MATHEMATICS AND STATISTICS

Statistical Topics and Stochastic Models for Dependent Data with Applications

Edited by
Vlad Stefan Barbu and Nicolas Vergne

ISTE

WILEY

Series Editor
Nikolaos Limnios

Statistical Topics and Stochastic Models for Dependent Data with Applications

Edited by
Vlad Stefan Barbu
Nicolas Vergne

iSTE

WILEY

First published 2020 in Great Britain and the United States by ISTE Ltd and John Wiley & Sons, Inc.

ISTE Ltd John
27-37 St George's Road
London SW19 4EU
UK
www.iste.co.uk

Wiley & Sons, Inc.
111 River Street
Hoboken, NJ 07030
USA
www.wiley.com

Library of Congress Control Number: 2020938718

British Library Cataloguing-in-Publication Data

A CIP record for this book is available from the British Library

ISBN 978-1-78630-603-6

Preface

This collective book stems from a workshop that took place in Rouen, October 3–5, 2018, within the framework of the project Random Models and Statistical Tools, Informatics and Combinatorics (MOUSTIC – Modèles aléatoires et Outils Statistiques, Informatiques et Combinatoires), financed by the Region of Normandy and the European Regional Development Fund. The main idea was to bring together leading scientists working on probabilistic and statistical topics for dependent data, as well as on associated applications.

This is the way this book was written, with the intention to offer to the scientific community a part of the latest advances in the field of stochastic modeling for dependent data, authored by leading experts in the field. This is a crucial aspect in a time when we face an increasing need for more and more complex models capable of capturing the main features of increasingly complex applications. From a technical point of view, our book is important for gathering theoretical developments and applications related to Markov type models (semi-Markov processes, autoregressive processes, piecewise deterministic Markov processes, and variable length Markov chains) as well as probabilistic/statistical techniques issued from information theory, based on divergence measures and entropies.

This volume is divided into three parts: the first one examines Markov and semi-Markov processes, the second one deals with autoregressive processes and the last one presents divergence measures and entropies. Particular attention is given to applications of these methods in various fields such as finance, DNA analysis, quantum physics and survival analysis.

The workshop in Rouen and, consequently, the present book, would not have been possible without the support of the Laboratory of Mathematics Raphaël Salem, the Laboratory of Mathematics Nicolas Oresme, the University of Rouen Normandy, the Region of Normandy, the French Statistical Society-SFdS, the MAS group of the French Society of Applied and Industrial Mathematics-SMAI, the Normandie-Mathématiques Research Federation and the Normastic Research Federation.

We would like to thank all the speakers of the workshop who contributed, although some of them indirectly, to the quality of the present volume. Our thanks go also to the anonymous reviewers for their valuable work: without their support, this volume could not have been successfully completed.

We would also like to thank Professor Nikolaos Limnios for proposing and encouraging us to elaborate this collective volume as well as the editorial staff of ISTE Ltd for their technical support.

Vlad Stefan BARBU
Nicolas VERGNE
Rouen, May 2020

PART 1
Markov and Semi-Markov Processes

1
Variable Length Markov Chains, Persistent Random Walks: A Close Encounter

We consider a walker on the line that at each step keeps the same direction with a probability that depends on the time already spent in the direction the walker is currently moving. These walks with memories of variable length can be seen as generalizations of directionally reinforced random walks (DRRWs) introduced in Mauldin *et al*. (1996). We give a complete and usable characterization of the recurrence or transience in terms of the probabilities to switch the direction. These conditions are related to some characterizations of existence and uniqueness of a stationary probability measure for a particular Markov chain: in this chapter, we define the general model for words produced by a variable length Markov chain (VLMC) and we introduce a key combinatorial structure on words. For a subclass of these VLMC, this provides necessary and sufficient conditions for existence of a stationary probability measure.

1.1. Introduction

This is the story of the encounter between two worlds: the world of random walks and the world of VLMCs. The meeting point turns around the semi-Markov property of underlying processes.

In a VLMC, unlike fixed-order Markov chains, the probability to predict the next symbol depends on a possibly unbounded part of the past, the length of which depends on the past itself. These relevant parts of pasts are called *contexts*. They are stored in *context tree*. With each context, a probability distribution is associated, prescribing the conditional probability of the next symbol, given this context.

VLMCs are now widely used as random models for character strings. They were introduced in Rissanen (1983) to perform data compression. When they have a finite memory, they provide a parsimonious alternative to fixed-order Markov chain models, in which the number of parameters to estimate grows exponentially fast with the order; they are also able to capture finer properties of character sequences. When they have infinite memory – this will be our case of study in this chapter – they are a tractable way to build non-Markov models and they may be considered as a subclass of "*chaînes à liaisons complètes*" (Doeblin and Fortet 1937) or "chains with infinite order" (Harris 1955)

VLMCs are used in bioinformatics, linguistics and coding theory to model how words grow or to classify words. In bioinformatics, both for protein families and DNA sequences, identifying patterns that have a biological meaning is a crucial issue. Using VLMC as a model enables one to quantify the influence of a meaning pattern by giving a transition probability on the following letter of the sequence. In this way, these patterns appear as contexts of a context tree. Note that their length may be unbounded (Bejerano and Yona 2001).

In addition, if the context tree is recognized to be a signature of a family (say, of proteins), this gives an efficient statistical method to test whether or not two samples belong to the same family (Busch *et al*. 2009).

