

UC Berkeley

UC Berkeley Previously Published Works

Title

Radix sort trees in the large

Permalink

<https://escholarship.org/uc/item/89r699t0>

Journal

Electronic Communications in Probability, 22(none)

ISSN

1083-589X

Authors

Evans, Steven N
Wakolbinger, Anton

Publication Date

2017

DOI

10.1214/17-ecp77

Peer reviewed

Radix sort trees in the large *

Steven N. Evans[†] Anton Wakolbinger[‡]

Abstract

The trie-based radix sort algorithm stores pairwise different infinite binary strings in the leaves of a binary tree in a way that the Ulam-Harris coding of each leaf equals a prefix (that is, an initial segment) of the corresponding string, with the prefixes being of minimal length so that they are pairwise different. We investigate the *radix sort tree chains* – the tree-valued Markov chains that arise when successively storing the finite collections of random infinite binary strings Z_1, \dots, Z_n , $n = 1, 2, \dots$ according to the trie-based radix sort algorithm, where the source strings Z_1, Z_2, \dots are independent and identically distributed. We establish a bijective correspondence between the full Doob–Martin boundary of the radix sort tree chain with a *symmetric Bernoulli source* (that is, each Z_k is a fair coin-tossing sequence) and the family of radix sort tree chains for which the common distribution of the Z_k is a diffuse probability measure on $\{0, 1\}^\infty$. In essence, our result characterizes all the ways that it is possible to condition such a chain of radix sort trees consistently on its behavior “in the large”.

Keywords: binary tree; tail σ -field; Doob–Martin kernel; harmonic function; bridge; exchangeability.

AMS MSC 2010: Primary 60J50, Secondary 60J10; 68W40.

Submitted to ECP on October 6, 2016, final version accepted on August 9, 2017.

1 Introduction

Various sorting algorithms proceed by storing the data in the leaves of a tree. If the data are infinite *binary strings* $z_1, \dots, z_n \in \{0, 1\}^\infty$, then a natural choice for the tree is the rooted binary tree with n leaves chosen such that the Ulam-Harris coding of each of the leaves coincides with a finite initial segment (otherwise called a prefix or left factor) of one of the z_j , and such that these initial segments are pairwise different and have minimal length (see below for a fuller description). This data structure is the basis of the *Radix Sort* algorithm. The tree $\mathbf{R}(z_1, \dots, z_n)$ in whose leaves the n strings are stored is sometimes called a *trie*, alluding to the word retrieval.

When the n strings are random, drawn i.i.d. from a diffuse probability distribution ν on $\{0, 1\}^\infty$, then this construction gives rise to a random tree ${}^\nu R_n := \mathbf{R}(Z_1, \dots, Z_n)$. In order to obtain a probabilistic analysis of the Radix Sort algorithm, asymptotic properties of these random trees as $n \rightarrow \infty$ have been considered for the *symmetric Bernoulli* or *unbiased memoryless source model*, where ν is the fair coin tossing measure, e.g. in [11] ch. 5 and [9] §5.2.2., and for more general inputs of random strings in [15]. The *density model*, where ν is the image under the binary expansion of an absolutely continuous

*SNE supported in part by NSF grant DMS-0907630, NSF grant DMS-1512933, and NIH grant 1R01GM109454-01. AW supported in part by DFG priority program 1590.

[†]University of California at Berkeley, United States of America. E-mail: evans@stat.berkeley.edu

[‡]Goethe-Universität, Germany. E-mail: wakolbinger@math.uni-frankfurt.de

probability measure on $[0, 1]$, was considered in [2]. *Dynamical sources* appear in [1]; these include *Markovian inputs*, where ν is the shift-invariant distribution of a Markov chain, see [14], [10].

In this paper we analyze the tree-valued Markov chains $({}^\nu R_n)_{n \in \mathbb{N}}$ from a more synoptic point of view. We show that any such chain is a harmonic transform of the Markov chain $({}^\gamma R_n)_{n \in \mathbb{N}}$, with γ the fair coin-tossing measure, and we prove that the family $({}^\nu R_n)_{n \in \mathbb{N}}$ as ν varies constitute the full Doob–Martin boundary of $({}^\gamma R_n)_{n \in \mathbb{N}}$. Loosely speaking, this means that all consistent ways of conditioning a chain of radix sort trees “in the large” are described by precisely the family $({}^\nu R_n)_{n \in \mathbb{N}}$.

In order to state our main result more formally, we first fix some notation. Denote by $\{0, 1\}^* := \bigsqcup_{k=0}^\infty \{0, 1\}^k$ the set of finite tuples or *words* drawn from the alphabet $\{0, 1\}$ (with the empty word \emptyset allowed) – the symbol \bigsqcup emphasizes that this is a disjoint union. Write an ℓ -tuple $v = (v_1, \dots, v_\ell) \in \{0, 1\}^*$ more simply as $v_1 \dots v_\ell$ and set $|v| = \ell$. Define a directed graph with vertex set $\{0, 1\}^*$ by declaring that if $u = u_1 \dots u_k$ and $v = v_1 \dots v_\ell$ are two words, then (u, v) is a directed edge (that is, $u \rightarrow v$) if and only if $\ell = k + 1$ and $u_i = v_i$ for $i = 1, \dots, k$. Call this directed graph the *complete rooted binary tree*. Say that $u < v$ for two words $u = u_1 \dots u_k$ and $v = v_1 \dots v_\ell$ in $\{0, 1\}^*$ if $k < \ell$ and $u_1 \dots u_k = v_1 \dots v_k$; that is, $u < v$ if there exist words $w_0, w_1, \dots, w_{\ell-k}$ with $u = w_0 \rightarrow w_1 \rightarrow \dots \rightarrow w_{\ell-k} = v$. This partial order extends to $\{0, 1\}^* \sqcup \{0, 1\}^\infty$ in the obvious way: if $u \in \{0, 1\}^*$ and $v \in \{0, 1\}^\infty$, then $u < v$ when $u = u_1 \dots u_k$ and $v = v_1 v_2 \dots$ with $u_1 \dots u_k = v_1 \dots v_k$ (and no two elements of $\{0, 1\}^\infty$ are comparable). It will be convenient to introduce the notation $\tau(y) := \{z \in \{0, 1\}^\infty : y < z\}$ for $y \in \{0, 1\}^*$.

A *finite rooted binary tree* is a non-empty subset \mathbf{t} of $\{0, 1\}^*$ with the property that if $v \in \mathbf{t}$ and $u \in \{0, 1\}^*$ is such that $u \rightarrow v$, then $u \in \mathbf{t}$. The vertex \emptyset (that is, the empty word) belongs to any such tree \mathbf{t} and is the *root* of \mathbf{t} . The *leaves* of \mathbf{t} are the elements $v \in \mathbf{t}$ such that if $v \rightarrow w$, then $w \notin \mathbf{t}$, and we use the notation $\mathbf{L}(\mathbf{t})$ for the leaves of \mathbf{t} . A finite rooted binary tree is uniquely determined by its leaves: it is the smallest rooted binary tree that contains the set of leaves and it consists of the leaves and the points $u \in \{0, 1\}^*$ such that $u < v$ for some leaf v . In general, write

$$\mathbf{T}(y_1, \dots, y_m) := \bigcup_{j=1}^m \{u \in \{0, 1\}^* : u \leq y_j\}$$

for the smallest finite rooted binary tree containing $y_1, \dots, y_m \in \{0, 1\}^*$; the leaves of this tree form a subset of $\{y_1, \dots, y_m\}$ and this subset is proper if and only if $y_i < y_j$ for some pair $1 \leq i \neq j \leq m$.

