

A Tight Lower Bound for Streett Complementation*

Yang Cai¹ and Ting Zhang²

- 1 MIT Computer Science and Artificial Intelligence Laboratory
The Stata Center, 32-G696, Cambridge, MA 02139 USA
ycai@csail.mit.edu
- 2 Department of Computer Science, Iowa State University
226 Atanasoff Hall, Ames, IA 50011 USA
tingz@iastate.edu

Abstract

Finite automata on infinite words (ω -automata) proved to be a powerful weapon for modeling and reasoning infinite behaviors of reactive systems. Complementation of ω -automata is crucial in many of these applications. But the problem is non-trivial; even after extensive study during the past two decades, we still have an important type of ω -automata, namely Streett automata, for which the gap between the current best lower bound $2^{\Omega(n \lg nk)}$ and upper bound $2^{\Omega(nk \lg nk)}$ is substantial, for the Streett index size k can be exponential in the number of states n . In [4] we showed a construction for complementing Streett automata with the upper bound $2^{O(n \lg n + nk \lg k)}$ for $k = O(n)$ and $2^{O(n^2 \lg n)}$ for $k = \omega(n)$. In this paper we establish a matching lower bound $2^{\Omega(n \lg n + nk \lg k)}$ for $k = O(n)$ and $2^{\Omega(n^2 \lg n)}$ for $k = \omega(n)$, and therefore showing that the construction is asymptotically optimal with respect to the $2^{\Theta(\cdot)}$ notation.

1998 ACM Subject Classification F.1.1 Models of Computation, F.4.1 Mathematical Logic, F.4.3 Formal Languages

Keywords and phrases ω -automata, Streett automata, complementation, lower bounds

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2011.339

1 Introduction

Complementation is a fundamental notion in automata theory. Given an automaton \mathcal{A} , the complementation problem asks to find an automaton \mathcal{B} that accepts exactly all words that \mathcal{A} does not accept. Complementation connects automata theory with mathematical logic due to the natural correspondence between language complementation and logical negation, and hence plays a pivotal role in solving many decision and definability problems in mathematical logic.

A fundamental connection between automata theory and the monadic second order logics was demonstrated by Büchi [1], who started the theory of finite automata on infinite words (ω -automata) [2]. The original ω -automata are now referred to as Büchi automata and Büchi complementation was a key to establish that the class of ω -regular languages (sets of ω -words generated by product \circ , union \cup , star $*$ and limit $^\omega$) is closed under complementation [2].

Büchi's discovery also has profound repercussions in applied logics. Since the '80s, with increasing demand of reasoning infinite computations of reactive and concurrent systems,

* This research has been supported by NSF CAREER Award CCF-0954132.



© Yang Cai and Ting Zhang;

licensed under Creative Commons License NC-ND

31st Int'l Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2011).

Editors: Supratik Chakraborty, Amit Kumar; pp. 339–350

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

ω -automata have been acknowledged as unifying representation for *programs* as well as for *specifications* [26]. Complementation of ω -automata is crucial in many of these applications.

But complementation of ω -automata is non-trivial. Only after extensive studies in the past two decades [23, 16, 18, 6, 27, 20] (also see survey [25]), do we have a good understanding of the complexity of Büchi complementation. But a question about a very important type of ω -automata remains unanswered, namely the complexity of Streett complementation, where the gap between the current lower bound and upper bound is substantial. Streett automata are ones of a kind, because Streett acceptance conditions naturally encode *strong fairness* that infinitely many requests are responded infinitely often, a necessary requirement for meaningful computations [5, 7].

1.1 Related Work

Obtaining nontrivial lower bounds has been difficult. The first nontrivial lower bound for Büchi complementation is $n! \approx (0.36n)^n$, obtained by Michel [16, 15]. In 2006, combining ranking with *full automaton* technique, Yan improved the lower bound of Büchi complementation to $\Omega(L(n))$ [27], which now is matched tightly by the upper bound $O(n^2(L(n)))$ [20], where $L(n) \approx (0.76n)^n$. Also established in [27] was a $(\Omega(nk))^n = 2^{\Omega(n \lg nk)}$ tight lower bound (where k is the number of Büchi indices) for generalized Büchi complementation, which also applies to Streett complementation because generalized Büchi automata are a subclass of Streett automata. In [3], we proved a tight lower bound $2^{\Omega(nk \lg n)}$ for Rabin complementation (where Rabin index size k can be as large as $2^{n-\epsilon}$ for any arbitrary but fixed $\epsilon > 0$). Several constructions for Streett complementation exist [24, 9, 19, 14, 17], but all involve at least $2^{O(nk \lg nk)}$ state blow-up, which is significantly higher than the current best lower bound $2^{\Omega(n \lg nk)}$, since the Streett index size k can reach 2^n . Determining the complexity of Streett complementation has been posed as an open problem since the late '80s [24, 14, 27, 25]. In [4] we showed a construction for Streett complementation with the upper bound $2^{O(n \lg n + nk \lg k)}$ for $k = O(n)$ and $2^{O(n^2 \lg n)}$ for $k = \omega(n)$. In this paper we establish a matching lower bound $2^{\Omega(n \lg n + nk \lg k)}$ for $k = O(n)$ and $2^{\Omega(n^2 \lg n)}$ for $k = \omega(n)$, and therefore showing that the construction in [4] is essentially optimal at the granularity of $2^{\Theta(\cdot)}$. This lower bound is obtained by applying two techniques: *fooling set* and *full automaton*.

1.2 Fooling Set

The fooling set technique is a classic way of obtaining lower bounds on nondeterministic finite automata on finite words (NFA). Let Σ be an alphabet and $\mathcal{L} \subseteq \Sigma^*$ a regular language. A set of pairs $P = \{(x_i, y_i) \mid x_i, y_i \in \Sigma^*, 1 \leq i \leq n\}$ is called a *fooling set* for \mathcal{L} , if $x_i y_i \in \mathcal{L}$ for $1 \leq i \leq n$ and $x_i y_j \notin \mathcal{L}$ for $1 \leq i, j \leq n$ and $i \neq j$. If \mathcal{L} has a fooling set P , then any NFA accepting \mathcal{L} has at least $|P|$ states [8]. The purpose of a fooling set is to identify runs with dual properties (called fooling runs): fragments of accepting runs of \mathcal{L} , when pieced together in certain ways, induce non-accepting runs. By an argument in the style of Pumping Lemma, a small automaton would not be able to distinguish how it arrives at a state, and hence it cannot differentiate between some accepting runs and some non-accepting ones.

In the setting of ω -automata, a similar technique exists, which we refer to as Michel's scheme [16]. A set $P = \{x_i \in \Sigma^* \mid 1 \leq i \leq n\}$ is called a *fooling set* for \mathcal{L} , if $(x_i)^\omega \in \mathcal{L}$ for $1 \leq i \leq n$ and $((x_i)^+(y_j)^+)^\omega \subseteq \overline{\mathcal{L}}$ for $1 \leq i, j \leq n$ and $i \neq j$ [16, 15].

