

From Complementation to Certification

Orna Kupferman
Hebrew University*

Moshe Y. Vardi
Rice University†

February 4, 2004

Abstract

In the automata-theoretic approach to model checking we check the emptiness of the product of a system S with an automaton $\mathcal{A}_{\neg\psi}$ for the complemented specification. This gives rise to two automata-theoretic problems: *complementation* of word automata, which is used in order to generate $\mathcal{A}_{\neg\psi}$, and the *emptiness* problem, to which model checking is reduced. Both problems have numerous other applications, and have been extensively studied for nondeterministic Büchi word automata (NBW). Nondeterministic *generalized* Büchi word automata (NGBW) have become popular in specification and verification and are now used in applications traditionally assigned to NBW. This is due to their richer acceptance condition, which leads to automata with fewer states and a simpler underlying structure.

In this paper we analyze runs of NGBW and use the analysis in order to describe a new complementation construction and a symbolic emptiness algorithm for NGBW. The complementation construction exponentially improves the best known construction for NGBW and is easy to implement. The emptiness algorithm is almost identical to a known variant of the Emerson-Lei algorithm, and our contribution is the strong relation we draw between the complementation construction and the emptiness algorithm – both naturally follow from the analysis of the runs, which easily implies their correctness. This relation leads to a new *certified* model-checking procedure, where a positive answer to the model-checking query is accompanied by a certificate whose correctness can be checked by methods independent of the model checker. Unlike certificates generated in previous works on certified model checking, our analysis enables us to generate a certificate that can be checked automatically and symbolically.

1 Introduction

In *model checking*, we check whether all the computations of a given system S satisfy a specification ψ . The system is usually given as a labeled state-transition graph and ψ is either a formula in LTL or a word automaton. In the automata-theoretic approach to model checking [Kur94b, VW94], one

*School of Engineering and Computer Science, Hebrew University, Jerusalem 91904, Israel Email: orna@cs.huji.ac.il, URL: <http://www.cs.huji.ac.il/~orna>. Supported in part by BSF grant 9800096, and by a grant from Minerva.

†Department of Computer Science, Houston, Rice University, TX 77251-1892, U.S.A. Email: vardi@cs.rice.edu, URL: <http://www.cs.rice.edu/~vardi>. Supported in part by NSF grants CCR-9988322, CCR-0124077, CCR-0311326, IIS-9908435, and IIS-9978135, by BSF grant 9800096, and by a grant from the Intel Corporation.

constructs an automaton $\mathcal{A}_{\neg\psi}$ for the negation of ψ and takes its product with S . The system S is correct with respect to ψ if this product is empty.

When ψ is given as an LTL formula, the construction of $\mathcal{A}_{\neg\psi}$ is relatively straightforward, we first negate ψ and then apply on $\neg\psi$ one of the many known translations of LTL formulas to word automata (cf. [VW94, GPVW95, SB00]). When the specification is given as an automaton \mathcal{A}_ψ , the task harder and one needs to complement the automaton. The product of $\mathcal{A}_{\neg\psi}$ and S can be viewed as a word automaton $\mathcal{A}_{S \times \neg\psi}$, and it is empty iff no computation of S violates ψ . Thus, the model-checking problem gives rise to two automata-theoretic problems: *complementation* of word automata, which is used in order to generate $\mathcal{A}_{\neg\psi}$ from \mathcal{A}_ψ , and the *emptiness* problem, to which model checking is reduced. Both problems have numerous other applications. First, refinement and optimization techniques that are based on language containment rather than simulation involve complementation and emptiness [Kur94a]. In addition, complementation is used in specification formalisms like ETL [Wol83, VW94], which have automata within the logic, and emptiness is used for satisfiability, planning and synthesis [Büc62, GV00, MW84].

Complementation and emptiness have been extensively studied for *nondeterministic Büchi word automata* (NBW, for short). The Büchi acceptance condition consists of a subset F of the state space, and a run of the automaton is accepting iff it visits F infinitely often. Consider an NBW \mathcal{A} with n states. In [Büc62], Büchi described a doubly-exponential complementation construction, which was improved in [SVW87] to a construction with $2^{O(n^2)}$ states. Only in [Saf88], Safra introduced an asymptotically optimal determinization construction, which also enabled a $2^{O(n \log n)}$ complementation construction, matching the known lower bound [Mic88]. Another $2^{O(n \log n)}$ construction was suggested in [Kla91], which circumvents the need for determinization. The optimal constructions in [Saf88, Kla91] are complicated, making their implementation very difficult [THB95]. In [KV01b], we suggested an optimal complementation construction that is based on alternating automata. This construction is considerably simpler, making it the first construction to be implemented [Mer00, GKS03]. The emptiness problem for NBW can be easily solved in linear time and NLOGSPACE [VW94]. The easy algorithms, however, are based on depth-first search, and cannot be implemented symbolically, which is very desirable in practice. Emerson and Lei’s algorithm for evaluation of μ -calculus formulas suggests a quadratic symbolic algorithm for the problem [EL86], and many variants of it have been suggested and studied (cf. [HTKB92, HKSV97, KPR98]). More involved algorithms with only $O(n \log n)$ [BGS00] and $O(n)$ [GPP03] symbolic steps are known too, but it is not clear that these algorithms are better in practice than the Emerson-Lei algorithm [RBS00].

The *generalized Büchi* acceptance condition consists of a set $\{F_1, \dots, F_k\}$ of subsets of the state space, and a run of the automaton is accepting iff it visits F_i infinitely often, for all $1 \leq i \leq k$. The number k of sets is the *index* of the automaton. The richer acceptance condition leads to automata with fewer states and simpler underlying structure. For example, the traditional translation of an LTL formula ψ to an NBW results in an automaton with state space $2^{cl(\psi)} \times 2^{cl(\psi)}$ [VW94]; the set $cl(\psi)$ is the set of ψ ’s subformulas and each state consists of a “local component”, which checks satisfaction of local requirements and an “eventuality component”, which checks satisfaction of eventualities. Using the generalized Büchi condition, it is easier to handle the different eventuality requirements, there is no need to the eventuality component, and the state space of the automaton is $2^{cl(\psi)}$ [GPVW95]. Nondeterministic generalized Büchi word automata (NGBW, for short) have become popular in specification and verification and are now used in applications traditionally assigned to NBW [Kur94b].

Once the NGBW is constructed, it is easy to translate it to an equivalent NBW, and then apply the known algorithms for NBW. For an NGBW with n states and index k , the constructed NBW has $O(nk)$ states [Cho74].

In this paper we analyze runs of NGBW, and use the analysis in order to suggest a new complementation construction and a symbolic emptiness algorithm for them. Recall that an NGBW \mathcal{A} rejects a word w if every run of \mathcal{A} has a set F_i in the acceptance condition that is visited only finitely often. The runs of \mathcal{A} can be arranged in a DAG (directed acyclic graph). We show that \mathcal{A} rejects w iff it is possible to label the vertices of the DAG by ranks so that some local conditions on the ranks of vertices and their successors are met. Intuitively, the ranks measure the distance from a position from which no states in F_i are visited.

