

RANDOM DETERMINISTIC AUTOMATA WITH ONE ADDED TRANSITION

ARNAUD CARAYOL ^a, PHILIPPE DUCHON ^b, FLORENT KOECHLIN ^c,
AND CYRIL NICAUD ^a

^a Univ Gustave Eiffel, CNRS UMR 8049, LIGM, F-77454 Marne-la-Vallée, France
e-mail address: arnaud.carayol@univ-eiffel.fr, cyril.nicaud@univ-eiffel.fr

^b Univ. Bordeaux, CNRS UMR 5800, LaBRI, F-33400 Talence, France
e-mail address: duchon@labri.fr

^c Univ. Sorbonne Paris Nord, LIPN, CNRS UMR 7030, F-93430 Villetaneuse, France
e-mail address: koechlin@lipn.fr

ABSTRACT. Every language recognized by a non-deterministic finite automaton can be recognized by a deterministic automaton, at the cost of a potential increase of the number of states, which in the worst case can go from n states to 2^n states. In this article, we investigate this classical result in a probabilistic setting where we take a deterministic automaton with n states uniformly at random and add just one random transition. These automata are almost deterministic in the sense that only one state has a non-deterministic choice when reading an input letter. In our model, each state has a fixed probability to be final. We prove that for any $d \geq 1$, with non-negligible probability the minimal (deterministic) automaton of the language recognized by such an automaton has more than n^d states; as a byproduct, the expected size of its minimal automaton grows faster than any polynomial. Our result also holds when each state is final with some probability that depends on n , as long as it is not too close to 0 and 1, at distance at least $\Omega(\frac{1}{\sqrt{n}})$ to be precise, therefore allowing models with a sublinear number of final states in expectation.

1. INTRODUCTION

A fundamental result in automata theory is that deterministic complete finite state automata recognize the same languages as non-deterministic finite state automata. This result can be established using the classical (accessible) subset construction [RS59, MF71, HU79]: starting with a non-deterministic automaton with n states, one can build a deterministic automaton with at most 2^n states that recognizes the same language. This upper bound is tight; there are regular languages recognized by an n -state non-deterministic automaton whose minimal automaton, i.e. the smallest deterministic and complete automaton that recognizes the language, has 2^n states. The number of states of the minimal automaton of a regular language is called its *state complexity*. Figure 1 shows two n -state non-deterministic

Key words and phrases: Random finite automata, average case complexity, discrete probabilities.

This article is the extended and revised version of the article [CDKN23] published in the proceedings of the conference STACS'23 .

automata with somewhat similar shape, but whose languages \mathcal{L}_ℓ and \mathcal{L}_r have very different state complexities. In both automata, there is only one non-deterministic choice, at the initial state.



Figure 1: On the left, a non-deterministic automaton with n states recognizing the language $\mathcal{L}_\ell = \Sigma^* a \Sigma^{n-2}$. On the right, a non-deterministic automaton with n states recognizing the language $\mathcal{L}_r = \Sigma^* a^{n-1}$. The minimal automaton of \mathcal{L}_ℓ has 2^{n-1} states, whereas the one of \mathcal{L}_r has n states.

In this article, we address the following question: if we take a random n -state deterministic automaton and add just one random transition, what can be said about the state complexity of the resulting recognized language? Does it hugely increase as for \mathcal{L}_ℓ , or does it remain small as for \mathcal{L}_r ?

From [BDS12], we know that with high probability, the state complexity of the language recognized by a n -state deterministic automaton taken uniformly at random is linear. This is important as it implies that the corresponding distribution on regular languages is not degenerated: this contrasts with the case of random regular expressions where the expected state complexity of the described regular languages is constant [KNR21], which means that the induced distribution on regular languages is concentrated on a finite number of languages.

To be more precise, our formal setting in this article is the following. Let $\Sigma = \{a, b, \dots\}$ be a finite alphabet with $k \geq 2$ letters. For any $n \geq 1$, we consider the uniform distribution on deterministic and complete automata on Σ , with $\{1, \dots, n\}$ as their set of states; the initial state is picked uniformly at random, and the action of the letters on the set of states are k uniform and independent random mappings. We also pick uniformly at random and independently two states p and q , and add a transition $p \xrightarrow{a} q$, if it is not already there. Finally each state is final with a given fixed probability $f \in (0, 1)$, independently. Hence in this model, the expected number of final states of an almost deterministic automaton is $f \times n$. Our results still hold if we allow the probability f of being final to depend on the size n of the automaton, provided that f_n has a distance to 0 and 1 in $\Omega(\frac{1}{\sqrt{n}})$. This allows us to consider probabilistic models in which random automata have an expected number of final states that is as low as $\Theta(\sqrt{n})$.

Our main result is that for any $d \geq 1$ there exists a constant $c_d > 0$ such that the state complexity of the language of such a random almost deterministic automaton is greater than n^d with probability at least c_d , for n sufficiently large. That is, for any polynomial P , there is a non-negligible probability that the state complexity of the language of a random automaton is greater than $P(n)$: we will say that the state complexity is *super-polynomial* with *visible probability*. As a direct consequence, the expected state complexity is super-polynomial.

It should be noted that with the same random models for deterministic automata, one cannot hope to replace visible probability in our results with a probability that converges to 1 (i.e., with high probability). Indeed random automata have, with high probability, a constant fraction of states that are not accessible from the initial state [Gru73]; if the source

of the added transition is not accessible from the initial state, the added transition does not impact the recognized language, whose state complexity is therefore at most equal to n . Thus, we make no effort in the present article to optimize our probabilistic lower bounds. See the conclusion for a more advanced discussion on this topic.

Related work. A subclass of almost deterministic automata has already been considered in [BM08] to study the worst-case complexity class for the non-deterministic minimization problem for NFAs.

The study of random deterministic automata can be traced back to the work of Grusho on the size of the accessible part [Gru73]: he established that, with high probability, a constant proportion of the states are accessible from the initial state. He also shows that with high probability there is a unique terminal strongly connected component¹ (SCC) of size approximately $\nu_k n$, for some $\nu_k > \frac{1}{2}$ that only depends on the size k of the alphabet. More structural results on the underlying graph of a random deterministic automaton were established in the work of Carayol and Nicaud [CN12], with a local limit law for the size of the accessible part and an application to random generation of accessible deterministic automata. More recently, Cai and Devroye gave in [CD17], a fine-grained analysis of what is happening outside the large strongly connected component. In [ABBPL20], Addario-Berry, Balle and Perarnau gave a precise analysis of the diameter of a random deterministic automaton, showing in particular that it is logarithmic. We will use some of these results in this article, for instance that there is a unique largest terminal strongly connected component with high probability. To deal with the restriction to states accessible from the initial state in the powerset construction, we also use a result of [CD17] which yields that, with high probability, there are no cycles of length $\Omega(\log n)$ outside this terminal strongly connected component.

All these results on random automata focus on the underlying graph of the transition structures, without saying much about the recognized languages, or on the average complexity of textbook algorithms on automata. Some results were established in this direction: the probability that a random accessible automaton is minimal was studied by Bassino, David and Sportiello [BDS12], the analysis of minimization algorithms by Bassino, David and Nicaud [BDN12, Dav12], etc.

In another direction, more recently, several articles studied the synchronization of random automata [Ber16, Nic19], until the very recent work of Chapuy and Perarnau [CP23], establishing that most deterministic automata are synchronizing, with a word of length $O(\sqrt{n} \log n)$ and Martinsson [Mar23] established that for any $\epsilon > 0$, there exists a synchronizing word of length $O(\epsilon^{-1} \sqrt{n} \log n)$ with probability at least $1 - \epsilon$. We refer the interested reader to the survey of Nicaud [Nic14] for an overview on random deterministic automata.

To our knowledge, there is no well-established random model for non-deterministic automata. It is not an easy task to obtain a satisfactory model: for instance, the uniform distribution is degenerated and produces languages with state complexity one or two with high probability (see the discussion in Section 10). Applying the powerset construction to the mirror of a random deterministic automaton was studied by De Felice and Nicaud [FN13, FN16], in order to analyze the average case complexity of Brzozowski's state minimization algorithm. As in the present article, they studied the determinization procedure of random automata, but for a model that is very different from ours: they consider the mirror of a uniform random deterministic automaton, obtained by reversing the transitions and swapping the initial and final states. In particular, with high probability, there is a linear number

¹A strongly connected component (SCC) is a set of vertices, maximal for inclusion, in which every vertex is reachable for each other vertex. A SCC is terminal if no edge leave the SCC.

of states having a non-deterministic choice in their setting. Another natural model would be to use a critical Erdős-Rényi [ER60] digraph for each letter, which would also result in a linear number of states having a non-deterministic choice. In this article, we choose a random model with the minimum amount of non-determinism by adding just one transition to a uniform deterministic automaton, and establish that we likely have a combinatorial explosion already in this case.

This article is the full version of the extended abstract published in the proceedings of the STACS conference [CDKN23]. It contains all the omitted proofs. We also introduce the notion of templates in Section 4 to simply handle computation under natural conditional properties on automata, and reshape completely the most technical part, Section 5, to exhibit an approximation by classical Galton-Watson processes. This last approach is interesting on its own for future works on random automata.

2. DEFINITIONS AND NOTATIONS

The cardinality of a finite set E is denoted by $|E|$. For any $n \geq 1$, let $[n] = \{1, \dots, n\}$. If $x, y \in \mathbb{R}$ with $x \leq y$, let $\llbracket x, y \rrbracket = [x, y] \cap \mathbb{Z}$ be the set of integers that are between x and y . Let \mathcal{E} be a set equipped with a size function s from \mathcal{E} to $\mathbb{Z}_{\geq 0}$, and let \mathcal{E}_n denote the elements of \mathcal{E} having size n . A property X on \mathcal{E} (that is, a subset of \mathcal{E} viewed as the set of elements for which the property holds) holds with *visible probability* if there exists some constant $c > 0$ such that, for n sufficiently large, \mathcal{E}_n is non-empty and $\mathbb{P}(X) \geq c$ for the uniform distribution on \mathcal{E}_n . By a slight abuse of notation, if X is a random variable $\mathcal{E} \rightarrow \mathbb{Z}_{\geq 0}$ we say that for the uniform distribution on \mathcal{E} , X is *super-polynomial with visible probability* when for any $d \geq 1$, there exists a constant $c_d > 0$, such that for n sufficiently large, $\mathcal{E}_n \neq \emptyset$ and $\mathbb{P}(X \geq n^d) \geq c_d$.

For any real number $\lambda > 0$, we denote by $\text{Poi}(\lambda)$ the Poisson random variable of parameter λ , whose support is $\mathbb{Z}_{\geq 0}$ and defined by $\mathbb{P}(\text{Poi}(\lambda) = k) = \frac{\lambda^k}{k!} e^{-\lambda}$, for all $k \in \mathbb{Z}_{\geq 0}$.

Recall that if u and v are two words on an ordered alphabet Σ , u is *smaller than* v for the *length-lexicographic order* if $|u| < |v|$ or they have same length and $u <_{\text{lex}} v$ for the lexicographic order.

Throughout the article, the set of states of an automaton with n states will always be $[n]$, with the exception of the powerset construction recalled just below. The alphabet will always be $\Sigma = \{a, b\}$, except in the statement of our main theorem, where we allow larger alphabets as it is trivially generalized to this case. Hence, in our setting, a *deterministic (and complete) automaton* is just a tuple (n, δ, F) , where $F \subseteq [n]$ is the *set of final states* and δ is the *transition function*, a mapping from $[n] \times \Sigma$ to $[n]$. We will often write $s \xrightarrow{\alpha} t$ instead of $\delta(s, \alpha) = t$, for $s, t \in [n]$ and $\alpha \in \Sigma$, and call this an α -transition or a transition. The transition function is classically extended to sets of states by setting $\delta(X, \alpha) = \{\delta(s, \alpha) : s \in X\}$, for $X \subseteq [n]$, and to words by setting inductively $\delta(s, w) = s$ if w is the empty word ε and $\delta(s, w\alpha) = \delta(\delta(s, w), \alpha)$. We will not need to specify the *initial state* until the end of the proof of Theorem 3.1; when we finally do, it will be generated uniformly at random and independently in $[n]$. Final states are only used in the last part of our proof, so to ease the presentation, we define a *deterministic (and complete) transition structure* as being an automaton with neither initial nor final states: it is given by a pair (n, δ) where n is the number of states and δ is the transition function.

An *almost deterministic automaton* $(n, \delta, F, p \xrightarrow{a} q)$ is a deterministic automaton (n, δ, F) in which we add the additional a -transition $p \xrightarrow{a} q$. Similarly, an *almost deterministic transition structure* $(n, \delta, p \xrightarrow{a} q)$ is a deterministic transition structure (n, δ) in which we add the additional a -transition $p \xrightarrow{a} q$. For any $\alpha \in \Sigma$ and any $r \in [n]$, the transition function γ of an almost deterministic automaton $(n, \delta, F, p \xrightarrow{a} q)$ (or almost deterministic transition structure) is therefore defined by $\gamma(r, \alpha) = \{\delta(r, \alpha)\}$ if $(r, \alpha) \neq (p, a)$ and $\gamma(p, a) = \{\delta(p, a), q\}$. These automata or transition structures can be deterministic, in the case where $\delta(p, a) = q$.

The *powerset automaton* \mathcal{B} of an almost deterministic automaton $\mathcal{A} = (n, \delta, F, p \xrightarrow{a} q)$, with a transition function γ , is a deterministic automaton \mathcal{B} with states in $2^{[n]}$ and transition function γ extended to sets, as defined above. If we add an initial state i_0 to \mathcal{A} , the initial state of \mathcal{B} is $\{i_0\}$ and \mathcal{B} recognizes the same language as \mathcal{A} when a state X of \mathcal{B} is final if and only if at least one of its element is final in \mathcal{A} . This construction can be restricted to the states of \mathcal{B} that are accessible from the initial state $\{i_0\}$ while still recognizing the same language; we call this automaton the *accessible powerset automaton* of \mathcal{A} .

Recall that two states r and s in a deterministic automaton \mathcal{A} are *equivalent* if the languages recognized by moving the initial state to r or to s are equal. The *minimal automaton* of a regular language \mathcal{L} is the deterministic complete automaton with the smallest number of states that recognizes \mathcal{L} . The number of states of the minimal automaton of \mathcal{L} is called the *state complexity* of \mathcal{L} . We will use the following classical property, which follows from the fact that the minimal DFA is obtained by merging equivalent states in any DFA accepting the language [HU79, Theorem 3.11]:

Proposition 2.1. *If there is a set of accessible states X in a deterministic automaton \mathcal{A} such that the states of X are pairwise non-equivalent, then \mathcal{A} has state complexity at least $|X|$.*

3. MAIN STATEMENT AND PROOF OUTLINE

Our main result is that the state complexity of the language recognized by a random almost deterministic automaton is super-polynomial with visible probability, when for each n , each state is final, independently, with some probability f_n that is not too close to either 0 or 1, as precised in the statement:

Theorem 3.1. *Let Σ be an alphabet with at least two letters. Let f_n be a map from $\mathbb{Z}_{\geq 1}$ to $(0, 1)$ such that there exists a constant $\alpha > 0$ such that $f_n \geq \frac{\alpha}{\sqrt{n}}$ and $1 - f_n \geq \frac{\alpha}{\sqrt{n}}$ for n sufficiently large. Consider an almost deterministic n -state transition structure \mathcal{A} on Σ taken uniformly at random. Each state of \mathcal{A} is then taken to be final with probability f_n , independently of everything else. Then, with visible probability, the language recognized by \mathcal{A} has super-polynomial state complexity.*

First, we observe that if $\Gamma \subseteq \Sigma$ are two non-empty alphabets and if \mathcal{L} is a regular language on Σ , then the state complexity of \mathcal{L} is at least the state complexity of $\mathcal{L} \cap \Gamma^*$. As a consequence, it is sufficient to establish Theorem 3.1 for a two-letter alphabet, and from now on, we fix $\Sigma = \{a, b\}$.

The proof of Theorem 3.1 consists in identifying a structure and several constraints (see Figure 2) that guarantee that when performing the accessible powerset construction and adding a random set of final states, we have sufficiently many pairwise non-equivalent

the most technical, we explore backward from p and forward from r until we reach a common state.

- (2) Assuming such a w is found, we iteratively generate the transitions starting from q and following the word $w(aw)^{d-1}$, and ask for the target of each such transition to be a state that was not previously seen in the whole process. This happens with visible probability.
- (3) Let $p_0 = p$ and $p_i = \delta(q, w(aw)^{i-1})$ for $i \in [d]$. At this point, if we add the transition $p \xrightarrow{a} q$, and denote $\hat{\delta}$ the resulting non-deterministic transition function, we have $\hat{\delta}(\{p\}, (aw)^d) = \{p_0, p_1, \dots, p_d\}$, and the outgoing b -transitions of the p_i 's are still unset. Then, for each p_i , we iteratively generate the b -transitions $\delta(p_i, b), \delta(p_i, bb), \dots$ until we cycle after λ_i steps. This process is considered successful if we do not use an already set b -transition and if the $d+1$ cycles are pairwise disjoint. We furthermore require that the λ_i 's are all in $\llbracket 2\sqrt{n}, 3\sqrt{n} \rrbracket$. All these properties happen with visible probability.
- (4) The set $\{p_0, \dots, p_d\}$ is now composed of $d+1$ different states, and reading b 's from each p_i eventually ends in a b -cycle of length ℓ_i . Given the λ_i 's, each ℓ_i is a uniform element of $[\lambda_i]$, and they are independent. Our precise requirements ensure that once met, the ℓ_i 's are uniform and independent elements of $\llbracket \sqrt{n}, 2\sqrt{n} \rrbracket$. We now require that the ℓ_i 's are pairwise coprime, which also happens with visible probability.
- (5) If everything worked so far, we can guarantee that $\{p\}$ is accessible in the subset construction, still with visible probability: we use the fact that with high probability, all cycles with length in $\Omega(\log n)$ are accessible in a random deterministic automaton [CD17]. By construction the cycle around p labeled aw built at Step (1) has length $\Theta(\log n)$, hence p is accessible with high probability.
- (6) Hence, at this stage, in the accessible powerset construction applied to the almost deterministic transition structure there is a b -cycle of length $\prod_{i=0}^d \ell_i = \Omega(n^{\frac{d+1}{2}})$. We now randomly determine which states are final. If we consider a b -cycle alone in the automaton, of length $\Omega(\sqrt{n})$, its states are pairwise non-equivalent with visible probability as soon as the probability f_n that a state is final is not too close to either 0 or 1, which we assumed in our model. This property happens to be preserved when building the product automaton for the union of one-letter cycles, provided their lengths are pairwise coprime. Consequently, the large b -cycles in the powerset construction are made of pairwise non-equivalent states with visible probability.