1.3 Full Automaton

Sakoda and Sipser introduced the *full automaton* technique [21] (the name was first coined in [27]) and used it to obtain several completeness and lower bound results on transformations involving 2-way finite automata [21]. In particular, they proved a classic result of automata theory: the lower bound of complementing an NFA with n states is 2^n .

To establish lower bounds for complementation, one starts with designing a class of automata \mathcal{A}_n and then a class of words \mathcal{W}_n such that \mathcal{W}_n are not contained in $\mathcal{L}(\mathcal{A}_n)$. Next one shows that runs of purported complementary automata \mathcal{C}_n on \mathcal{W}_n exhibit dual properties by application of the fooling set technique. However, some fooling runs can only be generated by long and sophisticated words, which are very difficult to be “guessed” right from the beginning. The ingenuity of the full automaton technique is to remove two levels of indirections: since the ultimate goal is to construct fooling runs, why should not one start with runs directly, and build \mathcal{W}_n and \mathcal{A}_n later?

Without a priori constraints imposed from \mathcal{A}_n or \mathcal{W}_n (they do not exist yet), full automata operate on all possible runs; for 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 run graphs. Removing the two levels of indirections proved to be powerful. By this technique, the 2^n lower bound proof for complementing NFA was surprisingly short and easy to understand [21] (a fooling set method was implicit in the proof).

We should note that full automata operate on large alphabets whose size grows exponentially with the state size, but this does not essentially limit its application to automata on conventional alphabets. By an encoding trick, a large alphabet can be mapped to a small alphabet with no compromise to lower bound results [22, 27, 3].

1.4 Ranking

For ω -automata, the power of fooling set and full automaton technique was further enhanced by the use of rankings on run graphs [27, 3]. Since first introduced in [9], rankings have been shown to a powerful tool to represent properties of run graphs; complementation constructions for various types of ω -automata were obtained by discovering respective rankings that precisely characterize those run graphs that contain no accepting path (with respect to source automata) [12, 13, 14, 6, 10]. With the help of rankings, constructing a fooling set amounts to designing certain type of rankings. In fact, as shown below, an explicit description of a fooling set might be very hard to find, but the essential properties the fooling set induce can be concisely represented by certain type of rankings.

1.5 Our Results

In this paper we establish a lower bound $L(n, k)$ for Streett complementation: $2^{\Omega(n \lg n + kn \lg k)}$ for $k = O(n)$ and $2^{\Omega(n^2 \lg n)}$ for $k = \omega(n)$, which matches the upper bound obtained in [4]. This lower bound applies to all Streett complementation constructions that output union-closed automata (see Section 2), which include Büchi, generalized Büchi and Streett automata. This bound considerably improves the current best bound $2^{\Omega(n \lg nk)}$ [27], especially in the case $k = \Theta(n)$.

Determinization is another fundamental concept in automata theory and it is closely related to complementation. A deterministic T -automaton can be easily complemented by switching from T -acceptance condition to the dual co- T condition (e.g., Streett vs. Rabin). Therefore, the lower bound $L(n, k)$ also applies to Streett determinization if the output

automata are the dual of union-closed automata. In particular, no construction for Streett determinization can output Rabin automata with state size asymptotically less than $L(n, k)$.

We can get a slightly weaker result for constructions that output Rabin automata (which are not union-closed): no construction for Streett complementation can output Rabin automata with state size $n' \leq L(n, k)$ and index size $k' = O(n')$, due to the fact that a Rabin automaton with state n' and index size k' can be translated to an equivalent Büchi automaton with $O(n'k')$ states. For the same reason, no construction for Streett determinization can output Streett automata with state size $n' \leq L(n, k)$ and index size $k' = O(n')$.

Even with the fooling set and full automaton techniques and the assistance of rankings, a difficulty remains: in the setting of Streett complementation, how large can a fooling set for a complementary automaton be? The challenge is two-fold. One is to implant potentially contradictory properties in each member of a fooling set so that complementary run graphs can be obtained by certain combinations of those members. The other is to avoid correlations between members of a fooling set so that each member has to be memorized by a distinct state in a purported complementary automaton. By exploiting the nature of Streett acceptance conditions, our fooling set is obtained via a type of multi-dimensional rankings, called Q -rankings, and members in the fooling set are called Q -words. To simultaneously accommodate potentially contradictory properties in multi-dimension requires handling nontrivial subtleties. We shall continue this discussion in Section 3 after presenting the definition of Q -rankings.

1.6 Paper Organization

Section 2 presents notations and basic terminology in automata theory. Section 3 introduces full Streett automata, Q -rankings and Q -words, and use them to establish the lower bound. Section 4 concludes with a discussion. Due to space limit, technical proofs are omitted, but they can be found in the full version of this paper at arXiv:1102.2963.

2 Preliminaries

2.1 Basic Notations

Let \mathbb{N} be the set of natural numbers. We write $[i..j]$ for $\{k \in \mathbb{N} \mid i \leq k \leq j\}$, $[i..j)$ for $[i..j - 1]$, $[n]$ for $[0..n)$. For an infinite sequence ϱ , we use $\varrho(i)$ to denote the i -th component for $i \in \mathbb{N}$, $\varrho[i..j]$ (resp. $\varrho[i..j)$) to denote the subsequence of ϱ from position i to position j (resp. $j - 1$). Similar notations for finite sequences and we use $|\varrho|$ to denote the length of ϱ . We assume readers are familiar with notations in language theory, such as $\alpha \circ \alpha'$, α^* , α^+ and α^ω where α and α' are sequences and α is finite, and similar ones such as $S \circ S'$, S^* , S^+ and S^ω where S is a set of finite sequences and S' is a set of sequences.

2.2 Automata and Runs

A finite (nondeterministic) automaton on infinite words (ω -automaton) is a 5-tuple $\mathcal{A} = \langle \Sigma, S, Q, \Delta, \mathcal{F} \rangle$, where Σ is an alphabet, S is a finite set of states, $Q \subseteq S$ is a set of initial states, $\Delta \subseteq S \times \Sigma \times S$ is a transition relation, and \mathcal{F} is an acceptance condition.

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 S such that $\varrho(0) \in Q$ and $\langle \varrho(i), w(i), \varrho(i+1) \rangle \in \Delta$ for $i \in \mathbb{N}$. Finite runs are defined similarly. Let $\text{Inf}(\varrho)$ the set of states that occur infinitely many times in ϱ . An automaton accepts w if there exists a run ϱ

over w that satisfies \mathcal{F} , which usually is defined as a predicate on $\text{Inf}(\varrho)$. We use $\mathcal{L}(\mathcal{A})$ to denote the set of ω -words accepted by \mathcal{A} and $\overline{\mathcal{L}(\mathcal{A})}$ the complement of $\mathcal{L}(\mathcal{A})$.

2.3 Acceptance Conditions and Automata Types

ω -automata are classified according their acceptance conditions. Below we list three types of ω -automata relevant to this paper. Let F be a subset of Q and G, B two functions $I \rightarrow 2^Q$ where $I = [1..k]$ is called the *index set*.