Therefore, estimating a context tree is an issue of interest and many authors (statisticians or not, applied or not) stress the fact that the height of the context tree should not be supposed to be bounded. This is the case in Galves and Leonardi (2008) where the algorithm CONTEXT is used to estimate an unbounded context tree, or in Garivier and Leonardi (2011). Furthermore, as explained in Csiszár and Talata (2006), the height of the estimated context tree grows with the sample size, so that estimating a context tree by assuming *a priori* that its height is bounded is not realistic.

There is extensive literature on the construction of efficient estimators of context trees, as well for finite or infinite context trees. This chapter is not a review of statistics issues, which would already be relevant for finite memory VLMC. This is a study of the probabilistic properties of infinite memory VLMC as random processes, and more specifically of the main property of interest for such processes: existence and uniqueness of a stationary measure.

As has already been said, VLMC are a natural generalization to infinite memory of Markov chains. It is usual to index a sequence of random variables forming a Markov chain with positive integers and to make the process grow to the right. The main drawback of this habit for an infinite memory process is that the sequence of the process is read from left to right, whereas the (possibly infinite) sequence giving the past needed to predict the next symbol is read in the context tree from right to left, thus giving rise to confusion and lack of readability. For this reason, in this chapter, the VLMC grows to the left. In this way, both the process sequence and the memory in the context tree are read from left to right.

Classical random walks have *independent* and identically distributed increments. In the literature, *persistent* random walks (PRMs), also called *Goldstein-Kac random walks* or *correlated random walks*, refer to random walks having a Markov chain of finite order as an increment process. For such walks, the dynamics of trajectories has a short memory of given length and the random walk itself is not Markovian anymore. What happens whenever the increments depend on a *non-bounded* past memory?

Consider a walker on $\mathbb{Z}$, allowed to increment its trajectory by –1 or 1 at each step of time. Assume that the probability to keep the current direction ±1 depends on the time already spent in the said direction – the distribution of increments thus acts as a reinforcement of the dependency from the past. More precisely, the process of increments of such a one-dimensional random walk is a Markov chain on the set of (right-)infinite words, with variable – and unbounded – length memory: a VLMC. The concerned VLMC is defined in section 1.3.1. It is based on a context tree called a *double comb*. Later, section 1.3.2 deals with a two-dimensional persistent random walk defined in an analogous manner on $\mathbb{Z}^2$ by a VLMC based on a context tree called a *quadruple comb*.

These random walks that have an unbounded past memory can be seen as a generalization of “directionally reinforced random walks (DRRW)” introduced by Mauldin *et al*. (1996), in the sense that the persistence times are anisotropic ones. For a one-dimensional random walk associated with a double comb, a complete characterization of recurrence and transience, in terms of changing (or not) direction probabilities, is given in section 1.3.1. More precisely, when one of the random times spent in a given direction (the so-called *persistence times*) is an integrable random variable, the recurrence property is equivalent to a classical drift-vanishing. In all other cases, the walk is transient unless the weight of the tail distributions of both persistent times are equal. In two-dimensional random walk, sufficient conditions of transience of recurrence are given in section 1.3.2.

Actually, because of the very specific form of the underlying driving VLMC, these PRWs turn out to be in one-to-one correspondence with so-called *Markov additive processes*. Section 1.5 examines the close links between PRWs, Markov additive processes, semi-Markov chains and VLMC.

In section 1.2, the definition of a general VLMC and a couple of examples are given. In section 1.3, the PRWs are defined and known results on their recurrence properties are collected. In view of section 1.5 where we show how PRW and VLMC meet through the world of semi-Markov chains, section 1.4 is devoted to results – together with a heuristic approach – on the existence and unicity of stationary measures for a VLMC.

1.2. VLMCs: definition of the model

Let A be a finite set, called the *alphabet*. Here A will most often be the standard alphabet $A = \{0, 1\}$, but also $A = \{d, u\}$ (for *down* and *up*) or $A = \{\mathrm{n, e, w, s}\}$ (for the cardinal directions). Let

$$\mathcal{R} = \{\alpha\beta\gamma\cdots : \alpha, \beta, \gamma, \cdots \in \mathcal{A}\}$$

be the set of *right-infinite* words over A. written by simple concatenation. A VLMC on A, defined below and most often denoted by $(U_n)_{n\in\mathbb{N}}$, is a particular type of R-valued discrete time Markov chain where:

– the process evolves between time n and time $n + 1$ by adding one letter *on the left* of U_n;

– the transition probabilities between time n and time $n + 1$ depend on a finite – but not bounded – prefix[1] of the current word U_n.

Giving a formal frame of such a process leads to the following definitions. For a complete presentation of VLMC, one can also refer to Cénac *et al*, (2012).

As usual, a *tree on* A is a set T of finite words – namely a subset of $\cup_{n\in\mathbb{N}} A^n$ – which contains the empty word $\emptyset$ (the *root* of T) and which is prefix-stable: for all finite words u, v, $uv \in T \Rightarrow u \in T$. A tree is made of *internal nodes* ($u \in T$ is internal when $\exists\alpha \in A$, $u\alpha \in T$) and of *leaves* ($u \in T$ is a leaf when it has no child: $\forall\alpha \in A$, $u\alpha \notin T$).

DEFINITION 1.1 (Context tree). – *A* context tree *on A is a saturated tree on A having an at most countable set of infinite branches*.