The complementation construction that follows from the analysis results in an NBW with $2^{O(n \log nk)}$ states. This exponentially improves current complementation constructions, which first translate the NGBW into an NBW with $O(nk)$ states, and ends up in an NBW with $2^{O(nk \log nk)}$ states. Like the construction in [KV01b], our construction is simple and easy to implement. The extension of the reasoning in [KV01b] to NGBW is not trivial and was left open in [GKSV03]. (A trivial extension of [KV01b] to NGBW does exist, but results in an NBW with $2^{O(nk \log nk)}$ states. The technical achievement of the construction here is a simultaneous handling of all the sets in the acceptance condition, which is the key to the improved complexity.) The emptiness algorithm that follows from the analysis is almost identical to the OWCTY algorithm of [FFK⁺01] for symbolic detection of bad cycles. Our contribution is the strong relation we draw between the complementation construction and the emptiness algorithm – both naturally follow from the analysis of the runs, which easily implies their correctness.

Beyond the theoretical contribution of the relation between complementation and emptiness, it gives rise to a new *certified* model-checking procedure. As discussed in [Nam01, PPZ01, PZ01], it is desirable to accompany a positive answer of a model checker by a proof whose correctness can be verified by methods that are independent of the model checker. As in the case of proof-carrying codes (cf. [Nec97]), such a proof certifies that the systems was verified, and checking the certificate is much easier than the original verification task. In addition, as in the case of a counterexample that is returned to the user when model checking fails, the proof explains why model checking succeeds and leads to a better understanding of the system. For a discussion of other applications of certificates, see [Nam01].

Recall that model checking is reduced to checking the emptiness of the product $\mathcal{A}_{S \times \neg\psi}$ of the system S and the complemented specification $\mathcal{A}_{\neg\psi}$. We show that the ranks we associate with the vertices in the run DAG of $\mathcal{A}_{S \times \neg\psi}$ constitute a certificate that the product of S and $\mathcal{A}_{\neg\psi}$ is indeed empty. Moreover, by adding to our symbolic emptiness algorithm an algebraic decision diagram (ADD) that maintains the ranks, the certificate is generated symbolically (ADDs extend OBDDs by allowing the leaves to have values from arbitrary domains, thus they maintain functions that are not necessarily Boolean. Thus, while OBDDs represent Boolean functions, ADDs represent pseudo-Boolean functions [BFG+97].) Once the certificate is generated, it can be easily checked, automatically and symbolically, or manually, and it involves only local checks and no fixed points. Unlike the certificates in [Nam01, PPZ01, PZ01], whose goal is to provide the user with a deductive proof to ponder, our goal is to generate a compact certificate that can be verified automatically. Since a deductive proof usually consists of a long list of assertions, we believe that machine-checkable cer-

tificates are more appropriate in the verification of large systems. The generation of certificates that can be checked automatically is possible thanks to the analysis of runs, which bounds the domain of the well-founded sets that are used in the deductive certificates generated in previous works. As explained in Section 6, our method can generate, for users that are interested in a manual check, also “list-based” proofs.

2 Preliminaries

A *word automaton* is $\mathcal{A} = \langle \Sigma, Q, \delta, Q_{in}, \alpha \rangle$, where Σ is the input alphabet, Q is a finite set of states, $\delta : Q \times \Sigma \rightarrow 2^Q$ is a transition function, $Q_{in} \subseteq Q$ is a set of initial states, and α is an acceptance condition that defines a subset of Q^ω .

Given an input word $w = \sigma_0 \cdot \sigma_1 \cdot \dots$ in Σ^ω , a *run* of \mathcal{A} on w is a word $r = q_0, q_1, \dots$ in Q^ω such that $q_0 \in Q_{in}$ and for every $i \geq 0$, we have $q_{i+1} \in \delta(q_i, \sigma_i)$; i.e., the run starts in the initial state and obeys the transition function. Since the transition function may specify many possible transitions for each state and letter, \mathcal{A} may have several runs on w . A run is accepting iff it satisfies the acceptance condition α . We consider here the *generalized Büchi* acceptance condition, where $\alpha = \{F_1, \dots, F_k\}$ is a set of subsets of Q . The number k of sets is the *index* of α (or \mathcal{A}). For a run r , let $inf(r)$ denote the set of states that r visits infinitely often. That is,

$$inf(r) = \{q \in Q : q_i = q \text{ for infinitely many } i \geq 0\}.$$

As Q is finite, it is guaranteed that $inf(r) \neq \emptyset$. A run r is accepting iff $inf(r) \cap F_j \neq \emptyset$ for all $1 \leq j \leq k$. That is, r is accepting if every set in α is visited infinitely often. The *generalized co-Büchi* acceptance condition dualizes the generalized Büchi condition. Thus, again $\alpha = \{F_1, F_2, \dots, F_k\}$ is a set of subsets of Q , but a run r is accepting if $inf(r) \cap F_j = \emptyset$ for some $1 \leq j \leq k$. Thus, r visits some set in α only finitely often.

If the automaton \mathcal{A} is *nondeterministic*, then it accepts an input word w iff it has an accepting run on w . If \mathcal{A} is *universal*, then it accepts w iff all its runs on w are accepting. The *language* of \mathcal{A} , denoted $\mathcal{L}(\mathcal{A})$ is the set of words that \mathcal{A} accepts. Dualizing a nondeterministic generalized Büchi automaton (NGBW, for short) amounts to viewing it as a universal generalized co-Büchi automaton (UGCW, for short). It is easy to see that by dualizing \mathcal{A} , we get an automaton that accepts its complementary language. Note that nondeterministic Büchi automata (NBW, for short) are a special case of NGBW, with $k = 1$.

In the linear-time approach to model checking, we check whether all the computations of a given system S satisfy a specification ψ . The system is usually given as a labeled state-transition graph and ψ is either a formula in LTL or a word automaton (traditionally, NBW or NGBW). LTL formulas can be translated to word automata. The original translation in [VW94] uses NBW. More recent translations use NGBW. For example, it is shown in [GPVW95] that an LTL formula ψ of length m can be translated to an NGBW \mathcal{A}_ψ that accepts exactly all the words that satisfy ψ . The automaton \mathcal{A}_ψ has $2^{O(m)}$ states and index m .

3 Ranks for UGCW

Let $\mathcal{A} = \langle \Sigma, Q, Q_{in}, \delta, \alpha \rangle$ be a universal co-Büchi automaton with $\alpha = \{F_1, \dots, F_k\}$. Let $|Q| = n$. The runs of \mathcal{A} on a word $w = \sigma_0 \cdot \sigma_1 \cdots$ can be arranged in an infinite DAG (directed acyclic graph) $G_r = \langle V, E \rangle$, where

- $V \subseteq Q \times \mathbb{N}$ is such that $\langle q, l \rangle \in V$ iff some run of \mathcal{A} on w has $q_l = q$. For example, the first level of G_r contains the vertices $Q_{in} \times \{0\}$.
- $E \subseteq \bigcup_{l \geq 0} (Q \times \{l\}) \times (Q \times \{l+1\})$ is such that $E(\langle q, l \rangle, \langle q', l+1 \rangle)$ iff $\langle q, l \rangle \in V$ and $q' \in \delta(q, \sigma_l)$.

Thus, G_r embodies exactly all the runs of \mathcal{A} on w . We call G_r the *run DAG* of \mathcal{A} on w . For a set $F \subseteq Q$, we say that a vertex $\langle q, l \rangle$ in G_r is an F -vertex iff $q \in F$. We say that G_r is *accepting* if each path π in G_r has an index $1 \leq j \leq k$ such that π contains only finitely many F_j -vertices. It is easy to see that \mathcal{A} accepts w iff G_r is accepting.

