

On Semantically-Deterministic Automata*

Bader Abu Radi and Orna Kupferman
School of Computer Science and Engineering
The Hebrew University, Jerusalem, Israel

May 23, 2023

Abstract

A nondeterministic automaton is *semantically deterministic* (SD) if different nondeterministic choices in the automaton lead to equivalent states. Semantic determinism is interesting as it is a natural relaxation of determinism, and as some applications of deterministic automata in formal methods can actually use automata with some level of nondeterminism, tightly related to semantic determinism.

In the context of finite words, semantic determinism coincides with determinism, in the sense that every pruning of an SD automaton to a deterministic one results in an equivalent automaton. We study SD automata on infinite words, focusing on Büchi, co-Büchi, and weak automata. We show that there, while semantic determinism does not increase the expressive power, the combinatorial and computational properties of SD automata are very different from these of deterministic automata. In particular, SD Büchi and co-Büchi automata are exponentially more succinct than deterministic ones (in fact, also exponentially more succinct than history-deterministic automata), their complementation involves an exponential blow up, and decision procedures for them like universality and minimization are PSPACE-complete. For weak automata, we show that while an SD weak automaton need not be pruned to an equivalent deterministic one, it can be determinized to an equivalent deterministic weak automaton with the same state space, implying also efficient complementation and decision procedures for SD weak automata.

1 Introduction

Automata are among the most studied computation models in theoretical computer science. Their simple structure has made them a basic formalism for the study of fundamental notions, such as *determinism* and *nondeterminism* [34]. While a deterministic computing machine examines a single action at each step of its computation, nondeterministic machines are allowed to examine several possible actions simultaneously. Understanding the power of nondeterminism is at the core of open fundamental questions in theoretical computer science (most notably, the P vs. NP problem).

A prime application of *automata on infinite words* is specification, verification, and synthesis of nonterminating systems. The automata-theoretic approach reduces questions about systems and their specifications to questions about automata [27, 40], and is at the heart of many algorithms and tools. A run of an automaton on infinite words is an infinite sequence of states, and acceptance is determined with respect to the set of states that the run visits infinitely often. For example, in *Büchi* automata, some of the states are designated as accepting states, and a run is accepting iff it visits states from the set α of accepting states infinitely often [8]. Dually, in *co-Büchi* automata, a run is

*A preliminary version appears in the Proceedings of the 50th International Colloquium on Automata, Languages, and Programming, 2023. Research supported by the Israel Science Foundation, Grant 2357/19, and the European Research Council, Advanced Grant ADVANSYNT.

accepting if it visits the set α only finitely often. Then, *weak* automata are a special case of both Büchi and co-Büchi automata in which every strongly connected component in the graph induced by the automaton is either contained in α or is disjoint from α . We use DBW and NBW to denote deterministic and nondeterministic Büchi word automata, respectively, and similarly for D/NCW, D/NWW, and D/NFW, for co-Büchi, weak, and automata on finite words, respectively.

For automata on infinite words, nondeterminism not only leads to exponential succinctness, but may also increase the expressive power. This is the case, for example, in Büchi and weak automata, thus NBWs are strictly more expressive than DBWs [28], and NWWs are strictly more expressive than DWWs [5]. On the other hand, NCWs are as expressive as DCWs [31], and in fact, also as NWWs [26]. In some applications of the automata-theoretic approach, such as model checking, algorithms can be based on nondeterministic automata, whereas in other applications, such as synthesis and reasoning about probabilistic systems, they cannot. There, the advantages of nondeterminism are lost, and algorithms involve a complicated determinization construction [35] or acrobatics for circumventing determinization [25, 20].

In a deterministic automaton, the transition function maps each state and letter to a single successor state. In recent years there is growing research on weaker types of determinism. This includes, for example, *unambiguous automata*, which may have many runs on each word, yet only one accepting run [9, 11], automata that are *deterministic in the limit*, where each accepting run should eventually reach a deterministic sub-automaton [39], and automata that are *determinizable by pruning* (DBP), thus embody an equivalent deterministic automaton [3].

In terms of applications, some weaker types of determinism have been defined and studied with specific applications in mind. Most notable are *history-deterministic* automata (HD), which can resolve their nondeterministic choices based on the history of the run [17, 7, 22]¹, and can therefore replace deterministic automata in algorithms for synthesis and control, and *good-for-MDPs* automata (GFM), whose product with Markov decision processes maintains the probability of acceptance, and can therefore replace deterministic automata when reasoning about stochastic behaviors [15, 38].

The different levels of determinism induce classes of automata that differ in their succinctness and in the complexity of operations and decision problems on them. Also, some classes are subclasses of others. For example, it follows quite easily from the definitions that every automaton that is deterministic in the limit is GFM, and every automaton that is DBP is HD.

In this paper we study the class of *semantically deterministic* automata (SD). An automaton \mathcal{A} is SD if its nondeterministic choices lead to equivalent states. Formally, if $\mathcal{A} = \langle \Sigma, Q, q_0, \delta, \alpha \rangle$, with a transition function $\delta : Q \times \Sigma \rightarrow 2^Q$, then \mathcal{A} is SD if for every state $q \in Q$, letter $\sigma \in \Sigma$, and states $q_1, q_2 \in \delta(q, \sigma)$, the set of words accepted from \mathcal{A} with initial state q_1 is equal to the set of words accepted from \mathcal{A} with initial state q_2 . Since all nondeterministic choices lead to equivalent states, one may be tempted to think that SD automata are DBP or at least have similar properties to deterministic automata. This is indeed the case for SD-NFWs, namely when one considers automata on finite words. There, it is not hard to prove that any pruning of an SD-NFW to a DFW results in an equivalent DFW. Thus, SD-NFWs are not more succinct than DFWs, and operations on them are not more complex than operations on DFWs.

Once, however, one considers automata on infinite words, the simplicity of SD automata is lost. In order to understand the picture in the setting of infinite words, let us elaborate some more on HD automata, which are strongly related to SD automata. Formally, a nondeterministic automaton \mathcal{A} is HD if there is a strategy g that maps each finite word $u \in \Sigma^*$ to a transition to be taken after u is read, and following g results in accepting all the words in the language of \mathcal{A} . Obviously, every DBP automaton is HD – the strategy g can suggest the same transition in all its visits in a state. On the other hand, while HD-NWWs are always DBP [24, 32], this is not the case for HD-NBW

¹The notion used in [17] is *good for games* (GFG) automata, as they address the difficulty of playing games on top of a nondeterministic automaton. As it turns out, the property of being good for games varies in different settings and HD is good for applications beyond games. Therefore, we use the term *history determinism*, introduced by Colcombet in the setting of quantitative automata with cost functions [10].

and HD-NCWs [6]. There, the HD strategy may need to suggest different transitions in visits (with different histories) to the same state.

It is easy to see that a strategy g as above cannot choose a transition to states whose language is strictly contained in the language of states that are reachable by other transitions. Thus, all HD automata can be pruned in polynomial time to SD automata [19, 4]. The other direction, however, does not hold: it is shown in [2] that an SD-NWW need not be HD (hence, SD-NBWs and SD-NCWs need not be HD too). Moreover, while all all HD-NWWs are DBP, this is not the case for all SD-NWWs.

In this work we study the succinctness of SD automata with respect to deterministic ones, as well as the complexity of operations on them and decision problems about them. Our goal is to understand the difference between determinism and semantic determinism, and to understand how this difference varies among different acceptance conditions. Our study is further motivated by the applications of automata with different levels of nondeterminism in algorithms for synthesis and for reasoning in a stochastic setting. In particular, beyond the connection to HD automata discussed above, as runs of an SD-NBW on words in the language are accepting with probability 1, all SD-NBWs are GFM [2].

We study semantic determinism for Büchi, co-Büchi, and weak automata. We consider automata with both *state-based* acceptance conditions, as defined above, and *transition-based* acceptance conditions. In the latter, the acceptance condition is given by a subset α of transitions, and a run is required to traverse transitions in α infinitely often (in Büchi automata, termed tD/tNBW) or finitely often (in co-Büchi automata, termed tD/tNCW). As it turns out, especially in the context of HD automata, automata with transition-based acceptance conditions may differ in their properties from automata with traditional state-based acceptance conditions. For example, while HD-tNCWs can be minimized in PTIME [1], minimization of HD-NCWs is NP-complete [37]. In addition, there is recently growing use of transition-based automata in practical applications, with evidences they offer a simpler translation of LTL formulas to automata and enable simpler constructions and decision procedures [13, 14, 12, 39, 29]. Our results for all types of acceptance conditions are summarized in Table 1 below, where we also compare them with known results about deterministic and HD automata.

	Deterministic	HD	SD
Succinctness	n (B,C,W)	n^2 (B), 2^n (C), n (W) [19, 24, 32]	2^n (B,C), n (W) Theorems 3.3, 4.3, 5.3
Complementation	n (B,C,W) [21]	n (B), 2^n (C), n (W) [19]	2^n (B,C), n (W) Theorems 3.4, 4.4, 5.4
Universality	NL (B,C,W) [21]	P (B,C,W) [16, 19]	PSPACE (B,C), P (W) Theorems 3.6, 4.5, 5.5
Minimization (state based)	NP (B,C), P (W) [36, 30]	NP (B,C), P (W) [37, 24, 30]	PSPACE (B,C), P (W) Theorems 3.7, 4.6, 5.5
Minimization (transition based)	open (B,C), P (W) [30]	open (B), P (C,W) [1, 24, 30]	PSPACE (B,C), P (W) Theorems 3.7, 4.6, 5.5

Table 1: Succinctness (determinization blow-up), complementation (blow-up in going from an automaton to a complementing one), universality (deciding whether an automaton accepts all words), and minimization (deciding whether an equivalent automaton of a given size exists, and the described results apply also for the case the given automaton is deterministic). All blow-ups are tight, except for HD-NBW determinization, where the quadratic bound has no matching lower bound; all NL, NP, and PSPACE bounds are complete.

Let us highlight the results we find the most interesting and surprising. While all three types of SD automata are not DBP, we are able to determinize SD-NWWs in polynomial time, and end up with a DWW whose state space is a subset of the state space of the original SD-NWW. Essentially, rather than pruning transitions, the construction redirects transitions to equivalent states in deep strongly

connected components of the SD-NWW, which we prove to result in an equivalent deterministic DWW². This suggests that despite the “not DBP anomaly” of SD-NWWs, they are very similar in their properties to DWWs. On the other hand, except for their expressive power, SD-NBWs and SD-NCWs are not at all similar to DBWs and DCWs: while HD-NBWs are only quadratically more succinct than DBWs, succinctness jumps to exponential for SD-NBWs. This also shows that SD-NBWs may be exponentially more succinct than HD-NBWs, and we show this succinctness gap also for co-Büchi automata, where exponential succinctness with respect to DCWs holds already in the HD level.

The succinctness results are carried over to the blow-up needed for complementation, and to the complexity of decision procedures. Note that this is not the case for HD automata. There, for example, complementation of HD-NBWs results in an automaton with the same number of states [19]. Moreover, even though HD-NCWs are exponentially more succinct than DCWs, language-containment for HD-NCWs can be solved in PTIME [16, 19], and HD-tNCWs can be minimized in PTIME [1]. For SD automata, we show that complementation involves an exponential blow-up. Also, universality and minimization of SD Büchi and co-Büchi automata, with either state-based or transition-based acceptance conditions, is PSPACE-complete, as it is for NBWs and NCWs. We also study the *D-to-SD minimization problem*, where we are given a deterministic automaton \mathcal{A} and a bound $k \geq 1$, and need to decide whether \mathcal{A} has an equivalent SD automaton with at most k states. Thus, the given automaton is deterministic, rather than SD. By [18], the D-to-N minimization problem for automata on finite words (that is, given a DFW, minimize it to an NFW) is PSPACE-complete. It is easy to see that the D-to-SD minimization problem for automata on finite words can be solved in PTIME. We show that while this is the case also for weak automata, D-to-SD minimization for Büchi and co-Büchi automata is PSPACE-complete.

Our results show that in terms of combinatorial and computational properties, semantic determinism is very similar to determinism in weak automata, whereas for Büchi and co-Büchi automata, semantic determinism is very similar to nondeterminism. The results are interesting also in comparison with history determinism, whose affect on Büchi and co-Büchi automata is different.

2 Preliminaries

2.1 Languages and Automata

For a finite nonempty alphabet Σ , an infinite *word* $w = \sigma_1 \cdot \sigma_2 \cdots \in \Sigma^\omega$ is an infinite sequence of letters from Σ . A *language* $L \subseteq \Sigma^\omega$ is a set of words. For $1 \leq i \leq j$, we use $w[i, j]$ to denote the infix $\sigma_i \cdot \sigma_{i+1} \cdots \sigma_j$ of w , use $w[i]$ to denote the letter σ_i , and use $w[i, \infty]$ to denote the infinite suffix $\sigma_i \cdot \sigma_{i+1} \cdots$ of w . We also consider languages $R \subseteq \Sigma^*$ of finite words, denote the empty word by ϵ , and denote the set of nonempty words over Σ by Σ^+ ; thus $\Sigma^+ = \Sigma^* \setminus \{\epsilon\}$. For a set S , we denote its complement by \bar{S} . In particular, for languages $R \subseteq \Sigma^*$ and $L \subseteq \Sigma^\omega$, we have $\bar{R} = \Sigma^* \setminus R$ and $\bar{L} = \Sigma^\omega \setminus L$.

A *nondeterministic automaton* is a tuple $\mathcal{A} = \langle \Sigma, Q, Q_0, \delta, \alpha \rangle$, where Σ is an alphabet, Q is a finite set of *states*, Q_0 is a set of *initial states*, $\delta : Q \times \Sigma \rightarrow 2^Q \setminus \emptyset$ is a *transition function*, and α is an *acceptance condition*, to be defined below. For states q and s and a letter $\sigma \in \Sigma$, we say that s is a σ -*successor* of q if $s \in \delta(q, \sigma)$. Note that the transition function of \mathcal{A} is total, thus for all states $q \in Q$ and letters $\sigma \in \Sigma$, q has at least one σ -successor. If $|Q_0| = 1$ and every state q has a single σ -successor, for all letters σ , then \mathcal{A} is *deterministic*. The transition function δ can be viewed as a transition relation $\Delta \subseteq Q \times \Sigma \times Q$, where for every two states $q, s \in Q$ and letter $\sigma \in \Sigma$, we have

²Our result implies that checking language-equivalence between states in an SD-NWW can be done in PTIME, as the check can be performed on the equivalent DWWs. We cannot, however, use this complexity result in our algorithm, as this involves a circular dependency. Consequently, our construction of the equivalent DWWs involves a language-approximation argument.

that $\langle q, \sigma, s \rangle \in \Delta$ iff $s \in \delta(q, \sigma)$. We define the *size* of \mathcal{A} , denoted $|\mathcal{A}|$, as its number of states, thus, $|\mathcal{A}| = |Q|$.

Given an input word $w = \sigma_1 \cdot \sigma_2 \cdots$, a *run* of \mathcal{A} on w is a sequence of states $r = r_0, r_1, r_2, \dots$, such that $r_0 \in Q_0$, and for all $i \geq 0$, we have that $r_{i+1} \in \delta(r_i, \sigma_{i+1})$, i.e., the run starts in some initial state and proceeds according to the transition function. If the word in the input is infinite, then so is the run. We sometimes view the run $r = r_0, r_1, r_2, \dots$ on $w = \sigma_1 \cdot \sigma_2 \cdots$ as a sequence of successive transitions $\langle r_0, \sigma_1, r_1 \rangle, \langle r_1, \sigma_2, r_2 \rangle, \dots$. Note that a deterministic automaton has a single run on an input w . We sometimes extend δ to sets of states and finite words. Then, $\delta : 2^Q \times \Sigma^* \rightarrow 2^Q$ is such that for every $S \in 2^Q$, finite word $u \in \Sigma^*$, and letter $\sigma \in \Sigma$, we have that $\delta(S, \epsilon) = S$, $\delta(S, \sigma) = \bigcup_{s \in S} \delta(s, \sigma)$, and $\delta(S, u \cdot \sigma) = \delta(\delta(S, u), \sigma)$. Thus, $\delta(S, u)$ is the set of states that \mathcal{A} may reach when it reads u from some state in S .

The acceptance condition α determines which runs are “good”. For automata on finite words, $\alpha \subseteq Q$, and a run is *accepting* if it ends in a state in α . For automata on infinite words, we consider *state-based* and *transition-based* acceptance conditions. Let us start with *state-based* conditions. Here, $\alpha \subseteq Q$, and we use the terms α -states and $\bar{\alpha}$ -states to refer to states in α and in $Q \setminus \alpha$, respectively. For a run $r \in Q^\omega$, let $\text{sinf}(r) \subseteq Q$ be the set of states that r visits infinitely often. Thus, $\text{sinf}(r) = \{q : q = r_i \text{ for infinitely many } i\}$. In *Büchi* automata, r is *accepting* iff $\text{sinf}(r) \cap \alpha \neq \emptyset$, thus if r visits states in α infinitely often. Dually, in *co-Büchi* automata, r is *accepting* iff $\text{sinf}(r) \cap \alpha = \emptyset$, thus if r visits states in α only finitely often.

We proceed to *transition-based* conditions. There, $\alpha \subseteq \Delta$ and acceptance depends on the set of transitions that are traversed infinitely often during the run. We use the terms α -transitions and $\bar{\alpha}$ -transitions to refer to transitions in α and in $\Delta \setminus \alpha$, respectively. For a run $r \in \Delta^\omega$, we define $\text{tinf}(r) = \{\langle q, \sigma, s \rangle \in \Delta : q = r_i, \sigma = \sigma_{i+1}, \text{ and } s = r_{i+1}, \text{ for infinitely many } i\}$. As expected, in transition-based Büchi automata, r is accepting iff $\text{tinf}(r) \cap \alpha \neq \emptyset$, and in transition-based co-Büchi automata, r is accepting iff $\text{tinf}(r) \cap \alpha = \emptyset$.

Consider an automaton $\mathcal{A} = \langle \Sigma, Q, Q_0, \delta, \alpha \rangle$. In all automata classes, a run of \mathcal{A} that is not accepting is *rejecting*. A word w is accepted by an automaton \mathcal{A} if there is an accepting run of \mathcal{A} on w . The language of \mathcal{A} , denoted $L(\mathcal{A})$, is the set of words that \mathcal{A} accepts.

