Fast Synchronization of Random Automata

Cyril Nicaud

Fast Synchronization of Random Automata

Cyril Nicaud*

Université Paris-Est, LIGM (UMR 8049), CNRS, ENPC, ESIEE Paris, UPEM, Marne-la-Vallée, France
cyril.nicaud@u-pem.fr

Abstract

A synchronizing word for an automaton is a word that brings that automaton into one and the same state, regardless of the starting position. Černý conjectured in 1964 that if a n-state deterministic automaton has a synchronizing word, then it has a synchronizing word of length at most $(n-1)^2$. Berlinkov recently made a breakthrough in the probabilistic analysis of synchronization: he proved that, for the uniform distribution on deterministic automata with n states, an automaton admits a synchronizing word with high probability. In this article, we are interested in the typical length of the smallest synchronizing word, when such a word exists: we prove that a random automaton admits a synchronizing word of length $\mathcal{O}(n \log^3 n)$ with high probability. As a consequence, this proves that most automata satisfy the Černý conjecture.

1998 ACM Subject Classification G.2.1 Combinatorics

Keywords and phrases random automata, synchronization, the Černý conjecture

Digital Object Identifier 10.4230/LIPIcs.APPROX-RANDOM.2016.43

1 Introduction

A *synchronizing word* (or a *reset word*) for an automaton is a word that brings that automaton into one and the same state, regardless of the starting position. This notion, first formalized by Černý in the sixties, arises naturally in automata theory and its extensions, and plays an important role in several application areas [17]. Perhaps one of the reasons synchronizing automata are still intensively studied in theoretical computer science is the following question asked by Černý [16] back in 1964: "*Does every synchronizing n-state automaton admits a synchronizing word of length at most $(n-1)^2$?*" The upper bound of $(n-1)^2$, as shown by Černý, is best possible. This question, known as *the Černý conjecture*, is now one of the most famous conjectures in automata theory. Though established for important subclasses of automata, the Černý conjecture remains open in the general case. The best known upper bound, established in the early eighties [13, 6], is $\frac{1}{6}(n^3-n)$. We refer the interested reader to Volkov's article [17] for a more detailed account on the Černý conjecture.

1.1 The probabilistic Cerný conjecture

In this article, we consider the Černý conjecture from a probabilistic point of view (as proposed, for example, by Cameron[1] in [4]). This leads to the following questions, for the uniform distribution on deterministic automata with n states, on a fixed alphabet:

* This work is supported by the French National Agency (ANR) through ANR-JCJC-12-JS02-012-01.

[1] Cameron studied the transformation monoid generated by a fixed number of mappings from a set Ω of size n to itself. This is the same as a deterministic automaton, where each mapping correspond to the action of a letter on the set of states Ω. In these settings, "Is the automaton synchronizing?" translates directly into "Does the monoid contain a constant mapping?".

Question 1: Is a random automaton synchronizing *with high probability*?
Question 2: Does a synchronizing automaton admits a synchronizing word of length at most $(n-1)^2$ *with high probability*?

Where *with high probability* means "with probability that tends to 1 as n goes to infinity".

1.2 Main related results

Berlinkov recently made a breakthrough [2] by giving a positive answer to Question 1: he proved that the probability that a random automaton is not synchronizing is $\mathcal{O}(n^{-\frac{1}{2}|A|})$, for an alphabet A with at least two letters.

Question 2 only received partial results so far: Skvortsov and Zaks [15] gave a positive answer for alphabets whose cardinality grows as n^β for $\beta > \frac{1}{2}$. They also proved that the probability of having a short reset word is non-negligible for alphabets with at least four letters [18].

Question 2 can also be simulated and experimental evidence suggests that most automata are synchronized by a short synchronizing word, of length sublinear in the number of states. Note that simulating the second question is nontrivial, as most problems related to the shortest reset word are hard [12] (for instance, deciding whether the shortest reset word as length ℓ is DP-complete, where DP is the closure of NP∪coNP for finite intersections); the best experimental results we are aware of were obtained by Kisielewicz, Kowalski, and Szykula [9]. According to these results, the expected length of the shortest reset word, when it exists, seems to grow in $\Theta(\sqrt{n})$.

Note finally that Berlinkov and Szykula [3] recently used the results of this paper to establish a bound of $n^{3/2+o(1)}$ for the expected value of the shortest reset word in a random synchronizing automaton.

1.3 Our results

In this paper we give a positive answer to Question 2 when the automaton is chosen uniformly among deterministic and complete n-state automata on an alphabet with at least two letters. More precisely, we show that with high probability, a random n-state automaton admits a synchronizing word of length $\mathcal{O}(n \log^3 n)$.

Even if the Černý conjecture is settled in the positive, our main result remains interesting, as it yields that most automata admit a synchronizing word of length almost linear.

Our proof also gives another way to show that automata are synchronizing with high probability, based a method that differ completely from Berlinkov's work. He used recent results on synchronization, as well as some advanced properties of random mappings. In our proof, we directly build words that iteratively shrink the set of states, using only basic discrete probabilities and variations on the probabilistic pigeonhole principle (also known as the Birthday Paradox). The proof proposed by Berlinkov is arguably more complicated, but also more precise, since it gives a sharp estimation of the probability of not being synchronizing[2].

Due to lack of space, the proofs are omitted in this extended abstract.

[2] Knowing the probability of not being synchronizing is important in many situations, especially for the average case analysis of algorithms, as illustrated in the conclusions of [2]. Berlinkov also replies precisely to a question asked by Cameron [4].

