

A NEW PROOF OF CAYLEY'S FORMULA

LOUIGI ADDARIO-BERRY, SERTE DONDERWINKEL, MICKAËL MAAZOUN,
AND JAMES B. MARTIN

ABSTRACT. We present a very simple new bijective proof of Cayley's formula. The bijection is useful for the analysis of random trees, and we explain some of the ways in which it can be used to derive probabilistic identities, bounds, and growth procedures for such trees. We also introduce a partial order on the degree sequences of rooted trees, and conjecture that it induces a stochastic partial order on heights of random rooted trees with given degrees.

1. INTRODUCTION

A rooted tree is a triple $t = (V, E, \rho)$ where (V, E) is a tree and $\rho \in V$. We write

$$\mathcal{T}_n = \{t = (V, E, \rho) : t \text{ is a rooted tree with } V = [n]\}$$

for the set of rooted trees with vertex set $[n] := \{1, \dots, n\}$. *Cayley's formula* (which to the best of current knowledge was first established by Borchardt [6]) provides a very simple formula for $|\mathcal{T}_n|$.

Theorem 1.1 (Cayley's formula). $|\mathcal{T}_n| = n^{n-1}$.

There are numerous proofs of Cayley's formula in the literature. The 1967 survey by Moon [13] presents ten such proofs, including the proof via so-called *Prüfer codes*, which is probably the one most frequently presented in undergraduate texts. More recent proofs include those discovered by Joyal [12, Section 2.2], which considers doubly-rooted trees; by Pitman [16], which analyzes a coalescent process for building rooted labeled trees; and by van der Hofstad [20, Section 3.6], which uses the connection between random labeled trees and conditioned Poisson branching processes.


In this work, we present a new proof of Cayley's formula, via a so-called *line-breaking* construction. We believe this is the simplest proof of Cayley's formula yet discovered. We discuss related constructions for unrooted trees, rooted trees with marked vertices and rooted forests in Section 2. The bijection is also extremely useful for the analysis of random combinatorial trees; Section 3 of this paper discusses some of its probabilistic implications. In that section, we also define an algorithm for growing random rooted trees with given vertex degrees. When applied to regular trees, the algorithm yields a sequence $(T_m, m \geq 1)$, where for each $m \geq 1$, T_m is a uniformly random rooted, leaf-labeled d -ary tree with m internal nodes, and T_{m+1} is generated from T_m by local regrafting. Our algorithm is somewhat similar in spirit to Rémy's algorithm [17], which generates a sequence of uniform binary leaf-labeled trees.

Before giving the proof of Cayley's formula, we introduce some terminology. Suppose t is a tree, S is a connected subset of its vertices, and x is a vertex. The *path from S to x* in t is the unique path in t which starts at a vertex of S , does not visit any other vertex of S , and ends at x . If P is a path, we may also write P to denote the set of vertices of the path. Finally, a *leaf* of t is a non-root vertex of t with no children.

Date: July 20, 2021.

2010 Mathematics Subject Classification. 05C05, 60C05.

Key words and phrases. Cayley's formula, random trees.

FIGURE 1. A tree t and the corresponding sequence $v(t)$.

Proof of Cayley's formula. The term n^{n-1} counts sequences $v = (v_1, \dots, v_{n-1}) \in [n]^{n-1}$. We prove the theorem by introducing a bijection between \mathcal{T}_n and $[n]^{n-1}$.

Bijection

For a tree t on $[n]$ with root ρ :

- Let $\ell_1 < \ell_2 < \dots < \ell_k$ be the leaves of t .
- Let $S_0 = \{\rho\}$.
- Recursively, for $i = 1, \dots, k$, let P_i be the path in t from S_{i-1} to ℓ_i , and let $S_i = S_{i-1} \cup P_i$. Let P_i^* be P_i omitting its final point.
- Let $v(t)$ be the concatenation of $P_1^*, P_2^*, \dots, P_k^*$.

We claim that $v(t) \in [n]^{n-1}$. Indeed, by definition, for each i , $P_i^* \subset [n]$. Moreover, note that P_1, P_2, \dots, P_k forms a partition of the edge set of t , and for each i , $|P_i^*|$ equals the number of edges in P_i . This implies that the length of $v(t)$ is equal to the number of edges in t , which equals $n - 1$.

To show that the above construction is a bijection, we describe its inverse. For a sequence $v = v_1, v_2, \dots, v_m$, we say that $j \in \{2, \dots, m\}$ is the location of a repeat of v if $v_j = v_i$ for some $i < j$.