Let $[2n]$ denote the set $\{0, 1, \dots, 2n\}$, and let $[2n]^{odd}$ and $[2n]^{even}$ denote the set of odd and even members of $[2n]$, respectively. Also, let $R = [2n]^{even} \cup ([2n]^{odd} \times \{1, \dots, k\})$, and \leq be the lexicographical order on the elements of R . We refer to the members of R in $[2n]^{even}$ as *even ranks* and refer to the members of R in $[2n]^{odd} \times \{j\}$ as *odd ranks with index j* .

A *ranking* for G_r is a function $f : V \rightarrow R$ that satisfies the following conditions:

1. For all vertices $\langle q, l \rangle \in V$, if $f(\langle q, l \rangle) = \langle 2i+1, j \rangle$, then $q \notin F_j$.
2. For all edges $\langle \langle q, l \rangle, \langle q', l+1 \rangle \rangle \in E$, we have $f(\langle q', l+1 \rangle) \leq f(\langle q, l \rangle)$.

Thus, a ranking associates with each vertex in G_r a rank in R so that ranks along paths decrease monotonically, and F_j -vertices cannot get an odd rank with index j . Note that each path in G_r eventually gets trapped in some rank. We say that the ranking f is an *odd ranking* if all the paths of G_r eventually get trapped in an odd rank. Formally, f is odd iff for all paths $\langle q_0, 0 \rangle, \langle q_1, 1 \rangle, \langle q_2, 2 \rangle, \dots$ in G_r , there is $l \geq 0$ such that $f(\langle q_l, l \rangle)$ is odd, and for all $l' \geq l$, we have $f(\langle q_{l'}, l' \rangle) = f(\langle q_l, l \rangle)$. Note that, equivalently, f is odd if every path of G_r has infinitely many vertices with odd ranks.

In the rest of this section we prove that G_r is accepting iff it has an odd ranking. Consider a (possibly finite) DAG $G \subseteq G_r$. We say that a vertex $\langle q, l \rangle$ is *finite* in G iff only finitely many vertices in G are reachable from $\langle q, l \rangle$. For a set $F \subseteq Q$, we say that a vertex $\langle q, l \rangle$ is F -free in G iff all the vertices in G that are reachable from $\langle q, l \rangle$ are not F -vertices. Note that, in particular, $\langle q, l \rangle$ is not an F -vertex.

We define an infinite sequence of DAGs $G_0 \supseteq G_1 \supseteq G_1^1 \supseteq \dots \supseteq G_1^k \supseteq G_3 \supseteq G_3^1 \supseteq \dots \supseteq G_3^k \supseteq G_5 \dots$ as follows. To simplify notations, we sometimes refer to G_{2i+1}^{k+1} as G_{2i+2} and refer to G_{2i+1}^1 as G_{2i+1}^1 . Thus, $G_1^1 = G_1$, $G_2 = G_1^{k+1}$, $G_3^1 = G_3$, $G_4 = G_3^{k+1}$, and so on.

- $G_0 = G_r$.
- $G_{2i+1} = G_{2i} \setminus \{\langle q, l \rangle \mid \langle q, l \rangle \text{ is finite in } G_{2i}\}$.
- $G_{2i+1}^{j+1} = G_{2i+1}^j \setminus \{\langle q, l \rangle \mid \langle q, l \rangle \text{ is } F_j\text{-free in } G_{2i+1}^j\}$, for $1 \leq j \leq k$.

Lemma 3.1 *For every $i \geq 0$, there exists l_i such that for all $l \geq l_i$, there are at most $n - i$ vertices of the form $\langle q, l \rangle$ in G_{2i} .*

Proof: We prove the lemma by an induction on i . The case where $i = 0$ follows from the definition of G_0 . Indeed, in G_r all levels $l \geq 0$ have at most n vertices of the form $\langle q, l \rangle$. Assume that the lemma's requirement holds for i , we prove it for $i + 1$. Consider the DAG G_{2i} . We distinguish between two cases. First, if G_{2i} is finite, then G_{2i+1} is empty, G_{2i+2} is empty as well, and we are done. Otherwise, we claim that there must be some F_j -free vertex in G_{2i+1}^j , for some $1 \leq j \leq k$. To see this, assume, by way of contradiction, that G_{2i} is infinite and all the vertices in G_{2i+1}^j are not F_j -free, for all $1 \leq j \leq k$. Note that then $G_{2i+1}^j = G_{2i+1}$ for all $1 \leq j \leq k$. Thus, all the vertices in G_{2i+1} are not F_j -free, for all $1 \leq j \leq k$. Since G_{2i} is infinite, G_{2i+1} is also infinite. Also, each vertex in G_{2i+1} has at least one successor. Consider some vertex $\langle q_0, l_0 \rangle$ in G_{2i+1} . Since, by the assumption, it is not F_1 -free, there exists an F_1 -vertex $\langle q'_0, l'_0 \rangle$ reachable from $\langle q_0, l_0 \rangle$. Let $\langle q_1, l_1 \rangle$ be a successor of $\langle q'_0, l'_0 \rangle$. By the assumption, $\langle q_1, l_1 \rangle$ is not F_2 -free. Hence, there exists an F_2 -vertex $\langle q'_1, l'_1 \rangle$ reachable from $\langle q_1, l_1 \rangle$. Let $\langle q_2, l_2 \rangle$ be a successor of $\langle q'_1, l'_1 \rangle$. By the assumption, $\langle q_2, l_2 \rangle$ is not F_3 -free. Thus, we can continue similarly and construct an infinite sequence of vertices $\langle q_h, l_h \rangle$ and $\langle q'_h, l'_h \rangle$ such that for all h , the vertex $\langle q'_h, l'_h \rangle$ is a $F_{(h \bmod k)+1}$ -vertex reachable from $\langle q_h, l_h \rangle$, and $\langle q_{h+1}, l_{h+1} \rangle$ is a successor of $\langle q'_h, l'_h \rangle$. Such a sequence, however, corresponds to a path in G_r that visits F_j infinitely often, for all $1 \leq j \leq k$, contradicting the assumption that G_r is accepting.

So, let j be the minimal index for which there is an F_j -free vertex in G_{2i+1}^j , and let $\langle q, l \rangle$ be such a vertex. By the minimality of j , we have that G_{2i+1}^j is equal to G_{2i+1} , and it contains no finite vertices. Hence, every F_j -free vertex in G_{2i+1}^j has a successor, which is also F_j -free, thus we can assume without loss of generality that $l \geq l_i$. We claim that taking $l_{i+1} = l$ satisfies the lemma's requirement. That is, we claim that for all $x \geq l$, there are at most $n - (i + 1)$ vertices of the form $\langle q, x \rangle$ in G_{2i+2} . Recall that $\langle q, l \rangle$ is not finite in G_{2i} . Thus, there are infinitely many vertices in G_{2i} that are reachable from $\langle q, l \rangle$. Hence, by König's Lemma, G_{2i} contains an infinite path $\langle q, l \rangle, \langle q_1, l + 1 \rangle, \langle q_2, l + 2 \rangle, \dots$. For all $x \geq 1$, the vertex $\langle q_x, l + x \rangle$ has infinitely many vertices reachable from it in G_{2i} and thus, it is not finite in G_{2i} . Therefore, the path $\langle q, l \rangle, \langle q_1, l + 1 \rangle, \langle q_2, l + 2 \rangle, \dots$ exists also in G_{2i+1} . Recall that $\langle q, l \rangle$ is F_j -free in G_{2i+1}^j . Hence, being reachable from $\langle q, l \rangle$, all the vertices $\langle q_x, l + x \rangle$ on the path are F_j -free as well. Therefore, they are not in G_{2i+1}^{j+1} . It follows that for all $x \geq l$, the number of vertices of the form $\langle q, x \rangle$ in G_{2i+1}^{j+1} (and hence also in G_{2i+2}) is strictly smaller than their number in G_{2i} . Hence, by the induction hypothesis, we are done. \square

Lemma 3.1 implies that G_{2n} is finite, and G_{2n+1} is empty.