2 Definitions and notations

For any integer $n \geq 1$, let $[n] = \{1, \ldots, n\}$ be the set of integers between 1 and n. The cardinality of a finite set E is denoted by $|E|$.

2.1 Automata

Let A be a finite alphabet, a *deterministic automaton* on A is a pair (Q, δ), where Q is a finite set of *states* and δ is the *transition function*, a (possibly partial) function from $Q \times A$ to Q. If $p, q \in Q$ and $a \in A$ are such that $\delta(p, a) = q$, then (p, a, q) is the *transition* from p to q labelled by a, and is denoted by $p \xrightarrow{a} q$. It is the *a-transition* outgoing from p. Since we only consider deterministic automata in this article, we simply call them *automata* in the sequel.

An automaton $\mathcal{A} = (Q, \delta)$ on A is classically seen as a labelled directed graph, whose set of vertices is Q and whose edges are the transitions of $\mathcal{A}$.

An automaton is *complete* when its transition function is a total function and *incomplete* otherwise. The transition function is extended inductively to $Q \times A^*$ by setting $\delta(p, \varepsilon) = p$ for every $p \in Q$ and, for every $u \in A^*$, $\delta(p, ua) = \delta(\delta(p, u), a)$ when everything is defined, and undefined otherwise. If $u \in A^*$, we denote by δ_u the (possibly partial) function from Q to Q defined by $\delta_u(p) = \delta(p, u)$, for all $p \in Q$.

If $\mathcal{A} = (Q, \delta)$ is an automaton on A, an *extension* of $\mathcal{A}$ is an automaton $\mathcal{B} = (Q, \lambda)$ on A such that for all $p \in Q$ and for all $a \in A$, if $\delta(p, a)$ is defined then $\lambda(p, a) = \delta(p, a)$. The automaton $\mathcal{B}$ is therefore obtained from $\mathcal{A}$ by adding some transitions. We denote by $\mathbf{Ext}(\mathcal{A})$ the set of all the extensions of an automaton $\mathcal{A}$. If $\mathcal{H}$ is a set of automata, we denote by $\mathbf{Ext}(\mathcal{H})$ the union of all the $\mathbf{Ext}(\mathcal{A})$ for $\mathcal{A} \in \mathcal{H}$.

2.2 Synchronization

Let $\mathcal{A}$ be an automaton on A. Two states p and q of $\mathcal{A}$ are *synchronized* by the word $w \in A^*$ when both $\delta_w(p)$ and $\delta_w(q)$ are defined and equal.

A *synchronizing word* for an automaton $\mathcal{A} = (Q, \delta)$ is a word $w \in A^*$ such that δ_w is a constant map: there exists a state $r \in Q$ such that for every p in Q, $\delta_w(p) = r$. An automaton that admits a synchronizing word is said to be *synchronizing*.

2.3 Mappings

A *mapping* on a set E is a total function from E to E. When E is finite, a mapping f on E can be seen as a directed graph with an edge $i \to j$ whenever $f(i) = j$. An example of such a graph is depicted in Figure 1.

Let f be a mapping on E. The element $x \in E$ is a *cyclic point*[3] of f when there exists an integer $i > 0$ such that $f^i(x) = x$. In the sequel, E will often be the set of states of an automaton, and we will therefore use the term "state" instead of "point".

If f is a mapping on E and $x \in E$, the *height* of x is the smallest $i \geq 0$ such that $f^i(x)$ is a cyclic point. The height of a cyclic point is therefore 0. The *height* of a mapping on E is the maximal height of an element of E. The mapping depicted in Figure 1 has height 3, and the maximal height is reached by the state 9.

[3] We will also say a *f-cyclic point* when the mapping under consideration is not clear in the context.

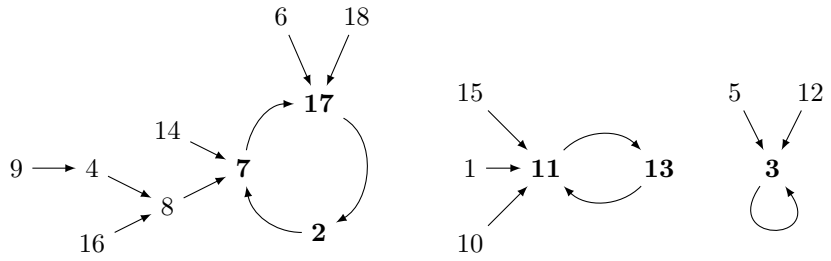

Figure 1 The graph of a mapping for $n = 18$. The cyclic points are indicated in bold.

2.4 Probabilities

Let (E, s) be a pair where E is a set and s is a *size function* from E to $\mathbb{Z}_{\geq 0}$. The pair (E, s) is a combinatorial set[4] when for every integer $n \geq 0$, the set E_n of size-n elements of E is finite. To simplify the definitions, we also assume that $E_n \neq \emptyset$ for every $n \geq 1$, which will always be the case in the following. Let $(\mathbb{P}_n)_{n\geq 1}$ be a sequence of total functions such that for each $n \geq 1$, $\mathbb{P}_n$ is a probability on E_n. We say that a property P holds *with high probability* for $(\mathbb{P}_n)_{n\geq 1}$ when $\mathbb{P}_n[P \text{ holds}] \to 1$ as $n \to \infty$.