A collection z_1, \dots, z_n of distinct elements of $\{0, 1\}^\infty$ determines a finite rooted binary tree in the following manner. For $n = 1$, put $H_{1,1}(z_1) := 0$ and $\zeta_{1,1}(z_1) := \emptyset$. For $n \geq 2$ and $1 \leq j \leq n$, let

$$H_{n,j}(z_1, \dots, z_n) := \min\{\ell : (z_{j,1}, \dots, z_{j,\ell}) \neq (z_{k,1}, \dots, z_{k,\ell}), k \neq j\}$$

be the minimal length at which a prefix of z_j differs from the prefixes of the same length of all the other z_k , $k \neq j$, and denote the corresponding prefix by

$$\zeta_{n,j}(z_1, \dots, z_n) := (z_{j,1}, \dots, z_{j,H_{n,j}(z_1, \dots, z_n)}) \in \{0, 1\}^*, \quad 1 \leq j \leq n. \quad (1.1)$$

The words $\zeta_{n,j}(z_1, \dots, z_n)$, $1 \leq j \leq n$, are distinct and $\zeta_{n,j}(z_1, \dots, z_n) < z_j$ for $1 \leq j \leq n$. Note that if σ is a permutation of $[n] := \{1, \dots, n\}$, then

$$\zeta_{n,\sigma(j)}(z_{\sigma(1)}, \dots, z_{\sigma(n)}) = \zeta_{n,j}(z_1, \dots, z_n). \quad (1.2)$$

The *radix sort tree determined by the input* z_1, \dots, z_n is defined as

$$\mathbf{R}(z_1, \dots, z_n) := \mathbf{T}(\zeta_{n,1}(z_1, \dots, z_n), \dots, \zeta_{n,n}(z_1, \dots, z_n)).$$

Thus, $\mathbf{R}(z_1, \dots, z_n)$ is the finite rooted binary tree whose n leaves are coded by the n finite strings of (1.1). Observe that

$$\mathbf{R}(z_1, \dots, z_n) = \mathbf{R}(z_{\sigma(1)}, \dots, z_{\sigma(n)}) \tag{1.3}$$

for any permutation σ of $[n]$.

Let Z_1, Z_2, \dots be i.i.d. $\{0, 1\}^\infty$ -valued random variables with common distribution some diffuse probability measure ν . Then Z_1, Z_2, \dots are a.s. pairwise distinct, and on this event we set ${}^\nu R_n := \mathbf{R}(Z_1, \dots, Z_n)$. When ν is fair coin-tossing measure γ (that is, γ is the infinite product of the uniform measure on $\{0, 1\}$), we drop the ν and simply write R_n for ${}^\gamma R_n$. It is not hard to see that $({}^\nu R_n)_{n \in \mathbb{N}}$ is a Markov chain; we call it a *radix sort tree chain*.

Note for $y \in \{0, 1\}^*$ and $n \geq k \geq 2$ that with probability one

$$\#\{1 \leq j \leq n : y \leq \zeta_{n,j}(Z_1, \dots, Z_n)\} = k$$

if and only if

$$\#\{1 \leq j \leq n : Z_j \in \tau(y)\} = k,$$

Thus,

$$\nu(\tau(y)) = \lim_{n \rightarrow \infty} \frac{1}{n} \#\{1 \leq j \leq n : y \leq \zeta_{n,j}(Z_1, \dots, Z_n)\} \quad \mathbb{P} - \text{a.s.}$$

and ν can be recovered almost surely from the tail σ -field of $({}^\nu R_n)_{n \in \mathbb{N}}$; in particular, different choices of ν result in different distributions for $({}^\nu R_n)_{n \in \mathbb{N}}$. It follows from (1.3) and the Hewitt–Savage zero–one law that the tail σ -field of $({}^\nu R_n)_{n \in \mathbb{N}}$ is \mathbb{P} -a.s. trivial.

In order to describe our results, we need to use some notions and facts from Doob–Martin boundary theory. A quick summary tailored to the sort of setting we are in of a process which “goes off to infinity” and never revisits states may be found in [4, 5], where there are also references to expositions of the general theory for arbitrary transient Markov chains following on from the seminal paper [3]. Analyses of binary-search-tree and digital-search-tree chains from the Doob–Martin point of view are presented in [4].

Let \mathbb{S}_n be the set of trees that can arise as $\mathbf{R}(z_1, \dots, z_n)$ for some choice of z_1, \dots, z_n and set $\mathbb{S} = \bigsqcup_{n \in \mathbb{N}} \mathbb{S}_n$. Of course, $\mathbb{S}_1 = \{\emptyset\}$. For $n \geq 2$, a finite rooted binary tree \mathbf{t} with n leaves belongs to \mathbb{S}_n if and only if whenever $u_1 u_2 \dots u_{m-1} u_m \in \mathbf{L}(\mathbf{t})$, then $u_1 u_2 \dots u_{m-1} \bar{u}_m \in \mathbf{t}$, where $\bar{0} := 1$ and $\bar{1} := 0$.

Given a binary tree $\mathbf{t} \in \mathbb{S}$ with $M(\mathbf{t})$ leaves (that is, $\mathbf{t} \in \mathbb{S}_{M(\mathbf{t})}$), write $R_1^{\mathbf{t}}, R_2^{\mathbf{t}}, \dots, R_{M(\mathbf{t})}^{\mathbf{t}}$ for the *bridge process* obtained by conditioning $R_1, \dots, R_{M(\mathbf{t})}$ on the event $\{R_{M(\mathbf{t})} = \mathbf{t}\}$. This Markov chain has the same backward transition probabilities as $(R_n)_{n \in \mathbb{N}}$; that is,

$$\mathbb{P}\{R_n^{\mathbf{t}} = \mathbf{r} \mid R_{n+1}^{\mathbf{t}} = \mathbf{s}\} = \mathbb{P}\{R_n = \mathbf{r} \mid R_{n+1} = \mathbf{s}\}$$

for $n + 1 \leq M(\mathbf{t})$.

An *infinite bridge* for $(R_n)_{n \in \mathbb{N}}$ is a Markov chain $(R_n^\infty)_{n \in \mathbb{N}}$ with $R_n^\infty \in \mathbb{S}_n$ for $n \in \mathbb{N}$ and the same backward transition probabilities as $(R_n)_{n \in \mathbb{N}}$. We show in Sec. 4 that each chain $({}^\nu R_n)_{n \in \mathbb{N}}$ is an infinite bridge for $(R_n)_{n \in \mathbb{N}}$. Any infinite bridge is a Doob h -transform of $(R_n)_{n \in \mathbb{N}}$; that is, it has forward transition probabilities of the form

$$\mathbb{P}\{R_{n+1}^\infty = \mathbf{t} \mid R_n^\infty = \mathbf{s}\} = h(\mathbf{s})^{-1} \mathbb{P}\{R_{n+1} = \mathbf{t} \mid R_n = \mathbf{s}\} h(\mathbf{t}),$$

where the nonnegative function h is given up to a constant multiple by

$$h(\mathbf{t}) = \frac{\mathbb{P}\{R_n^\infty = \mathbf{t}\}}{\mathbb{P}\{R_n = \mathbf{t}\}}.$$

The function h is harmonic for $(R_n)_{n \in \mathbb{N}}$; that is,

$$\sum_{\mathbf{t}} \mathbb{P}\{R_{n+1} = \mathbf{t} \mid R_n = \mathbf{s}\} h(\mathbf{t}) = h(\mathbf{s}).$$

Conversely, any Markov chain with initial state the trivial tree \emptyset and transition probabilities that arise from those of $(R_n)_{n \in \mathbb{N}}$ through the h -transform construction for some nonnegative harmonic function h (normalized, without loss of generality, so that $h(\emptyset) = 1$) is an infinite bridge.

The distribution of an infinite bridge is a mixture of distributions of infinite bridges with almost surely trivial tail σ -fields. Equivalently, the collection of nonnegative harmonic functions h with $h(\emptyset) = 1$ is a compact convex set (for the product topology on $\mathbb{R}_+^{\mathbb{S}}$) and any such function is a unique convex combination of the extreme points of this set. In particular, there is a bijective correspondence between the extreme points of these two sets; that is between the set of infinite bridges with trivial tail σ -fields and extremal normalized nonnegative harmonic functions.

One way to construct infinite bridges is to look for sequences $(\mathbf{t}_k)_{k \in \mathbb{N}}$ with $M(\mathbf{t}_k) \rightarrow \infty$ as $k \rightarrow \infty$ such that initial segments of the finite bridges $(R_1^{\mathbf{t}_k}, R_2^{\mathbf{t}_k}, \dots, R_{M(\mathbf{t}_k)}^{\mathbf{t}_k})$ converge in distribution as $k \rightarrow \infty$. A necessary condition for an infinite bridge to have an almost surely trivial tail σ -field is that it arises from such a construction.