The first steps of the proof sketch are depicted in Figure 2. More details and notations will be introduced in the next section.

We also state the following consequence for the expected state complexity of a random almost deterministic automaton. Note this corollary is weaker than the result of Theorem 3.1. Indeed, as its state complexity could be up to 2^n , a negligible proportion of automata could still contribute sufficiently to have a super-polynomial expected state complexity.

Corollary 3.2. *Under the conditions of Theorem 3.1, the expected state complexity of the language recognized by \mathcal{A} grows faster than any polynomial in n .*

Proof. Let S be the random variable that maps a random automaton to the state complexity of the language it recognizes. By Theorem 3.1, for n sufficiently large, $\mathbb{P}(S \geq n^d) \geq c_d n^d$ for some $c_d > 0$. Hence, $\mathbb{E}[S] \geq n^d \mathbb{P}(S \geq n^d) \geq c_d n^d$, concluding the proof. \square

4. TEMPLATES

In this section we introduce the notion of templates, which will be used to describe families of transition structures and will allow to handle the computations of conditional probabilities in a unified way.

A *template* \mathcal{T} is a pair (n, λ) , where λ is a function from $[n] \times \Sigma$ to $2^{[n]}$, the subsets of $[n]$, such that for every $(x, \alpha) \in [n] \times \Sigma$, $\lambda(x, \alpha)$ is non-empty.

From their definition, templates could be identified as non-deterministic and complete transition structures, but we only use them to encode restrictions. For $\alpha \in \Sigma$ and $x \in [n]$, $\lambda(x, \alpha)$ is seen as the set of allowed targets for the α -transition outgoing from x .

We say that a transition structure $\mathcal{A} = (n, \delta)$ *satisfies* a template $\mathcal{T} = (n, \lambda)$, denoted by $\mathcal{A} \models \mathcal{T}$, if for all $(x, \alpha) \in [n] \times \Sigma$, $\delta(x, \alpha) \in \lambda(x, \alpha)$. In particular, if $|\lambda(x, \alpha)| = 1$ for each $(x, \alpha) \in [n] \times \Sigma$, then $\mathcal{T} = (n, \lambda)$ can be identified with the unique transition structure that satisfies \mathcal{T} .

The next lemma is used throughout the article to work with the uniform distribution on transition structures conditioned on satisfying a template. It is obtained by direct counting.

Lemma 4.1. *Conditioned on satisfying the template $\mathcal{T} = (n, \lambda)$, the uniform distribution on n -state transition structures (n, δ) is distributed the same way as generating the image of each $\delta(x, \alpha)$ uniformly at random and independently in $\lambda(x, \alpha)$.*

In the sequel we define several sets of templates to capture the required properties at every step, following the informal description given in Section 3. Without going into the details at the moment, which will be given as needed, we will define:

- $\mathbf{B}_n(p)$ as the templates that force the required conditions on the backward exploration from state p , defined in Section 5;
- $\mathbf{C}_n(p)$ as the templates built from $\mathbf{B}_n(p)$ that furthermore force the existence of a small cycle around p , defined in Section 6;
- $\mathbf{T}_n(p, q)$ as the templates built from $\mathbf{C}_n(p)$ that furthermore force the existence of the path starting at q and of the b -threads, as depicted in Figure 2 and defined in Section 7.

5. BACKWARD SUBSTRUCTURE OF A RANDOM TRANSITION STRUCTURE

This section studies the shape of the backward substructure at logarithmic depth from a state $p \in [n]$ in a random transition structure of size n . The backward substructure at depth h from a vertex p , formally defined in Subsection 5.1, is obtained by restricting the transition structure to the vertices that can reach p in at most h steps.

More precisely, we study the backward substructure at depth $h := \lceil \log_2 \sqrt{n} \rceil$ and we want to show that with visible probability in a random transition structure, this backward substructure contains at most $d_1 \sqrt{n}$ vertices with at least $d_2 \sqrt{n}$ vertices at depth exactly h , for some positive constants d_1 and d_2 . This result is formally stated in Proposition 5.6 and its proof is the aim of this section.

For a fixed maximal depth m , Cai and Devroye [CD17] studied the stochastic process giving the number of vertices at depth at most m in a random transition structure of size n by providing a perfect coupling with a Galton-Watson process whose offspring distribution is $\text{Bin}(2n, 1/n)$. As we work at logarithmic depth $\lceil \log_2 \sqrt{n} \rceil$, we are not able to provide such a tight description. However, inspired by their approach, we introduce in Subsection 5.2 a stochastic process \mathcal{G}_h that builds what we called backward multi-trees. Under certain

natural restrictions, a backward multi-tree of height h can encode a backward structure of depth h . In Lemma 5.2, we show that for every backward-structure of size at most $\frac{n}{3}$, the probability that a random structure of size n has this backward structure at depth h is lower-bounded by the probability that the process \mathcal{G}_h produces a multi-tree that encodes this backward structure.

Thanks to this reduction, it only remains to show that with visible probability the process \mathcal{G}_h produces a multi-tree encoding a backward structure satisfying the restrictions on its total size and the number of nodes at depth h required in Proposition 5.6. A key ingredient in obtaining this result is that the process \mathcal{G}_h is constructed in such a way that if we only consider the number of vertices at each depth and not the structure of the multi-tree, we are in fact studying a Galton-Watson process with offspring distribution $\text{Poi}(2)$. This model is well-studied and we derive the results necessary for our purpose in Subsection 5.3, using standard techniques.

5.1. Backward substructure. Let \mathcal{A} be a n -state deterministic transition structure and $p \in [n]$. For any state $x \in [n]$, let $d_p^-(x)$ denote the *backward distance* from p to x , that is, the length of a shortest path from x to p in \mathcal{A} . By convention, $d_p^-(x) = +\infty$ if p is not reachable from x . For $t \geq 0$, the t -th *backward layer* of \mathcal{A} from p is the set $L_t(\mathcal{A}, p)$ (simply denoted by L_t or $L_t(\mathcal{A})$ when the context is clear) consisting of the states at backward distance t from p : $L_t(\mathcal{A}, p) = \{x : d_p^-(x) = t\}$.

For any integer $h \geq 0$, we define the h -backward substructure \mathcal{B}_h of \mathcal{A} from p as the incomplete deterministic transition structure obtained by selecting the states at backward distance at most h from p and the transitions that are part of their shortest paths to p . Formally, if $\mathcal{A} = (n, \delta)$, then $\mathcal{B}_h = (n, \gamma)$, where γ is the partial transition function defined by $\gamma(x, \alpha) = y$ if and only if $\delta(x, \alpha) = y$ and $d_p^-(x) \leq h$ and $d_p^-(y) = d_p^-(x) - 1$. The *backward support* of \mathcal{B}_h , denoted $\text{Support}_{\text{B}}(\mathcal{B}_h)$ is $\cup_{k=0}^h L_k$, the set of states of \mathcal{B}_h which are either the source or target of a transition, or just $\{p\}$ in the special case of $h = 0$. When there is no ambiguity on h and p , we sometimes write $\text{Support}_{\text{B}}(\mathcal{A})$ instead of $\text{Support}_{\text{B}}(\mathcal{B}_h)$.

For any $n \geq 1$, any $p \in [n]$ and any $h \geq 0$, denote by $\mathfrak{B}_{n,h}(p)$ the set of all possible h -backward substructures from p obtained from a n -state deterministic transition structure.

The template *associated to* $\mathcal{B}_h = (n, \gamma) \in \mathfrak{B}_{n,h}(p)$ is the template $\mathcal{T} = (n, \lambda)$ defined for every $x \in [n]$ and $\alpha \in \Sigma$ by:

- (1) If $\gamma(x, \alpha)$ is defined, then $\lambda(x, \alpha) = \{\gamma(x, \alpha)\}$.
- (2) If $\gamma(x, \alpha)$ is undefined and $x \in \text{Support}_{\text{B}}(\mathcal{B}_h)$, then $\lambda(x, \alpha) = [n] \setminus \cup_{k=0}^{d_p^-(x)-1} L_k$.
- (3) Otherwise, if $x \notin \text{Support}_{\text{B}}(\mathcal{B}_h)$, then $\lambda(x, \alpha) = [n] \setminus \cup_{k=0}^{h-1} L_k$.

This template is defined to precisely capture all n -state transition structures admitting \mathcal{B}_h as h -backward substructure.

Lemma 5.1. *If \mathcal{T} is the template associated to $\mathcal{B}_h = (n, \gamma) \in \mathfrak{B}_{n,h}(p)$, then a n -state transition structure \mathcal{A} satisfies \mathcal{T} if and only if its h -backward substructure is \mathcal{B}_h .*

Proof. Suppose that $\mathcal{A} = (n, \delta)$ satisfies \mathcal{T} . A direct induction shows that for all $t \in \{0, \dots, h\}$ we have $L_t(\mathcal{A}) = L_t(\mathcal{B}_h)$. Moreover, if $x \in L_t$ for $t \in \{1, \dots, h\}$, then the conditions of Case (2) in the definition of \mathcal{T} yields that $\delta(x, \alpha) \in L_{t-1}$ if and only if $\gamma(x, \alpha)$ is defined, yielding $\delta(x, \alpha) = \gamma(x, \alpha)$ as we must be in Case (1). Hence the h -backward substructure of \mathcal{A} is \mathcal{B}_h .

Conversely, if the h -backward structure of \mathcal{A} is \mathcal{B}_h , then for every $(x, \alpha) \in [n] \times \Sigma$:

- (1) If $\gamma(x, \alpha)$ is defined, then $\delta(x, \alpha) = \gamma(x, \alpha)$, thus $\delta(x, \alpha) \in \lambda(x, \alpha)$.
- (2) If $\gamma(x, \alpha)$ is undefined and $x \in \text{Support}_{\mathcal{B}}(\mathcal{B}_h)$, then $\delta(x, \alpha)$ cannot be in $\cup_{k=0}^{d-1} L_k$, where $d = d_p^-(x)$ in \mathcal{B}_h , otherwise it would create a shorter path from x to p in \mathcal{A} , and \mathcal{B}_h would not be its h -backward substructure.
- (3) If $x \notin \text{Support}_{\mathcal{B}}(\mathcal{B}_h)$, using the same argument, $\delta(x, \alpha)$ cannot be in $\cup_{k=0}^{h-1} L_k$.

This concludes the proof. □

Following the statement of Lemma 5.1, we write $\mathcal{A} \models \mathcal{B}_h$ when \mathcal{A} satisfies the template associated with \mathcal{B}_h . We further extend the notation to sets \mathfrak{X} of backward substructures: $\mathcal{A} \models \mathfrak{X}$ if \mathcal{A} satisfies at least one element of \mathfrak{X} . Lemma 5.1 also allows us to define $\text{Support}_{\mathcal{B}}(\mathcal{T}) = \text{Support}_{\mathcal{B}}(\mathcal{B}_h)$, when \mathcal{T} is the template associated with $\mathcal{B}_h \in \mathfrak{B}_{n,h}(p)$.

5.2. Backward multi-trees, trees and processes. We define backward multi-trees which are general trees and identify a subclass called backward trees that naturally encode backward substructures. Using this view, we can consider a simpler stochastic process that produces backward multi-trees instead of backward-substructures. In Lemma 5.2, we show that this stochastic process can be used to give a lower-bound on the probability that a given backward substructure appears in a random transition structure.

A *n*-backward multi-tree rooted at p is a finite general² tree whose root is labeled by p and whose other nodes are labeled by pairs $(x, \alpha) \in [n] \times \Sigma$. By a slight abuse of notation, if a node N is labeled by (x, α) , we say that x is the state that labels N or that x labels N . The name multi-tree underlines the fact that a state can label several nodes in the tree.

A *n*-backward tree is a *n*-backward multi-tree whose nodes are labeled by pairwise distinct states. In particular, its labels are pairwise distinct. The notions of backward multi-tree and of backward tree are illustrated in Figure 3, for $n = 8$.

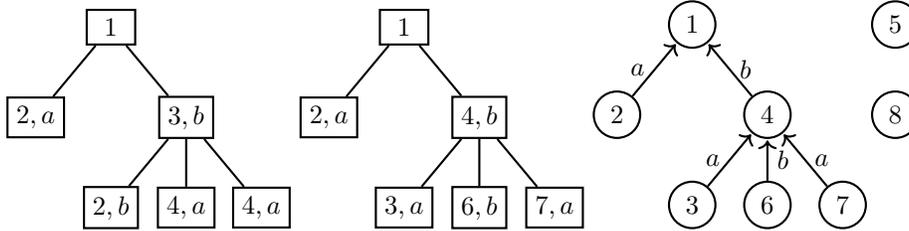


Figure 3: On the left, a backward multi-tree which is not a backward tree because, for instance, the state 2 labels two nodes. In the center, a backward multi-tree which is a tree with its corresponding backward substructure on the right.

A *n*-backward tree T of depth at most h can be very simply transformed into a h -backward substructure with n states, by starting with no transition then changing every child→parent relation $(x, \alpha) \rightarrow (y, \beta)$ into the transition $x \xrightarrow{\alpha} y$, and $(x, \alpha) \rightarrow p$ into $x \xrightarrow{\alpha} p$: each letter is moved from its node's label to the edge linking it to its parent, and the unused states are added with no outgoing transition. We denote by Λ this map from a *n*-backward tree to a *n*-state backward substructure.

²In a general tree, the children of each node are not ordered.

We say that a n -backward multi-tree T matches a backward substructure \mathcal{B} when T is a tree and $\Lambda(T) = \mathcal{B}$. Let \mathfrak{T}_n denote the set of n -backward multi-trees that are trees, and let $\mathfrak{T}_{n,h}(p)$ be the set of backward trees that match a h -backward substructure at p :

$$\mathfrak{T}_{n,h}(p) := \{T \in \mathfrak{T}_n : \Lambda(T) \in \mathfrak{B}_{n,h}(p)\}.$$

Observe that a backward substructure can be matched by a backward tree if and only if every state in its support has exactly one outgoing transition, except p which has none. We say that such a backward substructure is a *tree*.

For $n \geq 1$ and $p \in [n]$, we define a stochastic process $(\mathcal{G}_t)_{t \geq 0}$, called the *backward multi-tree process* of parameters n and p , which produces random n -backward multi-trees as follows. At time $t = 0$, the tree consists of a unique root node labeled p , which is at depth 0. For $t \geq 0$, the tree \mathcal{G}_{t+1} is built from \mathcal{G}_t in the following way. For every node N at depth t in \mathcal{G}_t , we draw independent Poisson random variables $X_{x,\alpha}$ of parameter $\frac{1}{n}$, one for every $x \in [n]$ and for every $\alpha \in \Sigma$. If $X_{x,\alpha} > 0$, then we add $X_{x,\alpha}$ children labeled (x, α) to N , at depth $t + 1$.

The following lemma justifies the introduction of the multi-tree process $(\mathcal{G}_t)_{t \geq 0}$ as we will use it to lift some visible properties of a backward tree generated by this process to the visible properties of the backward substructure of a random transition structures.

Lemma 5.2. *Let $\mathcal{B} \in \mathfrak{B}_{n,h}(p)$ such that $|\text{Support}_{\mathcal{B}}(\mathcal{B})| < \frac{n}{3}$ and the h -th layer of \mathcal{B} is not empty. Let \mathcal{A}_n denote a uniform random transition structure with n states, and $(\mathcal{G}_t)_{t \geq 0}$ be the backward tree process of parameters n and p . Then $\mathbb{P}(\mathcal{A}_n \models \mathcal{B}) \geq \mathbb{P}(\mathcal{G}_h \text{ matches } \mathcal{B})$.*

Proof. If \mathcal{B} is not a tree then $\mathbb{P}(\mathcal{G}_h \text{ matches } \mathcal{B}) = 0$ and the inequality trivially holds. So we assume that \mathcal{B} is a tree for the remainder of the proof.

Let \mathcal{T} denote the template associated with \mathcal{B} . We proceed by induction on the depth h of $\mathcal{B} = (n, \gamma)$. The claim trivially holds if $h = 0$. For the induction step, assume the claim is true for depth h and consider a backward substructure \mathcal{B} of depth $h + 1$. Let $\hat{\mathcal{B}}$ be the associated backward substructure of depth h , which just consists in removing the outgoing transitions from the states of $L_{h+1}(\mathcal{B})$. Observe that $\hat{\mathcal{B}}$ is also a tree, thus we have $\mathbb{P}(\mathcal{G}_h \text{ matches } \hat{\mathcal{B}}) > 0$.

