

Fast Synchronization of Random Automata*

Cyril Nicaud

LIGM, Université Paris-Est & CNRS,
5 bd Descartes, Champs-sur-Marne, 77454 Marne-la-Vallée Cedex 2, France
nicaud@univ-mlv.fr

April 29, 2014

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 size at most $(n - 1)^2$. Berlinkov recently made a breakthrough in the probabilistic analysis of synchronization by proving that with high probability, an automaton has a synchronizing word. In this article, we prove that with high probability an automaton admits a synchronizing word of length smaller than $n^{1+\epsilon}$, and therefore that the Černý conjecture holds with high probability.

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 [14]. Perhaps one of the reasons synchronizing automata are still intensively studied in theoretical computer science is a beautiful question asked by Černý [13] back in 1964: “*Does every synchronizing n -state automaton admits a synchronizing word of length at most $(n - 1)^2$?*” The bound of $(n - 1)^2$, as shown by Černý, is best possible. This question, known as *Černý’s conjecture*, is one of the most famous conjectures in automata theory. Though established for important subclasses of automata, Černý’s conjecture remains open in the general case. The best known upper bound, established in the early eighties [12, 6], is $\frac{1}{6}(n^3 - n)$. For a more detailed account on Černý’s conjecture, we refer the interested reader to Volkov’s article [14].

Probabilistic Černý conjecture. Considering Černý’s conjecture from a probabilistic point of view is natural (see for instance [3]), and leads to the following questions:

*This work is supported by the French National Agency (ANR) through ANR-10-LABX-58 and through ANR-2010-BLAN-0204.

Question 1: Is a random automaton synchronizing *with high probability*?

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

Here, *with high probability* means “with probability that tends to 1 as n goes to infinity”.

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

Question 2 can 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 finding the shortest reset word is hard [11]; the best experimental results we are aware of were obtained by Kisielewicz, Kowalski, and Szykula [9].

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 for any $\epsilon > 0$, the probability that a random n -state automaton has a synchronizing word of length smaller than $n^{1+\epsilon}$ tends to 1 when n goes to infinity. Of course, an immediate consequence is that if Černý’s conjecture is false, a counterexample will hardly be found by mere uniform random exploration.

Our proof also gives another way to show that automata are synchronizing with high probability, based on a completely different method: Berlinkov used advanced properties on the highest tree in a random mapping to study the *stable pairs* of the automaton, where we directly build words that iteratively shrink the set of states, using only basic discrete probabilities and variations on the probabilistic pigeonhole principle¹. The proof proposed by Berlinkov is arguably more complicated but also more precise since it gives the error term in $\mathcal{O}(n^{-\frac{1}{2}|A|})$ for the probability of not being synchronizing². There is little hope that the method presented below can be used to achieve such a precise estimation of the number of non-synchronizing automata.

2 Definitions and notations

Basic notations. For any integer $n \geq 1$, let $[n] = \{1, \dots, n\}$ be the set of integer from 1 to n . The cardinality of a finite set E is denoted by $|E|$.

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) mapping 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 an *a-transition* outgoing from p .

In this article, we are not interested in initial and final states since they do not matter for the synchronization. We will also focus on deterministic automata only, and therefore, throughout the article, we will simply call “automaton” a deterministic automaton with no initial and final states.

¹Also known as the Birthday Paradox.

²Knowing 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].

An automaton $\mathcal{A} = (Q, \delta)$ on A is classically seen as a labelled directed graph of set of vertices 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 $u \in A^*$, $\delta(p, ua) = \delta(\delta(p, a), u)$ 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 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 $\mathcal{A} \in \mathcal{H}$.

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 $\delta_w(q)$ exist and are 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 has a synchronizing word is said to be *synchronizing*.

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 $i \rightarrow j$ whenever $f(i) = j$. An example of such a graph is depicted in Figure 1 page 10.

If f is a mapping on E , $x \in E$ is a *cyclic point* of f (or *f-cyclic point* when there are several mappings) 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”: *f-cyclic state*, ...

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 the elements of E . The mapping depicted in Figure 1 page 10 has height 3, and the maximal height is reached by 9.

Probabilities. Let (E, s) be a pair where E is a set and s is a *size function* s from E to $\mathbb{Z}_{\geq 0}$. The pair (E, s) is a combinatorial set³ 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* (*whp*) for $(\mathbb{P}_n)_{n \geq 1}$ when $\mathbb{P}_n[P \text{ holds}] \rightarrow 1$ as $n \rightarrow \infty$.

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 *whp* 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.

Random mappings and random p -mappings. A *random mapping* of size

³The size is often clear in the context (number of nodes in a tree, ...) and can be omitted.

$n \geq 1$ 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 following the probability p .

A result stated as “a random p -mapping satisfies property P whp” 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 $(p_n)_{n \geq 1}$.

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 $|A|$, and we are therefore interested in the uniform distribution where each size- n complete automaton has probability $n^{-|A|n}$. Note 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 $a \in A$. This alternative way to look at random automata will be widely used in the sequel, especially in the following way: for a fixed incomplete automaton \mathcal{A} with n states, the uniform distribution on complete automata of $\mathbf{Ext}(\mathcal{A})$ can be seen as setting 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. The proofs are not new but given for completeness.

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 $\ell(n - 1)$.*

Proof: Assume we successfully synchronized i pairwise distinct states q_1, \dots, q_i using a word u of length smaller than or equal to $\ell(i - 1)$: for all $j, k \in \{1, \dots, i\}$, $\delta_u(x_j) = \delta_u(x_k)$. Let x_{i+1} be a state distinct from x_1, \dots, x_i and let v be a word of length at most ℓ that synchronizes $\delta_u(x_1)$ and $\delta_u(x_{i+1})$. Then the word uv synchronizes x_1, \dots, x_{i+1} and has length at most $\ell \cdot i$. The result follows by induction. \square

Random mappings and random p -mappings have been studied intensively in the literature [7, 4, 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 statements⁴, but we will only need the following results in the sequel⁵.

⁴For instance, limit distributions of some parameters [5] or even a notion of continuous limit for random mappings [1].

⁵The bound are not tight, we choose them for readability.