We will often consider the *uniform distribution* on E, which is the sequence $(\mathbb{P}_n)_{n\geq 1}$ defined by $\mathbb{P}_n[\{e\}] = \frac{1}{|E_n|}$ for any e in E_n: A sentence like "property P holds with high probability for the uniform distribution on E" therefore means that the probability that P holds tends to 1 as n tends to infinity, when for each n we consider the uniform distribution on E_n. The reader is referred to [5] for more information on combinatorial probabilistic models.

2.5 Random mappings and random p-mappings

A *random mapping* of size $n \geq 1$ is a mapping on $[n]$ taken with the uniform distribution. If p is a probability mass function on $[n]$, a random p-mapping is the distribution on the mappings on $[n]$ such that the probability of a mapping f is $\prod_{i\in[n]} p(f(i))$: the image of each $i \in [n]$ is chosen independently of the others, following the probability p.

A result stated as "a random p-mapping satisfies property P with high probability" means that for *any* sequence $(p_n)_{n\geq 1}$, where p_n is a probability on $[n]$, the probability that a p_n-random mapping on $[n]$ satisfies P tends to 1 as n tends to infinity. It is therefore a strong result that does not depend on the choice of $(p_n)_{n\geq 1}$.

2.6 Random automata

In the sequel, the set of states of an n-state automaton will always be $[n]$. With this condition, there are exactly $n^{|A|n}$ complete automata with n states on $|A|$. For the uniform distribution, each size-n complete automaton has therefore probability $n^{-|A|n}$. See [11] for a recent account on the typical properties of uniform random deterministic automata.

Remark that one can also see this distribution as drawing uniformly at random and independently in $[n]$ the image of each $\delta(p, a)$, for all $p \in [n]$ and $a \in A$. These alternative way to look at random automata will be widely used in the sequel, especially in the following way: Let $\mathcal{A}$ be a fixed incomplete automaton with n states. The uniform distribution on complete

[4] The size is often clear in the context (number of nodes in a tree, ...) and can be omitted.

automata of $\mathbf{Ext}(\mathcal{A})$ is obtained by choosing uniformly at random and independently in $[n]$ the transitions that are undefined in $\mathcal{A}$.

3 Preliminary classical results

In this section, we recall some classical results that will be useful in sequel. Though elementary, these results are the main ingredients of this article.

We start with the following property for synchronizing automata: an automaton is synchronizing if and only if every pair of states can be synchronized.

▶ **Lemma 1.** *Let $\mathcal{A}$ be an n-state automaton and ℓ be a non-negative integer. If for every pair of states (p, q) in $\mathcal{A}$ there exists a word u of length at most ℓ such that $\delta_u(p) = \delta_u(q)$, then $\mathcal{A}$ admits a synchronizing word of length at most $\ell(n-1)$.*

Random mappings and random p-mappings have been studied intensively in the literature [7, 14, 10], using probabilistic techniques or methods from analytic combinatorics. In this section, we only recall basic properties of the typical number of cyclic points and of the typical height of a random p-mapping. This can be achieved using variations on the probabilistic pigeonhole principle only; more advanced techniques can be used to obtain more precise statements[5], but we will only need the following results in the sequel.

▶ **Lemma 2.** *The probability that a random p-mapping of size n has more than $2\sqrt{n \log n}$ cyclic points or that it has height greater than $2\sqrt{n \log n}$ is $\mathcal{O}(\frac{1}{n})$.*

The proof of Lemma 2 consists of two steps. It is first established for uniform random mappings then extended to general p-random mappings, by proving that the uniform case is the worst possible distribution for a p-random mapping.

4 Main Result

The main result of this article is the following theorem.

▶ **Theorem 3.** *Let A be an alphabet with at least two letters. For the uniform distribution, an n-state deterministic and complete automaton on A admits a synchronizing word of length $\mathcal{O}(n \log^3 n)$ with high probability. More precisely, the probability that no such word exists is $\mathcal{O}(n^{-\frac{1}{8}} \log^4 n)$.*

The statement does not hold for alphabets with only one letter, since there are cycles of length greater than 1 in a random mapping with high probability [14]: two distinct states in such a cycle cannot be synchronized.

As a consequence of Theorem 3, a random deterministic and complete automaton is synchronizing with high probability; our proof therefore constitutes an alternative proof of [2] for that property. Our statement is weaker, since Berlinkov also obtained the upper bound $\mathcal{O}(n^{-\frac{1}{2}|A|})$ for the error term (the number of automata that are not synchronizing), which is tight for two-letter alphabets. On the other hand, it is arguably more elementary as we mostly rely on Lemma 2 and some basic discrete probabilities; in any cases, we hope our proof sheds a new light on the reasons why automata are often synchronizing.

[5] For instance, limit distributions of some parameters [5] or even a notion of continuous limit for random mappings [1].

If we consider the uniform distribution on synchronizing automata, we directly obtain that there exists a small synchronizing word with high probability, yielding the following corollary.

▶ **Corollary 4.** *For the uniform distribution on synchronizing deterministic and complete automata on an alphabet with at least two letters, the Černý conjecture holds true with high probability.*

We prove Theorem 3 in two main steps:

- We first construct a word $w_n \in \{a, b\}^*$ such that the image of δ_{w_n} for a random n-state automaton has size at most $n^{1/8} \log^{7/8} n$ with high probability. This is done by building a set $\mathcal{G}_n$ of incomplete automata that have this property, and showing that a random n-state automaton extends an element of $\mathcal{G}_n$ with high probability. Roughly speaking, $\mathcal{G}_n$ and w_n are built by three consecutive applications of Lemma 2, starting with incomplete automata with only a-transitions, which we then augment by b-transitions in two rounds.
- It remains to synchronize those $n^{1/8} \log^{7/8} n$ states. This is done by showing that for a random automaton that extends an element of $\mathcal{G}_n$, with high probability any two of those states can be synchronized by a word of the form $b^i w_n$, with $i \leq n^{1/4}$. Lemma 1 is then used to combine these words, and also w_n, into a synchronizing word for the automaton.

