



Preservation of normality by transducers

Olivier Carton, Elisa Orduna

► To cite this version:

Olivier Carton, Elisa Orduna. Preservation of normality by transducers. Information and Computation, 2022, 282, pp.104650. 10.1016/j.ic.2020.104650 . hal-03660658

HAL Id: hal-03660658

<https://u-paris.hal.science/hal-03660658>

Submitted on 15 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Preservation of normality by transducers

Olivier Carton Elisa Orduna

October 31, 2018

Abstract

We consider input-deterministic finite state transducers with infinite inputs and infinite outputs, and we consider the property of Borel normality on infinite words. When these transducers are given by a strongly connected set of states, and when the input is a Borel normal sequence, the output is an infinite word such that every word has a frequency given by a weighted automaton over the rationals. We prove that there is an algorithm that decides in cubic time whether an input-deterministic transducer preserves normality.

Keywords: transducers, weighted automata, normal sequences

1 Introduction

We start with the definition of *normality* for real numbers, given by Émile Borel [6] more than one hundred years ago. A real number is normal to an integer base if, in its infinite expansion expressed in that base, all blocks of digits of the same length have the same limiting frequency. Borel proved that almost all real numbers are normal to all integer bases. However, very little is known on how to prove that a given number has the property.

The definition of normality was the first step towards a definition of randomness. Normality formalizes the least requirements about a random sequence. It is indeed expected that in a random sequence, all blocks of symbols with the same length occur with the same limiting frequency. Normality, however, is a much weaker notion than the one of purely random sequences defined by Martin-Löf.

The motivation of this work is the study of transformations preserving randomness, hence preserving normality. The paper is focused on very simple transformations, namely those that can be realized by finite-state machines. We consider input-deterministic automata with outputs, also known as sequential transducers mapping infinite sequences of symbols to infinite sequences of symbols. The main result is that it can be decided in cubic time whether such a machine preserves or not normality. Preserving normality means that the output sequence is normal whenever the input sequence is.

The main result is obtained through a second result involving weighted automata. This second result states that if a sequential transducer is strongly connected, then the frequency of each block in the output of a run with a normal input is given by a weighted automaton on rational numbers. It implies, in particular, that the frequency of each block in the output sequence does not depend on the input as long as this input sequence is normal.

This is not the first result linking normality and automata. A fundamental theorem relates normality and finite automata: an infinite word is normal to a given alphabet if and only if it cannot be compressed by lossless finite transducers. These are deterministic finite automata with injective input-output behaviour. This result was first obtained by joining a theorem by Schnorr and Stimm [14] with a theorem by Dai, Lathrop, Lutz and Mayordomo [12]. Becher and Heiber gave a direct proof [3].

Agafonov's Theorem [1] is another striking result relating normality and automata. It establishes that oblivious selection of symbols with a regular set preserves normality. Oblivious selection with a regular set L means that a symbol a_i is selected whenever the prefix $a_1 \cdots a_{i-1}$ belongs to L . This oblivious selections can actually be realized by sequential transducers considered in this work.

The paper is organized as follows. Notions of normal sequences and transducers are introduced in Section 2. Main results are stated in Section 3. Proofs of the results and algorithms are given in Section 4.

2 Basic Definitions

2.1 Normality

Before giving the formal definition of normality, let us introduce some simple definitions and notation. Let A be a finite set of symbols that we refer to as the *alphabet*. We write A^ω for the set of all infinite words on the alphabet A and A^* for the set of all finite words. The length of a finite word w is denoted by $|w|$. The positions of finite and infinite words are numbered starting from 1. To denote the symbol at position i of a word w we write $w[i]$, and to denote the substring of w from position i to j inclusive we write $w[i \dots j]$. The empty word is denoted by λ . The cardinality of a finite set E is denoted by $\#E$.

Given two words w and v in A^* , the number $|w|_v$ of occurrences of v in w is defined by:

$$|w|_v = \#\{i : w[i \dots i + |v| - 1] = v\}.$$

For example, $|abbab|_{ab} = 2$.

Given a finite word $w \in A^+$ and an infinite word $x \in A^\omega$, we refer to the *frequency of w in x* as

$$\text{freq}(x, w) = \lim_{n \rightarrow \infty} \frac{|x[1 \dots n]|_w}{n}$$

when this limit is well-defined.