Let $\hat{\mathcal{T}}$ denote the template associated with $\hat{\mathcal{B}}$. As $|\text{Support}_{\mathcal{B}}(\hat{\mathcal{B}})| \leq |\text{Support}_{\mathcal{B}}(\mathcal{B})| \leq \frac{n}{3}$, by induction hypothesis we have $\mathbb{P}(\mathcal{A}_n \models \hat{\mathcal{T}}) \geq \mathbb{P}(\mathcal{G}_h \text{ matches } \hat{\mathcal{B}})$.

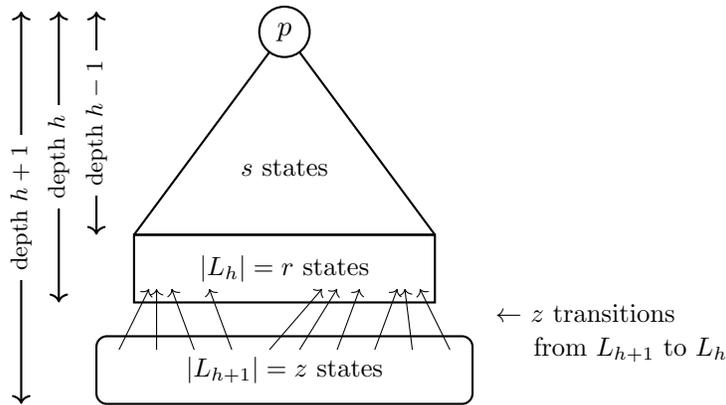


Figure 4: An illustration for the values of r , s and z in the proof of Lemma 5.2.

Let $r = |L_h(\mathcal{B})| = |L_h(\hat{\mathcal{B}})|$, $s = |\cup_{k=0}^{h-1} L_k|$ and $z = |L_{h+1}(\mathcal{B})|$. Observe that since \mathcal{B} is a tree, z is also the number of transitions from L_{h+1} to L_h in \mathcal{B} . They are the transitions we removed to build $\hat{\mathcal{B}}$.

A transition structure $\mathcal{A} = (n, \delta)$ that satisfies $\hat{\mathcal{T}} = (n, \hat{\lambda})$ also satisfies \mathcal{T} if and only if for every $x \notin \text{Support}_{\mathcal{B}}(\hat{\mathcal{B}})$ and every $\alpha \in \Sigma$ we have (i) if $\gamma(x, \alpha)$ is defined in \mathcal{B} then $\delta(x, \alpha) = \gamma(x, \alpha)$, and (ii) if $\gamma(x, \alpha)$ is undefined in \mathcal{B} then $\delta(x, \alpha) \notin \text{Support}_{\mathcal{B}}(\hat{\mathcal{B}})$. Since $\mathcal{A} \models \hat{\mathcal{T}}$, for every $x \notin \text{Support}_{\mathcal{B}}(\hat{\mathcal{B}})$ and every $\alpha \in \Sigma$, $\hat{\lambda}(x, \alpha) = [n] \setminus \cup_{k=0}^{h-1} L_k$, hence $|\hat{\lambda}(x, \alpha)| = n - s$. Out of these $n - s$ possible targets, only one is possible for the z transitions of Case (i) and $n - r - s$ are possible for the $2(n - r - s) - z$ transitions of Case (ii). Since the other transitions have the same constraints in \mathcal{T} and in $\hat{\mathcal{T}}$, by Lemma 4.1 we have

$$\mathbb{P}(\mathcal{A}_n \models \mathcal{T} \mid \mathcal{A}_n \models \hat{\mathcal{T}}) = \frac{1}{(n-s)^z} \times \left(\frac{n-r-s}{n-s} \right)^{2(n-s-r)-z},$$

We rewrite this probability the following way:

$$\mathbb{P}(\mathcal{A}_n \models \mathcal{T} \mid \mathcal{A}_n \models \hat{\mathcal{T}}) = \frac{1}{(n-r-s)^z} \times \left(1 - \frac{r}{n-s} \right)^{2(n-s-r)}.$$

On the other hand, by construction of the process and since $\mathbb{P}(\text{Poi}(\frac{1}{n}) = 0) = e^{-1/n}$ and $\mathbb{P}(\text{Poi}(\frac{1}{n}) = 1) = \frac{1}{n}e^{-1/n}$, considering the $2nr$ possible transitions that ends in one of the r vertices of the last layer in $\hat{\mathcal{B}}$:

$$\mathbb{P}(\mathcal{G}_{h+1} \text{ matches } \mathcal{B} \mid \mathcal{G}_h \text{ matches } \hat{\mathcal{B}}) = \left(\frac{1}{n}e^{-1/n} \right)^z \times \left(e^{-1/n} \right)^{2nr-z} = \frac{e^{-2r}}{n^z}.$$

Let R be the ratio of these two probabilities, we have

$$R = \frac{\mathbb{P}(\mathcal{A}_n \models \mathcal{T} \mid \mathcal{A}_n \models \hat{\mathcal{T}})}{\mathbb{P}(\mathcal{G}_{h+1} \text{ matches } \mathcal{B} \mid \mathcal{G}_h \text{ matches } \hat{\mathcal{B}})} \geq \left(1 - \frac{r}{n-s} \right)^{2(n-s-r)} e^{2r}.$$

As $-\log(1-x) \geq x$ for $x \in (0, 1)$, we have

$$\log R \geq 2(n-s) \log \left(1 - \frac{r}{n-s} \right) + \frac{2r^2}{n-s} + 2r.$$

A basic study of the function $x \mapsto \log(1-x) + x + \frac{3}{4}x^2$ yields that for all $x \in [0, \frac{1}{3}]$, we have $\log(1-x) \geq -x - \frac{3}{4}x^2$. Moreover, $r + \frac{s}{3} \leq r + s = |\text{Support}_{\mathcal{B}}(\mathcal{B})| \leq \frac{n}{3}$, hence $r \leq \frac{n-s}{3}$. Therefore, we have

$$\log R \geq -2(n-s) \left(\frac{r}{n-s} + \frac{3r^2}{4(n-s)^2} \right) + \frac{2r^2}{n-s} + 2r = \frac{r^2}{2(n-s)} \geq 0.$$

This yields that $\mathbb{P}(\mathcal{A}_n \models \mathcal{T} \mid \mathcal{A}_n \models \hat{\mathcal{T}}) \geq \mathbb{P}(\mathcal{G}_{h+1} \text{ matches } \mathcal{B} \mid \mathcal{G}_h \text{ matches } \hat{\mathcal{B}})$. Therefore,

$$\begin{aligned} \mathbb{P}(\mathcal{A}_n \models \mathcal{T}) &= \mathbb{P}(\mathcal{A}_n \models \mathcal{T} \text{ and } \mathcal{A}_n \models \hat{\mathcal{T}}) \\ &= \mathbb{P}(\mathcal{A}_n \models \mathcal{T} \mid \mathcal{A}_n \models \hat{\mathcal{T}}) \mathbb{P}(\mathcal{A}_n \models \hat{\mathcal{T}}) \\ &\geq \mathbb{P}(\mathcal{G}_{h+1} \text{ matches } \mathcal{B} \mid \mathcal{G}_h \text{ matches } \hat{\mathcal{B}}) \mathbb{P}(\mathcal{G}_h \text{ matches } \hat{\mathcal{B}}) \\ &= \mathbb{P}(\mathcal{G}_{h+1} \text{ matches } \mathcal{B} \text{ and } \mathcal{G}_h \text{ matches } \hat{\mathcal{B}}) \\ &= \mathbb{P}(\mathcal{G}_{h+1} \text{ matches } \mathcal{B}), \end{aligned}$$

concluding the proof by induction. \square

5.3. Results in Galton-Watson processes of offspring distribution $\text{Poi}(2)$. Let $(\mathcal{G}_t)_{t \geq 0}$ be a backward multi-tree process of parameters n and p . To each \mathcal{G}_t we associate the quantity Z_t , defined as its number of nodes at depth t . Since the sum of $|\Sigma| \cdot n$ independent Poisson random variables of parameter $\frac{1}{n}$ is a Poisson random variable of parameter $|\Sigma|$ and $|\Sigma| = 2$, each node at depth t in \mathcal{G}_t gives birth to $\text{Poi}(2)$ children at time $t + 1$. Hence $(Z_t)_{t \geq 0}$ is exactly a Galton-Watson process of offspring distribution $\text{Poi}(2)$, which is a well studied branching process [Har63]. We tailor classical results on Galton-Watson processes to our needs as follows.

Theorem 5.3. *Let $(Z_t)_{t \geq 0}$ be a Galton-Watson process of offspring distribution $\text{Poi}(2)$. There exist positive real numbers c, c_1, c_2, c_3 and a positive integer t_0 such that,*

$$\mathbb{P} \left(\forall t \geq t_0, c_1 2^t \leq Z_t \leq c_2 2^t \text{ and } \sum_{k=0}^{t-1} Z_k \leq c_3 2^t \right) \geq c.$$

Proof. As the expected value of the offspring law is 2 and its variance is finite, the sequence of random variables $(W_t)_{t \geq 0}$ defined by $W_t = 2^{-t} Z_t$ almost surely converges to a random variable W with $\mathbb{E}[W] = 1$ (see [Har63, Theorem 8.1, p.13]). Since $\mathbb{E}[W] > 0$, there must exist a constant $d > 0$ such that $\mathbb{P}(|W - d| < \frac{d}{4}) > 0$: we can cover $\mathbb{R}_{\geq 0}$ with a countable number of intervals $(\frac{3}{4}d, \frac{5}{4}d)$, and W is in at least one of them with positive probability.

For $t_0 \geq 1$ and $x > 0$, we consider the following events:

- event $A_{t_0} : \bigcap_{t=t_0}^{\infty} |W_t - W| < \frac{d}{4}$,
- event $B : |W - d| < \frac{d}{4}$,
- event $C_{t_0, x} : \sum_{t=0}^{t_0-1} Z_t \leq x(2^{t_0} - 1)$.

Our choice of d yields that $c := \mathbb{P}(B)$ is positive. As $(W_t)_{t \geq 0}$ converges almost surely to W , $\lim_{t_0 \rightarrow \infty} \mathbb{P}(A_{t_0}) = 1$. So we can fix a value for t_0 such that $\mathbb{P}(A_{t_0}) \geq 1 - \frac{c}{4}$. For this t_0 , $\lim_{x \rightarrow \infty} \mathbb{P}(C_{t_0, x}) = 1$, so we can choose some value for x such that $\mathbb{P}(C_{t_0, x}) \geq 1 - \frac{c}{4}$. For these choices of t_0 and x , this yields, by the union bound, using \bar{X} for the complement of X :

$$\mathbb{P}(A_{t_0} \cap B \cap C_{t_0, x}) \geq 1 - \mathbb{P}(\bar{A}_{t_0}) - \mathbb{P}(\bar{B}) - \mathbb{P}(\bar{C}_{t_0, x}) \geq \frac{c}{2} > 0.$$

We finalize the proof as follows. Assume that the events A_{t_0} , B and $C_{t_0, x}$ occur simultaneously. For all $t \geq t_0$, we have $|W_t - d| \leq |W_t - W| + |W - d| \leq \frac{d}{4} + \frac{d}{4} \leq \frac{d}{2}$. Hence, for all $t \geq t_0$, $\frac{d}{2} 2^t \leq Z_t \leq \frac{3d}{2} 2^t$, and we can take $c_1 = \frac{d}{2}$ and $c_2 = \frac{3d}{2}$. Furthermore, for $t > t_0$, we have:

$$\sum_{k=0}^{t-1} Z_k = \sum_{k=0}^{t_0-1} Z_k + \sum_{k=t_0}^{t-1} Z_k \leq x(2^{t_0} - 1) + \frac{3d}{2} \sum_{k=t_0}^{t-1} 2^k,$$

which is at most $c_3 2^t$ for $c_3 = \max\{x, 3d/2\}$, concluding the proof. \square

Lemma 5.4. *Let $(Z_t)_{t \geq 0}$ be a Galton-Watson process of offspring distribution $\text{Poi}(2)$. The probability that $\sum_{i=0}^t Z_i \geq 2^{3t/2}$ tends to 0 as t tends to infinity.*

Proof. For all $t \geq 0$, we have $\mathbb{E}[Z_t] = 2^t$. By linearity of expectation, $\mathbb{E}[\sum_{i=0}^t Z_i] = 2^{t+1} - 1$. Hence by Markov inequality, for all $i \geq 0$, $\mathbb{P}(\sum_{i=0}^t Z_i \geq 2^{3t/2}) \leq \frac{(2^{t+1} - 1)}{2^{3t/2}}$ which tends to 0 as t tends to infinity, concluding the proof. \square

5.4. Main result on backward substructures. The aim of this section is to conclude our work on backward substructures, by proving Proposition 5.6 stated below.

Let \mathfrak{D} denote the set of backward multi-trees in which no node has two children with the same label. For instance in Figure 3, the multi-tree on the left is not in \mathfrak{D} as $(3, b)$ has two children with same label $(4, a)$. We first show that with high probability the multi-tree produced by the process is in \mathfrak{D} .

Lemma 5.5. *Let $(\mathcal{G}_t)_{t \geq 0}$ be the backward multi-tree process of parameters n and p . With high probability, $\mathcal{G}_h \in \mathfrak{D}$ for $h := \lceil \log_2 \sqrt{n} \rceil$.*

Proof. A node has children with the same label if and only if at least one of the $\text{Poi}(1/n)$ used to construct its children values at least two.

The probability that such a random variable is at least two is $1 - e^{-1/n}(1 + 1/n) = 1/n^2 + o(1/n^2)$. Thus if we consider a sequence of $o(n^2)$ such random variables, the probability that at least one of them reaches two or more is $o(1)$ by the union bound.

Let $(\xi_i)_{i \geq 1}$ be a sequence of i.i.d. $\text{Poi}(1/n)$ random variables used for the construction of \mathcal{G}_h . By the union bound, as $\mathbb{P}(|\mathcal{G}_h| \geq 2^{3h/2} \text{ and } \mathcal{G}_h \notin \mathfrak{D}) \leq \mathbb{P}(|\mathcal{G}_h| \geq 2^{3h/2})$ we have

$$\mathbb{P}(\mathcal{G}_h \notin \mathfrak{D}) \leq \mathbb{P}\left(|\mathcal{G}_h| \geq 2^{3h/2}\right) + \mathbb{P}\left(\exists i \in \{1, \dots, 2n \lfloor 2^{3h/2} \rfloor\}, \xi_i \geq 2\right).$$

The first probability tends to 0 by Lemma 5.4. The second probability also tends to 0 by the above remark and the fact that $n \lfloor 2^{3h/2} \rfloor = \Theta(n^{7/4})$. This concludes the proof. \square

For a multi-tree in \mathfrak{D} , we define a notion of shape which intuitively corresponds to drawing the multi-tree in the plane with the children ordered lexicographically and then in removing all the labels. By definition of \mathfrak{D} , all children of a node have different labels, hence this embedding is uniquely defined. The proof of Proposition 5.6 below consists in showing that for a fixed shape S , given it has shape S , the probability that \mathcal{G}_h is a tree can be bounded from below by some constant τ independently of the shape S . We will then conclude by the law of total probabilities.

Proposition 5.6. *There exists $c_B, d_1, d_2 > 0$ such that, for any n sufficiently large and any $p \in [n]$, the h -backward substructure \mathcal{B} from p of depth $h = \lceil \log_2 \sqrt{n} \rceil$ of a uniform random n -states transition structure is such that $|\text{Support}_B(\mathcal{B})| \leq d_1 \sqrt{n}$ and $|L_h(\mathcal{B})| \geq d_2 \sqrt{n}$ with probability at least c_B .*

Proof. We choose some values for c, c_1, c_2, c_3 and t_0 that work with Theorem 5.3. We set $d_1 = 2(c_2 + c_3)$ and $d_2 = c_1$. Let $(\mathcal{G}_t)_{t \geq 0}$ be a backward multi-tree process of parameters n and p , and let $(Z_t)_{t \geq 0}$ denote the number of nodes at depth t in \mathcal{G}_t (and therefore in all $\mathcal{G}_{t'}$ for $t' \geq t$).

Let \mathfrak{G}_h denote the set defined by:

$$\mathfrak{G}_h = \left\{ \mathcal{G} \in \mathfrak{D} : c_1 2^h \leq Z_h(\mathcal{G}) \leq c_2 2^h \text{ and } \sum_{k=0}^{h-1} Z_k(\mathcal{G}) \leq c_3 2^h \right\}.$$

For n sufficiently large, $h \geq t_0$ and Theorem 5.3 applies. Hence, together with Lemma 5.5, for n sufficiently large we have

$$\mathbb{P}(\mathcal{G}_h \in \mathfrak{G}_h) \geq \frac{c}{2}, \tag{5.1}$$

as the intersection of a property with visible probability and a property with high probability.

We start by showing that there is a visible probability that \mathcal{G}_h belongs to \mathfrak{G}_h and is a tree. To do so, we introduce the notion of shape of a backward multi-tree and bound from below the probability at fixed shape.

To any $\mathcal{G} \in \mathfrak{D}$ we associate its *shape* $s(\mathcal{G})$, a tuple of non-negative integers which is inductively defined as follows:

- if \mathcal{G} is reduced to a root node, then $s(\mathcal{G}) = (0)$,
- otherwise, if $(\mathcal{G}_i)_{i=1,\dots,k}$ denote its k children ordered by their root labels (using the lexicographic order), then $s(\mathcal{G}) = (k) \otimes s(\mathcal{G}_1) \otimes s(\mathcal{G}_2) \otimes \dots \otimes s(\mathcal{G}_k)$, where \otimes denote the concatenation of tuples: $(x_1, \dots, x_m) \otimes (y_1, \dots, y_r) = (x_1, \dots, x_m, y_1, \dots, y_r)$.