The remainder of this section is devoted to a more detailed proof of Theorem 3. For the presentation, we will follow an idea used by Karp in his article on random direct graphs [8]: we start from an automaton with no transition, then add new random transitions during at each step of the construction, progressively improving the synchronization.

Since it is clearly sufficient to establish the result for a two-letter alphabet, we consider that $A = \{a, b\}$ from now on, except for the informal discussion at the beginning of Section 4.3.

4.1 Generating the a-transitions

The first step consists in generating all the a-transitions. This forms a mapping for δ_a that follows the uniformly distribution on size-n mappings. We can therefore apply Lemma 2, and obtain that words of the form a^i can already be used to reduce significantly the number of states to be synchronized.

Let $\alpha_n = \lfloor 2\sqrt{n \log n} \rfloor$ and let $\mathcal{E}_n$ denote the set of incomplete automata $\mathcal{A}$ with n states such that:

1. The defined transitions of $\mathcal{A}$ are exactly its a-transitions.
2. The action δ_a of a has at most α_n cyclic states.
3. The height of δ_a is at most α_n.

▶ **Example 5.** Let $\mathcal{A}$ be an automaton with 18 states, which has only a-transitions and such that δ_a is the mapping of Figure 1 page 4. Its set $\mathbf{Cyc_a}(\mathcal{A})$ is $\{\mathbf{2}, \mathbf{3}, \mathbf{7}, \mathbf{11}, \mathbf{13}, \mathbf{17}\}$. Since $\alpha_{18} = 14$, the word $u_{18} = a^{14}$ is used to start the synchronization:

$$\{6, 7, 9, 18\} \xrightarrow{u_{18}} \mathbf{2}; \quad \{3, 5, 12\} \xrightarrow{u_{18}} \mathbf{3}; \quad \{4, 16, 17\} \xrightarrow{u_{18}} \mathbf{7};$$
$$\{11\} \xrightarrow{u_{18}} \mathbf{11}; \quad \{1, 10, 13, 15\} \xrightarrow{u_{18}} \mathbf{13}; \quad \{2, 8, 14\} \xrightarrow{u_{18}} \mathbf{17}.$$

As there are $6 \leq \alpha_{18}$ cyclic states and since this mapping's height is $3 \leq \alpha_{18}$, the automaton $\mathcal{A}$ is an element of $\mathcal{E}_{18}$.

As the action of the letter a in a uniform random complete automaton is exactly a uniform random mapping, the following result is a direct consequence of Lemma 2.

▶ **Lemma 6.** *A random complete automaton with n states extends an element of $\mathcal{E}_n$ with high probability. More precisely, the probability that such an automaton does not extend an element of $\mathcal{E}_n$ is $\mathcal{O}(\frac{1}{n})$.*

For any automaton $\mathcal{A}$ whose a-transitions are all defined, let $\mathbf{Cyc_a}(\mathcal{A})$ denote its set of δ_a-cyclic states. They also are the δ_a-cyclic states of any automaton that extends $\mathcal{A}$.

Let $u_n = a^{\alpha_n}$. By Lemma 6, we can already start the synchronization using u_n, as the image of the set of states $[n]$ by δ_{u_n} is included in $\mathbf{Cyc_a}(\mathcal{A})$, which is much smaller than n with high probability. In the sequel, we therefore work on synchronizing the elements of $\mathbf{Cyc_a}(\mathcal{A})$.

4.2 Adding some random b-transitions

Let $\mathcal{A}$ be a fixed element of $\mathcal{E}_n$. We are now working on $\mathbf{Ext}(\mathcal{A})$ and we consider the process of adding a random b-transition starting from every state of $\mathbf{Cyc_a}(\mathcal{A})$.

Let $\mathcal{B} \in \mathbf{Ext}(\mathcal{A})$ be an automaton obtained this way and let $f_{\mathcal{B}}$ denote the restriction of δ_{bu_n} to $\mathbf{Cyc_a}(\mathcal{A})$. It is a total map, since all the needed b-transitions are defined. Moreover, the image of $f_{\mathcal{B}}$ is included in $\mathbf{Cyc_a}(\mathcal{A})$, as $f_{\mathcal{B}}(x) = \delta_{bu_n}(x) = \delta_{u_n}(\delta_b(x))$, for every $x \in \mathbf{Cyc_a}(\mathcal{A})$. Hence $f_{\mathcal{B}}$ is a total map from $\mathbf{Cyc_a}(\mathcal{A})$ to itself.