The nonnegative harmonic function corresponding to an infinite bridge constructed in this way (normalized to have $h(\emptyset) = 1$) is

$$h(\mathbf{s}) = \lim_{k \rightarrow \infty} K(\mathbf{s}, \mathbf{t}_k), \tag{1.4}$$

where

$$K(\mathbf{s}, \mathbf{t}) := \frac{\mathbb{P}\{R_{M(\mathbf{s})}^{\mathbf{t}} = \mathbf{s}\}}{\mathbb{P}\{R_{M(\mathbf{s})} = \mathbf{s}\}} = \frac{\mathbb{P}\{R_{M(\mathbf{t})} = \mathbf{t} \mid R_{M(\mathbf{s})} = \mathbf{s}\}}{\mathbb{P}\{R_{M(\mathbf{t})} = \mathbf{t}\}} \tag{1.5}$$

is the *Doob–Martin kernel*. A necessary condition for a normalized nonnegative harmonic function to be an extreme point is that it arises as such a limit.

The following is our main result characterizing all the ways that it is possible to condition the radix sort tree chain with inputs distributed according to fair coin-tossing measure. We prove this result in Section 6.

Theorem 1.1. *An infinite bridge for the radix sort tree chain with inputs distributed according to fair coin-tossing measure on $\{0, 1\}^\infty$ has an almost surely trivial tail σ -field if and only if it is a Markov chain with the same distribution as the radix sort tree chain with inputs distributed according to some diffuse probability measure on $\{0, 1\}^\infty$. Consequently, the distribution of an infinite bridge for the radix sort tree chain with inputs distributed according to fair coin-tossing measure is a unique mixture of distributions of radix sort tree chains with inputs distributed according to diffuse probability measures on $\{0, 1\}^\infty$. Moreover, an infinite bridge $(R_n^\infty)_{n \in \mathbb{N}}$ has an almost surely trivial tail σ -field if and only if there is a sequence $(\mathbf{t}_k)_{k \in \mathbb{N}}$ with $M(\mathbf{t}_k) \rightarrow \infty$ as $k \rightarrow \infty$ such that for all $n \in \mathbb{N}$ the initial segment $(R_1^{\mathbf{t}_k}, \dots, R_n^{\mathbf{t}_k})$ converges in distribution to $(R_1^\infty, \dots, R_n^\infty)$ as $k \rightarrow \infty$.*

We mention in passing that ideas similar to those in this paper apply to another Markov chain that is closely related to the radix sort tree chain. We describe this chain and make some comments about the counterpart of Theorem 1.1 for it in the next few paragraphs.

The out-degree 1 vertices in the radix sort tree $\mathbf{R}(z_1, \dots, z_n)$ for the input sequence z_1, \dots, z_n don't provide any information that is useful for the purpose of sorting z_1, \dots, z_n into lexicographical order. The associated *PATRICIA* tree $\bar{\mathbf{R}}(z_1, \dots, z_n) := \Phi \circ \mathbf{R}(z_1, \dots, z_n)$,

where Φ is the transformation that takes a binary tree, removes the out-degree 1 vertices, and “closes up the gaps,” doesn’t lose any information that is relevant for sorting the inputs (PATRICIA is an acronym for “Practical Algorithm To Retrieve Information Coded In Alphanumeric”). PATRICIA trees were invented independently in [12, 6]. Features of PATRICIA trees for random inputs were first studied in [8] and this topic has since been the subject of quite a large literature.

Given i.i.d. $\{0, 1\}^\infty$ -valued random variables Z_1, Z_2, \dots with common distribution some diffuse probability measure ν , set ${}^\nu\bar{R}_n := \bar{\mathbf{R}}(Z_1, \dots, Z_n)$. When ν is the fair coin-tossing measure γ , drop the ν and simply write R_n for ${}^\gamma\bar{R}_n$. As the composition of a Markov chain with a non-injective function, the process $({}^\nu\bar{R}_n)_{n \in \mathbb{N}}$ is not a *priori* Markovian. Indeed, it is not hard to see that Dynkin’s classical criterion for the composition of a Markov chain with a function to be Markovian (see, for example, Equation (8) of [13]) does not hold. However, it follows from the description of the backward transition probabilities of $({}^\nu R_n)_{n \in \mathbb{N}}$ in Section 3 that Dynkin’s criterion does apply to show that the time-reversal of $({}^\nu\bar{R}_n)_{n \in \mathbb{N}}$ is Markovian, so that $({}^\nu\bar{R}_n)_{n \in \mathbb{N}}$ is itself Markovian.

We remarked above that ν can be recovered with probability one from the sample paths of $({}^\nu R_n)_{n \in \mathbb{N}}$ and hence different choices of ν result in different distributions for $({}^\nu R_n)_{n \in \mathbb{N}}$. The latter is not true for the PATRICIA chain: as a trivial example, if $\nu' = \gamma$ and $\nu'' = \delta_0 \otimes \gamma$, then $({}^{\nu'}\bar{R}_n)_{n \in \mathbb{N}}$ and $({}^{\nu''}\bar{R}_n)_{n \in \mathbb{N}}$ have the same distribution. This observations suggest the definition of an equivalence relation on diffuse probability measures on $\{0, 1\}^\infty$, where we regard two probability measures ν' and ν'' to be equivalent if $({}^{\nu'} R_n)_{n \in \mathbb{N}}$ and $({}^{\nu''} R_n)_{n \in \mathbb{N}}$ have the same distribution – it is not difficult to characterize this equivalence relation more simply and concretely, but we don’t do so here.

If we correct for this lack of uniqueness, then it is natural to conjecture that a close analogue of Theorem 1.1 holds: namely, the extremal infinite bridges for the PATRICIA chain built from i.i.d. fair coin-tossing measure distributed inputs are exactly the chains of the form $({}^\nu\bar{R}_n)_{n \in \mathbb{N}}$ as ν ranges over some complete collection of equivalence class representatives. The family of extremal infinite bridges for the PATRICIA chain is actually much richer than this, as we note in Remark 6.2. We will present the details in a forthcoming paper.

The structure of the remainder of the paper is as follows. In Sections 2, 3, and 4 we obtain the forward transition probabilities, backward transition probabilities, and Doob–Martin kernels of the radix sort tree chains. In Section 4 we also show that each radix sort tree chain $({}^\nu R_n)_{n \in \mathbb{N}}$ is a Doob h -transform of the Markov chain $(R_n)_{n \in \mathbb{N}}$. We consider infinite bridges for the Markov chain $(R_n)_{n \in \mathbb{N}}$ in Section 5 and introduce an auxiliary consistent labeling of the leaves of the state of the bridge at each time n by $[n] := \{1, \dots, n\}$ such that, intuitively, these labelings determine a labeling of the limit of the bridge at time ∞ and the whole bridge path can be recovered from the limit and its labeling. We prove two results, Theorem 6.1 and Corollary 6.3, in Section 6 that together establish Theorem 1.1.

2 Forward transition probabilities

Recall that \mathbb{S}_n is the set of trees that can arise as $\mathbf{R}(z_1, \dots, z_n)$ for some choice of distinct $z_1, \dots, z_n \in \{0, 1\}^\infty$. It is clear that $\mathbf{R}(z_1, \dots, z_n)$ is the unique finite rooted binary tree $\mathbf{t} \in \mathbb{S}_n$ with the following property: if $\mathbf{L}(\mathbf{t}) = \{y_1, \dots, y_n\}$, then there is a permutation π of $[n]$ such that $z_i \in \tau(y_{\pi(i)})$ for $i \in [n]$.

For $n \in \mathbb{N}$, the distribution of ${}^\nu R_n$ is specified by

$$\mathbb{P}\{{}^\nu R_1 = \emptyset\} = 1$$

and, for $n \geq 2$ and $\mathbf{t} \in \mathbb{S}_n$ with $\{y_1, \dots, y_n\} = \mathbf{L}(\mathbf{t})$,

$$\begin{aligned} \mathbb{P}\{\nu R_n = \mathbf{t}\} &= \mathbb{P}\{\{\zeta_{n,1}(Z_1, \dots, Z_n), \dots, \zeta_{n,n}(Z_1, \dots, Z_n)\} = \{y_1, \dots, y_n\}\} \\ &= n! \prod_{k=1}^n \nu(\tau(y_k)). \end{aligned} \tag{2.1}$$

In particular,

$$\mathbb{P}\{R_n = \mathbf{t}\} = n! \prod_{k=1}^n \gamma(\tau(y_k)) = n! \prod_{k=1}^n 2^{-|y_k|}. \tag{2.2}$$