Consider a directed graph $G = \langle V, E \rangle$. A *strongly connected set* in G (SCS, for short) is a set $C \subseteq V$ such that for every two vertices $v, v' \in C$, there is a path from v to v' . An SCS is *maximal* if for every non-empty set $C' \subseteq V \setminus C$, it holds that $C \cup C'$ is not an SCS. The *maximal strongly connected sets* are also termed *strongly connected components* (SCCs, for short). An automaton $\mathcal{A} = \langle \Sigma, Q, Q_0, \delta, \alpha \rangle$ induces a directed graph $G_{\mathcal{A}} = \langle Q, E \rangle$, where $\langle q, q' \rangle \in E$ iff there is a letter $\sigma \in \Sigma$ such that $\langle q, \sigma, q' \rangle \in \Delta$. The SCSs and SCCs of \mathcal{A} are those of $G_{\mathcal{A}}$.

An automaton $\mathcal{A} = \langle \Sigma, Q, Q_0, \delta, \alpha \rangle$ with a state-based acceptance condition $\alpha \subseteq Q$ is *weak* [33] if for each SCC C of \mathcal{A} , either $C \subseteq \alpha$ in which case C is *accepting*, or $C \cap \alpha = \emptyset$ in which case C is *rejecting*. We view \mathcal{A} as a Büchi automaton, yet note that a weak automaton can be viewed as both a Büchi and a co-Büchi automaton. Indeed, a run of \mathcal{A} visits α infinitely often iff it gets trapped in an SCC that is contained in α iff it visits states in $Q \setminus \alpha$ only finitely often. Note also that when \mathcal{A} uses a transition-based acceptance condition, we can ignore the membership in α of transitions between SCCs (indeed, such transitions are traversed only finitely often), and say that \mathcal{A} is weak if the transitions in each SCC are all in α or all disjoint from α .

Consider two automata \mathcal{A}_1 and \mathcal{A}_2 . We say that \mathcal{A}_1 is *contained in* \mathcal{A}_2 if $L(\mathcal{A}_1) \subseteq L(\mathcal{A}_2)$. Then, \mathcal{A}_1 and \mathcal{A}_2 are *equivalent* if $L(\mathcal{A}_1) = L(\mathcal{A}_2)$, and \mathcal{A}_1 is *universal* if $L(\mathcal{A}_1) = \Sigma^\omega$ (or $L(\mathcal{A}_1) = \Sigma^*$, in case it runs on finite words). Finally, \mathcal{A}_1 is *minimal* (with respect to a class of automata, say state-based deterministic Büchi automata) if for all automata \mathcal{A}_2 equivalent to \mathcal{A}_1 , we have that $|\mathcal{A}_1| \leq |\mathcal{A}_2|$.

2.2 SD and HD Automata

Consider an automaton $\mathcal{A} = \langle \Sigma, Q, Q_0, \delta, \alpha \rangle$. For a state $q \in Q$ of \mathcal{A} , we define $\mathcal{A}^q = \langle \Sigma, Q, \{q\}, \delta, \alpha \rangle$, as the automaton obtained from \mathcal{A} by setting the set of initial states to be $\{q\}$. We say that two states $q, s \in Q$ are *equivalent*, denoted $q \sim_{\mathcal{A}} s$, if $L(\mathcal{A}^q) = L(\mathcal{A}^s)$. We say that q is *reachable* if there is a finite word $x \in \Sigma^*$ with $q \in \delta(Q_0, x)$, and say that q is *reachable from* s if q is reachable in \mathcal{A}^s . We say that \mathcal{A} is *semantically deterministic* (SD, for short) if different nondeterministic choices lead to equivalent states. Thus, all initial states are equivalent, and for every state $q \in Q$ and letter $\sigma \in \Sigma$, all the σ -successors of q are equivalent. Formally, if $s, s' \in \delta(q, \sigma)$, then $s \sim_{\mathcal{A}} s'$. The following proposition, termed the *SDness property*, follows immediately from the definitions and implies that in SD automata, for all finite words x , all the states in $\delta(Q_0, x)$ are equivalent. Intuitively, it means that a run of an SD automaton can take also bad nondeterministic choices, as long as it does so only finitely many times.

Proposition 2.1. *Consider an SD automaton $\mathcal{A} = \langle \Sigma, Q, Q_0, \delta, \alpha \rangle$, and states $q, s \in Q$. If $q \sim_{\mathcal{A}} s$, then for every $\sigma \in \Sigma$, $q' \in \delta(q, \sigma)$, and $s' \in \delta(s, \sigma)$, we have that $q' \sim_{\mathcal{A}} s'$.*

A nondeterministic automaton \mathcal{A} is *history deterministic* (HD, for short) if its nondeterminism can be resolved based on the past, thus on the prefix of the input word read so far. Formally, \mathcal{A} is HD if there exists a *strategy* $f : \Sigma^* \rightarrow Q$ such that the following hold:

1. The strategy f is consistent with the transition function. That is, $f(\epsilon) \in Q_0$, and for every finite word $u \in \Sigma^*$ and letter $\sigma \in \Sigma$, we have that $f(u \cdot \sigma) \in \delta(f(u), \sigma)$.
2. Following f causes \mathcal{A} to accept all the words in its language. That is, for every word $w = \sigma_1 \cdot \sigma_2 \cdots$, if $w \in L(\mathcal{A})$, then the run $f(\epsilon), f(w[1, 1]), f(w[1, 2]), \dots$ is an accepting run of \mathcal{A} on w .

We say that the strategy f *witnesses* \mathcal{A} 's HDness. Note that, by definition, we can assume that every SD and HD automaton \mathcal{A} has a single initial state. Thus, we sometimes abuse notation and write \mathcal{A} as $\mathcal{A} = \langle \Sigma, Q, q_0, \delta, \alpha \rangle$, where q_0 is the single initial state of the SD (or HD) automaton \mathcal{A} .

For an automaton \mathcal{A} , we say that a state q of \mathcal{A} is *HD*, if \mathcal{A}^q is HD. Note that every deterministic automaton is HD. Also, while not all HD automata can be pruned to deterministic ones [6], removing of transitions that are not used by a strategy f that witnesses \mathcal{A} 's HDness does not reduce the language of \mathcal{A} and results in an SD automaton. Moreover, since every state that is used by f is HD, the removal of non-HD states does not affect \mathcal{A} 's language nor its HDness. Accordingly, we have the following [19, 4].

Proposition 2.2. *Every HD automaton \mathcal{A} can be pruned to an equivalent SD automaton all whose states are HD.*

We use three-letter acronyms in $\{\text{D, N}\} \times \{\text{B, C, W, F}\} \times \{\text{W}\}$ to denote the different automata classes. The first letter stands for the branching mode of the automaton (deterministic or nondeterministic); the second for the acceptance condition type (Büchi, co-Büchi, weak, or an automaton that runs on finite inputs); and the third indicates that we consider automata on words. For transition-based automata, we start the acronyms with the letter “t”, and for HD or SD automata, we add an HD or SD prefix, respectively. For example, an HD-tNBW is a transition-based HD nondeterministic Büchi automaton, and an SD-NWW is a state-based weak SD automaton.

3 Semantically Deterministic Büchi Automata

In this section we examine SD-tNBWs and SD-NBW. Our results use the following definitions and constructions: For a language $R \subseteq \Sigma^*$ of finite words, we use ∞R to denote the language of infinite

words that contain infinitely many disjoint infixes in R . Thus, $w \in \infty R$ iff $\epsilon \in R$ or there are infinitely many indices $i_1 \leq i'_1 < i_2 \leq i'_2 < \dots$ such that $w[i_j, i'_j] \in R$, for all $j \geq 1$. For example, taking $\Sigma = \{a, b\}$, we have that $\infty\{ab\}$ is the language of words with infinitely many ab infixes, namely all words with infinitely many a 's and infinitely many b 's. We say that a finite word $x \in \Sigma^*$ is a *good prefix* for a language $R \subseteq \Sigma^*$ if for all finite words $y \in \Sigma^*$, we have that $x \cdot y \in R$. For example, while the language $(a + b)^* \cdot a$ does not have a good prefix, the word a is a good prefix for the language $a \cdot (a + b)^*$.

Theorem 3.1 below states that one can translate an NFW that recognizes a language $R \subseteq \Sigma^*$ to an SD-tNBW for ∞R . Theorem 3.2 is more complicated and suggests that this translation can be viewed as a reversible encoding, in the sense that one can obtain an NFW for R from an SD-tNBW for $\infty(\$ \cdot R \cdot \$)$, where $\$ \notin \Sigma$. The translations involve no blow up, and both theorems play a major role in the rest of this section.

Theorem 3.1. *Given an NFW \mathcal{N} , one can obtain, in polynomial time, an SD-tNBW \mathcal{A} such that $L(\mathcal{A}) = \infty L(\mathcal{N})$ and $|\mathcal{A}| = |\mathcal{N}|$.*

Proof. Let $\mathcal{N} = \langle \Sigma, Q, Q_0, \delta, F \rangle$. Then, $\mathcal{A} = \langle \Sigma, Q, Q_0, \delta', \alpha \rangle$ is obtained from \mathcal{N} by adding transitions to Q_0 from all states with all letters. A new σ -transition is in α if $Q_0 \cap F \neq \emptyset$ or when \mathcal{N} could transit with σ to a state in F . Formally, for all $s \in Q$ and $\sigma \in \Sigma$, we have that $\delta'(s, \sigma) = \delta(s, \sigma) \cup Q_0$, and $\alpha = \{ \langle s, \sigma, q \rangle : q \in Q_0 \text{ and } (Q_0 \cap F \neq \emptyset \text{ or } \delta(s, \sigma) \cap F \neq \emptyset) \}$.

It is easy to see that $|\mathcal{A}| = |\mathcal{N}|$. In order to prove that \mathcal{A} is SD and $L(\mathcal{A}) = \infty L(\mathcal{N})$, we prove next that for every state $q \in Q$, it holds that $L(\mathcal{A}^q) = \infty L(\mathcal{N})$. We distinguish between two cases. First, if $Q_0 \cap F \neq \emptyset$, then $\epsilon \in L(\mathcal{N})$, in which case $\infty L(\mathcal{N}) = \Sigma^\omega$. Indeed, every word in Σ^ω has infinitely many empty infixes. Recall that when $Q_0 \cap F \neq \emptyset$, the tNBW \mathcal{A} has Σ -labeled α transitions from every state to Q_0 . Then, the run $q, (q_0)^\omega$, where $q_0 \in Q_0$, is an accepting run of \mathcal{A}^q on any word $w \in \Sigma^\omega$. So, $L(\mathcal{A}^q) = \Sigma^\omega = \infty L(\mathcal{N})$, and we are done.

Now, if $Q_0 \cap F = \emptyset$, we first prove that for every state $q \in Q$, we have that $\infty L(\mathcal{N}) \subseteq L(\mathcal{A}^q)$. Consider a word $w \in \infty L(\mathcal{N})$, and consider a run of \mathcal{A}^q on w that upon reading an infix $\sigma_1 \cdot u \cdot \sigma_2$ such that $u \cdot \sigma_2 \in L(\mathcal{N}^{q_0}) \subseteq L(\mathcal{N})$, for some $q_0 \in Q_0$, moves to q_0 while reading the letter σ_1 , and then follows an accepting run of \mathcal{N} on $u \cdot \sigma_2$ except that, upon reading the letter σ_2 , it moves to a state in Q_0 , traversing an α transition, instead of traversing a transition to F . Note that the transitions in \mathcal{A} enable such a behavior. Also note that $\infty L(\mathcal{N}) \subseteq \infty \Sigma \cdot L(\mathcal{N})$, and so such a run traverses α infinitely often and is thus accepting.

It is left to prove that $L(\mathcal{A}^q) \subseteq \infty L(\mathcal{N})$ for every state $q \in Q$. Let $r = r_0, r_1, \dots$ be an accepting run of \mathcal{A}^q on a word $w = \sigma_1 \cdot \sigma_2 \cdot \dots$. We show that $w \in \infty L(\mathcal{N})$. For that, we show that for every $i \geq 0$, there are k_1, k_2 such that $i \leq k_1 \leq k_2$ and $w[k_1 + 1, k_2 + 1]$ is an infix of w in $L(\mathcal{N})$. Being accepting, the run r traverses infinitely many α -transitions. Given $i \geq 0$, let $k_2 \geq i$ be such that $t = \langle r_{k_2}, \sigma_{k_2+1}, r_{k_2+1} \rangle$ is an α -transition, but not the first α -transition that r traverses when it reads the suffix $w[i + 1, \infty]$, and let k_1 be the maximal index such that $i \leq k_1 \leq k_2$ and $r_{k_1} \in Q_0$. Since t is not the first α -transition that r traverses when it reads the suffix $w[i + 1, \infty]$, and since α -transitions lead to Q_0 , then k_1 exists. Note that the maximality of k_1 imply that the run $r_{k_1}, r_{k_1+1}, \dots, r_{k_2}$ is a run of \mathcal{N} on the infix $w[k_1 + 1, k_2]$. Now, as $\langle r_{k_2}, \sigma_{k_2+1}, r_{k_2+1} \rangle \in \alpha$ and $Q_0 \cap F = \emptyset$, it follows by the definition of δ' that there is a state $s \in \delta(r_{k_2}, \sigma_{k_2+1}) \cap F$. Hence, $r_{k_1}, r_{k_1+1}, \dots, r_{k_2}, s$ is an accepting run of \mathcal{N} on the infix $w[k_1 + 1, k_2 + 1]$, which is thus in $L(\mathcal{N})$. \square

We proceed to the other direction, translating an automaton for $\infty(\$ \cdot R \cdot \$)$ to an automaton for R . Note that the language R is well defined. That is, if $R_1 \subseteq \Sigma^*$ and $R_2 \subseteq \Sigma^*$ are such that $\infty(\$ \cdot R_1 \cdot \$) = \infty(\$ \cdot R_2 \cdot \$)$, then it must be that $R_1 = R_2$.

Theorem 3.2. *Consider a language $R \subseteq \Sigma^*$ and a letter $\$ \notin \Sigma$. For every SD-tNBW \mathcal{A} such that $L(\mathcal{A}) = \infty(\$ \cdot R \cdot \$)$, there exists an NFW \mathcal{N} such that $L(\mathcal{N}) = R$ and $|\mathcal{N}| \leq |\mathcal{A}| + 1$. In addition, if R has no good prefixes, then $|\mathcal{N}| \leq |\mathcal{A}|$.*

Proof. If R is trivial, then one can choose \mathcal{N} to be a one-state NFW. Assume that R is nontrivial. Let $\mathcal{A} = \langle \Sigma \cup \{\$\}, Q, q_0, \delta, \alpha \rangle$ be an SD-tNBW for $\infty(\$ \cdot R \cdot \$)$, W.l.o.g we assume that all the states of \mathcal{A} are reachable. For a nonempty set of states $S \in 2^Q \setminus \emptyset$, we define the *universal $\bar{\alpha}$ language* of S as $L_{u\bar{\alpha}}(S) = \{w \in (\Sigma \cup \{\$\})^\omega : \text{for all } q \in S, \text{ all the runs of } \mathcal{A}^q \text{ on } w \text{ do not traverse } \alpha\}$. We say that S is *hopeful* when $(\$ \cdot \bar{R})^\omega \subseteq L_{u\bar{\alpha}}$. Note that S is hopeful iff for every state $q \in S$, it holds that $\{q\}$ is hopeful. Also, if S is hopeful, $x \in \Sigma^*$ is a finite word, and there is a run from S on $\$ \cdot x$ that traverses α , then $x \in R$. Then, we say that S is *good* when for all words $x \in \Sigma^*$, it holds that $x \in \bar{R}$ iff all the runs from S on $\$ \cdot x$ do not traverse α , and the set $\delta(S, \$ \cdot x)$ is hopeful. Note that as R is nontrivial, there exists a word x in \bar{R} , and thus by definition, all the $\$$ -labeled transitions going out from a good set S are in $\bar{\alpha}$.

Note that good sets in \mathcal{A} characterize the language R . We first show that a good set exists. Assume towards contradiction that every nonempty set $S \in 2^Q \setminus \emptyset$ is not good. We define iteratively a sequence S_1, S_2, S_3, \dots of nonempty sets in 2^Q , and a sequence x_1, x_2, x_3, \dots of finite words in $(\Sigma \cup \{\$\})^+$, as follows. The set S_1 is chosen arbitrarily. For all $i \geq 1$, given the set S_i , we define x_i and S_{i+1} as follows. By the assumption, S_i is not good. Therefore at least one of the following holds:

- (a) There is a witness in R that S_i is not good: there is a word $a_i \in R$ such that all the runs from S_i on $\$ \cdot a_i$ do not traverse α , and the set $\delta(S_i, \$ \cdot a_i)$ is hopeful.
- (b) There is a witness in \bar{R} that S_i is not good: there is a word $b_i \in \bar{R}$ such that there is a run from S_i on $\$ \cdot b_i$ that traverses α , or the set $\delta(S_i, \$ \cdot b_i)$ is not hopeful.

If there is a witness $a_i \in R$ that S_i is not good, we take $x_i = \$ \cdot a_i$ and $S_{i+1} = \delta(S_i, x_i)$. Otherwise, there is a witness $b_i \in \bar{R}$ that S_i is not good. Then, as we argue below, there exists a word $z_i \in (\$ \cdot \bar{R})^+$, and there is a run r_i from S_i on z_i that traverses α . In this case, we take S_{i+1} to be the set that contains only the state that r_i ends in, and we take $x_i = z_i$. To see why such z_i exists, we distinguish between two cases. First, if there is a run from S_i on $\$ \cdot b_i$ that traverses α , then we take $z_i = \$ \cdot b_i$. Otherwise, the set $\delta(S_i, \$ \cdot b_i)$ is not hopeful. Therefore, it contains a state q that has a run that traverses α on a word in $(\$ \cdot \bar{R})^\omega$. The latter implies that there is a word $z' \in (\$ \cdot \bar{R})^+$ such that there is a finite run r' of q on z' that traverses α . Then, since q is reachable from S_i upon reading $\$ \cdot b_i$, we take $z_i = \$ \cdot b_i \cdot z'$. Indeed, a run r_i from S_i on $z_i = \$ \cdot b_i \cdot z'$ that traverses α , starts from S_i and reaches q upon reading $\$ \cdot b_i$, and then follows the run r' upon reading z' .

