Tightening the Exchange Rates Between Automata

Orna Kupferman

School of Computer Science and Engineering, Hebrew University,
Jerusalem 91904, Israel.

Abstract. Automata on infinite objects were the key to the solution of several fundamental decision problems in mathematics and logic. Today, automata on infinite objects are used for formal specification and verification of reactive systems. The practical importance of automata in formal methods has motivated a re-examination of the blow up that translations among different types of automata involve. For most translations, the situation is satisfying, in the sense that even if there is a gap between the upper and the lower bound, it is small. For some highly practical cases, however, the gap between the upper and the lower bound is exponential or even larger. The article surveys several such frustrating cases, studies features that they share, and describes recent efforts (with partial success) to close the gaps.

1 Introduction

Finite *automata on infinite objects* were first introduced in the 60's, and were the key to the solution of several fundamental decision problems in mathematics and logic [3, 25, 31]. Today, automata on infinite objects are used for *specification* and *verification* of reactive systems. The automata-theoretic approach to verification reduces questions about systems and their specifications to questions about automata [19, 41]. Recent industrial-strength property-specification languages such as Sugar, ForSpec, and the recent standard PSL 1.01 [7] include regular expressions and/or automata, making specification and verification tools that are based on automata even more essential and popular.

Early automata-based algorithms aimed at showing decidability. The complexity of the algorithm was not of much interest. For example, the fundamental automata-based algorithms of Büchi and Rabin, for the decidability of S1S and SnS (the monadic second-order theories of infinite words and trees, respectively) [3, 31] are of non-elementary complexity (i.e., the complexity can not be bounded by a stack of exponentials of a fixed height [26]). Proving the decidability of a given logic was then done by translating the logic to a monadic second-order theory, ignoring the fact that a direct algorithm could have been more efficient. Things have changed in the early 80's, when decidability of highly expressive logics became of practical importance in areas such as artificial intelligence and formal reasoning about systems. The change was reflected in the development

of two research directions: (1) direct and efficient translations of logics to automata [43, 37, 40], and (2) improved algorithms and constructions for automata on infinite objects [33, 10, 30].

Both research directions are relevant not only for solving the decidability problem, but also for solving other basic problems in formal methods, such as model checking [5] and synthesis [30]. Moreover, input from the industry continuously brings to the field new problems and challenges, requiring the development of new translations and algorithms.[1] For many problems and constructions, our community was able to come up with satisfactory solutions, in the sense that the upper bound (the complexity of the best algorithm or the blow-up in the best known construction) coincides with the lower bound (the complexity class in which the problem is hard, or the blow-up that is known to be unavoidable). For some problems and constructions, however, the gap between the upper bound and the lower bound is significant. This situation is especially frustrating, as it implies that not only we may be using algorithms that can be significantly improved, but also that something is missing in our understanding of automata on infinite objects.

Before turning to the frustrating cases, let us first describe one "success story" — the complementation construction for nondeterministic Büchi automata on infinite words (NBWs). Translating S1S into NBWs, Büchi had to prove the closure of NBWs under complementation. For that, Büchi suggested in 1962 a doubly-exponential construction. Thus, starting with an NBW with n states, the complementary automaton had $2^{2^{O(n)}}$ states [3]. The lower bound known then for NBW complementation was 2^n, which followed from the complementation of automata on finite words. Motivated by problems in formal methods, Sistla, Vardi, and Wolper developed in 1985 a better complementation construction with only a $2^{O(n^2)}$ blow-up [36]. Only in 1988, Safra introduced a determinization construction for NBWs that enabled a $2^{O(n \log n)}$ complementation construction [33], and Michel proved a matching lower bound [28]. The story, however, was not over. A careful analysis of the lower and upper bounds reveals an exponential gap hiding in the constants of the $O()$ notations. While the upper bound of Safra is n^{2n}, the lower bound of Michel is only $n!$, which is roughly $(n/e)^n$. Only recently, a new complementation construction, which avoids determinization, has led to an improved upper bound of $(0.97n)^n$ [11], and a new concept, of full automata, has led to an improved lower bound of $(0.76n)^n$ [44]. Thus, a gap still exists, but it is an acceptable one, and it probably does not point to a significant gap in our understanding of nondeterministic Büchi automata.

In the article, we survey two representative problems for which the gap between the upper and the lower bound is still exponential. In Section 3, we consider safety properties and the problem of translating safety properties to non-

[1] In fact, the practical importance of automata has lead to a reality in which the complexity of a solution or a construction is only one factor in measuring its quality. Other measures, such as the feasibility of a symbolic implementation or the effectiveness of optimizations and heuristics in the average case are taken into an account too. In this article, however, we only consider worst-case complexity.

deterministic automata on finite words. In Section 4, we consider the problem of translating nondeterministic Büchi word automata to nondeterministic co-Büchi word automata. Both problems have strong practical motivation, and in both progress has been recently achieved. We study the problems, their motivation, and their current status. The study is based on joint work with Moshe Vardi [16], Robby Lampert [14], and Benjamin Aminof [2].

2 Preliminaries

Word automata. An *infinite word* over an alphabet Σ is an infinite sequence $w = \sigma_1 \cdot \sigma_2 \cdots$ of letters in Σ. A *nondeterministic Büchi word automaton* (NBW, for short) is $\mathcal{A} = \langle \Sigma, Q, \delta, Q_0, F \rangle$, where Σ is the input alphabet, Q is a finite set of states, $\delta : Q \times \Sigma \to 2^Q$ is a transition function, $Q_0 \subseteq Q$ is a set of initial states, and $F \subseteq Q$ is a set of accepting states. If $|Q_0| = 1$ and δ is such that for every $q \in Q$ and $\sigma \in \Sigma$, we have that $|\delta(q, \sigma)| = 1$, then $\mathcal{A}$ is a *deterministic* Büchi word automaton (DBW).

