Recognizing ω-regular Languages with Probabilistic Automata

Christel Baier, Marcus Größer⋆,
Institut für Informatik I, Universität Bonn
Römerstrasse 164, D-53117 Bonn, Germany
baier|groesser@cs.uni-bonn.de

Abstract

Probabilistic finite automata as acceptors for languages over finite words have been studied by many researchers. In this paper, we show how probabilistic automata can serve as acceptors for ω-regular languages. Our main results are that our variant of probabilistic Büchi automata (PBA) is more expressive than non-deterministic ω-automata, but a certain subclass of PBA, called uniform PBA, has exactly the power of ω-regular languages. This also holds for probabilistic ω-automata with Streett or Rabin acceptance. We show that certain ω-regular languages have uniform PBA of linear size, while any nondeterministic Streett automaton is of exponential size, and vice versa. Finally, we discuss the emptiness problem for uniform PBA and the use of PBA for the verification of Markov chains against qualitative linear-time properties.

1 Introduction

Automata as acceptors for infinite words play a crucial role in logic, for verification purposes and other areas, see e.g. [33, 29, 19]. Several types of automata for languages over infinite words have been studied in the literature. They can be classified via their branching structure (deterministic, nondeterministic, universal, alternating) and acceptance condition (Büchi, Müller, Rabin, Streett, etc.).

In this paper, we introduce a probabilistic variant of ω-automata that serve as acceptors for languages over infinite words. Although probabilistic finite automata (PFA) have attracted many researchers, see e.g. [26, 25, 17, 8, 15, 5], we are not aware of any paper that discusses probabilism for the recognition of languages of infinite words.

As a first step towards probabilistic ω-acceptors we concentrate on some fundamental aspects such as expressiveness, efficiency and the emptiness-problem. Potential application areas might be any research topic where ω-regular languages are of importance, such as the verification of reactive systems [31] or reasoning with biological processes [18]. Given the wide range of applications of PFA (e.g., for speech recognition [27], Arthur-Merlin games [3], planning questions in Markov decision processes [23, 6], or prediction of climatic parameters [24]), probabilistic ω-automata might be of interest also for other areas. In addition, probabilistic ω-automata could serve as basis for "quantum ω-automata" (in analogy to quantum finite automata [21, 2] which can be regarded as an extension of PFA) or they might be useful in combination with costs as in [13, 14] where ω-variants of weighted automata are studied.

In this paper we show that probabilistic automata with Büchi acceptance *"the set of final states has to be visited infinitely often"* and the acceptance criteria *"the accepting runs have a positive probability measure"* are more expressive than nondeterministic ω-automata. This stands in contrast to the facts that (1) deterministic Büchi automata do not have the full power of ω-regular languages as their nondeterministic counterpart and (2) PFA with the acceptance criteria "the accepting runs have a positive probability measure" can be viewed as nondeterministic finite automata, and hence, have exactly the power of regular languages. However, a certain subclass of probabilistic Büchi automata (PBA), called uniform PBA, covers exactly the class of ω-regular languages. Dealing with acceptance thresholds in $]0,1[$ (as it is standard for PFA), the class of accepted languages by uniform PBA still subsumes the ω-regular languages, but also contains non-ω-regular languages.

Regarding the efficiency, uniform PBA are not comparable with nondeterministic ω-automata. On the one hand, there exists a family of ω-regular languages that can be recognized by uniform PBA of linear size, while even the smallest nondeterministic Streett automaton for them has at least $\Omega(2^n/n)$ states. On the other hand there are ω-regular languages with nondeterministic Büchi automata of linear size while any PBA has at most $\Omega(2^n)$ states. However, uniform PBA are at least as efficient as nondeterministic Büchi automata that are deterministic in limit [32, 9]. Moreover,

⋆ Supported by the DFG-Project "VERIAM" and the DFG-NWO-Project "VOSS".

there are polynomial transformations from uniform probabilistic Rabin automata to uniform PBA. In particular, up to polynomial transformations uniform PBA are no worse than deterministic Rabin automata. Surprisingly, for probabilistic Streett automata a polynomial transformation to (possibly non-uniform) PBA can be provided, while in the nondeterministic case the switch from Streett to Büchi acceptance can cause an exponential blowup [28].

Checking emptiness for uniform PBA is computationally harder than for nondeterministic ω-automata (NP-hard and solvable in exponential time), while intersection and union can be realized as in the nondeterministic case. We finally show that the qualitative analysis of Markov chains against automata-specifications can be realized with PBA in a quite simple way.

Organization of the paper. Section 2 briefly recalls the basic definition of (non)deterministic ω-automata. Probabilistic Büchi automata are introduced in Section 3. Section 4 discusses the power of PBA, while probabilistic automata with other acceptence conditions are considered in Section 5. The efficiency of PBA is studied in Section 6. Section 7 is concerned with operators on PBA. Section 8 briefly sketches how PBA can serve for verification purposes. The paper ends with a brief conclusion in Section 9.

2 Ordinary ω-automata

Throughout the paper, we assume some familiarity with formal languages, finite automata and ω-automata. We briefly recall the basic concepts and explain our notations concerning non-deterministic ω-automata with the Büchi, Rabin and Streett acceptance criterion. For further details see e.g. [29, 19].

Throughout the paper, Σ denotes a nonempty finite alphabet. Letters $a,b,c,\dots$ will be used to denote the elements of Σ. Σ^ω denotes the set of infinite words over Σ, while Σ^* stands for the set of finite words over Σ. A *nondeterministic* ω*-automaton* over Σ is a tuple $\mathcal{A} = (Q,\delta,Q_0,Acc)$ where Q is a finite state space, $\delta : Q\times\Sigma \to 2^Q$ the transition function and $Q_0 \subseteq Q$ the set of initial states. The last component Acc denotes the acceptance condition of $\mathcal{A}$. Its type depends on the type of ω-automata. For Büchi automata, Acc is a set of accepting states, that is, $Acc = F$ for some $F \subseteq Q$. For Rabin or Streett automata, Acc is a set $\{(H_1,K_1),\dots,(H_m,K_m)\}$ of pairs (H_i,K_i) consisting of sets $H_i,K_i \subseteq Q$. $|\mathcal{A}|$ denotes the number of states in $\mathcal{A}$ (i.e., $|\mathcal{A}| = |Q|$). $\mathcal{A}$ is called deterministic if $|Q_0| = 1$ and $|\delta(q,a)| \le 1$ for all $q \in Q$ and $a \in \Sigma$. We write NBA, NRA, NSA, DBA, DRA and DSA to denote the nondeterministic or deterministic version of Büchi, Rabin or Streett automata, respectively.

Given an infinite word $\rho = a_1a_2\dots$ over Σ a run for ρ in $\mathcal{A}$ denotes any finite or infinite state-sequence $\pi = q_0,q_1,\dots$ where $q_0 \in Q_0$ and $q_i \in \delta(q_{i-1},a_i)$, $i = 1,2,\dots$ and such that π is either infinite or $\pi = q_0,\dots,q_n$ where $\delta(q_n,a_{n+1}) = \emptyset$. We write $\inf(\pi)$ to denote the set of states that occur infinitely often in π. (We have $\inf(\pi) = \emptyset$ iff π is finite.) An infinite run π is called accepting with respect to the Büchi acceptance condition F if F is visited infinitely often in π, i.e., if $\inf(\pi)\cap F \neq \emptyset$. For the Rabin acceptance condition $\{(H_1,K_1),\dots,(H_m,K_m)\}$, π is called accepting if there exists an index $i \in \{1,\dots,m\}$ such that $\inf(\pi) \subseteq H_i$ and $\inf(\pi)\cap K_i \neq \emptyset$. For the Streett acceptance condition $\{(H_1,K_1),\dots,(H_m,K_m)\}$, π is called accepting if for all indices $i \in \{1,\dots,m\}$ we have either $\inf(\pi)\cap H_i = \emptyset$ or $\inf(\pi)\cap K_i \neq \emptyset$. Any finite run is non-accepting. Using LTL notation, Büchi acceptance can be described by the formula $\Box\Diamond F$ (infinitely often F), Rabin acceptance by the formula

$$\bigvee_{1\le i\le m} \Diamond\Box(H_i \wedge \Diamond K_i)$$

stating that, for at least one acceptance pair (H_i,K_i), almost all states belong to H_i and infinitely many states belong to K_i, while Streett acceptance can be read as strong fairness:

$$\bigwedge_{1\le i\le m} (\Box\Diamond H_i \;\to\; \Box\Diamond K_i).$$

The accepted language $\mathcal{L}(\mathcal{A})$ of a NBA (DBA, NRA, DRA, NSA, DSA) $\mathcal{A}$ is the set of all infinite words $\sigma \in \Sigma^\omega$ that have an accepting run in $\mathcal{A}$. It is well known that the classes of languages accepted by a NBA, NRA, DRA, NSA and DSA agree exactly with the class of ω-regular languages, while DBA are strictly less expressive.