The radix sort chain $(\nu R_n)_{n \in \mathbb{N}}$ has the following forward transition dynamics. Consider $\mathbf{s} \in \mathbb{S}_n$. There are two classes of trees $\mathbf{t} \in \mathbb{S}_{n+1}$ such that

$$\mathbb{P}\{\nu R_{n+1} = \mathbf{t} \mid \nu R_n = \mathbf{s}\} > 0.$$

Case I. Here $\mathbf{t} \in \mathbb{S}_{n+1}$ is a tree with $\mathbf{L}(\mathbf{t}) = \mathbf{L}(\mathbf{s}) \sqcup \{w\}$, where $w = x\bar{u}_m$ for some $x = u_1 u_2 \dots u_{m-1}$ with $xu_m \in \mathbf{s} \setminus \mathbf{L}(\mathbf{s})$. In this case,

$$\mathbb{P}\{\nu R_{n+1} = \mathbf{t} \mid \nu R_n = \mathbf{s}\} = \nu(\tau(w)). \tag{2.3}$$

In particular,

$$p(\mathbf{s}, \mathbf{t}) := \mathbb{P}\{R_{n+1} = \mathbf{t} \mid R_n = \mathbf{s}\} = 2^{-|w|} = 2^{-(|x|+1)}. \tag{2.4}$$

Case II. Here $\mathbf{t} \in \mathbb{S}_{n+1}$ is a tree with $\mathbf{L}(\mathbf{t}) = (\mathbf{L}(\mathbf{s}) \setminus \{y\}) \sqcup \{y', y''\}$, where $y = u_1 u_2 \dots u_{m-1} u_m \in \mathbf{L}(\mathbf{s})$, $y' = u_1 u_2 \dots u_{m-1} u_m v_1 \dots v_p$ and $y'' = u_1 u_2 \dots u_{m-1} u_m \bar{v}_1 \dots \bar{v}_p$ for some $p \geq 1$ and $v_1, \dots, v_p \in \{0, 1\}$. In this case,

$$\begin{aligned} \mathbb{P}\{\nu R_{n+1} = \mathbf{t} \mid \nu R_n = \mathbf{s}\} &= \nu(\tau(y')) \frac{\nu(\tau(y''))}{\nu(\tau(y))} + \nu(\tau(y'')) \frac{\nu(\tau(y'))}{\nu(\tau(y))} \\ &= 2 \frac{\nu(\tau(y')) \nu(\tau(y''))}{\nu(\tau(y))}. \end{aligned} \tag{2.5}$$

In particular,

$$p(\mathbf{s}, \mathbf{t}) := \mathbb{P}\{R_{n+1} = \mathbf{t} \mid R_n = \mathbf{s}\} = 2 \frac{2^{-|y'|} 2^{-|y''|}}{2^{-|y|}}. \tag{2.6}$$

For later use we note that, with $\mathbf{d} := \mathbf{T}(v_1 \dots v_{p-1} v_p, v_1 \dots v_{p-1} \bar{v}_p)$, this may be written as

$$p(\mathbf{s}, \mathbf{t}) = 2^{-|y|} \mathbb{P}\{R_2 = \mathbf{d}\}. \tag{2.7}$$

3 Backward transition probabilities

Note that if $\mathbf{s} \in \mathbb{S}_n$ and $\mathbf{t} \in \mathbb{S}_{n+1}$ are such that $\mathbb{P}\{\nu R_{n+1} = \mathbf{t} \mid \nu R_n = \mathbf{s}\} > 0$, then the leaf set of \mathbf{s} is obtained either by removing a leaf from the leaf set of \mathbf{t} that has a sibling which is not a leaf (corresponding to Case I above), in which case (1.2) implies that

$$\mathbb{P}\{\nu R_n = \mathbf{s} \mid \nu R_{n+1} = \mathbf{t}\} = \frac{1}{n+1}, \tag{3.1}$$

or by removing two sibling leaves from the leaf set of \mathbf{t} and replacing them by a single new leaf positioned at the start of the path that led from the rest of \mathbf{t} to their common parent (corresponding to Case II above), in which case (1.2) implies that

$$\mathbb{P}\{\nu R_n = \mathbf{s} \mid \nu R_{n+1} = \mathbf{t}\} = \frac{2}{n+1}. \tag{3.2}$$

These backward transition probabilities can also be obtained directly. We have, of course, that

$$\mathbb{P}\{\nu R_n = \mathbf{s} \mid \nu R_{n+1} = \mathbf{t}\} = \frac{\mathbb{P}\{\nu R_n = \mathbf{s}\} \mathbb{P}\{\nu R_{n+1} = \mathbf{t} \mid \nu R_n = \mathbf{s}\}}{\mathbb{P}\{\nu R_{n+1} = \mathbf{t}\}}. \quad (3.3)$$

Write $\mathbf{L}(\mathbf{s}) = \{y_1, \dots, y_n\}$. In Case I (using the notation that was introduced to first describe this case),

$$\mathbb{P}\{\nu R_n = \mathbf{s}\} = n! \prod_{k=1}^n \nu(\tau(y_k)) \quad (3.4)$$

by (2.1),

$$\mathbb{P}\{\nu R_{n+1} = \mathbf{t} \mid \nu R_n = \mathbf{s}\} = \nu(\tau(w)) \quad (3.5)$$

by (2.3), and

$$\mathbb{P}\{\nu R_{n+1} = \mathbf{t}\} = (n+1)! \left(\prod_{k=1}^n \nu(\tau(y_k)) \times \nu(\tau(w)) \right) \quad (3.6)$$

by (2.1). Substituting (3.4), (3.5), and (3.6) into (3.3) gives (3.1). In Case II (also using the notation that was introduced to first describe this case), (3.4) still holds,

$$\mathbb{P}\{\nu R_{n+1} = \mathbf{t} \mid \nu R_n = \mathbf{s}\} = \frac{2\nu(\tau(y'))\nu(\tau(y''))}{\nu(\tau(y))} \quad (3.7)$$

by (2.5),

$$\mathbb{P}\{\nu R_{n+1} = \mathbf{t}\} = (n+1)! \left(\frac{\prod_{k=1}^n \nu(\tau(y_k))}{\nu(\tau(y))} \times \nu(\tau(y')) \times \nu(\tau(y'')) \right), \quad (3.8)$$

by (2.1). Substituting (3.4), (3.7), and (3.8) into (3.3) gives (3.2).

The above observations are summarized in the following Definition and Remark.

Definition 3.1. Suppose that $\mathbf{t} \in \mathbb{S}_{n+1}$ and $v = v_1 \dots v_m$ is a leaf of \mathbf{t} . If $v_1 \dots \bar{v}_m$ is not a leaf of \mathbf{t} , let $\kappa(\mathbf{t}, v) \in \mathbb{S}_n$ be the tree $\mathbf{t} \setminus \{v\}$ (that is, $\kappa(\mathbf{t}, v)$ is the tree with the same leaf set as \mathbf{t} except that v has been removed).

If $v_1 \dots \bar{v}_m$ is also a leaf of \mathbf{t} , then there is a largest $\ell < m$ such that $v_1 \dots v_\ell$ and $v_1 \dots v_{\ell-1} \bar{v}_\ell$ are both vertices of \mathbf{t} , and in this case let $\kappa(\mathbf{t}, v) \in \mathbb{S}_n$ be the tree $\mathbf{t} \setminus (\{v_1 \dots v_p : \ell + 1 \leq p \leq m\} \cup \{v_1 \dots v_{m-1} \bar{v}_m\})$ (that is, $\kappa(\mathbf{t}, v)$ is the tree with the same leaf set as \mathbf{t} except that v and its sibling leaf $v_1 \dots v_{m-1} \bar{v}_m$ have both been removed and replaced by the single leaf $v_1 \dots v_\ell$).

Remark 3.2. Using Definition 3.1, we can describe the backward evolution of $(\nu R_n)_{n \in \mathbb{N}}$ by saying that conditional on $\{\nu R_{n+1}, \nu R_{n+2}, \dots\}$ one of the $n+1$ leaves of νR_{n+1} is chosen uniformly at random and, denoting this leaf by V_{n+1} , the random tree νR_n is constructed as $\kappa(\nu R_{n+1}, V_{n+1})$.