So, let S_1, S_2, S_3, \dots and x_1, x_2, x_3, \dots be sequences as above. We distinguish between three cases:

1. There exists $i \geq 1$ such that for all $j \geq i$, there is a witness $a_j \in R$ that S_j is not good: consider $j \geq i$. In this case, $x_j = \$ \cdot a_j \in \$ \cdot R$, $S_{j+1} = \delta(S_j, x_j)$, and all the runs of S_j on x_j do not traverse α . Hence, all the runs from S_i on the word $w = x_i \cdot x_{i+1} \cdot x_{i+2} \cdots \in (\$ \cdot R)^\omega$ do not traverse α , in particular, if we consider a state $q \in S_i$, we get that the word w is not accepted from q . As all the states of \mathcal{A} are reachable, there is a word $y \in (\Sigma \cup \{\$\})^*$ such that $q \in \delta(q_0, y)$. Then, the SDness property of \mathcal{A} implies that all the states in $\delta(q_0, y)$ do not accept w , and hence the word $y \cdot w$ is not accepted by \mathcal{A} even though it has infinitely many infixes in $\$ \cdot R \cdot \$$, and we have reached a contradiction.
2. There exists $i \geq 1$ such that for all $j \geq i$, no word in R witnesses that S_j is not good: consider $j \geq i$. In this case, as argued above, there is a witness $b_j \in \bar{R}$ that S_j is not good. So, $x_j = z_j \in (\$ \cdot \bar{R})^+$, and there is a run r_j from S_j on z_j that traverses α and ends in a state s_{j+1} , where $S_{j+1} = \{s_{j+1}\}$. The concatenation of the runs r_i, r_{i+1}, \dots is an accepting run of \mathcal{A}^q , for some $q \in S_i$, on the word $z_i \cdot z_{i+1} \cdot z_{i+2} \cdots \in (\$ \cdot \bar{R})^\omega$. Let $y \in (\Sigma \cup \{\$\})^*$ be such that $q \in \delta(q_0, y)$; then the word $w = y \cdot z_i \cdot z_{i+1} \cdot z_{i+2} \cdots$ is accepted by \mathcal{A} even though it has only finitely many infixes in $\$ \cdot R \cdot \$$, and we have reached a contradiction.
3. If the above two cases do not hold, then there are infinitely many indices $i \geq 1$ such that no word in R witnesses that S_i is not good, and there are infinitely many indices $i \geq 1$ such that

there is a witness in R that S_i is not good. Therefore, there must be $i \geq 1$ such that there is a witness $a_i \in R$ that S_i is not good, and no word in R witnesses that S_{i+1} is not good. On the one hand, the set $S_{i+1} = \delta(S_i, \$ \cdot a_i)$ is hopeful. On the other hand, the fact that no word in R witnesses that S_{i+1} is not good, implies that there is a witness $b_{i+1} \in \bar{R}$ that S_{i+1} is not good, and thus there is a run from S_{i+1} that traverses α on the word $x_{i+1} = z_{i+1} \in (\$ \cdot \bar{R})^+$. Thus, we have reached a contradiction to the fact that S_{i+1} is hopeful.

We show now that a good set in \mathcal{A} induces an NFW \mathcal{N} for R with the required properties. Let $S \in 2^Q \setminus \emptyset$ be a good set. We define the NFW $\mathcal{N} = \langle \Sigma, Q \cup \{q_{acc}\}, Q_0^S, \delta_S, F_S \rangle$, where $Q_0^S = \delta(S, \$)$, $F_S = \{q_{acc}\} \cup \{q \in Q : \text{the set } \{q\} \text{ is not hopeful}\}$, and the transition function δ_S is defined as follows. For every two states $q, s \in Q$ and letter $\sigma \in \Sigma$, it holds that $s \in \delta_S(q, \sigma)$ iff $\langle q, \sigma, s \rangle \in \bar{\alpha}$. Also, $q_{acc} \in \delta_S(q, \sigma)$ iff there is a σ -labeled α -transition going out from q in \mathcal{A} . Also, for all letters $\sigma \in \Sigma$, it holds that $\delta(q_{acc}, \sigma) = \{q_{acc}\}$; that is, q_{acc} is an accepting sink. Thus, \mathcal{N} behaves as the states in $\delta(S, \$)$ as long as it reads $\bar{\alpha}$ transitions of \mathcal{A} , moves to the accepting sink q_{acc} whenever an α -transition is encountered, and accepts also whenever it reaches a state in Q that is not hopeful.

We prove next that $L(\mathcal{N}) = R$. We first prove that $L(\mathcal{N}) \subseteq R$. Consider a word $x = \sigma_1 \cdot \sigma_2 \cdots \sigma_n \in L(\mathcal{N})$. If there is a state $q \in \delta(S, \$)$ such that \mathcal{A}^q has a run on x that traverses α , then as S is good, we get that $x \in R$. Otherwise, for all states $q \in \delta(S, \$)$, it holds that all the runs of \mathcal{A}^q on x do not traverse α . Let $r = r_0, r_1, \dots, r_n$ be an accepting run of \mathcal{N} on x . By the definition of \mathcal{N} , we have that $r_0 \in \delta(S, \$)$; in particular, all the runs of \mathcal{A}^{r_0} on x do not traverse α . Hence, since r follows a run of \mathcal{A}^{r_0} as long it does not visit q_{acc} , and since r visits q_{acc} only when it follows a run of \mathcal{A}^{r_0} that traverses α , we get that the run r exists in \mathcal{A} . Hence, $r_n \in Q$ and as r is an accepting run of \mathcal{N} , we get that the set $\{r_n\}$ is not hopeful. The fact that r exists in \mathcal{A} implies also that $r_n \in \delta(S, \$ \cdot x)$, and so $\delta(S, \$ \cdot x)$ is not hopeful. Hence, as S is a good set, we conclude that $x \in R$.

For the other direction, namely $R \subseteq L(\mathcal{N})$, consider a word $x \in R$. If there is a state $q \in \delta(S, \$)$ and a run r of \mathcal{A}^q on x that traverses α , then an accepting run r' of \mathcal{N} on x can be obtained by following the run r , and moving to the accepting sink q_{acc} when r traverses α for the first time. Indeed, $q \in Q_0^S$, $\bar{\alpha}$ transitions of \mathcal{A} exist in \mathcal{N} , and \mathcal{N} moves to q_{acc} whenever \mathcal{A} traverses α . Otherwise, for all $q \in \delta(S, \$)$, it holds that all the runs of \mathcal{A}^q on x do not traverse α . Also, recall that all the $\$$ -labeled transitions from a good set are in $\bar{\alpha}$. Hence, all the runs from S on $\$ \cdot x$ do not traverse α . In this case, as S is a good set, and $x \in R$, it follows that $\delta(S, \$ \cdot x)$ is not hopeful. Therefore, there is a state $s \in \delta(S, \$ \cdot x)$ such that $\{s\}$ is not hopeful. Hence, by the definition of \mathcal{N} , it holds that $s \in F_S$. Hence, a run in \mathcal{A} from $\delta(S, \$)$ that reaches s upon reading x and does not traverse α , is an accepting run of \mathcal{N} on x , and we are done.

Since the state space of \mathcal{N} is $Q \cup \{q_{acc}\}$, then $|\mathcal{N}| = |\mathcal{A}| + 1$. Moreover, as q_{acc} is an accepting sink, a word $x \in L(\mathcal{N})$ that has a run that ends in q_{acc} is a good prefix for $L(\mathcal{N})$. Hence, as $L(\mathcal{N}) = R$, if R has no good prefixes, then q_{acc} is not reachable in \mathcal{N} and thus can be removed without affecting \mathcal{N} 's language. Thus, in this case, we get an NFW \mathcal{N} for R whose size is at most $|\mathcal{A}|$. \square

3.1 Succinctness and Complementation

In this section we study the succinctness of SD Büchi automata with respect to deterministic ones, and the blow-up involved in their complementation. We show that SD-tNBWs are exponentially more succinct than tDBWs, matching the known upper bound [2], and in fact, also from HD-tNBWs. We also prove an exponential lower bound for complementation. Similar results for SD-NBWs follow, as the transition between the two types of acceptance conditions is linear.

Theorem 3.3. *There is a family L_1, L_2, L_3, \dots of languages such that for every $n \geq 1$, there is an SD-tNBW with $3n + 3$ states that recognizes L_n , yet every tDBW or HD-tNBW that recognizes L_n needs at least 2^n states.*

Proof. For $n \geq 1$, let $[n] = \{1, \dots, n\}$, and let $\Sigma_n = \{1, \dots, n, \$, \#\}$. We say that a word $z \in \Sigma_n^*$ is *good* if $z = \$ \cdot x \cdot \# \cdot i$, where $x \in [n]^+$ and i appears in x . Let $R_n \subseteq \Sigma_n^*$ be the language of all good words. We define $L_n = \infty R_n$. First, it is not hard to see that R_n can be recognized by an NFW \mathcal{N}_n with $3n + 3$ states. Essentially, \mathcal{N}_n guesses the last letter i in the input word and then checks that the guess is correct. A sketch of \mathcal{N}_n appears in Figure 1 below. By Theorem 3.1, there is an SD-tNBW for L_n with $3n + 3$ states.

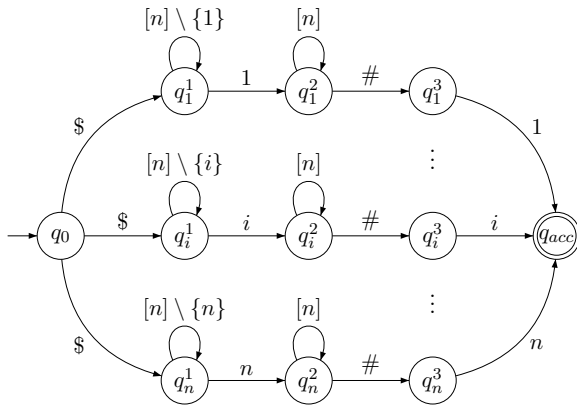


Figure 1: The NFW \mathcal{N}_n . Missing transitions lead to a rejecting-sink

Before we prove that every tDBW or HD-tNBW that recognizes L_n needs at least 2^n states, let us note that it is already known that going from a DFW for a language $R \subseteq \Sigma^*$ to a tDBW for ∞R , may involve a blow-up of $2^{n-2-\log_2(n)}$ [23]. Thus, Theorem 3.1 implies an exponential gap between SD-tNBWs and tDBWs. Moreover, as HD-tNBWs are at most quadratically more succinct than tDBWs [19], the above can be extended to a $2^{\frac{n-2-\log_2(n)}{2}}$ lower bound for the succinctness of SD-tNBWs with respect to HD-tNBWs. Our example here is tighter.

We now prove that every HD-tNBW that recognizes L_n needs at least 2^n states. Assume towards contradiction that $\mathcal{A} = \langle \Sigma_n, Q, q_0, \delta, \alpha \rangle$ is an HD-tNBW for L_n with $|Q| < 2^n$ states. Below, we iteratively define infinite sequences of finite words x_1, x_2, x_3, \dots and states q_0, q_1, \dots such that for all $k \geq 1$, the word x_k starts with $\$$, has no good infixes, and there is a run of the form $r_k = q_{k-1} \xrightarrow{x_k} q_k$ in \mathcal{A} on x_k that traverses α . To see why such sequences imply a contradiction, note that the concatenation of the runs r_1, r_2, \dots is an accepting run of \mathcal{A} on the word $x = x_1 \cdot x_2 \cdot \dots$. As \mathcal{A} recognizes $L_n = \infty R_n$, it follows that $x \in \infty R_n$. On the other hand, x has no good infixes, and so $x \notin \infty R_n$. Indeed, if there is a good infix in x , then it must contain letters from different x_k 's; in particular, it must contain at least two $\$$'s.

We proceed by describing how, given x_1, x_2, \dots, x_{k-1} and q_0, q_1, \dots, q_{k-1} , we define x_k and q_k . The challenging part in the construction is to make it valid also for HD (and not only deterministic) automata. For this, the definition of the words in the sequence is defined with respect to a strategy that attempts to witness the HDness of \mathcal{A} . First, recall that by Proposition 2.2, we can assume that \mathcal{A} is SD and all its state are HD. Also, note that $q_0 \xrightarrow{x_1} q_1 \xrightarrow{x_2} q_2 \cdots \xrightarrow{x_{k-1}} q_{k-1}$ is a run in \mathcal{A} ; in particular, q_{k-1} is reachable in \mathcal{A} . Hence, $L(\mathcal{A}) \subseteq L(\mathcal{A}^{q_{k-1}})$. Indeed, if y is a word in $L(\mathcal{A}) = \infty R_n$, then $x_1 \cdot x_2 \cdots x_{k-1} \cdot y \in \infty R_n$, and so the SDness of \mathcal{A} implies that all the states in $\delta(q_0, x_1 \cdot x_2 \cdots x_{k-1})$; in particular q_{k-1} , accept y . We abuse notation and use S , for a subset $S \subseteq [n]$, to denote a finite word over $[n]$ consisting of exactly the numbers appearing in S . We can now define x_k and q_k . Fix $u_0, d_0 = q_{k-1}$, and let f be a strategy witnessing $\mathcal{A}^{q_{k-1}}$'s HDness. We define inductively two infinite

runs r_u and r_d on distinct words, and of the form

$$\begin{aligned} r_u &= u_0 \xrightarrow{\$.2^{[n]}} p_0 \xrightarrow{\# \cdot [n]} u_1 \xrightarrow{\$.2^{[n]}} p_1 \xrightarrow{\# \cdot [n]} u_2 \xrightarrow{\$.2^{[n]}} \dots \\ r_d &= d_0 \xrightarrow{\$.2^{[n]}} p_0 \xrightarrow{\# \cdot [n]} d_1 \xrightarrow{\$.2^{[n]}} p_1 \xrightarrow{\# \cdot [n]} d_2 \xrightarrow{\$.2^{[n]}} \dots \end{aligned}$$

Thus, while the states u_i and d_i are not necessarily equal, the runs meet infinitely often at the states p_i upon reading infixes of the form $\$.2^{[n]}$. When we define the runs, we make sure that both are runs of \mathcal{A}^{qk-1} that are consistent with the strategy f . Consider $j \geq 0$, and assume that we have defined the following sub-runs $u_0 \xrightarrow{z_u^j} u_j$ and $d_0 \xrightarrow{z_d^j} d_j$ that are consistent with the strategy f . A finite run r of the state u_j on a word $z = z_1 \cdot z_2 \dots z_n$ is *consistent* if it is of the form $r = f(z_u^j), f(z_u^j \cdot z_1), \dots, f(z_u^j z_n)$; that is, r is consistent with the strategy f w.r.t the prefix z_u^j . Similarly, we consider runs of d_j that are consistent with the strategy f w.r.t the prefix z_d^j . As $|Q| < 2^n$, there are subsets $S_{j+1}, T_{j+1} \subseteq [n]$ with $i_{j+1} \in S_{j+1} \setminus T_{j+1}$ such that

$$\left((u_j \xrightarrow{\$.S_{j+1}} p_j \xrightarrow{\# \cdot i_{j+1}} u_{j+1}) \text{ and } (d_j \xrightarrow{\$.T_{j+1}} p_j \xrightarrow{\# \cdot i_{j+1}} d_{j+1}) \right)$$

or

$$\left((u_j \xrightarrow{\$.T_{j+1}} p_j \xrightarrow{\# \cdot i_{j+1}} u_{j+1}) \text{ and } (d_j \xrightarrow{\$.S_{j+1}} p_j \xrightarrow{\# \cdot i_{j+1}} d_{j+1}) \right)$$

where the runs are consistent. Indeed, if we consider all consistent runs of u_j on words of the form $\$.2^{[n]}$ and all consistent runs of d_j on words of the form $\$.2^{[n]}$, we get that there are two runs, of u_j and of d_j , corresponding to distinct subsets of $[n]$, that reach the same state p_j . So far, we have defined two infinite runs of \mathcal{A}^{qk-1} that are consistent with the strategy f :

$$\begin{aligned} r_u &= u_0 \xrightarrow{\$.U_1} p_0 \xrightarrow{\# \cdot i_1} u_1 \xrightarrow{\$.U_2} p_1 \xrightarrow{\# \cdot i_2} u_2 \xrightarrow{\$.U_3} \dots \\ r_d &= d_0 \xrightarrow{\$.D_1} p_0 \xrightarrow{\# \cdot i_1} d_1 \xrightarrow{\$.D_2} p_1 \xrightarrow{\# \cdot i_2} d_2 \xrightarrow{\$.D_3} \dots \end{aligned}$$

where for all $j \geq 1$, $i_j \in (U_j \setminus D_j) \cup (D_j \setminus U_j)$. In particular, exactly one of $\$.U_j \cdot \# \cdot i_j$ and $\$.D_j \cdot \# \cdot i_j$ is good. Therefore, at least one of the runs r_u and r_d runs on a word in ∞R_n . Since both runs are consistent with the strategy f and $\infty R_n \subseteq L(\mathcal{A}^{qk-1})$, we get that one of the runs traverses α infinitely often. W.l.o.g consider some $j > 0$ such that the sub-run $p_j \xrightarrow{\# \cdot i_{j+1}} u_{j+1} \xrightarrow{\$.U_{j+2}} p_{j+1}$ traverses α . We define

$$x_k = (\$.A_1 \cdot \# \cdot i_1) \cdot (\$.A_2 \cdot \# \cdot i_2) \dots (\$.A_j \cdot \# \cdot i_j) \cdot (\$.A_{j+1} \cdot \# \cdot i_{j+1}) \cdot \$.U_{j+2}$$

where for all $1 \leq l \leq j+1$, it holds that $A_l \in \{U_l, D_l\}$ and $i_l \notin A_l$. In particular, x_k starts with $\$$ and has no good infixes. To conclude the proof, we show that there is a run r_k of \mathcal{A}^{qk-1} on x_k that traverses an α -transition (in particular, we can choose q_k to be the last state that r_k visits). The run r_k can be obtained from the runs r_u and r_d as follows. First, r_k starts by following one of the sub-runs $u_0 \xrightarrow{\$.U_1} p_0$ or $d_0 \xrightarrow{\$.D_1} p_0$, depending on whether $A_1 = U_1$ or $A_1 = D_1$, respectively. Then, if the run r_k is in a state of the form p_m , for $m < j$, the run r_k proceeds, upon reading $\# \cdot i_{m+1}$, to one of the states u_{m+1} or d_{m+1} , depending on whether the next two letters to be read are $\$.U_{m+2}$ or $\$.D_{m+2}$, respectively. Finally, once r_k reaches p_j , it ends by traversing α while following the sub-run $p_j \xrightarrow{\# \cdot i_{j+1}} u_{j+1} \xrightarrow{\$.U_{j+2}} p_{j+1}$. \square