The tree T is *saturated* whenever any internal node has #(A) children: for any finite word u and for any $\alpha \in A$, $u\alpha \in T \Rightarrow (\forall\beta \in A, u\beta \in T)$. A right-infinite word on A is an *infinite branch* of T when all its finite prefixes belong to T.

Following the vocabulary introduced by Rissanen, a *context* of the tree is a leaf or an infinite branch. A finite or right-infinite word on A is an *external node* when it is neither internal nor a context. See Figure 1.1 which illustrates these definitions, as well as the `pref` function defined hereunder.

DEFINITION 1.2 (`pref` function).– *Let T be a context tree. If w is any external node or any context, the symbol* `pref` *w denotes the longest (finite or infinite) prefix of w that belongs to T*.

In other words, `pref` w is the only context c for which $w = c\cdots$ For a more visual presentation, hang w by its head (its left-most letter) and insert it into the tree; the only context through which the word goes out of the tree is its `pref`.

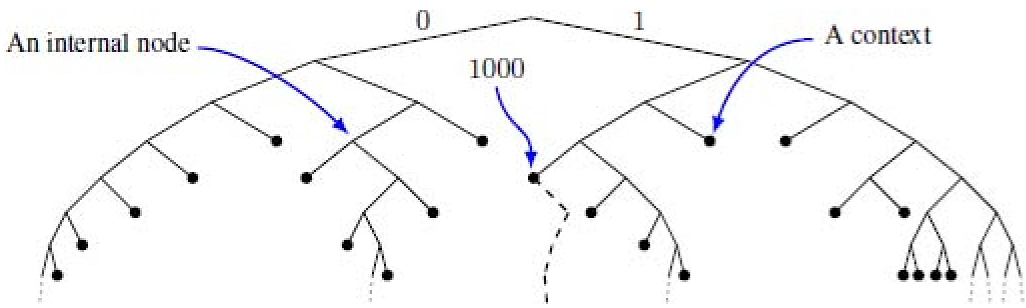

Figure 1.1. *A context tree on the alphabet $A = \{0,1\}$. The dotted lines are possibly the beginning of infinite branches. Any word that writes $1000\ \cdots$, like the one represented by the dashed line, admits 1000 as a* pref. *For a color version of this figure, see www.iste.co.uk/barbu/data.zip*

With these definitions, it is now possible to define a VLMC.

DEFINITION 1.3 (VLMC).– *Let T be a context tree. For every context c of T, let q_c be a probability measure on A. The* VLMC *defined by T and by the $(q_c)_c$ is the R-valued discrete-time Markov chain $(U_n)_{n\in\mathbb{N}}$ defined by the following transition probabilities:* $\forall n \in \mathbb{N}$, $\forall \alpha \in A$,

$$\mathbb{P}\left(U_{n+1} = \alpha U_n | U_n\right) = q_{\mathrm{pref}(U_n)}\left(\alpha\right). \qquad [1.1]$$

To get a realization of a VLMC as a process on R, take a (random) right infinite word

$$U_0 = X_0 X_{-1} X_{-2} X_{-3} \cdots$$

At each step of time $n \geq 0$, one gets U_{n+1} by adding a random letter X_{n+1} on the left of U_n:

$$\begin{aligned} U_{n+1} &= X_{n+1} U_n \\ &= X_{n+1} X_n \cdots X_1 X_0 X_{-1} X_{-2} \cdots \end{aligned}$$

under the conditional distribution [1.1].

REMARK 1.1.– *Probabilizing* a context tree consists, as in definition 1.3, of endowing it with a family of probability measures on the alphabet, indexed by the set of contexts. This vocabulary is used below.

REMARK 1.2.– Assume that the context tree is finite and denote its height by h ; in this condition, the VLMC is just a Markov chain of order h on A. On the contrary, when the context tree is infinite, and this is mainly our case of interest, the VLMC is generally *not* a Markov process on A.

EXAMPLE 1.1.– Take $A = \{\mathrm{n, e, w, s}\}$ as an (ordered) alphabet, so that the daughters of an internal node are represented, as shown on the left side of Figure 1.2. Making the transition probabilities $\mathbb{P}(U_{n+1} = \alpha U_n / U_n)$ depend only on the length of the largest prefix of the form n^k ($k \geq 0$) of U_n amounts to taking a comb as a context tree, as shown on the right side of Figure 1.2. Its finite contexts are the $n^k\ \alpha$ where $k \geq 0$ and $\alpha \in A \backslash \{\mathrm{n}\}$.

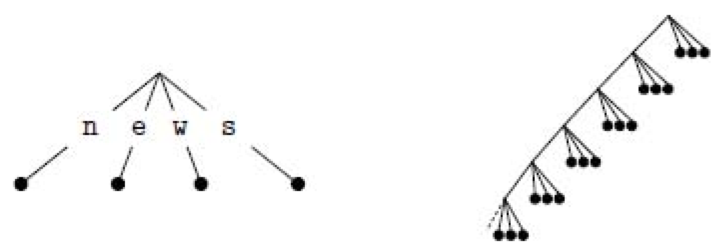

Figure 1.2. *On the left: how one can represent trees on A* = {n, e, w, s}. *On the right, the so-called* left comb *on A* = {n, e, w, s}

EXAMPLE 1.2.– Take again $A = \{\mathrm{n, e, w, s}\}$ as an alphabet. Making the transition probabilities $\mathbb{P}(U_{n+1} = \alpha \mathrm{U}_n | U_n)$ depend only on the length of the largest prefix of the form α^k ($k \geq 1$) of U_n, where α is *any* letter, amounts to taking a *quadruple comb* as a context tree, as shown on the right side of Figure 1.3. In the same vein, if one takes $A = \{u, d\}$, the *double comb* is the context tree, as shown on the left side of Figure 1.3. In the corresponding VLMC, the transitions depend only on the length of the last current run u^k or d^k, $k \geq 1$. The double comb and the quadruple comb are used below to define PRWs.