4 The Doob-Martin kernel and harmonic functions

Suppose that $\mathbf{s} \in \mathbb{S}_m$ and $\mathbf{t} \in \mathbb{S}_{m+n}$ are such that $\mathbb{P}\{R_{m+n} = \mathbf{t} \mid R_m = \mathbf{s}\} > 0$, a state of affairs which we denote by $\mathbf{s} \triangleleft \mathbf{t}$. Write x_1, \dots, x_p for the vertices of \mathbf{s} that have out-degree 1 and y_1, \dots, y_q for the leaves of \mathbf{s} . Of course, $q = m$, but it will be clearer to use this alternative notation. Then \mathbf{t} is obtained from \mathbf{s} by attaching subtrees to some of the vertices $\{x_1, \dots, x_p\} \cup \{y_1, \dots, y_q\}$. More precisely, $\mathbf{t} \setminus \mathbf{s} = (\bigsqcup_{i=1}^p \mathbf{a}_i) \sqcup (\bigsqcup_{j=1}^q \mathbf{b}_j)$ where the subtrees \mathbf{a}_i and \mathbf{b}_j are as follows. Suppose that $x_i = x_{i1} \dots x_{if_i}$ and $u_i \in \{0, 1\}$ is such that $x_{i1} \dots x_{if_i} u_i = x_i u_i \notin \mathbf{s}$, then either $\mathbf{a}_i = \emptyset$ (that is, no subtree is attached to x_i , in which case we set $\alpha_i = 0$) or there is an $\alpha_i \geq 1$ and $\mathbf{c}_i \in \mathbb{S}_{\alpha_i}$ such that $\mathbf{a}_i = \{x_i u_i w : w \in \mathbf{c}_i\}$. Suppose that $y_j = y_{j1} \dots y_{jg_j}$, then either $\mathbf{b}_j = \emptyset$ (that is, no

subtree is attached to y_j , in which case we set $\beta_j = 0$) or there is a $\beta_j \geq 1$ and $\mathbf{d}_j \in \mathbb{S}_{\beta_j+1}$ such that $\mathbf{b}_j = \{y_j w : w \in \mathbf{d}_j\} \setminus \{y_j\}$. We have $n = \sum_i \alpha_i + \sum_j \beta_j$. Given a tree $\mathbf{r} \in \mathbb{S}_h$ for some $h \in \mathbb{N}$, set $M(\mathbf{r}) = h$ (so that $M(\mathbf{r})$ is the number of leaves of \mathbf{r}) and $\pi(\mathbf{r}) := \mathbb{P}\{R_h = \mathbf{r}\}$.

Then, by iterating the arguments that lead to (2.4) and (2.7),

$$\begin{aligned} & \mathbb{P}\{R_{m+n} = \mathbf{t} \mid R_m = \mathbf{s}\} \\ &= \frac{n!}{\prod_i \alpha_i! \prod_j \beta_j!} \prod_i (2^{-(|x_i|+1)})^{\alpha_i} \prod_j (2^{-|y_j|})^{\beta_j} \prod_{\alpha_i \neq 0} \pi(\mathbf{c}_i) \prod_{\beta_j \neq 0} \pi(\mathbf{d}_j). \end{aligned}$$

Also, because of (2.2),

$$\begin{aligned} & \mathbb{P}\{R_{m+n} = \mathbf{t}\} \\ &= (m+n)! \prod_{\alpha_i \neq 0} (2^{-(|x_i|+1)})^{\alpha_i} \frac{1}{\alpha_i!} \pi(\mathbf{c}_i) \prod_{\beta_j=0} (2^{-|y_j|}) \prod_{\beta_j \neq 0} (2^{-|y_j|})^{(\beta_j+1)} \frac{1}{(\beta_j+1)!} \pi(\mathbf{d}_j) \\ &= \frac{(m+n)!}{\prod_i \alpha_i! \prod_j (\beta_j+1)!} \prod_i (2^{-(|x_i|+1)})^{\alpha_i} \prod_j (2^{-|y_j|})^{(\beta_j+1)} \prod_{\alpha_i \neq 0} \pi(\mathbf{c}_i) \prod_{\beta_j \neq 0} \pi(\mathbf{d}_j). \end{aligned}$$

Note also, that

$$\mathbf{s} \triangleleft \mathbf{t} \iff \{v : v \in \mathbf{L}(\mathbf{t}), y_j \leq v\} \neq \emptyset, 1 \leq j \leq m.$$

Therefore, the Doob-Martin kernel is

$$\begin{aligned} K(\mathbf{s}, \mathbf{t}) &= \frac{\mathbb{P}\{R_{m+n} = \mathbf{t} \mid R_m = \mathbf{s}\}}{\mathbb{P}\{R_{m+n} = \mathbf{t}\}} \\ &= \mathbb{1}_{\{\mathbf{s} \triangleleft \mathbf{t}\}} \frac{\prod_{j=1}^m 2^{|y_j|} (\beta_j + 1)}{(n+1) \cdots (n+m)} \\ &= \mathbb{1}_{\{\mathbf{s} \triangleleft \mathbf{t}\}} \frac{\prod_{j=1}^m 2^{|y_j|} \#\{v : v \in \mathbf{L}(\mathbf{t}), y_j \leq v\}}{M(\mathbf{t})(M(\mathbf{t}) - 1) \cdots (M(\mathbf{t}) - m + 1)} \\ &= \frac{\prod_{j=1}^m 2^{|y_j|} \#\{v : v \in \mathbf{L}(\mathbf{t}), y_j \leq v\}}{M(\mathbf{t})(M(\mathbf{t}) - 1) \cdots (M(\mathbf{t}) - m + 1)}. \end{aligned}$$

Remark 4.1. It follows that, for $\mathbf{s} \in \mathbb{S}_m$, $m \in \mathbb{N}$ with leaves $\mathbf{L}(\mathbf{s}) = \{y_1, \dots, y_m\}$ and a sequence $(\mathbf{t}_n)_{n \in \mathbb{N}}$ with $\lim_{n \rightarrow \infty} M(\mathbf{t}_n) = \infty$, the sequence $K(\mathbf{s}, \mathbf{t}_n)$ converges as $n \rightarrow \infty$ if and only if the limit of

$$\prod_{j=1}^m \left[\frac{\#\{v : v \in \mathbf{L}(\mathbf{t}_n), y_j \leq v\}}{M(\mathbf{t}_n)} \middle/ \gamma(\tau(y_j)) \right] \tag{4.1}$$

exists, in which case the limits coincide. Recall that for $y \in \{0, 1\}^*$ the cardinality $\#\{1 \leq j \leq n : y \leq \zeta_{n,j}(z_1, \dots, z_n)\}$ equals $\#\{1 \leq j \leq n : y \leq z_j\}$ if the latter cardinality is at least two and it is zero otherwise. Hence a sufficient condition for the limit as $n \rightarrow \infty$ of $K(\mathbf{s}, \mathbf{t}_n)$ (equivalently, of (4.1)) to exist for all $\mathbf{s} \in \mathbb{S}$ is that $\mathbf{t}_n = \mathbf{R}(z_1, \dots, z_n)$ for a sequence $(z_n)_{n \in \mathbb{N}}$ of distinct elements of $\{0, 1\}^\infty$ such that for some probability measure ν on $\{0, 1\}^\infty$ we have

$$\nu\{z \in \{0, 1\}^\infty : y \leq z\} = \lim_{n \rightarrow \infty} \frac{1}{n} \#\{1 \leq j \leq n : y \leq z_j\}$$

for all $y \in \{0, 1\}^*$; that is, the sequence of empirical probability distributions $(\frac{1}{n} \sum_{j=1}^n \delta_{z_j})_{n \in \mathbb{N}}$ converges weakly to ν (where we put the usual topology on $\{0, 1\}^\infty$ for which the sets $\tau(y)$ are both closed and open). In this case

$$\lim_{n \rightarrow \infty} K(\mathbf{s}, \mathbf{t}_n) = \nu h(\mathbf{s}) := \prod_{a \in \mathbf{L}(\mathbf{s})} \frac{\nu(\tau(a))}{\gamma(\tau(a))}. \tag{4.2}$$

The function ${}^\nu h$ is excessive as a pointwise limit of excessive functions. Moreover, if ν is diffuse, then

$$\lim_{n \rightarrow \infty} K(\mathbf{s}, {}^\nu R_n) = {}^\nu h(\mathbf{s}), \quad \mathbb{P} - \text{a.s.}$$

for all $\mathbf{s} \in \mathbb{S}$.