Remark 3.1. In order to get a slightly tighter bound, one can show that a minimal tDBW for L_n needs at least 2^{n+1} states and that the language L_n is not *HD-helpful*. That is, a minimal HD-tNBW for L_n is not smaller than a minimal tDBW for L_n , and so the 2^{n+1} bound holds also for a minimal

HD-tNBW. The proof starts with a tDBW for L_n that has 2^{n+1} states, considers its complement tDCW, and shows that the application of the polynomial minimization algorithm of [1] on it, namely the algorithm that returns an equivalent minimal HD-tNCW, does result in a smaller automaton. The result then follows from the fact that a minimal HD-tNCW for a language is smaller than a minimal HD-tNBW for its complement [19]. The proof is easy for readers familiar with [1], yet involves many notions and observations from there. We thus leave it for the appendix. Specifically, in Appendix A.1, we describe a tDBW with 2^{n+1} states for L_n , and readers familiar with [1] can observe that the application of the HD-tNCW minimization algorithm on its dual tDCW does not make it smaller. \square

Theorem 3.4. *There is a family L_1, L_2, L_3, \dots of languages such that for every $n \geq 1$, there is an SD-tNBW with $O(n)$ states that recognizes L_n , yet every SD-tNCW that recognizes $\overline{L_n}$ needs at least $2^{O(n)}$ states.*

Proof. Let $\Sigma = \{0, 1\}$. For $n \geq 1$, let $R_n = \{w : w \in 0 \cdot (0+1)^{n-1} \cdot 1 + 1 \cdot (0+1)^{n-1} \cdot 0\}$. Thus, R_n contains all words of length $n+1$ whose first and last letters are different. It is easy to see that R_n can be recognized by an NFW (in fact, a DFW) \mathcal{N}_n with $O(n)$ states. We define $L_n = \infty R_n$. First, by Theorem 3.1, there is an SD-tNBW with $O(n)$ states for L_n . In order to prove that an SD-tNCW for $\overline{L_n}$ needs at least $2^{O(n)}$ states, we prove that in fact every tNCW for $\overline{L_n}$ needs that many states. For this, note that $\overline{L_n}$ consists of all words w for which there is $u \in (0+1)^n$ such that $w \in (0+1)^* \cdot u^\omega$. Indeed, for such words w , the suffix u^ω contains no infix in R_n . Also, if a word contains only finitely many infixes in R_n , then it must have a suffix with no infixes in R_n , namely a suffix of the form u^ω for some $u \in (0+1)^n$. Then, the proof that a tNCW for $\overline{L_n}$ needs exponentially many states is similar to the proof that an NFW for $\{u \cdot u : u \in (0+1)^n\}$ needs exponentially many states. Indeed, it has to remember the last n letters read. \square

3.2 Decision Problems

We continue to decision problems about SD-tNBWs and SD-NBWs, and show that the exponential succinctness comes with a price: the complexity of all the problems we study coincides with the one known for tNBWs and NBWs. Accordingly, we only prove lower bounds for SD-tNBWs. Matching upper bounds follow from the known complexity for tNBWs, and same bounds for SD-NBWs follow from linear translations between SD-tNBWs and SD-NBWs. The problems we study are *language containment*: given two SD-tNBWs \mathcal{A}_1 and \mathcal{A}_2 , decide whether $L(\mathcal{A}_1) \subseteq L(\mathcal{A}_2)$, *universality*: given an SD-tNBW \mathcal{A}_1 , decide whether $L(\mathcal{A}_1) = \Sigma^\omega$, and *minimization*: given an SD-tNBW \mathcal{A}_1 and an integer $k \geq 1$, decide whether there is an SD-tNBW \mathcal{A}_2 such that $L(\mathcal{A}_1) = L(\mathcal{A}_2)$ and $|\mathcal{A}_2| \leq k$.

The exponential succinctness of SD automata motivates also the study of the *D-to-SD minimization problem*. Here, we are given a tDBW \mathcal{A}_1 and an integer $k \geq 1$, and we need to decide whether there is an SD-tNBW \mathcal{A}_2 such that $L(\mathcal{A}_1) = L(\mathcal{A}_2)$ and $|\mathcal{A}_2| \leq k$. For automata on finite words, the D-to-N minimization problem is known to be PSPACE-complete [18].

Note that a lower bound for universality implies a lower bound also for language containment. We still start with language containment, as it is much simpler.

Theorem 3.5. *The language-containment problem for SD-tNBWs is PSPACE-hard.*

Proof. We describe a reduction from the universality problem for NFWs. Given an NFW \mathcal{N} over Σ , let \mathcal{N}' be an NFW over $\Sigma \cup \{\$\}$ such that $L(\mathcal{N}') = \$ \cdot L(\mathcal{N}) \cdot \$$. Now, let \mathcal{A}_1 be a 1-state tDBW over $\Sigma \cup \{\$\}$ such that $L(\mathcal{A}_1) = \infty \$$, and let \mathcal{A}_2 be the SD-tNBW obtained by applying the operation from Theorem 3.1 on \mathcal{N}' . Note that $L(\mathcal{A}_2) = \infty (\$ \cdot L(\mathcal{N}) \cdot \$)$ and $|\mathcal{A}_2| = |\mathcal{N}'| + 3$.

We claim that \mathcal{N} is universal iff $L(\mathcal{A}_1) \subseteq L(\mathcal{A}_2)$. First, if $L(\mathcal{N}) = \Sigma^*$, then $L(\mathcal{A}_2) = \infty (\$ \cdot \Sigma^* \cdot \$) = \infty \$$ and so $L(\mathcal{A}_1) \subseteq L(\mathcal{A}_2)$. Conversely, if there is a word $x \in \Sigma^* \setminus L(\mathcal{N})$, then the word $w = (\$ \cdot x)^\omega$ is in $L(\mathcal{A}_1) \setminus L(\mathcal{A}_2)$. Indeed, w has infinitely many $\$$'s, yet for every infix $\$ \cdot y \cdot \$$ of w , we have that $y \notin L(\mathcal{N})$, and so $w \notin L(\mathcal{A}_2)$. \square

The proof in Theorem 3.5 uses $\infty\$$ as the “contained language”. For the universality problem, we cannot rely on hints from words in the contained language. In particular, taking the union of \mathcal{A}_2 there with an automaton for “only finitely many $\$$ ’s” results in an automaton that is not SD, and it is not clear how to make it SD. Consequently, we have to work much harder. Specifically, we prove PSPACE hardness by a generic reduction from polynomial space Turing machines. Such reductions associate with a Turing machine T an automaton \mathcal{A} that recognizes the language R of words that do not encode legal rejecting computations of T , and so $R = \Sigma^*$ iff the machine has no rejecting computations. The automaton \mathcal{A} is nondeterministic, as it has to guess violations of attempts to encode legal accepting computations. In order to replace \mathcal{A} by an SD automaton, we manipulate the Turing machine so that the language of the generated automaton is of the form ∞R , for which we can construct an SD-tNBW.

Theorem 3.6. *The universality problem for SD-tNBWs is PSPACE-hard.*

Proof. We do a reduction from polynomial-space Turing machines. Given a Turing machine T with space complexity $s : \mathbb{N} \rightarrow \mathbb{N}$, we construct in time polynomial in $|T|$ and $s(0)$, an SD-tNBW \mathcal{A} of size polynomial in T and $s(0)$, such that \mathcal{A} is universal iff T accepts the empty tape³. Let $n_0 = s(0)$. Thus, each configuration in the computation of T on the empty tape uses at most n_0 cells. We assume that T halts from all configurations (that is, not just from these reachable from an initial configuration of T); Indeed, by adding a polynomial-space counter to T , one can transform a polynomial-space Turing machine that need not halt from all configurations to one that does halt. We also assume, without loss of generality, that once T reaches a final (accepting or rejecting) state, it erases the tape, moves with its reading head to the leftmost cell, and moves to the initial state. Thus, all computations of T are infinite and after visiting a final configuration for the first time, they eventually consists of repeating the same finite computation on the empty tape that uses at most n_0 tape cells.

We define \mathcal{A} so that it accepts a word w iff (C1) no suffix of w is an encoding of a legal computation of T that uses at most n_0 tape cells, or (C2) w has infinitely many infixes that encode the accepting state of T .

It is not hard to see that T accepts the empty tape iff \mathcal{A} is universal. Indeed, if T accepts the empty tape, and there is a word w that does not satisfy (C1), thus w has a suffix that is an encoding of a legal computation of T that uses at most n_0 cells, then the encoded computation eventually reaches a final configuration, from which it eventually repeats the accepting computation of T on the empty tape infinitely many times, and so w satisfies (C2). Conversely, If T rejects the empty tape, then the word w that encodes the computation of T on the empty tape does not satisfy (C1) nor (C2), and so \mathcal{A} does not accept w .

Finally, the fact that T is a polynomial-space Turing machine enables us to define \mathcal{A} with polynomially many states, as we detail next. Let $T = \langle \Gamma, Q, \rightarrow, q_0, q_{acc}, q_{rej} \rangle$ be a Turing machine with polynomial space complexity $s : \mathbb{N} \rightarrow \mathbb{N}$, where Γ is the tape-alphabet, Q is the set of states, $\rightarrow \subseteq Q \times \Gamma \times Q \times \Gamma \times \{L, R\}$ is the transition relation (we use $(q, a) \rightarrow (q', b, \Delta)$ to indicate that when T is in state q and it reads the input a in the current tape cell, it moves to state q' , writes b in the current tape cell, and its reading head moves one cell to the left/right, according to Δ), q_0 is the initial state, q_{acc} is the accepting state, and q_{rej} is the rejecting one.

We encode a configuration of T that uses at most $n_0 = s(0)$ tape cells by a word of the form $\#\gamma_1\gamma_2 \dots (q, \gamma_i) \dots \gamma_{n_0}$. That is, the encoding of a configuration starts with a special letter $\#$, and all its other letters are in Γ , except for one letter in $Q \times \Gamma$. The meaning of such a configuration is that the j 'th cell in T , for $1 \leq j \leq n_0$, is γ_j -labeled, the head of T points at i 'th cell, and T is in the state

³This is sufficient, as one can define a generic reduction from every language L in PSPACE as follows. Let T_L be a Turing machine that decides L in polynomial space $f(n)$. On input w for the reduction, the reduction considers the machine T_w that on every input, first erases the tape, writes w on its tape, and then runs as T_L on w . Then, the reduction outputs an automaton \mathcal{A} , such that T_w accepts the empty tape iff \mathcal{A} is SD. Note that the space complexity of T_w is $s(n) = \max(n, f(|w|))$, and that w is in L iff T_w accepts the empty tape. Since \mathcal{A} is constructed in time polynomial in $s(0) = f(|w|)$ and $|T_w| = \text{poly}(|w|)$, it follows that the reduction is polynomial in $|w|$.

q . For example, the initial configuration of T on the empty tape is encoded as $\#(q_0, b)b \dots b$ (with $n_0 - 1$ occurrences of b 's) where b stands for an empty cell. We can now encode a computation of T as a sequence of configurations.

Let $\Sigma = \{\#\} \cup \Gamma \cup (Q \times \Gamma)$, and let $\#\sigma_1 \dots \sigma_{n_0} \#\sigma'_1 \dots \sigma'_{n_0}$ be two successive configurations of T . We also set σ_0, σ'_0 , and σ_{n_0+1} to $\#$. For each triple $\langle \sigma_{i-1}, \sigma_i, \sigma_{i+1} \rangle$ with $1 \leq i \leq n_0$, we know, by the transition relation of T , what σ'_i should be. In addition, the letter $\#$ should repeat exactly every $n_0 + 1$ letters. Let $next(\sigma_{i-1}, \sigma_i, \sigma_{i+1})$ denote our expectation for σ'_i . That is,

- $next(\sigma_{i-1}, \sigma_i, \sigma_{i+1}) = \sigma_i$ if $\sigma_i = \#$, or if non of σ_{i-1}, σ_i and σ_{i+1} are in $Q \times \Gamma$.
- $next((q, \gamma_{i-1}), \gamma_i, \gamma_{i+1}) = next((q, \gamma_{i-1}), \gamma_i, \#) =$

$$\begin{cases} \gamma_i & \text{If } (q, \gamma_{i-1}) \rightarrow (q', \gamma'_{i-1}, L) \\ (q', \gamma_i) & \text{If } (q, \gamma_{i-1}) \rightarrow (q', \gamma'_{i-1}, R) \end{cases}$$

- $next(\gamma_{i-1}, \gamma_i, (q, \gamma_{i+1})) = next(\#, \gamma_i, (q, \gamma_{i+1})) =$

$$\begin{cases} \gamma_i & \text{If } (q, \gamma_{i+1}) \rightarrow (q', \gamma'_{i+1}, R) \\ (q', \gamma_i) & \text{If } (q, \gamma_{i+1}) \rightarrow (q', \gamma'_{i+1}, L) \end{cases}$$

- $next(\gamma_{i-1}, (q, \gamma_i), \gamma_{i+1}) = next(\gamma_{i-1}, (q, \gamma_i), \#) = next(\#, (q, \gamma_i), \gamma_{i+1}) = \gamma'_i$ where $(q, \gamma_i) \rightarrow (q', \gamma'_i, \Delta)$.⁴

Note that every word that encodes a legal computation of T that uses at most n_0 tape cells must be consistent with $next$.

Consider now an infinite word w . Note that if w satisfies (C1), then w has infinitely many infixes that witness a violation of the encoding of a legal computation of T that uses at most n_0 tape cells. Accordingly, we can write the language of \mathcal{A} as ∞R , where R is the language of one-letter words of the form (q_{acc}, γ) , or words that describe a violation of a legal computation that uses at most n_0 cells. Words that describe a violation, either are words of length $n_0 + 1$ that describe a violation of the encoding of a single configuration, or are words of length $n_0 + 3$ that describe a violation $next$.

We can now define an NFW \mathcal{N} for R of size polynomial in T and n_0 , and thus the required SD-tNBW \mathcal{A} can be obtained from \mathcal{N} by applying the construction described in Theorem 3.1. The NFW \mathcal{N} is defined as follows. From its initial states, \mathcal{N} guesses to check whether the input word is a single-letter word of the form (q_{acc}, γ) , or guesses to check whether the input is a violation of the encoding of a legal computation of T that uses at most n_0 tape cells: this amounts to guessing that the input is a word of length $n_0 + 1$ that has no $\#$'s, or starts with $\#$, yet does not encode a legal configuration, or amounts to guessing that the input is a word of length $n_0 + 3$ that describes a violation of $next$, where \mathcal{N} checks whether the input start with three letters $\sigma_1, \sigma_2, \sigma_3$ describing three successive letters in the encoding and ends with a letter σ'_2 , which comes $n_0 + 1$ letters after σ_2 in the encoding, yet is different than $next(\sigma_1, \sigma_2, \sigma_3)$. \square

We continue to the minimization problem. Note that here, a PSPACE upper bound does not follow immediately from the known PSPACE upper bound for tNBWs, as the candidate automata need to be SD. Still, as SDness can be checked in PSPACE [2], a PSPACE upper bound follows. Also note that here, the case of SD-NBWs is easy, as a non-empty SD-NBW is universal iff it has an equivalent SD-NBW with one state. For transition-based acceptance, the language of a single-state SD-tNBW need not be trivial, and so we have to examine the specific language used for the universality PSPACE-hardness proof:

⁴We note that in the case where the head of T is at left most cell, and T moves its head to the left should be handled slightly differently, as we assume that the tape of T is left bounded.

Theorem 3.7. *The minimization problem for SD-tNBWs is PSPACE-hard.*

Proof. We argue that the SD-tNBW \mathcal{A} in the proof of Theorem 3.6 has an equivalent one-state SD-tNBW iff \mathcal{A} is universal. To see why, assume towards contradiction that T rejects the empty tape and there is a one-state SD-tNBW \mathcal{D} equivalent to \mathcal{A} . Recall that we encode computations of T as a sequence of successive configurations of length n_0 each, separated by the letter $\#$. On the one hand, as the word $\#^\omega$ has infinitely many violations, then it is in $L(\mathcal{A})$, and so the $\#$ -labeled self-loop at the state of \mathcal{D} must be an α -transition. On the other hand, the word w that encodes the computation of T on the empty tape is not in $L(\mathcal{A})$, yet it has infinitely many $\#$'s. \square

As we show below, the minimization problem stays hard even when we start from a deterministic automaton:

Theorem 3.8. *The D-to-SD minimization problem for Büchi automata is PSPACE-hard.*

Proof. We start with Büchi automata with transition-based acceptance, and describe a reduction from the D-to-N minimization problem for automata on finite words: given a DFW \mathcal{A}_1 , and an integer $k \geq 1$, decide whether there is an NFW \mathcal{A}_2 such that $L(\mathcal{A}_1) = L(\mathcal{A}_2)$ and $|\mathcal{A}_2| \leq k$. In [18], the authors prove that the problem is PSPACE-hard, in fact PSPACE-hard already for DFWs that recognize a language that has no good prefixes.

Consider a language $R \subseteq \Sigma^*$. The reduction is based on a construction that turns an NFW for R into an SD-tNBW for $\infty(\$ \cdot R \cdot \$)$, for a letter $\$ \notin \Sigma$. Note that by applying the construction from Theorem 3.1 on the language $\$ \cdot R \cdot \$$, we can get an SD-tNBW for $\infty(\$ \cdot R \cdot \$)$. The construction there, however, does not preserve determinism. Therefore, we need a modified polynomial construction, which takes advantage of the $\$$'s. We describe the modified construction below.

Given an NFW $\mathcal{A} = \langle \Sigma, Q, Q_0, \delta, F \rangle$, and a letter $\$ \notin \Sigma$, we construct the tNBW $\mathcal{A}' = \langle \Sigma \cup \{\$\}, Q, Q_0, \delta', \alpha \rangle$, where for all states $q \in Q$ and letters $\sigma \in \Sigma$, we have that $\delta'(q, \sigma) = \delta(q, \sigma)$. Also, $\delta'(q, \$) = Q_0$, and $\alpha = \{ \langle s, \$, q \rangle : q \in Q_0 \text{ and } s \in F \}$. Thus, \mathcal{A}' is obtained from \mathcal{A} by adding $\$$ -transitions from all states to Q_0 . A new transition is in α iff its source is an accepting state of \mathcal{A} . It is not hard to see that the construction of \mathcal{A}' preserves determinism. Indeed, if \mathcal{A} is deterministic, then all the $\$$ -transitions in \mathcal{A}' are deterministic since they are directed to the single initial state of \mathcal{A} . To show that \mathcal{A}' is SD and $L(\mathcal{A}') = \infty(\$ \cdot L(\mathcal{A}) \cdot \$)$, we prove below, similarly to Theorem 3.1, that for all states $q \in Q$, it holds that $L(\mathcal{A}'^q) = \infty(\$ \cdot L(\mathcal{A}) \cdot \$)$.