An infinite word $x \in A^\omega$ is *normal* on the alphabet A if for every word $w \in A^*$:

$$\text{freq}(x, w) = \frac{1}{(\#A)^{|w|}}$$

An alternative definition of normality can be given by counting *aligned* occurrences, and it is well-known that they are equivalent (see for example [2]). We refer the reader to [8, Chap.4] for a complete introduction to normality.

The most famous example of a normal word is due to Champernowne [10], who showed in 1933 that the infinite word obtained from concatenating all the

natural numbers (in their usual order):

0123456789101112131415161718192021222324252627282930...

is normal on the alphabet $\{0, 1, \dots, 9\}$.

2.2 Input-deterministic Transducers

In this paper we consider automata with outputs, also known as transducers. Such finite-state machines are used to realize functions mapping words to words and especially infinite words to infinite words. We only consider input-deterministic transducers, also known as *sequential* in the literature (See [13, Sec. V.1.2] and [4, Chap. IV]). Each transition of these transducers consumes exactly one symbol of their input and outputs a finite word which might be empty. Furthermore, ignoring the output label of each transition must yield a deterministic automaton.

More formally a *transducer* $\mathcal{T}$ is a tuple $\langle Q, A, B, \delta, q_0 \rangle$, where Q is a finite set of states, A and B are the input and output alphabets respectively, $\delta \subseteq Q \times A \times B^* \times Q$ is a finite transition relation and $q_0 \in Q$ is the initial state. A transition is a tuple $\langle p, a, v, q \rangle$ in $Q \times A \times B^* \times Q$ and it is written $p \xrightarrow{a|v} q$.




Figure 1: A deterministic complete transducer

A transducer $\mathcal{T}$ is *input-deterministic*, or *sequential* for short, if whenever $p \xrightarrow{a|v} q$ and $p \xrightarrow{a|v'} q'$ are two of its transitions, then $q = q'$ and $v = v'$. A transducer $\mathcal{T}$ is *complete* if for each symbol $a \in A$ and each state $p \in Q$ there is a transition from p and consuming a , that is, there exists a transition $p \xrightarrow{a|v} q$.

A finite (respectively infinite) *run* in $\mathcal{T}$ is a finite (respectively infinite) sequence of consecutive transitions,

$$q_0 \xrightarrow{a_1|v_1} q_1 \xrightarrow{a_2|v_2} \dots q_{n-1} \xrightarrow{a_n|v_n} q_n$$

Its *input* and *output labels* are the words $a_1 a_2 \dots a_n$ and $v_1 v_2 \dots v_n$ respectively. Note that there is no accepting condition and note also that the output label of an infinite run might be finite since the output label of some transitions might be empty. An infinite run is *accepting* if its first state is initial and its output label is infinite. If $\mathcal{T}$ is an input-deterministic transducer, each infinite word x is the input label of at most one accepting run in $\mathcal{T}$. When this run does exist, its output is denoted by $\mathcal{T}(x)$.

Each transducer $\mathcal{T}$ can be seen as a graph by ignoring the labels of its transitions. For this reason, we may consider *strongly connected components*

(SCC) of $\mathcal{T}$. Using the terminology of Markov chains, a strongly connected component is called *recurrent* if no transition leaves it.

We say that a input-deterministic transducer $\mathcal{T}$ *preserves normality* if for each normal word x , $\mathcal{T}(x)$ is also normal.

2.3 Weighted Automata

We now introduce weighted automata. In this paper we only consider weighted automata whose weights are rational numbers with the usual addition and multiplication (See [13, Chap. III] for a complete introduction).

A *weighted automaton* $\mathcal{A}$ is a tuple $\langle Q, B, \Delta, I, F \rangle$, where Q is the state set, B is the alphabet, $I : Q \rightarrow \mathbb{Q}$ and $F : Q \rightarrow \mathbb{Q}$ are the functions that assign to each state an initial and a final weight and $\Delta : Q \times B \times Q \rightarrow \mathbb{Q}$ is a function that assigns to each transition a weight.

As usual, the weight of a run is the product of the weights of its transitions times the initial weight of its first state and times the final weight of its last state. Furthermore, the weight of a word $w \in B^*$ is the sum of the weights of all runs with label w .