We often identify any ω-regular language $L \subseteq \Sigma^\omega$ with the ω-regular expressions that describe L. E.g., $(a+b)^*a^\omega$ is identified with the set of infinite words over $\Sigma = \{a,b\}$ that contain only finitely many b's.

For acceptors of languages over *finite* words, we use the abbreviation PFA for probabilistic finite automata, while NFA and DFA stands for nondeterministic or deterministic finite automata, respectively.

3 Probabilistic Büchi automata

We now introduce probabilistic Büchi automata which can be viewed as NBA where the nondeterminism is resolved by a probabilistic choice. That is, for any state q and letter $a \in \Sigma$ either q does not have any a-successor or there is a probability distribution for the a-successors of q.

Definition 1 (Probabilistic Büchi automata (PBA)). *A PBA over the alphabet Σ is a tuple $\mathcal{P} = (Q,\delta,\mu,F)$ where Q is a finite state space, $\delta : Q\times\Sigma\times Q \to [0,1]$ the transition probability function such that for all $q \in Q$ and $a \in \Sigma$,*

$$\sum_{p\in Q} \delta(q,a,p) \in \{0,1\},$$

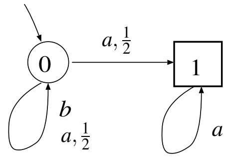

Fig. 1. **PBA for** $(a+b)^*a^\omega$

μ the initial distribution (i.e., μ is a function $Q \to [0,1]$ with $\sum_{q \in Q} \mu(q) = 1$) and $F \subseteq Q$ the set of accepting states. The states $q \in Q$ where $\mu(q) > 0$ are called initial.

The operational behaviour of $\mathcal{P}$ for a given input word $\rho = a_1 a_2 \ldots \in \Sigma^\omega$ is as follows. The automaton chooses at random an initial state q_0 according to the initial distribution μ. After having consumed the first i input symbols $a_1, \ldots, a_i$, $\mathcal{P}$ in state q_i moves with probability $\delta(q_i, a_{i+1}, p)$ to state $p = q_{i+1}$ and tries to read the next input symbol a_{i+2} from p. If there is no outgoing a_{i+1}-transition from the current state q_i, i.e., if $\sum_{p \in Q} \delta(q_i, a_{i+1}, p) = 0$, then $\mathcal{P}$ rejects. As for NBA, the resulting (finite or infinite) state-sequence $q_0, q_1, \ldots$ is called a run for ρ in $\mathcal{P}$. An accepting run in $\mathcal{P}$ denotes an infinite run π with $\inf(\pi) \cap F \neq \emptyset$.
Thus, for the input word $\rho = a_1 a_2 \ldots \in \Sigma^\omega$, $\mathcal{P}$'s behavior can be formalized by a (possibly infinite-state) *Markov chain* that arises by unfolding $\mathcal{P}$ into a tree[1] following the letters $a_1, a_2, \ldots$ as long as possible. This Markov chain might have terminal states, namely if there is a run $q_0, q_1, \ldots, q_n$ in $\mathcal{P}$ for a prefix $a_1 \ldots a_n$ of ρ such that q_n does not have an outgoing a_{n+1}-transition.

Definition 2 (Accepted language of a PBA). *Using the standard probability measure on Markov chains (see e.g. [20, 22]), we define the acceptance probability* $\Pr_{\mathcal{P}}(\rho)$*, or briefly* $\Pr(\rho)$*, for* ρ *in* $\mathcal{P}$ *by*

$$\Pr(\rho) \;=\; \Pr\{\pi : \pi \textit{ is an accepting run for } \rho\}.$$

The accepted language $\mathcal{L}(\mathcal{P})$ *consists of all words* $\rho \in \Sigma^\omega$ *where* $\Pr(\rho) > 0$.

Intuitively, $\Pr(\rho)$ denotes the probability for the event "infinitely often F" under the scheduling policy induced by ρ. By the results of [30, 9] the set of accepting runs for ρ is measurable. In the pictures of PBA, we use boxes to denote the accepting states and circles for the non-accepting states. We simply write a as label for a transition from q to p if $\delta(q,a,p) = 1$. Label a, x with $x \in]0,1[$ for a transition from q to p denotes that $\delta(q,a,p) = x$.

Example 3. Fig. 1 shows a PBA that accepts the language $(a+b)^*a^\omega$. To see this, we first notice that only the words in $(a+b)^*a^\omega$ have an accepting run, because the a-labelled self-loop at the accepting state 1 is the only outgoing transition of state 1. On the other hand, $\Pr(a^\omega) = 1$ (as the non-accepting run $0,0,0,\ldots$ has probability 0 while all other runs for a^ω are accepting). For any word $\rho \in (a+b)^n ba^\omega$, the acceptance probability is at least $1/2^n$. E.g., for $aaba^\omega$, there are three non-accepting runs: 0^ω and the finite runs $0,0,1$ and $0,1,1$, while all other runs $0,0,0,\ldots,1,1,\ldots$ for $aaba^\omega$ are accepting. The probability for the non-accepting runs is $0 + \frac{1}{4} + \frac{1}{2} = \frac{3}{4}$. Thus, $\Pr(\rho) = \frac{1}{4}$.
The PBA shown in Fig. 2 accepts the language $(ab + ac)^*(ab)^\omega$. The intuitive argument why any word ρ in $(ab+ac)^\omega$ with infinitely many c's is rejected relies on the observation that almost all runs for ρ are finite and end in state 1 (where the next input symbol is c that cannot be consumed in state 1).[2] □

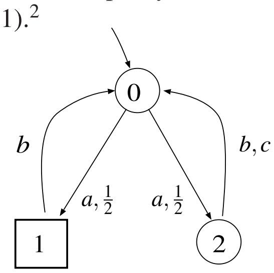

Fig. 2. **PBA for** $(ab+ac)^*(ab)^\omega$

Clearly, any DBA can be viewed as a PBA (where $\delta_{PBA}(q,a,p) = 1$ if $\delta_{DBA}(q,a) = \{p\}$). On the other hand, it is well known that the language $(a+b)^*a^\omega$ cannot be described by a DBA, while there is a PBA for $(a+b)^*a^\omega$ (see Fig. 1). Thus, PBA are strictly more expressive than DBA. It is worth mentioning that the qualitative criteria "accepting runs have positive probability" is different from the acceptance criteria "there is an accepting run" in the context of languages of *infinite* words, while they agree for probabilistic automata viewed as acceptors for *finite* words. In fact, the naive transformation from PBA to NBA which relies on ignoring the probabilities can fail to yield an equivalent NBA. E.g., for the PBA in Fig. 2 the underlying NBA that we obtain by ignoring the probabilities accepts the language $\big((ac)^*ab\big)^\omega$. This example also demonstrates that replacing the nondeterministic choices in a given NBA by probabilistic choices might yield a non-equivalent PBA. However, the accepted language of a PBA is included in the language that is accepted by the naive associated NBA, i.e. the NBA that stems from the given PBA by ignoring the probabilities.
Another example is the NBA in Fig. 3 for the language $(ab + ac)^\omega$. If we attach non-zero probabilities to both a-transitions then the resulting PBA accepts the empty language. (Note that any infinite word in $(ab+ac)^\omega$ has exactly one accepting run, but its probability is 0.)

[1] More precisely, for PBA with two or more initial states we obtain a forest rather than a single tree.

[2] The formulation "almost all runs have property x" means that the probability measure of the runs where property x does not hold is 0.