So we compute $s(\mathcal{G})$ by producing the sequence of the number of children in a depth-first traversal of \mathcal{G} , where the children of a node are taken in lexicographic order of their labels. There is no ambiguity as in an element of \mathfrak{D} , two siblings always have different labels. If $\mathcal{G} \notin \mathfrak{D}$, we set $s(\mathcal{G}) = \perp$ to indicate that no shape is defined.

The label sequence of $\mathcal{G} \in \mathfrak{D}$, denoted by $\ell(\mathcal{G})$ is computed as $s(\mathcal{G})$, except that we collect the labels of the nodes instead of their numbers of children. One can readily verify that if both $s(\mathcal{G})$ and $\ell(\mathcal{G})$ are given, there is a unique element of \mathfrak{D} that matches them, which is \mathcal{G} : informally $s(\mathcal{G})$ encodes the tree structure, and the nodes are labeled with $\ell(\mathcal{G})$ using a depth-first traversal. Let $\mathcal{S}_h = \{s(\mathcal{G}) : \mathcal{G} \in \mathfrak{G}_h\}$ be the set of shapes of elements of \mathfrak{G}_h . Observe that if $s(\mathcal{G}) \in \mathcal{S}_h$, then by definition $\mathcal{G} \in \mathfrak{D}$ (otherwise $s(\mathcal{G}) = \perp$); moreover, the shape of \mathcal{G} determines the size of its layers, and therefore $\mathcal{G} \in \mathfrak{G}_h$. Hence $s(\mathcal{G}) \in \mathcal{S}_h$ if and only if $\mathcal{G} \in \mathfrak{G}_h$.

The probability that $\mathcal{G}_h = \mathcal{G}$ is the same for any \mathcal{G} having the same shape $\mathcal{S} \in \mathcal{S}_h$. Indeed, by direct induction on the steps of the depth-first traversal, each time we generate the k children of the current node, the value of k being given by $s(\mathcal{G})$, we draw $2n$ independent $\text{Poi}(1/n)$ out of which exactly k must value 1 and $2n - k$ must value 0. This happens with probability that only depends on k , hence only on $s(\mathcal{G})$.

So we can compute the probability that \mathcal{G}_h is a backward tree, given its shape \mathcal{S} , by a counting argument. Indeed, to construct a backward tree with shape $\mathcal{S} = (s_1, \dots, s_t)$, we first need to choose the states labeling the s_1 children of the root (which is always labeled by p) and for each child assign a letter in Σ . This represents $\binom{n-1}{s_1} 2^{s_1}$ choices. Then we need to choose the s_2 states labeling the s_2 children of the first child of the root. These s_2 states must be chosen in amongst the $n - 1 - s_1$ not used so far. For each of these states, which is uniquely identified by its label, we must choose a letter in Σ for a total number of $\binom{n-1-s_1}{s_2} 2^{s_2}$ choices (observe that the formula is also correct if $s_2 = 0$). By repeating this argument in the form of a direct induction, we have that the number of ways to build a backward tree of shape $\mathcal{S} = (s_1, \dots, s_t)$ is

$$T_n(\mathcal{S}) := \binom{n-1}{s_1} 2^{s_1} \times \binom{n-1-s_1}{s_2} 2^{s_2} \times \dots \times \binom{n-1-\sum_{i=1}^{t-1} s_i}{s_t} 2^{s_t}.$$

On the other hand, the number of way of labeling \mathcal{S} to obtain an element of \mathfrak{D} is

$$D_n(\mathcal{S}) := \binom{2n}{s_1} \times \binom{2n}{s_2} \times \dots \times \binom{2n}{s_t},$$

as the only constraint is that the children of a node have different labels, amongst the $2n$ possible ones. Recalling that \mathfrak{T}_n denotes the set of n -backward trees, we just established that $\mathbb{P}(\mathcal{G}_h \in \mathfrak{T}_n \mid s(\mathcal{G}_h) = \mathcal{S}) = \frac{T_n(\mathcal{S})}{D_n(\mathcal{S})}$. We can easily bound from below as follows, using

$m = 1 + \sum_{i=1}^t s_i$ to denote the number of nodes:

$$T_n(\mathcal{S}) \geq 2^{m-1} \frac{(n-m)^{m-1}}{\prod_{i=1}^t s_i!} \text{ and } D_n(\mathcal{S}) \leq \frac{(2n)^{m-1}}{\prod_{i=1}^t s_i!}.$$

Therefore, we have

$$\mathbb{P}(\mathcal{G}_h \in \mathfrak{T}_n \mid s(\mathcal{G}_h) = \mathcal{S}) \geq \left(1 - \frac{m}{n}\right)^m.$$

In particular, if \mathcal{S} has at most $d_1\sqrt{n}$ nodes, then for n sufficiently large we have

$$\mathbb{P}(\mathcal{G}_h \in \mathfrak{T}_n \mid s(\mathcal{G}_h) = \mathcal{S}) \geq \frac{1}{2} \exp(-d_1^2), \quad (5.2)$$

as $\lim_n (1 - d_1/\sqrt{n})^{d_1\sqrt{n}} = e^{-d_1^2}$.

We can now finalize the proof by bounding from below the probability that \mathcal{G}_h is a n -backward tree and in \mathfrak{G}_h as follows, for n sufficiently large. We first partition the elements of \mathfrak{G}_h according to their shape:

$$\begin{aligned} \mathbb{P}(\mathcal{G}_h \in \mathfrak{G}_h \cap \mathfrak{T}_n) &= \sum_{\mathcal{S} \in \mathcal{S}_h} \mathbb{P}(s(\mathcal{G}_h) = \mathcal{S} \text{ and } \mathcal{G}_h \in \mathfrak{T}_n) \\ &= \sum_{\mathcal{S} \in \mathcal{S}_h} \mathbb{P}(\mathcal{G}_h \in \mathfrak{T}_n \mid s(\mathcal{G}_h) = \mathcal{S}) \mathbb{P}(s(\mathcal{G}_h) = \mathcal{S}). \end{aligned}$$

As \mathcal{S}_h is the set of shapes of the elements of \mathfrak{G}_h , all the \mathcal{G} such that $s(\mathcal{G}) \in \mathcal{S}_h$ has at most $(c_2 + c_3)2^h \leq 2(c_2 + c_3)\sqrt{n} = d_1\sqrt{n}$ nodes, Equation (5.2) applies and yields:

$$\mathbb{P}(\mathcal{G}_h \in \mathfrak{G}_h \cap \mathfrak{T}_n) \geq \frac{e^{-d_1^2}}{2} \sum_{\mathcal{S} \in \mathcal{S}_h} \mathbb{P}(s(\mathcal{G}_h) = \mathcal{S}) = \frac{e^{-d_1^2}}{2} \mathbb{P}(\mathcal{G}_h \in \mathfrak{G}_h).$$

By Equation (5.1), if we set $c_B = \frac{ce^{-d_1^2}}{4}$, then for n sufficiently large $\mathcal{G}_h \in \mathfrak{G}_h \cap \mathfrak{T}_n$ with probability at least c_B .

Recall that for $\mathcal{G} \in \mathfrak{T}$, $\Lambda(\mathcal{G})$ denote the associated backward substructure. Let $\mathfrak{B}_h \subseteq \mathfrak{B}_{n,h}(p)$ denote the set of backward substructures \mathcal{B} such that $|\text{Support}_B(\mathcal{B})| \leq d_1\sqrt{n}$ and $|L_h(\mathcal{B})| \geq d_2\sqrt{n}$. As every structure in $\Lambda(\mathfrak{G}_h \cap \mathfrak{T}_n)$ is in \mathfrak{B}_h , we have

$$\mathbb{P}(\mathcal{A}_n \models \mathfrak{B}_h) \geq \mathbb{P}(\mathcal{A}_n \models \Lambda(\mathfrak{G}_h \cap \mathfrak{T}_n)) = \sum_{\mathcal{G} \in \mathfrak{G}_h \cap \mathfrak{T}_n} \mathbb{P}(\mathcal{A}_n \models \Lambda(\mathcal{G})).$$

We used the fact that, by Lemma 5.1, two different elements of $\mathfrak{G}_h \cap \mathfrak{T}_n$ cannot be the h -backward substructures of the same transition structure.

Moreover, by Lemma 5.2, for $\mathcal{G} \in \mathfrak{T}_n$ we have

$$\mathbb{P}(\mathcal{A}_n \models \Lambda(\mathcal{G})) \geq \mathbb{P}(\mathcal{G}_h \text{ matches } \Lambda(\mathcal{G})) = \mathbb{P}(\mathcal{G}_h = \mathcal{G}).$$

Therefore,

$$\mathbb{P}(\mathcal{A}_n \models \mathfrak{B}_h) \geq \sum_{\mathcal{G} \in \mathfrak{G}_h \cap \mathfrak{T}_n} \mathbb{P}(\mathcal{G}_h = \mathcal{G}) = \mathbb{P}(\mathcal{G}_h \in \mathfrak{G}_h \cap \mathfrak{T}_n) \geq c_B.$$

This concludes the proof. \square

For any $n \geq 1$ and $p \in [n]$, let $\mathbf{B}_n(p)$ (or simply \mathbf{B}_n when there is no ambiguity on p) denote the set of templates associated with the $\lceil \log_2 \sqrt{n} \rceil$ -backward substructures \mathcal{B} from p of the n -state transition structures such that $|\text{Support}_B(\mathcal{B})| \leq d_1\sqrt{n}$ and $|L_h(\mathcal{B})| \geq d_2\sqrt{n}$.

6. FORWARD TREE AND SHORT CYCLE AROUND p

In this section, we continue the construction started in Section 5 by performing a forward exploration starting from $\delta(p, a)$ in a random transition structure that satisfies $\mathbf{B}_n(p)$.

More precisely, let $\mathcal{T} \in \mathbf{B}_n(p)$ of depth $h = \lceil \log_2 \sqrt{n} \rceil$. For a transition structure $\mathcal{A} = (n, \delta)$ that satisfies \mathcal{T} , we perform a breadth-first traversal of \mathcal{A} from $r := \delta(p, a)$: the states are discovered in order $r_0 = \delta(r, \varepsilon)$, $r_1 = \delta(r, a)$, $r_2 = \delta(r, b)$, $r_3 = \delta(r, aa)$, $r_4 = \delta(r, ab)$, \dots , where the words are taken in length-lexicographic order. This process halts as soon as one of the following events happens:

- (1) $r_i \in L_h$, where L_h is the h -th layer of \mathcal{T} , i.e. the states x of \mathcal{A} such that $d_p^-(x) = h$;
- (2) $r_i \in \text{Support}_{\mathbf{B}}(\mathcal{B}_h) \setminus L_h$, which is only possible for $r_0 = \delta(p, a)$, as \mathcal{A} satisfies \mathcal{T} ;
- (3) $r_i = r_j$ with $j < i$, i.e. we have a collision with a state already seen during the process;
- (4) $i = \lceil \sqrt{n} \rceil$, we force the process to halt after \sqrt{n} steps.

We call this process the *forward process* from p , and say that it is a *success* if it halts because Case (1) is triggered and it is a *failure* otherwise. It is deterministic for a given \mathcal{A} , and we will consider its probability of success for random transition structures.

Lemma 6.1. *There exists a constant $c_F > 0$ such that, for n sufficiently large, for any $p \in [n]$, and for any $\mathcal{T} \in \mathbf{B}_n(p)$, the forward process from p of a uniform random n -state transition structure conditioned to satisfy \mathcal{T} is a success with probability at least c_F .*

Proof. Let \mathcal{B}_h be the h -backward substructure from p associated to \mathcal{T} . By Lemma 4.1, in a random transition structure \mathcal{A} that satisfies \mathcal{T} , $\delta(p, a)$ is a uniform random element of $[n]$, and for every $x \notin \text{Support}_{\mathbf{B}}(\mathcal{T})$ and every $\alpha \in \Sigma$, $\delta(x, \alpha)$ is a uniform random element of $S_{\mathcal{T}} := L_h(\mathcal{B}_h) \cup ([n] \setminus \text{Support}_{\mathbf{B}}(\mathcal{B}_h))$, every choice being made independently.

We simulate the evolution of the process by generating the target of the transitions one by one according to the restrictions induced by \mathcal{T} , when needed, in the order of the traversal.

First, we want to upper bound the probability that the process fails because we reach the threshold of \sqrt{n} steps, i.e. it halts because of Case (4). As the first transition considered is $\delta(p, a)$, and the process immediately fails when $\delta(p, a) \in \text{Support}_{\mathbf{B}}(\mathcal{B}_h) \setminus L_h$, it continues with probability $1 - \frac{|\text{Support}_{\mathbf{B}}(\mathcal{B}_h) \setminus L_h|}{n}$.

For $1 \leq i \leq \lceil \sqrt{n} \rceil - 1$, if r_0, \dots, r_{i-1} avoided the halting conditions, then r_i , which is a uniform element of $S_{\mathcal{T}}$, also avoids the halting conditions with probability $1 - \frac{i + |L_h|}{|S_{\mathcal{T}}|}$, as it must be different from the previous r_j 's and not belong to L_h . By independence, it yields that the probability π_n that the process halts because of Case (4) is

$$\pi_n = \left(1 - \frac{|\text{Support}_{\mathbf{B}}(\mathcal{B}_h) \setminus L_h|}{n}\right) \prod_{i=1}^{\lceil \sqrt{n} \rceil - 1} \left(1 - \frac{i + |L_h|}{|S_{\mathcal{T}}|}\right) \leq \left(1 - \frac{|L_h|}{n}\right)^{\lceil \sqrt{n} \rceil - 1}.$$

As $\mathcal{T} \in \mathbf{B}_n(p)$, we have $|L_h| \geq d_2 \sqrt{n}$ and thus $\pi_n \leq (1 - d_2/\sqrt{n})^{\sqrt{n}-1}$. Moreover, since $\lim_n (1 - d_2/\sqrt{n})^{\sqrt{n}-1} = \exp(-d_2)$, for n large enough we have $\pi_n \leq \kappa$ by choosing, for instance, $\kappa := \frac{1}{2}(1 + \exp(-d_2)) \in (0, 1)$.

Secondly, if we condition the process to halt at some step $i < \lceil \sqrt{n} \rceil$ then r_i is chosen uniformly at random in $\{r_0, \dots, r_{i-1}\} \cup L_h$. Hence, given it halts at step i , the probability of success is

$$\frac{|L_h|}{i + |L_h|} \geq \frac{d_2 \sqrt{n}}{\sqrt{n} + d_1 \sqrt{n}} = \frac{d_2}{1 + d_1}.$$

Finally, since this lower bound does not depend on i , the probability that it halts because of Case (1), i.e. the process is a success, is at least $\frac{d_2}{1+d_1}$ times the probability that the process halts before $\lceil \sqrt{n} \rceil$ steps. Hence, it is at least $c_F = \frac{d_2}{1+d_1}(1 - \kappa) > 0$, concluding the proof. \square

Let $(u_i)_{i \geq 0}$ denote the sequence of words on Σ in the length lexicographic order, so that $u_0 = \varepsilon$, $u_1 = a$, $u_2 = b$, $u_3 = aa$, \dots . For any n -state transition structure $\mathcal{A} = (n, \delta)$ that satisfies a template $\mathcal{T} = (n, \gamma)$ of $\mathbf{B}_n(p)$ and whose forward process is successful and halts, at step i , we associate the template $\mathcal{T}_{\mathcal{A}} = (n, \lambda)$ as follows. For any $x \in [n]$ and $\alpha \in \Sigma$,

- for all $j < i$ and all $\alpha \in \{a, b\}$ such that $u_j \alpha$ belongs $\{u_0, \dots, u_i\}$, $\lambda(x, \alpha) = \{\delta(x, \alpha)\}$ with $x = \delta(p, au_j)$,
- otherwise, $\lambda(x, \alpha) = \gamma(x, \alpha)$.

In other words, starting from the template \mathcal{T} , we force the targets of all the transitions explored during the forward process. It is direct to establish that every transition structure that satisfies \mathcal{A} has the same $\lceil \log_2 \sqrt{n} \rceil$ -backward substructure than \mathcal{A} and exactly the same forward process, which is successful.

Let $\mathbf{C}_n(p)$ denote the set of templates $\mathcal{T}_{\mathcal{A}}$ for the transition structures \mathcal{A} that satisfy $\mathcal{T} \in \mathbf{B}_n(p)$ and that have a successful forward process from p . For such a transition structure \mathcal{A} whose process halts after i steps, define $\text{Support}_{\mathbf{C}}(\mathcal{A}) = \text{Support}_{\mathbf{B}}(\mathcal{A}) \cup \{\delta(p, au_j) : 0 \leq j < i\}$ be the set of states visited during the backward construction and the forward process.

The next lemma is a direct consequence of Lemma 5.1 and of the way we constrained the forward tree in the construction of $\mathcal{T}_{\mathcal{A}}$.