Each vertex $\langle q, l \rangle$ in G_r has a unique $i \geq 1$ such that $\langle q, l \rangle$ is either finite in G_{2i} or F_j -free in G_{2i+1}^j , for some $1 \leq j \leq k$. This induces a function $f : V \rightarrow R$ defined as follows.

$$f(\langle q, l \rangle) = \begin{cases} 2i & \text{If } \langle q, l \rangle \text{ is finite in } G_{2i}. \\ \langle 2i + 1, j \rangle & \text{If } \langle q, l \rangle \text{ is } F_j\text{-free in } G_{2i+1}^j. \end{cases}$$

For an odd rank $\eta = \langle 2i + 1, j \rangle$, we refer to G_{2i+1}^j as G_η .

Lemma 3.2 *For every vertex $\langle q, l \rangle$ in G_r and $\eta \in R$, if $\langle q, l \rangle \notin G_\eta$, then $f(\langle q, l \rangle) < r$.*

Proof: We prove the lemma by an induction on η . Since $G_0 = G_r$, the case where $\eta = 0$ is immediate. For the induction step, we distinguish between two cases. For the case $\eta = \langle 2i + 1, j \rangle$, consider a vertex $\langle q, l \rangle \notin G_{2i+1}^{j+1}$. If $\langle q, l \rangle \notin G_{2i+1}^j$, the lemma's requirement follows from the induction hypothesis. If $\langle q, l \rangle \in G_{2i+1}^j$, then $\langle q, l \rangle$ is F_j -free in G_{2i+1}^j . Accordingly, $f(\langle q, l \rangle) = \langle 2i + 1, j \rangle$, meeting the lemma's requirement. For the case $\eta = 2i$, consider a vertex $\langle q, l \rangle \notin G_{2i+1}$. If $\langle q, l \rangle \notin G_{2i}$, the lemma's requirement follows from the induction hypothesis. If $\langle q, l \rangle \in G_{2i}$, then $\langle q, l \rangle$ is finite in G_{2i} . Accordingly, $f(\langle q, l \rangle) = 2i$, meeting the lemma's requirement. \square

Lemma 3.3 *For every two vertices $\langle q, l \rangle$ and $\langle q', l' \rangle$ in G_r , if $\langle q', l' \rangle$ is reachable from $\langle q, l \rangle$, then $f(\langle q', l' \rangle) \leq f(\langle q, l \rangle)$.*

Proof: We distinguish between two cases. If $f(\langle q, l \rangle) = 2i$, then $\langle q, l \rangle$ is finite in G_{2i} . Hence, either $\langle q', l' \rangle$ is not in G_{2i} , in which case, by Lemma 3.2, we have that $f(\langle q', l' \rangle) < 2i$, or $\langle q', l' \rangle$ is in G_{2i} , in which case, being reachable from $\langle q, l \rangle$, it must be finite in G_i , with $f(\langle q', l' \rangle) = 2i$, and we are done. If $f(\langle q, l \rangle) = \langle 2i + 1, j \rangle$, then $\langle q, l \rangle$ is F_j -free in G_{2i+1}^j . Hence, either $\langle q', l' \rangle$ is not in G_{2i+1}^j , in which case, by Lemma 3.2, we have that $f(\langle q', l' \rangle) < \langle 2i + 1, j \rangle$, or $\langle q', l' \rangle$ is in G_{2i+1}^j , in which case, being reachable from $\langle q, l \rangle$, it must be F_j -free in G_{2i+1}^j , in which case $f(\langle q', l' \rangle) = \langle 2i + 1, j \rangle$, and we are done. \square

Lemma 3.4 *For every infinite path in G_r , there exists an index $1 \leq j \leq k$ and a vertex $\langle q, l \rangle$ with an odd rank with index j such that all the vertices $\langle q', l' \rangle$ on the path that are reachable from $\langle q, l \rangle$ have $f(\langle q', l' \rangle) = f(\langle q, l \rangle)$.*

Proof: By Lemma 3.3, in every infinite path in G_r , there exists a vertex $\langle q, l \rangle$ such that all the vertices $\langle q', l' \rangle$ on the path that are reachable from $\langle q, l \rangle$ have $f(\langle q', l' \rangle) = f(\langle q, l \rangle)$. We need to prove that $f(\langle q, l \rangle)$ is odd. Assume, by way of contradiction, that $f(\langle q, l \rangle)$ is some even i . Thus, $\langle q, l \rangle$ is finite in G_i . Then, all the vertices $\langle q', l' \rangle$ on the path that are reachable from $\langle q, l \rangle$ also have $f(\langle q', l' \rangle) = i$. By Lemma 3.2, they all belong to G_i . Since the path is infinite, there are infinitely many such vertices, contradicting the fact that $\langle q, l \rangle$ is finite in G_i . \square

We can now conclude with Theorem 3.5 below.

Theorem 3.5 *G_r is accepting iff it has an odd ranking.*

Proof: Assume first that there is an odd ranking for G_r . Then, every path in G_r eventually gets trapped in some odd rank $\langle 2i + 1, j \rangle$. Hence, as F_j -vertices cannot get this rank, the path visits F_j only finitely often, and we are done.

For the other direction, note that Lemma 3.3, together with the fact that a vertex gets an odd rank with index j only if it is F_j -free, imply that the function f described above is a ranking. Lemma 3.4 then implies that the ranking is odd. \square

We note that the reasoning above is similar to the one described for co-Büchi automata in [KV01b]. The extension to the case of generalized co-Büchi is not trivial and involves a refinement of the DAG G_{2i+2} . In particular, the minimality of j in the proof of Lemma 3.1 is crucial for its correctness.

4 Complementation of NGBW

Theorem 3.5 implies that a UGCW \mathcal{A} accepts a word w iff there is an odd ranking for the run DAG G_r of \mathcal{A} on w – a ranking in which every infinite path in G_r has infinitely many vertices with an odd rank. Intuitively, the theorem suggests that the requirements imposed by the generalized co-Büchi condition (finitely often, for some set in α) can be reduced to a new condition of a simpler type (infinitely often, for vertices with an odd rank). Recall that by dualizing an NGBW, we get a UGCW for the complementary language. Theorem 3.5 enables us to translate this UGCW to an NBW, resulting in the complementation construction described below.

Theorem 4.1 *Let \mathcal{A} be an NGBW with n states and index k . There is an NBW \mathcal{A}' with $2^{O(n \log kn)}$ states such that $\mathcal{L}(\mathcal{A}') = \Sigma^\omega \setminus \mathcal{L}(\mathcal{A})$.*

Proof: Let $\tilde{\mathcal{A}}$ be the UGCW that dualizes \mathcal{A} . The UGCW $\tilde{\mathcal{A}}$ accepts exactly all words rejected by \mathcal{A} . We obtain \mathcal{A}' by translating $\tilde{\mathcal{A}}$ to an NBW. When \mathcal{A}' reads a word w , it guesses a ranking for the run DAG G_r of $\tilde{\mathcal{A}}$ on w . At a given point of a run of \mathcal{A}' , it keeps in its memory a whole level of G_r and a guess for the ranks of the vertices at this level. In order to check that the ranking is odd, \mathcal{A}' keeps track of states that owe a visit to vertices with odd ranks.