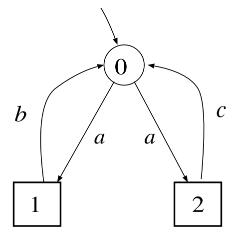

Fig. 3. **NBA for $(ab+ac)^\omega$, while PBA accepts $\emptyset$**

4 Expressiveness of PBA

We first establish the result stating that PBA are more expressive than nondeterministic ω-automata:

Theorem 4. *The class of languages that can be accepted by a PBA strictly contains the class of ω-regular languages.*

The proof of Theorem 4 splits into two parts. In Lemma 5, we show that for any NBA $\mathcal{A}$ there exists a PBA $\mathcal{P}$ such that $\mathcal{L}(\mathcal{P}) = \mathcal{L}(\mathcal{A})$. Second, we provide an example of a PBA for which the accepted language is not ω-regular.
Following [32, 9], we call a NBA $\mathcal{A}$ to be *deterministic in limit* if $|\delta(p,a)| \leq 1$ for any state p that is reachable from an accepting state $q \in F$ and any symbol $a \in \Sigma$. If we regard a NBA $\mathcal{A}$ that is deterministic in limit as a PBA $\mathcal{P}$ (with arbitrary probability distributions to resolve the nondeterministic choices) then $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{P})$. E.g., ignoring the probabilities in the PBA in Fig. 1 yields an equivalent NBA that is deterministic in limit, while the NBA in Fig. 3 is not deterministic in limit. [9] provided a transformation from a given NBA $\mathcal{A}$ into an equivalent NBA that is deterministic in limit and whose size is (single) exponential in $|\mathcal{A}|$. This yields the proof of the following lemma:

Lemma 5 (From NBA to PBA). *For any NBA $\mathcal{A}$ there is a PBA $\mathcal{P}$ such that $\mathcal{L}(\mathcal{P}) = \mathcal{L}(\mathcal{A})$ and $|\mathcal{P}| = O(\exp(|\mathcal{A}|))$.*

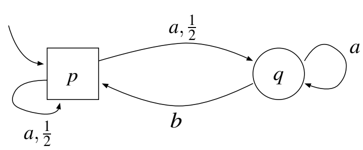

Fig. 4. **PBA that accepts a non-ω-regular language**

It remains to provide an example for a PBA that recognizes a language that is not ω-regular. For PBA $\mathcal{P}$ in Fig. 4 we have $\mathcal{L}(\mathcal{P}) = L$ where L is

$$\left\{a^{k_1}ba^{k_2}ba^{k_3}b\ldots : k_1,k_2,\ldots \in \mathbb{N}_{\geq 1} \text{ s.t. } \prod_{i=1}^{\infty}(1-(\tfrac{1}{2})^{k_i}) > 0\right\}.$$

Note that $\mathcal{L}(\mathcal{P}) \subseteq (a^+b)^\omega$ as every accepting run for an infinite word ρ that has only finitely many b's has to stay in state p from some point on. But such runs have probability 0. Moreover, $\Pr(a^{k_1}ba^{k_2}ba^{k_3}b\ldots) = \prod_{i=1}^{\infty}(1-(\frac{1}{2})^{k_i})$. This yields $\mathcal{L}(\mathcal{P}) = L$.
We now show that $\mathcal{L}(\mathcal{P})$ is not ω-regular. Otherwise, there exists a NBA $\mathcal{A}$ with $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{P})$. Let $q_0 \xrightarrow{a_1} \ldots \xrightarrow{a_n} q_n = q_0$ be a reachable accepting cycle of $\mathcal{A}$. (An accepting cycle means a cycle in $\mathcal{A}$ that contains at least one accepting state.) Then, there is at least one index $0 \leq i < n$ such that $a_i = b$. Hence, $a_0a_1\ldots a_{n-1}$ is a word of the form $a^{j_1}ba^{j_2}b\ldots a^{j_k}b$. But then, $\mathcal{A}$ accepts a word of the form $(a^+b)^*(a^{j_1}ba^{j_2}b\ldots a^{j_k}b)^\omega$, which contradicts the assumption $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{P})$ as no such word is accepted by $\mathcal{P}$. □

Although PBA are in general more expressive than ω-regular languages, we will now classify a subset of PBA, so called *uniform* PBA, which is as expressive as ω-regular languages; see Theorem 12 below.

Notation 6. Let $\mathcal{P} = (Q,\delta,\mu,F)$ be a PBA. $|\mathcal{P}| = |Q|$ denotes the number of states in $\mathcal{P}$. $\delta(q,a)$ stands for the set of states $p \in Q$ where $\delta(q,a,p) > 0$. Similarly, if $w = a_1\ldots a_k \in \Sigma^*$ then $\delta(q,w)$ denotes the set of states $p \in Q$ that are reachable from q via the word w with a nonzero probability. For $R \subseteq Q$, we put $\delta(R,w) = \bigcup_{r\in R}\delta(r,w)$.
If $\pi = q_0,q_1,q_2,\ldots$ is an infinite run for a word $\rho = a_1,a_2,\ldots \in \Sigma^\omega$ then $\mathsf{Lim}(\pi,\rho)$ denotes the pair (P,A) where $P = \inf(\pi)$ and $A : P \to 2^\Sigma$ is the function that assigns to any state $p \in P$ the set $A(p)$ of symbols $a \in \Sigma$ such that $(q_i = p) \wedge (a_{i+1} = a)$ for infinitely many i. If ρ is understood from the context then we briefly write $\mathsf{Lim}(\pi)$. □

In the sequel, we shall use the concepts of de Alfaro's end components [11, 12] that have been originally defined for Markov decision processes:

Definition 7 (End components, cf. [11,12]). *Let $\mathcal{P} = (Q,\delta,\mu,F)$ be a PBA. An end component of $\mathcal{P}$ is a pair (P,A) where $P \subseteq Q$ and $A : P \to 2^\Sigma$ is a function such that*

- $\sum_{q\in P}\delta(p,a,q) = 1$ *for any $p \in P$ and $a \in A(p)$,*
- *the underlying digraph of (P,A) is strongly connected.*

(P,A) is called accepting if $P \cap F \neq \emptyset$. In this case, $\Pr(\rho,(P,A))$ *stands for the probability of all accepting runs π for ρ with* $\mathsf{Lim}(\pi) = (P,A)$.

For each infinite word ρ, for almost all infinite runs π for ρ, $\mathsf{Lim}(\pi)$ is an end component [11, 12]. We obtain:

Lemma 8 (AEC-Lemma). *For any PBA $\mathcal{P}$ and $\rho \in \Sigma^\omega$, $\rho \in \mathcal{L}(\mathcal{P})$ iff $\Pr(\rho,(P,A)) > 0$ for some accepting end component (P,A).*

Definition 9 (Prob-fairness). *If π is an infinite run for $\rho = a_1a_2\ldots$ then π is called prob-fair if $\forall(p,a,q) \in Q\times\Sigma\times Q$*

$$\delta(p,a,q) > 0 \wedge \overset{\infty}{\exists}\, i \geq 1.(p,a) = (q_{i-1},a_i)$$
$$\Longrightarrow \overset{\infty}{\exists}\, i \geq 1 \text{ with } (p,a,q) = (q_{i-1},a_i,q_i).$$

Here, $\overset{\infty}{\exists}$ means "there exists infinitely many".

Prob-fairness can be understood as strong fairness for the probabilistic choices. By the results of [4], it can be derived that almost all runs are either finite or prob-fair. Thus $\Pr(\rho)$ is the probability of all prob-fair accepting runs for ρ.

Notation 10. We often consider the runs for a given (infinite or finite) word $\rho = a_1 a_2 \ldots \in \Sigma^* \cup \Sigma^\omega$ as labeled paths

$$\pi = q_0 \xrightarrow{a_1} q_1 \xrightarrow{a_2} q_2 \xrightarrow{a_3} \ldots$$

If (P,A) is an end component then $q_0 \xrightarrow{b_1} \ldots \xrightarrow{b_k} q_k$ is said to be a path fragment in (P,A) if $q_i \in P$ and $b_{i+1} \in A(q_i)$ for all $i \in \{0,\ldots,k-1\}$. We write

$$\Pr(q_0 \xrightarrow{b_1} \ldots \xrightarrow{b_k} q_k) = \prod_{0 \le i < k} \delta(q_i, b_{i+1}, q_{i+1})$$

to denote the probability of taking the run $q_0, q_1, \ldots, q_k$ for the word $b_1, \ldots, b_k$ when starting in state q_0. If $b_1 \ldots b_k$ are clear from the context then we simply write $\Pr(q_0, \ldots, q_k)$. If (P,A) is an end component then $\Pr_{(P,A)}(q \xrightarrow{b_1,\ldots,b_k} p)$ denotes the sum of the probabilities for all path fragments in (P,A) from q to p for the word $b_1 \ldots b_k$.

$$\Pr_{(P,A)}(q \xrightarrow{b_1,\ldots,b_k}) \; = \; \sum_{p \in P} \Pr_{(P,A)}(q \xrightarrow{b_1,\ldots,b_k} p)$$

denotes the probability for the runs in (P,A) for the word $b_1, \ldots, b_k$ when starting in state q. □

Definition 11 (Uniform PBA). *PBA $\mathcal{P}$ is called uniform if there exists $\theta > 0$ such that for all accepting end components (P,A), all states $p \in P$ and all finite words $w \in \Sigma^*$ one of the following conditions (i) or (ii) holds:*

(i) $\Pr_{(P,A)}(p \xrightarrow{w}) \le 1 - \theta$

(ii) $\Pr_{(P,A)}(p \xrightarrow{w}) = 1 \;\wedge\; \forall\, q \in \delta(p,w).\ \Pr_{(P,A)}(p \xrightarrow{w} q) \ge \theta$

The PBA shown in Fig. 1, 2 and 3 are uniform, while the PBA in Fig. 4 is not as we have:

$$\Pr(p \xrightarrow{a^k b}) = 1 - (\tfrac{1}{2})^k \text{ converges to 1 if } k \text{ tends to } \infty.$$

A more complex example for a uniform PBA will be discussed in Section 6. Any deterministic-in-limit NBA, viewed as a PBA, is uniform, since its accepting end components are deterministic. Thus, by the results of [9], for any ω-regular language L there exists a uniform PBA that accepts L. In Lemma 16, we will provide a transformation from uniform PBA to NSA. Thus:

Theorem 12 (Expressiveness of uPBA). *The class of accepted languages of uniform PBA agrees with the class of ω-regular languages.*

The remainder of this section is concerned with the transformation from uniform PBA to NSA. We start with a discussion of the acceptance structure of uniform PBA.