Lemma 6.2. *Two n -state transition structures $\mathcal{A} = (n, \delta)$ and $\mathcal{A}' = (n, \delta')$ satisfy the same template $\mathcal{T} \in \mathbf{C}_n(p)$ if and only if \mathcal{A} and \mathcal{A}' have same $\lceil \log_2 \sqrt{n} \rceil$ -backward substructure and the forward processes from p of both \mathcal{A} and \mathcal{A}' halt successfully at the same step i , with $\delta(p, au_j) = \delta'(p, au_j)$ for every $j \in \{0, \dots, i-1\}$.*

For $\mathcal{T}_{\mathcal{A}} \in \mathbf{C}_n(p)$ a template associated to a transition structure $\mathcal{A} = (n, \delta)$ whose forward process from p halts successfully after i steps, we let $\text{Support}_{\mathbf{C}}(\mathcal{T}_{\mathcal{A}}) = \text{Support}_{\mathbf{C}}(\mathcal{A})$ denote the set $\text{Support}_{\mathbf{B}}(\mathcal{A}) \cup \{\delta(p, au_j) : 0 \leq j < i\}$, which we call the *cycle support* of $\mathcal{T}_{\mathcal{A}}$ (or of \mathcal{A}). We added the states discovered during the process to $\text{Support}_{\mathbf{B}}(\mathcal{A})$.

Finally, observe that the last step of a successful process builds a cycle around p as $\delta(p, au_i) \in L_h$ and as there is a path from any state of L_h to p , by construction. Let v be the smallest word for the length-lexicographic order that labels a path from $\delta(p, au_i) \in L_h$ to p , then $au_i v$ labels a cycle around p in \mathcal{A} . Define $w_{\mathcal{A}}(p) := au_i v$, or just $w_{\mathcal{A}}$ if p is clear from the context.

Moreover, this word only depends on $\mathcal{T}_{\mathcal{A}}$ since it only uses transitions determined by the template, by construction of the backward and forward traversals: two transition structures \mathcal{A} and \mathcal{A}' satisfying the same $\mathcal{T} \in \mathbf{C}_n(p)$ are such that $w_{\mathcal{A}} = w_{\mathcal{A}'}$. So we can define this word for a given $\mathcal{T} \in \mathbf{C}_n(p)$ as $w_{\mathcal{T}} := w_{\mathcal{A}}$ with no ambiguity.

The properties we need in the sequel are summarized in the next statement.

Proposition 6.3. *There exists $c > 0$ such that for n sufficiently large and $p \in [n]$, a random n -state transition structure satisfies $\mathbf{C}_n(p)$ with probability at least c .*

Furthermore, any transition structure \mathcal{A} that satisfies a template $\mathcal{T} = (n, \lambda) \in \mathbf{C}_n(p)$ admits a cycle around p labeled by the word $aw_{\mathcal{T}}$, which only visits states in $\text{Support}_{\mathbf{C}}(\mathcal{T})$ and has length between h and $2h$, where $h = \lceil \log_2 \sqrt{n} \rceil$. We also have $|\text{Support}_{\mathbf{C}}(\mathcal{T})| \leq (d_1 + 1)\sqrt{n}$, and for any $x \notin \text{Support}_{\mathbf{C}}(\mathcal{T})$ and any $\alpha \in \Sigma$, $\lambda(x, \alpha) = L_h(\mathcal{T}) \cup ([n] \setminus \text{Support}_{\mathbf{B}}(\mathcal{T}))$.

Proof. As both the h -backward substructure construction and the forward process are deterministic, a given n -state transition structure can satisfy at most one template in $\mathbf{C}_n(p)$. Hence, if we partition according to the elements of $\mathbf{C}_n(p)$, the law of total probabilities and Lemma 6.1 yield the first statement of the proposition.

The second part is a consequence of the construction of $\mathcal{T}_{\mathcal{A}}$ and of the fact that if the forward process is successful, it halts before \sqrt{n} states are discovered, and thus after building a forward tree of depth at most h . \square

7. FORMING THE b -CYCLES

From now on, we fix some integer $d \geq 1$. In this section, we consider a random transition structure conditioned to satisfying a template $\mathcal{T} \in \mathbf{C}_n(p)$ with cycle around p labeled by aw , with $w := w_{\mathcal{T}}$ and $|w| \in \Theta(\log n)$. The statements of this part articulate as follows.

We consider the path $\mathcal{P} = q \xrightarrow{w} p_1 \xrightarrow{aw} p_2 \rightsquigarrow \dots \rightsquigarrow p_d$ depicted in Figure 2. We first establish in Lemma 7.1 that with high probability it does not intersect the support of the template \mathcal{T} and it does not go twice through the same state.

Starting at state x , we grow what we call a b -thread by successively drawing the outgoing b -transitions until we cycle back. Then in Lemma 7.2, we show that with visible probability we can grow a b -thread from p and each p_i such that these threads have size in $O(\sqrt{n})$, do not intersect the support of the template \mathcal{T} nor the path \mathcal{P} and are pairwise disjoint.

If we condition the cycles formed by the b -threads to have length in $[\sqrt{n}, 2\sqrt{n}]$, we furthermore establish that these lengths are uniformly drawn in this interval and that it still happens with visible probability in Proposition 7.4.

7.1. Growing the b -threads. Let $\mathcal{A} = (n, \delta)$ be a n -state transition structure that satisfies a template $\mathcal{T} \in \mathbf{C}_n(p)$, for some $p \in [n]$. For any $q \in [n]$, we associate to \mathcal{A} its *starting states tuple* $s(\mathcal{A}, q)$ which is either a tuple of $d + 1$ states, or \perp in case the construction of the tuple failed. To build $s(\mathcal{A}, q)$ consider the path labeled in \mathcal{A} that starts at state q and labeled by the word

$$u = w_{\mathcal{A}}(aw_{\mathcal{A}})^{d-1}$$

where $aw_{\mathcal{A}}$ labels the cycle around p in \mathcal{A} defined in the previous sections. The word u has length $|u| = d|w_{\mathcal{A}}| + d - 1$, and go through all the states $x_i := \delta(q, v_i)$, where v_i is the prefix of length i of u and $0 \leq i \leq |u|$. Letting $X_q(\mathcal{A}) = \{x_i : 0 \leq i \leq |u|\}$, the construction fails when $X_q(\mathcal{A}) \cap \text{Support}_{\mathcal{C}}(\mathcal{A}) \neq \emptyset$ or $|X_q(\mathcal{A})| < |u| + 1$; in other words, it fails if the path uses a state of $\text{Support}_{\mathcal{C}}(\mathcal{A})$ or if it goes twice through the same state. If the construction fails, we set $s(\mathcal{A}, q) = \perp$, otherwise

$$s(\mathcal{A}, q) := \left(p, \delta(q, w_{\mathcal{A}}), \delta(q, w_{\mathcal{A}}(aw_{\mathcal{A}})^1), \dots, \delta(q, w_{\mathcal{A}}(aw_{\mathcal{A}})^{d-1}) \right).$$

Lemma 7.1. *For any $\epsilon > 0$, for any n sufficiently large, for all $p \in [n]$, if q is chosen uniformly at random in $[n]$, then for any $\mathcal{T} \in \mathbf{C}_n(p)$ we have*

$$\mathbb{P}(s(\mathcal{A}, q) \neq \perp \mid \mathcal{A} \models \mathcal{T}) \geq 1 - \epsilon.$$

Proof. Observe that the construction immediately fails for $\mathcal{A} \models \mathcal{T}$ if the starting state of the path $x_0 = q$ is in $\text{Support}_{\mathcal{C}}(\mathcal{T})$, as $\text{Support}_{\mathcal{C}}(\mathcal{A}) = \text{Support}_{\mathcal{C}}(\mathcal{T})$. This happens with probability $|\text{Support}_{\mathcal{C}}(\mathcal{T})|/n$ as q is chosen uniformly at random in $[n]$. As stated in Proposition 6.3, in a random \mathcal{A} satisfying \mathcal{T} , the transitions starting from $x \notin \text{Support}_{\mathcal{C}}(\mathcal{T})$

must arrive in $S := L_h(\mathcal{T}) \cup ([n] \setminus \text{Support}_B(\mathcal{T}))$. Furthermore their targets are all chosen uniformly at random and independently in S , according to Lemma 4.1. For $i \in \{0, \dots, |u|\}$, let $Y_i := \text{Support}_C(\mathcal{T}) \cup \{x_0, \dots, x_i\}$. Observe that for any $i \in \{1, \dots, |u|\}$, if the path for the first $i-1$ letters of u did not produce a failure, the target of a transition starting from x_{i-1} is an element of S taken uniformly at random. Hence, we have, as $|\text{Support}_C(\mathcal{T}) \cap S| \leq |\text{Support}_C(\mathcal{T})|$:

$$\mathbb{P}(x_i \notin Y_{i-1} \mid \forall j \leq i-1, x_j \notin Y_{j-1}) \geq 1 - \frac{|\text{Support}_C(\mathcal{T})| + i}{|S|} \geq 1 - \frac{|\text{Support}_C(\mathcal{T})| + |u|}{|S|}.$$

Combining with the probability that $x_0 = q \notin \text{Support}_C(\mathcal{T})$ this yields

$$\mathbb{P}(s(\mathcal{A}, q) \neq \perp \mid \mathcal{A} \models \mathcal{T}) \geq \left(1 - \frac{|\text{Support}_C(\mathcal{T})| + |u|}{|S|}\right)^{|u|+1}.$$

By Proposition 6.3, $|u| \leq 2\lceil \log_2 \sqrt{n} \rceil$ and $|\text{Support}_B(\mathcal{T})| \leq |\text{Support}_C(\mathcal{T})| \leq (d_1 + 1)\sqrt{n}$, so that $|S| \geq n - (d_1 + 1)\sqrt{n}$ and we have

$$\mathbb{P}(s(\mathcal{A}, q) \neq \perp \mid \mathcal{A} \models \mathcal{T}) \geq \left(1 - \frac{(d_1 + 1)\sqrt{n} + 2\lceil \log_2 \sqrt{n} \rceil}{n - (d_1 + 1)\sqrt{n}}\right)^{2\lceil \log_2 \sqrt{n} \rceil + 1}.$$

This concludes the proof, as it tends to 1 when n tends to infinity. \square

If \mathcal{A} satisfies \mathcal{T} and $s(\mathcal{A}, q) \neq \perp$, we write $s(\mathcal{A}, q) = (p_0, p_1, \dots, p_d)$, with $p_0 = p$ by construction. Let u_i denote the i -th letter of the word u and let also $\mathcal{P}_q(\mathcal{A}) = x_0 \xrightarrow{u_1} x_1 \xrightarrow{u_2} \dots \xrightarrow{u_{|u|}} x_{|u|}$ denote the path labeled by $u = u_1 u_2 \dots u_{|u|}$ and starting from $x_0 = q$. Recall that $X_q(\mathcal{A})$ is the set of states of $\mathcal{P}_q(\mathcal{A})$.

We consider the *thread process* of \mathcal{A} which, if successful, consists in building in order the sets $E_i = \{\delta(p_i, b^j) : j \geq 0\}$ for $i \in \{0, \dots, d\}$ as follows. For i from 0 to d we start with $E_i = \{p_i\}$ and then iteratively add $\delta(p_i, b^j)$ for $j \geq 0$ until:

- (1) $\delta(p_i, b^j) \in \cup_{k=0}^{i-1} E_k$, in which case the process halts and is a failure;
- (2) or $\delta(p_i, b^j) \in \text{Support}_C(\mathcal{T}) \cup X_q(\mathcal{A})$, in which case the process halts and is a failure;
- (3) or $\delta(p_i, b^j) \in E_i$, in which case the process halts and is a failure if $|E_i| \notin \llbracket 2\sqrt{n}, 3\sqrt{n} \rrbracket$. If not there are two cases: if $i < d$ the process starts building E_{i+1} , and if $i = d$ it halts with a success, as all the b -threads are successfully built.

Define $\lambda(\mathcal{A}, q) = (|E_0|, |E_1|, \dots, |E_d|)$ the tuple of the b -threads' lengths if the process is successful and $\lambda(\mathcal{A}, q) = \perp$ if it fails. Observe that if successful, then the E_i 's are pairwise disjoint, and do not intersect $\text{Support}_C(\mathcal{T}) \cup X_q(\mathcal{A})$.

The following statement is a variation on the classical Birthday Problem.

Lemma 7.2. *There exists $c_\lambda > 0$ such that for any n sufficiently large, for any $p \in [n]$, for any $\mathcal{T} \in \mathbf{C}_n(p)$ and for any path $\mathcal{P} = x_0 \xrightarrow{u_1} x_1 \xrightarrow{u_2} \dots \xrightarrow{u_{|u|}} x_{|u|}$ labeled by $u = w_{\mathcal{T}}(aw_{\mathcal{T}})^{d-1}$ such that the x_i 's are pairwise distinct and not in $\text{Support}_C(\mathcal{T})$ we have*

$$\mathbb{P}(\lambda(\mathcal{A}, x_0) \neq \perp \mid \mathcal{A} \models \mathcal{T} \text{ and } \mathcal{P}_{x_0}(\mathcal{A}) = \mathcal{P}) \geq c_\lambda.$$

Proof. Conditioning by $\mathcal{A} \models \mathcal{T}$ and $\mathcal{P}_{x_0}(\mathcal{A}) = \mathcal{P}$ is exactly conditioning on satisfying the template $\mathcal{T}_{\mathcal{P}} = (n, \gamma_{\mathcal{P}})$ obtained from $\mathcal{T} = (n, \gamma)$ by setting $\gamma_{\mathcal{P}}(x_i, u_{i+1}) = \{x_{i+1}\}$ for all $i \in \{0, \dots, |u| - 1\}$, and $\gamma_{\mathcal{P}}(x, \alpha) = \gamma(x, \alpha)$ for all other transitions.

Let $S := L_h(\mathcal{T}) \cup ([n] \setminus \text{Support}_C(\mathcal{T}))$. For every $i \in \{1, \dots, d\}$, by construction $p_i = x_{i|aw_{\mathcal{T}}| - 1}$ and $\gamma_{\mathcal{P}}(p_i, b) = \gamma(p_i, b) = S$ since a labels the transition outgoing from p_i in

\mathcal{P} for $i < d - 1$ and p_d is the end of \mathcal{P} . Also, $\gamma_{\mathcal{P}}(p_0, b) = \gamma(p_0, b) = \gamma(p, b) = [n]$, as the template gives no constraint on $p \xrightarrow{b}$.

So during the thread process, as long as we are adding a new state $x = \delta(p_i, b^j)$ in E_i , the transition $\delta(x, b)$ is a uniform element of S (or of $[n]$ if $i = j = 0$), independently of the previous steps. Hence if we are not in the case $i = j = 0$, out of the $|S|$ possibilities, $r_i := |\cup_{k=0}^{i-1} E_k|$ possibilities trigger a failure because of Condition (1), $s := |S \cap \text{Support}_{\mathcal{C}}(\mathcal{T})| + |X_q(\mathcal{A})|$ possibilities trigger a failure because of Condition (2), and j possibilities complete the process for E_i because we cycle back on a previously seen element of E_i . So for $i \geq 1$ and $t < 3\sqrt{n}$, the probability $\pi_i(t)$ that neither Condition (1) nor Condition (2) were triggered and $|E_i| = t$, conditioned on the fact that E_j were successfully built for all $j < i$ is

$$\pi_i(t) = \underbrace{\frac{|S| - r_i - s}{|S|}}_{\text{choice of } \delta(p_i, b)} \times \underbrace{\frac{|S| - r_i - s - 1}{|S|}}_{\text{choice of } \delta(p_i, b^2)} \times \dots \times \underbrace{\frac{|S| - r_i - s - (t-2)}{|S|}}_{\text{choice of } \delta(p_i, b^{t-1})} \times \underbrace{\frac{t}{|S|}}_{\text{cycling back}}.$$

Hence

$$\pi_i(t) = \frac{t}{|S|} \prod_{j=0}^{t-2} \left(1 - \frac{r_i + s + j}{|S|}\right).$$

For $2\sqrt{n} \leq t \leq 3\sqrt{n}$ we have

$$\pi_i(t) \geq \frac{t}{|S|} \left(1 - \frac{r_i + s + 3\sqrt{n}}{|S|}\right)^{3\sqrt{n}}.$$

Moreover, as at most d threads have been completed so far, $r_i \leq 3d\sqrt{n}$, and for n sufficiently large, $|X_q(\mathcal{A})| \leq \sqrt{n}$ and thus $s \leq (d_1 + 2)\sqrt{n}$, by Proposition 5.6. Observe that we also have $|S| \geq n - (d_1 + 1)\sqrt{n}$, which is greater than $\frac{1}{2}n$ for n sufficiently large. Therefore, using $\mu := 2d_1 + 6d + 10$ and the fact that $\lim_m (1 - \mu/m)^m = \exp(-\mu)$, we have

$$\pi_i(t) \geq \frac{t}{|S|} \left(1 - \frac{\mu}{\sqrt{n}}\right)^{3\sqrt{n}} \geq \frac{te^{-3\mu}}{2n}.$$

This yields that

$$\sum_{t=\lceil 2\sqrt{n} \rceil}^{\lfloor 3\sqrt{n} \rfloor} \pi_i(t) \geq \frac{e^{-3\mu}}{2n} \sum_{t=\lceil 2\sqrt{n} \rceil}^{\lfloor 3\sqrt{n} \rfloor} t \geq \frac{e^{-3\mu}}{2n} (\lfloor 3\sqrt{n} \rfloor - \lceil 2\sqrt{n} \rceil + 1) \lfloor 2\sqrt{n} \rfloor.$$

As the limit of the right hand term is $e^{-3\mu} > 0$, there exists $\kappa > 0$ such that, for n sufficiently large we have $\sum_{t=\lceil 2\sqrt{n} \rceil}^{\lfloor 3\sqrt{n} \rfloor} \pi_i(t) \geq \kappa$.