Figure 2: A weighted automaton

A weighted automaton is pictured in Figure 2. A transition $p \xrightarrow{a} q$ with weight x is pictured $p \xrightarrow{a:x} q$. Non-zero initial and final weights are given over small incoming and outgoing arrows. The weight of the run $q_0 \xrightarrow{1} q_1 \xrightarrow{0} q_1 \xrightarrow{1} q_1 \xrightarrow{0} q_1$ is $1 \cdot 1 \cdot 2 \cdot 2 \cdot 2 \cdot 1 = 8$. The weight of the word $w = 1010$ is $8 + 2 = 10$. More generally the weight of a word $w = a_1 \cdots a_k$ is the integer $n = \sum_{i=1}^k a_i 2^{k-i}$ (w is a binary expansion of n with possibly some leading zeros).

3 Results

We now state the main results of the paper. The first one states that when a transducer is strongly connected, deterministic and complete, the frequency of each finite word w in the output of a run with a normal input label is given by a weighted automaton over $\mathbb{Q}$. The second one states that it can be checked in cubic time whether an input-deterministic transducer preserves normality.

Theorem 1. *Given a transducer $\mathcal{T} = \langle Q, A, B, \delta, q_0 \rangle$ which is strongly connected, deterministic and complete, there exists a weighted automaton $\mathcal{A}$ such that for any normal word x and for any finite word w , $\text{freq}(\mathcal{T}(x), w) = \text{weight}_{\mathcal{A}}(w)$.*

Furthermore, the weighted automaton $\mathcal{A}$ can be computed in cubic time with respect to the size of the transducer $\mathcal{T}$.

The hypothesis that the transducer is complete guarantees that each normal word x is the input label of an infinite run. Nevertheless, this run may not be accepting since its output label is not necessarily infinite. The really restrictive hypothesis is that the transducer must be input-deterministic.

To illustrate Theorem 1 we give in Figure 3 a weighted automaton $\mathcal{A}$ which computes the frequency of each finite word w in $\mathcal{T}(x)$ for a normal input x and the transducer $\mathcal{T}$ pictured in Figure 1. States 2 and 3 are useless and could be removed since their initial weight is zero and they have no incoming transitions. They have been kept because the automaton pictured in Figure 3 is the one given by the procedure described in the next section.




Figure 3: A weighted automaton for the transducer pictured in Fig. 1

Theorem 2. *Given a transducer $\mathcal{T} = \langle Q, A, B, \delta, q_0 \rangle$ which is complete and deterministic, it can be decided in cubic time with respect to the size of $\mathcal{T}$ whether $\mathcal{T}$ preserves normality.*

4 Algorithms and Proofs

In this section we provide the proofs for Theorems 1 and 2. The next proposition shows that it suffices to independently analyze each recurrent strongly connected component of the transducer.

Proposition 3. *A deterministic and complete transducer preserves normality if and only if each of its recurrent strongly connected components preserves normality.*

The previous proposition follows directly from the next lemma which is Satz 2.5 in [14].

Lemma 4. *A run labeled with a normal word in a deterministic and complete automaton always reaches a recurrent strongly connected component.*

By Proposition 3 it suffices to analyze preservation of normality in each recurrent strongly connected component. In what follows we mainly consider strongly connected transducers. If all transitions have an empty output label, the output of any run is empty and the transducer does not preserve normality.

Therefore, we assume that transducers have at least one transition with a non empty output label. By Lemmas 5 and 6, this transition is visited infinitely often if the input is normal because all entries of the stationary distribution are positive [15, Thm 1.1(b)]. This guarantees that each normal word x is the input label of an accepting run and that $\mathcal{T}(x)$ is well-defined.

Some frequencies are obtained as stationary distributions of Markov chains [7, Thm 4.1]. For that purpose, we associate a Markov chain $\mathcal{M}$ to each strongly connected automaton $\mathcal{A}$. For simplicity, we assume that the state set Q of $\mathcal{A}$ is the set $\{1, \dots, \#Q\}$. The state set of the Markov chain is the same set $\{1, \dots, \#Q\}$. The transition matrix of the Markov chain is the matrix $P = (p_{i,j})_{1 \leq i,j \leq \#Q}$ where each entry $p_{i,j}$ is equal to $\#\{a : i \xrightarrow{a} j\} / \#A$. Note that $\#\{a : i \xrightarrow{a} j\}$ is the number of transitions from i to j . Since the automaton is assumed to be deterministic and complete, the matrix P is stochastic. If the automaton $\mathcal{A}$ is strongly connected, the Markov chain is irreducible and it has therefore a unique stationary distribution π such that $\pi P = \pi$. Note that all entries of P and π are rational numbers which can be effectively computed from $\mathcal{A}$. The vector π is called the *distribution* of $\mathcal{A}$. This definition as well as Lemmas 5 and 6 below also apply to input-deterministic transducers by ignoring the output labels.