Definition 13 (Strong prob-fairness). *Let $\pi = q_0, q_1, \ldots$ be an infinite run for $\rho = a_1 a_2 \ldots \in \Sigma^\omega$ in a PBA $\mathcal{P}$. π is called strongly prob-fair if the following condition holds. For all runs $\pi' = q_0', q_1', \ldots$ for ρ and for all $p \in Q$ and $R', R \subseteq Q$ such that $p \in R' \subseteq R$ we have:*

$$\overset{\infty}{\exists} j \ge 0 \exists k > j \text{ s.t. } R = \delta(q_j, a_{j+1} \ldots a_k) \wedge R' = \delta(q_j', a_{j+1} \ldots a_k)$$
$$\Downarrow$$
$$\overset{\infty}{\exists} j \ge 0 \exists k > j \text{ s.t.}$$
$$q_k = p \wedge R = \delta(q_j, a_{j+1} \ldots a_k) \wedge R' = \delta(q_j', a_{j+1} \ldots a_k)$$

Note that with $\pi = \pi'$, $R = R' = \delta(p,a)$ for $p \in P$, $a \in A(p)$ where $\mathsf{Lim}(\pi) = (P,A)$ we get that strongly prob-fair runs are prob-fair.

Lemma 14. *Let $\mathcal{P}$ be a uniform PBA and $\rho \in \mathcal{L}(\mathcal{P})$. Then, almost all accepting runs for ρ are strongly prob-fair.*

We now show that in uniform PBA the existence of a certain accepting run for a word ρ is equivalent to the acceptance of ρ in $\mathcal{P}$.

Lemma 15 (Acceptance in uPBA). *Let $\mathcal{P}$ be a uniform PBA and $\rho = a_1 a_2 \ldots \in \Sigma^\omega$. Then $\rho \in \mathcal{L}(\mathcal{P})$ if and only if there exists an accepting end component (P,A) and a strongly prob-fair run $\pi = q_0, q_1, \ldots$ for ρ and an index $\ell \ge 0$ such that $\mathsf{Lim}(\pi) = (P,A)$, $q_\ell \in P$ and*

$$\Pr_{(P,A)}(q_\ell \xrightarrow{a_{\ell+1} \ldots a_{\ell+j}}) = 1 \textit{ for all } j \ge 1.^3$$

Lemma 16 (From uPBA to NSA). *For any uniform PBA $\mathcal{P}$ there is a NSA $\mathcal{A}$ with $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{P})$ and $|\mathcal{A}| = O(\exp(|\mathcal{P}|))$.*[4]

Proof. Given a uniform PBA $\mathcal{P} = (Q, \delta, \mu, F)$ we first transform $\mathcal{P}$ into an equivalent PBA $\mathcal{P}'$ which is in some stronger sense uniform and which we then transform into an equivalent NSA.

PBA $\mathcal{P}'$. The state space of $\mathcal{P}'$ is $Q' = Q \cup (Q \times 2^Q \times 2^Q)$. The initial distribution in $\mathcal{P}'$ is $\mu' = \mu$. Within the states Q, the PBA $\mathcal{P}'$ imitates $\mathcal{P}$, always having the option to jump to the powerset part with a positive probability. Within the powerset construction, given a state $\langle p, R, R' \rangle$ in Q', the automaton essentially behaves as p, but remembers the set R of states it has "to synchronize with". The third component R' is needed for the acceptance condition in the NSA and serves to check whether all states in R are visited infinitely often. Formally, we choose the transition probability function δ' as follows.

[3] Note that $\Pr_{(P,A)}(q_\ell \xrightarrow{a_{\ell+1} \ldots a_{\ell+j}}) = 1$ iff for any path fragment $q_\ell \xrightarrow{a_{\ell+1} \ldots a_{\ell+i}} p$ in (P,A) where $1 \le i < j$ we have $a_{\ell+i+1} \in A(p)$.

[4] Here, we assume the alphabet Σ to be fixed. Otherwise, the bound $|\mathcal{A}| = O(\exp(|\mathcal{P}|) \cdot |\Sigma|)$ has to be used.

$$\delta'(q,a,p) = \delta'(q,a,\langle p,\{p\},\{p\}\rangle) = \tfrac{1}{2}\cdot\delta(q,a,p)$$

and for $R,R' \subseteq Q$ such that $\delta(r,a) \neq \emptyset$ for all $r \in R$,

$$\delta'(\langle q,R,R'\rangle,a,\langle p,S,S'\rangle) \;=\; \delta'(\langle q,R,R'\rangle,a,\langle p,S,\{p\}\rangle)$$

$$= \tfrac{1}{2}\delta(q,a,p) \quad \text{if } S = \delta(R,a),\, S' = \delta(R',a)$$

In the above formula, we assume $S' \neq \{p\}$. For $S' = \{p\}$, we put $\delta'(\langle q,R,R'\rangle,a,\langle p,S,\{p\}\rangle) = \delta(q,a,p)$. In all remaining cases, we define $\delta'(\cdot) = 0$. The acceptance set F' of $\mathcal{P}'$ covers the states $\langle p,R,R'\rangle$ of the powerset construction whose first component p is an accepting state of $\mathcal{P}$ and where the synchronization set R agrees with R'. That is,

$$F' = \big\{\langle p,R,R\rangle : p \in F, R \in 2^Q\big\}.$$

An inspection of the transition relation yields the uniformity of $\mathcal{P}'$ (as $\mathcal{P}$ is uniform). It can be shown that $\mathcal{L}(\mathcal{P}) = \mathcal{L}(\mathcal{P}')$.
NSA $\mathcal{A}$. The states in the NSA $\mathcal{A}$ have the form (q',a) where $q' \in Q' = Q \cup (Q \times 2^Q \times 2^Q)$ is a state in $\mathcal{P}'$ and $a \in \Sigma$. Intuitively, (q',a) is a copy of state q' where only a-transitions can be taken and where the probabilistic choices according to a are viewed as nondeterministic choices. Formally, $\mathcal{A}' = (Q', \delta'_{\mathcal{A}}, \mu', Acc)$ where the transition relation in $\mathcal{A}$ is given by $\delta'_{\mathcal{A}}((q',a),b) = \emptyset$ if $a \neq b$ and

$$\delta'_{\mathcal{A}}((q',a),a) = \big\{(p',b) : p' \in \delta'(q',a), b \in \Sigma\big\}.$$

The Streett acceptance condition in $\mathcal{A}$ is given by:

$$\underbrace{\Box\Diamond \bigvee_{\substack{a\in\Sigma\\ p\in F}} (\langle p,R,R\rangle,a)}_{\text{unconditional fairness for } F'} \;\wedge \qquad \text{(I)}$$

$$\bigwedge_{\substack{q\in R'\subseteq R\\ a\in\Sigma}} \big(\Box\Diamond(\langle q,R,R'\rangle,a) \rightarrow \Box\Diamond(\langle q,R,\{q\}\rangle,a)\big) \;\wedge \qquad \text{(II)}$$

$$\bigwedge_{\substack{q,p\in R'\subseteq R\\ a\in\Sigma}} \big(\Box\Diamond(\langle q,R,R'\rangle,a) \rightarrow \Box\Diamond(\langle p,R,R'\rangle,a)\big) \qquad \text{(III)}$$