Consider a state $q \in Q$, we first show that $L(\mathcal{A}'^q) \subseteq (\$ \cdot L(\mathcal{A}) \cdot \$)$. Consider a word $w \in L(\mathcal{A}'^q)$ and let $r = r_0, r_1, r_2, \dots$ be an accepting run of \mathcal{A}'^q on w . We show that for every $i \geq 1$, there are k_1, k_2 such that $i \leq k_1 \leq k_2$ and $w[k_1, k_2 + 1]$ is an infix of w in $\$ \cdot L(\mathcal{A}) \cdot \$$. Given $i \geq 1$, let k_2 be such that $t = \langle r_{k_2}, w[k_2 + 1], r_{k_2+1} \rangle$ is an α -transition, but not the first $\$$ -transition that r reads when it runs on the suffix $w[i, \infty]$, and let k_1 be the maximal index such that $i \leq k_1 \leq k_2$ and $\langle r_{k_1-1}, w[k_1], r_{k_1} \rangle = \langle r_{k_1-1}, \$, r_{k_1} \rangle$. Since t is not the first $\$$ -transition that r traverses when it reads the suffix $w[i, \infty]$, then k_1 exists. Note that the fact that new transitions lead to Q_0 in \mathcal{A}' , and the maximality of k_1 , imply that the run $r_{k_1}, r_{k_1+1}, \dots, r_{k_2}$ is a run of \mathcal{A} on the infix $w[k_1 + 1, k_2]$. Also, as t is an α transition, the latter run is an accepting run of \mathcal{A} . Now, as the transitions $\langle r_{k_1-1}, w[k_1], r_{k_1} \rangle$ and t are $\$$ -transitions, we get that $r_{k_1-1}, r_{k_1}, \dots, r_{k_2}, r_{k_2+1}$ is a run on a word in $\$ \cdot L(\mathcal{A}) \cdot \$$. Thus, $w[k_1, k_2 + 1]$ is an infix of w in $\$ \cdot L(\mathcal{A}) \cdot \$$.

Next, we show that $\infty(\$ \cdot L(\mathcal{A}) \cdot \$) \subseteq L(\mathcal{A}'^q)$. Consider a word $w \in \infty(\$ \cdot L(\mathcal{A}) \cdot \$)$, and consider a run of \mathcal{A}'^q on w that upon reading an infix $\$ \cdot x \cdot \$$ in $\$ \cdot L(\mathcal{A}^{q_0}) \cdot \$$, for some $q_0 \in Q_0$, moves to q_0 while reading $\$$, and then follows an accepting run of \mathcal{A}^{q_0} on x , and finally moves to a state in Q_0 , traversing α while reading $\$$. Note that the transitions in \mathcal{A}' enable such a behavior. Thus, such a run traverses α infinitely often and is thus accepting.

We now describe the reduction. Given a DFW \mathcal{A} over Σ such that $L(\mathcal{A})$ has no good prefixes, and given an integer k , the reduction returns the polynomial tDBW \mathcal{A}' and the integer k . We prove next that the reduction is correct. Thus, the DFW \mathcal{A} has an equivalent NFW with at most k states iff

the tDBW \mathcal{A}' has an equivalent SD-tNBW with at most k states. For the first direction, if \mathcal{B} is an NFW equivalent to \mathcal{A} whose size is at most k , then by the above construction, it holds that \mathcal{B}' is an SD-tNBW whose size is at most k , and $L(\mathcal{B}') = \infty(\$ \cdot L(\mathcal{B}) \cdot \$) = \infty(\$ \cdot L(\mathcal{A}) \cdot \$) = L(\mathcal{A}')$.

Conversely, if \mathcal{B}' is an SD-tNBW for $\infty(\$ \cdot L(\mathcal{A}) \cdot \$)$ whose size is at most k , then as $L(\mathcal{A})$ has no good prefixes, we get by Theorem 3.2 that there is an NFW \mathcal{B} for $L(\mathcal{A})$ whose size is at most k , and we are done.

Finally, we show next that despite the fact that transitions between automata with transition-based and state-based acceptance may involve a linear blow-up, it is possible to adapt the above arguments to automata with state-based acceptance, thus PSPACE-completeness holds also for Büchi automata with state-based acceptance. Essentially, given a DFW \mathcal{A} and an integer k , the reduction returns a DBW for $\infty(\$ \cdot L(\mathcal{A}) \cdot \$)$ and the integer $k + 1$.

To see why such a reduction exists, note that the construction of the SD-tNBW \mathcal{A}' from the NFW \mathcal{A} can be adapted to the state-based setting by adding a single new α state to \mathcal{A}' that behaves as the initial states of \mathcal{A} , and captures infixes in $\$ \cdot L(\mathcal{A}) \cdot \$$. Formally, given an NFW $\mathcal{A} = \langle \Sigma, Q, Q_0, \delta, F \rangle$, and a letter $\$ \notin \Sigma$, we construct the NBW $\mathcal{A}' = \langle \Sigma \cup \{\$, \}, Q \cup \{q_{acc}\}, Q_0, \delta', \{q_{acc}\} \rangle$, where for all states $q \in Q$ and letters $\sigma \in \Sigma$, we have that $\delta'(q, \sigma) = \delta(q, \sigma)$. Also, $\delta'(Q \setminus F, \$) = Q_0$, $\delta'(F, \$) = \{q_{acc}\}$, $\delta(q_{acc}, \$) = Q_0$, and $\delta'(q_{acc}, \sigma) = \delta(Q_0, \sigma)$. Intuitively, the state q_{acc} mimics traversal of α transitions.

Then, Theorem 3.2 can also be adapted to SD-NBWs as follows. If \mathcal{A} is an SD-NBW for $\infty(\$ \cdot R \cdot \$)$ for some nontrivial language $R \subseteq \Sigma^*$ that has no good prefixes, then one can construct an NFW \mathcal{N} for R with at most $|\mathcal{A}| - 1$ states. Essentially, we view \mathcal{A} as a transition-based automaton \mathcal{A}_t by considering transitions that leave α states as α transitions. Then, we consider the proof of Theorem 3.2 when applied on \mathcal{A}_t . As R has no good prefixes, then the state q_{acc} of \mathcal{N} is unreachable. Hence, all runs from Q_0^S cannot reach α transitions in \mathcal{A}_t . In particular, α states of \mathcal{A} are unreachable in \mathcal{N} and they can be removed with q_{acc} as well. \square

4 Semantically Deterministic co-Büchi Automata

In this section we study SD co-Büchi automata. Here too, our results are based on constructions that involve encodings of NFWs by SD-tNCWs. Here, however, the constructions are more complicated, and we first need some definitions and notations.

Consider a tNCW $\mathcal{A} = \langle \Sigma, Q, Q_0, \delta, \alpha \rangle$. We refer to the SCCs we get by removing \mathcal{A} 's α -transitions as the $\bar{\alpha}$ -components of \mathcal{A} ; that is, the $\bar{\alpha}$ -components of \mathcal{A} are the SCCs of the graph $G_{\mathcal{A}^{\bar{\alpha}}} = \langle Q, E^{\bar{\alpha}} \rangle$, where $\langle q, q' \rangle \in E^{\bar{\alpha}}$ iff there is a letter $\sigma \in \Sigma$ such that $\langle q, \sigma, q' \rangle \in \bar{\alpha}$. We say that \mathcal{A} is *normal* if there are no $\bar{\alpha}$ -transitions connecting different $\bar{\alpha}$ -components. That is, for all states q and s of \mathcal{A} , if there is a path of $\bar{\alpha}$ -transitions from q to s , then there is also a path of $\bar{\alpha}$ -transitions from s to q . Note that an accepting run of \mathcal{A} eventually gets trapped in one of \mathcal{A} 's $\bar{\alpha}$ -components. In particular, accepting runs in \mathcal{A} traverse transitions that leave $\bar{\alpha}$ -components only finitely often. Hence, we can add transitions among $\bar{\alpha}$ -components to α without changing the language of the automaton. Accordingly, in the sequel we assume that given tNCWs are normal.

We proceed to encoding NFWs by SD-tNCWs. For a language $R \subseteq \Sigma^*$, we define the language $\bowtie_{\$}(R) \subseteq (\Sigma \cup \{\$, \})^\omega$, by

$$\bowtie_{\$}(R) = \{w : \text{if } w \text{ has infinitely many } \$, \text{ then it has a suffix in } (\$R)^\omega\}.$$

Also, we say that a finite word $x \in \Sigma^*$ is a *bad infix* for R if for all words $w \in \Sigma^*$ that have x as an infix, it holds that $w \in \bar{R}$.

Note that for every language $R \subseteq \Sigma^*$, we have that $\bowtie_{\$}(R) = \overline{\infty(\$ \cdot \bar{R} \cdot \$)}$. Thus, $\bowtie_{\$}(R)$ complements $\infty(\$ \cdot \bar{R} \cdot \$)$. Yet, unlike tDCWs and tDBWs, which dualize each other, SD-tNBWs and SD-tNCWs are not dual. Hence, adjusting Theorems 3.1 and 3.2 to the co-Büchi setting, requires different, in fact more complicated, constructions.

Theorem 4.1. *Given an NFW \mathcal{N} , one can obtain, in linear time, an SD-tNCW \mathcal{A} such that $L(\mathcal{A}) = \bowtie_{\$}(L(\mathcal{N}))$ and $|\mathcal{A}| = |\mathcal{N}|$.*

Proof. Given $\mathcal{N} = \langle \Sigma, Q, Q_0, \delta, F \rangle$, we obtain \mathcal{A} by adding $\$$ -transitions from all states to Q_0 . The new transitions are in α iff they leave a state in $Q \setminus F$. Formally, $\mathcal{A} = \langle \Sigma \cup \{\$\}, Q, Q_0, \delta', \alpha \rangle$, where for all $s \in Q$ and $\sigma \in \Sigma$, we have that $\delta'(s, \sigma) = \delta(s, \sigma)$, and $\delta'(s, \$) = Q_0$. Then, $\alpha = \{ \langle s, \$, q \rangle : q \in Q_0 \text{ and } s \in Q \setminus F \}$. It is easy to see that $|\mathcal{A}| = |\mathcal{N}|$. In order to prove that \mathcal{A} is SD and $L(\mathcal{A}) = \bowtie_{\$}(L(\mathcal{N}))$, we prove next that for every state $q \in Q$, it holds that $L(\mathcal{A}^q) = \bowtie_{\$}(L(\mathcal{N}))$.

First, we show that $\bowtie_{\$}(L(\mathcal{N})) \subseteq L(\mathcal{A}^q)$. Consider a word $w \in \bowtie_{\$}(L(\mathcal{N}))$. If w has finitely many $\$$'s, then as α transitions in \mathcal{A} are $\$$ -labeled, we have that all runs of \mathcal{A}^q on w traverse α finitely often and thus are accepting. Assume that $w = z \cdot y$, where y is of the form $\$x_1\$x_2\$x_3 \dots$, where for all i , $x_i \in L(\mathcal{N})$. An accepting run r of \mathcal{A}^q on w can be defined arbitrarily upon reading z . Then, upon reading an infix of the form $\$x_i$, r moves to q_0 upon reading $\$$, where $q_0 \in Q_0$ and $x_i \in L(\mathcal{N}^{q_0})$. Then, r proceeds by following an accepting run of \mathcal{N}^{q_0} on x_i . As accepting runs of \mathcal{N} end in states in F , and $\$$ -transitions from F are in $\bar{\alpha}$, we get that r does not traverse α upon reading the $\$$'s after x_1 . Also, as transitions of \mathcal{N} are $\bar{\alpha}$ transitions in \mathcal{A} , we get that r follows a run that does not traverse α when it reads x_i . Hence, r eventually visits only $\bar{\alpha}$ transitions and thus is accepting.

Next, we show that $L(\mathcal{A}^q) \subseteq \bowtie_{\$}(L(\mathcal{N}))$. Consider a word $w = \sigma_1 \cdot \sigma_2 \dots \in L(\mathcal{A}^q)$ and let $r = r_0, r_1, \dots$ be an accepting run of \mathcal{A}^q on w . If w has finitely many $\$$'s, then $w \in \bowtie_{\$}(L(\mathcal{N}))$ and we are done. Assume that w has infinitely many $\$$'s and let t be such that $r[t-1, \infty] = r_{t-1}, r_t, r_{t+1}, \dots$ is a run that does not traverse α on $w[t, \infty]$. To conclude the proof, we show that between every two consecutive $\$$'s in $w[t, \infty]$ there is a word in $L(\mathcal{N})$. That is, we show that for all $t \leq t_1 < t_2$ such that $\sigma_{t_1} = \sigma_{t_2} = \$$ and $w[t_1+1, t_2-1]$ has no $\$$'s, it holds that $w[t_1+1, t_2-1] \in L(\mathcal{N})$. As $\$$ -transitions lead to Q_0 , we have that $r_{t_1}, r_{t_2} \in Q_0$. Recall that $r[t-1, \infty]$ does not traverse α , in particular, we have that $r[t_1, t_2] = r_{t_1}, r_{t_1+1}, \dots, r_{t_2}$ is a run from Q_0 to Q_0 that traverses only $\bar{\alpha}$ transitions, on the infix $w[t_1+1, t_2]$. Also, $r[t_1, t_2]$ traverses $\$$ only in its last transition. As the transition $\langle r_{t_2-1}, \sigma_{t_2}, r_{t_2} \rangle$ is an $\bar{\alpha}$ $\$$ -transition that leads to Q_0 , then $r_{t_2-1} \in F$. Also, since all the transitions that $r[t_1, t_2-1]$ traverses are Σ -labeled, then they are transitions in \mathcal{N} , and so $r[t_1, t_2-1]$ is an accepting run of \mathcal{N} on $w[t_1+1, t_2-1]$. Hence $w[t_1+1, t_2-1] \in L(\mathcal{N})$ and we are done. \square

Theorem 4.2. *Consider a language $R \subseteq \Sigma^*$ and a letter $\$ \notin \Sigma$. For every tNCW \mathcal{A} such that $L(\mathcal{A}) = \bowtie_{\$}(R)$, there exists an NFW \mathcal{N} such that $L(\mathcal{N}) = R$ and $|\mathcal{N}| \leq |\mathcal{A}| + 1$. In addition, if R has bad infixes, then $|\mathcal{N}| \leq |\mathcal{A}|$.*

Proof. Let $\mathcal{A} = \langle \Sigma \cup \{\$\}, Q, Q_0, \delta, \alpha \rangle$ be a tNCW for $\bowtie_{\$}(R)$. We assume that \mathcal{A} is normal and all of its states are reachable. We define the NFW $\mathcal{N} = \langle \Sigma, Q \cup \{q_{rej}\}, Q_0^{\mathcal{N}}, \delta_{\mathcal{N}}, F_{\mathcal{N}} \rangle$, where $Q_0^{\mathcal{N}} = \{q \in Q : \text{there is a state } q' \text{ such that } \langle q', \$, q \rangle \in \bar{\alpha}\}$, $F_{\mathcal{N}} = \{q \in Q : \text{there is a state } q' \text{ such that } \langle q, \$, q' \rangle \in \bar{\alpha}\}$, and for every two states $q, s \in Q$ and letter $\sigma \in \Sigma$, it holds that $s \in \delta_{\mathcal{N}}(q, \sigma)$ iff $\langle q, \sigma, s \rangle \in \bar{\alpha}$. Also, if $\{q\} \times \{\sigma\} \times \delta(q, \sigma) \subseteq \alpha$, then $\delta_{\mathcal{N}}(q, \sigma) = q_{rej}$. Also, for all letters $\sigma \in \Sigma$, it holds that $\delta_{\mathcal{N}}(q_{rej}, \sigma) = q_{rej}$; that is, q_{rej} is a rejecting sink. Thus, \mathcal{N} tries to accept words $x \in \Sigma^*$ for which there is a run in \mathcal{A} that does not traverse α on the word $\$ \cdot x \cdot \$$.

We prove that $L(\mathcal{N}) = R$. We first prove that $R \subseteq L(\mathcal{N})$. Consider a word $x \in R$, and let $r = r_0, r_1, r_2, \dots$ be an accepting run of \mathcal{A} on $(\$ \cdot x)^{\omega}$. As r is accepting, there are $i < j$ such that r_i, r_{i+1}, \dots, r_j is a run that does not traverse α on $\$ \cdot x \cdot \$$. By the definition of $\delta_{\mathcal{N}}$, $r_{i+1} \in Q_0^{\mathcal{N}}$, $r_{j-1} \in F_{\mathcal{N}}$, and so $r_{i+1}, r_{i+2}, \dots, r_{j-1}$ is an accepting run of \mathcal{N} on x .

We prove next that $L(\mathcal{N}) \subseteq R$. Consider a word $x = \sigma_1 \cdot \sigma_2 \dots \sigma_n \in L(\mathcal{N})$ and let $r = r_0, r_1, \dots, r_n$ be an accepting run of \mathcal{N} on x . As r ends an accepting state of \mathcal{N} , then it does not visit q_{rej} . Hence, the definition of $\delta_{\mathcal{N}}$ implies that r is a run that does not traverse α in \mathcal{A} . Also, as $r_0 \in Q_0^{\mathcal{N}}$ and $r_n \in F_{\mathcal{N}}$, there are states $q_1, q_2 \in Q$ such that $\langle q_1, \$, r_0 \rangle, \langle r_n, \$, q_2 \rangle \in \bar{\alpha}$. Hence, $q_1, r_0, r_1, \dots, r_n, q_2$ is a run that does not traverse α in \mathcal{A} on the word $\$ \cdot x \cdot \$$. As \mathcal{A} is normal, there is a word $y \in (\Sigma \cup \{\$\})^*$ such that there is a run that does not traverse α of \mathcal{A}^{q_2} on y that reaches q_1 . Therefore, $(\$ \cdot x \cdot \$ \cdot y)^{\omega} \in L(\mathcal{A}^{q_1})$. As

all the states of \mathcal{A} are reachable, it follows that there is a word $z \in (\Sigma \cup \{\$\})^*$ such that $q_1 \in \delta(Q_0, z)$, and thus $z \cdot (\$ \cdot x \cdot \$ \cdot y)^\omega \in L(\mathcal{A})$. Hence, as $L(\mathcal{A}) = \text{=}\bowtie_{\$}(R)$, we get that $x \in R$.