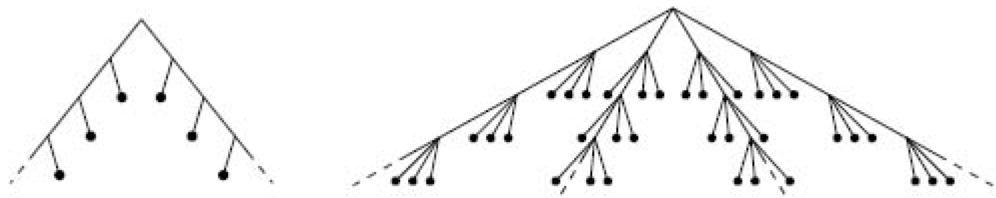

Figure 1.3. *The double comb and the quadruple comb*

EXAMPLE 1.3.– Take $A = \{0, 1\}$ (naturally ordered for the drawings). The left comb of right combs, shown on the left side of Figure 1.4, is the context tree of a VLMC that makes its transition probabilities depend on the largest prefix of U_n of the form 0^p1^q. If one has to take into consideration the largest prefix of the form 0^p1^q or 1^p0^q, one has to use the double comb of opposite combs, as shown on the right side of Figure 1.4.

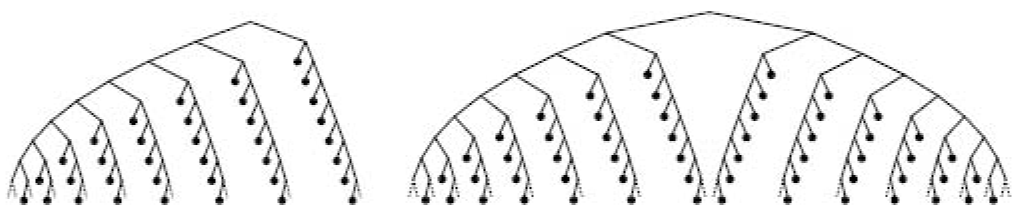

Figure 1.4. *Context trees on* $A = \{0, 1\}$; *the left comb of right combs (on the left) and a double comb of opposite combs (on the right)*

DEFINITION 1.4 (Non-nullness).– *A VLMC is called* non-null *when no transition probability vanishes, i.e. when* $q_c(\alpha) > 0$ *for every context* c *and for every* $\alpha \in A$.

Non-nullness appears below as an irreducibility-like assumption made on the driving VLMC of PRWs and for existence and unicity of an invariant probability measure for a general VLMC as well.

1.3. Definition and behavior of PRWs

In this section, the so-called *PRWs* are defined. A PRW is a random walk driven by some VLMC. In dimensions one and two, results on transience and the recurrence of PRW are given. These results are detailed and proven in Cénac *et al*. (2018b, 2013) in dimension one and in Cénac *et al*. (2020) in dimension two.

1.3.1. *PFWs in dimension one*

In this section, we deal with one-dimensional PRWs. Note that, contrary to the classical random walk, a PRW is generally not Markovian. Let $A: = \{d, u\} = \{-1,1\}$ (d for down and u for up) and consider the *double comb* on this alphabet as a context tree, probabilize it and denote by $(U_n)_n$ a realization of the associated VLMC. The n^{th} increment X_n of the PRW is given as the first letter of U_n: define the persistent random walk $S = (S_n)_{n \geq 0}$ by $S_0 = 0$ and, for $n \geq 1$,

$$S_n := \sum_{\ell=1}^{n} X_\ell, \tag{1.2}$$

so that for any $n \geq 1$, $m \geq 0$,

$$\begin{aligned} \mathbb{P}\left(S_{m+1} = S_m + 1 | U_m = d^n u \ldots\right) &= q_{d^n u}(u) \\ \mathbb{P}\left(S_{m+1} = S_m - 1 | U_m = u^n d \ldots\right) &= q_{u^n d}(d). \end{aligned}$$

Furthermore, for the sake of simplicity and without loss of generality, we condition the walk to start almost surely (a.s.) from $\{X_{-1} = u, X_0 = d\}$ – this amounts to changing the origin of time. In this model, a walker on a line keeps the same direction with a probability that depends on the discrete time already spent in the direction the walker is currently moving (see Figure 1.5). This model can be seen as a generalization of DRRWs introduced in Mauldin *et al*. (1996).