Lemma 2 *Let $\epsilon > 0$ be a fixed real number. For n large enough, the probability that a random p -mapping of size n has more than $n^{\frac{1}{2}+\epsilon}$ cyclic points or that its height is greater than $n^{\frac{1}{2}+\epsilon}$ is at most $\exp(-n^\epsilon)$.*

The proof of Lemma 2 consists in two steps. It is first established for uniform random mappings then extended to general p -random mappings using the following technical folklore lemma.

Lemma 3 *Let n and ℓ be two positive integers such that $\ell \leq n$. Let (E, \leq) be a totally ordered finite set of cardinality n . Let f be a map from E to $\mathbb{R}_{\geq 0}$, and denote by s the sum of the images by f : $s = \sum_{x \in E} f(x)$. The following result holds:*

$$\sum_{x_1 < x_2 < \dots < x_\ell} f(x_1)f(x_2) \cdots f(x_\ell) \leq \binom{n}{\ell} \left(\frac{s}{n}\right)^\ell,$$

where the sum range over all increasing ℓ -tuples of elements of E . The sum on the left is therefore maximal when $f(x) = \frac{s}{n}$, for every $x \in E$.

Proof: Let $\nu(f)$ denote the number of elements $x \in E$ such that $f(x)$ is different from $\frac{s}{n}$:

$$\nu(f) = \left| \left\{ x \in E : f(x) \neq \frac{s}{n} \right\} \right|.$$

We prove by induction on the value of ν that every map from E to $\mathbb{R}_{\geq 0}$ that sums up to s satisfies the inequality stated in the lemma.

► If $\nu(f) = 0$ then $f(x) = \frac{s}{n}$ for every $x \in E$. Thus, for any ℓ -tuple (x_1, \dots, x_ℓ) we have that

$$f(x_1) \cdots f(x_\ell) = \left(\frac{s}{n}\right)^\ell,$$

and therefore, since there are $\binom{n}{\ell}$ such increasing sequences,

$$\sum_{x_1 < \dots < x_\ell} f(x_1) \cdots f(x_\ell) = \binom{n}{\ell} \left(\frac{s}{n}\right)^\ell.$$

► Assume now that $\nu(f) \neq 0$. For the induction step, we build another map g , starting from f , such that g sums up to s , $\nu(g) < \nu(f)$ and

$$\sum_{x_1 < \dots < x_\ell} f(x_1) \cdots f(x_\ell) \leq \sum_{x_1 < \dots < x_\ell} g(x_1) \cdots g(x_\ell).$$

Let $y \in E$ be an element such that $|f(y) - \frac{s}{n}|$ is minimal amongst the y 's such that $f(y) \neq \frac{s}{n}$. We assume that $f(y) - \frac{s}{n} > 0$, the proof is almost the same if $f(y) - \frac{s}{n} < 0$. Since $f(y) > \frac{s}{n}$ and $\sum_{x \in E} f(x) = s$, there exists an element $z \neq y$ such that $f(z) < \frac{s}{n}$. Consider the new map g obtained from f by changing the value of y and z the following way:

$$\begin{cases} g(x) = f(x) \text{ if } x \neq y \text{ and } x \neq z, \\ g(y) = \frac{s}{n}, \\ g(z) = f(z) + f(y) - \frac{s}{n}. \end{cases}$$

It is direct to verify that g is always non-negative and sums up to s . Moreover, by construction $\nu(g) < \nu(f)$. We claim that

$$\sum_{x_1 < \dots < x_\ell} f(x_1) \cdots f(x_\ell) \leq \sum_{x_1 < \dots < x_\ell} g(x_1) \cdots g(x_\ell). \quad (1)$$

To prove this inequality, we consider three cases:

- If an increasing ℓ -tuple (x_1, \dots, x_ℓ) does not contain y nor z , then we have $g(x_1) \cdots g(x_\ell) = f(x_1) \cdots f(x_\ell)$.

- We sum the contributions of tuples containing exactly one of y or z : if $\{x_1, \dots, x_{\ell-1}\}$ are $\ell - 1$ elements of $[n] \setminus \{y, z\}$ we have

$$\begin{aligned} & g(x_1) \cdots g(x_{\ell-1}) g(y) + g(x_1) \cdots g(x_{\ell-1}) g(z) \\ &= g(x_1) \cdots g(x_{\ell-1}) (g(y) + g(z)) \\ &= f(x_1) \cdots f(x_{\ell-1}) \left(\frac{s}{n} + f(z) + f(y) - \frac{s}{n} \right) \\ &= f(x_1) \cdots f(x_{\ell-1}) f(y) + f(x_1) \cdots f(x_{\ell-1}) f(z). \end{aligned}$$

Hence the contributions of such tuples globally do not change the value of the sum when switching from f to g .

- If both y and z are in the tuple, then

$$\begin{cases} f(x_1) \cdots f(x_\ell) = f(y) f(z) \prod_{\substack{x_i \neq y \\ x_i \neq z}} f(x_i) \\ g(x_1) \cdots g(x_\ell) = g(y) g(z) \prod_{\substack{x_i \neq y \\ x_i \neq z}} g(x_i) = g(y) g(z) \prod_{\substack{x_i \neq y \\ x_i \neq z}} f(x_i). \end{cases}$$

Let α and β be the two positive real numbers defined by $\alpha = f(y) - \frac{s}{n}$ and $\beta = \frac{s}{n} - f(z)$. We have

$$g(z)g(y) = \frac{s}{n} \left(f(z) + f(y) - \frac{s}{n} \right) = \frac{s}{n} \left(\frac{s}{n} + \alpha - \beta \right) = \frac{s^2}{n^2} + \frac{s(\alpha - \beta)}{n},$$

whereas

$$f(y) f(z) = \left(\frac{s}{n} + \alpha \right) \left(\frac{s}{n} - \beta \right) = \frac{s^2}{n^2} + \frac{s(\alpha - \beta)}{n} - \alpha\beta,$$

therefore $f(y) f(z) \leq g(y) g(z)$ and $f(x_1) \cdots f(x_\ell) \leq g(x_1) \cdots g(x_\ell)$ for such a tuple.

This proves Equation (1) and concludes the proof by induction on the value of ν . \square

Proof of Lemma 2: We start with the uniform case.