Since the state space of \mathcal{N} is $Q \cup \{q_{rej}\}$, we have that $|\mathcal{N}| = |\mathcal{A}| + 1$. Moreover, if R has a bad infix x , then $x \in \bar{R}$. Hence, if we consider the word x^ω , then as it has no $\$$'s, it is accepted by \mathcal{A} . Hence, there is a state $q \in Q$ such that \mathcal{A}^q has a run on x^ω that does not leave the $\bar{\alpha}$ -component of q . As we detail below, the fact that x is a bad infix of R implies that the $\bar{\alpha}$ -component of q does not contain $\$$ -transitions. Hence, since the only states in Q that are reachable in \mathcal{N} lie in $\bar{\alpha}$ -components that contain a $\$$ -transition, we can remove the state q_{rej} from \mathcal{N} and make q a rejecting sink instead. Hence, in this case, we have that $|\mathcal{N}| \leq |\mathcal{A}|$, and we are done.

Let x be a bad infix of R , and let $q \in Q$ be such that \mathcal{A}^q has a run on x^ω that does not leave the $\bar{\alpha}$ -component of q . Let $S(q)$ denote the $\bar{\alpha}$ -component of q . To conclude the proof, we show that $S(q)$ does not contain a $\$$ -transition. Assume by contradiction that $S(q)$ contains a $\$$ -transition $s_1 \xrightarrow{\$} s_2$. As q has a run on x^ω that does not leave $S(q)$, we get that q has a finite run of the form $q \xrightarrow{x} q_x$ on x that does not leave $S(q)$. As \mathcal{A} is normal, we get that there is a run in $S(q)$ that starts at q_x , passes through the transition $s_1 \xrightarrow{\$} s_2$, and finally ends in q : $q_x \xrightarrow{y_1} s_1 \xrightarrow{\$} s_2 \xrightarrow{y_2} q$. Hence, we get that there is a cycle in $S(q)$ of the form $q \xrightarrow{x \cdot y} q$ on a word of the form $x \cdot y$, where y contains a $\$$. Hence, $(x \cdot y)^\omega$ is accepted from q in \mathcal{A} .

Note that the word $(x \cdot y)^\omega = x \cdot y \cdot x \cdot y \cdots$ has infinitely many infixes of the form $\$ \cdot u \cdot \$$, where $u \in \Sigma^*$ is a bad infix of R . To see why, consider the infix $z = y \cdot x \cdot y$ that appears infinitely often in $(x \cdot y)^\omega$, and consider the infix $\$ \cdot u \cdot \$$ of $z = y \cdot x \cdot y$ that starts at last $\$$ of the first y and ends at the first $\$$ of the second y . Clearly, such u is over Σ , and since it contains x as an infix, then it is also a bad infix. In particular, $u \in \bar{R}$. As all the states of \mathcal{A} are reachable, there is a word $v \in (\Sigma \cup \{\$\})^*$ such that $q \in \delta(q_0, v)$. Hence, the word $v \cdot (x \cdot y)^\omega$ is accepted in \mathcal{A} even though it has infinitely many infixes of the form $\$ \cdot u \cdot \$$, where $u \in \bar{R}$, and we have reached a contradiction to the fact that $L(\mathcal{A}) = \text{=}\bowtie_{\$}(R)$. \square

4.1 Succinctness and Complementation

In this section we study the succinctness of SD co-Büchi automata with respect to deterministic ones, and the blow-up involved in their complementation. Recall that unlike Büchi automata, for co-Büchi automata, HD automata are exponentially more succinct than deterministic ones. Accordingly, our results about the succinctness of SD automata with respect to HD ones are more surprising than those in the setting of Büchi automata.

Theorem 4.3. *There is a family L_1, L_2, L_3, \dots of languages such that for every $n \geq 1$, there is an SD-tNCW with $3n + 3$ states that recognizes L_n , yet every tDCW or HD-tNCW that recognizes L_n needs at least 2^n states.*

Proof. For $n \geq 1$, consider the alphabet $\Sigma_n = [n] \cup \{\#\}$, and the language $R_n = \{x \cdot \# \cdot i : x \in [n]^+ \text{ and } i \text{ appears in } x\} \subseteq \Sigma_n^*$. We define $L_n = \text{=}\bowtie_{\$}(R_n)$. Recall that a word over $z \in (\Sigma_n \cup \{\$\})^*$ is good if $z = \$ \cdot x \cdot \# \cdot i$, where $x \in [n]^+$ and i appears in x , and note that L_n consists of exactly the words with finitely many $\$$'s, or that have a suffix that is a concatenation of good words. First, it is not hard to see that R_n can be recognized by an NFW \mathcal{N}_n with $3n + 3$ states, for example, a candidate for \mathcal{N}_n can be obtained from the NFW in Figure 1 by removing the $\$$ -transitions from q_0 , and letting q_0 guess to behave as any state in $\bigcup_{i \in [n]} \{q_i^1\}$. By Theorem 4.1, there is an SD-tNCW for L_n with $3n + 3$ states.

We show next that an HD-tNCW for L_n needs at least 2^n states. Consider an HD-tNCW $\mathcal{A} = \langle \Sigma_n \cup \{\$\}, Q, q_0, \delta, \alpha \rangle$ for L_n . By [19, 1], we can assume that \mathcal{A} is $\bar{\alpha}$ -deterministic, thus for all states q and letters σ , we have that there is at most one state q' such that $\langle q, \sigma, q' \rangle \in \bar{\alpha}$. We also assume that \mathcal{A} is normal, and all its states are reachable. To prove that $|Q| \geq 2^n$, we first prove that there is a state $q \in Q$ such that $(\$ \cdot R_n)^\omega \subseteq L_{\bar{\alpha}}(q)$, where $L_{\bar{\alpha}}(q) = \{w :$

there is a run from q on w that does not traverse α . Then, we argue that the $\bar{\alpha}$ -component of q has at least 2^n states.

For a nonempty set $S \in 2^Q$, and a letter $\sigma \in \Sigma_n \cup \{\$\}$, define $\delta^{\bar{\alpha}}(S, \sigma) = \{q' \in Q : \text{there is a state } q \in S \text{ such that } \langle q, \sigma, q' \rangle \in \bar{\alpha}\}$. Consider the tDCW $\mathcal{D} = \langle \Sigma_n \cup \{\$\}, 2^Q \setminus \emptyset, Q, \delta_{\mathcal{D}}, \alpha_{\mathcal{D}} \rangle$, where for every nonempty set $S \in 2^Q \setminus \emptyset$, and letter $\sigma \in \Sigma_n \cup \{\$\}$, it holds that $\delta_{\mathcal{D}}(S, \sigma) = \delta^{\bar{\alpha}}(S, \sigma)$, if $\delta^{\bar{\alpha}}(S, \sigma) \neq \emptyset$, and $\delta_{\mathcal{D}}(S, \sigma) = Q$, otherwise. The former type of transitions are in $\bar{\alpha}_{\mathcal{D}}$, and the latter are in $\alpha_{\mathcal{D}}$. Thus, the state space of \mathcal{D} encodes nonempty sets $S \in 2^Q$ of states, and \mathcal{D} starts a run from the set of all states Q . Reading its input, \mathcal{D} keeps track of runs in \mathcal{A} that have not traversed α from the last traversal of $\alpha_{\mathcal{D}}$, and when all tracked runs in \mathcal{A} have traversed α , the tDCW \mathcal{D} resets again to the set of all states Q .

We prove next that⁵ $L(\mathcal{A}) \subseteq L(\mathcal{D})$. Consider a word $w \in L(\mathcal{A})$, and let $r = r_0, r_1, r_2, \dots$ be an accepting run of \mathcal{A} on w . As r is accepting, there is $i \geq 0$ such that $r[i, \infty] = r_i, r_{i+1}, r_{i+2}, \dots$ does not traverse α on the suffix $w[i+1, \infty]$. Let $r_S = S_0, S_1, S_2, \dots$ denote the run of \mathcal{D} on w , and assume towards contradiction that r_S traverses $\alpha_{\mathcal{D}}$ infinitely often. Then, consider an index $j \geq i$ such that $\langle S_j, w[j+1], S_{j+1} \rangle \in \alpha_{\mathcal{D}}$. By the definition of \mathcal{D} , we have that $S_{j+1} = Q$. Consider the run $r[j+1, \infty]$ on the suffix $w[j+2, \infty]$. As $r_{j+1} \in S_{j+1}$, $r[j+1, \infty]$ is a run of r_{j+1} that does not traverse α , and \mathcal{D} keeps track of runs in \mathcal{A} that do not traverse α from the last reset, then it must be the case that the run $r_S[j+1, \infty] = S_{j+1}, S_{j+2}, S_{j+3}, \dots$ does not traverse $\alpha_{\mathcal{D}}$, and we have reached a contradiction.

Next, we show that there is a state $q \in Q$ such that $(\$ \cdot R_n)^\omega \subseteq L_{\bar{\alpha}}(q)$. Assume towards contradiction that for all states $q \in Q$, it holds that there is a word $w_q \in (\$ \cdot R_n)^\omega \setminus L_{\bar{\alpha}}(q)$. We define a sequence S_0, S_1, S_2, \dots of states in \mathcal{D} , and a sequence w_1, w_2, w_3, \dots of finite words in $(\$ \cdot R_n)^+$ such that for all $i \geq 0$, it holds that $\delta_{\mathcal{D}}(S_i, w_{i+1}) = S_{i+1}$. We take $S_0 = Q$. For all $i \geq 0$, given S_i , we define w_{i+1} and S_{i+1} as follows. Consider a state $q_i \in S_i$ and the word $w_{q_i} \in (\$ \cdot R_n)^\omega \setminus L_{\bar{\alpha}}(q_i)$. Then, as we argue below, there exists a prefix h_q of w_{q_i} in $(\$ \cdot R_n)^+$ such that all runs of q_i on h_q traverse α . We take w_{i+1} to be h_q , and take $S_{i+1} = \delta_{\mathcal{D}}(S_i, w_{i+1})$. We explain now why such a prefix exists. Assume contrarily that there are indices $k_1 < k_2 < k_3 < \dots$ such that for all $j \geq 1$, the word $w_{q_i}[1, k_j] = \sigma_1 \cdot \sigma_2 \cdot \dots \cdot \sigma_{k_j} \in (\$ \cdot R_n)^+$ is a prefix of $w_{q_i} = \sigma_1 \cdot \sigma_2 \cdot \sigma_3 \cdot \dots$ such that q_i has a run r_{k_j} that does not traverse α on $w_{q_i}[1, k_j]$. Then, as \mathcal{A} is $\bar{\alpha}$ deterministic, we get that the run r_{k_j} extends the run $r_{k_{j_1}}$ for all $j_2 > j_1$. Hence, there is an infinite run of q_i that does not traverse α on w_{q_i} , and we have reached a contradiction to the fact that $w_{q_i} \notin L_{\bar{\alpha}}(q_i)$.

Consider sequences S_0, S_1, S_2, \dots and w_1, w_2, w_3, \dots as above. First, it holds that the word $w = w_1 \cdot w_2 \cdot w_3 \cdot \dots$ is in $(\$ \cdot R_n)^\omega \subseteq L_n$, and, by definition, $r_S = S_0 \xrightarrow{w_1} S_1 \xrightarrow{w_2} S_2 \cdot \dots$ is a run of \mathcal{D} on w . Recall that $L(\mathcal{A}) = L_n \subseteq L(\mathcal{D})$; in particular $w \in L(\mathcal{D})$. To reach a contradiction, we prove below that the run r_S of \mathcal{D} on w , is rejecting. We claim that it is sufficient to show that for all $i \geq 0$, either the run $r_{i \rightarrow i+1} = S_i \xrightarrow{w_{i+1}} S_{i+1}$ traverses $\alpha_{\mathcal{D}}$, or $|S_i| > |S_{i+1}|$. Indeed, if the latter holds, and there is $i \geq 0$ such that the run $S_i \xrightarrow{w_{i+1}} S_{i+1} \xrightarrow{w_{i+2}} S_{i+2} \cdot \dots$ does not traverse $\alpha_{\mathcal{D}}$, then we get that $|S_j| > |S_{j+1}|$ for all $j \geq i$, which is a impossibility.

So we prove next that for all $i \geq 0$, either the run $r_{i \rightarrow i+1} = S_i \xrightarrow{w_{i+1}} S_{i+1}$ traverses $\alpha_{\mathcal{D}}$, or $|S_i| > |S_{i+1}|$. Consider a state S of \mathcal{D} , a letter $\sigma \in \Sigma_n \cup \{\$\}$, and note the following. If for all $q \in S$ it holds that $(\{q\} \times \{\sigma\} \times \delta(q, \sigma)) \cap \bar{\alpha} \neq \emptyset$, then as S proceeds to the $\bar{\alpha}$ σ successors of the states in S upon reading σ , and as \mathcal{A} is $\bar{\alpha}$ deterministic, we get that $|S| \geq |\delta_{\mathcal{D}}(S, \sigma)|$ and $\langle S, \sigma, \delta_{\mathcal{D}}(S, \sigma) \rangle \in \bar{\alpha}_{\mathcal{D}}$. Otherwise, if there is a state $q \in S$ such that $(\{q\} \times \{\sigma\} \times \delta(q, \sigma)) \cap \bar{\alpha} = \emptyset$, then either $\langle S, \sigma, \delta_{\mathcal{D}}(S, \sigma) \rangle \in \alpha_{\mathcal{D}}$, or $|S| > |\delta_{\mathcal{D}}(S, \sigma)|$. Consider $i \geq 0$. Let $w_{i+1} = \sigma_1 \cdot \sigma_2 \cdot \dots \cdot \sigma_{n_i}$, and consider the run $r_{i \rightarrow i+1} = S_i \xrightarrow{w_{i+1}} S_{i+1} = T_0, T_1, \dots, T_{n_i}$. If, by contradiction, the run $r_{i \rightarrow i+1}$ does not traverse $\alpha_{\mathcal{D}}$ and $|S_i| = |S_{i+1}|$, then by the above, for all $0 \leq j < n_i$ and $q \in T_j$, we have that $(\{q\} \times \{\sigma_{j+1}\} \times \delta(q, \sigma_{j+1})) \cap \bar{\alpha} \neq \emptyset$. Hence, if we consider a state $t_0 \in T_0$, we get that there is a transition $\langle t_0, \sigma_1, t_1 \rangle \in \bar{\alpha}$, and by the defition of $\delta_{\mathcal{D}}$, we have that $t_1 \in T_1$. Then, proceeding iteratively from t_1 ,

⁵In fact, using the fact that all the states of \mathcal{A} are reachable, the fact that HD-tNCWs can be assumed to be SD [19], and the fact that \mathcal{A} recognizes L_n , one can prove that \mathcal{D} is equivalent to \mathcal{A} .

we get that there is a run of t_0 in \mathcal{A} that does not traverse α on the word $w_{i+1} = \sigma_1 \cdot \sigma_2 \cdots \sigma_{n_i}$. As the latter holds for all $t_0 \in T_0 = S_i$; in particular, for $t_0 = q_i$, we get a contradiction to the definition of w_{i+1} .

What we have shown so far is that there exists a state $q \in Q$ such that $(\$ \cdot R_n)^\omega \subseteq L_{\bar{\alpha}}(q)$. To complete the proof, we prove next that the $\bar{\alpha}$ -component of q has at least 2^n states. Assume towards contradiction that the $\bar{\alpha}$ -component of q has strictly less than 2^n states. For every subset $S \subseteq [n]$, we abuse notation and use S to also denote a word in $[n]^*$ consisting of exactly the numbers that appear in S . Then, consider the word $\$ \cdot S$. The word $\$ \cdot S$ can be extended to a word in $(\$ \cdot R_n)^\omega \subseteq L_{\bar{\alpha}}(q)$. Hence, q has a run r_S that does not traverse α on the word $\$ \cdot S$, and as \mathcal{A} is normal, we have that r_S ends in a state that belongs to the $\bar{\alpha}$ -component of q . As the $\bar{\alpha}$ -component of q has strictly less than 2^n states, there are sets $S, T \subseteq [n]$ and $i \in [n]$ such that $i \in S \setminus T$, and r_S and r_T end in the same state q' . As \mathcal{A} is $\bar{\alpha}$ deterministic, we get that the run r_S is the only run that does not traverse α from the state q on the word $\$ \cdot S$. Hence, as $\$ \cdot S \cdot \# \cdot i \cdot \$$ can be extended to a word in $(\$ \cdot R_n)^\omega \subseteq L_{\bar{\alpha}}(q)$ and r_S ends in q' , we get that there is a run of the form $q' \xrightarrow{\# \cdot i \cdot \$} q''$ that traverses only $\bar{\alpha}$ transitions, in particular, it does not leave the $\bar{\alpha}$ -component of q . Now consider the run $r = q \xrightarrow{\$ \cdot T} q' \xrightarrow{\# \cdot i \cdot \$} q''$. The run r is a run that does not traverse α from the state q on the word $\$ \cdot T \cdot \# \cdot i \cdot \$$. As \mathcal{A} is normal, there is a word $y \in (\Sigma_n \cup \{\$\})^*$ such that there is a run that does traverse α of $\mathcal{A}^{q''}$ on y that reaches q . Therefore, $(\$ \cdot T \cdot \# \cdot i \cdot \$ \cdot y)^\omega \in L_{\bar{\alpha}}(q) \subseteq L(\mathcal{A}^q)$. As all the states of \mathcal{A} are reachable, there is a word $z \in (\Sigma_n \cup \{\$\})^*$ such that $q \in \delta(q_0, z)$. Hence, the word $z \cdot (\$ \cdot T \cdot \# \cdot i \cdot \$ \cdot y)^\omega$ is accepted by \mathcal{A} even though it has infinitely many infixes in $\$ \cdot \overline{R_n} \cdot \$$, and we have reached a contradiction. \square

Remark 4.1. As has been the case with Büchi automata, here too, an analysis of the application of the HD-tNCW minimization algorithm on a tDCW for L_n leads to a slightly tighter bound. Specifically, in Appendix A.2, we describe a tDCW for L_n with $2^{n+1} + 1$ states, such that the application of the HD-tNCW minimization algorithm on it does not make it smaller. \square

Theorem 4.4. *There is a family L_1, L_2, L_3, \dots of languages such that for every $n \geq 1$, there is an SD-tNCW with $O(n)$ states that recognizes L_n , yet every SD-tNBW that recognizes $\overline{L_n}$ needs at least $2^{O(n)}$ states.*

Proof. Let $\Sigma = \{0, 1\}$, and let $R_n = \{w : w \in (0+1)^* \cdot (0 \cdot (0+1)^{n-1} \cdot 1 + 1 \cdot (0+1)^{n-1} \cdot 0) \cdot (0+1)^*\}$. Thus, R_n is the language of all words that contain two letters that are at distance n and are different. We define $L_n = \bowtie_{\$}(R_n)$.