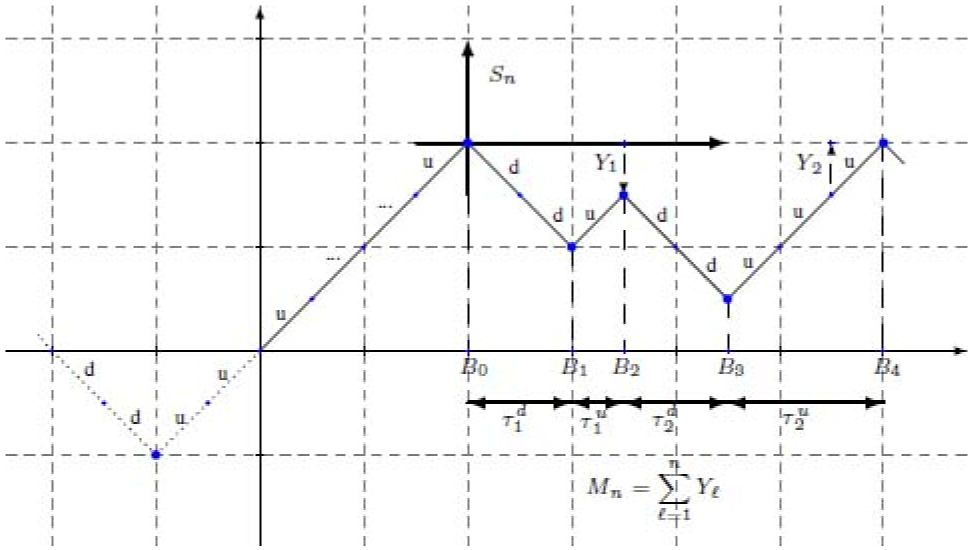

Figure 1.5. *A one-dimensional PRW. For a color version of this figure, see www.iste.co.uk/barbu/data.zip*

Taking different probabilized context trees would lead to different probabilistic impacts on the asymptotic behavior of resulting PRWs. Moreover, the characterization of the recurrent versus transient behavior is difficult in general. We state here exhaustive recurrence criteria for PRWs defined from a double comb.

In order to avoid trivial cases, we assume that S cannot be frozen in one of the two directions with a positive probability. Therefore, we make the following assumption.

ASSUMPTION 1.1 (Finiteness of the length of runs).– For any $\alpha, \beta \in \{u, d\}$, $\alpha \neq \beta$,

$$\lim_{n \to +\infty} \left(\prod_{k=1}^{n} q_{\alpha^k \beta}(\alpha) \right) = 0. \qquad [1.3]$$

Let τ_n^u and τ_n^d be, respectively, the length of the n^{th} rixe and of the n^{th} descent.

Then, by a renewal-type property (see Cénac *et al*. 2013, proposition 2.3), $(\tau_n^d)_{n \geq 1}$ and $(\tau_n^u)_{n \geq 1}$ are independent sequences of i.d.d. random variables. Their distribution tails are straightforwardly given by: for any $\alpha, \beta \in \{u, d\}$, $\alpha \neq \beta$ and $n \geq 1$,

$$\mathbb{P}(\tau_1^\alpha \geq n) = \prod_{k=1}^{n-1} q_{\alpha^k \beta}(\alpha). \qquad [1.4]$$

Note that assumption 1.1 amounts to supposing that the *persistence times* τ_n^d and τ_n^u are a.s. finite. The *jump times* (or breaking times) are: $B_0 = 0$ and, for $n \geq 1$,

$$B_{2n} := \sum_{k=1}^{n} \left(\tau_k^d + \tau_k^u\right) \text{ and } B_{2n+1} := B_{2n} + \tau_{n+1}^d. \qquad [1.5]$$

In order to deal with a more tractable random walk built with the possibly unbounded but i.d.d. increments $Y_n := \tau_n^u - \tau_n^d$, we introduce the underlying *skeleton* random walk $(M_n)_{n \geq 1}$, which is the original walk observed at the random times of up-to-down turns:

$$M_n := \sum_{k=1}^{n} Y_k = S_{B_{2n}}. \qquad [1.6]$$

Two main quantities play a key role in the asymptotic behavior, namely the expectations of the lengths of runs: with formula [1.4], let

$$\Theta_d := \mathbf{E}[\tau_1^d] = \sum_{n\geq 1} \prod_{k=1}^{n-1} q_{d^k u}(d) \text{ and } \Theta_u := \mathbf{E}[\tau_1^u] = \sum_{n\geq 1} \prod_{k=1}^{n-1} q_{u^k d}(u). \qquad [1.7]$$

Actually, Θ_d and Θ_u are already discussed in Cénac *et al*. (2013, proposition B1), where it is shown that the driving VLMC of a one-dimensional PRW admits a unique invariant probability measure if and only if $\Theta_d < \infty$ and $\Theta_u < \infty$,

Note that the expectation of Y_1 is well defined in $[-\infty, +\infty]$ whenever at least one of the persistence times τ_1^u or τ_1^d is integrable. Thus, as soon as $\Theta_d < \infty$ or $\Theta_u < \infty$, let

$$\mathbf{d}_M := \mathbf{E}[Y_1] = \underbrace{\Theta_u - \Theta_d}_{\in[-\infty,+\infty]} \qquad [1.8]$$

and

$$\mathbf{d}_S := \frac{\mathbf{E}[\tau_1^u] - \mathbf{E}[\tau_1^d]}{\mathbf{E}[\tau_1^u] + \mathbf{E}[\tau_1^d]} = \frac{\Theta_u - \Theta_d}{\Theta_u + \Theta_d} \in [-1, 1]. \qquad [1.9]$$

