

Büchi Complementation Made Tight

Sven Schewe

► **To cite this version:**

Sven Schewe. Büchi Complementation Made Tight. Susanne Albers and Jean-Yves Marion. 26th International Symposium on Theoretical Aspects of Computer Science - STACS 2009, Feb 2009, Freiburg, Germany. IBFI Schloss Dagstuhl, pp.661-672, 2009, Proceedings of the 26th Annual Symposium on the Theoretical Aspects of Computer Science. <inria-00360108>

HAL Id: inria-00360108

<https://hal.inria.fr/inria-00360108>

Submitted on 10 Feb 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

BÜCHI COMPLEMENTATION MADE TIGHT

SVEN SCHEWE

University of Liverpool

E-mail address: Sven.Schewe@liverpool.ac.uk

URL: <http://www.csc.liv.ac.uk/research/logics/>

ABSTRACT. The precise complexity of complementing Büchi automata is an intriguing and long standing problem. While optimal complementation techniques for finite automata are simple – it suffices to determinize them using a simple subset construction and to dualize the acceptance condition of the resulting automaton – Büchi complementation is more involved. Indeed, the construction of an EXPTIME complementation procedure took a quarter of a century from the introduction of Büchi automata in the early 60s, and stepwise narrowing the gap between the upper and lower bound to a simple exponent (of $(6e)^n$ for Büchi automata with n states) took four decades. While the distance between the known upper ($O((0.96n)^n)$) and lower ($\Omega((0.76n)^n)$) bound on the required number of states has meanwhile been significantly reduced, an exponential factor remains between them. Also, the upper bound on the size of the complement automaton is not linear in the bound of its state space. These gaps are unsatisfactory from a theoretical point of view, but also because Büchi complementation is a useful tool in formal verification, in particular for the language containment problem. This paper proposes a Büchi complementation algorithm whose complexity meets, modulo a quadratic ($O(n^2)$) factor, the known lower bound for Büchi complementation. It thus improves over previous constructions by an exponential factor and concludes the quest for optimal Büchi complementation algorithms.

1. Introduction

The precise complexity of Büchi complementation is an intriguing problem for two reasons: First, the quest for optimal algorithms is a much researched problem (c.f., [Büc62, SS78, Péc86, SVW87, Saf88, Mic88, Tho99, Löd99, KV01, GKSV03, FKV06, Pit07, Var07, Yan08]) that has defied numerous approaches to solving it. And second, Büchi complementation is a valuable tool in formal verification (c.f., [Kur94]), in particular when studying language inclusion problems of ω -regular languages. In addition to this, complementation is useful to check the correctness of other translation techniques [Var07, TCT⁺08]. The GOAL tool [TCT⁺08], for example, provides such a test suite and incorporates four of the more recent algorithms [Saf88, Tho99, KV01, Pit07] for Büchi complementation.

1998 ACM Subject Classification: F.4 Mathematical Logic and Formal Languages.

Key words and phrases: Automata and Formal Languages, Büchi Complementation, Automata Theory, Nondeterministic Büchi Automata.

This work was partly supported by the EPSRC through the grand EP/F033567/1 *Verifying Interoperability Requirements in Pervasive Systems*.

While devising optimal complementation algorithms for nondeterministic finite automata is simple—nondeterministic *finite* automata can be determinized using a simple subset construction, and deterministic finite automata can be complemented by complementing the set of final states [RS59, SS78]—devising optimal complementation algorithms for nondeterministic Büchi automata is hard, because simple subset constructions are not sufficient to determinize or complement them [Mic88, Löd99].

Given the hardness and importance of the problem, Büchi complementation enjoyed much attention [Büc62, Péc86, SVW87, Mic88, Saf88, Löd99, Tho99, KV01, GKSV03, FKV06, Var07, TCT⁺08, Yan08], resulting in a continuous improvement of the upper and lower bounds.

The first complementation algorithm dates back to the introduction of Büchi automata in 1962. In his seminal paper “On a decision method in restricted second order arithmetic” [Büc62], Büchi develops a complementation procedure that comprises a doubly exponential blow-up. While Büchi’s result shows that nondeterministic Büchi automata (and thus ω -regular expressions) are closed under complementation, complementing an automaton with n states may, when using Büchi’s complementation procedure, result in an automaton with $2^{2^{O(n)}}$ states, while an $\Omega(2^n)$ lower bound [SS78] is inherited from finite automata.

In the late 80s, these bounds have been improved in a first sequence of results, starting with establishing an EXPTIME upper bound [Péc86, SVW87], which matches the EXPTIME lower bound [SS78] inherited from finite automata. However, the early EXPTIME complementation techniques produce automata with up to $2^{O(n^2)}$ states [Péc86, SVW87]; hence, these upper bounds are still exponential in the lower bounds.

This situation changed in 1988, when Safra introduced his famous determinization procedure for nondeterministic Büchi automata [Saf88], resulting in an $n^{O(n)}$ bound for Büchi complementation, while Michel [Mic88] established a seemingly matching $\Omega(n!)$ lower bound in the same year. Together, these results imply that Büchi complementation is in $n^{\theta(n)}$, leaving again the impression of a tight bound.

As pointed out by Vardi [Var07], this impression is misleading, because the $O()$ notation hides an $n^{\theta(n)}$ gap between both bounds. This gap has been narrowed down in 2001 to $2^{\theta(n)}$ by the introduction of an alternative complementation technique that builds on level rankings and a cut-point construction [KV01]. (Level rankings are functions from the states Q of a nondeterministic Büchi automaton to $\{0, 1, \dots, 2|Q|+1\}$.) The complexity of the plain method is approximately $(6n)^n$ [KV01], leaving a $(6e)^n$ gap to Michel’s lower bound [Mic88].

Recently, tight level rankings [FKV06, Yan08]—a special class of level rankings that is onto a predefined subset—have been exploited by Friedgut, Kupferman, and Vardi [FKV06] to improve the upper complexity bound to $O((0.96n)^n)$, and by Yan [Yan08] to improve the lower complexity bound to $\Omega((0.76n)^n)$.

In the remainder of this paper, we first recapitulate the basic complementation technique of Kupferman and Vardi [KV01], and discuss the core ideas of the improved complexity analysis of Friedgut, Kupferman, and Vardi [FKV06] and Yan [Yan08]. We then show how to improve the complementation technique of Friedgut, Kupferman, and Vardi [FKV06] such that the resulting complementation algorithm meets the known lower bound [Yan08] modulo a small polynomial factor (quadratic in the size of the automaton that is to be complemented), and show that, different to older constructions [KV01, GKSV03], we can achieve an equivalent bound on the number of edges.