Before we define \mathcal{A}' , we first need some notations. A *level ranking* for \mathcal{A} is a function $g : Q \rightarrow R$, such that if $g(q)$ is odd with index j , then $q \notin F_j$. Let \mathcal{R} be the set of all level rankings¹. For $S \subseteq Q$ and a letter σ , let $\delta(S, \sigma) = \bigcup_{s \in S} \delta(s, \sigma)$. Note that if level $l-1$ in G_r contains the states in S , and the l -th letter in w is σ , then level l of G_r contains the states in $\delta(S, \sigma)$. For two level rankings g and g' in \mathcal{R} and a letter σ , we say that g' *covers* $\langle g, \sigma \rangle$ if for all q and q' in Q , if $q' \in \delta(q, \sigma)$, then $g'(q') \leq g(q)$. Thus, if g describes the ranks of the vertices of level $l-1$, and the l -th letter in w is σ , then g' is a possible level ranking for level l . Finally, for $g \in \mathcal{R}$, let $\text{odd}(g) = \{q : g(q) \in [2n]^{\text{odd}} \times \{1, \dots, k\}\}$. Thus, $\text{odd}(g)$ contains states to which g gives an odd rank.

Let $\mathcal{A} = \langle \Sigma, Q, q_{in}, \delta, \alpha \rangle$. Then $\mathcal{A}' = \langle \Sigma, Q', q'_{in}, \delta', \alpha' \rangle$, where

- $Q' = 2^Q \times 2^Q \times \mathcal{R}$. A state $\langle S, O, g \rangle \in Q'$ indicates that the current level of the DAG contains the states in S and the guessed level ranking for the current level is g . The set $O \subseteq S$ contains states along paths that have not visited a vertex with an odd rank since the last time O has been empty.
- $q'_{in} = \langle \{q'_{in}\}, \emptyset, f_{in} \rangle$, where $f_{in}(q) = 2n$ for all $q \in Q$.
- δ' is defined, for all $\langle S, O, g \rangle \in Q'$ and $\sigma \in \Sigma$, as follows.
 - If $O \neq \emptyset$, then $\delta'(\langle S, O, g \rangle, \sigma) = \{\langle \delta(S, \sigma), \delta(O, \sigma) \setminus \text{odd}(g'), g' \rangle : g' \text{ covers } \langle g, \sigma \rangle\}$.
 - If $O = \emptyset$, then $\delta'(\langle S, O, g \rangle, \sigma) = \{\langle \delta(S, \sigma), \delta(S, \sigma) \setminus \text{odd}(g'), g' \rangle : g' \text{ covers } \langle g, \sigma \rangle\}$.
- $\alpha' = 2^Q \times \{\emptyset\} \times \mathcal{R}$.

¹When we refer to level rankings, we only care for the ranks of a subset of Q (the set of states that appear in the corresponding level). For technical convenience, we let g range on all states.

Thus, when \mathcal{A}^l reads the l -th letter in the input, it guesses the level ranking for level l in the run DAG. This level ranking should cover the level ranking of level $l - 1$. In addition, in the O component, \mathcal{A}^l keeps track of states along paths that owe a visit to a vertex with an odd rank. When all the paths of the DAG have visited a vertex with an odd rank, the set O becomes empty, and is initiated according to the states in the current level and its ranking. The acceptance condition then checks that there are infinitely many levels in which O become empty.

Since there are at most $(k(2n + 1))^n$ level rankings, the number of states in \mathcal{A}^l is at most $2^{2n} \cdot (k(2n + 1))^n = 2^{O(n \log kn)}$. \square

Note that the previous complementation construction for NGBW involves a $2^{O(nk \log nk)}$ blow up, as they first translate the NGBW into an NBW with $O(nk)$ states, and complementing an NBW with m states results in an NBW with $2^{O(m \log m)}$ states [Saf88, Mic88]. Thus, our construction exponentially improves the previous construction.

5 Model Checking

Recall that the model-checking problem is reduced to the emptiness problem of an NGBW $\mathcal{A}_{S \times \neg \psi}$ over a single-letter alphabet. Equivalently, we can check the nonemptiness of the UGCW $\tilde{\mathcal{A}}_{S \times \neg \psi}$ that dualizes $\mathcal{A}_{S \times \neg \psi}$. Indeed, since $\mathcal{A}_{S \times \neg \psi}$ has a single-letter alphabet, it is empty iff $\tilde{\mathcal{A}}_{S \times \neg \psi}$ is not empty (see also [MP87]).

In this section we describe a symbolic algorithm for UGCW non-emptiness. The algorithm is induced by the analysis in Section 3, and its correctness follows immediately from Theorem 3.5. The algorithm is a variant of Emerson-Lei algorithm and is similar to the OWCTY algorithm of [FFK⁺01] for detecting bad-cycles (see also [HTKB92, KPR98]). The tight relation between the complementation construction and the emptiness procedure is very interesting, as it shows that progress in the emptiness procedure can be measured by means of the ranks that are used in the construction of the complementary automaton. As we describe in the next section, this observation gives rise to a new *certified* model-checking procedure.

Consider a single-letter UGCW $\mathcal{A} = \langle \{a\}, Q, \delta, Q_{in}, \alpha \rangle$. The analysis in Section 3 associates ranks with members of the infinite set $Q \times \mathbb{N}$. On the other hand, nonemptiness algorithms handle the finite state set Q . Accordingly, we first associate ranks with states:

Lemma 5.1 *Consider a UGCW \mathcal{A} over a single-letter alphabet. Then, for every state q of \mathcal{A} , all the vertices in $\{q\} \times \mathbb{N}$ have the same rank.*

Proof: Consider a state q and two levels l_1 and l_2 such that $\langle q, l_1 \rangle$ and $\langle q, l_2 \rangle$ are vertices in G_r . Recall that the DAG G_r embodies all the runs of \mathcal{A} on the input word. Since \mathcal{A} is a single-letter UGCW, the sub-DAG with root $\langle q, l_1 \rangle$ coincides with the sub-DAG with root $\langle q, l_2 \rangle$. Indeed, both embody exactly all the runs of \mathcal{A} with initial state q on a^ω . Thus, all the sub-DAGs of G_r with roots in $\{q\} \times \mathbb{N}$ coincide. Hence, it is easy to prove by an induction on $r \in R$ that for all states q and $r \in R$, either all vertices in $\{q\} \times \mathbb{N}$ get rank r , or no vertex in $\{q\} \times \mathbb{N}$ gets rank r . Indeed, again by an induction on i and j , for each DAG G_{2i} or G_{2i+1}^j in the sequence of DAGs, either all vertices in $\{q\} \times \mathbb{N}$ are in the DAG, or no vertex is in it. \square

For a state q of \mathcal{A} , the *rank of q* , denoted $\text{rank}(q)$, is the rank of the vertices $\{q\} \times \mathbb{N}$ in G_r . We are now ready to describe the nonemptiness procedure that follows. The procedure, described in Figure 1, gets as input the UGCW \mathcal{A} and calculates the set b of all the states q such that \mathcal{A} with initial state q is empty. The UGCW \mathcal{A} is then not empty iff $Q_{in} \cap b = \emptyset$. The algorithm uses the following set-based operations (all easily implemented by means of OBDDs).