It is immediate from the expressions for the forward transition probabilities derived in Section 2 that

$$\mathbb{P}\{{}^\nu R_{n+1} = \mathbf{t} \mid {}^\nu R_n = \mathbf{s}\} = {}^\nu h(\mathbf{s})^{-1} \mathbb{P}\{R_{n+1} = \mathbf{t} \mid R_n = \mathbf{s}\} {}^\nu h(\mathbf{t}).$$

Thus, the nonnegative function ${}^\nu h$ is harmonic, the Markov chain $({}^\nu R_n)_{n \in \mathbb{N}}$ is the h -transform of $(R_n)_{n \in \mathbb{N}}$ with the harmonic function ${}^\nu h$, and hence $({}^\nu R_n)_{n \in \mathbb{N}}$ is an infinite bridge for $(R_n)_{n \in \mathbb{N}}$. Recall that the tail σ -field of $({}^\nu R_n)_{n \in \mathbb{N}}$ is \mathbb{P} -a.s. trivial. It follows that the normalized nonnegative harmonic function ${}^\nu h$ is extremal. We show in Theorem 6.1 and Corollary 6.3 that the extremal normalized nonnegative harmonic functions are precisely those of this form and that they are, in turn, precisely the harmonic functions that arise as a limit of the form $\mathbf{r} \mapsto \lim_{k \rightarrow \infty} K(\mathbf{r}, \mathbf{t}_k)$, where $(\mathbf{t}_k)_{k \in \mathbb{N}}$ is such that $M(\mathbf{t}_k) \rightarrow \infty$ as $k \rightarrow \infty$. In the language of Doob–Martin theory, this shows that the *minimal Doob–Martin boundary* of the radix sort tree chain $(R_n)_{n \in \mathbb{N}}$ coincides with the full *Doob–Martin boundary*. It may be feasible to prove this fact “bare–hands”, but the simpler indirect route we take is, we believe, more informative.

5 Labeled infinite bridges

Recall that the backward transition dynamics of any finite bridge $(R_n^{\mathbf{t}})_{n=1}^{M(\mathbf{t})}$ and any infinite bridge $(R_n^\infty)_{n \in \mathbb{N}}$ may be described in terms of the “pruning” operation κ from Definition 3.1 and Remark 3.2:

- Suppose that the value of the process at time $n + 1$ is $\mathbf{t} \in \mathbb{S}_{n+1}$.
- Pick a leaf v uniformly at random.
- Replace \mathbf{t} by $\kappa(\mathbf{t}, v) \in \mathbb{S}_n$ to produce the value of the process at time n .

Consider a binary tree $\mathbf{t}'' \in \mathbb{S}_{n+1}$. Label the $n + 1$ leaves of \mathbf{t}'' with $[n + 1]$ uniformly at random (that is, all $(n + 1)!$ labelings are equally likely). Let V be the leaf labeled $n + 1$. Set $\mathbf{t}' := \kappa(\mathbf{t}'', V)$. If the sibling of V was not a leaf in \mathbf{t}'' , then the leaves of \mathbf{t}' were also leaves of \mathbf{t}'' and we maintain their labels. If the sibling of V was also a leaf of \mathbf{t}'' , labeled, say, $k \in [n]$, then in passing from \mathbf{t}'' to \mathbf{t}' we remove V and its sibling along with some vertices on the path leading to their parent, thereby creating a new leaf which we label k while leaving the labels of the remaining leaves (which are common to both \mathbf{t}'' and \mathbf{t}') unchanged. The distribution of \mathbf{t}' is that arising from one step starting from \mathbf{t}'' of the backward radix sort dynamics (that is, the common backward dynamics of all infinite bridges). Moreover, the labeling of \mathbf{t}' by $[n]$ is uniformly distributed over the $n!$ possible labelings.

Now suppose that $(R_n^\infty)_{n \in \mathbb{N}}$ is an infinite bridge. For some N , let S_N be a random binary tree with the same distribution as R_N^∞ . Label S_N uniformly at random with $[N]$ to produce a leaf-labeled binary tree \tilde{S}_N . The pruning procedure described above is deterministic once the labeling is given and applying it successively for $n = N - 1, \dots, 1$ produces leaf-labeled binary trees $\tilde{S}_{N-1}, \dots, \tilde{S}_1$, where \tilde{S}_n has n leaves labeled by $[n]$ for $1 \leq n \leq N - 1$. Write S_n for the underlying binary tree obtained by removing the labels of \tilde{S}_n . It follows from the observations above that the sequence (S_1, \dots, S_N) has the same joint distribution as $(R_1^\infty, \dots, R_N^\infty)$. Note that the joint distribution of the sequence $(\tilde{S}_1, \dots, \tilde{S}_N)$ is uniquely determined by the distribution of R_N^∞ and hence, *a fortiori*, by the joint distribution of $(R_n^\infty)_{n \in \mathbb{N}}$. Note also that if we perform this construction for two

different values of N , say $N' < N''$, to produce, with the obvious notation, sequences $(\tilde{S}'_1, \dots, \tilde{S}'_{N'})$ and $(\tilde{S}''_1, \dots, \tilde{S}''_{N''})$, then $(\tilde{S}'_1, \dots, \tilde{S}'_{N'})$ has the same joint distribution as $(\tilde{S}''_1, \dots, \tilde{S}''_{N'})$.

By Kolmogorov's extension theorem we may therefore suppose that there is a Markov process $(\tilde{R}_n^\infty)_{n \in \mathbb{N}}$ such that for each $n \in \mathbb{N}$ the random element \tilde{R}_n is a leaf-labeled binary tree with n leaves labeled by $[n]$ and the following hold.

- The binary tree obtained by removing the labels of \tilde{R}_n is R_n .
- For every $n \in \mathbb{N}$, the conditional distribution of \tilde{R}_n^∞ given R_n^∞ is uniform over the $n!$ possible labelings of R_n^∞ .
- In going backward from time $n+1$ to time n , \tilde{R}_{n+1}^∞ is transformed into \tilde{R}_n^∞ according to the deterministic procedure described above.

The distribution of the labeled infinite bridge $(\tilde{R}_n)_{n \in \mathbb{N}}$ is uniquely specified by the distribution of $(R_n^\infty)_{n \in \mathbb{N}}$ and the above requirements. Because of this distributional uniqueness, we refer to $(\tilde{R}_n^\infty)_{n \in \mathbb{N}}$ as **the labeled version** of $(R_n^\infty)_{n \in \mathbb{N}}$ and $(R_n^\infty)_{n \in \mathbb{N}}$ as the **unlabeled version** of $(\tilde{R}_n^\infty)_{n \in \mathbb{N}}$ and speak of the "leaf of R_n^∞ labeled with $i \in [n]$ in \tilde{R}_n^∞ ."

Definition 5.1. Given $i \in [n]$, let $\langle i \rangle_n \in \{0, 1\}^*$ be the leaf of R_n^∞ labeled i in \tilde{R}_n^∞ . Observe that $\langle i \rangle_i \leq \langle i \rangle_{i+1} \leq \dots$ and so $\langle i \rangle_\infty = \lim_{n \rightarrow \infty} \langle i \rangle_n \in \{0, 1\}^* \sqcup \{0, 1\}^\infty$ is well-defined. Moreover, for distinct $i, j \in \mathbb{N}$, $\langle i \rangle_n \wedge \langle j \rangle_n$ is the same for all $n \geq i \wedge j$ and coincides with $\langle i \rangle_\infty \wedge \langle j \rangle_\infty$.

Remark 5.2. We have $R_1^\infty \subset R_2^\infty \subset \dots$ and

$$R_\infty^\infty := \bigcup_{n \in \mathbb{N}} R_n^\infty = \bigcup_{i \in \mathbb{N}} \{v \in \{0, 1\}^* \sqcup \{0, 1\}^\infty : v \leq \langle i \rangle_\infty\}.$$

That is, R_∞^∞ is the subtree of $\{0, 1\}^* \sqcup \{0, 1\}^\infty$ with leaves $\{\langle i \rangle_\infty : i \in \mathbb{N}\}$ and we define \tilde{R}_∞^∞ to be the tree R_∞^∞ with the leaf $\langle i \rangle_\infty$ labeled i , $i \in \mathbb{N}$. We will drop the subscripts and write $\langle i \rangle$ for $\langle i \rangle_\infty$, $i \in \mathbb{N}$.

6 Proof of Theorem 1.1

Theorem 1.1 is an immediate consequence of Theorem 6.1 and Corollary 6.3 below.

Theorem 6.1. Consider an infinite bridge $(R_n^\infty)_{n \in \mathbb{N}}$ and its associated labeled version $(\tilde{R}_n^\infty)_{n \in \mathbb{N}}$.