2. Preliminaries

2.1. Büchi Automata

Nondeterministic Büchi automata [Büc62] are used to represent ω -regular languages $L \subseteq \Sigma^\omega = \omega \rightarrow \Sigma$ over a finite alphabet Σ . A nondeterministic Büchi automaton $\mathcal{A} = (\Sigma, Q, I, \delta, F)$ is a five tuple, consisting of a finite alphabet Σ , a finite set Q of states with a non-empty subset $I \subseteq Q$ of initial states, a transition function $\delta : Q \times \Sigma \rightarrow 2^Q$ that maps states and input letters to sets of successor states, and a set $F \subseteq Q$ of final states.

Nondeterministic Büchi automata are interpreted over infinite sequences $\alpha : \omega \rightarrow \Sigma$ of input letters. An infinite sequence $\rho : \omega \rightarrow Q$ of states of \mathcal{A} is called a *run* of \mathcal{A} on an input word α if the first letter $\rho(0) \in I$ of ρ is an initial state, and if, for all $i \in \omega$, $\rho(i + 1) \in \delta(\rho(i), \alpha(i))$ is a successor state of $\rho(i)$ for the input letter $\alpha(i)$.

A run $\rho : \omega \rightarrow Q$ is called *accepting* if some finite state appears infinitely often in ρ ($\text{inf}(\rho) \cap F \neq \emptyset$ for $\text{inf}(\rho) = \{q \in Q \mid \forall i \in \omega \exists j > i \text{ such that } \rho(j) = q\}$). A word $\alpha : \omega \rightarrow \Sigma$ is *accepted* by \mathcal{A} if \mathcal{A} has an accepting run on α , and the set $\mathcal{L}(\mathcal{A}) = \{\alpha \in \Sigma^\omega \mid \alpha \text{ is accepted by } \mathcal{A}\}$ of words accepted by \mathcal{A} is called its *language*.

For technical convenience we also allow for finite runs $q_0q_1q_2 \dots q_n$ with $\delta(q_n, \alpha(n)) = \emptyset$. Naturally, no finite run satisfies the Büchi condition; all finite runs are therefore rejecting, and have no influence on the language of an automaton.

The two natural complexity measures for a Büchi automaton are the size $|Q|$ of its state space, and its size $\sum_{q \in Q, \sigma \in \Sigma} 1 + |\delta(q, \sigma)|$, measured in the size of its transition function.

2.2. Run DAG and Acceptance

In [KV01], Kupferman and Vardi introduce a Büchi complementation algorithm that uses level rankings as witnesses for the absence of an accepting run.

The set of all runs of a nondeterministic Büchi automaton $\mathcal{A} = (\Sigma, Q, I, \delta, F)$ on a word $\alpha : \omega \rightarrow \Sigma$ can be represented by a directed acyclic graph (DAG) $\mathcal{G}_\alpha = (V, E)$ with

- vertices $V \subseteq Q \times \omega$ such that $(q, p) \in V$ is in the set V of vertices if and only if there is a run ρ of \mathcal{A} on α with $\rho(p) = q$, and
- edges $E \subseteq (Q \times \omega) \times (Q \times \omega)$ such that $((q, p), (q', p')) \in E$ if and only if $p' = p + 1$ and $q' \in \delta(q, \alpha(p))$ is a successor of q for the input letter $\alpha(p)$.

We call $\mathcal{G}_\alpha = (V, E)$ the *run DAG* of \mathcal{A} for α , and the vertices $V \cap (Q \times \{i\})$ of $\mathcal{G}_\alpha = (V, E)$ that refer to the i^{th} position of runs the i^{th} *level* of $\mathcal{G}_\alpha = (V, E)$.

The run DAG \mathcal{G}_α is called *rejecting* if no path in \mathcal{G}_α satisfies the Büchi condition. That is, \mathcal{G}_α is rejecting if and only if \mathcal{A} rejects α . \mathcal{A} can therefore be complemented to a nondeterministic Büchi automaton \mathcal{B} that checks if \mathcal{G}_α is rejecting.

The property that \mathcal{G}_α is rejecting can be expressed in terms of ranks. We call a vertex $(q, p) \in V$ of a DAG $\mathcal{G} = (V, E)$ *finite*, if the set of vertices reachable from (q, p) in \mathcal{G} is finite, and *endangered*, if no vertex reachable from (q, p) is accepting (that is, in $F \times \omega$).

Based on these definitions, *ranks* can be assigned to the vertices of a rejecting run DAG. We set $\mathcal{G}_\alpha^0 = \mathcal{G}_\alpha$, and repeat the following procedure until a fixed point is reached, starting with $i = 0$:

- Assign all finite vertices of \mathcal{G}_α^i the rank i , and set \mathcal{G}_α^{i+1} to \mathcal{G}_α^i minus the states with rank i (that is, minus the states finite in \mathcal{G}_α^i).

- Assign all endangered vertices of \mathcal{G}_α^{i+1} the rank $i + 1$, and set \mathcal{G}_α^{i+2} to \mathcal{G}_α^{i+1} minus the states with rank $i + 1$ (that is, minus the states endangered in \mathcal{G}_α^{i+1}).
- Increase i by 2.

A fixed point is reached in $n + 1$ steps, and the ranks can be used to characterize the complement language of a nondeterministic Büchi automaton:

Proposition 2.1. [KV01] *A nondeterministic Büchi automaton \mathcal{A} with n states rejects a word $\alpha : \omega \rightarrow \Sigma$ if and only if $\mathcal{G}_\alpha^{2n+1}$ is empty.* ■

To see that a fixed point is reached after $n + 1$ iterations, note that deleting all finite or endangered vertices leaves a DAG without finite or endangered vertices, respectively. Hence a fixed point is reached as soon as we do not assign a rank i to any vertex. By construction, no DAG $\mathcal{G}_\alpha^{2i+1}$ contains finite vertices. If it contains an endangered vertex v , then all vertices reachable from v are endangered, too. This implies that some vertex of almost all levels is assigned the rank $2i + 1$. Hence, if no fixed point is reached earlier, some rank is assigned to all vertices of almost all levels after n iterations (there cannot be more than n vertices in a level), and all vertices in \mathcal{G}_α^{2n} must be finite.