For the case $i = 0$, as $r_0 = 0$ and $\gamma_{\mathcal{P}}(p_0, b) = [n]$, we have almost the same formula:

$$\pi_0(t) = \frac{t}{|S|} \left(1 - \frac{s}{n}\right) \prod_{j=1}^{t-2} \left(1 - \frac{s+j}{|S|}\right).$$

With the same technique as before, we can find some positive constant κ_0 such that $\sum_{t=\lceil 2\sqrt{n} \rceil}^{\lfloor 3\sqrt{n} \rfloor} \pi_0(t) \geq \kappa_0$, for n sufficiently large. Combining the results we obtain that

$$\mathbb{P}(\lambda_{x_0}(\mathcal{A}) \neq \perp \mid \mathcal{A} \models \mathcal{T} \text{ and } \mathcal{P}_{x_0}(\mathcal{A}) = \mathcal{P}) \geq \kappa_0 \kappa^d,$$

concluding the proof. \square

We now write the conditions in terms of templates, by naming the states encountered during the thread process and forcing the associated transitions. For $n \geq 1$ and $p, q \in [n]$, let $\mathbf{T}_n(p, q)$ denote the set of templates $\mathcal{T} = (n, \gamma)$ such that:

- \mathcal{T} satisfies a template $\hat{\mathcal{T}} = (n, \tilde{\gamma})$ of $\mathbf{C}_n(p)$, with a cycle around p labeled by the word $aw_{\mathcal{T}}$. Hence $\text{Support}_{\mathbf{C}}(\mathcal{T}) = \text{Support}_{\mathbf{C}}(\hat{\mathcal{T}})$ and, to simplify the notations, define $w := w_{\mathcal{T}} = w_{\hat{\mathcal{T}}}$.
- We write $x_0 = q$ and for $u = w(aw)^{d-1} = u_1 u_2 \cdots u_{|u|}$, there exist pairwise distinct states $x_1, x_2, \dots, x_{|u|}$ in $[n] \setminus \text{Support}_{\mathbf{C}}(\mathcal{T})$ such that for all $i \in \{0, \dots, |u| - 1\}$, $\gamma(x_i, u_{i+1}) = \{x_{i+1}\}$. This determines the path starting at q and labeled by u of any \mathcal{A} that satisfies \mathcal{T} . Let $X_q(\mathcal{T}) = \{x_i : 0 \leq i \leq |u|\}$.
- We write $p_{0,1} = p$ and $p_{i,1} = x_{i|aw|-1}$ for $1 \leq i \leq d$. For any $i \in \{0, \dots, d\}$, there exists an integer λ_i such that $2\sqrt{n} \leq \lambda_i \leq 3\sqrt{n}$ and states $p_{i,2}, \dots, p_{i,\lambda_i}$ in $[n] \setminus (\text{Support}_{\mathbf{C}}(\mathcal{T}) \cup X_q(\mathcal{T}))$ such that:
 - The $p_{i,j}$'s are pairwise distinct.
 - For all $i \in \{0, \dots, d\}$, for all $j \in \{1, \dots, \lambda_j - 1\}$, $\gamma(p_{i,j}, b) = \{p_{i,j+1}\}$. This determines the b -threads starting from $p_{0,1}, \dots, p_{d,1}$. Let $E_i = \{p_{i,j} : j \in [\lambda_i]\}$.
 - For all $i \in \{0, \dots, d\}$, $\gamma(p_{i,\lambda_i}, b) = E_i$, to ensure we cycle back in each b -thread.
- For every other transition (x, α) we have $\gamma(x, \alpha) = \hat{\gamma}(x, \alpha)$.

As the construction follows the constructions of the path and of the b -threads, we readily have that $\mathcal{A} \models \mathbf{T}_n(p, q)$ if and only if $\mathcal{A} \models \mathbf{C}_p(n)$ and $\lambda(\mathcal{A}, q) \neq \perp$. As the constructions are deterministic, for given $p, q \in [n]$, a given n -state transition structure \mathcal{A} cannot satisfy more than one template of $\mathbf{T}_n(p, q)$. Moreover, all the transition structures \mathcal{A} that satisfy a given $\mathcal{T} \in \mathbf{T}_n(p, q)$ have same $\lambda(\mathcal{A}, q) = (\lambda_0, \lambda_1, \dots, \lambda_d)$, which we can therefore write $\lambda(\mathcal{T}, q)$ with no ambiguity.

Proposition 7.3. *There exists $c_T > 0$ such that, for n sufficiently large and for $p \in [n]$, if q is a uniform random element of $[n]$ and \mathcal{A} is a uniform n -state transition structure, taken independently, then $\mathbb{P}(\mathcal{A} \models \mathbf{T}_n(p, q)) \geq c_T$.*

Proof. We partition $\mathbf{T}_n(p, q)$ according to the possible valid paths as follows. For given $\mathcal{T} \in \mathbf{T}_n(p, q)$, the path $\mathcal{P}_q(\mathcal{A})$ is the same for every \mathcal{A} that satisfies \mathcal{T} , so we can define $\mathcal{P}_q(\mathcal{T}) := \mathcal{P}_q(\mathcal{A})$ for any such \mathcal{A} . Let $\mathfrak{P}_q(n)$ denote the set of all possible paths:

$$\mathfrak{P}_q(n) = \{\mathcal{P}_q(\mathcal{T}) : \mathcal{T} \in \mathbf{T}_n(p, q)\}.$$

We have

$$\mathbb{P}(\mathcal{A} \models \mathbf{T}_n(p, q)) = \sum_{\mathcal{P} \in \mathfrak{P}_q(n)} \sum_{\mathcal{T} \in \mathbf{T}_n(p, q)} \mathbb{P}(\mathcal{A} \models \mathcal{T} \text{ and } \mathcal{P}_q(\mathcal{T}) = \mathcal{P}).$$

Recall that $\mathcal{A} \models \mathbf{T}_n(p, q)$ if and only if $\mathcal{A} \models \mathbf{C}_p(n)$ and $\lambda(\mathcal{A}, q) \neq \perp$. This yields

$$\mathbb{P}(\mathcal{A} \models \mathbf{T}_n(p, q)) = \sum_{\mathcal{P} \in \mathfrak{P}_q(n)} \sum_{\mathcal{T} \in \mathbf{C}_p(n)} \mathbb{P}(\mathcal{A} \models \mathcal{T} \text{ and } \lambda(\mathcal{A}, q) \neq \perp \text{ and } \mathcal{P}_q(\mathcal{T}) = \mathcal{P}).$$

For any $\mathcal{T} \in \mathbf{C}_p(n)$ and any $\mathcal{P} \in \mathfrak{P}_q(n)$, by Lemma 7.2, for n sufficiently large we have

$$\mathbb{P}(\mathcal{A} \models \mathcal{T} \text{ and } \lambda(\mathcal{A}, q) \neq \perp \text{ and } \mathcal{P}_q(\mathcal{T}) = \mathcal{P}) \geq c_\lambda \mathbb{P}(\mathcal{A} \models \mathcal{T} \text{ and } \mathcal{P}_q(\mathcal{T}) = \mathcal{P}).$$

By definition of $s_q(\mathcal{A})$, a given transition structure that satisfies $\mathcal{T} \in \mathbf{C}_p(n)$ has a path in $\mathfrak{P}_q(n)$ if and only if $s_q(\mathcal{A}) \neq \perp$. Thus

$$\sum_{\mathcal{P} \in \mathfrak{P}_q(n)} \mathbb{P}(\mathcal{A} \models \mathcal{T} \text{ and } \mathcal{P}_q(\mathcal{T}) = \mathcal{P}) = \mathbb{P}(\mathcal{A} \models \mathcal{T} \text{ and } s_q(\mathcal{A}) \neq \perp).$$

Moreover, taking $\epsilon = \frac{1}{2}$ in Lemma 7.1, for n sufficiently large we have

$$\mathbb{P}(\mathcal{A} \models \mathcal{T} \text{ and } s_q(\mathcal{A}) \neq \perp) \geq \frac{1}{2} \mathbb{P}(\mathcal{A} \models \mathcal{T}).$$

Putting all together yields

$$\mathbb{P}(\mathcal{A} \models \mathbf{T}_n(p, q)) \geq \frac{c_\lambda}{2} \sum_{\mathcal{T} \in \mathbf{C}_p(n)} \mathbb{P}(\mathcal{A} \models \mathcal{T}) = \frac{c_\lambda}{2} \mathbb{P}(\mathcal{A} \models \mathbf{C}_p(n)).$$

This concludes the proof by Proposition 6.3. \square

Let $\mathcal{T} \in \mathbf{T}_n(p, q)$ with $\lambda(\mathcal{T}, q) = (\lambda_0, \dots, \lambda_d)$. If $\mathcal{A} \models \mathcal{T}$ then at the end of each b -thread from $p_{i,1}$, the b -transition outgoing from p_{i,λ_i} ends in an element $p_{i,j}$ of E_i , forming a b -cycle of length $\ell_i = \lambda_i - j + 1$. Let $\ell(\mathcal{A}, p, q)$ be the $d+1$ -tuple of the b -cycle lengths (ℓ_0, \dots, ℓ_d) if all its coordinates are in $\llbracket \sqrt{n}, 2\sqrt{n} \rrbracket$, and $\ell(\mathcal{A}, p, q) = \perp$ otherwise. We also set $\ell(\mathcal{A}, p, q) = \perp$ if \mathcal{A} does not satisfy $\mathbf{T}_n(p, q)$. If $\ell(\mathcal{A}, p, q) = (\ell_0, \dots, \ell_d)$, we define $\ell_i(\mathcal{A}) := \ell_i$ to directly access its coordinates.

Proposition 7.4. *There exists $c_\ell > 0$ such that for n sufficiently large and $p \in [n]$, if q is taken uniformly in $[n]$ and \mathcal{A} is a uniform n -state transition structure taken independently then $\mathbb{P}(\ell(\mathcal{A}, p, q) \neq \perp) \geq c_\ell$. Moreover, conditioned on $\ell(\mathcal{A}, p, q) \neq \perp$, the random variables ℓ_i are independent uniform integers of $\llbracket \sqrt{n}, 2\sqrt{n} \rrbracket$.*

Proof. By definition, we have $\ell(\mathcal{A}, p, q) \neq \perp$ only for transition structures \mathcal{A} such that $\mathcal{A} \models \mathbf{T}_n(p, q)$. Let $\mathcal{T} \in \mathbf{T}_n(p, q)$, let $s_q(\mathcal{T}) = (p_0, p_1, \dots, p_d)$ denote its tuple of starting states and let $\lambda_q(\mathcal{T}) = (\lambda_0, \dots, \lambda_p)$ denote its tuple of b -thread lengths. If $\mathcal{A} = (n, \delta)$ is a random transition structure conditioned on satisfying \mathcal{T} , then, by definition of $\mathbf{T}_n(p, q)$, for every $i \in \{0, \dots, d\}$, $\delta(p_i, b^{\lambda_i})$ is a uniform element of $E_i = \{\delta(p_i, b^j) : j \geq 0\}$, with $|E_i| = \lambda_i$, and they are all independent. Hence each b -cycle length $\hat{\ell}_i$ is a uniform random integer in $[\lambda_i]$, and they are all independent.

By definition of $\mathbf{T}_n(p, q)$, for every $i \in \{0, \dots, d\}$, we have $2\sqrt{n} \leq \lambda_i \leq 3\sqrt{n}$. In particular $\llbracket \sqrt{n}, 2\sqrt{n} \rrbracket \subseteq \llbracket 1, \lambda_i \rrbracket$, and the probability that a uniform element of $\llbracket 1, \lambda_i \rrbracket$ is in $\llbracket \sqrt{n}, 2\sqrt{n} \rrbracket$ is at least $\frac{1}{4}$ (it is lower-bounded by a quantity that tends to $\frac{1}{3}$). Hence the probability that all the $\hat{\ell}_i$'s are in the valid range $\llbracket \sqrt{n}, 2\sqrt{n} \rrbracket$ is at least 4^{-d-1} . Since this lower bound does not depend on $\mathcal{T} \in \mathbf{T}_n(p, q)$, this proves by the law of total probabilities that

$$\mathbb{P}(\ell(\mathcal{A}, p, q) \neq \perp) = \sum_{\mathcal{T} \in \mathbf{T}_n(p, q)} \mathbb{P}(\ell(\mathcal{A}, p, q) \neq \perp \mid \mathcal{A} \models \mathcal{T}) \mathbb{P}(\mathcal{A} \models \mathcal{T}) \geq 4^{-d-1} \mathbb{P}(\mathcal{A} \models \mathbf{T}_n(p, q)).$$

By Proposition 7.3, this yields $\mathbb{P}(\ell(\mathcal{A}, p, q) \neq \perp) \geq c_\ell$, with $c_\ell = 4^{-d-1} c_T$.

The second part of the statement is just a consequence of the fact that a uniform element of $[\lambda_i]$ conditioned to be in $\llbracket \sqrt{n}, 2\sqrt{n} \rrbracket$ is a uniform element of $\llbracket \sqrt{n}, 2\sqrt{n} \rrbracket$. And as direct computation shows, the independence is preserved if we consider all the ℓ_i 's together. \square

Proposition 7.4 is our main probabilistic result on random transition structures. It states that our global construction succeeds with visible probability and produces $d+1$ b -cycles of uniform and independent length in $\llbracket \sqrt{n}, 2\sqrt{n} \rrbracket$. These cycles are linked to p in a way that can be exploited during the accessible subset construction when we add the transition $p \xrightarrow{a} q$.

8. SUPER-POLYNOMIAL GROWTH OF THE SUBSET CONSTRUCTION

A uniform n -state almost deterministic transition structure $\mathcal{A} = (n, \delta, p \xrightarrow{a} q)$ is obtained by choosing δ , p and q uniformly at random and independently. If we furthermore choose the starting state i_0 uniformly at random and independently, we can use Proposition 7.4 to establish that the accessible subset construction has a super-polynomial number of states with visible probability. This is the focus of this section.

First, Grusho [Gru73] established that with high probability a uniform n -state transition structure has a unique terminal strongly connected component. When there is only one such strongly connected component, it is necessarily accessible from any state. We also rely on a result of Cai and Devroye [CD17, Theorem 2], which ensures that with high probability there are no cycles of length larger than $\frac{1}{4} \log_2 n$ outside this unique terminal strongly connected component. More precisely, they show that the length of the longest cycle outside the unique terminal strongly connected component of a random n -state transition structure is in $O_p(1)$ with the notations of [JLR00]. By [JLR00, Remark 1.3], this in particular implies that for any $\omega(n) \rightarrow \infty$, all the cycles outside the accessible part have length at most $\omega(n)$ with high probability.

Observe that if $\mathcal{A} \models \mathbf{T}_n(p, q)$, then it has a word $aw_{\mathcal{A}}$ that labels a cycle around p of length at most $h = \lceil \log_2 \sqrt{n} \rceil$. By Proposition 7.4, this happens with visible probability. The conjunction of a high-probability event with a visible event being a visible event, we directly have the following lemma.

Lemma 8.1. *If \mathcal{A} is a n -state transition structure, and p, q, i_0 are states in $[n]$, all being taken uniformly at random and independently in their respective sets, then with visible probability, $\ell(\mathcal{A}, p, q) \neq \perp$ and p is accessible from i_0 .*

At this point, to prove our first theorem, on the accessible powerset construction applied to a random transition structure with an added transition, we need a result of probabilistic number theory. Tóth in [Tót02] generalizes the folklore result that two independent random numbers in $[N]$ are coprime with probability that tends to $\frac{6}{\pi^2}$ to a fixed number of independent random numbers.

Theorem 8.2 [Tót02]. *For any $k \geq 2$, there exists some constant $A_k > 0$ such that k integers taken uniformly at random and independently in $[N]$ are pairwise coprime with probability that tends to A_k as N tends to infinity, with*

$$A_k = \prod_{p \text{ prime}} \left(1 + \frac{k}{p-1}\right) \left(1 - \frac{1}{p}\right)^k.$$

This theorem was proven using probabilistic arguments in [CB03, Theorem 3.3]. We adapt this proof to obtain the following corollary.

Corollary 8.3. *For any $d \geq 1$, $d+1$ integers taken uniformly at random and independently in $[\sqrt{n}, 2\sqrt{n}]$ are pairwise coprime with probability that tends to A_{d+1} as n tends to infinity.*

Proof. For any real $x > 0$, let P_x be the set of prime numbers smaller than or equal to x : $P_x = \{p \in \mathbb{Z}_{\geq 1} : p \leq x \text{ and } p \text{ prime}\}$. For any $n \geq 2$, let i_n and x_n be defined by

$$i_n = \max\{i \in \mathbb{Z}_{\geq 0} : \prod_{p \in P_i} p \leq \sqrt[3]{n}\} \quad \text{and} \quad x_n = \prod_{p \in P_{i_n}} p.$$

Observe that $(i_n)_{n \geq 2}$ and $(x_n)_{n \geq 2}$ both tend to infinity as n tends to infinity.

We now follow the idea of Cai and Bach [CB03]. We consider the uniform distribution on $[\sqrt{n}, 2\sqrt{n}]^{d+1}$ and define the following events:

- E_n : the $d + 1$ integers are pairwise coprime,
- for all prime p , G_n^p : at most one of the $d + 1$ integers is divisible by p ,
- for all prime p , B_n^p : at least two of the $d + 1$ integers are divisible by p .