▶ **Example 7.** This is the automaton of Example 5, where the b-transitions originating from the elements of $\mathbf{Cyc_a}(\mathcal{A})$ have been added (in bold):

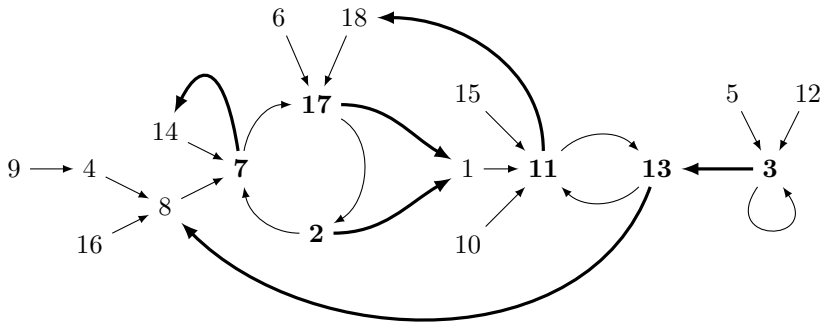

The map $f_{\mathcal{B}}$, which is the restriction of δ_{bu_n} to $\mathbf{Cyc_a}(\mathcal{A})$, is depicted below. An edge $\mathbf{p} = x \Longrightarrow \mathbf{q}$ means that $\delta_b(p) = x$ and $\delta_{u_n}(x) = q$, so that $f_{\mathcal{B}}(p) = q$:

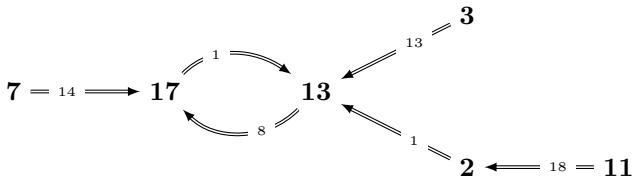

From a probabilistic point of view, if we fix $\mathcal{A}$ and build $\mathcal{B}$ by adding uniformly at random and independently the b-transitions that start from the states of $\mathbf{Cyc_a}(\mathcal{A})$, the induced distribution for the mapping $f_{\mathcal{B}}$ is usually not the uniform distribution on the mappings of $\mathbf{Cyc_a}(\mathcal{A})$. More precisely, for any $q \in \mathbf{Cyc_a}(\mathcal{A})$ the probability that the image by $f_{\mathcal{B}}$ of an element of $\mathbf{Cyc_a}(\mathcal{A})$ is q is proportional to the number of preimages of q by δ_{u_n}. It is exactly

$\frac{1}{n}|\delta_{u_n}^{-1}(\{q\})|$, the probability that a random state is mapped to q when reading u_n. For any word $\omega \in A^*$, let $\mathbb{P}_{\mathcal{A},\omega}$ be the function from $[n]$ to $[0,1]$ defined by

$$\mathbb{P}_{\mathcal{A},\omega}(q) = \frac{|\delta_\omega^{-1}(\{q\})|}{n}, \text{ for all } q \in [n]. \tag{1}$$

From the observations above, we get that once $\mathcal{A}$ is fixed, $f_\mathcal{B}$ is a random p-mapping, where the distribution on $\mathbf{Cyc_a}(\mathcal{A})$ is given by the restriction of $\mathbb{P}_{\mathcal{A},u_n}$ to $\mathbf{Cyc_a}(\mathcal{A})$.

Let $\beta_n = \lfloor 3\, n^{1/4} \log^{3/4} n \rfloor$. Applying Lemma 2 to $f_\mathcal{B}$ yields the following result.

▶ **Lemma 8.** *Let $\mathcal{A}$ be a fixed automaton of $\mathcal{E}_n$. Consider the random process of building $\mathcal{B}$ by adding a b-transition to every element of $\mathbf{Cyc_a}(\mathcal{A})$, choosing the target uniformly and independently in $[n]$. For n sufficiently large, the probability that $f_\mathcal{B}$ has more than β_n cyclic states or that it has height greater than β_n is smaller than $\frac{M}{n^{1/4}}$, for some positive constant M that does not depends on $\mathcal{A}$.*

For any automaton $\mathcal{B}$ whose a-transitions are all defined and whose b-transitions starting from an element of $\mathbf{Cyc_a}(\mathcal{B})$ are also all defined, let $\mathbf{Cyc_f}(\mathcal{B})$ denote the set of $f_\mathcal{B}$-cyclic states of $\mathcal{B}$.

Let $v_n = u_n(bu_n)^{\beta_n}$. At this point, the number of states to be synchronized has been reduced to less than β_n with high probability, since the image of δ_{v_n} is usually included in $\mathbf{Cyc_f}(\mathcal{B})$. It has been achieved by generating all the a-transitions, but using only the b-transitions that start from the δ_a-cyclic states: there still remain at least $n - \alpha_n$ undefined b-transitions that can be used to continue the synchronization. Nonetheless, before going on, we first refine the construction of $\mathcal{B}$ introduced in this section by forbidding some cases, for technical reasons explained in the next section.

4.3 Forbidding correlated shapes

The number of states to be synchronized has been reduced to no more than β_n states with high probability, but this quantity is still too large. For the technique used at the end of the proof, we need to shrink this set once more. Should the alphabet contain one more letter c, we could use the same kind of construction as in Section 4.2, and be left with at most, roughly, $n^{1/8}$ states to synchronize. This is because c-transitions can be generated independently of what has been done during the previous steps.

Some care is required to adapt this idea for a two-letter alphabet. We aim at using the word bb instead of the letter c in the informal description above. Let $\mathcal{B}$ be an incomplete automaton that extends $\mathcal{A} \in \mathcal{E}_n$ and whose defined transitions are all the a-transitions and also the b-transitions that start from the δ_a-cyclic states. We are interested in building an automaton $\mathcal{C}$ from $\mathcal{B}$, by adding some new random b-transitions, in a way such that δ_{bbv_n} is totally defined on $\mathbf{Cyc_f}(\mathcal{B})$. It means that for every $q \in \mathbf{Cyc_f}(\mathcal{B})$, the state $\delta_b(q)$ must have an outgoing b-transition in $\mathcal{C}$. For such an extension $\mathcal{C}$ of $\mathcal{B}$, let $g_\mathcal{C}$ denote the restriction of δ_{bbv_n} to $\mathbf{Cyc_f}(\mathcal{B})$.