Given an input word $w = \sigma_0 \cdot \sigma_1 \cdots$ in Σ^ω, a *run* of $\mathcal{A}$ on w is a sequence $r_0, r_1, \ldots$ of states in Q such that $r_0 \in Q_0$ and for every $i \geq 0$, we have $r_{i+1} \in \delta(r_i, \sigma_i)$; i.e., the run starts in one of the initial states and obeys the transition function. Note that a nondeterministic automaton can have many runs on w. In contrast, a deterministic automaton has a single run on w. For a run r, let $inf(r)$ denote the set of states that r visits infinitely often. That is, $inf(r) = \{q \in Q : r_i = q \text{ for infinitely many } i \geq 0\}$. As Q is finite, it is guaranteed that $inf(r) \neq \emptyset$. The run r is *accepting* iff $inf(r) \cap F \neq \emptyset$. That is, a run r is accepting iff there exists a state in F that r visits infinitely often. A run that is not accepting is *rejecting*. An NBW $\mathcal{A}$ accepts an input word w iff there exists an accepting run of $\mathcal{A}$ on w. The *language* of an NBW $\mathcal{A}$, denoted $\mathcal{L}(\mathcal{A})$, is the set of words that $\mathcal{A}$ accepts. We assume that a given NBW $\mathcal{A}$ has no empty states (that is, at least one word is accepted from each state – otherwise we can remove the state).

Linear Temporal Logic. The logic *LTL* is a linear temporal logic. Formulas of LTL are constructed from a set AP of atomic propositions using the usual Boolean operators and the temporal operators G ("always"), F ("eventually"), X ("next time"), and U ("until"). Formulas of LTL describe computations of systems over AP. For example, the LTL formula $G(req \rightarrow F ack)$ describes computations in which every position in which *req* holds is eventually followed by a position in which *ack* holds. For the detailed syntax and semantics of LTL, see [29]. The *model-checking problem* for LTL is to determine, given an LTL formula ψ and a system M, whether all the computations of M satisfy ψ.

General methods for LTL model checking are based on translation of LTL formulas to nondeterministic Büchi word automata:

Theorem 1. [41] *Given an LTL formula ψ, one can construct an NBW $\mathcal{A}_\psi$ that accepts exactly all the computations that satisfy ψ. The size of $\mathcal{A}_\psi$ is exponential in the length of ψ.*

Given a system M and a property ψ, model checking of M with respect to ψ is reduced to checking the emptiness of the product of M and $\mathcal{A}_{\neg\psi}$ [41]. This check can be performed on-the-fly and symbolically [6, 12, 38], and the complexity of model checking that follows is PSPACE, with a matching lower bound [35].

3 Translating Safety Properties to Automata

Of special interest are properties asserting that the system always stays within some allowed region, in which nothing "bad" happens. For example, we may want to assert that two processes are never simultaneously in the critical section. Such properties of systems are called *safety properties*. Intuitively, a property ψ is a safety property if every violation of ψ occurs after a finite execution of the system. In our example, if in a computation of the system two processes are in the critical section simultaneously, this occurs after some finite execution of the system.

In this section we study the translation of safety properties to nondeterministic automata on finite words (NFWs). We first define safety and co-safety languages, define bad and good prefixes, and motivate the above construction. We then describe a recent result that describes a construction of an NFW for bad and good prefixes for the case the safety or the co-safety property is given by means on an LTL formula.

3.1 Safety properties and their verification

We refer to computations of a nonterminating system as infinite words over an alphabet Σ. Typically, $\Sigma = 2^{AP}$, where AP is the set of the system's atomic propositions. Consider a language $L \subseteq \Sigma^\omega$ of infinite words over the alphabet Σ. A finite word $x \in \Sigma^*$ is a *bad prefix* for L iff for all $y \in \Sigma^\omega$, we have $x \cdot y \notin L$. Thus, a bad prefix is a finite word that cannot be extended to an infinite word in L. Note that if x is a bad prefix, then all the finite extensions of x are also bad prefixes. A language L is a *safety* language iff every infinite word $w \notin L$ has a finite bad prefix.[2] For a safety language L, we denote by *bad-pref*(L) the set of all bad prefixes for L. For example, if $\Sigma = \{0, 1\}$, then $L = \{0^\omega, 1^\omega\}$ is a safety language. To see this, note that every word not in L contains either the sequence 01 or the sequence 10, and a prefix that ends in one of these sequences cannot be extended to a word in L. Thus, *bad-pref*(L) is the language of the regular expression $(0^* \cdot 1 + 1^* \cdot 0) \cdot (0 + 1)^*$.

For a language $L \subseteq \Sigma^\omega$ (Σ^*), we use *comp*(L) to denote the complement of L; i.e., *comp*$(L) = \Sigma^\omega \setminus L$ ($\Sigma^* \setminus L$, respectively). We say that a language $L \subseteq \Sigma^\omega$ is a *co-safety* language iff *comp*(L) is a safety language. (The term used in [23] is *guarantee* language.) Equivalently, L is co-safety iff every infinite word

[2] The definition of safety we consider here is given in [1], it coincides with the definition of limit closure defined in [9], and is different from the definition in [20], which also refers to the property being closed under stuttering.

$w \in L$ has a *good prefix* $x \in \Sigma^*$: for all $y \in \Sigma^\omega$, we have $x \cdot y \in L$. For a co-safety language L, we denote by *good-pref*(L) the set of good prefixes for L. Note that for a safety language L, we have that *good-pref*$(comp(L)) = $ *bad-pref*(L). Thus, in order to construct the set of bad prefixes for a safety property, one can construct the set of good prefixes for its complementary language.

We say that an NBW $\mathcal{A}$ is a safety (co-safety) automaton iff $\mathcal{L}(\mathcal{A})$ is a safety (co-safety) language. We use *bad-pref*$(\mathcal{A})$, *good-pref*$(\mathcal{A})$, and *comp*$(\mathcal{A})$ to abbreviate *bad-pref*$(\mathcal{L}(\mathcal{A}))$, *good-pref*$(\mathcal{L}(\mathcal{A}))$, and *comp*$(\mathcal{L}(\mathcal{A}))$, respectively.