► Number of cyclic points (uniform case): For any integer ℓ such that $1 \leq \ell \leq n$, the probability that there is a cyclic part of size ℓ in a uniform random mapping is at most $P(n, \ell) = \binom{n}{\ell} \ell! n^{-\ell}$, since we need to choose the ℓ elements that form the cyclic part, the way they are mapped to each other, and this forces ℓ images of the map which are correctly set with probability $\frac{1}{n}$ each. This is an upper bound since we do not prevent the formation of other cycles in our counting.

Moreover,

$$\begin{aligned} P(n, \ell) &= \frac{n!}{(n-\ell)!} n^{-\ell} = \left(1 - \frac{1}{n}\right) \left(1 - \frac{2}{n}\right) \cdots \left(1 - \frac{\ell-1}{n}\right) = \prod_{i=1}^{\ell-1} \left(1 - \frac{i}{n}\right) \\ &\leq \exp\left(-\sum_{i=1}^{\ell-1} \frac{i}{n}\right) \leq \exp\left(-\frac{\ell(\ell-1)}{2n}\right). \end{aligned}$$

Hence, the probability that there is a cyclic part of length greater than or equal to $n^{\frac{1}{2}+\epsilon}$ is at most, for n large enough,

$$\sum_{\ell=\lceil n^{\frac{1}{2}+\epsilon} \rceil}^n P(n, \ell) \leq n \cdot \exp\left(-\frac{\lceil n^{1/2+\epsilon} \rceil (\lceil n^{1/2+\epsilon} \rceil - 1)}{2n}\right) \leq \frac{1}{2} \exp(-n^\epsilon).$$

► **Height (uniform case):** Consider an element $i \in [n]$. For any integer ℓ such that $0 < \ell < n$, the probability that a uniform random mapping f on $[n]$ is such that $f(i), f^2(i) = f(f(i)), \dots, f^\ell(i)$ are all distinct is classically

$$\left(1 - \frac{1}{n}\right) \left(1 - \frac{2}{n}\right) \cdots \left(1 - \frac{\ell-1}{n}\right) = P(n, \ell).$$

We can therefore use the previous computations. If f has height greater than or equal to ℓ , then there exists a i with more than ℓ distinct iterates. Hence, using the union bound by summing the contribution of all $i \in [n]$, we get that the probability that f has height greater than $\lceil n^{1/2+\epsilon} \rceil$ is at most, for n large enough,

$$n \cdot P(n, \lceil n^{1/2+\epsilon} \rceil) \leq n \cdot \exp\left(-\frac{\lceil n^{1/2+\epsilon} \rceil (\lceil n^{1/2+\epsilon} \rceil - 1)}{2n}\right) \leq \frac{1}{2} \exp(-n^\epsilon).$$

This concludes the proof for the uniform case.

We now consider the case of random p -mappings.

► **Number of cyclic points (non-uniform case):** We start as in the proof for the uniform case. The difference is that if the points involved in the cyclic part of length ℓ are x_1, x_2, \dots, x_ℓ then the upper bound for the probability is not $\ell! n^{-\ell}$ anymore⁶ but

$$P_n(x_1, \dots, x_\ell) = \ell! p(x_1)p(x_2) \cdots p(x_\ell).$$

If we sum this quantity on all possible ℓ -subsets of $[n]$ we obtain an upper bound of

$$P_n(\ell) = \ell! \sum_{1 \leq x_1 < x_2 < \dots < x_\ell \leq n} p(x_1)p(x_2) \cdots p(x_\ell).$$

At this point we can apply Lemma 3 with $f = p$, $E = [n]$ and $s = 1$ to obtain a uniform bound for $P_n(\ell)$:

$$P_n(\ell) \leq \ell! \binom{n}{\ell} n^{-\ell},$$

⁶The $\ell!$ term is for counting the number of way to map bijectively the x_i 's to themselves, forming the cyclic part (or a part of it).

and this is the same bound as for the uniform case, yielding the same result.

► **Height (non-uniform case):** Similarly, we start with the same idea as in the proof for the uniform case. Fix some $x \in [n]$. Let (x_1, \dots, x_ℓ) be a ℓ -tuple of distinct elements in $[n] \setminus \{x\}$. The probability that a map f is such that $x_i = f^i(x)$, for all $1 \leq i \leq \ell$, is simply $p(x_1)p(x_2) \cdots p(x_\ell)$. Hence the probability that x has ℓ distinct iterates that are different from x when applying f is $p(x_1)p(x_2) \cdots p(x_\ell)$ where (x_1, \dots, x_ℓ) ranges over all ℓ -tuples of pairwise distinct elements of $[n] \setminus \{x\}$. We obtain an upper bound by allowing one of the x_i 's to be equal to x , which simplifies the writing; the bound is:

$$\ell! \sum_{1 \leq x_1 < x_2 < \dots < x_\ell \leq n} p(x_1)p(x_2) \cdots p(x_\ell),$$

since there are $\ell!$ ways to permute the x_i 's. This is the same quantity as $P_n(\ell)$ for the number of cyclic points just above, yielding the same result and concluding the proof. \square

4 Main Result

The main result of this article is the following theorem.

Theorem 4 *Let ϵ be a positive real number smaller than $\frac{1}{8}$, and 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 smaller than $n^{1+\epsilon}$ with high probability. More precisely, the probability it has no such word is in $\mathcal{O}(n^{-\frac{1}{8}+\epsilon})$.*

The statement does not hold for alphabets with only one letter, since there are cycles of length greater than 1 in a random mapping *whp* [4]: two distinct states in such a cycle cannot be synchronized.

As a consequence of Theorem 4, a random deterministic and complete automaton is synchronizing *whp*; our proof therefore constitutes an alternative proof of [2] for that property. Our statement is weaker since Berlinkov also obtained bounds in $\mathcal{O}(n^{-\frac{1}{2}|A|})$ 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 to shed a new light on the reasons why automata are often synchronizing.

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

Corollary 5 *For the uniform distribution on synchronizing deterministic and complete automata on an alphabet with at least two letters, Černý's conjecture holds with high probability.*

We prove Theorem 4 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 $n^{1/8+4\epsilon}$ *whp*. This is done by building a set \mathcal{G}_n of incomplete automata, *all* of which have the desired property,

and showing that a random n -state automaton extends an element of \mathcal{G}_n *whp*. 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 $n^{1/8+4\epsilon}$ states. This is done by showing that for a random automaton that extends an element of \mathcal{G}_n , any two among those $n^{1/8+4\epsilon}$ states can be synchronized by a word of the form $b^i w_n$ *whp*, with $i \leq n^{1/8+5\epsilon}$. Lemma 1 is then used to combine these words, and also w_n , into a synchronizing word for that automaton.