If the reached fixed point $\mathcal{G}_\alpha^\infty$ is non-empty, then it contains only infinite and non-endangered vertices, and constructing an accepting run from $\mathcal{G}_\alpha^\infty$ is simple. Vice versa, an accepting run ρ on an ω -word α can be viewed as an infinite sub-graph of \mathcal{G}_α that does not contain finite or endangered nodes. By a simple inductive argument, the sub-graph identified by ρ is therefore a sub-graph of \mathcal{G}_α^i for all $i \in \omega$, and hence of $\mathcal{G}_\alpha^\infty$.

2.3. Büchi Complementation

The connection between Büchi complementation, run DAGs and ranks leads to an elegant complementation technique. We call the maximal rank of a vertex in a level the rank of this level, the rank of almost all vertices $(\rho(i), i)$, $i \in \omega$ of a run ρ (or: path in \mathcal{G}) the rank of ρ , and the rank of almost all levels of a DAG \mathcal{G} the rank of \mathcal{G} . (Note that level ranks can only go down, and that vertex ranks can only go down along a path.)

For a given nondeterministic Büchi automaton $\mathcal{A} = (\Sigma, Q, I, \delta, F)$ with n states, we call a function $f : Q \rightarrow \{0, 1, \dots, 2n\}$ that maps all accepting states to odd numbers ($f(F) \cap 2\omega = \emptyset$) a *level ranking*.

Proposition 2.2. [KV01] *For a given nondeterministic Büchi automaton $\mathcal{A} = (\Sigma, Q, I, \delta, F)$, the nondeterministic Büchi automaton $\mathcal{B} = (\Sigma, Q', I', \delta', F')$ with*

- $Q' = 2^Q \times 2^Q \times \mathcal{R}$,
- $I' = \{I\} \times \{\emptyset\} \times \mathcal{R}$,
- $\delta'((S, O, f), \sigma) = \{(\delta(S, \sigma), \delta(O, \sigma) \setminus \text{odd}(f'), f') \mid f' \leq_\sigma^S f, O \neq \emptyset\} \cup \{(\delta(S, \sigma), \delta(S, \sigma) \setminus \text{odd}(f'), f') \mid f' \leq_\sigma^S f, O = \emptyset\}$, and
- $F' = 2^Q \times \{\emptyset\} \times \mathcal{R}$,

where

- \mathcal{R} is the set of all level rankings of \mathcal{A} ,
- $\text{odd}(f) = \{q \in Q \mid f(q) \text{ is odd}\}$, and
- $f' \leq_\sigma^S f \Leftrightarrow \forall q \in S, q' \in \delta(q, \sigma). f'(q') \leq f(q)$,

accepts the complement $\mathcal{L}(\mathcal{B}) = \overline{\mathcal{L}(\mathcal{A})} = \Sigma^\omega \setminus \mathcal{L}(\mathcal{A})$ of the language of \mathcal{A} . ■

The first element S of a triple $(S, O, f) \in Q'$ reflects the set of states of \mathcal{A} reachable upon the input seen so far (the states in the respective level of the run DAG), and the third element f is a mapping that intuitively maps reachable states to their rank.

The condition $f' \leq_{\sigma}^S f$ ensures that the rank of the vertices (or rather: the value assigned to them) is decreasing along every path of the run DAG. The second element is used for a standard cut-point construction, comparable to the cut-point constructions in the determinization of Co-Büchi automata or the nondeterminization of alternating Büchi automata. It contains the positions whose rank (or rather: the value assigned to them) has been even ever since the last cut-point ($O = \emptyset$) was reached; it intuitively ensures that the respective vertices are finite.

2.4. Tight Level Rankings

Friedgut, Kupferman, and Vardi [FKV06] improved this complementation technique by exploiting the observation that the true ranks of the run DAG \mathcal{G}_{α} of a rejected ω -word α are eventually always tight. A level ranking $f : Q \rightarrow \omega$ is called *tight*, if it has an odd rank r , and is onto the odd numbers $\{1, 3, \dots, r\}$ up to its rank r , and *S-tight*, if its restriction to S is tight and if it maps all states not in S to 1 ($f(q) = 1 \ \forall q \in Q \setminus S$).

Proposition 2.3. [FKV06] *For every run DAG \mathcal{G}_{α} with finite rank r , it holds that*

- r is odd, and
- there is a level $l \geq 0$ such that, for all levels $l' \geq l$ and all odd ranks $o \leq r$, there is a node $(q, l') \in \mathcal{G}_{\alpha}$ with rank o in \mathcal{G}_{α} . ■

This immediately follows from what was said in Subsection 2.2 on reaching a fixed point after $n + 1$ iterations: If the rank \mathcal{G}_{α} is r then, for every odd number $o \leq r$, almost all levels contain a vertex with rank o , and assuming that r is even implies that all vertices of \mathcal{G}_{α}^r are finite, which in turn implies that only finitely many levels contain a vertex with rank r and hence leads to a contradiction.

Using this observation, the construction from Proposition 2.2 can be improved by essentially replacing \mathcal{R} by the set $\mathcal{T} = \{f \in \mathcal{R} \mid f \text{ is tight}\}$ of tight level rankings. While the size $|\mathcal{R}| = (2n + 1)^n$ of \mathcal{R} is in $\theta((2n)^n)$, the size of \mathcal{T} ,

$$\text{tight}(n) = |\mathcal{T}|,$$

is much smaller. Building on an approximation of Stirling numbers of the second kind by Temme [Tem93], Yan and Friedgut, Kupferman, and Vardi [FKV06, Yan08] showed that $\text{tight}(n)$ can be approximated by $(\kappa n)^n$ for a constant $\kappa \approx 0.76$, that is, they showed

$$\kappa = \lim_{n \rightarrow \infty} \frac{\sqrt[n]{\text{tight}(n)}}{n} \approx 0.76.$$

Friedgut, Kupferman, and Vardi [FKV06] use this observation—together with other improvements—for an improved complementation algorithm that produces a complement automaton with approximately $(0.96 n)^n$ states [FKV06].

Yan [Yan08] showed for full automata—a family of automata that has exactly one accepting state, and an alphabet that encodes the possible transitions between the states of the automaton—that every nondeterministic Büchi automaton that accepts the complement language of a full automaton with $n + 1$ states must have $\Omega(\text{tight}(n))$ states.