In addition to proof-based methods for the verification of safety properties [23, 24], there is extensive work on model checking of safety properties. Recall that general methods for model checking of linear properties are based on a construction of an NBW $\mathcal{A}_{\neg\psi}$ that accepts exactly all the infinite computations that violate the property ψ and is of size exponential in ψ [22, 41]. Verification of a system M with respect to ψ is then reduced to checking the emptiness of the product of M and $\mathcal{A}_{\neg\psi}$ [39].

When ψ is a safety property, the NBW $\mathcal{A}_{\neg\psi}$ can be replaced by *bad-pref*$(\mathcal{A}_\psi)$ – an NFW that accepts exactly all the bad prefixes of ψ [16]. This has several advantages, as reasoning about finite words is simpler than reasoning about infinite words: symbolic reasoning (in particular, bounded model checking procedures) need not look for loops and can, instead, apply backward or forward reachability analysis [4]. In fact, the construction of *bad-pref*$(\mathcal{A}_\psi)$ reduces the model-checking problem to the problem of invariance checking [23], which is amenable to both model-checking techniques and deductive verification techniques. In addition, using *bad-pref*$(\mathcal{A}_\psi)$, we can return to the user a finite error trace, which is a bad prefix, and which is often more helpful than an infinite error trace.

Consider a safety NBW $\mathcal{A}$. The construction of *bad-pref*$(\mathcal{A})$ was studied in [16]. If $\mathcal{A}$ is deterministic, we can construct a *deterministic* automaton on finite words (DFW) for *bad-pref*$(\mathcal{A})$ by defining the set of accepting states to be the set of states s for which $\mathcal{A}$ with initial state s is empty. Likewise, if $\mathcal{A}$ is a co-safety automaton, we can construct a DFW for *good-pref*$(\mathcal{A})$ by defining the set of accepting states to be the set of states s for which $\mathcal{A}$ with initial state s is universal.

When $\mathcal{A}$ is nondeterministic, the story is more complicated. Even if we are after a nondeterministic, rather than a deterministic, automaton for the bad or good prefixes, the transition from infinite words to finite words involves an exponential blow-up. Formally, we have the following.

Theorem 2. [16] *Consider an NBW $\mathcal{A}$ of size n.*

1. *If $\mathcal{A}$ is a safety automaton, the size of an NFW for bad-pref*$(\mathcal{A})$ *is* $2^{\Theta(n)}$.
2. *If $\mathcal{A}$ is a co-safety automaton, the size of an NFW for good-pref*$(\mathcal{A})$ *is* $2^{\Theta(n)}$.

The lower bound in Theorem 2 for the case $\mathcal{A}$ is a safety automaton is not surprising. Essentially, it follows from the fact that *bad-pref*$(\mathcal{A})$ refers to words that are not accepted by $\mathcal{A}$. Hence, it has the flavor of complementation, and complementation of nondeterministic automata involves an exponential blow-up [27]. The second blow up, however, in going from a co-safety automaton to a

nondeterministic automaton for its good prefixes is surprising. Its proof in [16] highlights the motivation behind the definition of a *fine automaton* for safety properties, and we describe it below.

For $n \geq 1$, let $\Sigma_n = \{1, \ldots, n, \&\}$. We define L_n as the language of all words $w \in \Sigma_n^\omega$ such that w contains at least one & and the letter after the first & is either & or it has already appeared somewhere before the first &. The language L_n is a co-safety language. Indeed, each word in L_n has a good prefix (e.g., the one that contains the first & and its successor). We can recognize L_n with an NBW with $O(n)$ states (the NBW guesses the letter that appears after the first &). Obvious good prefixes for L_n are 12&&, 123&2, etc. That is, prefixes that end one letter after the first &, and their last letter is either & or has already appeared somewhere before the &. We can recognize these prefixes with an NFW with $O(n)$ states. But L_n also has some less obvious good prefixes, like $1234 \cdots n\&$ (a permutation of $1 \ldots n$ followed by &). These prefixes are indeed good, as every suffix we concatenate to them would start with & or with a letter in $\{1, \ldots, n\}$, which has appeared before the &. To recognize these prefixes, an NFW needs to keep track of subsets of $\{1, \ldots, n\}$, for which it needs 2^n states. Consequently, an NFW for *good-pref*(L_n) must have at least 2^n states.

It is also shown in [16] that the language L_n can be encoded by an LTL formula of length quadratic in n. This implies that the translation of safety and co-safety LTL formulas to NFWs for their bad and good prefixes is doubly exponential. Formally, we have the following.

Theorem 3. [16] *Given a safety LTL formula ψ of size n, the size of an NFW for bad-pref(ψ) is $2^{2^{\Omega(\sqrt{n})}}$.*

3.2 Fine automata and their construction

As described in the proof of Theorem 2, some good prefixes for L_n (the "obvious prefixes") can be recognized by a small NFW. What if we give up the non-obvious prefixes and construct an NFW $\mathcal{A}'$ that accepts only the "obvious subset" of L_n? It is not hard to see that each word in L_n has an obvious prefix. Thus, while $\mathcal{A}'$ does not accept all the good prefixes, it accepts at least one prefix of every word in L. This useful property of $\mathcal{A}'$ is formalized below.

Consider a safety language L. We say that a set $X \subseteq$ *bad-pref*(L) is a *trap* for L iff every word $w \notin L$ has at least one bad prefix in X. Thus, while X need not contain all the bad prefixes for L, it must contain sufficiently many prefixes to "trap" all the words not in L. Dually, a trap for a co-safety language L is a set $X \subseteq$ *good-pref*(L) such that every word $w \in L$ has at least one good prefix in X. We denote the set of all the traps, for an either safety or co-safety language L, by *trap*(L).