The remainder of this section is devoted to a more detailed proof of Theorem 4. 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.

From now on, we fix a real $\epsilon > 0$ small enough⁷. Since it is clearly sufficient to establish the result for a two-letter alphabet, we consider that $A = \{a, b\}$ in the sequel.

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 uniform 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 the number of states to be synchronized.

Let $\alpha_n = \lfloor n^{\frac{1}{2}+\epsilon} \rfloor$ 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 .

An example of an element of \mathcal{E}_n is given below.

Example 1 Let \mathcal{A} be a mapping with 18 states, which has only a -transitions and such that δ_a is the mapping of Figure 1.

The set $\mathbf{Cyc}_a(\mathcal{A})$ is made of the bold labels $\{\mathbf{2}, \mathbf{3}, \mathbf{7}, \mathbf{11}, \mathbf{13}, \mathbf{17}\}$. Assume that for our ϵ we have $\alpha_{18} = 6$, then $u_n = aaaaaa$ is used to start the synchronization:

$$\begin{aligned} \{2, 8, 14\} &\xrightarrow{u_n} \mathbf{2}; & \{3, 5, 12\} &\xrightarrow{u_n} \mathbf{3}; & \{6, 7, 9, 18\} &\xrightarrow{u_n} \mathbf{7}; \\ \{11\} &\xrightarrow{u_n} \mathbf{11}; & \{1, 10, 13, 15\} &\xrightarrow{u_n} \mathbf{13}; & \{4, 16, 17\} &\xrightarrow{u_n} \mathbf{17}. \end{aligned}$$

Since there are $6 \leq \alpha_{18}$ cyclic states and since the height of the mapping is $3 \leq \alpha_{18}$, \mathcal{A} is in \mathcal{E}_n . ◊

⁷To simplify the writing, we will prove that there is a synchronizing word of length $\mathcal{O}(n^{1+11\epsilon})$ *whp*, and then get the statement of Theorem 4 by changing ϵ into $\epsilon/13$.

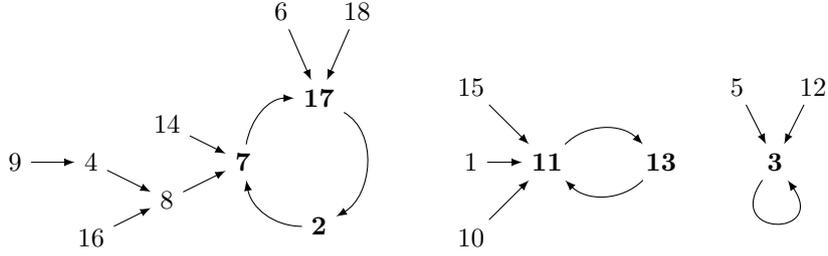


Figure 1: This mapping represents the action of a in the automaton \mathcal{A} .

By independency, the action of letter a in a uniform random complete automaton is exactly a uniform random mapping, yielding the following consequence of Lemma 2.

Lemma 6 *A random complete automaton with n states extends an element of \mathcal{E}_n whp. More precisely, the probability that such an automaton does not extend an element of \mathcal{E}_n is at most $\exp(-n^\epsilon)$.*

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} .

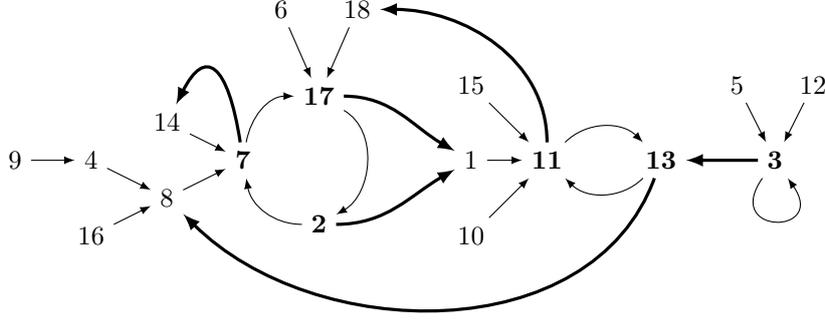
Let $u_n = a^{\alpha_n}$. By Lemma 6, we can already start the synchronization using u_n , since whp the image of the set of states $[n]$ by δ_{u_n} is included in $\mathbf{Cyc}_a(\mathcal{A})$, which has size at most α_n . In the sequel, we therefore work on synchronizing the elements of $\mathbf{Cyc}_a(\mathcal{A})$.

4.2 Adding random b -transitions that start from the δ_a -cyclic states

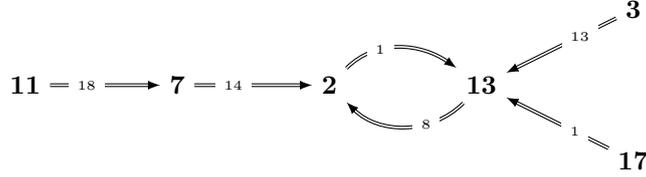
Let \mathcal{A} be a fixed element of \mathcal{E}_n . We are now working on $\mathbf{Ext}(\mathcal{A})$ and we consider the process of adding a random b -transition starting from every state of $\mathbf{Cyc}_a(\mathcal{A})$.

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 2 This is the automaton of Example 1, where the b -transitions that start from the elements of $\mathbf{Cyc}_a(\mathcal{A})$ have been added (in bold):



Below is depicted the map $f_{\mathcal{B}}$, which is the restriction of δ_{bu_n} to $\mathbf{Cyc}_{\mathbf{a}}(\mathcal{A})$; an edge $\mathbf{p} = x \implies \mathbf{q}$ 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 $\mathbf{Cyc}_{\mathbf{a}}(\mathcal{A})$, the induced distribution for the mapping $f_{\mathcal{B}}$ is usually not the uniform distribution on mappings of $\mathbf{Cyc}_{\mathbf{a}}(\mathcal{A})$. More precisely, for any $q \in \mathbf{Cyc}_{\mathbf{a}}(\mathcal{A})$ the probability that the image by $f_{\mathcal{B}}$ of an element of $\mathbf{Cyc}_{\mathbf{a}}(\mathcal{A})$ 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

$$\mathbb{P}_{\mathcal{A}, \omega}(q) = \frac{|\delta_{\omega}^{-1}(\{q\})|}{n}, \text{ for all } q \in [n]. \quad (2)$$

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

Let $\beta_n = \lfloor n^{\frac{1}{4} + 2\epsilon} \rfloor$. Applying Lemma 2 to $f_{\mathcal{B}}$ yields the following result.