The main point here is that for a fixed $\mathcal{B}$, we want $g_\mathcal{C}$ to be defined as a random p-mapping, so that we can use Lemma 2 once more. There are, *a priori*, two kind of issues that can prevent this from happening:

1. When there exists a state $q \in \mathbf{Cyc_f}(\mathcal{B})$ such that the b-transition starting from $\delta_b(q)$ is already defined in $\mathcal{B}$, that is, when $\delta_b(q) \in \mathbf{Cyc_a}(\mathcal{B})$.
2. When two distinct states q and q' in $\mathbf{Cyc_f}(\mathcal{B})$ are such that $\delta_b(q) = \delta_b(q')$.

Fortunately, the second case cannot occur: if $\delta_b(q) = \delta_b(q')$ then $f_\mathcal{B}(q) = f_\mathcal{B}(q')$, which is not possible for two distinct $f_\mathcal{B}$-cyclic states.

The first case can occur, and then the image of $\delta_b(q)$ by b is already defined in $\mathcal{B}$ and therefore $g_{\mathcal{C}}$ does not follow a p-distribution when we build $\mathcal{C}$ by generating the missing transitions uniformly at random[6].

Conversely, if for every $q \in \mathbf{Cyc_f}(\mathcal{B})$, $\delta_b(q) \notin \mathbf{Cyc_a}(\mathcal{B})$, then it is easy to verify that $g_{\mathcal{C}}$ is a random p-mapping: the image of $q \in \mathbf{Cyc_f}(\mathcal{B})$ by $g_{\mathcal{C}}$ is a given x when $\delta_{bbv_n}(q) = x$, which is equivalent to $\delta_b(\delta_b(q)) \in \delta_{v_n}^{-1}(\{x\})$. Since $\delta_b(\delta_b(q))$ is chosen uniformly at random in $[n]$, it happens with probability $\mathbb{P}_{\mathcal{B},v_n}(x)$, using the notation of Equation (1).

We therefore forbid the bad cases and define the set $\mathcal{F}_n$ of incomplete automata $\mathcal{B}$ with n states such that (we add the last condition to what was done in the previous section):

1. $\mathcal{B}$ extends an element of $\mathcal{E}_n$.
2. The defined transitions of $\mathcal{B}$ are all the a-transitions and the b-transitions starting from the states of $\mathbf{Cyc_a}(\mathcal{B})$.
3. The map $f_{\mathcal{B}}$ has height at most β_n and has at most β_n cyclic states.
4. For every $q \in \mathbf{Cyc_f}(\mathcal{B})$, $\delta_b(q) \notin \mathbf{Cyc_a}(\mathcal{B})$.

▶ **Example 9.** The automaton of Example 7 is in $\mathcal{F}_n$. For the fourth condition, observe that the $f_{\mathcal{B}}$-cyclic states are **13** and **17**. Their images by δ_b, which are 8 and 1 respectively, are not in $\mathbf{Cyc_a}(\mathcal{B})$. The fact that $\delta_b(\mathbf{3})$ is in $\mathbf{Cyc_a}(\mathcal{B})$ is not a problem here, since **3** is not an $f_{\mathcal{B}}$-cyclic state.

If we forget the last condition in the definition of $\mathcal{F}_n$, the other requirements hold with high probability for every fixed $\mathcal{A} \in \mathcal{E}_n$, as a consequence of Lemma 8. Lemma 10 below states that after our additional restriction, the set we obtain is still sufficiently large.

▶ **Lemma 10.** *With high probability a random complete automaton with n states extends an element of $\mathcal{F}_n$. More precisely, the probability that it does not extend an element of $\mathcal{F}_n$ is at most $n^{-1/4}\log^2 n$, for n sufficiently large.*

4.4 Adding more random b-transitions

Starting from an element of $\mathcal{B} \in \mathcal{F}_n$, we can now use the idea explained at the beginning of Section 4.3, and add the random b-transitions that are needed for δ_{bb} to be totally defined on $\mathbf{Cyc_f}(\mathcal{B})$. For such an extension $\mathcal{C}$ of $\mathcal{B}$, recall that the mapping $g_{\mathcal{C}}$ is the restriction of δ_{bbv_n} to $\mathbf{Cyc_f}(\mathcal{B})$. Let $\mathbf{Cyc_g}(\mathcal{C})$ denote the set of $g_{\mathcal{C}}$-cyclic states in $\mathcal{C}$. Thanks to the last condition of the definition of $\mathcal{F}_n$, we need to randomly choose the b-transitions starting from the images by δ_b of $\mathbf{Cyc_f}(\mathcal{B})$, which are all distinct since two distinct states of $\mathbf{Cyc_f}(\mathcal{B})$ cannot have the same image by δ_b.