An NFW $\mathcal{A}$ is *fine* for a safety or a co-safety language L iff $\mathcal{A}$ accepts a trap for L. For example, an NFW that accepts $0^* \cdot 1 \cdot (0 + 1)$ does not accept all the bad prefixes of the safety language $\{0^\omega\}$; in particular, it does not accept the minimal bad prefixes in $0^* \cdot 1$. Yet, such an NFW is fine for $\{0^\omega\}$. Indeed, every infinite word that is different from 0^ω has a prefix in $0^* \cdot 1 \cdot (0 + 1)$. Likewise,

the NFW is fine for the co-safety language $0^* \cdot 1 \cdot (0+1)^\omega$. In practice, almost all the benefit that one obtains from an NFW that accepts all the bad/good prefixes can also be obtained from a fine automaton. It is shown in [16] that for natural safety formulas ψ, the construction of an NFW fine for ψ is as easy as the construction of $\mathcal{A}_{\neg\psi}$. In more details, if we regard $\mathcal{A}_{\neg\psi}$ as an NFW, with an appropriate definition of the set of accepting states, we get an automaton fine for ψ. For general safety formulas, the problem of constructing small fine automata was left open in [16] and its solution in [14] has led to new mysteries in the context of safety properties. Let us first describe the result in [14].

Recall that the transition from a safety NBW to an NFW for its bad prefixes is exponential, and that the exponential blow up follows from the fact that a complementing NBW can be construction from a tight NFW. When we consider fine automata, things are more complicated, as the fine NFW need not accept all bad prefixes. As we show below, however, a construction of fine automata still has the flavor of complementation, and must involve an exponential blow up.

Theorem 4. [14] *Given a safety NBW $\mathcal{A}$ of size n, the size of an NFW fine for $\mathcal{A}$ is $2^{\Theta(n)}$.*

We now move on to consider co-safety NBWs. Recall that, as with safety properties and bad prefixes, the transition from a co-safety NBW to an NFW for its good prefixes is exponential. We show that a fine NFW for a co-safety property can be constructed from the NBWs for the property and its negation. The idea is that it is possible to bound the number of times that a run of $\mathcal{A}$ visits the set of accepting states when it runs on a word not in $\mathcal{L}(\mathcal{A})$. Formally, we have the following:

Lemma 1. [14] *Consider a co-safety NBW $\mathcal{A}$. Let F be the set of accepting states of $\mathcal{A}$ and let $\overline{\mathcal{A}}$ be an NBW with $\overline{n}$ states such that $\mathcal{L}(\overline{\mathcal{A}}) = comp(\mathcal{L}(\mathcal{A}))$. If a run of $\mathcal{A}$ on a finite word $h \in \Sigma^*$ visits F more than $|F| \cdot \overline{n}$ times, then h is a good prefix for $\mathcal{L}(\mathcal{A})$.*

Consider a co-safety NBW $\mathcal{A}$ with n states, m of them accepting. Let $\overline{\mathcal{A}}$ be an NBW with $\overline{n}$ states such that $\mathcal{L}(\overline{\mathcal{A}}) = comp(\mathcal{L}(\mathcal{A}))$. Following Lemma 1, we can construct an NFW fine for $\mathcal{A}$ by taking $(m \cdot \overline{n}) + 1$ copies of $\mathcal{A}$, and defining the transition function such that when a run of $\mathcal{A}'$ visits F in the j-th copy of $\mathcal{A}$, it moves to the $(j+1)$-th copy. The accepting states of $\mathcal{A}'$ are the states of F in the $(m \cdot \overline{n} + 1)$-th copy. This implies the following theorem.

Theorem 5. [14] *Consider a co-safety NBW $\mathcal{A}$ with n states, m of them accepting. Let $\overline{\mathcal{A}}$ be an NBW with $\overline{n}$ states such that $\mathcal{L}(\overline{\mathcal{A}}) = comp(\mathcal{L}(\mathcal{A}))$. There exists an NFW $\mathcal{A}'$ with $n \cdot (m \cdot \overline{n} + 1)$ states such that $\mathcal{A}'$ is fine for $\mathcal{L}(\mathcal{A})$.*

Given a safety NBW, its complement NBW is co-safety. Thus, dualizing Theorem 5, we get the following.

Theorem 6. [14] *Consider a safety NBW with n states. Let $\overline{\mathcal{A}}$ be an NBW with $\overline{n}$ states, $\bar{m}$ of them accepting, such that $\mathcal{L}(\overline{\mathcal{A}}) = comp(\mathcal{L}(\mathcal{A}))$. There exists an NFW $\mathcal{A}'$ with $\bar{n} \cdot (\bar{m} \cdot n + 1)$ states such that $\mathcal{A}'$ is fine for $\mathcal{L}(\mathcal{A})$.*

By Theorem 1, given an LTL formula ψ, we can construct NBWs $\mathcal{A}_\psi$ and $\mathcal{A}_{\neg\psi}$ for ψ and $\neg\psi$, respectively. The number of states in each of the NBWs is at most $2^{O(|\psi|)}$. Hence, by Theorem 5, we can conclude:

Theorem 7. [14] *Consider a safety LTL formula φ of length n. There exists an NFW fine for φ with at most $2^{O(n)}$ states.*

It follows from Theorem 7 that the transition from a tight NFW (one that accepts exactly all bad or good prefixes) to a fine NFW is significant, as it circumvents the doubly exponential blow-up in Theorem 3.

3.3 Discussion

The work in [14] has answered positively the question about the existence of exponential fine automata for general safety LTL formulas, improving the doubly-exponential construction in [16]. Essentially, the construction adds a counter on top of the NBW for the formula. The counter is increased whenever the NBW visits an accepting state, and a computation is accepted after the counter reaches a bound that depends on the size of the formula. For a discussion on the application of the result in the context of bounded model checking and run-time verification see [14]. Here, we discuss the theoretical aspects of the result.