Lemma 7 *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 $\mathbf{Cyc}_{\mathbf{a}}(\mathcal{A})$, choosing the target uniformly and independently in $[n]$. The probability that $f_{\mathcal{B}}$ has more than β_n cyclic states or that it is of height greater than β_n is at most $\exp(-n^{\epsilon/4})$.*

Proof: Let c denote the number of δ_a -cyclic states in \mathcal{A} . Recall that $f_{\mathcal{B}}$ is the map from $\mathbf{Cyc}_{\mathbf{a}}(\mathcal{A})$ to itself defined by $f_{\mathcal{B}}(x) = \delta_{bu_n}(x)$, for every $x \in \mathbf{Cyc}_{\mathbf{a}}(\mathcal{A})$. We consider two cases:

- If $c \leq \beta_n$ then there is nothing to prove as $\mathbf{Cyc}_{\mathbf{a}}(\mathcal{A})$ already has at most β_n states. The probability is therefore 0 in this case.
- If $c > \beta_n$, we can apply Lemma 2 when n is large enough, since $f_{\mathcal{B}}$ is a p -random mapping on a set with c elements. Therefore, the probability that

there are more than $c^{\frac{1}{2}+\epsilon}$ cyclic points or that the height is greater than $c^{\frac{1}{2}+\epsilon}$ is at most $\exp(-c^\epsilon)$. Moreover, observe that since $\mathcal{A} \in \mathcal{E}_n$,

$$c^{\frac{1}{2}+\epsilon} \leq \alpha_n^{\frac{1}{2}+\epsilon} \leq n^{\frac{1}{4}+\epsilon+\epsilon^2} \leq \beta_n.$$

This concludes the proof, as $c \geq n^{\frac{1}{4}}$ gives the announced upper bound. \square

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 $n^{\frac{1}{4}+2\epsilon}$ *whp*, since the image of δ_{v_n} is included in $\mathbf{Cyc}_f(\mathcal{B})$, which has size at most β_n . It has been achieved by generating all the a -transitions, but using only the b -transitions that start from the δ_a -cyclic states: *whp*, there still are at least $n - \alpha_n$ unset b -transitions that can be used to continue the synchronization. Nonetheless, before going on we will 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

We have reduced the number of states to be synchronized to no more than $n^{1/4+2\epsilon}$ states *whp*, but this quantity is still too large for the idea used at the end of the proof, we need to shrink this set once more. For an alphabet with at least one more letter c , we could use the same kind of construction as in Section 4.2, considering the restriction of δ_{cv_n} to the cyclic states of $f_{\mathcal{B}}$ and would obtain less than roughly $n^{1/8}$ states to be synchronized. 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:

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 issue 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 random⁸.

On the other hand, if for every $q \in \mathbf{Cyc}_{\mathbf{f}}(\mathcal{B})$, $\delta_b(q) \notin \mathbf{Cyc}_{\mathbf{a}}(\mathcal{B})$, then it is easy to verify that $g_{\mathcal{C}}$ is a random p -mapping: the image of $q \in \mathbf{Cyc}_{\mathbf{f}}(\mathcal{B})$ 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 (2).

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 $\mathbf{Cyc}_{\mathbf{a}}(\mathcal{B})$,
3. the map $f_{\mathcal{B}}$ has height at most β_n and has at most β_n cyclic states,
4. for every $q \in \mathbf{Cyc}_{\mathbf{f}}(\mathcal{B})$, $\delta_b(q) \notin \mathbf{Cyc}_{\mathbf{a}}(\mathcal{B})$.

Example 3 Assume that for our ϵ , $\beta_{18} = 3$. The automaton of Example 2 is in \mathcal{F}_n : looking at the mapping $f_{\mathcal{B}}$, we can see that the $f_{\mathcal{B}}$ -cyclic states are **2** and **13**, and their images by δ_b , which are 1 and 8 respectively, are not in $\mathbf{Cyc}_{\mathbf{a}}(\mathcal{B})$. The fact that $\delta_b(\mathbf{3})$ is in $\mathbf{Cyc}_{\mathbf{a}}(\mathcal{B})$ is not a problem, since **3** is not a $f_{\mathcal{B}}$ -cyclic state. \diamond

If we forget the last condition in the definition of \mathcal{F}_n , the other requirements hold *whp* for every fixed $\mathcal{A} \in \mathcal{E}_n$, as a consequence of Lemma 7. Lemma 8 below states that after our additional restriction, we still have a set large enough.

Lemma 8 *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 $2n^{-1/4+3\epsilon}$, for n large enough.*

Proof: Fix $\mathcal{A} \in \mathcal{E}_n$, and consider the extensions \mathcal{B} of \mathcal{A} obtained by adding b -transitions to the δ_a -cyclic states. A state x of $\mathbf{Cyc}_{\mathbf{f}}(\mathcal{B})$ is a *bad state* when there exists $y \in \mathbf{Cyc}_{\mathbf{f}}(\mathcal{B})$ and $z \in \mathbf{Cyc}_{\mathbf{a}}(\mathcal{A})$ such that $y \xrightarrow{b} z$ and $z \xrightarrow{u_n} x$. In such a case, y is the cyclic predecessor of x for the mapping $f_{\mathcal{B}}$ and it does not satisfy the last condition of \mathcal{F}_n 's definition. Clearly, if \mathcal{B} is not in \mathcal{F}_n then Condition 3 is not satisfied or there is at least one bad state in \mathcal{B} .