We now show that $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{P})$. If $\rho \in \mathcal{L}(\mathcal{P}) = \mathcal{L}(\mathcal{P}')$ then there exists a strongly prob-fair, accepting run $\pi = q'_0, q'_1, \ldots$ for $\rho = a_1a_2\ldots$ in $\mathcal{P}'$ (Lemma 14). But then,

$$\pi_{\mathcal{A}} = (q'_0,a_1),(q'_1,a_2),\ldots$$

is an accepting run in $\mathcal{A}$, and therefore, $\rho \in \mathcal{L}(\mathcal{A})$. Note that (I) holds as π is accepting, (II) holds as π is prob-fair (cf. (+)) and (III) follows from the condition in Def. 13.
Vice versa, if $\rho = a_1a_2,\ldots \in \mathcal{L}(\mathcal{A})$ and $\pi_{\mathcal{A}}$ is an accepting run for ρ in $\mathcal{A}$ then $\pi_{\mathcal{A}}$ has the form $(q_0,a_1),\ldots,(q_{j-1},a_j)$, $(\langle q_j,R_j,R'_j\rangle,a_{j+1})$, $(\langle q_{j+1},R_{j+1},R'_{j+1}\rangle,a_{j+2}),\ldots$ where $q_0,q_1,\ldots \in Q$ are states in $\mathcal{P}$, $R_i = \delta(R_{i-1},a_i)$ and $R'_i = \{q_i\}$ or $R'_i = \delta(R'_{i-1},a_i)$ where $R_{j-1} = R'_{j-1} = \{q_{j-1}\}$. Let

$$\pi = q_0,\ldots,q_j,\langle q_j,R_j,R'_j\rangle,\langle q_{j+1},R_{j+1},R'_{j+1}\rangle,\ldots$$

be the induced run in $\mathcal{P}'$ and let $\mathsf{Lim}(\pi) = (P',A')$. The acceptance condition in $\mathcal{A}$ ensures that (1) π visits infinitely often states in $F' = \{\langle p,R,R\rangle : p \in F\}$ (because of (I)) and (2) whenever $\langle q,R,R'\rangle \in P'$ and $a \in A'(\langle q,R,R'\rangle)$ then (because of (II) and (III)) for all states $p \in R'$:

- $\langle p,R,\{p\}\rangle \in P'$ and $\langle p,R,R'\rangle \in P'$,
- $a \in A'(\langle p,R,\{p\}\rangle) \cap A'(\langle p,R,R'\rangle)$

π needs not to be prob-fair as we may have $\langle q,R,R'\rangle \in P'$, $a \in A'(\langle q,R,R'\rangle)$ and $\langle p,S,\{p\}\rangle \in P'$ for all $p \in S = \delta(R,a)$ but $\langle p,S,S'\rangle \notin P'$ for some $p \in S' = \delta(R',a)$. However, we may modify π (by skipping certain "resets" of the third component) to obtain a strongly prob-fair run π' for ρ in $\mathcal{P}'$. Note that if

$$\langle q,R,R'\rangle \xrightarrow{a} \langle p,S,\{p\}\rangle \xrightarrow{w} \langle r,R,R\rangle$$

then $\langle q,R,R'\rangle \xrightarrow{a} \langle p,S,S'\rangle \xrightarrow{w} \langle r,R,R\rangle$ where $S' = \delta(R',a)$. The resulting run fulfills the condition of Lemma 15. Thus, $\rho \in \mathcal{L}(\mathcal{P}') = \mathcal{L}(\mathcal{P})$. □

5 Other probabilistic ω-automata

In generalized Büchi automata, the acceptance condition is a set $\mathcal{F} = \{F_0,\ldots,F_{m-1}\}$ consisting of subsets F_i of the state space Q. Acceptance of a run requires that each of the F_i's is visited infinitely often. In the nondeterministic case, it is well-known that the Büchi and generalized Büchi acceptance condition have equal power, in the sense that any generalized NBA $\mathcal{G}$ can be transformed into an equivalent NBA $\mathcal{A}$. (The converse is obvious, as any NBA can be viewed as a GNBA.) The idea behind the transformation GNBA $\mathcal{G} \rightsquigarrow$ NBA $\mathcal{A}$ is to work with m copies $\mathcal{G}_0,\ldots,\mathcal{G}_{m-1}$ of $\mathcal{G}$ where the outgoing transitions from a F_i-state in the i-th copy are redirected to the $(i+1)$-st copy (modulo m), while the other transitions stay in the same copy. The Büchi acceptance condition in the modified NBA $\mathcal{A}$ consists of the F_0-states in $\mathcal{G}_0$. The same technique yields a polynomial transformation from GPBA to PBA that preserves uniformity. Thus:

Theorem 17 (From GPBA to PBA). *For any (uniform) generalized probabilistic Büchi automaton $\mathcal{P}$ there is a (uniform) PBA $\mathcal{P}'$ with $\mathcal{L}(\mathcal{P}) = \mathcal{L}(\mathcal{P}')$ and $|\mathcal{P}'| = O(m|\mathcal{P}'|)$ where m is the number of acceptance-sets.*

We now consider probabilistic ω-automata with Streett or Rabin acceptance condition. Formally, a PSA or PRA is a tuple (Q,δ,μ,Acc) such that Q,δ,μ are as in PBA and $Acc = \{(H_1,K_1),\ldots,(H_m,K_m)\}$ where $H_i,K_i \subseteq Q$, $i = 1,\ldots,m$. The accepted language $\mathcal{L}_{Streett}(\mathcal{P}_S)$ of a PSA $\mathcal{P}_S$ is the set of all infinite words $\rho \in \Sigma^\omega$ where the probability for the Streett-accepting runs is positive. Similarly, for a PRA $\mathcal{P}_R$, $\mathcal{L}_{Rabin}(\mathcal{P}_R)$ denotes the set of all infinite words $\rho \in \Sigma^\omega$ where the set of Rabin-accepting runs has non-zero probability measure. (Streett-acceptance and Rabin-acceptance of a run are defined as in the non-deterministic case, see Section 2.) The notions of accepting end components and uniformity for PSA and PRA are defined in the obvious way.

Any PBA can be viewed as a PSA or PRA by replacing the Büchi-acceptance-set F with the singleton Streett/Rabin acceptence set $Acc = \{(Q,F)\}$. Thus, (uniform) PBA are special instances of (uniform) PSA and PRA. The classes of languages accepted by uniform PRA and uniform PRA agree with the class of ω-regular languages. More precisely:

Theorem 18 (From PRA/PSA to PBA).

(a) For any (uniform) PRA $\mathcal{P}_R$ there exists a (uniform) PBA $\mathcal{P}$ with $\mathcal{L}_{Rabin}(\mathcal{P}_R) = \mathcal{L}(\mathcal{P})$ and $|\mathcal{P}| = O(m|\mathcal{P}_R|)$.

(b) For any uniform PSA $\mathcal{P}_S$ there exists a uniform PBA $\mathcal{P}$ with $\mathcal{L}_{Streett}(\mathcal{P}_S) = \mathcal{L}(\mathcal{P})$ and $|\mathcal{P}| = O(m2^m|\mathcal{P}_S|)$.

(c) For any PSA $\mathcal{P}_S$ there exists a (possibly non-uniform) PBA $\mathcal{P}$ with $\mathcal{L}_{Streett}(\mathcal{P}_S) = \mathcal{L}(\mathcal{P})$ and $|\mathcal{P}| = O(m^2|\mathcal{P}_S|)$.

Here, m is the number of acceptance-pairs in $\mathcal{P}_R$ resp. $\mathcal{P}_S$.

Proof. The transformations in (a) and (b) are similar to the non-deterministic case and omitted here.

We now provide the proof for (c). Let $\mathcal{P}_S = (Q_S, \delta_S, \mu_S, \{(H_1,K_1),\ldots,(H_m,K_m)\})$ be a PSA. For simplicity, we may assume that $H_i \cap K_i = \emptyset$ as otherwise H_i could be replaced with $H_i \setminus K_i$. Intuitively, the PBA $\mathcal{P}$ arises from $\mathcal{P}_S$ by making several copies of $\mathcal{P}_S$: a subautomaton $\mathcal{P}_{init}$ in which the process starts, a subautomaton $\mathcal{P}_{accept}$ which has to be visited infinitely often and which is reachable with positive probability via any outgoing transition from the states in $\mathcal{P}_{init}$, and subautomata $\mathcal{P}_i$ and $\mathcal{P}_{i,j}$ for $i,j \in \{1,\ldots,m\}$, $i \neq j$, that are reached from $\mathcal{P}_{accept}$ whenever a state in H_i is visited in $\mathcal{P}_{accept}$. Subautomaton $\mathcal{P}_i$ can only be left via transitions from a K_i-state in $\mathcal{P}_i$ from which we move back to $\mathcal{P}_{accept}$. Subautomaton $\mathcal{P}_{i,j}$ can behave as $\mathcal{P}_i$, but in addition it also serves to take care about the Streett-acceptance pair (H_j,K_j). When we reach a H_j-state in $\mathcal{P}_{i,j}$ then we randomly choose to stay in $\mathcal{P}_{i,j}$ or to move to $\mathcal{P}_j$ or one of the subautomata $\mathcal{P}_{j,k}$. Formally, PBA $\mathcal{P} = (Q,\delta,\mu,F)$ is defined as follows. The state space is

$$Q = Q_{init} \cup Q_{accept} \cup \bigcup_{0 \le i < n} Q_i \cup \bigcup_{\substack{0 \le i,j < n \\ i \neq j}} Q_{i,j}$$

where $Q_* = \{\langle q,*\rangle : q \in Q_S\}$. The set of accepting states is $F = Q_{accept}$. The initial distribution is given by $\mu(\langle q,init\rangle) = \mu_S(q)$ and $\mu(\langle q,*\rangle) = 0$ for all other states $\langle q,*\rangle \in Q$. For the states $q \in Q_S$ and $\delta_S(q,a,p) > 0$, we have:

