ for $i \in [1..k]$.

The idea underlying the construction is the same as before. The desired graph $\mathcal{G}_{f,g}$ is a concatenation of subgraphs $\mathcal{G}_1, \dots, \mathcal{G}_k$ where each \mathcal{G}_i fulfills the requirements with respect to the i -th index. In each \mathcal{G}_i , two kinds of bypasses help preserve properties of all subgraphs under concatenation. T (*Tunnel*) is intended to build paths that satisfy Property (3.3) for this index i , while R (*Refuge*) is to let pass through \mathcal{G}_i those paths that are obliged to satisfy Property (3.3) for indices other than i .

To simplify the construction of each \mathcal{G}_i , we introduce a sequence of $k-1$ transitional Q_k -rankings that gradually bridges the difference between f and g , in such a way that any two adjacent rankings differ at exactly one coordinate. In each \mathcal{G}_i , the Q_k -rankings of all levels but the last one are the same and they differ from the Q_k -ranking of the last level only at the i -th coordinate. Formally, we define, for each $i \in [k]$,

$$\langle f, g \rangle_i = (g^{(1)}, \dots, g^{(i)}, f^{(i+1)}, \dots, f^{(k)}).$$

Note that $\langle f, g \rangle_0 = f$, $\langle f, g \rangle_k = g$, and $\langle f, g \rangle_{i-1} = \langle f, g \rangle_i^{(i)}$ for $j \neq i$. We will assign $\langle f, g \rangle_i$ to the last level of \mathcal{G}_i and $\langle f, g \rangle_{i-1}$ to all other levels of \mathcal{G}_i .

We also need to satisfy other properties. Property (3.4) is easily seen satisfied once the whole construction is given. Property (3.1a) is also easy to accommodate because no other properties require that the G -vertex be visited. So we simply do not use the G -vertex and Property (3.1a) trivially holds. Property (3.1b) indeed needs more care (see below).

Now we begin the formal construction. For $i, j \in [1..k]$, we define the following letters.

$$\begin{aligned} S' &= I \cup T \cup R & \text{Id}(S') &= \{\langle p, p \rangle \mid p \in S'\} \\ S_i \text{to} B_j &= \text{Id}(S') \cup \{\langle s_i, b_j \rangle\} & S_i \text{to} R_j &= \text{Id}(S') \cup \{\langle s_i, r_j \rangle\} \\ S_i \text{to} T_j &= \text{Id}(S') \cup \{\langle s_i, t_j \rangle\} & B_i \text{to} B_j &= \text{Id}(S') \cup \{\langle b_i, b_j \rangle\} \\ B_i \text{to} R_j &= \text{Id}(S') \cup \{\langle b_i, r_j \rangle\} & B_i \text{to} T_j &= \text{Id}(S') \cup \{\langle b_i, t_j \rangle\} \\ \text{Converge} &= \{\langle t_j, s_j \rangle \mid j \in [n]\} \cup \{\langle r_j, s_j \rangle \mid j \in [n]\} \end{aligned}$$

As mentioned before, each \mathcal{G}_i is to fulfill Property (3.3) for index i . That means that for every pair of states (p, q) ($p, q \in I$) such that $\langle f, g \rangle_{i-1}^{(i)}(p) = \langle f, g \rangle_{i-1}^{(i)}(q)$, \mathcal{G}_i should contain a full path from $(p, 0)$ to $(q, |\mathcal{G}_i|)$ such that no B_i -vertex appears on the path. This is done by using *Tunnel*. For each $j \in [n]$, there exists one and exactly one $j' \in [n]$ such that $\langle f, g \rangle_{i-1}^{(i)}(s_j) = \langle f, g \rangle_{i-1}^{(i)}(s_{j'})$.

So let $\tilde{h} : [n] \rightarrow [n]$ be such a bijective function and then let $toTunnel_j^{(i)}$ be

$$S_j to B_{i_1} \circ B_{i_1} To B_{i_2} \circ \dots \circ B_{i_{t-1}} To B_{i_t} \circ B_{i_t} to T_{\tilde{h}_i(j)},$$

where i_1, \dots, i_t is a decreasing enumeration of

$$I_j^{(i)} = \{l \in [n] \mid \langle f, g \rangle_{i-1}^{(l)}(s_j) < \langle f, g \rangle_i^{(l)}(s_{\tilde{h}_i(j)})\}.$$