An elementary computation shows that $\mathbb{E}(M_n) = n\mathbf{d}_M$ and $\mathbb{E}(S_n) \sim n\mathbf{d}_S$ when n tends to infinity. Thus, d_M and d_S appear as asymptotic drifts when the walks $(M_n)_n$ and $(S_n)_n$ respectively, turn out to be transient (see Table 1.1). The behavior of the walk also depends on quantities $J_{\alpha|\beta}$, defined for α and $\beta \in A$, $\alpha \neq \beta$ by:

$$J_{\alpha|\beta} := \sum_{n=1}^{\infty} \frac{n\mathbb{P}(\tau_1^\alpha = n)}{\sum_{k=1}^{n} \mathbb{P}(\tau_1^\beta \geq k)}.$$

A complete and usable characterization of the recurrence and the transience of the PRW in terms of the probabilities to persist in the same direction or to switch is given in proposition 1.1. Its proof relies on a criterion of Erickson (1973), applied to the skeleton walk $(M_n)_n$, which is simpler to deal with because its increments are independent.

PROPOSITION 1.1.– Under the non-nullness assumption and assumption 1.1, the random walk $(S_n)_n$ is recurrent or transient as described in Table 1.1.

Table 1.1. *Recurrence versus transience (drifting) for $(S_n)_n$ in dimension one*

<table>
<tr><td></td><td colspan="2">$\Theta_u < \infty$</td><td colspan="2">$\Theta_u = \infty$</td></tr>
<tr><td rowspan="2">$\Theta_d < \infty$</td><td rowspan="2">Recurrent $d_S = 0$</td><td>Drifting $+\infty$ $d_S > 0$</td><td colspan="2" rowspan="2">Drifting $+\infty$</td></tr>
<tr><td>Drifting $-\infty$ $d_S < 0$</td></tr>
<tr><td rowspan="2">$\Theta_d = \infty$</td><td colspan="2" rowspan="2">Drifting $-\infty$</td><td rowspan="2">Recurrent $J_{u|d} = J_{d|u} = \infty$</td><td>Drifting $+\infty$ $\infty = J_{u/d} > J_{d/u}$</td></tr>
<tr><td>Drifting $-\infty$ $\infty = J_{d/u} > J_{u/d}$</td></tr>
</table>

The most fruitful situation emerges when both running times τ_1^u and τ_1^d have infinite means. In that case, the recurrence properties of $(S_n)_n$ are related to the behavior of the skeleton random walk $(M_n)_n$ defined in [1.6], the drift of which, d_M, is not defined. Thus, the behavior of $(S_n)_n$ depends on the comparison between the distribution tails of τ_1^u and τ_1^d defined in [1.4], expressed by the quantities $J_{\alpha \mid \beta}$. Note that the case when both $J_{u|d}$ and $J_{d|u}$ are finite does not appear in the table since it would imply that $\Theta_u < \infty$ and $\Theta_d < \infty$ (see Erickson 1973).

In all three other cases, the drift d_S is well defined and the PRW is recurrent if and only if $ds = 0$. In that case, $\lim_{n\to\infty} \frac{S_n}{n} = d_S = 0$. Note that modifying one transition q_c transforms a recurrent PRW into a transient one, since d_S becomes non-zero.

1.3.2. *PRWs In dimension two*

Take the alphabet $A := \{\mathrm{n}, \mathrm{e}, \mathrm{w}, \mathrm{s}\}$. Here, (e, n) stands for the canonical basis of $\mathbb{Z}^2$, w = –e and s = –n. Hence, the letters e, n, w and s stand for moves to the east, north, west and south, respectively. Having in mind a random walk with increments in A. any word of the form $\alpha\beta$, $\alpha, \beta \in A$, $\alpha \neq \beta$ is called a *bend*. For the sake of simplicity, we condition the walk to start a.s. with an ne bend: $\{X_{-1} = \mathrm{n}, X_0 = \mathrm{e}\}$.

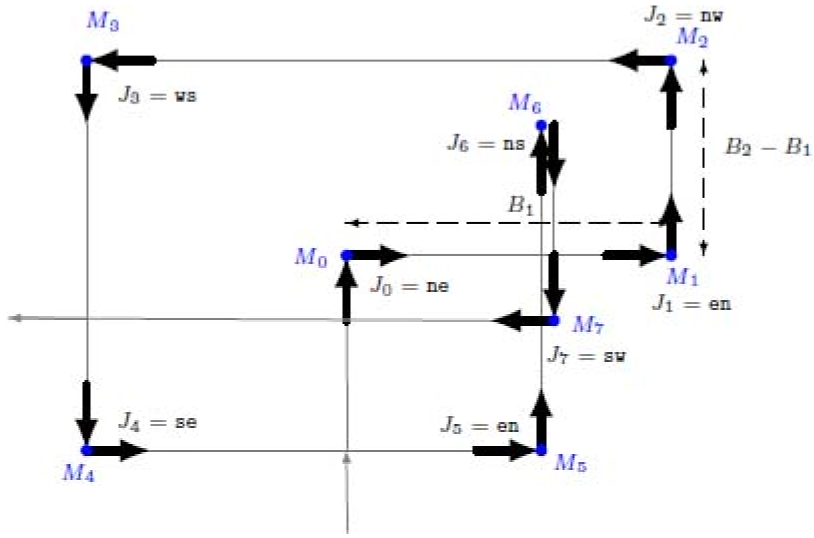

Figure 1.6. *A walk in dimension two. For a color version of this figure, see www.iste.co.uk/barbu/data.zip*

Take a non-null VLMC associated with a quadruple comb on A, as shown in Figure 1.3 the contexts are $\alpha^n \beta$ for $\alpha, \beta, \in A$, $\alpha \neq \beta$, $n \geq 1$ and the attached probability distributions are denoted by $q_{\alpha^n\beta}$. The two-dimensional PRW $(S_n)_n$ is defined, using this VLMC, as in formula [1.2].