Proposition 2.4. [Yan08] *A nondeterministic Büchi automaton that accepts the complement language of a full Büchi automaton with n states has $\Omega(\text{tight}(n - 1))$ states. ■*

3. Efficient Büchi Complementation

To optimize the construction from Proposition 2.2, we turn not only to tight functions (c.f. [FKV06]), but also refine the cut-point construction. While the cut-point construction of Proposition 2.2 tests *concurrently* for all even ranks if a path has finite even rank, we argue that it is much cheaper to test this property *turn wise* for all even ranks individually. As a result, the overall construction becomes more efficient and meets, modulo a small polynomial factor in $O(n^2)$, the lower bound recently established by Yan [Yan08].

3.1. Construction

The obtained state space reduction of the proposed construction compared to [FKV06] is due to an efficient cut-point construction in combination with the restriction to tight rankings. The improved cut-point construction is inspired by the efficient translation from generalized to ordinary Büchi automata. Indeed, the acceptance condition that no trace has an *arbitrary* even rank, which is reflected by the straight-forward acceptance condition of previous algorithms [KV01, FKV06], can be replaced by an acceptance condition, which only rules out that some trace has a particular even rank, but does so for all potential even ranks.

Checking the condition for a particular even rank allows for focusing on exactly this rank in the cut-point construction, which led to a significant cut in the size of the resulting state space. While this approach cannot be taken if we literally use a generalized Büchi condition, the idea of cyclically considering the relevant even ranks proves to be feasible.

Construction: For a given nondeterministic Büchi automaton $\mathcal{A} = (\Sigma, Q, I, \delta, F)$ with $n = |Q|$ states, let $\mathcal{C} = (\Sigma, Q', I', \delta', F')$ denote the nondeterministic Büchi automaton with

- $Q' = Q_1 \cup Q_2$ with
 - $Q_1 = 2^Q$ and
 - $Q_2 = \{(S, O, f, i) \in 2^Q \times 2^Q \times \mathcal{T} \times \{0, 2, \dots, 2n - 2\} \mid$
 $f \text{ is } S\text{-tight, } O \subseteq S \text{ and } \exists i \in \omega. O \subseteq f^{-1}(2i)\},$
- $I' = \{I\},$
- $\delta' = \delta_1 \cup \delta_2 \cup \delta_3$ for
 - $\delta_1 : Q_1 \times \Sigma \rightarrow 2^{Q_1}$ with $\delta_1(S, \sigma) = \{\delta(S, \sigma)\},$
 - $\delta_2 : Q_1 \times \Sigma \rightarrow 2^{Q_2}$ with $(S', O, f, i) \in \delta_2(S, \sigma) \Leftrightarrow S' = \delta(S, \sigma), O = \emptyset, \text{ and } i = 0,$
 - $\delta_3 : Q_2 \times \Sigma \rightarrow 2^{Q_2}$ with $(S', O', f', i') \in \delta_3((S, O, f, i), \sigma)$
 $\Leftrightarrow S' = \delta(S, \sigma), f' \leq_{\sigma}^S f, \text{ rank}(f) = \text{rank}(f'), \text{ and}$
 - * $i' = (i + 2) \bmod (\text{rank}(f') + 1)$ and $O' = f'^{-1}(i')$ if $O = \emptyset$ or
 - * $i' = i$ and $O' = \delta(O, \sigma) \cap f'^{-1}(i)$ if $O \neq \emptyset$, respectively, and
- $F' = \{\emptyset\} \cup (2^Q \times \{\emptyset\} \times \mathcal{T} \times \omega) \cap Q_2.$

The complement automaton \mathcal{C} operates in two phases. In a first phase it only traces the states reachable in \mathcal{A} upon a finite input sequence. In this phase, only the states in Q_1 and the transition function δ_1 are used. In the special case that \mathcal{A} rejects an ω -word α because \mathcal{A} has no run on α , \mathcal{C} accepts by staying forever in phase one, because $\{\emptyset\}$ is final.

\mathcal{C} intuitively uses its nondeterministic power to guess a point in time where all successive levels are tight and have the same rank. At such a point, \mathcal{C} traverses from Q_1 to Q_2 , using a transition from δ_2 . Staying henceforth in Q_2 (using the transitions from δ_3), \mathcal{C} intuitively verifies turn wise for all potential even ranks e that no path has this particular even rank e . For a particular rank e , it suffices to trace the positions on traces with unchanged rank e (hence $O \subseteq f^{-1}(e)$), and to cyclically update the designated even rank after every cut-point.

3.2. Correctness

To show that the automaton \mathcal{C} from the construction introduced in the previous subsection recognises the complement language of \mathcal{A} , we first show that the complement language of \mathcal{C} contains the language of \mathcal{A} ($\mathcal{L}(\mathcal{A}) \subseteq \overline{\mathcal{L}(\mathcal{C})}$), and then that the complement language of \mathcal{A} is contained in the language of \mathcal{C} ($\overline{\mathcal{L}(\mathcal{A})} \subseteq \mathcal{L}(\mathcal{C})$).

Lemma 3.1. *If a given nondeterministic Büchi automaton $\mathcal{A} = (\Sigma, Q, I, \delta, F)$ accepts an ω -word $\alpha : \omega \rightarrow \Sigma$, then α is rejected by the automaton $\mathcal{C} = (\Sigma, Q', I', \delta', F')$. ($\mathcal{L}(\mathcal{A}) \subseteq \overline{\mathcal{L}(\mathcal{C})}$)*

Proof. Let $\rho : \omega \rightarrow Q$ be an accepting run of \mathcal{A} on α . First, $\rho' = S_0, S_1, S_2 \dots$ with $S_0 = I$ and $S_{i+1} = \delta(S_i, \alpha(i))$ for all $i \in \omega$ is no accepting run of \mathcal{C} , because $\rho(i) \in S_i$ and hence no state of ρ' is accepting ($\emptyset \neq S_i \not\subseteq F'$ for all $i \in \omega$). Let now

$$\rho' = S_0, S_1, S_2, \dots, S_p, (S_{p+1}, O_{p+1}, f_{p+1}, i_{p+1}), (S_{p+2}, O_{p+2}, f_{p+2}, i_{p+2}), \dots$$

be a run of \mathcal{C} on α . Again, we have that $\rho(j) \in S_j$ for all $j \in \omega$. Furthermore, the construction guarantees that $f_{j+1} \leq_{\alpha(j)}^S f_j$ holds for all $j > p$. The sequence

$$f_{p+1}(\rho(p+1)) \geq f_{p+2}(\rho(p+2)) \geq f_{p+3}(\rho(p+3)) \geq \dots$$

is therefore decreasing, and stabilizes eventually. That is, there is a $k > p$ and a $v \leq 2n$ such that $f_l(\rho(l)) = v$ for all $l \geq k$. Since ρ is accepting, there is a position $l \geq k$ with $\rho(l) \in F$. Taking into account that f_l is a level ranking, this implies that $f_l(\rho(l))$ —and hence v —is even. Assuming that ρ' is accepting, we can infer that, for some position $l > k$ which follows one of the first n accepting states of ρ' after position k , $i_l = v$ and $O_l = f_l^{-1}(v) \ni \rho(l)$. It is now easy to show by induction that, for all $m \geq l$, $i_m = v$ and (using $f_m(\rho(m)) = v$) $\rho(m) \in O_m \neq \emptyset$ hold true, which contradicts the assumption that ρ' is accepting. ■