- (a) The sequence $(\langle i \rangle)_{i \in \mathbb{N}}$ is exchangeable.
- (b) The tail σ -field of $(R_n^\infty)_{n \in \mathbb{N}}$ is \mathbb{P} -a.s. trivial if and only if $(\langle i \rangle)_{i \in \mathbb{N}}$ is an independent identically distributed sequence.
- (c) If $(\langle i \rangle)_{i \in \mathbb{N}}$ is independent and identically distributed with common distribution ν , then ν is concentrated on $\{0, 1\}^\infty$ and diffuse.
- (d) The tail σ -field of $(R_n^\infty)_{n \in \mathbb{N}}$ is \mathbb{P} -a.s. trivial if and only if $(R_n^\infty)_{n \in \mathbb{N}}$ has the same distribution as $({}^\nu R_n)_{n \in \mathbb{N}}$ for some diffuse probability measure ν on $\{0, 1\}^\infty$.

Proof. (a) It is clear by construction that $(\langle i \rangle_n)_{i \in [n]}$ is (finitely) exchangeable and the claim follows upon taking limits as $n \rightarrow \infty$.

(b) The bijective correspondence between the distributions of the infinite bridges $(R_n^\infty)_{n \in \mathbb{N}}$ and the distributions of their labeled versions $(\tilde{R}_n^\infty)_{n \in \mathbb{N}}$ is compatible with convex combinations, and hence preserves extremality. Therefore the tail σ -field of the infinite bridge $(R_n^\infty)_{n \in \mathbb{N}}$ is \mathbb{P} -a.s. trivial if and only if the exchangeable sequence $(\langle i \rangle)_{i \in \mathbb{N}}$ is ergodic. (This situation closely parallels one appearing in the analysis of Rémy's tree growth chain in [5], and we refer to the more detailed argument in Proposition 5.19 of

[5].) Finally, a well-known consequence of de Finetti's theorem is that an exchangeable sequence is ergodic if and only if it is independent and identically distributed.

(c) For any $u \in \{0, 1\}^*$, the sequence $(\mathbb{1}\{u = \langle k \rangle\})_{k \in \mathbb{N}}$ is independent and identically distributed, and hence $\#\{k \in \mathbb{N} : u = \langle k \rangle\} = 0$ \mathbb{P} -a.s. or $\#\{k \in \mathbb{N} : u = \langle k \rangle\} = \infty$ \mathbb{P} -a.s. Now, if $\mathbb{P}\{\langle i \rangle \in \{0, 1\}^*\} > 0$ there would be a $u \in \{0, 1\}^*$ such that with positive probability $\langle i \rangle_n = \langle i \rangle = u$ for all n sufficiently large. Then, on the event $\{\langle i \rangle = u\}$ we would have $\#\{k \in \mathbb{N} : \langle k \rangle = u\} = 1$, since it follows from the construction in Definition 5.1 that $\langle j \rangle \neq \langle i \rangle$ for $j \neq i$ when $\langle i \rangle \in \{0, 1\}^*$. This shows that $\mathbb{P}\{\langle i \rangle \in \{0, 1\}^*\} = 0$.

We therefore have that $(\langle k \rangle)_{k \in \mathbb{N}}$ is an independent identically distributed sequence of $\{0, 1\}^\infty$ -valued random variables. Because $\langle i \rangle \wedge \langle j \rangle = \langle i \rangle_n \wedge \langle j \rangle_n \in \{0, 1\}^*$ for all $n \geq i \vee j$ \mathbb{P} -a.s. when $i \neq j$, it follows that $\langle i \rangle \neq \langle j \rangle$ \mathbb{P} -a.s. for $i \neq j$ and the common distribution of $(\langle k \rangle)_{k \in \mathbb{N}}$ is diffuse.

(d) We have already seen that when ν is a diffuse probability measure on $\{0, 1\}^\infty$ the process $({}^\nu R_n)_{n \in \mathbb{N}}$ is an infinite bridge which, by the Hewitt-Savage zero-one law, has a trivial tail σ -field.

Conversely, suppose that the infinite bridge $(R_n^\infty)_{n \in \mathbb{N}}$ has a trivial tail σ -field. Let ν be the common diffuse distribution of the independent, identically distributed sequence of $\{0, 1\}^\infty$ -valued random variables $(\langle i \rangle)_{i \in \mathbb{N}}$. In the notation of the Introduction, it is clear that $R_n^\infty = \mathbf{R}(\langle 1 \rangle, \dots, \langle n \rangle)$, $n \in \mathbb{N}$, and so $(R_n^\infty)_{n \in \mathbb{N}}$ has the same distribution as $({}^\nu R_n)_{n \in \mathbb{N}}$. \square

Remark 6.2. One might expect that the analogue of Theorem 6.1 holds, with a similar proof, for the infinite bridges of the PATRICIA chain with i.i.d. fair coin-tossing inputs. Note, however, that the proof cannot be carried through *mutatis mutandis* because the monotonicity of the sequences $(\langle i \rangle_n)_{n \geq i}$, $i \in \mathbb{N}$, observed in Definition 5.1 does not hold for the counterpart of these sequences for an infinite PATRICIA bridge. However, it follows from the identification of the backward transition probabilities of $({}^\nu R_n)_{n \in \mathbb{N}}$ in Section 3 that the conditional distribution of ${}^\nu \bar{R}_n$ given $\{\bar{R}_{n+1} = \mathbf{t}\}$ is that of the random tree obtained by picking one of the $n + 1$ leaves of \mathbf{t} uniformly at random, deleting it and its sibling, and “closing up the gap” if there is one. These backward transition probabilities are exactly the same as those for the Rémy chain in [5], and so a description of the extremal infinite bridges of the PATRICIA chain along with a characterization of its Doob–Martin boundary can be read from there; in particular, although each process $({}^\nu \bar{R}_n)_{n \in \mathbb{N}}$ is an extremal infinite bridge for $(\bar{R}_n)_{n \in \mathbb{N}}$, there are other extremal infinite bridges which are not of this type. We will lay out the details in a forthcoming paper.

Corollary 6.3. *The extremal normalized nonnegative harmonic functions are precisely those that arise as $\mathbf{s} \mapsto \lim_{k \rightarrow \infty} K(\mathbf{s}, \mathbf{t}_k)$ for a sequence $(\mathbf{t}_k)_{k \in \mathbb{N}}$ with $M(\mathbf{t}_k) \rightarrow \infty$ as $k \rightarrow \infty$. There is a bijective correspondence between diffuse probability measures on $\{0, 1\}^\infty$ and such functions: the measure ν corresponds to the normalized nonnegative harmonic function ${}^\nu h$ of (4.2) and, conversely, if h is an extremal normalized nonnegative harmonic function and $(R_n^\infty)_{n \in \mathbb{N}}$ is the infinite bridge constructed as the Doob h -transform of $(R_n)_{n \in \mathbb{N}}$ using the function h , then $h = {}^\nu h$, where ν is the common distribution of the independent identically distributed sequence $(\langle i \rangle)_{i \in \mathbb{N}}$ associated with the labeled infinite bridge $(\tilde{R}_n^\infty)_{n \in \mathbb{N}}$.*

Proof. We know from Theorem 6.1 that the extremal normalized nonnegative harmonic functions correspond to infinite bridges of the form $({}^\nu R_n)_{n \in \mathbb{N}}$ where ν is a diffuse probability measure on $\{0, 1\}^\infty$, and hence they are the harmonic functions ${}^\nu h$. In order to see that the correspondence between ν and the distribution of $({}^\nu R_n)_{n \in \mathbb{N}}$ is bijective, we observe that ν is determined uniquely by the distribution of the labeled version of $({}^\nu R_n)_{n \in \mathbb{N}}$ and hence by the distribution of $({}^\nu R_n)_{n \in \mathbb{N}}$ itself.

It remains to check that if the normalized nonnegative harmonic function h is given by $h(s) = \lim_{k \rightarrow \infty} K(s, \mathbf{t}_k)$ for a sequence $(\mathbf{t}_k)_{k \in \mathbb{N}}$ with $M(\mathbf{t}_k) \rightarrow \infty$ as $k \rightarrow \infty$, then h is extremal. We will follow an argument similar to the proof of Corollary 5.21 in [5]. Writing $(R_n^\infty)_{n \in \mathbb{N}}$ for the infinite bridge given by the Doob h -transform of $(R_n)_{n \in \mathbb{N}}$ associated with h , we recall that extremality of h is equivalent to the tail σ -field of $(R_n^\infty)_{n \in \mathbb{N}}$ being \mathbb{P} -a.s. trivial. By Theorem 6.1, this is in turn equivalent to showing that the exchangeable sequence $(\langle i \rangle)_{i \in \mathbb{N}}$ has the equivalent properties of being ergodic or independent and identically distributed.