- The operator $pre : 2^Q \rightarrow 2^Q$. Given a set of states γ , the set $pre(\gamma)$ contains all states that have an immediate successor in γ . Formally, $q \in pre(\gamma)$ iff $\delta(q, a) \cap \gamma \neq \emptyset$ (in temporal logic, $q \models EX\gamma$).
- The operator $until : 2^Q \times 2^Q \rightarrow 2^Q$. Given two sets of states η and γ , the set $until(\eta, \gamma)$ contains all states that reach a state in $\gamma \cap \eta$ via states in η . Formally, $q \in until(\eta, \gamma)$ iff there are q_0, \dots, q_l such that $q_0 = q$, for all $0 \leq i < l$, we have that $q_{i+1} \in \delta(q_i, a)$ and $q_i \in \eta$, and $q_l \in \eta \cap \gamma$ (in temporal logic, $q \models E\eta U(\eta \wedge \gamma)$). Note that the operator $until$ can be implemented by repeatedly applying the pre and intersection operators, until a fixpoint is reached.

```

procedure NonEmpty( $\mathcal{A}$ )
   $b := Q$ ;
  while  $b$  changes do
    while  $b$  changes do  $b := b \cap pre(b)$ ;
    for  $j = 1 \dots k$  do  $b := b \cap until(b, F_j)$ ;
  if  $Q_{in} \cap b = \emptyset$  then return("not empty") else return("empty");

```

Figure 1: A nonemptiness procedure.

Note that the set b is monotonically decreasing during the execution of the procedure NonEmpty. Intuitively, b contains all states that have not yet been ranked. At initialization, b contains all the states, and in each iteration it is intersected with some set. We say that a state q is removed from b in iteration $i[0]$ if q is removed from b during the internal while loop of the i -th external while loop. We say that q is removed from b in iteration $(i[1], j)$ if q is removed from b during the j -th internal for loop of the i -th external while loop. Lemma 5.2 then follows directly from the definition of ranks.

Lemma 5.2 *Consider a state q in \mathcal{A} .*

- *The state q is removed from b in iteration $i[0]$ iff $\text{rank}(q) = 2i$.*
- *The state q is removed from b in iteration $(i[1], j)$ iff $\text{rank}(q) = \langle 2i + 1, j \rangle$.*

By Lemma 5.2, a state q is removed from b during the execution of the procedure if it has a well-defined rank, which holds, by Theorem 3.5, if \mathcal{A} with initial state q is not empty. Thus, Lemma 5.2, together with the analysis in Section 3, naturally induce the algorithm and immediately imply its correctness. (The only, minor, difference between our algorithm and the OWCTY algorithm [FFK⁺01], is that in our algorithm the internal while loop precedes the internal for loop, rather than the other way around. Since the purpose of the internal while loop is to eliminate quickly states that cannot be on a cycle, it makes sense to apply it as soon as possible.)

Remark 5.3 When ψ is a safety property, the automaton \mathcal{A}_ψ is a looping automaton (all infinite runs are accepting), and the automaton $\mathcal{A}_{\neg\psi}$ can be defined as a nondeterministic automaton on finite words [Sis94, KV01a]. Thus, a runs of $\mathcal{A}_{\neg\psi}$ is accepting iff it reaches a set F of accepting states. Accordingly (assuming that the system has no fairness conditions), the automaton $\tilde{\mathcal{A}}_{S \times \neg\psi}$ is a universal automaton in which all runs except these that reach F are accepting. As a result, we need a much simpler nonemptiness procedure, which corresponds to backwards traversal. Thus, it initializes b with $Q \setminus F$, follows with the single while loop **while** b changes **do** $b := b \cap \text{pre}(b)$, and returns “not empty” when $Q_{in} \subseteq b$. \square

6 Certified Model Checking

Recall that $\tilde{\mathcal{A}}_{S \times \neg\psi}$ with initial state q is not empty iff $\text{rank}(q)$ is well defined, and hence belongs to R . Thus, beyond a correctness proof, the analysis in Section 3 can be used in order to accompany the output of the procedure described in Figure 1 by a certificate, namely the odd ranking, that $\tilde{\mathcal{A}}_{S \times \neg\psi}$ is indeed not empty, and S satisfies ψ . In this section we describe how to generate and check such a certificate.

As we showed, a function $f : V \rightarrow R$ is an odd ranking if F_j -vertices do not get an odd rank with index j , ranks along paths decrease monotonically, and all the paths of G_r eventually get trapped in an odd rank. The number of vertices along a path that get an even rank depends on the input word and is in general unbounded. Accordingly, checking whether a given function f is an odd ranking involves, in addition to local checks, also a check for eventualities, which involves a fixed-point computation. We now show that when \mathcal{A} is a single-letter automaton, as is the case with $\tilde{\mathcal{A}}_{S \times \neg\psi}$, it is possible to bound the number of vertices that get even ranks. Let $\mathcal{A} = \langle \{a\}, Q, \delta, Q_{in}, \alpha \rangle$ be a single-letter UGCW. Consider a vertex $\langle q, l \rangle$ that is finite in G_{2i} . Let $\text{height}(q, l)$ be the length of the longest path from $\langle q, l \rangle$ to a leaf of G_{2i} .

Lemma 6.1 *For all vertices $\langle q, l \rangle$, we have that $\text{height}(q, l) \in \{0, \dots, n - 1\}$.*

Proof: Assume by way of contradiction that G_{2i} contains a vertex $\langle q, l \rangle$ such that $\text{height}(q, l) \geq n$. Then, the longest path from $\langle q, l \rangle$ to a leaf of G_{2i} contains at least one state q' that repeats at least twice. Thus, there is a path in G_{2i} that starts in $\langle q, l \rangle$, reaches a vertex $\langle q', l_1 \rangle$ with $l_1 \geq l$ and continues to a vertex $\langle q', l_2 \rangle$ with $l_2 > l_1$. As argued in the proof of Lemma 5.1, the sub-DAG with root $\langle q', l_1 \rangle$ coincides with the sub-DAG with root $\langle q', l_2 \rangle$. Hence, there is a path in G_{2i} that starts in $\langle q', l_2 \rangle$ and reaches a vertex $\langle q', l_3 \rangle$ with $l_3 > l_2$. By repeating this argument, we obtain an infinite path in G_{2i} that starts in $\langle q, l \rangle$, contradicting the fact that $\langle q, l \rangle$ is finite. \square

Following Lemma 6.1, we refine our set of ranks to $R = ([2n]^{\text{even}} \times \{1, \dots, n\}) \cup ([2n]^{\text{odd}} \times \{1, \dots, k\})$. We say that a function $f : V \rightarrow R$ is a *bounded odd ranking* if the following hold.

1. For all vertices $\langle q, l \rangle \in V$, if $f(\langle q, l \rangle) = \langle 2i + 1, j \rangle$, then $q \notin F_j$.
2. Consider an edge $\langle \langle q, l \rangle, \langle q', l + 1 \rangle \rangle \in E$.
 - (a) If $f(\langle q, l \rangle)$ is odd, then $f(\langle q', l + 1 \rangle) \leq f(\langle q, l \rangle)$.

(b) If $f(\langle q, l \rangle)$ is even, then $f(\langle q', l + 1 \rangle) < f(\langle q, l \rangle)$.

Thus, in a bounded odd ranking, the rank of successors of a vertex $\langle q, l \rangle$ with an even rank must be strictly smaller than the rank of $\langle q, l \rangle$.