To prove the second lemma, $\overline{\mathcal{L}(\mathcal{A})} \subseteq \mathcal{L}(\mathcal{C})$, we use Propositions 2.1 and 2.3 to infer that the run DAG \mathcal{G}_α of an ω -word rejected by \mathcal{A} is either finite or has odd bounded rank and only finitely many non-tight level rankings. We use this to build an accepting run of \mathcal{C} on α .

Lemma 3.2. *For a nondeterministic Büchi automaton $\mathcal{A} = (\Sigma, Q, I, \delta, F)$, the automaton $\mathcal{C} = (\Sigma, Q', I', \delta', F')$ accepts an ω -word $\alpha : \omega \rightarrow \Sigma$ if α is rejected by \mathcal{A} . ($\overline{\mathcal{L}(\mathcal{A})} \subseteq \mathcal{L}(\mathcal{C})$)*

Proof. If $\alpha : \omega \rightarrow \Sigma$ is rejected by \mathcal{A} , then the run DAG \mathcal{G}_α has bounded rank by Proposition 2.1, and by Proposition 2.3 almost all levels of \mathcal{G}_α have tight level rankings with the same rank r . For the special case that the rank of all vertices of \mathcal{G}_α is 0, that is, if all vertices of \mathcal{G}_α are finite, $\rho' = S_0, S_1, S_2 \dots$ with $S_0 = I$ and $S_{i+1} = \delta(S_i, \alpha(i))$ for all $i \in \omega$ is an accepting run of \mathcal{C} on α .

If \mathcal{G}_α contains an infinite vertex, then we fix a position $p \in \omega$ such that the rank of all levels $p' \geq p$ of \mathcal{G}_α is r and tight for some (odd) $r \geq 1$. We now consider a run

$$\rho' = S_0, S_1, S_2, \dots, S_p, (S_{p+1}, O_{p+1}, f_{p+1}, i_{p+1}), (S_{p+2}, O_{p+2}, f_{p+2}, i_{p+2}), \dots \text{ of } \mathcal{C} \text{ on } \alpha \text{ with}$$

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

ρ' is obviously a run of \mathcal{C} on α . To show that ρ' is accepting, we have to show that O_j is empty infinitely many times. Let us assume that this is not the case; that is, let us assume that there is a last element $\rho'(j)$ with $O_j = \emptyset$ and $O_k \neq \emptyset$ for all $k > j$. (Note that $O_{p+1} = \emptyset$ is empty.) Then we have that $i_k = i_{j+1}$ for all $k > j$, and it is easy to show by induction for all $k > j$ that $O_k \times \{k\}$ is the set of states reachable in $\mathcal{G}_\alpha^{i_{j+1}}$ from some state in $O_{j+1} \times \{j+1\}$. But since the rank of all states in $O_{j+1} \times \{j+1\}$ is i_{j+1} (and thus even), all of these states are finite in $\mathcal{G}_\alpha^{i_{j+1}}$, which implies that there are only finitely many states reachable in $\mathcal{G}_\alpha^{i_{j+1}}$ from $O_{j+1} \times \{j+1\}$, and thus contradicts the assumption that $O_k \neq \emptyset$ is non-empty for all $k > j$. ■

The two lemmata of this subsection immediately imply the claimed language complementation:

Corollary 3.3. *For a given nondeterministic Büchi automaton $\mathcal{A} = (\Sigma, Q, I, \delta, F)$, the automaton \mathcal{C} resulting from the construction introduced in Subsection 3.1 recognises the complement language of \mathcal{A} . ($\mathcal{L}(\mathcal{C}) = \overline{\mathcal{L}(\mathcal{A})}$)* ■

3.3. Complexity

The costly part in previous approaches [KV01, FKV06] that the proposed method avoids (de facto, although not technically), is a subset construction in addition to the level rankings. Avoiding the subset construction results in a state space reduction by a factor exponential in the size of the automaton \mathcal{A} , and even to upper bounds comparable to the established lower bounds [Yan08].

The subset construction is de facto avoided, because it can be encoded into the ranking function once we allow for a slightly enlarged set of output values. For example, we could map all states not in S to -1 , and all states in O to -2 . Following this convention, the first two elements of every tuple in Q_2 could be pruned. (They remain explicit in the construction because this representation is more comprehensible, and outlines the connection to the older constructions of Friedgut, Kupferman, and Vardi [KV01, FKV06].

Theorem 3.4. *For a given nondeterministic Büchi automaton \mathcal{A} with n states, the automaton \mathcal{C} has $O(\text{tight}(n+1))$ states.*

Proof. Q_1 is obtained by a simple subset construction, and hence $Q_1 \in O(2^n)$, which is a small subset of $O(\text{tight}(n+1))$. For a fixed i , a state (S, O, f, i) with an S -tight level ranking f of rank r can be represented by a function $g: Q \rightarrow \{-2, -1, \dots, r\}$ that maps every state $q \in O$ to $g(q) = -1$, every state $q \in Q \setminus S$ not in S to $g(q) = -2$, and every other state $q \in S \setminus O$ to $g(q) = f(q)$. Every such function g either has a domain $g(Q) \subseteq \{0, 1, \dots, r\}$ and is onto $\{1, 3, \dots, r\}$, or is a function to $\{-2, -1, \dots, r\}$ and onto $\{-1\} \cup \{1, 3, \dots, r\}$ or $\{-2\} \cup \{1, 3, \dots, r\}$. The size of all three groups of functions is hence in $O(\text{tight}(n))$ (for a fixed i), which results in an overall size in $O(\text{tight}(n+1)) = O(n \cdot \text{tight}(n))$. ■