Note that if $I_j^{(i)} = \emptyset$, we have $toTunnel_j^{(i)} = S_j to T_{\tilde{h}_i(j)}$. In $toTunnel_j^{(i)}$, no B_i -vertex is visited, and before a path jumps to the $\tilde{h}_i(j)$ -th track in $toTunnel_j^{(i)}$, the path first visits a B_l -vertex for every $l \in I_j^{(i)}$, which is to respect Property (3.1b) for indices other than i . These visits surely do not violate Property (3.3) for this i because $i \notin I_j^{(i)}$. Now let

$$ToTunnel^{(i)} = toTunnel_0^{(i)} \circ toTunnel_1^{(i)} \circ \dots \circ toTunnel_{n-1}^{(i)},$$

which contains a tunnel for each $j \in [n]$.

We are done if we only consider Properties (3.1) and (3.3) for \mathcal{G}_i . However, to guarantee these properties hold for $\mathcal{G}_{f,g}$, we need to create a bypass in each \mathcal{G}_i such that a full path in $\mathcal{G}_{f,g}$ that is obliged to satisfy these properties for coordinate i will take the corresponding bypass in every \mathcal{G}_j with $j \neq i$. This is achieved by using *Refuge*. We define $toRefuge_j^{(i)}$ to be

$$\begin{array}{ll} S_j to B_i \circ B_i to R_j & \text{if } \langle f, g \rangle_{i-1}^{(i)}(s_j) < \langle f, g \rangle_i^{(i)}(s_j) \\ S_j to R_j & \text{otherwise} \end{array}$$

which makes sure that there is a bypass from s_j to s_j that visits B_i if $\langle f, g \rangle_{i-1}^{(i)}(s_j) < \langle f, g \rangle_i^{(i)}(s_j)$. Therefore, Property (3.1b) for this i is respected. Obviously Property (3.1b) for other indices are not violated because $\langle f, g \rangle_{i-1}^{(j)} = \langle f, g \rangle_i^{(j)}$ for $i \neq j$. Now let $ToRefuge^{(i)}$ be

$$toRefuge_0^{(i)} \circ toRefuge_1^{(i)} \circ \dots \circ toRefuge_{n-1}^{(i)},$$

which contains a refuge for each $j \in [n]$.

Put all together we define

$$\mathcal{G}_i = Id(I) \circ ToRefuge^{(i)} \circ ToTunnel^{(i)} \circ Converge,$$

where $Id(I)$ is to force paths to enter \mathcal{G}_i only through I -vertices, and $Converge$ is to synchronize both T -vertices and R -vertices with the corresponding I -vertices before they leave \mathcal{G}_i . Note that in each \mathcal{G}_i , $Converge$ forces that any full path from s_j to s_j either goes through *Tunnel* or goes through *Refuge*. In the former case, we have $\langle f, g \rangle_{i-1}^{(i)}(s_j) = \langle f, g \rangle_i^{(i)}(s_j)$, and in the latter case, we have $j = j'$. \square

We show how to select a path $l_{f,g}$ from $(s_j, 0)$ to $(s_{j'}, |\mathcal{G}_{f,g}|)$ that satisfies Property (3.3) for a fixed index i . Beginning from $(p, 0)$, $l_{f,g}$ uses *Refuge* to pass through $\mathcal{G}_0, \dots, \mathcal{G}_{i-1}$ until it enters \mathcal{G}_i . Inside \mathcal{G}_i , $l_{f,g}$ takes the horizontal track to reach $toTunnel_j^{(i)}$ and from there it takes a path to enter *Tunnel*. It continues to move through *Tunnel* until it reaches $s_{j'}$ at the last level of \mathcal{G}_i . Then again $l_{f,g}$ uses *Refuge* to pass through the remaining $\mathcal{G}_{i+1}, \dots, \mathcal{G}_k$ until reaching $(s_{j'}, |\mathcal{G}_{f,g}|)$.

Example 3 (Υ -Graph $\mathcal{G}_{f,g}$). Let us revisit Example 1 to have a close look at $\mathcal{G}_{f,g}$, which has two parts, \mathcal{G}_1 , shown in Figure 2, and \mathcal{G}_2 , shown in Figure 3. All levels but the last one of \mathcal{G}_1 and \mathcal{G}_2 are ranked, respectively, by f and $\langle f, g \rangle_1$. The last levels of \mathcal{G}_1 and \mathcal{G}_2 are ranked, respectively, by $\langle f, g \rangle_1$ and g . In \mathcal{G}_1 , a path starting from $(s_i, 18)$ ($i \in [1..k]$) has to go through either *Refuge* or *Tunnel*. In the former case, the path ends at $(s_i, 32)$ and in the latter case, the path ends at $(s_j, 32)$ for some $j \in [1..k]$ such that $f^{(1)}(s_i) = \langle f, g \rangle_1^{(1)}(s_j)$. Similarly for paths in \mathcal{G}_2 .