Let $\gamma_n = \lfloor 2\, n^{1/8} \log^{7/8} n \rfloor$ and let $X_{\mathcal{B}}$ denote the set of images of $\mathbf{Cyc_f}(\mathcal{B})$ by δ_b, i.e. $X_{\mathcal{B}} = \{\delta_b(x) : x \in \mathbf{Cyc_f}(\mathcal{B})\}$. We define the set $\mathcal{G}_n$ of incomplete automata $\mathcal{C}$ with n states that satisfy the following conditions:

1. $\mathcal{C}$ extends an automaton $\mathcal{B}$ of $\mathcal{F}_n$.
2. The only b-transitions of $\mathcal{C}$ are those starting from $\mathbf{Cyc_a}(\mathcal{B})$ and from $X_{\mathcal{B}}$.
3. The map $g_{\mathcal{C}}$ has no more than γ_n cyclic states and has height at most γ_n.
4. For every $q \in \mathbf{Cyc_g}(\mathcal{C})$ the b-transition of $\delta_{bb}(q)$ is undefined.

The last condition in the definition of $\mathcal{G}_n$ is useful for the same kind of reasons than the last condition of $\mathcal{F}_n$ is. It ensures some independence for the final step of the synchronization, which is presented in Section 4.5.

[6] Except in the very degenerate case where the restriction of δ_{bb} to $\mathbf{Cyc_f}(\mathcal{B})$ is already a totally defined and constant map in $\mathcal{B}$.

▶ **Lemma 11.** *With high probability, a random complete automaton with n states extends an element of $\mathcal{G}_n$. More precisely, the probability it does not extends an element of $\mathcal{G}_n$ is $\mathcal{O}(\frac{1}{\gamma_n})$.*

Let $w_n = v_n(bbv_n)^{\gamma_n}$. Lemma 11 ensures that for a random complete automaton $\mathcal{A}$, the image of δ_{w_n} is usually included in $\mathbf{Cyc_g}(\mathcal{A})$, which has size at most γ_n. This concludes the first part of the synchronization: with high probability, the word w_n maps the set of states of $\mathcal{A}$ to the much smaller set of states $\mathbf{Cyc_g}(\mathcal{A})$.

4.5 Synchronizing the remaining states

Let $\lambda_n = \lfloor n^{1/4} \rfloor$ and let $\mathcal{C}$ be a fixed automaton of $\mathcal{G}_n$. Starting from $\mathcal{C} \in \mathcal{G}_n$, we now prove that the elements of $\mathbf{Cyc_g}(\mathcal{C})$ can be synchronized with high probability when setting randomly the undefined b-transitions. We follow the idea given at the beginning of Section 4 and first prove that with high enough probability, two states of $\mathbf{Cyc_g}(\mathcal{C})$ can be synchronized by a word of the form $b^j w_n$, for some integer $j \geq 0$.

Let q and r be two states of $\mathbf{Cyc_g}(\mathcal{C})$. By definition of $\mathcal{G}_n$, the states $q_2 = \delta_{bb}(q)$ and $r_2 = \delta_{bb}(r)$ have no outgoing b-transitions in $\mathcal{C}$. For $i \geq 3$, we iteratively build a sequence of uniform and independent pairs (q_i, r_i) of $[n] \times [n]$, and set $\delta_b(q_{i-1}) = q_i$ and $\delta_b(r_{i-1}) = r_i$. We stop this random process if at any step either $\delta_{w_n}(q_i) = \delta_{w_n}(r_i)$, or q_i (or r_i) already has a defined b-transition, or $i = \lambda_n$. By studying the probability that this random process halts because of the condition $\delta_{w_n}(q_i) = \delta_{w_n}(r_i)$, we obtain the following Lemma.

▶ **Lemma 12.** *Let $\mathcal{C} \in \mathcal{G}_n$ and let q and r be two distinct states of $\mathbf{Cyc_g}(\mathcal{C})$. If we add all the missing b-transitions to $\mathcal{C}$ by drawing them uniformly at random and independently, then the probability that for all $j \in \{0, \dots, \lambda_n\}$ we have $\delta_{b^j \cdot w_n}(r) \neq \delta_{b^j \cdot w_n}(r)$ is at most $n^{-3/8} \log^2 n$, for n sufficiently large.*

To conclude the proof of Theorem 3, we use the union bound: for any automaton $\mathcal{A}$ that extends an element of $\mathcal{G}_n$, which happens with high probability, there are less than γ_n^2 pairs of states in $\mathbf{Cyc_g}(\mathcal{A})$; the probability that one of these pairs (q, r) cannot be synchronized using a word of the form $b^j \cdot w_n$ is therefore at most $\gamma_n^2 \cdot n^{-3/8} \log^2 n$, which is $\mathcal{O}(n^{-\frac{1}{8}} \log^4 n)$.

To obtain the length of the synchronizing word, we apply Lemma 1 to the elements of $\mathbf{Cyc_g}(\mathcal{A})$: with high probability there are at most γ_n such states, which can be pairwise synchronized using words of the form $b^j w_n$, of length at most $|w_n| + \lambda_n$. Hence, the set $\mathbf{Cyc_g}(\mathcal{A})$ can be synchronized using a word z of length at most $(\gamma_n - 1)(|w_n| + \lambda_n)$, which is asymptotically equivalent to $n \log^3 n$. This conclude the proof, as $w_n z$ is synchronizing and has length which is also asymptotically equivalent to $n \log^3 n$.

5 Conclusion

In this article we proved that most complete automata are synchronizing and admit a synchronizing word of length $\mathcal{O}(n \log^3 n)$.

Our proof can be turned into an heuristic that try to find a short synchronizing word, which succeeds with high probability for uniform random automata: δ_{w_n} and $\mathbf{Cyc_g}(\mathcal{A})$ can be computed by just verifying some conditions on the height and cycle length of three mappings; once it is done, checking whether the property of Lemma 12 holds for every pair of elements of the image of δ_{w_n} can be achieved in sublinear time, as it is very small with high probability. Experiments seem to indicate that this algorithm behaves better in practice than its theoretical analysis: it looks like an important proportion of automata that fail to fulfill every step of our construction are still detected as synchronizing by the combination of computing δ_{w_n} and synchronizing the states of its image with the b^j's.