Together with Proposition 2.4, this establishes tight bounds for Büchi complementation:

Corollary 3.5. *The minimal size of the state space of a nondeterministic Büchi automaton that accept the complement language of a nondeterministic Büchi automaton with n states is in $\Omega(\text{tight}(n-1))$ and $O(\text{tight}(n+1))$.* ■

4. Reduced Average Outdegree

A flaw in the construction presented in Section 3 is that it is optimal only with respect to the state space of the automata. In [KV01], Kupferman and Vardi discuss how to reduce the number of edges such that the bound on the number of edges becomes trilinear in the alphabet size, the bound on the number of states, and the rank of the resulting automaton (see also [GKSV03]). In this section we improve the construction from the previous section such that the bound on the number of edges is merely bilinear in the bound on the number of states and the size of the input alphabet. The technique can be adapted to generally restrict the outdegree to $|\delta(q, \sigma)| \leq 2$ when level rankings are not required to be tight.

4.1. Construction

The automaton \mathcal{C} obtained from the construction described in Section 3 operates in two phases. In a first phase, it stays in Q_1 and only tracks the reachable states of the Büchi automata \mathcal{A} it complements. It then guesses a point $p \in \omega$ such that all levels $j > p$ of \mathcal{G}_α have a tight level ranking to transfer to Q_2 .

We improve over the construction of Section 3 by restricting the number of entry points to Q_2 from $O(\text{tight}(n))$ to $O(n!)$, and by restricting the number of outgoing transitions $|\delta(q, \sigma)| \leq 2$ for all states $q \in Q_2$ and input letters $\sigma \in \Sigma$ to two. The latter is achieved by allowing only the successor $(S, O, f, i) \in \delta_3(q, \sigma)$ with a point wise maximal function f (the γ_3 -transitions) or with a function f that is maximal among the final states $(S, O, f, i) \in \delta_3(q, \sigma) \cap F$ among them (the γ_4 -transitions). If such elements exist, then they are unique.

The first restriction is achieved by restricting δ_2 to states (S, O, f, i) for which f is maximal with respect to S . We call an S -tight level ranking f with rank r *maximal with respect to S* if it maps all final states $q \in F \cap S$ in S to $r - 1$, exactly one state to every odd number $o < r$ smaller than r ($|f^{-1}(o)| = 1$) and all remaining states of S to r , and denote the set of tight rankings that are maximal with respect to S by $\mathcal{M}_S = \{f \in \mathcal{T} \mid f \text{ is maximal with respect to } S\}$.

As there are only $|Q_1| \leq 2^n$ states in Q_1 , the impact of their high outdegree is outweighed by the small outdegree ($|\delta(q, \sigma)| \leq 2$) of the remaining $|Q_2| \in O(\text{tight}(n + 1))$ states.

Construction: For a given nondeterministic Büchi automaton $\mathcal{A} = (\Sigma, Q, I, \delta, F)$ with $n = |Q|$ states, let $\mathcal{D} = (\Sigma, Q', I', \gamma, F')$ denote the nondeterministic Büchi automaton with Q' , I' , F' , Q_1, Q_2 , and δ_1 as in the construction from Section 3, and with $\gamma = \delta_1 \cup \gamma_2 \cup \gamma_3 \cup \gamma_4$ for

- $\gamma_2 : Q_1 \times \Sigma \rightarrow 2^{Q_2}$ with $(S', O, g, i) \in \gamma_2(S, \sigma) \Leftrightarrow (S', O, g, i) \in \delta_2(S, \sigma)$ and $g \in \mathcal{M}_{S'}$,
- $\gamma_3 : Q_2 \times \Sigma \rightarrow 2^{Q_2}$ with $\gamma_3((S, O, f, i), \sigma) = \{ \max_g \{ (S', O', g, i') \in \delta_3((S, O, f, i), \sigma) \} \}$,
- $\gamma_4 : Q_2 \times \Sigma \rightarrow 2^{Q_2}$ with $(S', O'', g', i') \in \gamma_4((S, O, f, i), \sigma)$ if

$$(S', O', g, i') \in \gamma_3((S, O, f, i), \sigma), O'' = \emptyset, i' \neq 0 \vee O' = \emptyset, \text{ and}$$

$$g'(q) = g(q) - 1 \text{ for all } q \in O' \text{ and } g'(q) = g(q) \text{ otherwise,}$$

where $\max_g \{ (S', O', g, i') \in \delta_3((S, O, f, i), \sigma) \}$ selects the unique element with a (point wise) maximal function g . The supremum over all function obviously exists, but it is not necessarily tight. (In this case, \max_g returns the empty set.) Since δ_3 is strict with respect to the selection of the other three elements S', O' and i' for a fixed ranking function, the mapping of γ_3 consists only of singletons and the empty set.

While γ_3 selects a maximal successor, γ_4 selects a maximal final successor, which only requires to decrease the value assigned to the states in O' by the ranking function g by one.

(Which cannot be done if their value is already 0, hence the restriction $i \neq 0$ or $O' = \emptyset$.) The tightness of g' is then inherited from the tightness of g .

4.2. Correctness

While it is clear that the language of \mathcal{D} is contained in the language of \mathcal{C} , the converse is less obvious. To prove $\mathcal{L}(\mathcal{C}) \subseteq \mathcal{L}(\mathcal{D})$, we show that any ω -word α rejected by \mathcal{A} will be accepted by \mathcal{D} by exploiting the “standard” run ρ' of \mathcal{C} on α from the proof of Lemma 3.2 to build an accepting run ρ'' of \mathcal{D} on α .

Proposition 4.1. *For a given nondeterministic Büchi automaton $\mathcal{A} = (\Sigma, Q, I, \delta, F)$, the automaton \mathcal{D} resulting from the construction introduced in Subsection 4.1 accepts an ω -word $\alpha : \omega \rightarrow \Sigma$ if and only if α is rejected by \mathcal{A} . ($\mathcal{L}(\mathcal{D}) = \overline{\mathcal{L}(\mathcal{A})}$)*

Proof. We show $\mathcal{L}(\mathcal{D}) = \overline{\mathcal{L}(\mathcal{A})}$ by demonstrating $\overline{\mathcal{L}(\mathcal{A})} \subseteq \mathcal{L}(\mathcal{D}) \subseteq \mathcal{L}(\mathcal{C}) \subseteq \overline{\mathcal{L}(\mathcal{A})}$, where the second inclusion is implied by the fact that every (accepting) run of \mathcal{D} is also an (accepting) run of \mathcal{C} , and the third inclusion is shown in Lemma 3.1.