While [14] has solved the problem of constructing exponential fine automata for LTL formulas, the problem of constructing polynomial fine automata for co-safety NBW is still open. The challenge here is similar to other challenges in automata-theoretic constructions in which one needs both the NBW and its complementing NBW — something that is easy to have in the context of LTL, but difficult in the context of NBW. More problems in this status are reported in [18]. For example, the problem of deciding whether an LTL formula ψ can be translated to a DBW can be solved by reasoning about the NBWs for ψ and $\neg\psi$. This involves an exponential blow up in the length of ψ, but, as in our case, no additional blow-up for complementation. The problem of deciding whether an NBW can be translated to a DBW cannot be solved using the same lines, as here complementation does involve an exponential blow up. From a practical point of view, however, the problem of going from a co-safety automaton to a fine NFW is of less interest, as users that use automata as their specification formalism are likely to start with an automaton for the bad or the good prefixes anyway. Thus, the problem about the size of fine automata is interesting mainly for the specification formalism of LTL, which [14] did solve.

4 From Büchi to co-Büchi Automata

The second open problem we describe is the problem of translating, when possible, a nondeterministic Büchi word automaton to an equivalent nondeterministic

co-Büchi word automaton (NCW). The *co-Büchi* acceptance condition is dual to the Büchi acceptance condition. Thus, $F \subseteq Q$ and a run r is accepting if it visits F only finitely many times. Formally, $inf(r) \cap F = \emptyset$. NCWs are less expressive than NBWs. For example, the language $\{w : w \text{ has only finitely many } 0\text{s}\} \subseteq \{0,1\}^\omega$ cannot be recognized by an NCW. In fact, NCWs are as expressive as deterministic co-Büchi automata (DCWs). Hence, as DBWs are dual to DCWs, a language can be recognized by an NCW iff its complement can be recognized by a DBW.

The best translation of NBW to NCW (when possible) that is currently known actually results in a deterministic co-Büchi automaton (DCW), and it goes as follows. Consider an NBW $\mathcal{A}$ that has an equivalent NCW. First, co-determinize $\mathcal{A}$ and obtain a deterministic Rabin automaton (DRW) $\tilde{\mathcal{A}}$ for the complement language. By [13], DRWs are *Büchi type*. That is, if a DRW has an equivalent DBW, then the DRW has an equivalent DBW on the same structure. Let $\tilde{\mathcal{B}}$ be the DBW equivalent to $\tilde{\mathcal{A}}$ (recall that since $\mathcal{A}$ can be recognized by an NCW, its complement can be recognized by a DBW). By dualizing $\tilde{\mathcal{B}}$ one gets a DCW equivalent to $\mathcal{A}$. The co-determinization step involves an exponential blowup in the number of states [33]. Hence, starting with an NBW with n states, we end up with an NCW with $2^{O(n \log n)}$ states. This is particularly annoying as even a lower bound showing that an NCW needs one more state is not known. As we discuss below, the translation of NBW to an equivalent NCW is of practical importance because of its relation to the problem of translating LTL formulas to equivalent alternation-free μ-calculus (AFMC) formulas (when possible).

It is shown in [17] that given an LTL formula ψ, there is an AFMC formula equivalent to $\forall\psi$ iff ψ can be recognized by a DBW. Evaluating specifications in the alternation-free fragment of μ-calculus can be done with linearly many symbolic steps. In contrast, direct LTL model checking reduces to a search for bad-cycles and its symbolic implementation involves nested fixed-points, and is typically quadratic [32]. Hence, identifying LTL formulas that can be translated to AFMC, and coming up with an optimal translation, is a problem of great practical importance. The best known translations of LTL to AFMC first translates the LTL formula ψ to a DBW, which is then linearly translated to an AFMC formula for $\forall\psi$. The translation of LTL to DBW, however, is doubly exponential, thus the overall translation is doubly-exponential, with only an exponential matching lower bound [17].

The reason that current translations go through an intermediate deterministic automaton is the need to run this automaton on all the computations of the system in a way that computations with the same prefix follow the same run. A similar situation exists when we expand a word automaton to a tree automaton [8] — the word automaton cannot be nondeterministic, as different branches of the tree that have the same prefix u may be accepted by runs of the word automaton that do not agree on the way they proceed on u. A promising direction for coping with this situation was suggested in [17]: Instead of translating the LTL formula ψ to a DBW, one can translate $\neg\psi$ to an NCW. This can be done either directly, or by translating the NBW for $\neg\psi$ to an equivalent

NCW. Then, the NCW can be linearly translated to an AFMC formula for $\exists\neg\psi$, whose negation is equivalent to $\forall\psi$. The fact that the translation can go through a nondeterministic rather than a deterministic automaton is very promising, as nondeterministic automata are typically exponentially more succinct than deterministic ones.[3] Nevertheless, the problem of translating LTL formulas to NCWs of exponential size is still open.[4] The best translation that is known today involves a doubly-exponential blow up, and it actually results in a DCW, giving up the idea that the translation of LTL to AFMC can be exponentially more efficient by using intermediate nondeterministic automata. Note that a polynomial translation of NBW to NCW will imply a singly-exponential translation of LTL to AFMC, as the only exponential step in the procedure will be the translation of LTL to NBW.[5]

Recall that while the best upper bound for an NBW to NCW translation is $2^{O(n \log n)}$, we do not even have a single example to a language whose NBW is smaller than its NCW. In fact, it was only recently shown that NBWs are not *co-Büchi-type*. That is, there is an NBW $\mathcal{A}$ such that $L(\mathcal{A})$ can be recognized by an NCW, but an NCW for $L(\mathcal{A})$ must have a different structure than $\mathcal{A}$. We describe such an NBW in the proof below (the original proof, in [15], has a different example).