Notice that G_n^p is the complement of B_n^p , and that $E_n = \bigcap_{\substack{p \leq 2\sqrt{n} \\ p \text{ prime}}} G_n^p$. By separating the cases $p \leq i_n$ and $p > i_n$, let us introduce the events $R_n = \bigcap_{\substack{p \leq i_n \\ p \text{ prime}}} G_n^p$ and $S_n = \bigcap_{\substack{p > i_n \\ p \text{ prime}}} G_n^p$, so that $E_n = R_n \cap S_n$, and $\overline{S_n} = \bigcup_{\substack{p > i_n \\ p \text{ prime}}} B_n^p$. Immediately, $\mathbb{P}(E_n) \leq \mathbb{P}(R_n)$, and as $\mathbb{P}(R_n \cap S_n) \geq \mathbb{P}(R_n) - \mathbb{P}(\overline{S_n})$, we have, using the union bound:

$$\mathbb{P}\left(\bigcap_{\substack{p \leq i_n \\ p \text{ prime}}} G_n^p\right) - \sum_{\substack{p > i_n \\ p \text{ prime}}} \mathbb{P}(B_n^p) \leq \mathbb{P}(E_n) \leq \mathbb{P}\left(\bigcap_{\substack{p \leq i_n \\ p \text{ prime}}} G_n^p\right). \quad (8.1)$$

The end of the proof consists in two steps. First, we show that $\sum_{\substack{p > i_n \\ p \text{ prime}}} \mathbb{P}(B_n^p)$ tends to 0.

Then we show that $\mathbb{P}\left(\bigcap_{\substack{p \leq i_n \\ p \text{ prime}}} G_n^p\right)$ tends to A_{d+1} . In an integer interval of length m there are at most $m/p + 1$ multiples of $p \geq 2$. Hence, by the union bound, if $m_n = \lfloor 2\sqrt{n} \rfloor - \lceil \sqrt{n} \rceil + 1$ denote the length of $[\sqrt{n}, 2\sqrt{n}]$:

$$\sum_{\substack{p > i_n \\ p \text{ prime}}} \mathbb{P}(B_n^p) \leq \binom{d+1}{2} \sum_{\substack{p=i_n+1 \\ p \text{ prime}}}^{\lfloor 2\sqrt{n} \rfloor} \left(\frac{1}{p} + \frac{1}{m_n}\right)^2 \leq \binom{d+1}{2} \sum_{p \geq i_n+1} \frac{9}{p^2},$$

as $2m_n \geq p$. The sum $\sum_{p \geq i_n+1} p^{-2}$ is the remainder of order $i_n \rightarrow \infty$ of a convergent series, hence

$$\lim_{n \rightarrow \infty} \sum_{p > i_n} \mathbb{P}(B_n^p) = 0.$$

For $n \geq 1$, let $I_n := \llbracket a_n, b_n \rrbracket$ be the largest interval of the form $\llbracket ix_n + 1, jx_n \rrbracket$ that is included in $[\sqrt{n}, 2\sqrt{n}]$:

$$a_n = \min\{ix_n + 1 : i \in \mathbb{Z}_{\geq 0} \text{ and } ix_n + 1 \geq \sqrt{n}\},$$

$$b_n = \max\{jx_n : j \in \mathbb{Z}_{\geq 0} \text{ and } jx_n \leq 2\sqrt{n}\}.$$

Observe that $I_n \neq \emptyset$ for n sufficiently large, as after some point $\sqrt[3]{n}$ is less than $\sqrt{n}/3$.

For any non-empty integer interval I , we denote by X_I the random variable uniformly distributed on I . Observe that for any set of prime numbers $\{p_1, \dots, p_t\} \in P_{i_n}$, the number of multiples in I_n of the product $p_1 \dots p_t$ is exactly $\frac{b_n - a_n + 1}{p_1 \dots p_t}$. Then we have, by direct counting

$$\mathbb{P}(\forall i \in [t], p_i \text{ divides } X_{I_n}) = \mathbb{P}(p_1 \dots p_t \text{ divides } X_{I_n}) = \frac{1}{p_1 \dots p_t} = \prod_{i=1}^t \mathbb{P}(p_i \text{ divides } X_{I_n}).$$

Hence we built I_n so that there is no rounding effects, to avoid slight alterations of the probabilities and get the independency. Moreover, for given $p \leq i_n$, the probability that none of the coordinates of a uniform random element of I_n is divisible by p is $(1 - 1/p)^{d+1}$, and

that exactly one of them is divisible by p is $(d+1)(1-1/p)^d/p = (1-1/p)^{d+1}(d+1)/(p-1)$. Hence, if we condition the integer vector to be in I_n^{d+1} we have:

$$\mathbb{P}\left(\bigcap_{\substack{p \leq i_n \\ p \text{ prime}}} G_n^p \mid I_n^{d+1}\right) = \prod_{\substack{p \leq i_n \\ p \text{ prime}}} \left(1 - \frac{1}{p}\right)^{d+1} \left(1 + \frac{d+1}{p-1}\right).$$

As $\left(1 - \frac{1}{p}\right)^{d+1} \left(1 + \frac{d+1}{p-1}\right) = 1 + O(p^{-2})$, the product has a limit as n , hence i_n , tends to infinity, which we call A_{d+1} as in [Tót02]. Finally, observe that since we chose $x_n \leq \sqrt[3]{n}$, for a uniform vector of $[\sqrt{n}, 2\sqrt{n}]^{d+1}$ we have

$$\mathbb{P}\left(I_n^{d+1}\right) \geq \left(\frac{m_n - 2x_n}{m_n}\right)^{d+1} \geq \left(1 - \frac{2}{n^{1/6}}\right)^{d+1}.$$

Therefore, $\mathbb{P}\left(I_n^{d+1}\right) = 1 - O(n^{-1/6})$. By the law of total probabilities we have

$$\mathbb{P}\left(\bigcap_{\substack{p \leq i_n \\ p \text{ prime}}} G_n^p\right) = \mathbb{P}\left(\bigcap_{\substack{p \leq i_n \\ p \text{ prime}}} G_n^p \mid I_n^{d+1}\right)\mathbb{P}\left(I_n^{d+1}\right) + \mathbb{P}\left(\bigcap_{\substack{p \leq i_n \\ p \text{ prime}}} G_n^p \cap \overline{I_n^{d+1}}\right).$$

This concludes the proof by Equation (8.1) since

$$\mathbb{P}\left(\bigcap_{\substack{p \leq i_n \\ p \text{ prime}}} G_n^p \mid I_n^{d+1}\right)\mathbb{P}\left(I_n^{d+1}\right) \xrightarrow{n \rightarrow \infty} A_{d+1},$$

and

$$\mathbb{P}\left(\bigcap_{\substack{p \leq i_n \\ p \text{ prime}}} G_n^p \cap \overline{I_n^{d+1}}\right) \leq \mathbb{P}\left(\overline{I_n^{d+1}}\right) = O(n^{-1/6}) \xrightarrow{n \rightarrow \infty} 0.$$

This concludes the proof. □

If $\mathcal{A} = (n, \delta, p \xrightarrow{a} q)$ is an almost deterministic transition structure, let $\mathcal{D}(\mathcal{A})$ denote its powerset transition structure, and let $\mathcal{D}(\mathcal{A}, i_0)$ denote its accessible powerset transition structure from the unique initial state $i_0 \in [n]$. Observe that if $\ell(\mathcal{A}, p, q) \neq \perp$, then $\mathcal{A} \models \mathbf{C}_p(n)$, and the word $aw_{\mathcal{A}}$ labels a cycle around p . Hence, using the word $au = (aw_{\mathcal{A}})^d$, as in Section 6, we have $\delta(\{p\}, au) = \{p_0, p_1, \dots, p_d\}$, where $(p_0, \dots, p_d) = s(\mathcal{A}, q)$ is the starting states tuple of \mathcal{A} . Hence, if \mathcal{A} satisfies a template in $\mathbf{T}_{p,q}(n)$, for every $j \geq 0$, we have $\delta(\{p_0, p_1, \dots, p_d\}, b^j) = \{\delta(p_0, b^j), \dots, \delta(p_d, b^j)\}$. For j large enough, all the $\delta(p_i, b^j)$'s are in the cyclic part of their respective b -thread, forming a b -cycle of length $\text{lcm}(\ell_0(\mathcal{A}), \dots, \ell_d(\mathcal{A}))$ in $\mathcal{D}(\mathcal{A})$. Let $\mathcal{C}_b(\mathcal{A})$ denote this b -cycle, which is defined when $\ell(\mathcal{A}, p, q) \neq \perp$.

Theorem 8.4. *Let ν be a positive integer. For the uniform distribution of almost deterministic transition structures with n -states with an added initial state taken uniformly at random and independently, the accessible powerset transition structure has more than n^ν states with visible probability.*

As a consequence, an algorithm that builds the accessible powerset automaton of a random almost deterministic automaton has super-polynomial average time complexity.

Proof. Let $\mathcal{A} = (n, \delta, p \xrightarrow{a} q)$ be an almost deterministic transition structure such that $\ell(\mathcal{A}, p, q) \neq \perp$, and p is accessible from $i_0 \in [n]$, which acts as its initial state. Then there exists a word $v \in \Sigma^*$ such that $\delta(\{i_0\}, v) = \{p\}$, as v can be chosen so that the path

labeled by v starting from i_0 is simple and does not encounter p twice, and as p is the only state with a non-deterministic transition. As a consequence, for $u = w_{\mathcal{A}}(aw_{\mathcal{A}})^{d-1}$, we have $\delta(\{i_0\}, vau) = \{p_0, p_1, \dots, p_d\}$, and therefore, $\mathcal{C}_b(\mathcal{A})$ is in $\mathcal{D}(\mathcal{A}, i_0)$. Recall that $|\mathcal{C}_b(\mathcal{A})| = \text{lcm}(\ell_0(\mathcal{A}), \dots, \ell_d(\mathcal{A}))$, we have

$$\begin{aligned} \mathbb{P}\left(|\mathcal{D}(\mathcal{A}, i_0)| \geq n^{(d+1)/2}\right) &\geq \mathbb{P}\left(\ell(\mathcal{A}, p, q) \neq \perp \text{ and } |\mathcal{D}(\mathcal{A})| \geq n^{(d+1)/2}\right) \\ &\geq \mathbb{P}\left(\ell(\mathcal{A}, p, q) \neq \perp \text{ and } \text{lcm}(\ell_0(\mathcal{A}), \dots, \ell_d(\mathcal{A})) \geq n^{(d+1)/2}\right) \end{aligned}$$

By Proposition 7.4, given $\ell(\mathcal{A}, p, q) \neq \perp$, $(\ell_0(\mathcal{A}), \dots, \ell_d(\mathcal{A}))$ is a uniform random element of $[\sqrt{n}, 2\sqrt{n}]^{d+1}$. Hence Corollary 8.3 applies and yields that for some $c > 0$ and n sufficiently large,

$$\mathbb{P}\left(\text{lcm}(\ell_0(\mathcal{A}), \dots, \ell_d(\mathcal{A})) \geq n^{(d+1)/2} \mid \ell(\mathcal{A}, p, q) \neq \perp\right) \geq c.$$

Therefore, we have

$$\mathbb{P}\left(|\mathcal{D}(\mathcal{A})| \geq n^{(d+1)/2}\right) \geq c \mathbb{P}(\ell(\mathcal{A}, p, q) \neq \perp),$$

Hence by Proposition 7.4, $|\mathcal{D}(\mathcal{A})| \geq n^{(d+1)/2}$ with visible probability. Since p is accessible with high probability, we have, as in Lemma 8.1, that both properties hold with visible probability. This concludes the proof. \square

9. STATE COMPLEXITY

We are now ready to randomly select which states are final, and to state a result on languages instead of on transition structures, by studying the state complexity of a language recognized by a random almost deterministic automaton.

In our model, for every n , each state is final with fixed probability f_n , which may depend on n as long as it is not too close to either 0 or 1: we require that f_n and $1 - f_n$ are in $\Omega(\frac{1}{\sqrt{n}})$. Using a variant of the Birthday Problem again, this ensures that a uniform random set of states of size $\Theta(\sqrt{n})$ contains both final and non-final states with visible probability.

In the proof of Theorem 8.4, we exhibited the existence with visible probability of $d + 1$ occurrences of b -cycles in a random almost deterministic transition structure, yielding a large b -cycle when applying the powerset construction. We will focus on b -cycles in the sequel, as it turns out to be sufficient to prove our main result. We rely on the notion of primitive words, which we now recall.

Let Γ be a nonempty finite alphabet. If $w \in \Gamma^\ell$ is a word of length ℓ , we write $w = w_0 \cdots w_{\ell-1}$ and use the convention that all indices are taken modulo ℓ : for instance w_ℓ is the letter w_0 . A nonempty word w is *primitive* if it is not a non-trivial power of another word: it cannot be written $w = z^k$ for some word z and some $k \geq 2$. If w is primitive, it is easily seen that every circular permutation of w is also primitive. See [Lot97] for a more detailed account on primitive words.

Primitive words appear in our proof with the following observation. If $\mathcal{C} = (c_0, \dots, c_{\ell-1})$ is a b -cycle of states starting at c_0 , its *associated binary word* is the size- ℓ word $v = v_0 \dots v_{\ell-1}$ of $\{0, 1\}^\ell$ where $v_i = 1$ if and only if c_i is a final state. Recall that if we start the same cycle elsewhere, at c_i , the associated word $v' = v_i \cdots v_\ell v_0 \cdots v_{i-1}$ is primitive if and only if v is primitive: reading the associated binary word from any starting state preserves primitivity.

A b -cycle is said to be *primitive* if one (equivalently, all) of its associated words is (are) primitive. Our study is based on the following statement.

Lemma 9.1. *Let \mathcal{A} be a deterministic automaton on Σ and $\alpha \in \Sigma$. If \mathcal{C} is a primitive α -cycle of \mathcal{A} , then the states of \mathcal{C} are pairwise non-equivalent: the state complexity of the language recognized by \mathcal{A} is at least $|\mathcal{C}|$, the number of states in \mathcal{C} .*

Proof. Let $\mathcal{A} = (n, \delta)$ and let p and q be two different states of \mathcal{C} . Let x and y be the associated binary words of \mathcal{C} starting at p and q , respectively. Let k be the smallest positive integer such that $\delta(p, \alpha^k) = q$ and let u be the prefix of length k of x , and v be the associated suffix: $x = uv$. Then $y = vu$. Assume by contradiction that p and q are equivalent. This implies that $x = y$, as the automata obtained by placing the initial states either on p or q recognize the same elements of $\{\alpha\}^*$. Hence $uv = vu$, and therefore u and v are the power of the same word by a classical result on primitive words [Lot97, Prop. 1.3.2 page 8]. This is in contradiction with the fact that \mathcal{C} is primitive. \square

9.1. A composition of binary primitive words. If $w^{(1)}$ and $w^{(2)}$ are two non-empty words of respective lengths ℓ_1 and ℓ_2 on the binary alphabet $\{0, 1\}$, we denote by $w^{(1)} \odot w^{(2)}$ the word w of length $\ell = \text{lcm}(\ell_1, \ell_2)$ given by $w_i = 1$ if and only if $w_i^{(1)} = 1$ or $w_i^{(2)} = 1$ (recall that the indices are taken modulo the length of the word). We will see in the sequel that this operation naturally happens when extending the notion of state equivalence from each b -cycle to the corresponding b -cycle in the powerset construction.

Lemma 9.2. *Let $w^{(1)}$ and $w^{(2)}$ be two primitive words on $\{0, 1\}$ of lengths at least 2 that are coprime. Then the word $w^{(1)} \odot w^{(2)}$ is primitive.*

Proof. Let $w = w^{(1)} \odot w^{(2)}$. Assume by contradiction that there exists some word z and some $k \geq 2$ such that $w = z^k$. Letting p be a prime number that divides k , we have $w = (z^{k/p})^p$. This yields that p divides $\ell = \text{lcm}(\ell_1, \ell_2) = \ell_1 \times \ell_2$ and that for every non-negative integer i , $w_i = w_{i+\ell/p}$ (indices taken modulo ℓ). Obviously, p divides either ℓ_1 or ℓ_2 , but not both as they are coprime. By symmetry, assume that it divides ℓ_1 : $\ell_1 = pr$ and $\ell/p = r\ell_2$.

Since $w^{(2)}$ has length at least 2 and is primitive, there exists an index $i_0 \in \{0, \dots, \ell_2 - 1\}$ such that $w_{i_0}^{(2)} = 0$. Define $i_j = i_0 + j\ell_2$, for any $j \geq 0$. As indices in $w^{(2)}$ are taken modulo ℓ_2 , we have $w_{i_j}^{(2)} = 0$ for all $j \geq 0$. Therefore, $w_{i_j}^{(1)} = 1$ if and only if $w_{i_j} = 1$, so that $w_{i_j} = w_{i_j}^{(1)}$ for all $j \geq 0$. As $w_{i_j} = w_{i_j+\ell/p}$ and $\ell/p = r\ell_2$, this yields that $w_{i_j}^{(1)} = w_{i_j+r\ell_2}^{(1)}$ for all $j \geq 0$. Moreover, $r\ell_2$ is not a multiple of ℓ_1 : let $\alpha \geq 1$ be the largest integer such that p^α divides $\ell_1 = pr$, then p^α does not divide $r\ell_2$, since p does not divide ℓ_2 .