Contrary to the one-dimensional PRWs, as detailed below, the probability to change direction depends on the time spent in the current direction but also on the previous direction. As in dimension one, we intend to avoid that S remains frozen in one of the four directions with a positive probability. Therefore, we make the following assumption, analogous to assumption 1.1 in dimension two.

ASSUMPTION 1.2 (Finiteness of the length of runs).– For any α, $\beta \in$ {n, e, w, s} $\alpha \neq \beta$,

$$\lim_{n \to +\infty} \left(\prod_{k=1}^{n} q_{\alpha^k \beta}(\alpha) \right) = 0. \qquad [1.10]$$

Let $(B_n)_{n \geq 0}$ be the *breaking times* defined inductively by

$$B_0 = 0 \quad \text{and} \quad B_{n+1} = \inf \{ k > B_n : X_k \neq X_{k-1} \}. \qquad [1.11]$$

As in dimension one, assumption 1.2 implies that the breaking times B_n are a.s. finite.

Define the so-called *internal chain* $(J_n)_{n \geq 0}$ by J_0 = ne and, for all $n \geq 1$,

$$J_n := X_{B_{n-1}} X_{B_n}. \qquad [1.12]$$

Let us illustrate these random variables with a small example, in which: $B_1 = 4$, $B_2 = 7$, $J_0 = X_{-1}X_0$, $J_1 = X_{B_0}X_{B_1} = X_0X_4$, $J_2 = X_{B_1}X_{B_2} = X_4X_7$.

n:	-1	0	1	2	3	4	5	6	7
X_n:	n	e	e	e	e	n	n	n	w
		$B_0 = 0$				$B_1 = 4$			$B_2 = 7$
		$J_0 = \mathtt{ne}$				$J_1 = \mathtt{en}$			$J_2 = \mathtt{nw}$

The process $(J_n)_{n \geq 0}$ is an irreducible Markov chain on the set of bends $S := \{\alpha\beta | \alpha \in A, \beta \in A, \alpha \neq \beta\}$. Its Markov kernel is defined by: for every β, α, $\gamma \in A$ with $\beta \neq \alpha$ and $\alpha \neq \gamma$,

$$P(\beta\alpha; \alpha\gamma) := \sum_{n=1}^{\infty} \left(\prod_{k=1}^{n-1} q_{\alpha^k \beta}(\alpha) \right) q_{\alpha^n \beta}(\gamma), \qquad [1.13]$$

the numbers $P(\alpha\beta, \gamma\delta)$ being 0 for every couple of bends not of the previous form. Remark that the non-nullness assumption (see definition 1.4) implies the irreducibility of $(J_n)_n$ and its aperiodicity. The state space S is finite so that $(J_n)_n$ is positive recurrent: it admits a unique invariant probability measure π_J.

Denote $T_0 = 0$ and $T_{n+1} := B_{n+1} - B_n$ for every $n \geq 0$. These waiting times (also called *persistence times*) are not independent, contrary to the one-dimensional case. The *skeleton random walk* $(M_n)_n \geq 0$ on $\mathbb{Z}^2$ – which is the PRW observed at the breaking times – is then defined as

$$M_n := S_{B_n} = \sum_{i=1}^{n} \left(\sum_{k=B_{i-1}+1}^{B_i} X_k \right) = \sum_{i=1}^{n} (B_i - B_{i-1}) X_{B_i}. \qquad [1.14]$$

Note that $(M_n)_n$ is generally not a classical RW with i.d.d. increments. Nevertheless, taking into account the additional information given by the internal Markov chain $(J_n)_n$, then $(J_n, M_n)_n$ is a Markov additive process (see Çinlar 1972) as it will appear in section 1.5.

Here, $(J_n)_n$ is positive recurrent but this does not imply the recurrence of $(S_n)_n$ or $(M_n)_n$. Moreover, $(S_n)_m$ and $(M_n)_n$ may have different behaviors. Explicit, necessary and sufficient conditions for the recurrence of $(M_n)_n$ in terms of characteristic functions and convergence of suitable series are given in Cénac *et al*. (2020, theorem 2. 1). The following proposition states a dichotomy between some recurrence versus transience phenomenon.

THEOREM 1.1.– Under non-nullness assumption, the following dichotomy holds:

i) the series $\sum_n \mathbb{P}(M_n = 0)$ diverges if and only if the process $(M_n)_n$ is recurrent in the following sense:

$$\exists r > 0, \ \mathbb{P}\left(\liminf_{n\to\infty} \|M_n\| < r\right) = 1.$$

ii) the senies $\sum_n \mathbb{P}(M_n = 0)$ converges if and only if the process $(M_n)_n$ is transient in the following sense:

$$\mathbb{P}\left(\lim_{n\to\infty} \|M_n\| = \infty\right) = 1.$$

Does the recurrence (respectively, the transience) of $(M_n)_n$ and $(S_n)_n$ occur at the same time? The answer to this 20-year-old question is no:

THEOREM 1.2 (Definitive invalidation of the conjecture in Mauldin *et al*. 1996).– There exist recurrent PRWs $(S_n)_n$ having an associated transient skeleton $(M_n)_n$.