For a given $x \in \mathbf{Cyc}_{\mathbf{a}}(\mathcal{A})$ and $\ell \in \{0, \dots, n-1\}$ let us bound from above the probability that x is a bad state and in a $f_{\mathcal{B}}$ -cycle of length $\ell + 1$ when adding the b -transitions: there must exist ℓ distinct states x_1, \dots, x_{ℓ} of $\mathbf{Cyc}_{\mathbf{a}}(\mathcal{A})$, all distinct from x , such that $x \xrightarrow{f_{\mathcal{B}}} x_1$, $x_1 \xrightarrow{f_{\mathcal{B}}} x_2$, \dots , $x_{\ell} \xrightarrow{f_{\mathcal{B}}} x$, and the image by b of x_{ℓ} must be in $\mathbf{Cyc}_{\mathbf{a}}(\mathcal{A})$. Hence $\delta_b(x_{\ell})$ must belong to $\mathbf{Cyc}_{\mathbf{a}}(\mathcal{A}) \cap \delta_{u_n}^{-1}(\{x\})$. Consequently, the probability that such a cycle exists is

$$\mathbb{P}_{\mathcal{A}, u_n}(x_1) \mathbb{P}_{\mathcal{A}, u_n}(x_2) \cdots \mathbb{P}_{\mathcal{A}, u_n}(x_{\ell}) \cdot \frac{|\mathbf{Cyc}_{\mathbf{a}}(\mathcal{A}) \cap \delta_{u_n}^{-1}(\{x\})|}{n}.$$

⁸Except in the very degenerate case where the restriction of δ_{bb} to $\mathbf{Cyc}_{\mathbf{f}}(\mathcal{B})$ is already a totally defined and constant map in \mathcal{B} .

We sum this quantity for every possible tuple (x_1, \dots, x_ℓ) of distinct elements of $E_x = [n] \setminus \{x\}$. We obtain, since there are $\ell!$ ways to order each $\{x_1, \dots, x_\ell\}$:

$$\begin{aligned} \ell! \sum_{\substack{x_1 < \dots < x_\ell \\ x_i \in E_x}} \mathbb{P}_{\mathcal{A}, u_n}(x_1) \mathbb{P}_{\mathcal{A}, u_n}(x_2) \cdots \mathbb{P}_{\mathcal{A}, u_n}(x_\ell) \cdot \frac{|\mathbf{Cyc}_{\mathbf{a}}(\mathcal{A}) \cap \delta_{u_n}^{-1}(\{x\})|}{n} \\ \leq \frac{|\mathbf{Cyc}_{\mathbf{a}}(\mathcal{A}) \cap \delta_{u_n}^{-1}(\{x\})|}{n} \ell! \binom{n-1}{\ell} \left(\frac{1 - \mathbb{P}_{\mathcal{A}, u_n}(x)}{n-1} \right)^\ell, \end{aligned}$$

by applying Lemma 3 with $f = \mathbb{P}_{\mathcal{A}, u_n}$, $E = E_x$ and $s = 1 - \mathbb{P}_{\mathcal{A}, u_n}(x)$. But an easy computation shows that $\ell! \binom{n-1}{\ell} \leq (n-1)^\ell$, hence the probability that x is a bad state in a $f_{\mathcal{B}}$ -cycle of length $\ell + 1$ is at most $\frac{1}{n} |\mathbf{Cyc}_{\mathbf{a}}(\mathcal{A}) \cap \delta_{u_n}^{-1}(\{x\})|$.

We now use the union bound and sum the contribution of all $x \in \mathbf{Cyc}_{\mathbf{a}}(\mathcal{A})$. Since the $\delta_{u_n}^{-1}(\{x\})$ are disjoint, we obtained that the probability that there is a bad state in a cycle of length $\ell + 1$ is at most $\frac{1}{n} |\mathbf{Cyc}_{\mathbf{a}}(\mathcal{A})|$. Hence, since $\mathcal{A} \in \mathcal{E}_n$, this probability is at most $n^{-1/2+\epsilon}$.

By Lemma 7, the probability that Condition 3 of the definition of \mathcal{F}_n is not satisfied is smaller than $\exp(-n^{\epsilon/4})$. Hence, for every fixed $\mathcal{A} \in \mathcal{E}_n$, the probability that there is a bad state or that Condition 3 does not hold is at most

$$\underbrace{\sum_{\ell=0}^{\lceil n^{1/4+2\epsilon} \rceil - 1} \frac{n^{\frac{1}{2}+\epsilon}}{n}}_{\text{bad state for typical case}} + \underbrace{\exp(-n^{\epsilon/4})}_{\text{Condition 3 does not hold}} \leq \frac{3}{2} \cdot n^{-1/4+3\epsilon}.$$

Note that we do not need to consider the cases where $\ell + 1 > n^{1/4+2\epsilon}$ in the first sum, since they do not satisfy Condition 3.

We therefore obtained a uniform upper bound of $\frac{3}{2} \cdot n^{-1/4+3\epsilon}$ for every $\mathcal{A} \in \mathcal{E}_n$. Since a complete automaton can extend at most one element of \mathcal{E}_n , the law of total probabilities applies and yields that: the probability that a complete automaton with n states that extends an element of \mathcal{E}_n does not satisfy Condition 3 or Condition 4 is at most $\frac{3}{2} \cdot n^{-1/4+3\epsilon}$, for n large enough. This concludes the proof since the probability of not being in \mathcal{E}_n is smaller than $\exp(-n^\epsilon)$. \square

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 $\mathbf{Cyc}_{\mathbf{f}}(\mathcal{B})$. For such an extension \mathcal{C} of \mathcal{B} , recall that the mapping $g_{\mathcal{C}}$ is the restriction of δ_{bbv_n} to $\mathbf{Cyc}_{\mathbf{f}}(\mathcal{B})$. Let $\mathbf{Cyc}_{\mathbf{g}}(\mathcal{C})$ 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 $\mathbf{Cyc}_{\mathbf{f}}(\mathcal{B})$, which are all distinct since two distinct states of $\mathbf{Cyc}_{\mathbf{f}}(\mathcal{B})$ cannot have the same image by δ_b .

Let $\gamma_n = \lfloor n^{1/8+4\epsilon} \rfloor$ and let \mathcal{G}_n denote the set of incomplete automata \mathcal{C} with n states such that:

1. \mathcal{C} extends an automaton \mathcal{B} of \mathcal{F}_n ,
2. if $X_{\mathcal{B}}$ denote the set of images of $\mathbf{Cyc}_{\mathbf{f}}(\mathcal{B})$ by δ_b , i.e. $X_{\mathcal{B}} = \{\delta_b(x) : x \in \mathbf{Cyc}_{\mathbf{f}}(\mathcal{B})\}$, then the only b -transitions of \mathcal{C} are those starting from $\mathbf{Cyc}_{\mathbf{a}}(\mathcal{B})$ and from $X_{\mathcal{B}}$;

3. the map g_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 here for the same kind of reasons than the last condition of \mathcal{F}_n : it is used to ensure some independency for the final step of the synchronization.