Note that $\langle i \rangle$ is the unique $v \in \{0, 1\}^\infty$ such that $\langle i \rangle \wedge \langle j \rangle \leq v$ for all $j \neq i$. It follows that there is a measurable bijection mapping the sequence $(\langle i \rangle)_{i \in \mathbb{N}}$ to the jointly exchangeable $\{0, 1\}^*$ -valued array $\{\langle i \rangle \wedge \langle j \rangle : i, j \in \mathbb{N}, i \neq j\}$ in such a way that the sequence will be ergodic if and only if the array is ergodic. By a result of Aldous (see, for example, [7, Lemma 7.35]), the array is ergodic if and only if for any disjoint finite subsets H_1, \dots, H_s of \mathbb{N} the finite subarrays $\{\langle i \rangle \wedge \langle j \rangle : i, j \in H_r, i \neq j\}$, $1 \leq r \leq s$, are independent.

Recall that $(R_1^{\mathbf{t}_k}, \dots, R_{M(\mathbf{t}_k)}^{\mathbf{t}_k})$ denotes the bridge to \mathbf{t}_k . For any $\ell \in \mathbb{N}$, $R_\ell^{\mathbf{t}_k}$ converges in distribution to R_ℓ^∞ as $k \rightarrow \infty$. We can build a labeled version $(\tilde{R}_1^{\mathbf{t}_k}, \dots, \tilde{R}_{m(\mathbf{t}_k)}^{\mathbf{t}_k})$ of $(R_1^{\mathbf{t}_k}, \dots, R_{m(\mathbf{t}_k)}^{\mathbf{t}_k})$ in much the same way that we built a labeled version of an infinite bridge: $\tilde{R}_{m(\mathbf{t}_k)}^{\mathbf{t}_k}$ consists of the tree $R_{m(\mathbf{t}_k)}^{\mathbf{t}_k} = \mathbf{t}_k$ with its $M(\mathbf{t}_k)$ leaves labeled uniformly at random with the set $[M(\mathbf{t}_k)]$ and the backward evolution of such a labeled finite bridge is the same as that of the labeled infinite bridge. It is clear that $\tilde{R}_\ell^{\mathbf{t}_k}$ converges in distribution to \tilde{R}_ℓ^∞ as $k \rightarrow \infty$ for all $\ell \in \mathbb{N}$: indeed, $\tilde{R}_\ell^{\mathbf{t}_k}$ and \tilde{R}_ℓ^∞ are just $R_\ell^{\mathbf{t}_k}$ and R_ℓ^∞ , respectively, equipped with uniform random labelings of their ℓ leaves by the set $[\ell]$.

Write $\langle i \rangle_\ell^k$ for the element of $\{0, 1\}^*$ labeled i in $\tilde{R}_\ell^{\mathbf{t}_k}$ for $1 \leq i \leq \ell \leq M(\mathbf{t}_k)$. The finite array $\{\langle i \rangle_\ell^k \wedge \langle j \rangle_\ell^k : 1 \leq i \neq j \leq \ell\}$ converges in distribution to the finite array $\{\langle i \rangle_\ell \wedge \langle j \rangle_\ell : 1 \leq i \neq j \leq \ell\} = \{\langle i \rangle \wedge \langle j \rangle : 1 \leq i \neq j \leq \ell\}$ as $k \rightarrow \infty$.

Write $u_1^k, \dots, u_{M(\mathbf{t}_k)}^k$ for the leaves of \mathbf{t}_k . Suppose that $I_1^k, \dots, I_{M(\mathbf{t}_k)}^k$ is a listing of $[M(\mathbf{t}_k)]$ in uniform random order and $J_1^k, \dots, J_{M(\mathbf{t}_k)}^k$ is a sequence of independent random variables uniformly distributed on $[M(\mathbf{t}_k)]$. By definition, $(\langle i \rangle_\ell^k)_{1 \leq i \leq \ell}$ has the same distribution as $(u_{I_i^k}^k)_{1 \leq i \leq \ell}$. We may couple $I_1^k, \dots, I_{M(\mathbf{t}_k)}^k$ and $J_1^k, \dots, J_{M(\mathbf{t}_k)}^k$ together on the same probability space in such a way that $\lim_{k \rightarrow \infty} \mathbb{P}\{\exists 1 \leq i \leq \ell : I_i^k \neq J_i^k\} = 0$ and hence $\lim_{k \rightarrow \infty} \mathbb{P}\{\exists 1 \leq i \neq j \leq \ell : u_{I_i^k}^k \wedge u_{J_j^k}^k \neq u_{J_i^k}^k \wedge u_{J_j^k}^k\} = 0$. If H_1, \dots, H_s is a collection of disjoint subsets of $[\ell]$, and k is so large that $M(\mathbf{t}_k) \geq \ell$, then it is clear that the arrays $\{u_{J_i^k}^k \wedge u_{J_j^k}^k : i, j \in H_r, i \neq j\}$, $1 \leq r \leq s$, are independent and hence the arrays $\{\langle i \rangle \wedge \langle j \rangle : i, j \in H_r, i \neq j\}$, $1 \leq r \leq s$, are also independent, as required. \square

References

- [1] J. Clément, P. Flajolet, and B. Vallée, *Dynamical sources in information theory: a general analysis of trie structures*, *Algorithmica* **29** (2001), no. 1-2, 307–369, Average-case analysis of algorithms (Princeton, NJ, 1998). MR-1887308
- [2] Luc Devroye, *A study of trie-like structures under the density model*, *Ann. Appl. Probab.* **2** (1992), no. 2, 402–434. MR-1161060
- [3] J. L. Doob, *Discrete potential theory and boundaries*, *J. Math. Mech.* **8** (1959), 433–458; erratum 993. MR-0107098
- [4] Steven N. Evans, Rudolf Grübel, and Anton Wakolbinger, *Trickle-down processes and their boundaries*, *Electron. J. Probab.* **17** (2012), no. 1, 58. MR-2869248
- [5] Steven N. Evans, Rudolf Grübel, and Anton Wakolbinger, *Doob-Martin boundary of Rémy's tree growth chain*, *Ann. Probab.* **45** (2017), no. 1, 225–277. MR-3601650
- [6] G. Gwehenberger, *Anwendung einer binären Verweiskettenmethode beim Aufbau von Listen*, *Elektronische Rechenanlagen* **10** (1968), 223–226.

- [7] Olav Kallenberg, *Probabilistic symmetries and invariance principles*, Probability and its Applications (New York), Springer, New York, 2005. MR-2161313
- [8] Donald E. Knuth, *The art of computer programming. Volume 3*, Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont., 1973, Sorting and searching, Addison-Wesley Series in Computer Science and Information Processing. MR-0445948
- [9] _____, *The art of computer programming. Vol. 3*, Addison-Wesley, Reading, MA, 1998, Sorting and searching, Second edition [of MR0445948]. MR-3077154
- [10] Kevin Leckey, Ralph Neininger, and Wojciech Szpankowski, *A limit theorem for Radix Sort and tries with Markovian input*, 2015, Available at arXiv:1505.07321 [math.PR].
- [11] Hosam M. Mahmoud, *Evolution of random search trees*, Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., New York, 1992, A Wiley-Interscience Publication. MR-1140708
- [12] Donald R. Morrison, *Patricia—Practical Algorithm To Retrieve Information Coded in Alphanumeric*, Journal of the ACM **15** (1968), 514–534.
- [13] L. C. G. Rogers and J. W. Pitman, *Markov functions*, Ann. Probab. **9** (1981), no. 4, 573–582. MR-624684
- [14] W. Szpankowski and Ph. Jacquet, *Analysis of digital tries with Markovian dependency*, IEEE Trans. Information Theory **27** (1991), 1470–1475.
- [15] Wojciech Szpankowski, *Average case analysis of algorithms on sequences*, Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley-Interscience, New York, 2001, With a foreword by Philippe Flajolet. MR-1816272

Acknowledgments. We thank Kevin Leckey and Ralph Neininger for valuable information about the literature around radix sort algorithms. We thank the anonymous referee for a very careful reading of the paper and for suggesting that our results for the radix sort chain may have counterparts for the PATRICIA chain.