

NOTE: RANDOM-TO-FRONT SHUFFLES ON TREES

ANDERS BJÖRNER¹

Royal Institute of Technology, Department of Mathematics, SE-100 44 Stockholm, Sweden
and

Institut Mittag-Leffler, Auravägen 17, SE-182 60 Djursholm, Sweden

email: bjonner@math.kth.se, bjonner@mittag-leffler.se

Submitted September 2, 2008, accepted in final form January 8, 2009

AMS 2000 Subject classification: 60J10; 60C05; 05E99

Keywords: Markov chain, shuffle, random-to-front, random walk, tree, semigroup

Abstract

A Markov chain is considered whose states are orderings of an underlying fixed tree and whose transitions are local “random-to-front” reorderings, driven by a probability distribution on subsets of the leaves. The eigenvalues of the transition matrix are determined using Brown’s theory of random walk on semigroups.

1 Introduction

The random-to-front shuffle of a linear list (known in the card-game model also as “inverse riffle shuffle”) is a well-known and much studied finite-state Markov chain. Its states are the linear orderings of an underlying finite set, and a step of the chain results from selecting a subset (often a singleton) and moving it to the front of the current list in the induced order. See e.g. [2, 5, 7] and the references given there. In this note we consider a slight generalization, namely to shuffles on trees.

Consider a fixed rooted tree T whose leaves L are all at the same depth. The following shows a such a tree of depth 3.




Figure 1.

¹ RESEARCH SUPPORTED BY THE KNUT AND ALICE WALLENBERG FOUNDATION, GRANT KAW.2005.0098.

Suppose that at each inner node (i.e., node that is not a leaf) a total ordering of its children is given. For instance, it can be the left-to-right ordering given by a planar drawing of the tree, such as in Figure 1. Now, a subset E of the set of leaves L is chosen with some probability. Then the ordering is rearranged locally at each inner node so that the children having some descendant in E come first, and otherwise the induced order is preserved. The process is illustrated in Figures 3 and 4.

In this note the eigenvalues of the transition matrix of this Markov chain are determined. This is a straight-forward application of Brown’s theory of random walks on semigroups [4].

Note that if $\text{depth}(T) = 1$ the Markov chain we describe amounts to the classical linear random-to-front shuffle. For $\text{depth}(T) > 1$ we perform such a linear shuffle locally at each inner node, in each case moving the set of E -related nodes to the front.




Figure 2.

If $\text{depth}(T) = 2$ we obtain the "library with several shelves" model considered in [3], as indicated in Figure 2. This case was derived in [3] via geometric considerations, ultimately relying on Brown’s theory of random walks on semigroups. If one cares only about the library result, and not about random walks on complex hyperplane arrangements, there is of course no need to mix in geometric considerations. This note can be seen as a self-contained appendix to [3] whose modest purpose is to fill in the details on how to obtain the general dynamic library model in the simplest and most direct way, avoiding geometry.

Another “tree analogue” of the classical linear random-to-front shuffle, different from the one considered here, has been studied in the literature. This is the random-to-root shuffle on binary trees, see e.g. [1, 6].

2 Shuffles on trees

We begin by establishing notation. For any finite set A , let

$$\begin{aligned} S(A) &\stackrel{\text{def}}{=} \{\text{linear orderings of } A\} \\ \Pi(A) &\stackrel{\text{def}}{=} \{\text{partitions of } A\} \\ \Pi^{\text{ord}}(A) &\stackrel{\text{def}}{=} \{\text{ordered partitions of } A\} \end{aligned}$$

The sets $\Pi(A)$ and $\Pi^{\text{ord}}(A)$ are partially ordered by refinement, meaning that $\alpha \leq \beta$ if and only if every block of the partition (or ordered partition) α is a union of blocks from β . Direct products (of sets, posets, ...) are denoted by \otimes .