Theorem 6.2 *Let \mathcal{A} be a single-letter automaton. Then, G_r is accepting iff it has a bounded odd ranking.*

Proof: Each bounded odd ranking is also an odd ranking (with the height component being ignored). Thus, the direction from right to left follows from Theorem 3.5. For the other direction, we refine the function $rank : V \rightarrow R$ to account for heights of vertices. Thus, $rank(q, l)$, for $\langle q, l \rangle$ that is finite with height h in G_{2i} is $\langle 2i, h \rangle$. It is easy to see that, as with odd ranking, the first two conditions on rank being a bounded odd ranking hold. For the third condition, consider a vertex $\langle q, l \rangle$ with rank $\langle 2i, h \rangle$. By the definition of height, the successors of $\langle q, l \rangle$ in G_{2i} have heights that are strictly smaller than h . By Lemma 3.2, the successors of $\langle q, l \rangle$ that are not in G_{2i} have rank that is strictly smaller than $2i$. Thus, all the successors of $\langle q, l \rangle$ have ranks that are strictly smaller than $\langle 2i, h \rangle$, and we are done. \square

It turns out that the nonemptiness procedure actually accounts for heights too: we say that a state q is removed from b in iteration $(i[0], h)$ if q is removed from b during the i -th external while loop and its h -th internal while loop (we start to count iterations from 0). Lemma 6.3 then follows directly from the definition of ranks.

Lemma 6.3 *Consider a state q in \mathcal{A} .*

- *The state q is removed from b in iteration $(i[0], h)$ iff $rank(q) = \langle 2i, h \rangle$.*
- *The state q is removed from b in iteration $(i[1], j)$ iff $rank(q) = \langle 2i + 1, j \rangle$.*

In a symbolic implementation of the procedure NonEmpty, we maintain b in an OBDD. By maintaining in addition an ADD that maps states to ranks, we generate a certificate that can be used to certify the model-checking procedure. Let $f : Q \rightarrow R$ be a partial function from Q to R . The procedure Certified_Nonempty described in Figure 2 gets as input a single letter UGCW \mathcal{A} and calculates, in addition to the set b , also a function f that describes the odd ranking, which is returned in case no state of Q_{in} is in b^2 . The procedure uses the operator *assign*, which given a set $\gamma \subseteq Q$ and a rank $r \in R$, returns a function in which all the states in γ are assigned r . In addition, initializing f to \emptyset corresponds to an empty function, and \cup between two functions with disjoint domains returns their union.

The procedure can be easily implemented symbolically, with f being maintained in an ADD. Note that R consists of pairs, thus in some ADD implementations, where the domain of the ADD is restricted to single values, we have to encode R , which is easy.

²In case the intersection of Q_{in} and b is not empty, it is possible to enhance the procedure to return an evidence to the emptiness of \mathcal{A} ; this is similar to the known generation of counterexamples and we do not discuss it here.

```

procedure Certified_NonEmpty( $\mathcal{A}$ )
   $b := Q; f := \emptyset; i := -1;$ 
  while  $b$  changes do
     $i := i + 1;$ 
     $h := 0;$ 
    while  $b$  changes do
       $f := f \cup \text{assign}(b \setminus \text{pre}(b), (2i, h));$ 
       $b := b \cap \text{pre}(b);$ 
       $h := h + 1;$ 
    for  $j = 1 \dots k$  do
       $f := f \cup \text{assign}(b \setminus \text{until}(b, F_j), (2i + 1, j));$ 
       $b := b \cap \text{until}(b, F_j);$ 
  if  $Q_{in} \cap b = \emptyset$  then return (“not empty with certificate”  $f$ ) else
  return (“empty”);

```

Figure 2: A nonemptiness procedure that generates a certificate.

Once the procedure `Certified_NonEmpty` terminates and f is returned to the user, she can check that f represent a bounded odd ranking and that all the states in Q_{in} have a rank. Note how the use of heights, which enables us to consider bounded odd rankings, is essential here, as checking f involves only local checks (a comparison of a rank of vertices and their successors) and no fixed points. As described in the procedure `Check_Certificate` in Figure 3, the check can be done automatically and symbolically. The procedure uses the following Boolean functions.

- $undef : 2^Q \rightarrow \{\mathbf{true}, \mathbf{false}\}$. Given a set γ of states, $undef(\gamma)$ is true iff f does not assign a rank to some state in γ ; i.e., $\gamma \cap \text{comp}(f^{-1}(R)) \neq \emptyset$.
- $oops : \{1, \dots, k\} \times 2^Q \rightarrow \{\mathbf{true}, \mathbf{false}\}$. Given an index $1 \leq j \leq k$ and a set γ of states, $oops(j, \gamma)$ is true iff there is $q \in \gamma$ with an odd rank with index j ; i.e., $f(q) \in [2n]^{odd} \times \{j\}$. The symbolic implementation of $oops$ checks whether the intersection of $f^{-1}([2n]^{odd} \times \{j\})$ and b is empty.

The correctness of the procedure `Check_Certificate` follows immediately from Theorem 6.2.

In the case of LTL model checking, the automaton \mathcal{A} is $\tilde{\mathcal{A}}_{S \times \neg\psi}$ and its state space consists of pairs $\langle s, P \rangle$, where s is a state of S and $P \subseteq \text{cl}(\psi)$ is a set of LTL formulas. The automaton $\tilde{\mathcal{A}}_{S \times \neg\psi}$ with initial state $\langle s, P \rangle$ is not empty if each path that starts in s violates at least one formula in P . Each set in the generalized Büchi acceptance condition corresponds to a formula of the form $\varphi U \theta$ and consists of all the states $\langle s, P \rangle$ in which P contains θ or does not contain $\varphi U \theta$. The rank of a state $\langle s, P \rangle$ explains how one of the formulas in P is not satisfied. If the rank is even, the explanation is transferred, via local conditions, to the successors of s . If the rank is odd, the particular formula that is not satisfied is recorded (by means of the index of the odd rank). Hence, for users that prefer to get a certificate that is similar to deductive proofs generated by proof-theoretic approaches to verification [MP92] (as in [PPZ01, PZ01]), it is possible to present the certificate as a list of states and how they satisfy (that is, do not satisfy the negation of) relevant subformulas of ψ .

```

procedure Check_Certificate( $\mathcal{A}, f$ )
  if  $undef(Q_{in})$  then return (“incorrect certificate”);
  for  $j = 1 \dots k$  do
    if  $oops(j, F_j)$  then return (“incorrect certificate”);
  for all  $q \in Q$  and  $q' \in \delta(q, a)$  do
    if  $f(q)$  is odd and  $f(q') > f(q)$  then return (“incorrect certificate”);
    if  $f(q)$  is even and  $f(q') \geq f(q)$  then return (“incorrect certificate”);
  return (“correct certificate”);

```

Figure 3: Verifying that a certificate is correct.

Remark 6.4 As discussed in Remark 5.3, when ψ is a safety property, nonemptiness of $\tilde{\mathcal{A}}_{S \times \neg\psi}$ can be checked by backwards traversal and the nonemptiness procedure consists of the single loop **while** b changes **do** $b := b \cap pre(b)$. In this case, we can take b itself as the certificate, and there is no need to compute ranks. To check that b is a correct certificate, one has to check that $Q_{in} \subseteq b$, $b \cap F = \emptyset$, and $b = b \cap pre(b)$. Thus, while the computation of b involves a fixed-point, checking that it is indeed a fixed-point involves only local checks. \square

7 Discussion

We suggested a new complementation construction for NGBW. The analysis behind the construction led to a symbolic certified model-checking procedure. For an NGBW with n states and index k , our complementation results in an NBW with $2^{O(n \log nk)}$ states. This exponentially improves current complementation construction that first translate the NGBW into an NBW with $O(nk)$ states, and end up in an complementary automaton with $2^{O(nk \log nk)}$ states. Using our analysis it can also be shown that an alternating generalized co-Büchi word automaton with n states and index k can be translated to an alternating weak word automaton with $O(n^2k)$ states, rather than $O(n^2k^2)$, which would be the result of a translation that first translates the automaton into an alternating co-Büchi automaton. These improvements suggest that when it is possible to use both NGBW and NBW, one should prefer an NGBW with fewer states, even if its index is large. The above give rise to the following problem:

Given an NBW, find an equivalent NGBW with fewer states.