Lemma 2. [15] *NBWs are not co-Büchi-type.*

Proof: Consider the NBW described in Figure 1. Note that the NBW has two initial states. The NBW recognizes the language L of all words with at least one a and at least one b. This language can be recognized by an NCW, yet it is easy to see that there is no way to define F on top of $\mathcal{A}$ such that the result is an NCW that recognizes L. □

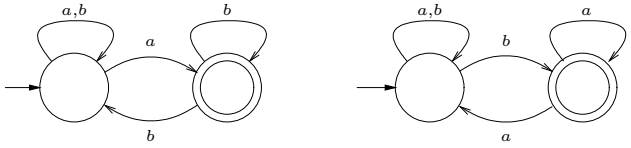

Fig. 1. An NBW for "at last one a and at least one b"

[3] Dually, we can translate the LTL formula to a *universal* Büchi automaton and translate this automaton to an AFMC formula. The universal Büchi automaton for ψ is dual to the nondeterministic co-Büchi automaton for $\neg\psi$.

[4] As mentioned above, not all LTL formulas can be translated to NCWs. When we talk about the blow up in a translation, we refer to formulas for which a translation exists.

[5] Wilke [42] proved an exponential lower-bound for the translation of an NBW for an LTL formula ψ to and AFMC formula equivalent to $\forall\psi$. This lower-bound does not preclude a polynomial upper-bound for the translation of an NBW for $\neg\psi$ to an AFMC formula equivalent to $\exists\neg\psi$, which is our goal.

During our efforts to solve the NBW to NCW problem, we have studied the related problem of translating NBWs to NFWs. In the next section we describe the problem, its relation to the NBW to NCW problem, and our partial success in this front.

4.1 From Büchi to limit finite automata

Recall that DBWs are less expressive than NBWs. Landweber characterizes languages $L \subseteq \Sigma^\omega$ that can be recognized by a DBW as those for which there is a regular language $R \subseteq \Sigma^*$ such that L is the *limit* of R. Formally, w is in the limit of R iff w has infinitely many prefixes in R [21]. It is not hard to see that a DBW for L, when viewed as a DFW, recognizes a language whose limit is L, and vice versa – a DFW for R, when viewed as a DBW, recognizes the language that is the limit of R. What about the case R and L are given by nondeterministic automata? It is not hard to see that the simple transformation between the two formalisms no longer holds. For example, the NBW $\mathcal{A}$ in Figure 2 recognizes the language L of all words with infinitely many bs, yet when viewed as an NFW, it recognizes $(a+b)^+$, whose limit is $(a+b)^\omega$. As another example, the language of the NBW $\mathcal{A}'$ is empty, yet when viewed as an NFW, it recognizes the language $(a+b)^* \cdot b$, whose limit is L. As demonstrated by the examples, the difficulty of the nondeterministic case originates from the fact that different prefixes of the infinite word may follow different accepting runs of the NFW, and there is no guarantee that these runs can be merged into a single run of the NBW. Accordingly, the best translation that was known until recently for going from an NFW to an NBW accepting its limit, or from an NBW to a limit NFW, is to first determinize the given automaton. This involves a $2^{O(n \log n)}$ blow up and gives up the potential succinctness of the nondeterministic model. On the other hand, no lower bound above $\Omega(n \log n)$ is known.

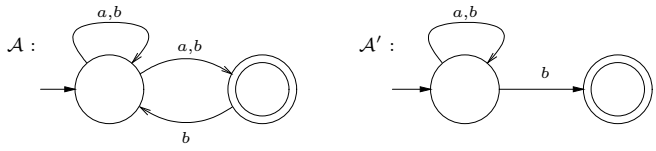

Fig. 2. Relating NBWs and limit NFWs.

In [2] we study this exponential gap and tried to close it. In addition to the limit operator introduced by Landweber, we introduce and study two additional ways to induce a language of infinite words from a language of finite words: the *co-limit* of R is the set of all infinite words that have only finitely many prefixes in R. Thus, co-limit is dual to Landweber's limit. Also, the *persistent limit* of R is the set of all infinite words that have only finitely many prefixes not in R. Thus, eventually all the prefixes are in R. Formally, we have the following.

Definition 1. *Consider a language $R \subseteq \Sigma^*$. We define three languages of infinite words induced by R.*

1. **[limit]** *$lim(R) \subseteq \Sigma^\omega$ is the set of all words that have infinitely many prefixes in R. I.e., $lim(R) = \{w \mid w[1,i] \in R \text{ for infinitely many } i\text{'s}\}$* [21].
2. **[co-limit]** *$co\text{-}lim(R) \subseteq \Sigma^\omega$ is the set of all words that have only finitely many prefixes in R. I.e., $co\text{-}lim(R) = \{w \mid w[1,i] \in R \text{ for finitely many } i\text{'s}\}$.*
3. **[persistent limit]** *$plim(R) \subseteq \Sigma^\omega$ is the set of all words that have only finitely many prefixes not in R. I.e., $plim(R) = \{w \mid w[1,i] \in R \text{ for almost all } i\text{'s}\}$.*

For example, for $R = (a+b)^*b$, the language $lim(R)$ consists of all words that have infinitely many b's, $co\text{-}lim(R)$ is the language of words that have finitely many b's, and $plim(R)$ is the language of words that have finitely many a's. For an NFW $\mathcal{A}$, we use $lim(\mathcal{A})$, $co\text{-}lim(\mathcal{A})$, and $plim(\mathcal{A})$, to denote $lim(L(\mathcal{A}))$, $co\text{-}lim(L(\mathcal{A}))$, and $plim(L(\mathcal{A}))$, respectively. The three limit operators are dual in the sense that for all $R \subseteq \Sigma^*$, we have $comp(lim(R)) = co\text{-}lim(R) = plim(comp(R))$.

Below we describe the main results of [2], which studies the relative succinctness of NBWs, NCWs, and NFWs whose limit, co-limit, and persistent limit correspond to the NBW and NCW.

We first need some notations. Consider an NFW $\mathcal{A} = \langle \Sigma, Q, \delta, Q_0, F \rangle$. For two sets of states $P, S \subseteq Q$, we denote by $L_{P,S}$ the language of $\mathcal{A}$ with initial set P and accepting set S. Theorem 8 is a key theorem in beating the "determinize first" approach. It implies that the transition from an NFW $\mathcal{A}$ to an NBW for $lim(\mathcal{A})$ need not involve a determinization of $\mathcal{A}$. Indeed, we can specify $lim(\mathcal{A})$ as the union of languages that are generated by automata with a structure similar to the structure of $\mathcal{A}$. Formally, we have the following.