- *Büchi*: $\langle F \rangle$: $\text{Inf}(\varrho) \cap F \neq \emptyset$.
- *Streett*: $\langle G, B \rangle_I$: $\forall i \in I, \text{Inf}(\varrho) \cap G(i) \neq \emptyset \rightarrow \text{Inf}(\varrho) \cap B(i) \neq \emptyset$.
- *Rabin*: $[G, B]_I$: $\exists i \in I, \text{Inf}(\varrho) \cap G(i) \neq \emptyset \wedge \text{Inf}(\varrho) \cap B(i) = \emptyset$.

Note that Streett and Rabin are dual to each other. An automaton \mathcal{A} is called *union-closed* if when two runs ϱ and ϱ' are accepting, so is any run ϱ'' if $\text{Inf}(\varrho'') = \text{Inf}(\varrho) \cup \text{Inf}(\varrho')$. It is easy to verify that both Büchi and Streett automata are union-closed while Rabin automata are not. Let $J \subseteq I$. We use $\langle G, B \rangle_J$ to denote the Streett condition with respect to only indices in J . When J is a singleton, say $J = \{j\}$, we simply write $\langle G(j), B(j) \rangle$ for $\langle G, B \rangle_J$. We can assume that B is injective and the index size k is bound by 2^n , because if $B(i) = B(i')$ for two different $i, i' \in I$, then we can shrink the index set I by replacing $\langle G, B \rangle_{\{i, i'\}}$ by $\langle G(i) \cup G(i'), B(i) \rangle$. The same convention and assumption are used for Rabin condition.

2.4 Δ -Graphs

A Δ -graph (run graph) of an ω -word w under \mathcal{A} is a directed graph $\mathcal{G}_w = (V, E)$ where $V = S \times \mathbb{N}$ and $E = \{ \langle \langle s, l \rangle, \langle s', l+1 \rangle \rangle \in V \times V \mid s, s' \in S, l \in \mathbb{N}, \langle s, w(l), s' \rangle \in \Delta \}$. By the l -th level, we mean the vertex set $S \times \{l\}$. Let $S = \{s_0, \dots, s_{n-1}\}$. By s_l -track we mean the vertex set $\{s_l\} \times \mathbb{N}$. For a subset X of S , we call a vertex $\langle s, l \rangle$ an X -vertex if $s \in X$. We simply use s for $\langle s, l \rangle$ when the index is irrelevant.

A Δ -graph \mathcal{G}_w of a finite word w is defined similarly. By $|\mathcal{G}_w|$ we denote the length of \mathcal{G}_w , which is the same as $|w|$. \mathcal{G}_σ for $\sigma \in \Sigma$ is called a *unit* Δ -graph. A path in \mathcal{G}_w is called 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 $\mathcal{G}_w \circ \mathcal{G}_{w'} = \mathcal{G}_{w \circ w'}$.

Let w be a finite word. For $l, l' \in \mathbb{N}$, $s, s' \in S$ we write $\langle s, l \rangle \xrightarrow{w} \langle s', l' \rangle$ to mean that there exists a run ϱ of \mathcal{A} such that $\varrho[l..l']$, the subsequence $\varrho(l)\varrho(l+1)\cdots\varrho(l')$ of ϱ , is a finite run of \mathcal{A} from s to s' over w . We simply write $s \xrightarrow{w} s'$, when omitting level indices causes no confusion.

2.5 Full Automata

A full automaton $\langle \Sigma, S, Q, \Delta, \mathcal{F} \rangle$ is a finite automaton with the following conditions: $\Sigma = 2^{S \times S}$, $\Delta \subseteq S \times 2^{S \times S} \times S$, and for all $s, s' \in S$, $\sigma \in \Sigma$, $\langle s, \sigma, s' \rangle \in \Delta$ if and only if $\langle s, s' \rangle \in \sigma$ [21, 27, 3]. For full automata, the alphabet Σ and the transition relation Δ are completely determined by S . As stated in the introduction, the essence of full automaton technique is to use run graphs as free as possible, without worrying which word generates which run graph. Let the functional version of Δ be $\delta : \Sigma \rightarrow 2^{S \times S}$, where for every $s, s' \in S$ and every $\sigma \in \Sigma$, $\langle s, s' \rangle \in \delta(\sigma)$ if and only if $\langle s, \sigma, s' \rangle \in \Delta$. The function δ maps a letter σ to a unit Δ -graph \mathcal{G}_σ , which represents the complete behavior of \mathcal{A} over σ (technically speaking, \mathcal{G}_σ , with index dropped, is the graph of $\delta(\sigma)$). In the setting of full automata, δ is

simply the identity function on $2^{S \times S}$. Words and run graphs are essentially the same thing. From now on we use the two terms interchangeably. For example, for a word w , $s \xrightarrow{w} s'$ is equivalent to say that a full path in \mathcal{G}_w goes from s to s' .

3 Lower Bound

In this section we define full Streett automata, and related Q -rankings and Q -words, and use them to establish the lower bound. From now on, we reserve n and k , respectively, for the effective state size and index size in our construction (except in Theorem 9 and Section 4 where n and k , respectively, mean the state size and index size of a complementation instance). All related notions are in fact parameterized with n and k , but we do not list them explicitly unless required for clarity. Let I be $[1..k]$. We first describe the plan of proof.

For each $k, n > 0$, we define a full Streett automaton $\mathcal{S} = (\Sigma, S, Q, \Delta, \mathcal{F})$ and a set of Q -rankings $f : Q \rightarrow [1..n] \times I^k$. For each Q -ranking f , we define a finite Δ -graph \mathcal{G}_f , called a Q -word. We then show that for each f , $(\mathcal{G}_f)^\omega \notin \mathcal{L}(\mathcal{S})$, yet $((\mathcal{G}_f)^+(\mathcal{G}_{f'})^+)^\omega \subseteq \mathcal{L}(\mathcal{S})$ for every distinct pair of Q -rankings f and f' , that is, Q -words constitute a fooling set for $\mathcal{L}(\mathcal{S})$. Using Michel's scheme [16, 15, 27], we show that if a union-closed automaton \mathcal{C} complements \mathcal{S} , then its state size is no less than the number of Q -rankings, because otherwise we can “weave” the runs of $(\mathcal{G}_f)^\omega$ and $(\mathcal{G}_{f'})^\omega$ in such a way that \mathcal{C} would accept a word in $((\mathcal{G}_f)^+(\mathcal{G}_{f'})^+)^\omega$, contradicting $((\mathcal{G}_f)^+(\mathcal{G}_{f'})^+)^\omega \subseteq \mathcal{L}(\mathcal{S})$.

► **Definition 1 (Full Streett Automata).** Let $\{\mathcal{S} = \langle \Sigma, S, Q, \Delta, \mathcal{F} \rangle\}_{n,k>0}$ be a family of full Streett automata such that

- 1.1 $S = Q \cup P_G \cup P_B \cup T$ where Q, P_G, P_B and T are pairwise disjoint sets of the following forms: $Q = \{q_0, \dots, q_{n-1}\}$, $P_G = \{g_1, \dots, g_k\}$, $T = \{t\}$, and $P_B = \{b_1, \dots, b_k\}$.
- 1.2 $\mathcal{F} = \langle G, B \rangle_I$ such that $G(i) = \{g_i\}$ and $B(i) = \{b_i\}$ for $i \in I$.