$\delta(\langle q,init\rangle,a,\langle p,init\rangle) > 0$	
$\delta(\langle q,init\rangle,a,\langle p,accept\rangle) > 0$	
$\delta(\langle q,accept\rangle,a,\langle p,accept\rangle) > 0$	if $q \notin H_1 \cup \ldots \cup H_m$
$\delta(\langle q,accept\rangle,a,\langle p,i\rangle) > 0$	if $q \in H_i$
$\delta(\langle q,accept\rangle,a,\langle p,i,j\rangle) > 0$	if $q \in H_i$
$\delta(\langle q,i\rangle,a,\langle p,accept\rangle) > 0$	if $q \in K_i$
$\delta(\langle q,i\rangle,a,\langle p,i\rangle) > 0$	if $q \notin K_i$
$\delta(\langle q,i,j\rangle,a,\langle p,accept\rangle) > 0$	if $q \in K_i$
$\delta(\langle q,i,j\rangle,a,\langle p,i,j\rangle) > 0$	if $q \notin K_i$
$\delta(\langle q,i,j\rangle,a,\langle p,j\rangle) > 0$	if $q \in H_j$
$\delta(\langle q,i,j\rangle,a,\langle p,j,k\rangle) > 0$	if $q \in H_j$

Here, i,j,k range over all indices in $\{1,\ldots,m\}$ with $i \neq j$ and $j \neq k$ (but possibly $i = k$). In the above we did not specify the precise transition probabilities, but we may assume that for any state $\langle q,*\rangle$ the probabilities for the transitions that arise through the lifting of $q \xrightarrow{a} p$ in $\mathcal{P}_S$ sum up to $\delta_S(q,a,p)$ and that uniform distributions are used for the switch from $\mathcal{P}_{accept}$ or $\mathcal{P}_{i,j}$ to the subautomata $\mathcal{P}_j$ and $\mathcal{P}_{j,k}$. In the sequel, we refer to the fragment of the Q_*-states as the $\mathcal{P}_*$-subautomaton.

Note that this construction ensures that whenever a H_j-state is visited in $\mathcal{P}_{accept}$ or $\mathcal{P}_{i,j}$ for some $i \neq j$ then with equal positive probability one of the subautomaton $\mathcal{P}_j$ or $\mathcal{P}_{j,k}$ is entered. Hence, if infinitely often a H_j-state is visited and the process does not stay forever in one of the subautomaton $\mathcal{P}_i$ (for some $i \neq j$) or some of the subautomata $\mathcal{P}_{k,\ell}$ then almost surely $\mathcal{P}_j$ is entered (via $\mathcal{P}_{accept}$ or one of the subautomata $\mathcal{P}_{i,j}$).[5] But $\mathcal{P}_j$ can only be left via a K_j-state. This yields that almost all accepting runs in $\mathcal{P}$ induce accepting runs in $\mathcal{P}_S$ (by throwing away the *-component of any state $\langle q,*\rangle$). Thus, $\rho \in \mathcal{L}(\mathcal{P})$ implies $\rho \in \mathcal{L}_{Streett}(\mathcal{P}_S)$.

Vice versa, let $\rho = a_1a_2\ldots \in \mathcal{L}_{Streett}(\mathcal{P}_S)$ and $\pi_S = q_0,q_1,\ldots$ an accepting run for ρ in $\mathcal{P}_S$. There is an index $r \geq 0$ such that $q_\ell \in H_i$ for some $\ell \geq r$ implies $q_h \in K_i$ for infinitely many $h \geq r$. But then, all liftings of π_S to runs for ρ in $\mathcal{P}$ that stay in $\mathcal{P}_{init}$ for the first r input symbols and eventually enter $\mathcal{P}_{accept}$ are accepting. (By a lifting of π_S we mean any run in $\mathcal{P}$ for ρ whose projection to the Q_S-components agrees with π_S.) This yields $\rho \in \mathcal{L}(\mathcal{P})$.

Note that $\mathcal{P}$ needs not to be uniform as the probability to stay in $\mathcal{P}_{i,j}$ tends to 0, if infinitely many H_j-states are visited, but no K_i-state. Nevertheless, $\mathcal{P}_{i,j}$ can be part of an accepting end component. □

Theorem 18 also applies to DRA and DSA as they arise as special instances of uniform PRA/PSA.

Following the concept of PFA [26], we may also look for PBA with a threshold $\lambda \in]0,1[$ for the accepted words. Although the uniformity condition has some similarities with

[5] By "process" we mean the stochastic process induced by $\mathcal{P}$ and a given input word ρ. Assuming the process visits infinitely often a state in $\mathcal{P}_{accept}$ and a H_j-state then with probability 1 each of the subautomata $\mathcal{P}_j$ and $\mathcal{P}_{j,k}$ are visited infinitely often.

the "isolated cutpoints" that are known to ensure that the accepted language of a PFA is regular [26], the class of languages recognized by uniform PBA with an acceptance-threshold strictly subsumes the ω-regular languages:

Theorem 19 (Acceptance beyond a threshold).

(a) *For any ω-regular language L and any $\lambda \in]0,1[$ there is a uniform PBA $\mathcal{P}$ with $L = \{\rho \in \Sigma^\omega : \Pr(\rho) > \lambda\}$.*
(b) *There exists a uniform PBA and threshold $\lambda \in]0,1[$ such that $\{\rho \in \Sigma^\omega : \Pr(\rho) > \lambda\}$ is not ω-regular.*

6 Efficiency of PBA

The transformations between nondeterministic ω-automata and PBA presented in Section 4 lead to an exponential blow-up in both directions. We now study the efficiency of PBA in more detail and show that for some languages, uniform PBA can be exponentially better than nondeterministic ω-automata, and vice versa.

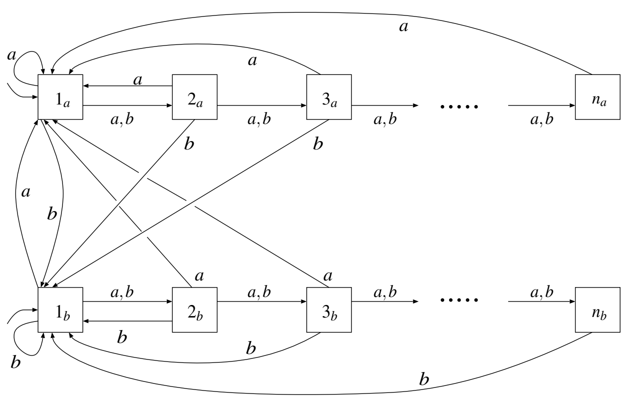

Fig. 5. **uPBA for L_n as in Theorem 20**

Theorem 20 (uPBA can be exp. better than NSA). *There exist languages L_n which are accepted by a uniform PBA with $2n$ states, while any NSA for L_n has at least $\frac{2^n}{n}$ states.*