Lemma 9 *A random complete automaton with n states extends an element of \mathcal{G}_n whp. More precisely, the probability it does not is at most $2n^{-1/4+3\epsilon}$.*

Proof: We only explain why the last condition holds *whp*, since one can easily prove that the probability that Condition 3 does not hold is at most $\exp(-n^{\epsilon/8})$ using the same technique as for Lemma 7.

We even prove the stronger result that *whp*, for every $q \in \mathbf{Cyc}_f(\mathcal{B})$ the b -transition of $\delta_{bb}(q)$ is undefined. Once \mathcal{C} is built by adding the needed transitions, the defined b -transitions start from $\mathbf{Cyc}_a(\mathcal{B})$ or from $X_{\mathcal{B}}$. Since $\mathcal{B} \in \mathcal{F}_n$, the cardinality of $\mathbf{Cyc}_a(\mathcal{B})$ and $X_{\mathcal{B}}$ are at most α_n and β_n , respectively. To build \mathcal{C} , we iteratively add new random outgoing b -transitions for each element of $X_{\mathcal{B}}$; when we add the i -th such transition, the probability it ends in a state that has a defined outgoing b -transition is therefore at most $\frac{1}{n}(\alpha_n + \beta_n + i - 1)$, which is smaller than $p_n = \frac{4}{3}n^{-1/2+\epsilon}$ for n large enough, as $i \leq |X_{\mathcal{B}}| \leq \beta_n$. Hence the probability that Condition 4 holds is at least $(1 - p_n)^{|X_{\mathcal{B}}|}$, which is greater than $1 - \frac{5}{3}n^{-1/4+3\epsilon}$ for n large enough, by basic computations. This concludes the proof after handling the probability that the other conditions do not hold. \square

Let $w_n = v_n(bbv_n)^{\gamma_n}$. Lemma 9 ensures that in a random complete automaton \mathcal{A} , the image of δ_{w_n} is included in $\mathbf{Cyc}_g(\mathcal{A})$, which has size at most $n^{1/8+4\epsilon}$. This concludes the first part of the synchronization: the word w_n maps the set of states of \mathcal{A} to the much smaller set of states $\mathbf{Cyc}_g(\mathcal{A})$ *whp*. We will use another technique to finalize the synchronization, which only works because $\mathbf{Cyc}_g(\mathcal{A})$ is small enough *whp*.

4.5 Synchronizing the states of $\mathbf{Cyc}_g(\mathcal{C})$

Let $\lambda_n = \lfloor n^{1/8+5\epsilon} \rfloor$ 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 *whp*, using the remaining randomness of 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 $b^j w_n$.

Lemma 10 *Let $\mathcal{C} \in \mathcal{G}_n$ and let p and q be in $\mathbf{Cyc}_g(\mathcal{C})$. 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 $\delta_{b^j \cdot w_n}(p) \neq \delta_{b^j \cdot w_n}(q)$ is at most $6n^{-3/8+5\epsilon}$, for n large enough.*

Proof: By definition of \mathcal{G}_n , the states $p_2 = \delta_{bb}(p)$ and $q_2 = \delta_{bb}(q)$ have no outgoing b -transitions. If $p_2 = q_2$, then p and q does not satisfy the property for $j = 2$. Otherwise we consider the sequence of pairs of states (p_i, q_i) that is generated using the following random process, starting from $i = 3$:

1. generate (p_i, q_i) uniformly at random and set $\delta_b(p_{i-1}) = p_i$ and $\delta_b(q_{i-1}) = q_i$ in the automaton;
2. if $\delta_{w_n}(p_i) = \delta_{w_n}(q_i)$ then stop the process and return a **success**⁹ (and the property does not hold for $j = i$);
3. otherwise, if p_i or q_i already have an outgoing b -transition, stop the process and return a **failure**;
4. in other cases, iterate the process for the next value of i by going back to step 1, until $i = \lambda_n$. When $i = \lambda_n$, the process halts and return a **failure**.

Hence we iteratively and in parallel create a sequence of b -transitions, starting from p_2 and q_2 . If the process returns a **success**, then clearly the property of the statement does not hold. Thus the probability of returning a **failure** is an upper bound for the probability that it holds.

Given that the process has not halted after building up to (p_{i-1}, q_{i-1}) for $1 \leq i < \lambda_n$, the probability that it halts at next step and returns a **success** is the probability that two randomly chosen elements of $[n]$ have the same image by δ_{w_n} , which is exactly

$$s_i = \sum_{x \in \mathbf{Cyc}_{\mathbf{g}}(C)} \mathbb{P}_{\mathcal{C}, w_n}(x)^2.$$

Therefore $s_i \geq \frac{1}{|\mathbf{Cyc}_{\mathbf{g}}(C)|} \geq n^{-1/8-4\epsilon}$ by Cauchy-Schwarz inequality.

Let $Y_{\mathcal{C}}$ denote the set of states of \mathcal{C} that have a defined b -transition. Since $\mathcal{C} \in \mathcal{G}_n$, the only b -transitions of \mathcal{C} start from elements of $\mathbf{Cyc}_{\mathbf{a}}(C)$ and from their images by b , and thus $|Y_{\mathcal{C}}| \leq 2\alpha_n$. Hence, given that the process has not halted after building up to (p_{i-1}, q_{i-1}) for $1 \leq i < \lambda_n$, the probability that it halts at next step and returns a **failure** only depends on i and satisfies

$$f_i \leq 1 - \left(1 - \frac{|Y_{\mathcal{C}}| + 2(i-3)}{n}\right)^2,$$

because it is not a failure when both p_i and q_i are not in $Y_{\mathcal{C}}$ and are not one of the $2(i-3)$ states that get a b -transition during the previous iterations of the process. This is an upper bound, since some cases yield a **success** (for instance when $p_i = q_i$). Therefore,

$$\begin{aligned} f_i &\leq 1 - \left(1 - \frac{|Y_{\mathcal{C}}| + 2(i-3)}{n}\right)^2 = 2 \frac{|Y_{\mathcal{C}}| + 2(i-3)}{n} - \frac{(|Y_{\mathcal{C}}| + 2(i-3))^2}{n^2} \\ &\leq 2 \frac{|Y_{\mathcal{C}}| + 2i}{n} \leq 2 \frac{|Y_{\mathcal{C}}| + 2\lambda_n}{n} \leq 5n^{-\frac{1}{2} + \epsilon}, \end{aligned}$$

for n large enough. Note that this bound does not depend on i , and we can therefore bound the probability of a failure given that the process halts when building (p_i, q_i) by

$$\mathbb{P}(\text{failure at step } i \mid \text{has not halted before}) = \frac{f_i}{f_i + s_i} \leq \frac{f_i}{s_i} \leq 5n^{-\frac{3}{8} + 5\epsilon}.$$