Q is intended to be the domain of Q -rankings. P_G and P_B are pools from which singletons $G(i)$'s and $B(i)$'s are formed. T is to be used for building a *bypass* track that makes graph concatenation behaves like a parallel composition so that properties associated with each subgraph are all preserved in the final concatenation.

► **Definition 2 (Q -Ranking).** A Q -ranking for \mathcal{S} is a function $f : Q \rightarrow [1..n] \times I^k$, which is identified with a pair of functions $\langle r, h \rangle$, where $r : Q \rightarrow [1..n]$ is one-to-one, and $h : Q \rightarrow I^k$ maps a state to a permutation of I .

For a Q -ranking $f = \langle r, h \rangle$, we call r (resp. h) the R -ranking or numeric ranking (resp. H -ranking or index ranking) of f . We use Q -ranks (resp. R -ranks, H -ranks) to mean values of Q -rankings (resp. R -rankings, H -rankings). For $q \in Q$, we write $h(q)[i]$ ($i \in I$) to denote the i -th component of $h(q)$. Let \mathcal{D}^Q be the set of all Q -rankings and $|\mathcal{D}^Q|$ be the size of \mathcal{D}^Q . Clearly, we have $n!$ R -rankings and $(k!)^n$ H -rankings, and so $|\mathcal{D}^Q| = (n!)(k!)^n = 2^{\Omega(n \lg n + nk \lg k)}$.

As stated in the introduction, Q -rankings are essential for obtaining the lower bound. It turns out that H -rankings are the core of Q -rankings, for $(k!)^n$ already begins to dominate $n!$ when k is larger than $\lg n$. Now we explain the idea behind the design of H -rankings. Recall that our goal is to have $(\mathcal{G}_f)^\omega \notin \mathcal{L}(\mathcal{S})$ for any Q -ranking f as well as $((\mathcal{G}_f)^+(\mathcal{G}_{f'})^+)^\omega \subseteq \mathcal{L}(\mathcal{S})$ for any two different Q -rankings f and f' . For simplicity, we ignore R -rankings and assume Q -rankings are just H -rankings. We say that a finite path *discharges* obligation j if the path visits $B(j)$ and a finite path *owes* obligation j if the path visits $G(j)$ but does not visit $B(j)$. As shown below, for each $i \in [n]$, q_i -track in \mathcal{G}_f is associated with the k -tuple $f(q_i)$, which is a permutation of I , and exactly k full paths in \mathcal{G}_f goes from the beginning of q_i -track to

the end of q_i -track. We say that those paths *on* q_i -track. For each $i \in [n]$ and $j \in I$, the j -th full path on q_i -track owes exactly the obligation $f(q_i)[j]$. Let $\varrho = \varrho_0 \circ \varrho_1 \circ \dots$ be an infinite path in $(\mathcal{G}_f)^\omega$ where ϱ_t ($t \geq 0$) is a full path in the t -th \mathcal{G}_f . Without R -rankings, our construction prescribes that all ϱ_t start and end at a specific track, say q_i -track, and hence are associated with $f(q_i)$. Obligations associated with all ϱ_t simply form a subset I' of I . However, we impose an ordering $\prec_{f,i}$ on I' (different from the standard numeric ordering) such that $f(q_i)[j] \prec_{f,i} f(q_i)[j']$ if and only if $j < j'$. The ordering $\prec_{f,i}$ is total thanks to $f(q_i)$ being a permutation of I . Then a condition in our construction guarantees that the minimum obligation with respect to $\prec_{f,i}$ will never be discharged on ϱ , and therefore ϱ violates $\langle G, B \rangle_I$. Since this ϱ is chosen arbitrarily, we have $(\mathcal{G}_f)^\omega \notin \mathcal{L}(\mathcal{S})$.

Now let $\mathcal{G} \in ((\mathcal{G}_f)^+(\mathcal{G}_{f'})^+)^\omega$. To show $\mathcal{G} \in \mathcal{L}(\mathcal{S})$, we construct an infinite path $\varrho = \varrho_0 \circ \varrho_1 \circ \dots$ in \mathcal{G} that satisfies $\langle G, B \rangle_I$, where ϱ_t ($t \geq 0$) is a full path in the t -th subgraph (which is either \mathcal{G}_f or $\mathcal{G}_{f'}$). Let i be such that $f(q_i) \neq f'(q_i)$ (it is always possible by the assumption $f \neq f'$). Different from before, q_i -track in \mathcal{G}_f is associated with $f(q_i)$ and q_i -track in $\mathcal{G}_{f'}$ is associated with $f'(q_i)$. Since $f(q_i)$ and $f'(q_i)$ are different permutations of I , a condition in our construction ensures that a full path ϱ_f in \mathcal{G}_f and a full path $\varrho_{f'}$ in $\mathcal{G}_{f'}$, both on q_i -track, mutually discharge each other's obligations. So we let all ϱ_t in \mathcal{G}_f be ϱ_f and all ϱ_t in $\mathcal{G}_{f'}$ be $\varrho_{f'}$. Since there are infinitely many ϱ_f and $\varrho_{f'}$ in ϱ , ϱ satisfies $\langle G, B \rangle_I$, giving us $\mathcal{G} \in \mathcal{L}(\mathcal{S})$. Since \mathcal{G} is chosen arbitrarily, we have $((\mathcal{G}_f)^+(\mathcal{G}_{f'})^+)^\omega \subseteq \mathcal{L}(\mathcal{S})$. Now we are read to formally define Q -words.

► **Definition 3** (Q -Word). A finite Δ -graph \mathcal{G} is called a Q -word if every level of \mathcal{G} is ranked by the same Q -ranking $f = \langle r, h \rangle$ and \mathcal{G} satisfies the following additional conditions.

- 3.1 For every $q, q' \in Q$, if $r(q) > r(q')$, there exists a full path ϱ from $\langle q, 0 \rangle$ to $\langle q', |\mathcal{G}| \rangle$ such that ϱ visits all of $B(1), \dots, B(k)$.
- 3.2 For every $q \in Q$, there exist *exactly* k full paths $\varrho_1, \dots, \varrho_k$ from $\langle q, 0 \rangle$ to $\langle q, |\mathcal{G}| \rangle$ such that for every $i \in I$, ϱ_i does not visit $B(h(q)[j])$ for $j \leq i$, but visits $B(h(q)[j])$ for $i < j$, and ϱ_i does not visit $G(h(q)[j])$ for $j < i$, but visits $G(h(q)[i])$.
- 3.3 Only Q -vertices have outgoing edges at the first level and incoming edges at the last level.
- 3.4 For every $q, q' \in Q$, there exists no full path from $\langle q, 0 \rangle$ to $\langle q', |\mathcal{G}| \rangle$ if $r(q) < r(q')$.