Each run of either an automaton or a transducer can be seen as a sequence of transitions. Therefore, the notion of the *frequency* $\text{freq}(\rho, \gamma)$ of a finite run γ in an infinite run ρ is defined as in Section 2. Note that $\text{freq}(\rho, \gamma)$ is a limit and might not exist. This notion applies also to states seen as runs of length 0.

The following lemma which is Lemma 4.5 in [14] states that if the automaton $\mathcal{A}$ is strongly connected, then the run on a normal input visits each state with a frequency. Moreover, these frequencies are independent of the input as long as it is normal.

Lemma 5. *Let $\mathcal{A}$ be a deterministic and complete automaton which is strongly connected and let ρ be a run in $\mathcal{A}$ labeled by a normal word. Then the frequency $\text{freq}(\rho, q)$ is equal to π_q for each state q where π is the stationary distribution of the Markov chain associated to $\mathcal{A}$.*

Let γ be a finite run whose first state is p and let ρ be an infinite run. We call *conditional frequency* of γ in ρ the ratio $\text{freq}(\rho, \gamma) / \text{freq}(\rho, p)$. It is defined as soon as both frequencies $\text{freq}(\rho, \gamma)$ and $\text{freq}(\rho, p)$ do exist.

Lemma 6. *Let $\mathcal{A}$ be a deterministic and complete automaton which is strongly connected and let ρ be a run in $\mathcal{A}$ labeled by a normal word. The conditional frequency of each run of length k is $1/(\#A)^k$.*

Proof. Let $p * w$ denote the unique run starting in state p with label w . We now define an automaton whose states are the runs of length n in $\mathcal{A}$. We let $\mathcal{A}^n$ denote the automaton whose state set is $\{p * w : p \in Q, w \in A^n\}$ and whose set of transitions is defined by

$$\{(p * bw) \xrightarrow{a} (q * wa) : p \xrightarrow{b} q \text{ in } \mathcal{A}, a, b \in A \text{ and } w \in A^{n-1}\}$$

The Markov chain associated with the automaton $\mathcal{A}^n$ is called the *snake* Markov chain. See Problems 2.2.4, 2.4.6 and 2.5.2 (page 90) in [7] for more details. It is pure routine to check that the stationary distribution ξ of $\mathcal{A}^n$ is given by $\xi_{p*w} = \pi_p / (\#A)^n$ for each state p and each word w of length n and where π

is the stationary distribution of $\mathcal{A}$. To prove the statement, apply Lemma 5 to the automaton $\mathcal{A}^n$. $\square$

The output labels of the transitions in $\mathcal{T}$ may have arbitrary lengths. We first describe the construction of an equivalent transducer $\mathcal{T}'$ such that all output labels in $\mathcal{T}'$ have length at most 1. We call this transformation *normalization* and it consists in *replacing* each transition $p \xrightarrow{a|v} q$ in $\mathcal{T}$ such that $|v| \geq 2$ by n transitions:

$$p \xrightarrow{a|b_1} q_1 \xrightarrow{\lambda|b_2} q_2 \cdots q_{n-1} \xrightarrow{\lambda|b_n} q$$

where $q_1, q_2, \dots, q_{n-1}$ are new states and $v = b_1 \cdots b_n$. We refer to p as the parent of $q_1, \dots, q_{n-1}$.

To illustrate the construction, the normalized transducer obtained from the transducer of Figure 1 is pictured in Figure 4.