Theorem 8. [2] *For every NFW $\mathcal{A} = \langle \Sigma, Q, \delta, Q_0, F \rangle$,*

$$lim(\mathcal{A}) = \bigcup_{p \in Q} L_{Q_0,\{p\}} \cdot (L_{\{p\},\{p\}} \cap L_{\{p\},F})^\omega.$$

Given $\mathcal{A}$ with n states, Theorem 8 implies that an NBW accepting $lim(\mathcal{A})$ can be constructed by intersection, application of ω, concatenation, and union, starting with NFWs with n states. Exploiting the the similarity in the structure of the involved NFWs, the resulting NBW has $O(n^2)$ states.

Corollary 1. *Given an NFW $\mathcal{A}$ with n states, there is an NBW $\mathcal{A}'$ with $O(n^2)$ states such that $L(\mathcal{A}') = lim(L(\mathcal{A}))$.*

Corollary 1 implies that going from an NFW to an NBW for its limit, it is possible to do better than determinize the NFW. On the other hand, it is shown in [2] that going from an NFW to an NCW for its co-limit or persistent limit, an exponential blow-up cannot be avoided, and determinization is optimal.

Further results of [2] study succinctness among NFWs to which different limit operators are applied. For example, in Theorem 9 below we prove that going from a persistent limit NFW to a limit NFW involves an exponential blow up. In other words, given an NFW $\mathcal{A}$ whose persistent limit is L, translating

$\mathcal{A}$ to an NFW whose limit is L may involve an exponential blow up. Note that persistent limit and limit are very similar – both require the infinite word to have infinitely many prefixes in $L(\mathcal{A})$, only that the persistent limit requires, in addition, that only finitely many prefixes are not in $L(\mathcal{A})$. This difference, which is similar to the difference between NBW and NCW, makes persistent limit exponentially more succinct. Technically, it follows from the fact that persistent limit NFWs inherit the power of alternating automata. In a similar, though less surprising way, co-limit NFWs inherit the power of complementation, and are also exponentially more succinct.

Theorem 9. [2] *For every $n \geq 1$, there is a language $L_n \subseteq \Sigma^\omega$ such that there are NFWs $\mathcal{A}$ with $O(n)$ states, and $\mathcal{A}'$ with $O(n^2)$ states, such that co-lim$(\mathcal{A})$ = plim$(\mathcal{A}') = L_n$ but an NFW $\mathcal{A}''$ such that lim$(\mathcal{A}'') = L_n$ must have at least 2^n states.*

Proof: Consider the language $L_n \subseteq \{0,1\}^\omega$ of all words w such that $w = uuz$, with $|u| = n$. We prove that an NFW $\mathcal{A}''$ such that lim$(\mathcal{A}'') = L_n$ must remember subsets of size n, and thus must have at least 2^n states. In order to construct small NFW for the co-limit and persistent limit operators, we observe that a word w is in L_n iff $\bigwedge_{i=1}^{n}(w[i] = w[n+i])$. In the case of co-limit, we can check that only finitely many (in fact, 0) prefixes h of an input word are such that $h[i] \neq h[i+n]$ for some $1 \leq i \leq n$. The case of persistent limit is much harder, as we cannot use the implicit complementation used in the co-limit case. Instead, we use the universal nature of persistence. We define the NFW $\mathcal{A}'$ as a union of n NFWs $\mathcal{A}'_1, \ldots, \mathcal{A}'_n$. The NFW $\mathcal{A}'_i$ is responsible for checking that $w[i] = w[n+i]$. In order to make sure that the conjunction on all $1 \leq i \leq n$ is satisfied, we further limit $\mathcal{A}'_i$ to accept only words of length i mod n. Hence, $\mathcal{A}'_i$ accepts a word $u \in \Sigma^*$ iff $u[i] = u[n+i] \wedge |u| = i \bmod n$. Thus, plim$(\mathcal{A}') = L_n$. □

4.2 Discussion

The exponential gap between the known upper and lower bounds in the translation of NBW to NCW is particularly annoying: the upper bound is $2^{O(n \log n)}$ and for the lower bound we do not even have an example of a language whose NCW needs one more state than the NBW. The example in the proof of Lemma 2 shows an advantage of the Büchi condition. In a recent work with Benjamin Aminof and Omer Lev, we hope to turn this advantage into a lower bound. The idea is as follows. NCWs cannot recognize the language of all words that have infinitely many occurrences of some letter. Indeed, DBWs cannot recognize the complement language [21]. Thus, a possible way to obtain a lower bound for the translation of NBW to NCW is to construct a language that is recognizable by an NCW, but for which an NCW needs more states than an NBW due to its inability to recognize infinitely many occurrences of a letter. One such candidate is the family of languages $L_1, L_2, \ldots$ over the alphabet $\{a, b\}$, where L_k contains exactly all words that have at least k occurrences of the letter a and at least k occurrences of the letter b. An NBW can follow the idea of the NBW in Figure 1:

since every infinite word has infinitely many a's or infinitely many b's, the NBW for L can guess which of the two letters occurs infinitely often, and count k occurrences of the second letter. Thus, the NBW is the union of two components, one looking for k occurrences of a followed by infinitely many b's and the other looking for k occurrences of b followed by infinitely many a's. This can be done with $2k+1$ states. We conjecture that an NCW needs more than two counters. The reason is that an NCW with less than k states accepting all words with infinitely many a's, inevitably also accepts a word with less than k a's.