It is easy to see that R_n can be recognized by an NFW with $O(n)$ states. Hence, by Theorem 4.1, there is an SD-tNCW with $O(n)$ states for L_n . We prove next that every SD-tNBW for $\overline{L_n}$ needs at least $2^{O(n)}$ states. For this, note that $\overline{L_n}$ consists of all words w such that w contains infinitely many $\$$'s yet has no suffix in $(\$R_n)^\omega$. Thus, w contains infinitely many infixes in $\$ \overline{R_n} \$$. Therefore, $\overline{L_n} = \infty(\$ \cdot \overline{R_n} \cdot \$)$. It is not hard to prove that an NFW for $\overline{R_n}$ needs at least $2^{O(n)}$ states. Then, by Theorem 3.2, this bound is carried over to a $2^{O(n)}$ lower bound on an SD-tNBW for $\overline{L_n}$. \square

4.2 Decision Problems

We continue to decision problems for SD-tNCWs and SD-NCWs. As in Section 3.2, we state only the lower bounds, and for SD-tNCWs. Matching upper bounds and similar results for SD-NCWs follow from the known upper bounds for tNCWs and the known linear translations between state-based and transition-based automata.

We start with language-containment and universality.

Theorem 4.5. *The language-containment and universality problems for SD-tNCWs are PSPACE-hard.*

Proof. We describe a reduction from universality of NFWs to universality of SD-tNCWs. Given an NFW \mathcal{N} over Σ , the reduction returns the SD-tNCW \mathcal{A} over $\Sigma \cup \{\$\}$ that is obtained by applying the operation from Theorem 4.1 on \mathcal{N} . Thus, $L(\mathcal{A}) = \bowtie_{\$}(L(\mathcal{N}))$.

First, if $L(\mathcal{N}) = \Sigma^*$, then $(\Sigma \cup \{\$\})^* \cdot (\$L(\mathcal{N}))^\omega = \infty \$$, and so $L(\mathcal{A}) = (\Sigma \cup \{\$\})^\omega$. Conversely, if there is a word $x \in \Sigma^* \setminus L(\mathcal{N})$, then the word $w = (\$x)^\omega$ is not in $L(\mathcal{A})$. Indeed, w has infinitely many $\$$'s, yet it has a word not in $L(\mathcal{N})$ between every two consecutive $\$$'s. \square

We continue to the minimization problem. Note that while minimization is PSPACE-complete for NCWs, it is NP-complete for DCWs and in PTIME for HD-tNCWs. Thus, our PSPACE lower bound suggests a significant difference between SD and HD automata. Note also that, as has been the case with Büchi automata, the case of SD-NCWs is easy, as a non-empty SD-NCW is universal iff it has an equivalent SD-NCW with one state, which is not the case for SD-tNCWs:

Theorem 4.6. *The minimization problem for SD-tNCWs is PSPACE-hard.*

Proof. We analyze the reduction from Theorem 4.5 and show that an NFW \mathcal{N} over Σ is universal iff the SD-tNCW \mathcal{A} over $\Sigma \cup \{\$\}$ that is constructed in Theorem 4.1 for $\bowtie_{\$}(L(\mathcal{N}))$ has an equivalent SD-tNCW with a single state.

First, note that \mathcal{N} is universal iff $\bowtie_{\$}(L(\mathcal{N}))$ is universal. Hence, if \mathcal{N} is universal, then \mathcal{A} has an equivalent SD-tNCW with a single state. For the other direction, assume that \mathcal{A} has an equivalent SD-tNCW \mathcal{D} with a single state, and assume towards contradiction that \mathcal{N} is not universal. Then, there is a word $x \notin L(\mathcal{N})$. Consider the word $(\$x)^\omega$. By the definition of $\bowtie_{\$}(L(\mathcal{N}))$, we have that $(\$x)^\omega \notin L(\mathcal{A})$. Hence, $(\$x)^\omega$ is not accepted by \mathcal{D} and so at least one of the letters σ in the word $\$x$ is such that the σ -labeled self-loop at the state of \mathcal{D} is an α -transition. Now as $L(\mathcal{A}) = \bowtie_{\$}(L(\mathcal{N}))$, we get that \mathcal{D} accepts all words that have no $\$$'s; in particular, the α -self-loop at the state of \mathcal{D} must be $\$$ -labeled.

Recall that the universality problem for NFWs is PSPACE-hard already for NFWs \mathcal{N} that accept ϵ . For example, this is valid for the NFW that is generated in the generic reduction from a PSPACE Turing machine, and which accepts all words that do not encode a valid computation of the Turing machine. Hence, for such an NFW \mathcal{N} , the word $\$^\omega$ is in $\bowtie_{\$}(L(\mathcal{N}))$, and should be accepted by \mathcal{D} , and we have reached a contradiction. \square

We continue to the D-to-SD minimization problem, showing it stays PSPACE-hard.

Theorem 4.7. *The D-to-SD minimization problem for co-Büchi automata is PSPACE-hard.*

Proof. We reduce from the D-to-N minimization problem for automata on finite words, which is already PSPACE-hard for languages that have bad infixes [18]. We start with co-Büchi automata with transition-based acceptance.

Consider the construction from Theorem 4.1. Recall it takes an NFW \mathcal{A} as input, returns an SD-tNCW \mathcal{A}' of the same size for $\bowtie_{\$}(L(\mathcal{A}))$, and preserves determinism.

Consider a DFW \mathcal{A} over Σ such that $L(\mathcal{A})$ has a bad infix, and an integer k . The reduction returns the SD-tNCW \mathcal{A}' constructed from \mathcal{A} in Theorem 4.1, and the integer k . Recall that $L(\mathcal{A}') = \bowtie_{\$}(L(\mathcal{A}))$, and that the construction in Theorem 4.1 preserves determinism. Thus, the automaton \mathcal{A}' is really a tDCW.

We prove next that the reduction is correct. That is, the DFW \mathcal{A} has an equivalent NFW with at most k states iff the tDCW \mathcal{A}' has an equivalent SD-tNCW with at most k states. For the first direction, if \mathcal{B} is an NFW equivalent to \mathcal{A} whose size is at most k , then, by applying to it the construction from Theorem 4.1, we get an SD-tNCW \mathcal{B}' whose size is at most k , and $L(\mathcal{B}') = \bowtie_{\$}(L(\mathcal{B})) = \bowtie_{\$}(L(\mathcal{A})) = L(\mathcal{A}')$.

Conversely, if \mathcal{B}' is an SD-tNCW for $\bowtie_{\$}(L(\mathcal{A}))$ whose size is at most k , then as $L(\mathcal{A})$ has bad infixes, we get by Theorem 4.2 that there is an NFW \mathcal{B} for $L(\mathcal{A})$ whose size is at most k , and we are done.

Finally, as in Büchi automata, it is not hard to see that the above arguments can be adapted to state-based automata, and so the D-to-SD minimization problem is PSPACE-hard also for co-Büchi automata with state-based acceptance. Essentially, given a DFW \mathcal{A} and an integer k , the reduction returns a DCW for $\bowtie_{\$}(L(\mathcal{A}))$ and the integer $k + 1$.

The above reduction exists since the construction from Theorem 4.1 can be adapted to state-based automata, essentially, by letting the co-Büchi automaton have one additional α state that behaves as the initial states, and that is reachable from every state in $Q \setminus F$ upon reading $\$$.

Also, Theorem 4.2 can be adapted to the state-based setting. Indeed, if \mathcal{A} is an NCW for $\bowtie_{\$}(R)$ for some language $R \subseteq \Sigma^*$ that has a bad infix x , then one can construct an NFW \mathcal{N} for R with at most $|\mathcal{A}| - 1$ states. To see why, we view \mathcal{A} as a transition-based automaton \mathcal{A}_t by considering transitions that touch α states as α transitions. Then, we consider the proof of Theorem 4.2 when applied on \mathcal{A}_t . We claim that all states p that are α states of \mathcal{A} can be removed from \mathcal{N} without affecting its language. Indeed, the $\bar{\alpha}$ -component of such a state p in \mathcal{A}_t is $\{p\}$ and contains no transitions, in particular, no α -transitions. Also, note that such states p are distinct from the state q that has a run that does not traverse α transitions on x^ω . Thus, they can be removed from \mathcal{N} along with the state q_{rej} . Hence, we get in this case that $|\mathcal{N}| \leq |\mathcal{A}| - 1$. \square

5 Semantically Deterministic Weak Automata

By [2], SD-NWWs need not be DBP or even HD. For completeness we describe here the example from [2], as it highlights the challenges in SD-NWW determinization. Consider the automaton \mathcal{A} in Figure 2.




Figure 2: An SD-NWW that is not DBP.

It is easy to see that \mathcal{A} is weak, and all its states are universal, and so it is SD. On the other hand, \mathcal{A} is not HD as every strategy has a word with which it does not reach q_{acc} – a word that forces every visit in q_a and q_b to be followed by a visit in q_0 . Below we show that despite not being DBP, SD-NWWs can be determinized in polynomial time. Essentially, our proof is based on redirecting transitions of the SD-NWW to deep components in the automaton. In our example, note that while the SD-NWW \mathcal{A} is not HD, it has a deterministic state q_{acc} that recognizes $L(\mathcal{A})$.

Consider an SD-NWW $\mathcal{A} = \langle \Sigma, Q, q_0, \delta, \alpha \rangle$. We denote the set of \mathcal{A} 's SCCs by $\mathcal{C}(\mathcal{A})$, and the SCC containing a state q by $C(q)$. Let $C_1 \leq C_2 \leq \dots \leq C_m$ be a total order on the SCCs of \mathcal{A} , extending the partial order induced by δ . That is, if $q' \in \delta(q, \sigma)$, for some letter σ , then $C(q) \leq C(q')$. When $C \leq C'$, we say that C' is *deeper than* C . Thus, states along runs of a weak automaton proceed from SCCs to deeper ones, and eventually get stuck in some SCC.

If we had an algorithm that checks language equivalence between states in \mathcal{A} in polynomial time, we could have a polynomial determinization algorithm that defines the σ -successor of a state as the deepest state among all states equivalent to its σ -successors (since \mathcal{A} is SD, all these successors agree on their language). Since we still do not have such an algorithm (in fact, it would follow from our construction), we approximate language equivalence by an equivalence that follows from the semantic determinism of \mathcal{A} .

For two states $s_1, s_2 \in Q$, we say that s_1 and s_2 are δ -close, if there is a state $q \in Q$ and a word $w \in \Sigma^*$ such that $s_1, s_2 \in \delta(q, w)$. Note that the δ -close relation refines equivalence, yet the converse does not hold. Indeed, the SDness property implies that $s_1 \sim_{\mathcal{A}} s_2$ for all δ -close states s_1 and s_2 .

Lemma 5.1. *If s_1 and s_2 are δ -close, then there is a state q and word w of length at most $|Q|^2$ such that $s_1, s_2 \in \delta(q, w)$.*

Proof. Consider δ -close states s_1 and s_2 . Let q and w be such that $s_1, s_2 \in \delta(q, w)$. If $|w| > |Q|^2$, then in the runs from q to s_1 and from q to s_2 there must be a pair of states $\langle p_1, p_2 \rangle$ that repeats twice, and we can shorten w by removing the subword between the two occurrences of $\langle p_1, p_2 \rangle$. Formally, there is a partition of $w = x \cdot y \cdot z$ such that $p_1, p_2 \in \delta(q, x)$, $p_1 \in \delta(p_1, y)$, $p_2 \in \delta(p_2, y)$, $s_1 \in \delta(p_1, z)$, and $s_2 \in \delta(p_2, z)$. Then, $s_1, s_2 \in \delta(q, x \cdot z)$. This process can repeat until w is of length at most $|Q|^2$. \square

In order to calculate the δ -close relation in polynomial time, we define a sequence $H_0, H_1, H_2, \dots \subseteq Q \times Q$ of relations, where $\langle s_1, s_2 \rangle \in H_i$ iff there is a state q and word w of length at most i such that $s_1, s_2 \in \delta(q, w)$. By Lemma 5.1, we are guaranteed to reach a fixed point after at most $|Q|^2$ iterations. The relations H_i are defined as follows.

- $H_0 = \{\langle q, q \rangle : q \in Q\}$.
- For $i \geq 0$, we define $H_{i+1} = H_i \cup \{\langle s_1, s_2 \rangle : \text{there is } \langle q_1, q_2 \rangle \in H_i \text{ and letter } \sigma \in \Sigma \text{ such that } s_1 \in \delta(q_1, \sigma) \text{ and } s_2 \in \delta(q_2, \sigma)\}$.

Let $j \geq 0$ be such that $H_{j+1} = H_j$. It is not hard to see that H_j is the δ -close relation. While the δ -close relation is reflexive and symmetric, it is not transitive. Now, let $H \subseteq Q \times Q$ be the closure of H_j under transitivity. That is, $\langle s_1, s_2 \rangle \in H$ iff there is $k \geq 2$ and states q_1, q_2, \dots, q_k such that $q_1 = s_1$, $q_k = s_2$ and $\langle q_i, q_{i+1} \rangle \in H_j$ for all $1 \leq i < k$.

The following lemma implies that H propagates to successor states.

Lemma 5.2. *If $H(s, s')$, then for every letter $\sigma \in \Sigma$ and states $q \in \delta(s, \sigma)$ and $q' \in \delta(s', \sigma)$, we have that $H(q, q')$.*

Proof. Consider the fixed point relation H_j . By definition, if $H_j(s, s')$, then for every letter $\sigma \in \Sigma$ and states $q \in \delta(s, \sigma)$ and $q' \in \delta(s', \sigma)$, we have that $H_{j+1}(q, q')$, and since H_j is the fixed point, then $H_j(q, q')$. Then, the lemma follows from the definition of H as the closure of H_j under transitivity. Indeed, if $H(s, s')$, then there is $k \geq 2$ and states q_1, q_2, \dots, q_k such that $q_1 = s, q_k = s'$ and $H_j(q_i, q_{i+1})$ for all $1 \leq i < k$. Now if s_1, s_2, \dots, s_k are such that $s_i \in \delta(q_i, \sigma)$, then by the above considerations, we have that $H_j(s_i, s_{i+1})$ for all $1 \leq i < k$. Hence, the definition of H implies that $H(s_1, s_k)$, and we are done. \square

It is easy to see that H is an equivalence relation. Let $\mathcal{P} = \{P_1, \dots, P_k\}$ be the set of the equivalence classes of H . For each equivalence class $P \in \mathcal{P}$, we fix the *representative* of P as some state in $P \cap C$, where $C \in \mathcal{C}(\mathcal{A})$ is the deepest SCC that intersects P . Let p_1, \dots, p_k be the representatives of the sets in \mathcal{P} . For a state $q \in Q$, let \tilde{q} denote the representative of the set $P \in \mathcal{P}$ with $q \in P$. Note that as H refines $\sim_{\mathcal{A}}$, we have that $q \sim_{\mathcal{A}} \tilde{q}$.

Theorem 5.3. *Given an SD-NWW \mathcal{A} with state space Q , we can construct, in polynomial time, an equivalent DWW \mathcal{D} with state space Q' , for $Q' \subseteq Q$.*

Proof. Given $\mathcal{A} = \langle \Sigma, Q, q_0, \delta, \alpha \rangle$, we define $\mathcal{D} = \langle \Sigma, Q', q'_0, \delta', \alpha \cap Q' \rangle$, where

- $Q' = \{p_1, \dots, p_k\}$. Note that indeed $Q' \subseteq Q$.
- $q'_0 = \tilde{q}_0$.

- For $p \in Q'$ and $\sigma \in \Sigma$, we define $\delta'(p, \sigma) = \tilde{q}$, for some $q \in \delta(p, \sigma)$. Note that all the states in $\delta(p, \sigma)$ are δ -close, and thus belong to the same set in \mathcal{P} . Hence, the choice of q is not important, as $\delta'(p, \sigma)$ is the representative of this set.

We prove that \mathcal{D} is a DWW equivalent to \mathcal{A} . First, in order to see that \mathcal{D} is weak, consider states p, p' such that $p' \in \delta'(p, \sigma)$. Thus, p' is the representative of a state $q \in \delta(p, \sigma)$. As \mathcal{A} is weak, we have that $C(p) \leq C(q)$. As $p' = \tilde{q}$, we have that $C(q) \leq C(p')$. Hence $C(p) \leq C(p')$, and we are done.

We continue and prove that $L(\mathcal{A}) = L(\mathcal{D})$. We first prove that $L(\mathcal{A}) \subseteq L(\mathcal{D})$. Consider a word w and assume that $w \in L(\mathcal{A})$. Let $r = q_0, q_1, q_2, \dots$ be an accepting run of \mathcal{A} on w , and let $r' = s_0, s_1, s_2, \dots$ be the run of \mathcal{D} on w . If r' is accepting, we are done. Otherwise, namely if r' is rejecting, we point to a contradiction. Let $j \geq 0$ be the index in which r' visits only states in $\text{sinf}(r')$. Note that all the states $s_j, s_{j+1}, s_{j+2}, \dots$ belong to some SCC C of \mathcal{A} . Since we assume that r' is rejecting, and acceptance in \mathcal{D} is inherited from \mathcal{A} , we get that C is a rejecting SCC of \mathcal{A} . We claim that all the runs of \mathcal{A}^{s_j} on the suffix $w[j+1, \infty]$ of w are stuck in C . Thus, $w[j+1, \infty] \notin L(\mathcal{A}^{s_j})$. On the other hand, we claim that for all $i \geq 0$, we have that $q_i \sim_{\mathcal{A}} s_i$. Then, however, $w[j+1, \infty] \notin L(\mathcal{A}^{q_j})$, contradicting the fact that r is accepting:

The easy part is to prove that for all $i \geq 0$, we have that $q_i \sim_{\mathcal{A}} s_i$. Indeed, it follows from the fact that for all $i \geq 0$, we have that $H(q_i, s_i)$ and the fact that H refines $\sim_{\mathcal{A}}$. The proof proceeds by an induction on i . First, as for every state $q \in Q$, we have that $H(q, \tilde{q})$, then $H(q_0, q'_0)$, and so the claim follows from s_0 being q'_0 . Assume now that $H(q_i, s_i)$. Recall that s_{i+1} is \tilde{q} for some $q \in \delta(s_i, w[i+1])$. Also, $q_{i+1} \in \delta(q_i, w[i+1])$. Then, by Lemma 5.2, we have that $H(q_{i+1}, q)$. Since, in addition, we have that $H(q, \tilde{q})$, then the transitivity of H implies that $H(q_{i+1}, s_{i+1})$, and we are done.