Figure 4: The normalized transducer of the transducer pictured in Fig. 1

From the normalized transducer $\mathcal{T}'$ we construct a weighted automaton $\mathcal{A}$ with the same state set as $\mathcal{T}'$. We define transitions between every pair of states p, q for each symbol b in B , that is the transition $p \xrightarrow{b} q$ is defined for all states p, q and for every symbol $b \in B$. To assign weight to transitions in $\mathcal{A}$, we first assign weights to transitions in $\mathcal{T}'$ as follows. Each transition starting from a state in $\mathcal{T}$ (and having a symbol as input label) has weight $1/\#A$ and each transition starting from a newly added state (and having the empty word as input label) has weight 1. Note that for each state p in $\mathcal{T}'$, the sum of weights of transitions starting from p is 1. We now consider separately transitions that generate empty output from those that do not. Consider the $Q \times Q$ matrix E whose (p, q) -entry is given for each pair (p, q) of states by

$$E_{p,q} = \sum_{a \in A} \text{weight}_{\mathcal{T}'}(p \xrightarrow{a|\lambda} q).$$

Let E^* be the matrix defined by $E^* = \sum_{k \geq 0} E^k$. Hence the entry $E^*_{p,q}$ is the sum of weights of all finite runs with empty output going from p to q . The matrix E^* can be computed because it is the solution of the system $E^* = EE^* + I$ where I is the identity matrix. This proves in particular that all its entries are rational numbers.

For each symbol $b \in B$ consider the $Q \times Q$ matrix N_b whose (p, q) -entry is given for each pair (p, q) of states by

$$(N_b)_{p,q} = \sum_{a \in A \cup \{\lambda\}} \text{weight}_{\mathcal{T}'}(p \xrightarrow{a|b} q).$$

We define the weight of a transition $p \xrightarrow{b} q$ in $\mathcal{A}$ as

$$\text{weight}_{\mathcal{A}}(p \xrightarrow{b} q) = (E^* N_b)_{p,q}.$$

To assign initial weight to states we consider the Markov chains $\mathcal{M}$ whose transition matrix is the stochastic matrix $P = \sum_{b \in B} E^* N_b$. The fact that this matrix is indeed stochastic follows from the observation that, for each state p , the set of input labels of the runs in $\bigcup_{q \in Q, b \in B} \Gamma_{p,b,q}$ (see below for the definition of $\Gamma_{p,b,q}$) is a maximal prefix code and Proposition 3.8 in [5, Chap. II.3]. The initial vector of $\mathcal{A}$ is the stationary distribution π of $\mathcal{M}$, that is, the line vector π such that $\pi P = \pi$. We assign to each state q the initial weight π_q . Finally we assign final weight 1 to all states.

We give below the matrices E , E^* , N_a , N_b and P and the initial vector π of the weighted automaton obtained from the transducer pictured in Figure 4.

$$\begin{aligned} E &= \begin{bmatrix} 0 & 1/2 & 0 & 0 & 0 \\ 0 & 0 & 1/2 & 0 & 0 \\ 0 & 0 & 1/2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} & \text{and} & E^* &= \begin{bmatrix} 1 & 1/2 & 1/2 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \\ N_a &= \begin{bmatrix} 1/2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \end{bmatrix} & \text{and} & N_b &= \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1/2 \\ 0 & 0 & 0 & 1/2 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \\ P &= \begin{bmatrix} 1/2 & 0 & 0 & 1/4 & 1/4 \\ 0 & 0 & 0 & 1/2 & 1/2 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \end{bmatrix} & \text{and} & \pi &= [2/3 \quad 0 \quad 0 \quad 1/6 \quad 1/6] \end{aligned}$$

Proposition 7. *The automaton $\mathcal{A}$ computes frequencies, that is, for every normal word x and any finite word w in B^* , $\text{weight}_{\mathcal{A}}(w) = \text{freq}(\mathcal{T}(x), w)$.*

The proof of the proposition requires some preliminary results. The following lemma is straightforward and follows directly from the normalization of $\mathcal{T}$ into $\mathcal{T}'$.