Property (3.1) concerns with only R -rankings. It says that for every two tracks with different R -ranks, a path exists that goes from the track with higher rank to the track with the lower rank, and such a path discharges all obligations in I . So if those (finite) paths occur infinitely often as fragments of an infinite path ϱ , then ϱ clearly satisfies the Streett condition $\langle G, B \rangle_I$. Property (3.2) concerns with only H -rankings. It says that exactly k full “parallel” paths exist between the two ends of every track, and each owes exactly one distinct obligation in I . As shown in Theorem 9, Property (3.2) is the core of the whole construction and proof, because with k increasing, H -rankings contribute more and more to the overall complexity. Properties (3.3) and (3.4) are merely technical; they ensure that no other *full paths* exist besides those prescribed by Properties (3.1) and (3.2). Note that in general more than one Q -word could exist for a Q -ranking f . We simply pick an arbitrary one and call it the Q -word of f , denoted by \mathcal{G}_f .

► **Theorem 4** (Q -Word). A Q -word exists for every Q -ranking.

► **Example 5** (Q -Word). Let us consider a full Streett automaton \mathcal{S} where $n = 3, k = 2, Q = \{q_0, q_1, q_2\}, T = \{t\}, P_B = \{b_1, b_2\}, P_G = \{g_1, g_2\}$, and the following Q -ranking $f = \langle r, h \rangle$: $r(q_0) = 2, r(q_1) = 1, r(q_2) = 3, h(q_0) = \langle 1, 2 \rangle, h(q_1) = \langle 1, 2 \rangle, h(q_2) = \langle 2, 1 \rangle$. Figure 1 shows a Q -word \mathcal{G}_f , which consists of two subgraphs \mathcal{G}_r and \mathcal{G}_h , where \mathcal{G}_r in turn

consists of two parts: $\mathcal{G}_r^{(1)}$ (level 0 to level 3) and $\mathcal{G}_r^{(2)}$ (level 3 to 6), and \mathcal{G}_h in turn consists of three parts: $\mathcal{G}_h^{(0)}$ (level 6 to level 12), $\mathcal{G}_h^{(1)}$ (level 12 to level 18), and $\mathcal{G}_h^{(2)}$ (level 18 to level 24). \mathcal{G}_r and \mathcal{G}_h are aimed to satisfy Properties (3.1) and (3.2), respectively.

The R -rank (numeric rank) of every level of \mathcal{G}_r is $(2, 1, 3)$. In $\mathcal{G}_r^{(1)}$, a full path ϱ_r starts from $\langle q_2, 0 \rangle$ whose R -rank is the highest. The path visits $\langle b_1, 1 \rangle$, $\langle b_2, 2 \rangle$ and then $\langle q_0, 3 \rangle$ whose R -rank is one less than that of q_2 . Similarly in $\mathcal{G}_r^{(2)}$, the path continues from $\langle q_2, 3 \rangle$, visits $\langle b_1, 4 \rangle$, $\langle b_2, 5 \rangle$ and ends at $\langle q_1, 6 \rangle$ whose R -rank is one less than that of q_0 .

The H -rank (index rank) of every level of \mathcal{G}_h is $(\langle 1, 2 \rangle, \langle 1, 2 \rangle, \langle 2, 1 \rangle)$. Let us take a look at $\mathcal{G}_h^{(1)}$. A full path ϱ_h (marked green except the last edge) starts at $\langle q_1, 12 \rangle$, visits $\langle b_2, 13 \rangle$ and $\langle g_1, 14 \rangle$ (because of $h(q_1)[1] = 1$), and enters t -track (the bypass track $\{t\} \times \mathbb{N}$) at $\langle t, 15 \rangle$, from where it stays on t -track till reaching $\langle t, 17 \rangle$. Another full path ϱ'_h (marked red except the last edge) starts at $\langle q_1, 12 \rangle$ too, takes q_1 -track to $\langle q_1, 15 \rangle$, and then visits $\langle g_2, 16 \rangle$ (because of $h(q_1)[2] = 2$), and enters t -track at $\langle t, 17 \rangle$. Both ϱ_h and ϱ'_h return to q_1 -track at $\langle q_1, 18 \rangle$ using the edge $\langle \langle t, 17 \rangle, \langle q_1, 18 \rangle \rangle$ (marked blue). By $\varrho_{0 \rightarrow 6}$, $\varrho_{6 \rightarrow 12}$ and $\varrho_{18 \rightarrow 24}$ (all marked blue) we denote the q_1 -tracks in \mathcal{G}_r , in $\mathcal{G}_h^{(0)}$ and in $\mathcal{G}_h^{(2)}$, respectively. It is easy to verify that Property (3.1) with respect to q_2 and q_1 is satisfied by both $\varrho_r \circ \varrho_{6 \rightarrow 12} \circ \varrho_h \circ \varrho_{18 \rightarrow 24}$ and $\varrho_r \circ \varrho_{6 \rightarrow 12} \circ \varrho_{h'} \circ \varrho_{18 \rightarrow 24}$. Also easily seen is that Property (3.2) with respect to q_1 is satisfied by $\varrho_{0 \rightarrow 6} \circ \varrho_{6 \rightarrow 12} \circ \varrho_h \circ \varrho_{18 \rightarrow 24}$ and $\varrho_{0 \rightarrow 6} \circ \varrho_{6 \rightarrow 12} \circ \varrho_{h'} \circ \varrho_{18 \rightarrow 24}$.

We are ready for the lower bound proof. Let $J \subseteq I$. We use $\langle G, B \rangle_J$ to denote the Streett condition with respect to only indices in J . The corresponding Rabin condition $[G, B]_J$ is similarly defined. When J is a singleton, say $J = \{j\}$, we simply write $\langle G(j), B(j) \rangle$ for $\langle G, B \rangle_J$ and $[G(j), B(j)]$ for $[G, B]_J$. Obviously, if an infinite run satisfies $\langle G, B \rangle_J$ (resp. $[G, B]_J$), then the run also satisfies $\langle G, B \rangle_{J'}$ (resp. $[G, B]_{J'}$) for $J' \subseteq J$ (resp. $J \subseteq J' \subseteq I$).

► **Lemma 6.** *For every Q -ranking f , $(\mathcal{G}_f)^\omega \notin \mathcal{L}(\mathcal{S})$.*

Proof. Let $f = \langle r, h \rangle$, $\mathcal{G} = (\mathcal{G}_f)^\omega$ and ϱ an infinite path in \mathcal{G} . For simplicity, we assume ϱ only lists states appearing on the boundaries of \mathcal{G}_f fragments; for any $j \geq 0$, $\varrho(j)$ (resp. $\varrho(j+1)$) is a state in the first (resp. last) level of the j -th \mathcal{G}_f fragment. Let $\varrho[j, j+1]$ denote the finite fragment from $\varrho(j)$ to $\varrho(j+1)$. Let $\varrho[j, \infty]$ denote the suffix of ϱ beginning from $\varrho(j)$.