⁹We call it a **success** because we have successfully synchronized p and q .

The probability that the process halts at a given step is greater than the probability it halts and returns a **success**. Hence, the probability that the process has build $(p_{\lambda_n}, q_{\lambda_n})$ without halting is smaller than or equal to

$$\prod_{i=3}^{\lambda_n} (1 - s_i) \leq \left(1 - n^{-1/8-4\epsilon}\right)^{\lambda_n-2} = \mathcal{O}(\exp(-n^\epsilon)).$$

Putting all together, we get that the probability of a failure is at most

$$5n^{-\frac{3}{8}+5\epsilon} + \mathcal{O}(\exp(-n^\epsilon)) \leq 6n^{-\frac{3}{8}+5\epsilon},$$

for n large enough, which concludes the proof. \square

To conclude the proof of Theorem 4, we use the union bound: for any automaton \mathcal{A} that extends an element of \mathcal{G}_n , which happens *whp*, there are less than γ_n^2 pairs of states in $\mathbf{Cyc}_{\mathbf{g}}(\mathcal{A})$; the probability that one of these pairs (p, q) cannot be synchronized using a word of the form $b^j \cdot w_n$ is therefore at most $\gamma_n^2 \cdot 4n^{-3/8+5\epsilon}$, which tends to 0; more precisely, it is in $\mathcal{O}(n^{-\frac{1}{8}+13\epsilon})$.

To obtain the length of the synchronizing word, we apply Lemma 1 to the elements of $\mathbf{Cyc}_{\mathbf{g}}(\mathcal{A})$: *whp* 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 $\mathbf{Cyc}_{\mathbf{g}}(\mathcal{A})$ can be synchronized using a word z of length at most $(\gamma_n - 1)(|w_n| + \lambda_n)$, which is asymptotically equivalent to $n^{1+11\epsilon}$. To conclude, observe that $w_n z$ is synchronizing, and also of length asymptotically equivalent to $n^{1+11\epsilon}$. Changing ϵ into $\epsilon/13$ yields the result.

5 Conclusion

In this article we proved that most complete automata are synchronizing, since they admit a synchronizing word of length smaller than $n^{1+\epsilon}$ with high probability.

Note that our proof can be turned into a probabilistic algorithm to try to quickly find a synchronizing word: Compute the action of δ_{u_n} , δ_{v_n} and then δ_{w_n} in linear time. Once it is done, check whether the property of Lemma 10 holds for every pair of elements of the image of δ_{w_n} , which is small with high probability. Experiments seem to indicate that the algorithm behaves way 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 $n^{1+\epsilon}$. Experiments have been done [9], and seem to indicate that the expected length of the smallest synchronizing word is often sublinear, probably in \sqrt{n} . There is plenty of room to improve our construction, as the synchronizing words we obtain have very specific shapes, but it might be quite difficult to have a proof that matches what was observed during the experiments of [9].

Acknowledgments: the author would like to thank Marie-Pierre Béal and Dominique Perrin for their interest in this work since the very beginning.

References

- [1] D. Aldous, G. Miermont, and J. Pitman. Brownian bridge asymptotics for random p-mappings. *Electron. J. Probab.*, 9:37–56, 2004.
- [2] M. V. Berlinkov. On the probability to be synchronizable. *arXiv*, abs/1304.5774, 2013. <http://arxiv.org/abs/1304.5774>.
- [3] P. J. Cameron. Dixon’s theorem and random synchronization. *Discrete Mathematics*, 313(11):1233–1236, 2013.
- [4] P. Flajolet and A. M. Odlyzko. Random mapping statistics. In *EUROCRYPT*, volume 434 of *LNCS*, pages 329–354. Springer, 1989.
- [5] P. Flajolet and R. Sedgewick. *Analytic Combinatorics*. Cambridge University Press, 2009.
- [6] P. Frankl. An extremal problem for two families of sets. *Eur. J. Comb.*, 3:125–127, 1982.
- [7] B. Harris. Probability distributions related to random mappings. *The Annals of Mathematical Statistics*, 31(4):1045–1062, 1960.
- [8] R. M. Karp. The transitive closure of a random digraph. *Random Struct. Algorithms*, 1(1):73–94, 1990.
- [9] A. Kisielewicz, J. Kowalski, and M. Szykula. A fast algorithm finding the shortest reset words. In D.-Z. Du and G. Zhang, editors, *COCOON*, volume 7936 of *Lecture Notes in Computer Science*, pages 182–196. Springer, 2013.
- [10] V. Kolčín. *Random Mappings: Translation Series in Mathematics and Engineering*. Translations series in mathematics and engineering. Springer London, Limited, 1986.
- [11] J. Olschewski and M. Ummels. The complexity of finding reset words in finite automata. In P. Hliněný and A. Kucera, editors, *MFCS*, volume 6281 of *Lecture Notes in Computer Science*, pages 568–579. Springer, 2010.
- [12] J.-E. Pin. On two combinatorial problems arising from automata theory. *Annals of Discrete Mathematics*, 17:535–548, 1983.
- [13] J. Černý. Poznámka k. homogénnym experimentom s konečnými automatmi. *Matematicko-fyzikálny Časopis Slovensk*, 14, 1964.
- [14] M. V. Volkov. Synchronizing automata and the Černý conjecture. In *LATA’08*, volume 5196 of *LNCS*, pages 11–27. Springer, 2008.