A short path from $(s_3, 23)$ to $(r_3, 25)$ via $(b_1, 24)$ (in $toRefuge_3^{(1)}$) is intended for any path passing through *Refuge* to respect Property (3.1b) because $f^{(1)}(s_3, 20) < \langle f, g \rangle_1^{(1)}(s_3, 32)$. Similarly for the path from $(s_1, 20)$ to $(r_1, 22)$ via $(b_1, 21)$ (in $toRefuge_1^{(1)}$) and the path from $(s_2, 35)$ to $(r_2, 37)$ via $(b_2, 36)$ (in $toRefuge_2^{(2)}$).

The red (double dotted) path from $(s_1, 18)$ to $(s_2, 44)$ satisfies Property (3.3) for index 1 and the green (double lined) path from $(s_3, 18)$ to $(s_2, 44)$ satisfies Property (3.3) for index 2. The former path uses the *Tunnel* in \mathcal{G}_1 and *Refuge* in \mathcal{G}_2 while the latter uses *Refuge* in \mathcal{G}_1 and *Tunnel* in \mathcal{G}_2 . Property (3.1b) holds for the reason stated above. Property (3.1a) holds vacuously because no G -vertex is visited. Property (3.4) is obvious as usual.

We are ready to prove the key theorem of this paper.

Theorem 1. For any pair of Q_k -rankings $\langle f, g \rangle$, there exists a Υ -graph $\mathcal{G}_{\langle f, g \rangle}$ that is R -compatible with $\langle f, g \rangle$.

Proof. Let \mathcal{G}_f be an Υ -graph that satisfies Lemma 5, and let $\mathcal{G}_{f,g}$ be a Q_k -ranked Δ -graph that satisfies Lemma 6. Let $\mathcal{G}_{\langle f, g \rangle} = \mathcal{G}_f \circ \mathcal{G}_{f,g}$. We show that $\mathcal{G}_{\langle f, g \rangle}$ satisfies Properties (3.1)-(3.4). Let $p, q \in I$, $i \in [1..k]$.

Suppose that $f^{(i)}(p) = g^{(i)}(q)$. By Lemma 5 (Property (3.3)), $p \xrightarrow[\sim B_i]{\mathcal{G}_f} p$ and by Lemma 6 (Property (3.3)), $p \xrightarrow[\sim B_i]{\mathcal{G}_{f,g}} q$. Therefore, $p \xrightarrow[\sim B_i]{\mathcal{G}_{\langle f, g \rangle}} q$, which gives us Property (3.3) for $\mathcal{G}_{\langle f, g \rangle}$.

Suppose that $f^{(i)}(p) > g^{(i)}(q)$. Because $f^{(i)}$ is bijective, there exists $r \in I$ such that $f^{(i)}(p) > f^{(i)}(r) = g^{(i)}(q)$. By Lemma 5 (Property (3.2)), $p \xrightarrow[\langle G, B_i \rangle]{\mathcal{G}_f} r$ and by Lemma 6

(Property (3.3)), $r \xrightarrow[\sim_{B_i}]{\mathcal{G}_{f,g}} q$. Therefore, $p \xrightarrow[\langle G, B_i \rangle]{\mathcal{G}_{f,g}} q$, which proves Property (3.2) for $\mathcal{G}_{f,g}$.

Property (3.4) is easily seen as it is proved both in Lemma 5 and in Lemma 6.

If a path starts and ends with I -vertices and has no other I -vertices in between, then it must be confined either in \mathcal{G}_f or in $\mathcal{G}_{f,g}$. Property (3.1) then follows as it is proved both in Lemma 5 and in Lemma 6. \square

5 Conclusion

In this paper we generalized the full automata technique with multidimensional ranking functions. Using the improved technique we obtained an almost tight lower bound for the complementation of Rabin automata. We also showed that the same lower bound holds for the determinization of Rabin automata. Note that our lower bounds can be further improved. In the proof, a Q_k -ranking is defined to be a sequence of k independent bijective functions from $[n]$ to $[n]$, and hence the number of Q_k -rankings is $(n!)^k \approx (0.36n)^{kn}$. Our proof can be adapted to use a new type of Q_k -rankings, each of which is a sequence of k independent *tight level rankings* as defined in [FKV06, Yan06]. In this way, the number of new Q_k -rankings is at least $(0.76n)^{kn}$. But this change does not affect the lower bounds expressed in the big-O notation, and hence we chose the current definition for simplicity of the proof.