By Property (3.3), $\varrho(i) \in Q$ for $i \geq 0$. By Property (3.4), ϱ eventually stabilizes on R -ranks in the sense that there exists a j_0 such that for any $j \geq j_0$, $r(\varrho(j)) = r(\varrho(j+1))$. Because every level of \mathcal{G} has the same rank, ϱ stabilizes on a (horizontal) track after j_0 , i.e., there exists $i \in [n]$ such that $\varrho(j) = q_i$ for $j \geq j_0$. Property (3.2) says that there are *exactly* k full paths $\varrho_1, \dots, \varrho_k$ from $\langle q_i, 0 \rangle$ to $\langle q_i, |\mathcal{G}_f| \rangle$ in \mathcal{G}_f . Therefore, $\varrho[j_0, \infty]$ can be divided into the infinite sequence $\varrho[j_0, j_0+1], \varrho[j_0+1, j_0+2], \dots$, each of which is one of $\varrho_1, \dots, \varrho_k$. Let $k_0 \in I$ be the smallest index such that ϱ_{k_0} appears infinitely often in this sequence, i.e., for some $j_1 \geq j_0$, none of $\varrho_1, \dots, \varrho_{k_0-1}$ appears in $\varrho[j_1, \infty]$. By Property (3.2) again, $\varrho[j_1, \infty]$ visits none of $B(h(q_i)[1]), \dots, B(h(q_i)[k_0])$, but visits $G(h(q_i)[k_0])$ infinitely often (because ϱ_{k_0} appears infinitely often). In particular, ϱ satisfies $[G(t), B(t)]$ for $t = h(q_i)[k_0]$ and hence $[G, B]_I$. Because ϱ is chosen arbitrarily, we have $\mathcal{G} \notin \mathcal{L}(\mathcal{S})$. ◀

► **Lemma 7.** *For every two different Q -rankings f and f' , $((\mathcal{G}_f)^+ \circ (\mathcal{G}_{f'})^+)^\omega \subseteq \mathcal{L}(\mathcal{S})$.*

Proof. Let $\mathcal{G} \in ((\mathcal{G}_f)^+ \circ (\mathcal{G}_{f'})^+)^\omega$ be an ω -word where both \mathcal{G}_f and $\mathcal{G}_{f'}$ occur infinitely often in \mathcal{G} . Let $f = \langle r, h \rangle$ and $f' = \langle r', h' \rangle$. We have two cases: either $r \neq r'$ or $h \neq h'$.

If $r \neq r'$. Since both r and r' are one-to-one functions from Q to $[1..n]$, there must be $i, j \in [n]$ such that $r(q_i) > r(q_j)$ and $r'(q_j) > r'(q_i)$. By Property (3.1), \mathcal{G}_f contains a full path $\varrho_{i \rightarrow j}$ from $\langle q_i, 0 \rangle$ to $\langle q_j, |\mathcal{G}_f| \rangle$ that visits all of $B(1), \dots, B(k)$. By the same property,

$\mathcal{G}_{f'}$ contains a path $\varrho'_{j \rightarrow i}$ from $\langle q_j, 0 \rangle$ to $\langle q_i, |\mathcal{G}_{f'}| \rangle$ that also visits all of $B(1), \dots, B(k)$. Then $\varrho_{i \rightarrow j} \circ \varrho'_{j \rightarrow i}$ is a path in $\mathcal{G}_f \circ \mathcal{G}_{f'}$ that visits all of $B(1), \dots, B(k)$. Also by Property (3.2), \mathcal{G}_f (resp. $\mathcal{G}_{f'}$) contains a path $\varrho_{i \rightarrow i}$ (resp. $\varrho'_{i \rightarrow i}$) from $\langle q_i, 0 \rangle$ to $\langle q_i, |\mathcal{G}_f| \rangle$ (resp. from $\langle q_i, 0 \rangle$ to $\langle q_i, |\mathcal{G}_{f'}| \rangle$).

Now we define an infinite path $\hat{\varrho}$ in \mathcal{G} as follows. We pick the finite path $\varrho_{i \rightarrow i}$ in every \mathcal{G}_f fragment and $\varrho'_{i \rightarrow i}$ in every $\mathcal{G}_{f'}$ fragment, except that in the case where a \mathcal{G}_f fragment is followed immediately by a $\mathcal{G}_{f'}$ fragment, we pick $\varrho_{i \rightarrow j}$ in the preceding \mathcal{G}_f and $\varrho'_{j \rightarrow i}$ in the following $\mathcal{G}_{f'}$. It is easily seen that $\hat{\varrho}$, in the form $((\varrho_{i \rightarrow i})^* \circ (\varrho_{i \rightarrow j} \circ \varrho'_{j \rightarrow i})^+ \circ (\varrho'_{i \rightarrow i})^*)^\omega$, visits all of $B(1), \dots, B(k)$ infinitely often, and hence it satisfies the Streett condition $\langle G, B \rangle_I$.

If $h \neq h'$. Then there exist $i \in [n], j \in I$ such that $h(q_i)[j] \neq h'(q_i)[j]$ and $h(q_i)[j^*] = h'(q_i)[j^*]$ for $j^* \in [1..j-1]$. Since both $h(q_i)$ and $h'(q_i)$ are permutations of I , we have $j < k$ and $\{h(q_i)[j^*] \mid j^* \in [j..k]\} = \{h'(q_i)[j^*] \mid j^* \in [j..k]\}$. By Property (3.2), in \mathcal{G}_f there exists a path $\varrho_{i \rightarrow i}$ from $\langle q_i, 0 \rangle$ to $\langle q_i, |\mathcal{G}_f| \rangle$ that visits none of $G(h(q_i)[j^*])$ for $j^* \in [1..j-1]$, but visits all of $B(h(q_i)[j^*])$ for $j^* \in [j+1..k]$. Similarly, in $\mathcal{G}_{f'}$ there exists a path $\varrho'_{i \rightarrow i}$ from $\langle q_i, 0 \rangle$ to $\langle q_i, |\mathcal{G}_{f'}| \rangle$ that visits none of $G(h'(q_i)[j^*])$ for $j^* \in [1..j-1]$, but visits all of $B(h'(q_i)[j^*])$ for $j^* \in [j+1..k]$. Because $h(q_i)$ and $h'(q_i)$ are different permutations of I , $h'(q_i)[j] = h(q_i)[j_0]$ for some $j_0 \in [j+1..k]$ and $h(q_i)[j] = h'(q_i)[j_1]$ for some $j_1 \in [j+1..k]$. It follows that both $\{h(q_i)[j^*] \mid j^* \in [j..k]\}$ and $\{h'(q_i)[j^*] \mid j^* \in [j..k]\}$ are equal to $\{h(q_i)[j^*] \mid j^* \in [j+1..k]\} \cup \{h'(q_i)[j^*] \mid j^* \in [j+1..k]\}$. Therefore $\varrho_{i \rightarrow i} \circ \varrho'_{i \rightarrow i}$ (in $\mathcal{G}_f \circ \mathcal{G}_{f'}$) visits all of $B(h(q_i)[j^*])$ for $j^* \in [j..k]$.