Lemma 8. *Both transducers $\mathcal{T}$ and $\mathcal{T}'$ realize the same function, that is, $\mathcal{T}(x) = \mathcal{T}'(x)$ for any infinite word x .*

Let us recall that a set of words L (resp. runs) is called *prefix-free* if no word in L is a proper prefix of another word in L . Let Γ be a set of finite runs and ρ be an infinite run. The limit $\text{freq}(\rho, \Gamma)$ is defined as the limit as n goes

to infinity of the ratio between the number of occurrences of runs in Γ in the prefix of length n of ρ and n . If Γ is prefix-free (not to count twice the same start), the following equality holds

$$\text{freq}(\rho, \Gamma) = \sum_{\gamma \in \Gamma} \text{freq}(\rho, \gamma)$$

assuming that each term of the right-hand sum does exist. If Γ is a set of finite runs starting from the same state p , the *conditional frequency* of Γ in a run ρ is defined as the ratio between the frequency of Γ in ρ and the frequency of p in ρ , that is, $\text{freq}(\rho, \Gamma) / \text{freq}(\rho, p)$. Furthermore if Γ is prefix-free, the conditional frequency of Γ is the sum of the conditional frequencies of its elements.

Let x be a fixed normal word and let ρ and ρ' be respectively the runs in $\mathcal{T}$ and $\mathcal{T}'$ with label x . By Lemma 5, the frequency $\text{freq}(\rho, q)$ of each state q is π_q where π is the stationary distribution of $\mathcal{T}$. The following lemma gives the frequency of states in ρ' .

Lemma 9. *There exists a constant C such that if q is a state of $\mathcal{T}$, then $\text{freq}(\rho', q) = \text{freq}(\rho, q)/C$ and if q is newly created, then $\text{freq}(\rho', q) = \text{freq}(\rho', p)/\#A$ where p is the parent of q .*

Proof. Observe that there is a one-to-one relation between runs labeled with normal words in $\mathcal{T}$ and in $\mathcal{T}'$. More precisely, each transition τ in ρ is replaced by $\max(1, |v_\tau|)$ transitions in ρ' (where v_τ is the output label of τ).

By combining Lemmas 5 and 6, each transition of $\mathcal{T}$ has a frequency in ρ . The first result follows by taking $C = \sum_{\tau} \text{freq}(\rho, \tau) \cdot \max(1, |v_\tau|)$ where the summation is taken over all transitions τ of $\mathcal{T}$ and v_τ is implicitly the output label of τ . The second result follows from Lemma 6 stating that transitions have a conditional frequency of $1/\#A$ in ρ . $\square$

For each pair (p, q) of states and each symbol $b \in B$, consider the set $\Gamma_{p,b,q}$ of runs from p to q in $\mathcal{T}'$ that have empty output labels for all their transitions but the last one, which has b as output label.

$$\Gamma_{p,b,q} = \{p \xrightarrow{a_1|\lambda} \dots \xrightarrow{a_n|\lambda} q_n \xrightarrow{a_{n+1}|b} q : n \geq 0, q_i \in Q, a_i \in A \cup \{\lambda\}\}$$

and let Γ be the union $\bigcup_{p,q \in Q, b \in B} \Gamma_{p,b,q}$. Note that the set Γ is prefix-free. Therefore, the run ρ' has a unique factorization $\rho' = \gamma_0 \gamma_1 \gamma_2 \dots$ where each γ_i is a finite run in Γ and the ending state of γ_i is the starting state of γ_{i+1} . Let $(p_i)_{i \geq 0}$ and $(b_i)_{i \geq 0}$ be respectively the sequence of states and the sequence of symbols such that γ_i belongs to $\Gamma_{p_i, b_i, p_{i+1}}$ for each $i \geq 0$. Let us call ρ'' the sequence $p_0 p_1 p_2 \dots$ of states of $\mathcal{T}'$.

Lemma 10. *For each state q of $\mathcal{T}'$, the frequency $\text{freq}(\rho'', q)$ does exist.*

Proof. The sequence ρ'' is a subsequence of the sequence of states in the run ρ' . An occurrence of a state q in ρ' is removed whenever the output of the previous transition is empty.

Consider the transducer $\hat{\mathcal{T}}$ obtained by splitting each state q of $\mathcal{T}$ into two states q^λ and q° in such a way that transitions with an empty output label end in a state q^λ and other transitions end in a state q° . Then each transition $p \xrightarrow{a|v} q$ is replaced by either the two transitions $p^\lambda \xrightarrow{a|v} q^\lambda$ and