To demonstrate $\overline{\mathcal{L}(\mathcal{A})} \subseteq \mathcal{L}(\mathcal{D})$, we reuse the proof of Lemma 3.2 to obtain an accepting run $\rho' = S_0, S_1, S_2, \dots, S_p, (S_{p+1}, O_{p+1}, f_{p+1}, i_{p+1}), (S_{p+2}, O_{p+2}, f_{p+2}, i_{p+2}), \dots$ for \mathcal{C} , where $f_j(q)$ is the rank of (q, j) in \mathcal{G}_α for all $j > p$ and $q \in S_j$. (If \mathcal{A} has no run on α , then \mathcal{D} has the same accepting standard run on α that stays in Q_1 as \mathcal{C} .)

Let us pick an S_{p+1} -tight ranking function g_{p+1} with the same rank as f_{p+1} that is maximal with respect to S_{p+1} . We show that we then can construct the run

$$\rho'' = S_0, S_1, S_2, \dots, S_p, (S_{p+1}, O'_{p+1}, g_{p+1}, i'_{p+1}), (S_{p+2}, O'_{p+2}, g_{p+2}, i'_{p+2}), \dots$$

of \mathcal{D} on α that satisfies $O_{p+1} = \emptyset$, and $i_{p+1} = 0$ and, for all $j > p$,

- $(S_{j+1}, O'_{j+1}, g_{j+1}, i'_{j+1}) \in \gamma_4((S_j, O'_j, g_j, i'_j), \alpha(j))$ if $O_{j+1} = \emptyset$ and $i_{j+1} = i'_{j+1}$ (note that i'_{j+1} does not depend on taking the transition from γ_3 or γ_4), and
- $(S_{j+1}, O'_{j+1}, g_{j+1}, i'_{j+1}) \in \gamma_3((S_j, O'_j, g_j, i'_j), \alpha(j))$ otherwise.

To show by induction that $g_j \geq f_j$ holds true for all $j > p$ (where \geq is the point wise comparison), we strengthen the claim by claiming additionally that if, for some position $k > p$, $\rho'(k)$ is a final state, $i'_{k+1} = i_{k+1}$ holds true, and $k' > k$ is the next position for which $\rho''(k')$ is final, then $q \in O'_j$ implies $q \in O_j \vee g_j(q) > f_j(q)$ for all $k < j \leq k'$.

For $j = p + 1$ this holds trivially (basis). For the induction step, let us first consider the case of γ_3 -transitions. Then $g_{j+1} \geq f_{j+1}$ is implied, because g_{j+1} is maximal among the S_{j+1} -tight level rankings $\leq_{\alpha(j)}^{S_j} g_j$, and $g_j \geq f_j$ holds by induction hypothesis. If $\rho'(j)$ is a final state and $i'_{j+1} = i_{j+1}$ holds true, then $O_{j+1} = f_j^{-1}(i_{j+1})$, and hence $q \in O'_j$ implies $g_{j+1}(q) = i'_{j+1} = i_{j+1}$, which implies $q \in O_{j+1} \vee g_{j+1}(q) > f_{j+1}(q)$ (using $g_{j+1} \geq f_{j+1}$).

If a γ_4 -transition is taken, then taking a γ_3 -transition implied $g_{j+1} \geq f_{j+1}$ and $g_{j+1}(q) = i'_{j+1} = i_{j+1} \Rightarrow g_{j+1}(q) > f_{j+1}(q)$ (note that O_{j+1} is empty) by the previous argument. This immediately implies $g_{j+1} \geq f_{j+1}$ for the γ_4 -transition. Consequently, $O_{j+1} = f_j^{-1}(i_{j+1})$ (which holds as $\rho'(j)$ is final) entails that $q \in O'_j$ implies $q \in O_{j+1} \vee g_{j+1}(q) > f_{j+1}(q)$.

It remains to show that all functions g_j are S_j -tight level rankings. To demonstrate this, let $q \in S_{p+1}$ be a state of the automaton \mathcal{A} such that $g_{p+1}(q) = o$ is the rank of $(q, p + 1)$ in \mathcal{G}_α for an odd number $o \leq r$. (Such a state exists for every odd number $o \leq r$ by construction.) Since $(q, p + 1)$ is endangered but not finite in \mathcal{G}_α^o , all nodes (q', j) with $j > p$ reachable from $(q, p + 1)$ in \mathcal{G}_α^o form an infinite connected sub-DAG of \mathcal{G}_α^o , all of whose

nodes have rank o . (Which, by the proof in Lemma 3.2, entails that $f_j(q') = o$ holds for every vertex (q', j) of this sub-DAG.) By definition of ρ'' , it is easy to show by induction that $g_j(q') \leq o$ holds for all of these nodes (q', j) . As we have just demonstrated $g_j(q') \geq f_j(q')$, this entails $g_j(q') = o$. Since o can be any odd number less or equal to the rank r of \mathcal{G}_α , and since there is, for every $j > p$, some vertex (q', j) reachable from $(q, p + 1)$ in \mathcal{G}_α^o , g_j is an S_j -tight level ranking for all $j > p$.

Finally, the assumption that O'_j is empty only finitely many times implied that there was a last position k such that O'_k is empty. But this implies that i_j is stable for all $j > k$, and within the next n visited fixed points in ρ' there is one that refers to this i_{k+1} . By construction of ρ'' , this position is a final state in ρ'' , too. ζ ■

4.3. Complexity

Extending the tight bound of Section 3 for the state space to a tight bound on the size of the complement automaton is simple: The mappings of δ_1 , γ_3 , and γ_4 consist of singletons or the empty set, such that only the size of γ_2 needs to be considered more closely.

Theorem 4.2. *For a given nondeterministic Büchi automaton \mathcal{A} with n states and an alphabet of size s , the automaton \mathcal{D} has size $O(s \text{tight}(n + 1))$.*

Proof. For all $S \in Q_1$ and $\sigma \in \Sigma$, we have that $\gamma_2(S, \sigma) = \{S'\} \times \{\emptyset\} \times \mathcal{M}_{S'} \times \{0\}$ for $S' = \delta(S)$. Thus, $|\gamma_2(S, \sigma)| = |\mathcal{M}_{S'}|$, which can be estimated by $\sum_{i=1}^m \frac{m!}{i!}$ for $m = |S' \setminus F|$, which is in $O(n!)$. Thus $\sum_{S \in Q_1, \sigma \in \Sigma} |\gamma_2(S, \sigma)| \in O(s 2^n n!) \subsetneq o(s \text{tight}(n))$ holds true.