We consider rooted trees T that are *pure*, meaning that all leaves are at the same depth d . Let V_j denote the set of nodes at depth j . So, $V_0 = \{\text{root}\}$, $I \stackrel{\text{def}}{=} \cup_{j=0}^{d-1} V_j = \{\text{inner nodes}\}$, and $L \stackrel{\text{def}}{=} V_d = \{\text{leaves}\}$.

Definition 2.1. Let $E \subseteq L$. A node $x \in T$ is E -related if some descendant of x belongs to E .

For each inner node $x \in I$, let C_x denote the set of its children.

Definition 2.2. A local ordering of T is a choice of linear order for the set of children C_x at each inner node $x \in I$. Denote by $\mathcal{O}(T)$ the set of all local orderings of T . Thus, $\mathcal{O}(T) \cong \otimes_{x \in I} S(C_x)$.

The subsets of L act on $\mathcal{O}(T)$ in the following way.

Definition 2.3. Let $\pi = (\pi_x)_{x \in I}$ be a local ordering, and let $E \subseteq 2^L$. Then $E(\pi) = (E_x(\pi_x))_{x \in I}$, where $E_x(\pi_x)$ is the linear ordering of C_x in which the E -related elements come first, in the order induced by π_x , followed by the remaining elements, also in the induced order.

The following figure shows a local ordering π of a tree T , which coincides with left-to-right order in the planar drawing of T .




Figure 3.

(iv) These are all the eigenvalues of P_w .

For clarity's sake, let us point out that $\varepsilon_\alpha = \varepsilon_\beta$, for $\alpha \neq \beta$, and $\varepsilon_\alpha = 0$ are possible.

Proof. As mentioned in the introduction, this is a special case of Brown's theory for walks on semigroups [4], with which we now assume familiarity.

Let $\text{Part}^{\text{ord}}(T) \stackrel{\text{def}}{=} \bigotimes_{x \in I} \Pi^{\text{ord}}(C_x)$. So, an element $\beta \in \text{Part}^{\text{ord}}(T)$ is a choice of *ordered* partition β_x of the set of children of x , for each inner node x . In particular, for each subset $E \subseteq L$ there is an element $\beta^E \in \text{Part}^{\text{ord}}(T)$ whose component β_x^E at $x \in I$ is the two-block ordered partition of C_x whose first block consists of the E -related elements of C_x , and second block of the remainder. (If one of these blocks is empty we forget about it and let β_x^E have only one block.)

Now, introduce the following probability distribution on $\text{Part}^{\text{ord}}(T)$:

$$\text{Prob}(\beta) = \begin{cases} w_E, & \text{if } \beta = \beta^E, E \subseteq L \\ 0, & \text{for all other ordered partitions.} \end{cases} \quad (2.1)$$

Given this set-up, the proof consists of verifying each of the following claims for $\text{Part}^{\text{ord}}(T)$, and then referring to [4].

1. $\text{Part}^{\text{ord}}(T)$ is an LRB (left regular band) semigroup with component-wise composition. The composition in each factor $\Pi^{\text{ord}}(A)$ has the following description. If $X = \langle X_1, \dots, X_p \rangle$ and $Y = \langle Y_1, \dots, Y_q \rangle$ are ordered partitions of A , then $X \circ Y = \langle X_i \cap Y_j \rangle$ with the blocks ordered by the lexicographic order of the pairs of indices (i, j) .

2. Its support lattice is $\text{Part}(T)$ and support map

$$\text{supp} : \text{Part}^{\text{ord}}(T) \rightarrow \text{Part}(T),$$

whose component at each $x \in I$ is the map $\Pi^{\text{ord}}(C_x) \rightarrow \Pi(C_x)$ that sends an ordered partition of C_x to an unordered partition by forgetting the ordering of its blocks.