References

1. B. Alpern and F.B. Schneider. Defining liveness. *IPL*, 21:181–185, 1985.
2. B. Aminof and O. Kupferman. On the succinctness of nondeterminizm. In *Proc. 4th ATVA*, LNCS 4218, pages 125–140. 2006.
3. J.R. Büchi. On a decision method in restricted second order arithmetic. In *Proc. Int. Congress on Logic, Method, and Philosophy of Science. 1960*, pages 1–12. Stanford University Press, 1962.
4. J.R. Burch, E.M. Clarke, K.L. McMillan, D.L. Dill, and L.J. Hwang. Symbolic model checking: 10^{20} states and beyond. *I&C*, 98(2):142–170, 1992.
5. E.M. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, 1999.
6. C. Courcoubetis, M.Y. Vardi, P. Wolper, and M. Yannakakis. Memory efficient algorithms for the verification of temporal properties. *FMSD*, 1:275–288, 1992.
7. C. Eisner and D. Fisman. *A Practical Introduction to PSL*. Springer, 2006.
8. A.E. Emerson and A.P. Sistla. Deciding full branching time logics. *I&C*, 61(3):175–201, 1984.
9. E.A. Emerson. Alternative semantics for temporal logics. *TCS*, 26:121–130, 1983.
10. E.A. Emerson and C. Jutla. The complexity of tree automata and logics of programs. In *Proc. 29th FOCS*, pages 328–337, 1988.
11. E. Friedgut, O. Kupferman, and M.Y. Vardi. Büchi complementation made tighter. In *Proc. 2nd ATVA*, LNCS 3299, pages 64–78, 2004.
12. 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, 1995.
13. S.C. Krishnan, A. Puri, and R.K. Brayton. Deterministic ω-automata vis-a-vis deterministic Büchi automata. In *Algorithms and Computations*, LNCS 834, pages 378–386, 1994.
14. O. Kupferman and R. Lampert. On the construction of fine automata for safety properties. In *Proc. 4th ATVA*, LNCS 4218, pages 110–124, 2006.
15. O. Kupferman, G. Morgenstern, and A. Murano. Typeness for ω-regular automata. In *Proc. 2nd ATVA*, LNCS 3299, pages 324–338, 2004.
16. O. Kupferman and M.Y. Vardi. Model checking of safety properties. *FMSD*, 19(3):291–314, 2001.
17. O. Kupferman and M.Y. Vardi. From linear time to branching time. *ACM TOCL*, 6(2):273–294, 2005.
18. O. Kupferman and M.Y. Vardi. Safraless decision procedures. In *Proc. 46th FOCS*, pages 531–540, 2005.
19. R.P. Kurshan. *Computer Aided Verification of Coordinating Processes*. Princeton Univ. Press, 1994.

20. L. Lamport. Logical foundation. In *Distributed systems - methods and tools for specification*, LNCS 190, 1985.
21. L.H. Landweber. Decision problems for ω–automata. *Mathematical Systems Theory*, 3:376–384, 1969.
22. O. Lichtenstein and A. Pnueli. Checking that finite state concurrent programs satisfy their linear specification. In *Proc. 12th POPL*, pages 97–107, 1985.
23. Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer, 1992.
24. Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Safety*. Springer, 1995.
25. R. McNaughton. Testing and generating infinite sequences by a finite automaton. *I&C*, 9:521–530, 1966.
26. A. R. Meyer. Weak monadic second order theory of successor is not elementary recursive. In *Proc. Logic Colloquium*, LNM 453, pages 132–154. Springer, 1975.
27. A.R. Meyer and M.J. Fischer. Economy of description by automata, grammars, and formal systems. In *Proc. 12th SWAT*, pages 188–191, 1971.
28. M. Michel. Complementation is more difficult with automata on infinite words. CNET, Paris, 1988.
29. A. Pnueli. The temporal semantics of concurrent programs. *TCS*, 13:45–60, 1981.
30. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proc. 16th POPL*, pages 179–190, 1989.
31. M.O. Rabin. Decidability of second order theories and automata on infinite trees. *Transaction of the AMS*, 141:1–35, 1969.
32. K. Ravi, R. Bloem, and F. Somenzi. A comparative study of symbolic algorithms for the computation of fair cycles. In *Proc. 3rd FMCAD*, LNCS 1954, pages 143–160, 2000.
33. S. Safra. On the complexity of ω-automata. In *Proc. 29th FOCS*, pages 319–327, 1988.
34. A.P. Sistla. Safety, liveness and fairness in temporal logic. *Formal Aspects of Computing*, 6:495–511, 1994.
35. A.P. Sistla and E.M. Clarke. The complexity of propositional linear temporal logic. *JACM*, 32:733–749, 1985.
36. A.P. Sistla, M.Y. Vardi, and P. Wolper. The complementation problem for Büchi automata with applications to temporal logic. In *Proc. 12th ICALP*, LNCS 194, pages 465–474, 1985.
37. R.S. Street and E.A. Emerson. An elementary decision procedure for the μ-calculus. In *Proc. 11th ICALP*, volume 172, pages 465–472. Springer, 1984.
38. H.J. Touati, R.K. Brayton, and R. Kurshan. Testing language containment for ω-automata using BDD's. *I&C*, 118(1):101–109, 1995.
39. M.Y. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification. In *Proc. 1st LICS*, pages 332–344, 1986.
40. M.Y. Vardi and P. Wolper. Automata-theoretic techniques for modal logics of programs. *JCSS*, 32(2):182–221, 1986.
41. M.Y. Vardi and P. Wolper. Reasoning about infinite computations. *I&C*, 115(1):1–37, 1994.
42. T. Wilke. CTL$^+$ is exponentially more succinct than CTL. In *Proc. 19th FST&TCS*, LNCS 1738, pages 110–121, 1999.
43. P. Wolper, M.Y. Vardi, and A.P. Sistla. Reasoning about infinite computation paths. In *Proc. 24th FOCS*, pages 185–194, 1983.
44. Q. Yan. Lower bounds for complementation of ω-automata via the full automata technique. In *Proc. 33rd ICALP*, LNCS 4052, pages 589–600, 2006.