($2^n n! \approx (\frac{2n}{e})^n \approx (0.74 n)^n$, whereas $\text{tight}(n) \approx (0.76 n)^n$.) The claim thus follows with Theorem 3.4, and $|\delta_1(q_1, \sigma)| = 1$ and $|\gamma_3(q_2, \sigma)|, |\gamma_4(q_2, \sigma)| \leq 1$ for all $q_1 \in Q_1, q_2 \in Q_2$ and $\sigma \in \Sigma$. ■

Together with Proposition 2.4, this establishes tight complexity bounds for Büchi complementation:

Corollary 4.3. *The complexity of complementing nondeterministic Büchi automata with n states is in $\Omega(\text{tight}(n - 1))$ and $O(\text{tight}(n + 1))$. The discussed complementation technique is therefore optimal modulo a small polynomial factor in $O(n^2)$.* ■

5. Discussion

This paper marks the end of the long quest for the precise complexity of the Büchi complementation problem. It shows that the previously known lower bound is sharp, which is on one hand surprising, because finding tight lower bounds is generally considered the harder problem, and seems on the other hand natural, because Yan’s lower bound builds on the concept of tight level rankings alone [Yan08], while the previously known upper bound [FKV06] incorporates an additional subset construction and builds on estimations on top of this, leaving the estimations of the lower bound the simpler concept of the two.

Similar to the complexity gap in Büchi complementation twenty years ago, the complexity of Büchi determinization is known to be in $n^{\theta(n)}$, but there is also an $n^{\theta(n)}$ gap between the upper [Sch09] and lower [Yan08] bound. Tightening the bounds for Büchi determinization appears to be the natural next step after the introduction of an optimal Büchi complementation algorithm.

References

- [Büc62] J. Richard Büchi. On a decision method in restricted second order arithmetic. In *Proceedings of the International Congress on Logic, Methodology, and Philosophy of Science, 1960, Berkeley, California, USA*, pages 1–11. Stanford University Press, 1962.
- [FKV06] Ehud Friedgut, Orna Kupferman, and Moshe Y. Vardi. Büchi complementation made tighter. *International Journal of Foundations of Computer Science*, 17(4):851–867, 2006.
- [GKSV03] Sankar Gurumurthy, Orna Kupferman, Fabio Somenzi, and Moshe Y. Vardi. On complementing nondeterministic Büchi automata. In *Proceedings of the 12th Advanced Research Working Conference on Correct Hardware Design and Verification Methods (CHARME 2003), 21–24 October, L'Aquila, Italy*, volume 2860 of *Lecture Notes in Computer Science*, pages 96–110. Springer-Verlag, 2003.
- [Kur94] Robert P. Kurshan. *Computer-aided verification of coordinating processes: the automata-theoretic approach*. Princeton University Press, 1994.
- [KV01] Orna Kupferman and Moshe Y. Vardi. Weak alternating automata are not that weak. *ACM Transactions on Computational Logic*, 2(2):408–429, July 2001.
- [Löd99] Christof Löding. Optimal bounds for transformations of ω -automata. In *Proceedings of the 19th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 1999), 13–15 December, Chennai, India*, volume 1738 of *Lecture Notes in Computer Science*, pages 110–121. Springer-Verlag, 1999.
- [Mic88] M. Michel. Complementation is more difficult with automata on infinite words. Technical report, CNET, Paris (Manuscript), 1988.
- [Péc86] Jean-Pierre Pécuchet. On the complementation of Büchi automata. *Theoretical Computer Science*, 47(3):95–98, 1986.
- [Pit07] Nir Piterman. From nondeterministic Büchi and Streett automata to deterministic parity automata. *Journal of Logical Methods in Computer Science*, 3(3:5), 2007.
- [RS59] Michael O. Rabin and Dana S. Scott. Finite automata and their decision problems. *IBM Journal of Research and Development*, 3:115–125, 1959.
- [Saf88] Shmuel Safra. On the complexity of ω -automata. In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science (FOCS 1988), 24–26 October*, pages 319–327, White Plains, New York, USA, 1988. IEEE Computer Society Press.
- [Sch09] Sven Schewe. Tighter bounds for the determinization of Büchi automata. In *Proceedings of the Twelfth International Conference on Foundations of Software Science and Computation Structures (FoSSaCS 2009), 22–29 March, York, England, UK*. (to appear), 2009.
- [SS78] William J. Sakoda and Michael Sipser. Non-determinism and the size of two-way automata. In *Proceedings of the 10th Annual ACM Symposium on Theory of Computing (STOC 1978), 1–3 May, San Diego, California, USA*, pages 274–286. ACM Press, 1978.
- [SVW87] A. Prasad Sistla, Moshe Y. Vardi, and Pierre Wolper. The complementation problem for Büchi automata with applications to temporal logic. *Theoretical Computer Science*, 49(3):217–239, 1987.
- [TCT⁺08] Yih-Kuen Tsay, Yu-Fang Che, Ming-Hsien Tsai, Wen-Chin Chan, and Chi-Jian Luo. GOAL extended: Towards a research tool for omega automata and temporal logic. In *Proceedings of the 14th International Conference On Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2008), 31 March–3 April, Budapest, Hungary*, volume 4963 of *Lecture Notes in Computer Science*, pages 346–350. Springer-Verlag, 2008.
- [Tem93] Nico M. Temme. Asymptotic estimates of Stirling numbers. *Studies in Applied Mathematics*, 89:233–243, 1993.
- [Tho99] Wolfgang Thomas. Complementation of Büchi automata revisited. In *Jewels are Forever, Contributions on Theoretical Computer Science in Honor of Arto Salomaa*, pages 109–122. Springer-Verlag, 1999.
- [Var07] Moshe Y. Vardi. The Büchi complementation saga. In *Proceedings of the 24th Annual Symposium on Theoretical Aspects of Computer Science (STACS 2007), 22–24 February, Aachen, Germany*, volume 4393 of *Lecture Notes in Computer Science*, pages 12–22. Springer-Verlag, 2007.
- [Yan08] Qiqi Yan. Lower bounds for complementation of *omega*-automata via the full automata technique. *Journal of Logical Methods in Computer Science*, 4(1:5), 2008.