3. The maximal elements of $\text{Part}^{\text{ord}}(T)$ are the local orderings $\mathcal{O}(T)$.
4. The steps of the semigroup random walk on $\mathcal{O}(T)$, induced as in [4] by the probability assignment (2.1), are precisely the steps described in Definition 2.4.
5. For each $E \subseteq L$ and $\alpha \in \text{Part}(T)$:

$$\text{supp}(\beta^E) \leq \alpha \iff E \text{ is } \alpha\text{-compatible.}$$

6. The number of maximal elements of $\text{Part}^{\text{ord}}(T)$ above some $\beta \in \text{Part}^{\text{ord}}(T)$ is by Zaslavsky's theorem the sum of Möbius function absolute values

$$\sum_{\alpha \geq \text{supp}(\beta)} |\mu(\alpha, \hat{1})|$$

computed on the product partition lattice $\text{Part}(T)$. From this follows, via Brown's theory [4], that

$$m_\alpha = |\mu(\alpha, \hat{1})|,$$

for all $\alpha \in \text{Part}(T)$. By the product property of the Möbius function and its well-known explicit evaluation on the partition lattice (see [8]), this quantity equals

$$|\mu(\alpha, \hat{1})| = \prod_{x \in I} \prod_{B \in \alpha_x} (|B| - 1)!$$

In view of these facts the theorem is obtained by specializing Theorem 1 on page 880 of [4] to the semigroup $\text{Part}^{\text{ord}}(T)$. \square

3 Remarks

3.1. The random walk of Theorem 2.6 has a unique stationary distribution π if and only if $\{E \in 2^L : w_E > 0\}$ is *separating*, meaning that for every inner node $x \in I$ and every pair of siblings $y, z \in C_x$, $y \neq z$, there is a subset $E \subseteq L$ with $w_E > 0$ for which one of y and z is E -related and the other is not.

This follows from Theorem 2 of Brown and Diaconis [5], using the fact that the random walk we consider can be realized as a walk on the complement of a product of real braid arrangements. Theorem 2 of [5] also gives additional information about the stationary distribution.

3.2. One easily checks that the subset $\{\beta^E : E \subseteq L\}$ generates the full semigroup $\text{Part}^{\text{ord}}(T)$, and that the set of its maximal elements $\mathcal{O}(T)$ is generated by $\{\beta^{\{e\}} : e \in L\}$.

3.3. Suppose that $w_E \neq 0$ only if $|E| = 1$. Then Theorem 2.6 implies that the eigenvalues are indexed by $\bigotimes_{x \in I} 2^{C_x}$, and that their multiplicities are products of derangement numbers, thus generalizing the well-known result of Donnelly, Kapoor-Reingold and Phatarfod for the Tsetlin library (the $\text{depth}(T) = 1$ case); see the references for this given in [2, 4, 5].

References

- [1] B. Allen and I. Munro, *Self-organizing binary search trees*, J. Assoc. Comput. Mach. **25** (1978), 526–535. MR0508699
- [2] P. Bidigare, P. Hanlon and D. Rockmore, *A combinatorial description of the spectrum for the Tsetlin library and its generalization to hyperplane arrangements*, Duke Math J. **99** (1999), 135–174. MR1700744
- [3] A. Björner, *Random walks, arrangements, cell complexes, greedoids, and self-organizing libraries*, in “Building Bridges” (eds. M. Grötschel and G. O. H. Katona), Bolyai Soc. Math. Studies **19** (2008), Springer (Berlin) and Janos Bolyai Math. Soc. (Budapest), pp.165–203.
- [4] K. S. Brown, *Semigroups, rings and Markov chains*, J. Theor. Probab. **13** (2000), 871–938. MR1785534
- [5] K. S. Brown and P. Diaconis, *Random walks and hyperplane arrangements*, Ann. Probab. **26** (1998), 1813–1854. MR1675083
- [6] R. P. Dobrow and J. A. Fill, *On the Markov chain for the move-to-root rule for binary search trees*, Ann. Applied Probab. **5** (1995), 1–19. MR1325037
- [7] J. A. Fill and L. Holst, *On the distribution of search cost for the move-to-front rule*, Random Structures and Algorithms **8** (1996), 179–186. MR1603279
- [8] R. P. Stanley, *Enumerative Combinatorics, Vol. 1*, Cambridge Univ. Press, 1997. MR1442260