A natural continuation of this work is to prove that with high probability automata are synchronized by words that are way shorter than $n \log^3 n$. Experiments have been done [9], and seem to indicate that the expected length of the smallest synchronizing word is often sublinear, probably in $\Theta(\sqrt{n})$. There is plenty of room to improve our construction, as the synchronizing words we obtain have very specific shapes. It still might be quite difficult to match the bounds predicted in [9].

Acknowledgments. The author would like to thank Marie-Pierre Béal and Dominique Perrin for their interest in this work since the very beginning, and Mikhail Berlinkov for our fruitful discussions on this topic.

References

1 David Aldous, Grégory Miermont, and Jim Pitman. Brownian bridge asymptotics for random p-mappings. *Electron. J. Probab*, 9:37–56, 2004.

2 Mikhail V. Berlinkov. On the probability of being synchronizable. In Sathish Govindarajan and Anil Maheshwari, editors, *Algorithms and Discrete Applied Mathematics – Second International Conference, CALDAM 2016, Thiruvananthapuram, India, February 18-20, 2016, Proceedings*, volume 9602 of *Lecture Notes in Computer Science*, pages 73–84. Springer, 2016. doi:10.1007/978-3-319-29221-2_7.

3 Mikhail V. Berlinkov and Marek Szykula. Algebraic synchronization criterion and computing reset words. In Giuseppe F. Italiano, Giovanni Pighizzini, and Donald Sannella, editors, *Mathematical Foundations of Computer Science 2015 – 40th International Symposium, MFCS 2015, Milan, Italy, August 24-28, 2015, Proceedings, Part I*, volume 9234 of *Lecture Notes in Computer Science*, pages 103–115. Springer, 2015. doi:10.1007/978-3-662-48057-1_8.

4 Peter J Cameron. Dixon's theorem and random synchronization. *Discrete Mathematics*, 313(11):1233–1236, 2013.

5 Philippe Flajolet and Robert Sedgewick. *Analytic Combinatorics*. Cambridge University Press, 2009.

6 Peter Frankl. An extremal problem for two families of sets. *Eur. J. Comb.*, 3:125–127, 1982.

7 Bernard Harris. Probability distributions related to random mappings. *The Annals of Mathematical Statistics*, 31(4):1045–1062, 1960.

8 Richard M. Karp. The transitive closure of a random digraph. *Random Struct. Algorithms*, 1(1):73–94, 1990. doi:10.1002/rsa.3240010106.

9 Andrzej Kisielewicz, Jakub Kowalski, and Marek Szykula. A fast algorithm finding the shortest reset words. In Ding-Zhu Du and Guochuan Zhang, editors, *COCOON*, volume 7936 of *Lecture Notes in Computer Science*, pages 182–196. Springer, 2013. doi:10.1007/978-3-642-38768-5_18.

10 Valentin F. Kolčin. *Random Mappings: Translation Series in Mathematics and Engineering*. Translations series in mathematics and engineering. Springer London, Limited, 1986.

11 Cyril Nicaud. Random deterministic automata. In Erzsébet Csuhaj-Varjú, Martin Dietzfelbinger, and Zoltán Ésik, editors, *Mathematical Foundations of Computer Science 2014 – 39th International Symposium, MFCS 2014, Budapest, Hungary, August 25-29, 2014. Proceedings, Part I*, volume 8634 of *Lecture Notes in Computer Science*, pages 5–23. Springer, 2014. doi:10.1007/978-3-662-44522-8_2.

12 Jörg Olschewski and Michael Ummels. The complexity of finding reset words in finite automata. In Petr Hlinený and Antonín Kucera, editors, *MFCS*, volume 6281 of *Lecture Notes*

in Computer Science, pages 568–579. Springer, 2010. doi:10.1007/978-3-642-15155-2_50.

13 Jean-Eric Pin. On two combinatorial problems arising from automata thery. *Annals of Discrete Mathematics*, 17:535–548, 1983.

14 Jean-Jacques Quisquater and Joos Vandewalle, editors. *Advances in Cryptology – EUROCRYPT'89, Workshop on the Theory and Application of of Cryptographic Techniques, Houthalen, Belgium, April 10-13, 1989, Proceedings*, volume 434 of *Lecture Notes in Computer Science*. Springer, 1990.

15 Evgeny S. Skvortsov and Yulia Zaks. Synchronizing random automata. *Discrete Mathematics & Theoretical Computer Science*, 12(4):95–108, 2010.

16 J. Černý. Poznámka k. homogénnym experimentom s konecnymi automatmi. *Matematicko-fyzikalny Časopis Slovensk*, 14, 1964.

17 Mikhail V. Volkov. Synchronizing Automata and the Cerny Conjecture. In Carlos Martín-Vide, Friedrich Otto, and Henning Fernau, editors, *Language and Automata Theory and Applications, Second International Conference, LATA 2008, Tarragona, Spain, March 13-19, 2008. Revised Papers*, volume 5196 of *Lecture Notes in Computer Science*, pages 11–27. Springer, 2008. doi:10.1007/978-3-540-88282-4_4.

18 Yulia Zaks and Evgeny S. Skvortsov. Synchronizing random automata on a 4-letter alphabet. *Journal of Mathematical Sciences*, 192:303–306, 2013.