Proof. The language $L_n = \{xy^\omega : x \in \{a,b\}^*, y \in \{a,b\}^n\}$ is accepted by the PBA $\mathcal{P}$ shown in Fig. 5, where we assume uniform distribution of the transition probabilities. Note that all states of $\mathcal{P}$ are accepting and that all states except n_a have a b-transition to the state 1_b and all states except n_b have an a-transition to the state 1_a.
Let $\rho = a_1a_2\ldots \notin L_n$. Then there are infinitely many indices i such that $a_i = a \wedge a_{i+n} = b$. Let I_ρ be the set of such indices. (Then, $|I_\rho| = \infty$.) Since every state except the state n_b has an a-transition to state 1_a, for all prob-fair runs $\pi = q_0, q_1, \ldots$ of ρ in $\mathcal{P}$ there is an infinite subset $I_\pi \subseteq I_\rho$ of indices such that $q_i = 1_a$ for all $i \in I_\pi$. Remember that $a_i = a$ and $a_{i+n} = b$ for all $i \in I_\pi$. We now consider the stochastic process induced by $\mathcal{P}$ and ρ. The above assures that with probability one the process will try to read b in state n_a and will therefore reject. Thus, almost all runs for ρ are rejecting which gives $\mathcal{L}(\mathcal{P}) \subseteq L_n$.
On the other hand, given a word $\rho \in L_n$, we can write ρ as xy^ω with $x \in \{a,b\}^*$ and $y \in \{a,b\}^n$. Then, $\hat{\pi} = 1_{c_1}, \ldots, 1_{c_k}, 1_{d_1}$ is a run for xd_1, where $x = c_1c_2\ldots c_k$ and $y = d_1d_2\ldots d_n$. (The c's and d's are symbols in $\{a,b\}$). The probability for this run is strictly greater than zero. Since from that state 1_{d_1} on, the process will never reject while reading the remaining suffix of ρ and since every infinite run is accepting, we have that ρ will be accepted with probability greater than zero. This yields $L_n \subseteq \mathcal{L}(\mathcal{P})$.
The uniformity of $\mathcal{P}$ can be derived from the fact that – as uniform distributions are assumed – for all accepting end components (P,A), all states p, all finite words $w \in \{a,b\}^*$ and all $k \in \{1,\ldots,n\}$:

$$\Pr_{(P,A)}(p \xrightarrow{w}) = 1 \vee \Pr_{(P,A)}(p \xrightarrow{w}) \leq 1 - (\tfrac{1}{2})^n$$

and either $\Pr_{(P,A)}(p \xrightarrow{w} k_a) = 0$ or $\Pr_{(P,A)}(p \xrightarrow{w} k_a) \geq (\frac{1}{2})^k \geq (\frac{1}{2})^n$. (The same holds for state k_b instead of k_a.)
It remains to show that any NSA for L_n has $\geq 2^n/n$ states. Let $\mathcal{A}$ be a NSA with $\mathcal{L}(\mathcal{A}) = L_n$. Let $x = c_1\ldots c_n$, $y = d_1\ldots d_n \in \{a,b\}^n$ such that

$$c_1\ldots c_n \neq d_i\ldots d_n d_1\ldots d_{i-1} \text{ for all } i = 1,\ldots,n \qquad (*)$$

Then, any two accepting cycles for $(c_1\ldots c_n)^\omega$ and $(d_1\ldots d_n)^\omega$ are disjoint. Otherwise, $\mathcal{A}$ would accept a word of the form $(a+b)^*(c_1\ldots c_j d_i\ldots d_n d_1\ldots d_{i-1} c_{j+1}\ldots c_n)^\omega$. But such a word is not in L_n because of (*). Thus, $\mathcal{A}$ has at least $2^n/n$ disjoint acceptance cycles, which proves the claim. □

Another example that illustrates the efficiency of PBA is the language L_n consisting of all infinite words $\rho = a_1a_2\ldots \in \{a,b,c\}^\omega$ such that, for all $0 \leq i < n$, if $a_{kn+i} = a$ for infinitely many k then $a_{kn+i} = b$ for infinitely many k, and vice versa. L_n is accepted by a PBA with $O(n^2)$ states, while any NBA for L_n has $\Omega(2^n)$ states.
Although the above results illustrate that PBA can be quite efficient, there are ω-regular languages where NBA are exponentially better than uniform PBA.

Theorem 21 (NBA can be exp. better than uPBA). *For the ω-regular languages $L_n = ((a+b)^*a(a+b)^{n-1}c)^\omega$ there exist accepting NBAs of size $O(n)$, while any uniform PBA for L_n has $\Omega(2^n)$ states.*

7 Emptiness, union and intersection

Proposition 22 (Checking emptiness). *Emptiness can be checked in time $O(\exp(|\mathcal{P}|)$ for uniform PBA. The question whether $\mathcal{L}(\mathcal{P}) \neq \emptyset$ for a given uniform PBA $\mathcal{P}$ is NP-hard.*

Proof. The question whether $\mathcal{L}(\mathcal{P}) \neq \emptyset$ for a given uniform PBA $\mathcal{P}$ can be answered via the transformation "PBA $\rightsquigarrow$ NSA" described in Section 4 and applying an emptiness-test for NSA [16]. The proof for the NP-hardness of the nonemptiness-problem can be given by a polynomial reduction from 3SAT to the nonemptiness-problem for uniform generalized PBA. (By Theorem 17 there is a polynomial reduction from the nonemptiness-problem for uniform GPBA to the nonemptiness-problem for uPBA.) □

Union and intersection. For union and intersection we may roughly apply the same techniques as for NBA. Let $\mathcal{P}_1 = (Q_1, \delta_1, \mu_1, F_1)$ and $\mathcal{P}_2 = (Q_2, \delta_2, \mu_2, F_2)$ be two PBA over the same alphabet Σ. A PBA $\mathcal{P}$ that accepts the language $\mathcal{L}(\mathcal{P}_1) \cup \mathcal{L}(\mathcal{P}_2)$ can be obtained by taking the disjoint union of the state spaces Q_1 and Q_2, equipped with the transitions in $\mathcal{P}_1$ and $\mathcal{P}_2$. The initial distribution μ in $\mathcal{P}$ assigns probability $\frac{1}{2}\mu_i(q)$ to any state $q \in Q_i$ $(i = 1,2)$. The set of accepting states of $\mathcal{P}$ is $F_1 \cup F_2$.
The intersection operator can be realized as for NBA. Using a product construction, we first generate a generalized PBA $\mathcal{P}_1 \bowtie \mathcal{P}_2$ that accepts the intersection language $\mathcal{L}(\mathcal{P}_1) \cap \mathcal{L}(\mathcal{P}_2)$, which can then be turned into an equivalent PBA (Theorem 17). The definition of the product automata $\mathcal{P}_1 \bowtie \mathcal{P}_2$ is fairly standard and omitted here.

8 Verifying probabilistic systems

Although the emptiness problem for uniform PBA is much harder than for nondeterministic ω-automata, in certain applications the use of PBA instead of nondeterministic ω-automata might be simpler, or even more efficient. We mention here the qualitative verification problem for Markov chains against ω-regular specifications.
Let $\mathsf{M} = (\mathsf{S}, \mathsf{P}, \mathsf{AP}, \mathsf{L}, \mathsf{s}_0)$ be a finite-state Markov chain with state-labels, i.e., S is a finite state space, $\mathsf{s}_0 \in \mathsf{S}$ the initial state, $\mathsf{P} : \mathsf{S} \times \mathsf{S} \to [0,1]$ a transition probability matrix such that $\sum_{\mathsf{t} \in \mathsf{S}} \mathsf{P}(\mathsf{s},\mathsf{t}) = 1$ for all states s, AP a finite set of atomic propositions and $\mathsf{L} : \mathsf{S} \to 2^{\mathsf{AP}}$ a labeling function. Let L be an ω-regular language over the alphabet $\Sigma = 2^{\mathsf{AP}}$ that specifies the "bad behaviours". M is viewed to be correct if for almost all infinite paths $\sigma = \mathsf{s}_0, \mathsf{s}_1, \ldots$ in M the induced words $\mathsf{trace}(\sigma) = \mathsf{L}(\mathsf{s}_0), \mathsf{L}(\mathsf{s}_1), \ldots$ do not belong to L, i.e., if $\Pr_{\mathsf{M}}(L) = 0$ where

$$\Pr_{\mathsf{M}}(L) = \Pr_{\mathsf{M}}\big\{\sigma : \sigma \text{ is an initial, infinite path in } \mathsf{M} \text{ with } \mathsf{trace}(\sigma) \in L \big\}.$$

In the literature, several methods to check whether $\Pr_{\mathsf{M}}(L) = 0$ have been proposed, such as special algorithms for LTL-specifications [9, 10], algorithms that assume a representation of L by a NBA [32, 9] or algorithms that operate with alternating automata [7]. As an alternative, we mention here a PBA-based approach which relies on the same, rather simple technique as it is known for deterministic ω-automata (or NBA that are deterministic in limit), but does not require the modification of the automaton, as it is the case for the approaches that start with an arbitrary nondeterministic or alternating automata.
Let $\mathcal{P} = (Q, \delta, \mu, F)$ be a PBA over the alphabet $\Sigma = 2^{\mathsf{AP}}$ with $\mathcal{L}(\mathcal{P}) = L$ and $\delta(q,a) \neq \emptyset$ for all $q \in Q$ and $a \in \Sigma$. To check whether $\Pr_{\mathsf{M}}(L) > 0$, we consider the product Markov chain $\mathsf{M} \times \mathcal{P} = (\mathsf{S} \times Q, \mathsf{P}', \mu')$ where the initial distribution μ' is given by $\mu'(\langle \mathsf{s}, q\rangle) = 0$ if $\mathsf{s} \neq \mathsf{s}_0$ and

- $\mu'(\langle \mathsf{s}_0, q_0\rangle) = \sum_{q \in Q} \mu(q) \cdot \delta(q, \mathsf{L}(\mathsf{s}_0), q_0)$,
- $\mathsf{P}'(\langle \mathsf{s}, q\rangle, \langle \mathsf{t}, p\rangle) = \mathsf{P}(\mathsf{s},\mathsf{t}) \cdot \delta(q, \mathsf{L}(\mathsf{t}), p)$.