$p^o \xrightarrow{a|v} q^\lambda$ if v is empty or by the two transitions $p^\lambda \xrightarrow{a|v} q^o$ and $p^o \xrightarrow{a|v} q^o$ otherwise. The state q_0^λ becomes the new initial state and non reachable states are removed. Let $\hat{\rho}$ be the run in $\hat{\mathcal{T}}$ labeled by x . By Lemma 5, the frequencies $\text{freq}(\hat{\rho}, q^\lambda)$ and $\text{freq}(\hat{\rho}, q^o)$ do exist. Now consider the normalization $\hat{\mathcal{T}}'$ of $\hat{\mathcal{T}}$ and the run $\hat{\rho}'$ in $\hat{\mathcal{T}}'$ labeled by x . By a lemma similar to Lemma 9, the frequencies $\text{freq}(\hat{\rho}', q^\lambda)$ and $\text{freq}(\hat{\rho}', q^o)$ do exist. The sequence ρ'' is obtained from $\hat{\rho}'$ by removing each occurrence of states q^λ and keeping occurrences of states q^o . It follows that the frequency of each state does exist in ρ'' . $\square$

Proof of Proposition 7. By Lemma 6, the conditional frequency of each finite run γ of length n in ρ is $1/(\#A)^n$. It follows that the conditional frequency of each finite run γ' in ρ' is equal to its weight in $\mathcal{T}'$ as defined auxiliary when defining the weights of transitions in the weighted automaton $\mathcal{A}$. This proves that the weight of the transition $p \xrightarrow{b} q$ in $\mathcal{A}$ is exactly the conditional frequency of the set $\Gamma_{p,b,q}$ for each triple (p, b, q) in $Q \times B \times Q$. More generally, the product of the weights of the transitions $p_0 \xrightarrow{b_1} p_1 \cdots p_{n-1} \xrightarrow{b_n} p_n$ is equal to the conditional frequency of the set $\Gamma_{p_0,b_1,p_1} \cdots \Gamma_{p_{n-1},b_n,p_n}$ in ρ' .

It remains to prove that the frequency of each state q in ρ'' is indeed its initial weight in the automaton $\mathcal{A}$. Let us recall that the initial vector of $\mathcal{A}$ is the stationary distribution of the stochastic matrix P whose (p, q) -entry is the sum $\sum_{b \in B} \text{weight}_{\mathcal{A}}(p \xrightarrow{b} q)$, which is the conditional frequency of pq (as a word of length 2) in ρ'' . It follows that the frequencies of states in ρ'' must be the stationary of the matrix P .

Since the frequency of a word $v = b_1 \cdots b_n$ in $\mathcal{T}'(x)$ is the same as the sum over all sequences $p_0, p_1, \dots, p_{n+1}$ of the frequencies of $\Gamma_{p_0,b_1,p_1} \cdots \Gamma_{p_n,b_n,p_{n+1}}$ in ρ' , it is the weight of the word v in the automaton $\mathcal{A}$. $\square$

Proofs of Theorems 1 and 2. To complete the proof of Theorems 1 and 2, we exhibit an algorithm deciding in cubic time whether an input-deterministic transducer preserves normality. Let $\mathcal{T}$ be the transducer $\langle Q, A, B, \delta, \{q_0\} \rangle$. By definition, its size is the sum $\sum_{\tau \in \delta} |\tau|$, where the size of a single transition $\tau = p \xrightarrow{a|v} q$ is $|\tau| = |av|$. We consider the alphabets to be fixed so they are not taken into account when calculating complexity.




Figure 5: Weighted automaton $\mathcal{B}$ such that $\text{weight}_{\mathcal{B}}(w) = 1/(\#A)^{|w|}$

By Proposition 3, the algorithm decomposes the transducer into strongly connected components and checks that each recurrent one does preserve normality. This is achieved by computing the weighted automaton $\mathcal{A}$ and checking that the weight of each word w is $1/(\#A)^{|w|}$. This latter step is performed by comparing $\mathcal{A}$ with the weighted automaton $\mathcal{B}$ such that $\text{weight}_{\mathcal{B}}(w) = 1/(\#A)^{|w|}$. The automaton $\mathcal{B}$ is pictured in Figure 5.

Input: $\mathcal{T} = \langle Q, A, B, \delta, q_0 \rangle$ an input-deterministic complete transducer.

Output: **True** if $\mathcal{T}$ preserves normality and **False** otherwise.