Now let $\hat{\varrho}$ be defined as follows: $\hat{\varrho}$ takes $\varrho_{i \rightarrow i}$ in every \mathcal{G}_f fragment and $\varrho'_{i \rightarrow i}$ in every $\mathcal{G}_{f'}$ fragment. That is, $\hat{\varrho}$ takes the following form $((\varrho_{i \rightarrow i})^+ \circ (\varrho'_{i \rightarrow i})^+)^{\omega}$. Recall that $h(q_i)[j^*] = h'(q_i)[j^*]$ for $j^* \in [1..j-1]$. It follows that $\hat{\varrho}$ does not visit any of $G(h(q_i)[j^*])$ for $j^* \in [1..j-1]$ because neither $\varrho_{i \rightarrow i}$ nor $\varrho'_{i \rightarrow i}$ does. Also since both \mathcal{G}_f and $\mathcal{G}_{f'}$ occur infinitely often in \mathcal{G} , $\hat{\varrho}$ contains infinitely many $\varrho_{i \rightarrow i} \circ \varrho'_{i \rightarrow i}$, which implies that $\hat{\varrho}$ visits all of $B(h(q_i)[j^*])$ for $j^* \in [j..k]$ infinitely often. Since $h(q_i)$ is a permutation of I , $\hat{\varrho}$ satisfies $\langle G, B \rangle_I$.

In either case (whether $r \neq r'$ or $h \neq h'$), \mathcal{G} contains a path that satisfies $\langle G, B \rangle_I$, which means $\mathcal{G} \in \mathcal{L}(\mathcal{S})$. Because \mathcal{G} is arbitrarily chosen, we have $((\mathcal{G}_f)^+ \circ (\mathcal{G}_{f'})^+)^{\omega} \subseteq \mathcal{L}(\mathcal{S})$. ◀

The following lemma is the core of Michel's scheme [16, 15], recast in the setting of full automata with rankings [27, 3]. Recall that \mathcal{D}^Q denotes the set of all Q -rankings and $|\mathcal{D}^Q|$ denotes the cardinality of \mathcal{D}^Q .

► **Lemma 8.** *A union-closed automaton that complements \mathcal{S} must have at least $|\mathcal{D}^Q|$ states.*

Proof. Let \mathcal{C} be a union-closed automaton that complements \mathcal{S} . By Lemma 6, for every Q -ranking f , $(\mathcal{G}_f)^\omega \in \mathcal{L}(\mathcal{C})$. Let f, f' be two different Q -rankings and \mathcal{G}_f and $\mathcal{G}_{f'}$ the corresponding Q -words. Let ϱ and ϱ' be the corresponding accepting runs of $(\mathcal{G}_f)^\omega$ and $(\mathcal{G}_{f'})^\omega$, respectively. Also let ϱ_0 and ϱ'_0 , respectively, be the accepting runs of $(\mathcal{G}_f)^\omega$ and $(\mathcal{G}_{f'})^\omega$ when we treat \mathcal{G}_f and $\mathcal{G}_{f'}$ as *atomic* letters, that is, ϱ_0 (resp. ϱ'_0) only records states visited at the boundary of \mathcal{G}_f (resp. $\mathcal{G}_{f'}$) and is a subsequence of ϱ (resp. ϱ'). Obviously, $\text{Inf}(\varrho_0) \subseteq \text{Inf}(\varrho)$, $\text{Inf}(\varrho'_0) \subseteq \text{Inf}(\varrho')$, $\text{Inf}(\varrho_0) \neq \emptyset$ and $\text{Inf}(\varrho'_0) \neq \emptyset$. If $\text{Inf}(\varrho_0) \cap \text{Inf}(\varrho'_0) = \emptyset$ for any pair of f and f' , then clearly \mathcal{C} has at least $|\mathcal{D}^Q|$ states because the state set of \mathcal{C} contains $|\mathcal{D}^Q|$ pairwise disjoint nonempty subsets.

Therefore we can assume that $\text{Inf}(\varrho_0) \cap \text{Inf}(\varrho'_0) \neq \emptyset$ for a fixed pair of f and f' . Let q be a state in $\text{Inf}(\varrho_0) \cap \text{Inf}(\varrho'_0)$. Because q occurs infinitely often in ϱ , then for some $m > 0$, there exists a path in $(\mathcal{G}_f)^m$ that goes from q to q and visits exactly all states in $\text{Inf}(\varrho)$ (or equivalently speaking, \mathcal{C} , upon reading the input word $(\mathcal{G}_f)^m$, runs from state q to q , visiting exactly all states in $\text{Inf}(\varrho)$ during the run). By $q \xrightarrow[\text{Inf}(\varrho)]{(\mathcal{G}_f)^m} q$ we denote the existence of such a

path. Similarly, we have $q \xrightarrow[\text{!Inf}(\varrho')]{(\mathcal{G}_{f'})^{m'}} q$ for some $m' > 0$. Also we have $q_0 \xrightarrow{(\mathcal{G}_f)^{m_0}} q$ where q_0 is an initial state of \mathcal{C} . Now consider the following infinite run ϱ^* in the form

$$q_0 \xrightarrow{(\mathcal{G}_f)^{m_0}} q \xrightarrow[\text{!Inf}(\varrho)]{(\mathcal{G}_f)^m} q \xrightarrow[\text{!Inf}(\varrho')]{(\mathcal{G}_{f'})^{m'}} q \xrightarrow[\text{!Inf}(\varrho)]{(\mathcal{G}_f)^m} q \xrightarrow[\text{!Inf}(\varrho')]{(\mathcal{G}_{f'})^{m'}} q \xrightarrow[\text{!Inf}(\varrho)]{(\mathcal{G}_f)^m} q \xrightarrow[\text{!Inf}(\varrho')]{(\mathcal{G}_{f'})^{m'}} q \xrightarrow[\text{!Inf}(\varrho)]{(\mathcal{G}_f)^m} q \cdots$$

which is an accepting run of \mathcal{C} for $(\mathcal{G}_f)^{m_0} \circ ((\mathcal{G}_f)^m \circ (\mathcal{G}_{f'})^{m'})^\omega$ because $\text{Inf}(\varrho^*) = \text{Inf}(\varrho) \cup \text{Inf}(\varrho')$. However, by Lemma 7, $(\mathcal{G}_f)^{m_0} \circ ((\mathcal{G}_f)^m \circ (\mathcal{G}_{f'})^{m'})^\omega \in ((\mathcal{G}_f)^+ \circ (\mathcal{G}_{f'})^+)^\omega \subseteq \mathcal{L}(\mathcal{S})$, a contradiction. \blacktriangleleft

► **Theorem 9.** *Streett complementation is in $2^{\Omega(n \lg n + kn \lg k)}$ for $k = O(n)$ and in $2^{\Omega(n^2 \lg n)}$ for $k = \omega(n)$, where n and k are the state size and index size of a complementation instance.*