In particular, it'd be interesting to look for a variant of the translation in [MH84], of an alternating Büchi word automaton with state space Q into an NBW with state space $2^Q \times 2^Q$, that will end up in an NGBW with state space 2^Q .

Another open problem refers to the particular nonemptiness algorithm we were able to relate to the complementation construction and to augment with a certificate. Recall that our algorithm is similar to the OWCTY algorithm, which is a variant of the quadratic Emerson-Lei algorithm. As discussed in Section 1, more recent algorithms solve the nonemptiness problem with a sub-quadratic

number of symbolic steps [BGS00, GPP03]. It'd be interesting to consider whether these algorithms can be related to complementation as well.

References

- [BFG+97] R. Bahar, E. Frohm, C. Gaona, G. Hachtel, E. Macii, A. Pardo, and F. Somenzi. Algebraic decision diagrams and their applications. *FMSD* 10(2/3):171–206, 1997.
- [BGS00] R. Bloem, H.N. Gabow, and F. Somenzi. An algorithm for strongly connected component analysis in $n \log n$ symbolic steps. In Proc. *FMCAD*, LNCS 1954, pages 37–54, 2000.
- [Büc62] J.R. Büchi. On a decision method in restricted second order arithmetic. In Proc. *Internat. Congr. Logic, Method. and Philos. Sci. 1960*, pages 1–12, Stanford, 1962. Stanford University Press.
- [Cho74] Y. Choueka. Theories of automata on ω -tapes: A simplified approach. *Journal of Computer and System Sciences*, 8:117–141, 1974.
- [EL86] E.A. Emerson and C.-L. Lei. Efficient model checking in fragments of the propositional μ -calculus. In Proc. *1st LICS*, pages 267–278, 1986.
- [FFK⁺01] K. Fisler, R. Fraer, G. Kamhi, M.Y. Vardi, and Z. Yang. Is there a best symbolic cycle-detection algorithm? In Proc. *7th TACAS*, LNCS 2031, pages 420–434, 2001.
- [GKSV03] S. Gurumurthy, O. Kupferman, F. Somenzi, and M.Y. Vardi. On complementing nondeterministic Büchi automata. In Proc. *12th CHARME*, LNCS, 2003.
- [GPP03] R. Gentilini, C. Piazza, and A. Policriti. Computing strongly connected components in a linear number of symbolic steps. In Proc. *14th SODA*, pages 573–582, 2003.
- [GPVW95] R. Gerth, D. Peled, M.Y. Vardi, and P. Wolper. Simple on-the-fly automatic verification of linear temporal logic. In *Protocol Specification, Testing, and Verification*, pages 3–18. Chapman & Hall, August 1995.
- [GV00] G. De Giacomo and M.Y. Vardi. Automata-theoretic approach to planning for temporally extended goals. In Proc. *5th ECP*, LNAI 1809, pages 226–238, 2000.
- [HKSV97] R.H. Hardin, R.P. Kurshan, S.K. Shukla, and M.Y. Vardi. A new heuristic for bad cycle detection using BDDs. In Proc. *9th CAV*, LNCS 1254, pages 268–278, 1997.
- [HTKB92] R. Hojati, H. Touati, R. Kurshan, and R. Brayton. Efficient ω -regula language containment. In Proc. *4th CAV*, LNCS 663, 1992.
- [Kla91] N. Klarlund. Progress measures for complementation of ω -automata with applications to temporal logic. In Proc. *32nd FOCS*, pages 358–367, 1991.
- [KPR98] Y. Kesten, A. Pnueli, and L. Raviv. Algorithmic verification of linear temporal logic specifications. In Proc. *25th ICALP*, LNCS 1443, pages 1–16, 1998.
- [Kur94a] R.P. Kurshan. The complexity of verification. In *26th STOC*, pages 365–371, 1994.
- [Kur94b] R.P. Kurshan. *Computer Aided Verification of Coordinating Processes*. Princeton Univ. Press, 1994.
- [KV01a] O. Kupferman and M.Y. Vardi. Model checking of safety properties. *Formal methods in System Design*, 19(3):291–314, November 2001.
- [KV01b] O. Kupferman and M.Y. Vardi. Weak alternating automata are not that weak. *ACM Trans. on Computational Logic*, 2001(2):408–429, July 2001.

- [Mer00] S. Merz. Weak alternating automata in Isabelle/HOL. In *Theorem Proving in Higher Order Logics: 13th International Conference*, LNCS 1869, pages 423–440, 2000.
- [MH84] S. Miyano and T. Hayashi. Alternating finite automata on ω -words. *Theoretical Computer Science*, 32:321–330, 1984.
- [Mic88] M. Michel. Complementation is more difficult with automata on infinite words. CNET, Paris, 1988.
- [MP87] Z. Manna and A. Pnueli. Specification and verification of concurrent programs by \forall -automata. In *Proc. Proc. 14th POPL*, pages 1–12, 1987.
- [MP92] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Berlin, January 1992.
- [MW84] Z. Manna and P. Wolper. Synthesis of communicating processes from temporal logic specifications. *ACM TOPLAS*, 6(1):68–93, January 1984.
- [Nam01] K.S. Namjoshi. Certifying model checkers. In *Proc. 13th CAV*, LNCS 2102, pages 2–13, 2001.
- [Nec97] G.C. Necula. Proof-carrying code. In *Proc. 24th POPL*, pages 106–119, 1997.
- [PPZ01] D. Peled, A. Pnueli, and L.D. Zuck. From falsification to verification. In *Proc. 21th FST&TCS*, LNCS 2245, pages 292–304, 2001.
- [PZ01] D. Peled and L.D. Zuck. From model checking to a temporal proof. In *Proc. 8th SPIN Workshop on Model Checking of Software*, LNCS 2057, pages 1–14, 2001.
- [RBS00] K. Ravi, R. Bloem, and F. Somenzi. A comparative study of symbolic algorithms for the computation of fair cycles. In *Proc. FMCAD*, LNCS 1954, pages 143–160, 2000.
- [Saf88] S. Safra. On the complexity of ω -automata. In *29th FOCS*, pages 319–327, 1988.
- [SB00] F. Somenzi and R. Bloem. Efficient Büchi automata from LTL formulae. In *Proc. 12th CAV*, LNCS 1855, pages 248–263, 2000.
- [Sis94] A.P. Sistla. Safety, liveness and fairness in temporal logic. *Formal Aspects of Computing*, 6:495–511, 1994.
- [SVW87] A.P. Sistla, M.Y. Vardi, and P. Wolper. The complementation problem for Büchi automata with applications to temporal logic. *Theoretical Computer Science*, 49:217–237, 1987.
- [THB95] S. Tasiran, R. Hojati, and R.K. Brayton. Language containment using non-deterministic ω -automata. In *Proc. 8th CHARME*, LNCS 987, pages 261–277, 1995.
- [VW94] M.Y. Vardi and P. Wolper. Reasoning about infinite computations. *Information and Computation*, 115(1):1–37, November 1994.
- [Wol83] P. Wolper. Temporal logic can be more expressive. *Information and Control*, 56(1–2):72–99, 1983.