We plan to use the improved technique to tighten bounds for other types of automata transformations. In particular, we are interested in investigating the complementation problem of Streett automata. The current best lower bound for this problem is $(\Omega(kN))^N$ ([Yan06]) while the best upper bound is $2^{O(kN \lg kN)}$ ([KV05a]), where k , the number of Streett pairs, can be as large as 2^N .

Acknowledgements. We are grateful to Professor Tian Liu of Peking University for reviewing the early draft of this paper and providing many valuable suggestions to improve it. We thank Qiqi Yan for comments concerning the upper bound results in [FKV06, Sch09] and for pointing out an error. We also thank anonymous reviewers for their valuable comments. Part of the work was done when the second author was at Microsoft Research Asia, and the first and the third author were at Peking University.

References

[Büc62] J.R. Büchi. On a decision method in restricted second order arithmetic. In *Proc. Internat. Congr. Logic,*

Method. and Philos. Sci. 1960, pages 1-12, Stanford, 1962. Stanford University Press.

[FK84] N. Francez and D. Kozen. Generalized fair termination. In *Proc. 11th POPL*, pages 46-53, 1984.

[FKV06] E. Friedgut and O. Kupferman and M.Y. Vardi. Büchi complementation made tighter. *International Journal of Foundations of Computer Science*, Vol. 17, No. 4 (2006) 851-867.

[Fra86] N. Francez. Fairness. *Texts and Monographs in Computer Science*. Springer-Verlag, 1986.

[KK91] N. Klarlund and D. Kozen. Rabin measures and their applications to fairness and automata theory. In *Proc. 6th LICS*, pages 256-265, 1991.

[Kla91] N. Klarlund. Progress measures for complementation of omega-automata with applications to temporal logic. In *Proc. 32th FOCS*, pages 358-367, 1991.

[Kur94] R.P. Kurshan. Computer aided verification of coordinating processes: an automata theoretic approach. Princeton University Press, 1994.

[KV01] O. Kupferman and M.Y. Vardi. Weak alternating automata are not that weak. *ACM Transactions on Computational Logic*, 2(3): 408-429, 2001.

[KV05a] O. Kupferman and M.Y. Vardi. Complementation constructions for nondeterministic automata on infinite words. In *Proc. 11th TACAS*, pages 206-221, 2005.

[Löd99] C. Löding. Optimal bounds for transformations of omega-automata. In *Proc. 19th FSTTCS*, volume 1738 of *LNCS*, pages 97-109, 1999.

[Mic88] M. Michel. Complementation is more difficult with automata on infinite words. *CNET*, Paris, 1988.

[Rob55] H. Robbins. A remark on Stirling's formula. *The American Mathematical Monthly*, 62(1):26-29, 1955.

[RS59] M. O. Rabin and D.Scott. Finite automata and their decision problems. *IBM Journal of Research and Development*, 3:115-125, 1959.

[Saf88] S. Safra. On the complexity of ω -automata. In *Proc. 29th FOCS*, pages 319-327, 1988.

[Sch09] S. Schewe. Büchi complementation made tight. In *Proc. 26th STACS*, pages 661-672, 2009.

[SS78] W. Sakoda, M. Sipser. Nondeterminism and the size of two way finite automata. In *Proc. 10th STOC*, pages 275-286, 1978.

[Sip79] M. Sipser. Lower bounds on the size of sweeping automata. In *Proc. 11th STOC*, pages 360-364, 1979.

[SVW87] A. P. Sistla, M.Y. Vardi, and P.Wolper. The complementation problem for Büchi automata with applications to temporal logic. *Theoretical Computer Science*, 49:217-327, 1987.

[Var07] M.Y. Vardi. The Büchi complementation saga. In *Proc. 24th STACS*, pages 12-22, 2007.

[VW94] M.Y. Vardi. and P. Wolper. Reasoning about infinite computations. *Information and Computation*, 115(1):1-37, November 1994.

[Yan06] Q. Yan. Lower bound for complementation of ω -automata via the full automata technique. In *Proc. 33th ICALP*, volume 4052 of *LNCS*, pages 589-600, 2006.