We proceed to the other part, namely, proving that all the runs of \mathcal{A}^{s_j} on the suffix $w[j+1, \infty]$ of w are stuck in C . Let u_{j+1}, u_{j+2}, \dots be such that for all $i \geq 1$, it holds that $u_{j+i} \in \delta(s_{j+i-1}, w[j+i])$ and $H(u_{j+i}, s_{j+i})$. Note that such states exist, as $s_{j+i-1} \xrightarrow{w[j+i]} s_{j+i}$ is a transition of \mathcal{D} and so $s_{j+i} = \tilde{q}$ for all $q \in \delta(s_{j+i-1}, w[j+i])$. Consider a run $r'' = v_0, v_1, v_2, \dots$ of \mathcal{A}^{s_j} on $w[j+1, \infty]$. Note that $v_0 = s_j$, and so $H(v_0, s_j)$. Therefore, by Lemma 5.2, we have that $H(v_1, u_{j+1})$, implying $H(v_1, s_{j+1})$. By repeated applications of Lemma 5.2, we get that $H(v_i, u_{j+i})$, implying $H(v_i, s_{j+i})$, for all $i \geq 1$. Assume now by way of contradiction that the run r'' leaves the SCC C , and so there is $i \geq 1$ such that $C(v_i) \not\subseteq C(s_{j+i})$. As $H(v_i, s_{j+i})$, the states v_i and s_{j+i} are in the same equivalence class $P \in \mathcal{P}$. Then, the definition of Q' implies that $C(s_{j+i}) \geq C(v_i)$, and we have reached a contradiction.

It is left to prove that $L(\mathcal{D}) \subseteq L(\mathcal{A})$. The proof follows similar considerations and is detailed next. Consider a word $w \in L(\mathcal{D})$. Let $r' = s_0, s_1, s_2, \dots$ be the run of \mathcal{D} on w , and let $j \geq 0$ be the index in which r' visits only states in $\text{sinf}(r')$; in particular, as argued above, $\text{sinf}(r')$ is included in some SCC C of \mathcal{A} . Thus, all the states $s_j, s_{j+1}, s_{j+2}, \dots$ are in C . Since $w \in L(\mathcal{D})$, then r' is accepting, and so C is an accepting SCC of \mathcal{A} . As in the previous direction, it holds that s_j is \mathcal{A} -equivalent to all the states in $\delta(q_0, w[1, j])$. Hence, to conclude that $w \in L(\mathcal{A})$, we show that there is an accepting run of \mathcal{A}^{s_j} on $w[j+1, \infty]$. Assume by way of contradiction that all the runs of \mathcal{A}^{s_j} on $w[j+1, \infty]$ are rejecting. In particular, all these runs leave the SCC C . As in the previous direction, this contradicts the choice of the states $s_j, s_{j+1}, s_{j+2}, \dots$ as representatives of equivalence classes in \mathcal{P} . \square

Since DWWs can be complemented by dualization (that is, by switching α and $\bar{\alpha}$), Theorem 5.3 implies the following.

Theorem 5.4. *Given an SD-NWW \mathcal{A} with n states, we can construct, in polynomial time, an SD-NWW (in fact, a DWW) that complements \mathcal{A} .*

Since the language-containment problem for DWW can be solved in NLOGSPACE, and minimization for DWW is similar to minimization of DFWs and can be solved in polynomial time [30], we have the following.

Theorem 5.5. *The language-containment, universality, and minimality problems for SD-NWWs can be solved in polynomial time.*

References

- [1] B. Abu Radi and O. Kupferman. Minimizing GFG transition-based automata. In *Proc. 46th Int. Colloq. on Automata, Languages, and Programming*, volume 132 of *LIPICs*, pages 100:1–100:16. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2019.
- [2] B. Abu Radi, O. Kupferman, and O. Leshkowitz. A hierarchy of nondeterminism. In *46th Int. Symp. on Mathematical Foundations of Computer Science*, volume 202 of *LIPICs*, pages 85:1–85:21, 2021.
- [3] B. Aminof, O. Kupferman, and R. Lampert. Reasoning about online algorithms with weighted automata. *ACM Transactions on Algorithms*, 6(2), 2010.
- [4] M. Bagnol and D. Kuperberg. Büchi good-for-games automata are efficiently recognizable. In *Proc. 38th Conf. on Foundations of Software Technology and Theoretical Computer Science*, volume 122 of *LIPICs*, pages 16:1–16:14, 2018.
- [5] B. Boigelot, S. Jodogne, and P. Wolper. On the use of weak automata for deciding linear arithmetic with integer and real variables. In *Proc. Int. Joint Conf. on Automated Reasoning*, volume 2083 of *Lecture Notes in Computer Science*, pages 611–625. Springer, 2001.
- [6] U. Boker, D. Kuperberg, O. Kupferman, and M. Skrzypczak. Nondeterminism in the presence of a diverse or unknown future. In *Proc. 40th Int. Colloq. on Automata, Languages, and Programming*, volume 7966 of *Lecture Notes in Computer Science*, pages 89–100, 2013.
- [7] U. Boker and K. Lehtinen. When a little nondeterminism goes a long way: an introduction to history-determinism. *ACM SIGLOG News*, 10(1):177–196, 2023.
- [8] 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.
- [9] Olivier Carton and Max Michel. Unambiguous büchi automata. *Theoretical Computer Science*, 297(1):37–81, 2003.
- [10] T. Colcombet. The theory of stabilisation monoids and regular cost functions. In *Proc. 36th Int. Colloq. on Automata, Languages, and Programming*, volume 5556 of *Lecture Notes in Computer Science*, pages 139–150. Springer, 2009.
- [11] T. Colcombet, K. Quaas, and M. Skrzypczak. Unambiguity in automata theory (dagstuhl seminar 21452). *Dagstuhl Reports*, 11(10):57–71, 2021.
- [12] A. Duret-Lutz, A. Lewkowicz, A. Fauchille, Th. Michaud, E. Renault, and L. Xu. Spot 2.0 — a framework for LTL and ω -automata manipulation. In *14th Int. Symp. on Automated Technology for Verification and Analysis*, volume 9938 of *Lecture Notes in Computer Science*, pages 122–129. Springer, 2016.
- [13] P. Gastin and D. Oddoux. Fast LTL to Büchi automata translation. In *Proc. 13th Int. Conf. on Computer Aided Verification*, volume 2102 of *Lecture Notes in Computer Science*, pages 53–65. Springer, 2001.
- [14] D. Giannakopoulou and F. Lerda. From states to transitions: Improving translation of LTL formulae to Büchi automata. In *Proc. 22nd International Conference on Formal Techniques for Networked and Distributed Systems*, volume 2529 of *Lecture Notes in Computer Science*, pages 308–326. Springer, 2002.

- [15] E.M. Hahn, M. Perez, S. Schewe, F. Somenzi, A. Trivedi, and D. Wojtczak. Good-for-MDPs automata for probabilistic analysis and reinforcement learning. In *Proc. 26th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*, volume 12078 of *Lecture Notes in Computer Science*, pages 306–323. Springer, 2020.
- [16] T.A. Henzinger, O. Kupferman, and S. Rajamani. Fair simulation. *Information and Computation*, 173(1):64–81, 2002.
- [17] T.A. Henzinger and N. Piterman. Solving games without determinization. In *Proc. 15th Annual Conf. of the European Association for Computer Science Logic*, volume 4207 of *Lecture Notes in Computer Science*, pages 394–410. Springer, 2006.
- [18] T. Jiang and B. Ravikumar. Minimal NFA problems are hard. *SIAM Journal on Computing*, 22(6):1117–1141, 1993.
- [19] D. Kuperberg and M. Skrzypczak. On determinisation of good-for-games automata. In *Proc. 42nd Int. Colloq. on Automata, Languages, and Programming*, pages 299–310, 2015.
- [20] O. Kupferman. Avoiding determinization. In *Proc. 21st IEEE Symp. on Logic in Computer Science*, pages 243–254, 2006.
- [21] O. Kupferman. Automata theory and model checking. In *Handbook of Model Checking*, pages 107–151. Springer, 2018.
- [22] O. Kupferman. Using the past for resolving the future. *Frontiers in Computer Science*, 4, 2023.
- [23] O. Kupferman and O. Leshkowitz. On repetition languages. In *45th Int. Symp. on Mathematical Foundations of Computer Science*, Leibniz International Proceedings in Informatics (LIPIcs), 2020.
- [24] O. Kupferman, S. Safra, and M.Y. Vardi. Relating word and tree automata. *Ann. Pure Appl. Logic*, 138(1-3):126–146, 2006.
- [25] O. Kupferman and M.Y. Vardi. Safraless decision procedures. In *Proc. 46th IEEE Symp. on Foundations of Computer Science*, pages 531–540, 2005.
- [26] R.P. Kurshan. Complementing deterministic Büchi automata in polynomial time. *Journal of Computer and Systems Science*, 35:59–71, 1987.
- [27] R.P. Kurshan. *Computer Aided Verification of Coordinating Processes*. Princeton Univ. Press, 1994.
- [28] L.H. Landweber. Decision problems for ω -automata. *Mathematical Systems Theory*, 3:376–384, 1969.
- [29] W. Li, Sh. Kan, and Z. Huang. A better translation from LTL to transition-based generalized Büchi automata. *IEEE Access*, 5:27081–27090, 2017.
- [30] C. Löding. Efficient minimization of deterministic weak ω -automata. *Information Processing Letters*, 79(3):105–109, 2001.
- [31] S. Miyano and T. Hayashi. Alternating finite automata on ω -words. *Theoretical Computer Science*, 32:321–330, 1984.
- [32] G. Morgenstern. Expressiveness results at the bottom of the ω -regular hierarchy. M.Sc. Thesis, The Hebrew University, 2003.

- [33] D.E. Muller, A. Saoudi, and P. E. Schupp. Weak alternating automata give a simple explanation of why most temporal and dynamic logics are decidable in exponential time. In *Proc. 3rd IEEE Symp. on Logic in Computer Science*, pages 422–427, 1988.
- [34] M.O. Rabin and D. Scott. Finite automata and their decision problems. *IBM Journal of Research and Development*, 3:115–125, 1959.
- [35] S. Safra. On the complexity of ω -automata. In *Proc. 29th IEEE Symp. on Foundations of Computer Science*, pages 319–327, 1988.
- [36] S. Schewe. Beyond Hyper-Minimisation—Minimising DBAs and DPAs is NP-Complete. In *Proc. 30th Conf. on Foundations of Software Technology and Theoretical Computer Science*, volume 8 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 400–411, 2010.
- [37] S. Schewe. Minimising good-for-games automata is NP-complete. In *Proc. 40th Conf. on Foundations of Software Technology and Theoretical Computer Science*, volume 182 of *LIPIcs*, pages 56:1–56:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [38] S. Schewe, Q. Tang, and T. Zhanabekova. Deciding what is good-for-MDPs. *CoRR*, abs/2202.07629, 2022. URL: <https://arxiv.org/abs/2202.07629>, arXiv:2202.07629.
- [39] S. Sickert, J. Esparza, S. Jaax, and J. Křetínský. Limit-deterministic Büchi automata for linear temporal logic. In *Proc. 28th Int. Conf. on Computer Aided Verification*, volume 9780 of *Lecture Notes in Computer Science*, pages 312–332. Springer, 2016.
- [40] M.Y. Vardi and P. Wolper. Reasoning about infinite computations. *Information and Computation*, 115(1):1–37, 1994.

A Missing Details

A.1 Missing details in Remark 3.1

Consider the tDBW $\mathcal{D}_n = \langle \Sigma_n, 2^{[n]} \times \{a, c\}, \langle \emptyset, c \rangle, \delta, \alpha \rangle$, defined as follows (see Figure 3). The state space of \mathcal{D}_n consists of two copies of subsets of $[n]$: the a -copy $2^{[n]} \times \{a\}$, with “a” standing for “accumulate”, and the c -copy $2^{[n]} \times \{c\}$, with “c” standing for “check”. Essentially, the state $\langle S, a \rangle$ in the a -copy remembers that we have read only numbers in $[n]$ after the last \$ and also remembers the set $S \subseteq [n]$ of these numbers. The state $\langle S, c \rangle$ in the c -copy remembers, in addition, that we have just read #. Accordingly, $\langle S, c \rangle$ checks whether the next letter is in S in order to traverse α . Effectively, \mathcal{D}_n traverses α when a good infix is detected. Formally, for every state $\langle S, o \rangle \in 2^{[n]} \times \{a, c\}$, and letter $\sigma \in \Sigma_n$, we define:

$$\delta(\langle S, a \rangle, \sigma) = \begin{cases} \langle S \cup \{\sigma\}, a \rangle, & \text{if } \sigma \in [n] \\ \langle S, c \rangle, & \text{if } \sigma = \# \\ \langle \emptyset, a \rangle, & \text{if } \sigma = \$ \end{cases} \quad \delta(\langle S, c \rangle, \sigma) = \begin{cases} \langle \emptyset, c \rangle, & \text{if } \sigma \neq \$ \\ \langle \emptyset, a \rangle, & \text{if } \sigma = \$ \end{cases}$$

Also, $\alpha = \{ \langle \langle S, c \rangle, \sigma, \langle \emptyset, c \rangle \rangle : \sigma \in S \}$.

Note that reading \$ from any state we move to $\langle \emptyset, a \rangle$ to detect a good infix, and if at some point we read an unexpected letter, we move to $\langle \emptyset, c \rangle$, where we wait for \$ to be read. It is easy to see that \mathcal{D}_n recognizes L_n and has 2^{n+1} states.




Figure 3: The tDBW D_n . For a finite word $x \in [n]^+$, we denote the set of numbers that appear in x by $S(x)$. Dashed transitions are α -transitions.

A.2 Missing details in Remark 4.1

Consider the tDCW $\mathcal{D}_n = \langle \Sigma_n \cup \{\#\}, (2^{[n]} \times \{a, c\}) \cup \{q_{pass}\}, \langle \emptyset, c \rangle, \delta, \alpha \rangle$ defined as follows (see Figure 4). The state space of \mathcal{D}_n consists of a state q_{pass} and two copies of subsets of $[n]$: the a -copy $2^{[n]} \times \{a\}$, with “a” standing for “accumulate”, and the c -copy $2^{[n]} \times \{c\}$, with “c” standing for “check”. Essentially, the state $\langle S, a \rangle$ in the a -copy remembers that we have read only numbers in $[n]$ after the last $\$$ and also remembers the set $S \subseteq [n]$ of these numbers. The state $\langle S, c \rangle$ in the c -copy remembers, in addition, that we have just read $\#$. Accordingly, the state $\langle S, c \rangle$, for a nonempty set S , checks whether the next letter is in S and moves to the state q_{pass} via a transition in $\bar{\alpha}$ only when the check passes. The fact that the above transitions are in $\bar{\alpha}$, and we can move from q_{pass} back to $\langle \emptyset, a \rangle$ upon reading $\$$ and without traversing α , imply that q_{pass} can trap runs that do not traverse α on words of the form $z_1 \cdot z_2 \cdots$, where for all $l \geq 1$ z_l is good. Formally, for every state $\langle S, o \rangle \in 2^{[n]} \times \{a, c\}$, and letter $\sigma \in \Sigma_n \cup \{\#\}$, we define:

$$\delta(\langle S, o \rangle, \sigma) = \begin{cases} \langle S \cup \{\sigma\}, a \rangle, & \text{if } o = a \text{ and } \sigma \in [n] \\ \langle S, c \rangle, & \text{if } o = a, S \neq \emptyset \text{ and } \sigma = \# \\ q_{pass}, & \text{if } o = c \text{ and } \sigma \in S \\ \langle \emptyset, c \rangle, & \text{if } \langle S, o \rangle = \langle \emptyset, c \rangle \text{ and } \sigma = \neg \$ \\ \langle \emptyset, c \rangle, & \text{if } o = c \text{ and } \sigma \in ([n] \cup \{\#\}) \setminus S \\ \langle \emptyset, a \rangle, & \text{if } \sigma = \$ \\ \langle \emptyset, c \rangle, & \text{if } \langle S, o \rangle = \langle \emptyset, a \rangle \text{ and } \sigma = \# \end{cases}$$

where all types transitions are in $\bar{\alpha}$ except for the last three. In addition, by reading a letter σ from q_{pass} we move to $\langle \emptyset, a \rangle$ via a transition in $\bar{\alpha}$ when $\sigma = \$$, and otherwise, we move to $\langle \emptyset, c \rangle$ via an α -transition. Note that by reading $\$$ from any state we move to $\langle \emptyset, a \rangle$ to detect a suffix of the form $z_1 \cdot z_2 \cdots$, where for all $l \geq 1$ z_l is good, and if at some point we read an unexpected letter, we move to $\langle \emptyset, c \rangle$, where we wait for $\$$ to be read.

It is easy to see that \mathcal{D}_n recognizes L_n and has $2^{n+1} + 1$ states.

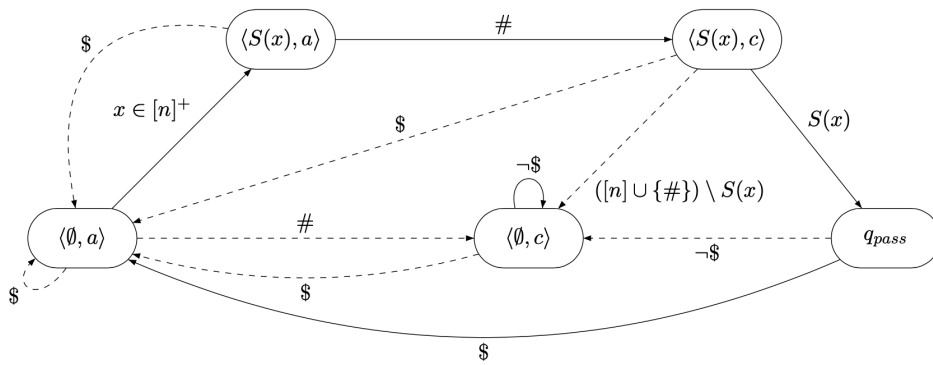


Figure 4: The tDCW D_n . For a finite word $x \in [n]^+$, we denote the set of numbers that appear in x by $S(x)$. Dashed transitions are α -transitions.