Supposing that the persistence time distributions are horizontally and vertically symmetric is a natural necessary condition for the random walk $(S_n)_n$ to be recurrent. One example is given by the DRRW, originally introduced in Mauldin *et al*. (1996) (see Figure 1.7). Some particular values of the transition probabilities $q_{\alpha^n\beta}$ provide counterexamples. It is shown in Cénac *et al*. (2020) that the corresponding distributions of the persistence times must be non-integrable. In section 1.5, this non-integrability will be related to the non-existence of any invariant probability measure for the driving VLMC.

1.4. VLMC: existence of stationary probability measures

Consider a VLMC denoted by $U = (U_n)_{n \geq 0}$, defined by a pair (T, q) where T is a context tree on an alphabet A and $q = (q_c)_{c \in C}$ a family of probability measures on A, indexed by the contexts of T. A probability measure π on R is *stationary* or invariant (with regard to U) whenever π is the distribution of every U_n as soon as it is the distribution of U_0. The question of interest consists here of finding conditions on (T, q) for the process to admit at least one – or a unique one – stationary probability measure. The heuristic presentation aims to show how combinatoric objects, namely the α-*LIS* of contexts, and conditional probabilities, the *cascades*, naturally emerge.

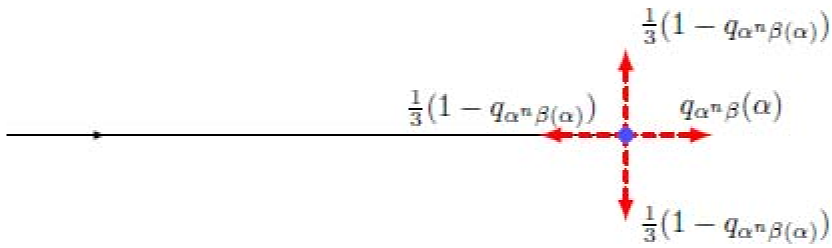

Figure 1.7. *The original directionally reinforced random walk (DRRW). For a color version of this figure, see www.iste.co.uk/barbu/data.zip*

Assume that π is a stationary probability measure on R:

– First step: finite words. Since R is endowed with the cylinder σ-algebra, π is determined by its values $\pi(wR)$ on the cylinders wR, where w runs over all finite words on A.

– Second step: longest internal suffixes of words. Assume that e is a finite non-internal word and take $a \in A$. Then, its `pref` is well defined and, because of formula [1.1], since π is stationary,

$$\pi\left(\alpha e\mathcal{R}\right)=q_{\mathrm{pref}(e)}(\alpha)\times\pi\left(e\mathcal{R}\right). \qquad [1.15]$$

Iterating this formula as far as possible leads to the following definitions. Consider any non-empty finite word w. It is uniquely decomposed as $w = p\alpha s = \beta_1\beta_2\beta_3 \cdots \beta_{\ell}\alpha s$, where the α and the β are letters and s is the *longest internal suffix* of w. The integer ℓ is non-negative and $p =$

$\beta_1\beta_2 \cdots \beta_\ell$ is a prefix of w that may be empty – in which case $\ell = 0$.

DEFINITION 1.5 (Lis and α-LIS).– *With these notations, the longest internal suffix s is shortened as the* lis *of w. The word αs is called the α*-LIS of w.

DEFINITION 1.6 (Cascade).– *With the notation above, the* cascade *of w is the product*

$$\mathrm{casc}(w) = q_{\mathrm{pref}(\beta_2\cdots\beta_\ell\alpha s)}(\beta_1)q_{\mathrm{pref}(\beta_3\cdots\beta_\ell\alpha s)}(\beta_2)\cdots q_{\mathrm{pref}(\alpha s)}(\beta_\ell). \quad [1.16]$$

Note that this definition makes sense because all the $\beta_k \cdots \beta_{\ell\alpha s}$ are non-internal words, $k \geq 2$. Moreover, if $w = \alpha s$ where s is internal, then $\ell = 0$ and casc $(w) = 1$.

With these definitions, iterating formula [1.15] leads to the following equality, known as the *cascade formula*: for every non-empty finite word w having αs as an α-LIS,

$$\pi(w\mathcal{R}) = \mathrm{casc}(w) \times \pi(\alpha s\mathcal{R}). \quad [1.17]$$

This shows that π is determined by its values on words of the form αs where s is internal and $\alpha \in A$.

– Third step: finite contexts. Assume that s is an internal word and that $\alpha \in A$. It is shown in Cénac *et al.* (2018a) that a stationary probability measure never charges infinite words so that, by disjoint union,

$$\pi(\alpha s\mathcal{R}) = \sum_{\substack{c:\text{ finite context}\\ c=s\cdots}} \pi(\alpha c\mathcal{R}) = \sum_{\substack{c:\text{ finite context}\\ c=s\cdots}} q_c(\alpha)\pi(c\mathcal{R}). \quad [1.18]$$

Note that the set of indices may be infinite but the family is summable because π is a finite measure. This shows that π is entirely determined by its values $\pi(cR)$ on the finite contexts.

List of Illustrations

List of Tables