Let $\Pr_{\mathsf{M}\times\mathcal{P}}(\Box\Diamond F)$ be the probability in $\mathsf{M} \times \mathcal{P}$ for the set of infinite paths $\langle \mathsf{s}_0, q_0\rangle, \langle \mathsf{s}_1, q_1\rangle, \ldots$ such that $q_i \in F$ for infinitely many indices i. We then have $\Pr_{\mathsf{M}}(L) > 0$ iff $\Pr_{\mathsf{M}\times\mathcal{P}}(\Box\Diamond F) > 0$. Hence, if we are given an ω-regular specification for M in form of a PBA then the qualitative model checking problem "Is $\Pr_{\mathsf{M}}(L) > 0$?" can be solved in polynomial time by building the product Markov chain and applying known graph-methods to check whether $\Pr_{\mathsf{M}\times\mathcal{P}}(\Box\Diamond F) > 0$.

9 Conclusion

In this paper, we introduced probabilistic ω-automata as acceptors for languages over infinite words. We mainly concentrated on some fundamental properties (expressiveness, efficiency, the emptiness-problem) and presented some results which we did not expect when starting to work in this area. First, we were surprised about the fact and simplicity of the proof that PBA are more expressive than NBA. The analogue result is known for PFA and NFA, but PFA are equipped with a threshold $\lambda \in]0,1[$ for acceptance, while we work with "positive acceptance probability" which on the level of finite automata leads to ordinary NFA. Secondly, we established a polynomial transformation from probabilistic Streett to probabilistic Büchi automata which is impossible in the nondeterministic case [28]. The third non-expected result was the observation that uniform PBA can be exponentially more efficient than nondeterministic Streett automata. We are not aware of such a result for finite automata.[6]
Many other theoretical questions are still open, like the decidability of whether a given PBA is uniform, the decidability (or complexity) of the emptiness-problem for arbitrary PBA, the precise complexity of the emptiness-problem for uniform PBA, a complementation operator that does not use

[6] As far as we know, [1] which shows the existence of a PFA with $O(n)$ states while any equivalent DFA has $\Omega\big(2^{n/\log n}\big)$ states, provides the best known result to illustrate the efficiency of PFA in contrast to (non)-deterministic finite automata.

nondeterministic automata as intermediate step, or the question of other criteria that ensure ω-regularity of the accepted language.

Potential application areas were mentioned in the introduction. The benefits of PBA for verification purposes (as we sketched in Section 8) will depend on whether there are reasonable logics (e.g., a sublogic of LTL) where efficient transformations into PBA – without using NBA as intermediate step – can be provided.[7]

References

1. A. Ambainis. The complexity of probabilistic versus deterministic finite automata. In *Proc. ISAAC*, LNCS 1178:233–237, 1996.
2. A. Ambainis and R. Freivalds. 1-way quantum finite automata: strengths, weaknesses and generalizations. In Proc. FOCS, 1998.
3. L. Babai and S. Moran. Arthur-Merlin games: a randomized proof system, and a hierarchy of complexity classes. *Journal of Computer Science*, 36:254–276, 1988.
4. C. Baier and M. Kwiatkowska. On the verification of qualitative properties of probabilistic processes unter fairness constraints. *Information Processing Letters*, 66:71–79, 1998.
5. V. Blondel and V. Canterini. Undecidable problems for probabilistic finite automata. *Theory of Computer Systems*, 36:231–245, 2003.
6. V. Blondel and J. Tsisiklis. A survey of computational complexity results in systems and control. *Automatica*, 36(9):1249–1274, 2000.
7. D. Bustan, S. Rubin, and M. Vardi. Verifying ω-regular properties of Markov chains. In Proc. CAV, LNCS 3114:189–201, 2004.
8. A. Condon, L. Hellerstein, S. Pottle, and A. Wigderson. On the power of finite automata with both nondeterministic and probabilistic states. *SIAM Journal on Computing*, 27(3):739–762, 1998.
9. C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *Journal of the ACM*, 42(4):857–907, July 1995.
10. J. Couvreur, N. Saheb, and G. Sutre. An optimal automata approach to LTL model checking of probabilistic systems. In Proc. LPAR, LNAI 2850:361–375, 2003.
11. L. de Alfaro. *Formal Verification of Probabilistic Systems*. PhD thesis, Stanford University, 1997.
12. L. de Alfaro. Stochastic transition systems. In Proc. CONCUR, LNCS 1466:423–438, 1998.

[7] Any LTL-formula with atoms in AP can be represented by a uniform PBA over the alphabet $\Sigma = 2^{\mathsf{AP}}$, using transformations "LTL $\rightsquigarrow$ NBA $\rightsquigarrow$ NBA which is deterministic in limit = PBA". However, this technique might lead to a double-exponential blow-up. For next-free formulas of restricted linear temporal logic uniform PBA of single exponential size can be obtained by applying the construction of Vardi and Wolper [32] which yields a NBA that is deterministic in limit of single-exponential size.

13. M. Droste and D. Kuske. Skew and infinitary formal power series. In Proc. ICALP, LNCS 2719:426–438, 2003.
14. M. Droste and U. Püschmann. Weighted Büchi automata. In *Workshop on Weighted Automata–Theory and Applications,*, page 56, Dresden, 2004.
15. C. Dwork and L. Stockmeyer. A time-complexity gap for two-way probabilistic finite state automata. *SIAM Journal of Computing*, 19:1011–1023, 1990.
16. E. Emerson and C. Lei. Modalities for model checking: Branching time logic strikes back. *Science of Computer Programming*, 8(3):275–306, 1987.
17. R. Freivalds. Probabilistic two-way machines. In *Proceedings of the International Symposium on Mathematical Foundations of Computer Science*, LNCS 188:33–45, 1981.
18. R. Freund, M. Oswald, and L. Staiger. Omega-P automata with communication rules. In *Workshop on Membrane Computing*, LNCS 2933:203–217, 2003.
19. E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, LNCS 2500, 2002.
20. J. Kemeny and J. Snell. *Denumerable Markov Chains*. Springer-Verlag, New York, 1976.
21. A. Kondacs and J. Watrous. On the power of quantum finite state automata. In Proc. FOCS, pages 66–75, 1997.
22. V. Kulkarni. *Modeling and Analysis of Stochastic Systems*. Chapman & Hall, 1995.
23. O. Madani, S. Hanks, and A. Condon. On the undecidability of probabilistic planning and infinite-horizon partially observable markov chains. In *Proc. 16th National Conference on Artifical Intelligence*, pages 541–548, 1999.
24. L. Mora-Lopez, R. Morales, M. Sidrach de Cardona, and F. Triguero. Probabilistic finite automata and randomness in nature: a new approach in the modelling and prediction of climatic parameters. In *Proc. International Environmental Modelling and Software Congress*, pages 78–83, 2002.
25. A. Paz. Some aspects of probabilistic automata. *Information and Control*, 9, 1966.
26. M. O. Rabin. Probabilistic automata. *Information and Control*, 6:230–245, 1963.
27. D. Ron, Y. Singer, and N. Tishby. The power of amnesia: Learning probabilistic automata with variable memory length. *Machine Learning*, 25(2–3):117–149, 1996.
28. S. Safra and M. Vardi. On ω-automata and temporal logic. In Proc. STOC, pages 127–137, 1989.
29. W. Thomas. Languages, automata, and logic. *Handbook of formal languages*, 3:389–455, 1997.
30. M. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In Proc. FOCS, pages 327–338, 1985.
31. M. Vardi. An automata-theoretic approach to linear temporal logic. In *Proc. Banff Higher Order Workshop on Logics for Concurrency: Structure versus Automata*, pages 238–266. Springer-Verlag, 1996.
32. M. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification. In Proc. LICS'86, pages 332–345, IEEE CS Press, 1986.
33. M. Vardi. Nontraditional applications of automata theory. In *Theoretical Aspects of Computer Software*, LNCS 789:575–597, 1994.