Procedure:

- I. Compute the strongly connected components of $\mathcal{T}$
- II. For each recurrent strongly connected component S_i of $\mathcal{T}$:
 1. Compute the normalized transducer $\mathcal{T}'$, equivalent to S_i .
 2. Use $\mathcal{T}'$ to build the weighted automaton $\mathcal{A}$:
 - a. Compute the weights of the transitions of $\mathcal{A}$.
Compute the matrix E
Compute the matrix E^* solving $(I - E)X = I$
For each $b \in B$, for each $p, q \in Q$:
compute the matrix N_b
define the transition $p \xrightarrow{b} q$ with weight $(E^* N_b)_{p,q}$.
 - b. Compute the stationary distribution π of the Markov chain induced by $\mathcal{A}$.
 - c. Assign initial weight $\pi[i]$ to each state i , and let final weight be 1 for all states.
 3. Compare $\mathcal{A}$ against the automaton $\mathcal{B}$ using Schützenberger’s algorithm [9] to check whether they realize the same function.
 4. If they do not compute the same function, return **False**.
- III. Return **True**

Now we analyze the complexity of the algorithm. Computing recurrent strongly connected components can be done in $O(|Q|^2) \leq O(n^2)$ using Kosaraju’s algorithm if the transducer is implemented with an adjacency matrix [11, Section 22.5].

We refer to the size of the component $|S_i|$ as n_i . The cost of normalizing the component is $O(n_i^2)$, mainly from filling the new adjacency matrix. The most expensive step when computing the transitions and their weight is to compute E^* . The cost is $O(n_i^3)$ to solve the system of linear equations. To compute the weights of the states we have $O(n_i^3)$ to solve the system of equations to find the stationary distribution. Comparing the automaton to the one computing the expected frequencies can be done in time $O(n_i^3)$ [9] since the coefficients of both automata are in $\mathbb{Q}$.

□

Acknowledgements

The authors would like to thank Verónica Becher for many fruitful discussions and suggestions. Both authors are members of the Laboratoire International Associé INFINIS, CONICET/Universidad de Buenos Aires–CNRS/Université Paris Diderot and they are supported by the ECOS project PA17C04. Carton is also partially funded by the DeLTA project (ANR-16-CE40-0007).

References

- [1] V. N. Agafonov. Normal sequences and finite automata. *Soviet Mathematics Doklady*, 9:324–325, 1968.
- [2] V. Becher and O. Carton. Normal numbers and computer science. In *Sequences, Groups, and Number Theory*, pages 233–269. Springer, 2018.
- [3] V. Becher and P. A. Heiber. Normal numbers and finite automata. *Theoretical Computer Science*, 477:109–116, 2013.
- [4] J. Berstel. *Transductions and Context-Free Languages*. B.G. Teubner, 1979.
- [5] J. Berstel and D. Perrin. *Theory of Codes*. Academic Press, Inc., 1985.
- [6] É. Borel. Les probabilités dénombrables et leurs applications arithmétiques. *Rend. Circ. Mat. Palermo*, 27(2):247–271, 1909.
- [7] P. Brémaud. *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer, 2008.
- [8] Y. Bugeaud. *Distribution modulo one and Diophantine approximation*, volume 193 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 2012.
- [9] A. Cardon and M. Crochemore. Détermination de la représentation standard d’une série reconnaissable. *ITA*, 14(4):371–379, 1980.
- [10] D. G. Champernowne. The construction of decimals normal in the scale of ten. *Journal of the London Mathematical Society*, 1(4):254–260, 1933.
- [11] T. H. Cormen, Ch. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms, Third Edition*. The MIT Press, 3rd edition, 2009.
- [12] J. Dai, J. Lathrop, J. Lutz, and E. Mayordomo. Finite-state dimension. *Theoretical Computer Science*, 310:1–33, 2004.
- [13] J. Sakarovitch. *Elements of Automata Theory*. Cambridge University Press, 2009.
- [14] C. P. Schnorr and H. Stimm. Endliche Automaten und Zufallsfolgen. *Acta Informatica*, 1:345–359, 1972.
- [15] E. Seneta. *Non-negative Matrices and Markov Chains*. Springer, 2006.

Olivier Carton
Institut de Recherche en Informatique Fondamentale
Université Paris Diderot, France
Olivier.Carton@irif.fr

Elisa Orduna
Universidad de Buenos Aires, Argentina
eorduna@dc.uba.ar