Proof. Here we switch to use n_0 and k_0 , respectively, for the effective state size and index size in our construction \mathcal{S} . We have $n = 2k_0 + n_0 + 1$. By Lemma 8, the complementation of \mathcal{S} requires $|\mathcal{D}^Q| = 2^{\Omega(n_0 \lg n_0 + n_0 k_0 \lg k_0)}$ states. If $k_0 \leq k$, we can construct a full Streett automaton \mathcal{S}' with state size n and index size k as follows. \mathcal{S}' is almost identical to \mathcal{S} except that its acceptance condition is defined as $\mathcal{F}' = \langle G', B' \rangle_{I'}$ (for $I' = [1..k]$) such that for $i \in [1..k_0]$, $G'(i) = G(i)$ and $B'(i) = B(i)$ and for $i \in [k_0 + 1, k]$, $G'(i) = B'(i) = \emptyset$. It is easily seen that \mathcal{S}' is equivalent to \mathcal{S} and hence the complementation lower bound for \mathcal{S} also applies to that for \mathcal{S}' . Now when $k = O(n)$, we can always find n_0 and k_0 such that $k_0 \leq k$, yet $n_0 = \Omega(n)$ and $k_0 = \Omega(k)$, and hence we have the lower bound $2^{\Omega(n \lg n + kn \lg k)}$. When $k = \omega(n)$, we set $k_0 = n_0$ so that $k_0 \leq k$, $n_0 = \Omega(n)$ and $k_0 = \Omega(n)$, and hence we have the lower bound $2^{\Omega(n^2 \lg n)}$. \blacktriangleleft

4 Concluding Remarks

In this paper we proved a tight lower bound $L(n, k)$ for Streett complementation. We note that we can improve the lower bound by two modifications. First, we allow $G(i)$ (resp. $B(i)$) to be arbitrary subsets of P_G (resp. P_B). Second, we also use multi-dimensional R -rankings; the range of r is a set of k -tuples of integers in $[1..n]$. As a result, both R -ranks and H -ranks are k -tuples of integers where k can be as large as 2^n (the current effective k is bounded by n). These two modifications require much more sophisticated definition of Q -rankings and construction of Q -words, but they have no asymptotic effect on $L(n, k)$. The situation is different from Rabin complementation [3], where Q -rankings are also multi-dimensional (though different terms other than Q -rankings and Q -words were used), and each component in a k -tuple (the value of a Q -ranking) is independent from one another, and hence each can impose an independent behavior on Q -words. Put it in another way, no matter how large the index set is (the maximum size can be 2^n), all dual properties, each of which is parameterized with an index, can be realized in one Q -word. For Streett complementation, the diminishing gain when pushing up k made us realize that with increasing number of Q -rankings, more and more correlations occur between Q -rankings. Exploiting these correlations leads us to the discovery of the corresponding upper bound.

Acknowledgment

We would like to thank anonymous reviewers for many useful comments, and we are grateful to Laurel Tweed and Wanwu Wang for carefully proofreading the paper.



■ Figure 1 Q-word \mathcal{G}_f ($f = (r, h)$)

References

- 1 J.R. Büchi. Weak Second-order Arithmetic and Finite Automata. *Mathematical Logic Quarterly*, 6(1-6): 66-92, 1960.
- 2 J.R. Büchi. On a decision method in restricted second order arithmetic. In *Proc. 1960 Internat. Congr. Logic, Method. and Philos. Sci.*, pp. 1-11, 1966.
- 3 Y. Cai, T. Zhang, and H. Luo. An improved lower bound for the complementation of Rabin automata. In *Proc. 24th LICS*, pp. 167-176, 2009.
- 4 Y. Cai and T. Zhang. Tight upper bounds for Streett and parity complementation. In *Proc. 20th CSL*, pp. 112-128, 2011.
- 5 N. Francez and D. Kozen. Generalized fair termination. In *Proc. 11th POPL*, pp. 46-53, 1984.
- 6 E. Friedgut and O. Kupferman and M.Y. Vardi. Büchi complementation made tighter. *International Journal of Foundations of Computer Science*, 17(4): 851-867, 2006.
- 7 N. Francez. Fairness. *Texts and Monographs in Computer Science*. Springer-Verlag, 1986.
- 8 Ian Glaister and Jeffrey Shallit. A lower bound technique for the size of nondeterministic finite automata. *Information Processing Letters*. 59(2): 75-77, 1996.
- 9 N. Klarlund. Progress measures for complementation of omega-automata with applications to temporal logic. In *Proc. 32th FOCS*, pp. 358-367, 1991.
- 10 O. Kupferman. Avoiding Determinization. In *Proc. 21th LICS*, pp. 243-254, 2006.
- 11 R.P. Kurshan. Computer aided verification of coordinating processes: an automata theoretic approach. Princeton University Press, 1994.
- 12 O. Kupferman and M.Y. Vardi. Weak alternating automata are not that weak. *ACM Transactions on Computational Logic*, 2(3): 408-429, 2001.
- 13 O. Kupferman and M.Y. Vardi. From complementation to certification. In *10th TACAS, LNCS 2988*, pp. 591-606, 2004.
- 14 O. Kupferman and M.Y. Vardi. Complementation constructions for nondeterministic automata on infinite words. In *Proc. 11th TACAS*, pp. 206-221, 2005.
- 15 C. Löding. Optimal bounds for transformations of omega-automata. In *Proc. 19th FSTTCS, LNCS 1738*, pp. 97-109, 1999.
- 16 M. Michel. Complementation is more difficult with automata on infinite words. *CNET*, Paris, 1988.
- 17 N. Piterman. From Nondeterministic Büchi and Streett Automata to Deterministic Parity Automata. In *Proc. 21th LICS*, pp. 255-264, 2006.
- 18 S. Safra. On the complexity of ω -automata. In *Proc. 29th FOCS*, pp. 319-327, 1988.
- 19 S. Safra. Exponential Determinization for ω -Automata with Strong-Fairness Acceptance Condition. In *Proc. 24th STOC*, pp. 275-327, 1992.
- 20 S. Schewe. Büchi complementation made tight. In *Proc. 26th STACS*, pp. 661-672, 2009.
- 21 W. Sakoda, M. Sipser. Nondeterminism and the size of two way finite automata. In *Proc. 10th STOC*, pp. 275-286, 1978.
- 22 M. Sipser. Lower bounds on the size of sweeping automata. In *Proc. 11th STOC*, pp. 360-364, 1979.
- 23 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.
- 24 S. Safra and M.Y. Vardi. On ω -Automata and Temporal Logics. In *Proc. 29th STOC*, pp. 127-137, 1989.
- 25 M.Y. Vardi. The Büchi complementation saga. In *Proc. 24th STACS*, pp. 12-22, 2007.
- 26 M.Y. Vardi. and P. Wolper. An automata-theoretic approach to automatic program verification. In *Proc. 1st LICS*, pp. 332-334, 1986.
- 27 Q. Yan. Lower bound for complementation of ω -automata via the full automata technique. In *Proc. 33th ICALP, LNCS 4052*, pp. 589-600, 2006.