Let $s := r\ell_2 \pmod{\ell_1}$, we just established that $s \neq 0$, so we have the non-trivial relation $w_{i_j}^{(1)} = w_{i_j+s}^{(1)}$ for all $j \geq 0$. Recall that $i_j = i_0 + j\ell_2$. As ℓ_1 and ℓ_2 are coprime, the i_j 's take all the values modulo ℓ_1 when j ranges from 0 to $\ell_1 - 1$. Hence, for all $k \in \{0, \dots, \ell_1 - 1\}$, $w_k^{(1)} = w_{k+s}^{(1)}$, for some $0 < s < \ell_1$. This is a contradiction with the fact that $w^{(1)}$ is primitive, concluding the proof. \square

Remark 9.3. Lemma 9.2 does not hold if the lengths are not coprime. For instance, if $w^{(1)} = 011111$ and $w^{(2)} = 1011$, then $w^{(1)} \odot w^{(2)} = \underbrace{1 \dots 1}_{12 \text{ times}}$, which is not primitive.

9.2. Primitivity of random words. From a probabilistic point of view, it is well known [Lot97] that a uniform random word is primitive with exponentially high probability. We rely on the following finer result.

Lemma 9.4 [FN16]. *Let μ be a probability measure on $\{0, 1\}^n$ such that $\mu(0^n) = \mu(1^n) = 0$ and such that two words with the same number of 0's have same probability. Then the probability that a word is not primitive under μ is at most $\frac{2}{n}$.*

We adapt it to our needs as follows:

Corollary 9.5. *Let f_n be a sequence of real numbers in $(0, 1)$ such that $f_n = \Omega(\frac{1}{\sqrt{n}})$ and $1 - f_n = \Omega(\frac{1}{\sqrt{n}})$. Consider the distribution on random words of length n where each letter is chosen independently to be 1 with probability f_n and 0 with probability $1 - f_n$. There exists $c > 0$ such that, for n sufficiently large and for any integer $\ell \geq \sqrt{n}$, a random binary word of length ℓ is primitive with probability at least c .*

Proof. Let X be the event that $w = 0^\ell$ or $w = 1^\ell$. We have $\mathbb{P}(X) = f_n^\ell + (1 - f_n)^\ell$. Since changing the 0's in 1's and the 1's in 0's preserves primitivity, we can assume by symmetry that $f_n \leq \frac{1}{2}$. By hypothesis, there exists some constant $\beta > 0$ such that $\frac{\beta}{\sqrt{n}} \leq f_n$ and $\frac{\beta}{\sqrt{n}} \leq 1 - f_n$ hence, as $f_n \leq \frac{1}{2}$, we have

$$f_n^\ell \leq \frac{1}{2^{\sqrt{n}}} \text{ and } (1 - f_n)^\ell \leq \left(1 - \frac{\beta}{\sqrt{n}}\right)^{\sqrt{n}}.$$

Since $(1 - \frac{\beta}{\sqrt{n}})^{\sqrt{n}} = e^{-\beta} + O(\frac{1}{\sqrt{n}})$, there exists some constant $\delta < 1$ such that $\mathbb{P}(X) \leq \delta$, for n sufficiently large.

Let W be a random word under our distribution. For any $w \in \{0, 1\}^\ell$, the conditional probability that W values w given that $W \notin \{0^\ell, 1^\ell\}$ is

$$\mathbb{P}(W = w \mid \bar{X}) = \begin{cases} 0 & \text{if } w = 0^\ell \text{ or } w = 1^\ell, \\ \frac{\mathbb{P}(W=w)}{1-\mathbb{P}(X)} & \text{otherwise.} \end{cases}$$

Hence we are in the settings of Lemma 9.4, and the probability that w is not primitive, given that $w \notin \{0^\ell, 1^\ell\}$ is at most $\frac{2}{\ell}$. Moreover, we have:

$$\mathbb{P}(w \text{ not primitive}) = \mathbb{P}(X) + \mathbb{P}(w \text{ not primitive} \mid \bar{X})\mathbb{P}(\bar{X}) \leq \delta + \frac{2}{\ell},$$

where $\delta + \frac{2}{\ell} < 1$, for n large enough. This concludes the proof. \square

9.3. Finalizing the proof of Theorem 3.1. If $\mathcal{A} = (n, \delta)$ is a transition structure such that $\ell(\mathcal{A}, p, q) \neq \perp$, recall that $\mathcal{C}_b(\mathcal{A})$ is the b -cycle of its powerset construction $\mathcal{D}(\mathcal{A})$ starting from its set of starting states $P = \{p_0, \dots, p_d\}$. Let j_0 be a positive integer such that $\delta(p_i, b^{j_0})$ is in the b -cycle of its b -thread for all $i \in \{0, \dots, d\}$, and let $P_0 = (\delta(p_0, b^{j_0}), \dots, \delta(p_d, b^{j_0}))$. We use a tuple for P_0 as it is easier for the upcoming arguments, but since the b -threads of \mathcal{A} are pairwise disjoint, and since we only read b 's from P_0 in the sequel, they can be identified with the associated sets of $\mathcal{D}(\mathcal{A})$. Let us write $c_j^{(i)} = \delta(p_i, b^{j_0+j})$ for all $i \in \{0, \dots, d\}$ and all $j \geq 0$. The b -cycle \mathcal{C}_i of the i -th b -thread has length ℓ_i , so if $j = j' \pmod{\ell_i}$, then $c_j^{(i)} = c_{j'}^{(i)}$. Moreover, the b -cycle $\mathcal{C}_b(\mathcal{A})$ is made of the $(c_j^{(0)}, \dots, c_j^{(d)})$ for $j \geq 0$.

We now consider a set of final state $F \subseteq [n]$. For this F , and for every $i \in \{0, \dots, d\}$, let $w^{(i)}$ denote the binary word associated to \mathcal{C}_i , starting at $c_0^{(i)}$. Let w denote the binary word associated with $\mathcal{C}_b(\mathcal{A})$, starting at P_0 . Since a state X in the powerset construction is final if and only if one of its state is in F , by a direct induction we have

$$w = w^{(0)} \odot w^{(1)} \odot \dots \odot w^{(d)}.$$

This is depicted in Figure 5 for two b -threads.

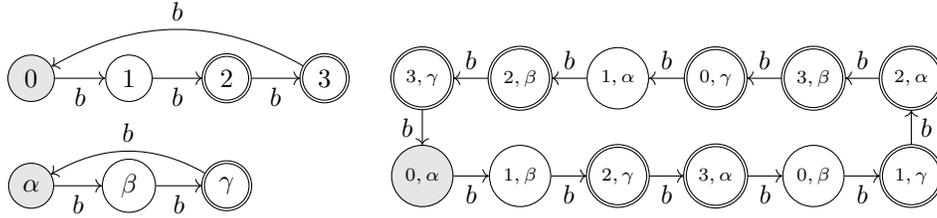


Figure 5: On the left, two primitive b -cycles (accepting states are denoted by double circles) whose associated words are 0011 (top) and 001 (bottom), starting at 0 and α , respectively. On the right, the b -cycle of $\{0, \alpha\}$ of associated word $0011 \odot 001 = 001101111011$, which is primitive by Lemma 9.2.

To finish the proof of Theorem 3.1, we follow the proof of Theorem 8.4 and add the independent choice of which states are final. By Corollary 9.5, all the b -cycles used to build $\mathcal{C}_b(\mathcal{A})$ are also primitive, still with visible probability. Hence by Lemma 9.2, the product cycle $\mathcal{C}_b(\mathcal{A})$ is also primitive with visible probability. This concludes the proof by Lemma 9.1.

10. REMARK ON DENSE RANDOM DFAS

In this short section, we illustrate the degenerate nature of *dense* random NFAs by considering the model with a unique initial state and in which each transition is added independently with a fixed probability $p \in (0, 1)$. A uniform dense random NFA (with one initial state) corresponds to $p = \frac{1}{2}$. The following proof is very similar to the analysis of the diameter of random dense undirected graphs [Bol01, Corollary 10.11].

Consider a random n -state NFA \mathcal{A} under this model, with initial state i_0 and transition function δ . Let $w = w_0w_1$ be any word of length 2 on Σ and let q be any state of \mathcal{A} . We bound from above the probability that $q \notin \delta(i_0, w)$ as follows :

$$\begin{aligned} \mathbb{P}(q \notin \delta(i_0, w)) &= \mathbb{P}(\forall r \in [n], r \notin \delta(i_0, w_0) \text{ or } q \notin \delta(r, w_1)) \\ &\leq \mathbb{P}(\forall r \in [n] \setminus \{i_0\}, r \notin \delta(i_0, w_0) \text{ or } q \notin \delta(r, w_1)) \\ &= \prod_{r \in [n] \setminus \{i_0\}} \mathbb{P}(r \notin \delta(i_0, w_0) \text{ or } q \notin \delta(r, w_1)) \\ &= \prod_{r \in [n] \setminus \{i_0\}} (1 - \mathbb{P}(r \in \delta(i_0, w_0) \text{ and } q \in \delta(r, w_1))) \\ &= \prod_{r \in [n] \setminus \{i_0\}} (1 - p^2) = (1 - p^2)^{n-1}. \end{aligned}$$

Thus, by the union bound we obtain that with high probability, the image of $\{i_0\}$ by any word of length 2 is the set of all states of \mathcal{A} . We therefore obtain the following statement, which establishes the degeneracy of the dense model.

Proposition 10.1. *The accessible subset automaton of a random n -state dense NFA with one initial state has at most $|\Sigma| + 2$ states with high probability. Hence, with high probability, the state complexity of the language is bounded by a constant independent of n and p .*

11. CONCLUSION AND DISCUSSION

Our main theorem states that the state complexity of a random almost deterministic automaton is greater than n^d with probability at least $c_d > 0$ for n sufficiently large. One can wonder how small the constant c_d is and for which sizes the lower-bound holds. As we said in the introduction, we did not try to estimate c_d nor did we try to optimize its value in this article. Since the powerset construction quickly generates very large automata which would need to be minimized, a proper experimental study does not seem feasible. However, we did generate 1000 almost deterministic transition structures with $n = 100$ states and apply the accessible powerset construction: in 78.6% of the 1000 cases the output had more than n^3 states. This would lead us to guess that even if the constant c_3 that can be derived from our proof is very small, combinatorial explosion does occur frequently in practice.

Also, as noticed above, in our settings it is certain that the property does not hold with high probability, as there is an asymptotically constant probability that the source of the added transition is not accessible. However, this probability is roughly 20.4%, not too far from what we obtained in our experiment on size-100 structures: it is very possible that if we condition the source of the added transition to be accessible, then our result holds with high probability. However, our proof techniques, based on an intensive use of the Birthday Problem cannot prove this: completely new ideas are necessary to establish such a result.

It is natural to ask whether our result could be strengthened from super-polynomial size to exponential size. We do not know if this generalization holds, but it seems out of reach of the techniques developed in this paper.

Another natural direction is to consider the case when there are *few* final states, as $\Theta(\sqrt{n})$ final states may be considered too large for a random deterministic automaton. The extreme case is to allow exactly one final state by choosing it uniformly at random. If we do so, our analysis using primitive words fails: with high probability the b -cycles we built have no final state at all, and neither has the associated b -cycle \mathcal{C} in the powerset construction. However, we are confident that our techniques can be used to capture this distribution: by studying the paths ending in this final state, we should be able to find for each b -cycle \mathcal{C}_i a word w_i that maps exactly one state to the final state, and such that the w_i 's are all different. This would be enough to establish that the states of \mathcal{C} are pairwise non-equivalent and prove the conjecture. Completely formalizing and proving this idea is an ongoing work.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their valuable feedback. This work was partially supported by the French ANR grant ASPAG (ANR-17-CE40-0017).

REFERENCES

- [ABBPL20] Louigi Addario-Berry, Borja Balle, and Guillem Perarnau Llobet. Diameter and stationary distribution of random r-out digraphs. *Electronic journal of combinatorics*, 27(P3. 28):1–41, 2020. doi:10.37236/9485.
- [BDN12] Frédérique Bassino, Julien David, and Cyril Nicaud. Average case analysis of Moore’s state minimization algorithm. *Algorithmica*, 63(1-2):509–531, 2012. doi:10.1007/s00453-011-9557-7.
- [BDS12] Frédérique Bassino, Julien David, and Andrea Sportiello. Asymptotic enumeration of minimal automata. In Christoph Dürr and Thomas Wilke, editors, *29th International Symposium on Theoretical Aspects of Computer Science, STACS 2012, February 29th - March 3rd, 2012, Paris, France*, volume 14 of *LIPICs*, pages 88–99. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2012. doi:10.4230/LIPICs.STACS.2012.88.
- [Ber16] 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.
- [BM08] Henrik Björklund and Wim Martens. The tractability frontier for NFA minimization. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir, and Igor Walukiewicz, editors, *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7-11, 2008, Proceedings, Part II - Track B: Logic, Semantics, and Theory of Programming & Track C: Security and Cryptography Foundations*, volume 5126 of *Lecture Notes in Computer Science*, pages 27–38. Springer, 2008. doi:10.1007/978-3-540-70583-3_3.
- [Bol01] Béla Bollobás. *Random graphs*. Number 73 in Cambridge studies in advanced mathematics. Cambridge University Press, 2 edition, 2001. doi:10.1017/CB09780511814068.
- [CB03] Jin-yi Cai and Eric Bach. On testing for zero polynomials by a set of points with bounded precision. *Theor. Comput. Sci.*, 296(1):15–25, 2003. doi:10.1016/S0304-3975(02)00429-2.
- [CD17] Xing Shi Cai and Luc Devroye. The graph structure of a deterministic automaton chosen at random. *Random Structures & Algorithms*, 51(3):428–458, 2017. doi:10.1002/rsa.20707.
- [CDKN23] Arnaud Carayol, Philippe Duchon, Florent Koechlin, and Cyril Nicaud. One drop of non-determinism in a random deterministic automaton. In Petra Berenbrink, Patricia Bouyer, Anuj Dawar, and Mamadou Moustapha Kanté, editors, *40th International Symposium on Theoretical Aspects of Computer Science, STACS 2023, March 7-9, 2023, Hamburg, Germany*, volume 254 of *LIPICs*, pages 19:1–19:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICs.STACS.2023.19.
- [CN12] Arnaud Carayol and Cyril Nicaud. Distribution of the number of accessible states in a random deterministic automaton. In Christoph Dürr and Thomas Wilke, editors, *29th International Symposium on Theoretical Aspects of Computer Science, STACS 2012, February 29th - March 3rd, 2012, Paris, France*, volume 14 of *LIPICs*, pages 194–205. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2012. doi:10.4230/LIPICs.STACS.2012.194.
- [CP23] Guillaume Chapuy and Guillem Perarnau. Short synchronizing words for random automata. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 581–604. SIAM, 2023. doi:10.1137/1.9781611977554.CH26.
- [Dav12] Julien David. Average complexity of Moore’s and Hopcroft’s algorithms. *Theor. Comput. Sci.*, 417:50–65, 2012. doi:10.1016/j.tcs.2011.10.011.
- [ER60] Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.*, 5(1):17–60, 1960.
- [FN13] Sven De Felice and Cyril Nicaud. Brzozowski algorithm is generically super-polynomial for deterministic automata. In Marie-Pierre Béal and Olivier Carton, editors, *Developments in Language Theory - 17th International Conference, DLT 2013, Marne-la-Vallée, France, June 18-21, 2013. Proceedings*, volume 7907 of *Lecture Notes in Computer Science*, pages 179–190. Springer, 2013. doi:10.1007/978-3-642-38771-5_17.
- [FN16] Sven De Felice and Cyril Nicaud. Average case analysis of Brzozowski’s algorithm. *Int. J. Found. Comput. Sci.*, 27(2):109–126, 2016. doi:10.1142/S0129054116400025.

- [Gru73] Aleksandr Aleksandrovich Grusho. Limit distributions of certain characteristics of random automaton graphs. *Mathematical Notes of the Academy of Sciences of the USSR*, 14(1):633–637, 1973.
- [Har63] Theodore E. Harris. *The theory of branching processes*. Die Grundlehren der Mathematischen Wissenschaften, Bd. 119. Springer-Verlag, Berlin, 1963.
- [HU79] J. Hopcroft and J. Ullman. *Introduction to Automata Theory, Languages and Computation*. Addison-Wesley, 1979.
- [JLR00] Svante Janson, Tomasz Luczak, and Andrzej Rucinski. *Random Graphs*. John Wiley & Sons, 2000. doi:10.1002/9781118032718.
- [KNR21] Florent Koechlin, Cyril Nicaud, and Pablo Rotondo. Simplifications of uniform expressions specified by systems. *Int. J. Found. Comput. Sci.*, 32(6):733–760, 2021. doi:10.1142/S0129054121420065.
- [Lot97] Lothaire. *Combinatorics on Words*. Cambridge Mathematical Library. Cambridge University Press, 2 edition, 1997. doi:10.1017/CB09780511566097.
- [Mar23] Anders Martinsson. Synchronizing random automata through repeated 'a' inputs, 2023. arXiv:2306.09040.
- [MF71] Albert R Meyer and Michael J Fischer. Economy of description by automata, grammars, and formal systems. In *12th Annual Symposium on Switching and Automata Theory (SWAT 1971)*, pages 188–191. IEEE Computer Society, 1971.
- [Nic14] Cyril Nicaud. Random deterministic automata. In Erzsébet Csuhaaj-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.
- [Nic19] Cyril Nicaud. The Černý Conjecture Holds with High Probability. *J. Autom. Lang. Comb.*, 24(2-4):343–365, 2019. doi:10.25596/jalc-2019-343.
- [RS59] Michael O. Rabin and Dana S. Scott. Finite automata and their decision problems. *IBM J. Res. Dev.*, 3(2):114–125, 1959. doi:10.1147/RD.32.0114.
- [Tót02] László Tóth. The probability that k positive integers are pairwise relatively prime. *Fibonacci Quart.*, 40:13–18, 2002. doi:10.1080/00150517.2002.12428676.