Inverse of the bijection

Given a sequence $v = (v_1, v_2, \dots, v_{n-1}) \in [n]^{n-1}$:

- Let $j(0) = 1$, let $j(1) < j(2) < \dots < j(k-1)$ be the locations of the repeats of the sequence v , and let $j(k) = n$.
- Let $\ell_1 < \ell_2 < \dots < \ell_k$ be the elements of $[n]$ not occurring in v .
- For $i = 1, \dots, k$, let P_i be the path $(v_{j(i-1)}, \dots, v_{j(i)-1}, \ell_i)$ with $j(i) - j(i-1)$ edges.
- Let $t(v)$ be the graph with vertex set $[n]$, with root v_1 and with edge set given by the union of the edges of the paths P_1, P_2, \dots, P_k .

We claim that $t(v)$ is a rooted tree with vertex set $[n]$. By construction, for $i = 2, \dots, k$, the set of vertices $\cup_{m < i} P_m$ and the path P_i intersect at a single point (namely $v_{j_{i-1}}$, the first point of P_i). Hence, by induction, for each i the union of the edges of the paths P_1, \dots, P_i is a tree, and in particular taking $i = k$ we have that $t(v)$ is a tree. Since for $i = 1, \dots, k$, the final point of P_i is ℓ_i , and since no element ℓ_i appears as the first point of any other path P_m , we have that the leaves of $t(v)$ are ℓ_1, \dots, ℓ_k .

We will argue that if $v \in [n]^{n-1}$, then $v(t(v)) = v$. Let $k, P_1, \dots, P_k, \ell_1, \dots, \ell_k$, and $v_1, v_{j_1}, \dots, v_{j_k}$ be defined as in the inverse of the bijection applied to v . By definition, v_1 is the root in $t(v)$. Note that P_1 is a path contained in $t(v)$ with endpoints v_1 and ℓ_1 . The fact that $t(v)$ is a tree implies that P_1 is the unique path from v_1 to ℓ_1 in $t(v)$. Similarly, for $1 < i \leq k$, P_i is a path contained in $t(v)$ with endpoints $v_{j_{i-1}}$ and ℓ_i , and

with $P_i \cap \cup_{m < i} P_m = \{v_{j_{i-1}}\}$, so that P_i is the unique path in $t(v)$ from the union of P_1, \dots, P_{i-1} to ℓ_i . Hence $v(t(v))$ is the concatenation of P_1^*, \dots, P_k^* , giving $v(t(v)) = v$ as desired. \square

2. EXTENSIONS AND RELATED CONSTRUCTIONS

Unrooted trees. A simple trick allows us to directly study unrooted trees with our bijection by considering vertex 1 as the root and removing the leading 1 of the corresponding coding sequence. This gives a proof of the more well-known form of Cayley's formula which states that there are n^{n-2} unrooted trees of size n .

There is another way to adapt our bijection to encode unrooted labeled trees of size n by sequences in $[n]^{n-2}$. It consists of letting P_1^* be the path between the first and second-lowest-labeled vertices of degree 1 (excluding both endpoints) and continuing as in the bijection for rooted trees.

Trees with given degrees. Our bijection has the property that for a given tree t , in the associated sequence v the number of times a given integer k appears is precisely the number of children of the vertex with label k in t .

The *type* of a rooted tree t is the vector $\mathbf{n} = (n_c, c \geq 0)$, where n_c is the number of vertices of t with exactly c children. Writing $n = \sum_{c \geq 0} n_c$ and

$$\mathcal{T}(\mathbf{n}) = \{t \in \mathcal{T}_n : t \text{ has type } \mathbf{n}\},$$

then

$$|\mathcal{T}(\mathbf{n})| = \binom{n}{n_c, c \geq 0} \cdot \frac{(n-1)!}{\prod_{c \geq 0} (c!)^{n_c}}.$$

This identity (also found in [15, Theorem 1.5] and [19, Corollary 3.5]) follows from the specializations of our bijection to trees of a given type: the right-hand side counts the number of ways to first choose the labels of the vertices with c children for each $c \geq 0$, then choose a sequence of length $n-1$ in which the label of each vertex with c children appears exactly c times.

Rooted trees with marked vertices. Our bijection can also be modified to encode rooted labeled trees with $r \geq 1$ distinguishable marks on the vertices, $(t, m_1, \dots, m_r) \in \mathcal{T}_n \times [n]^r$, by sequences in $[n]^{n+r-1}$. The modification consists of changing the definition of P_i^* in the recursive step slightly when constructing the sequence from the tree: for $i = 1, \dots, r$, P_i^* is the path from S_{i-1} to the i -th marked point (including the final point of this path), and for $i = 1, \dots, k$, P_{r+i}^* is the path from S_{r+i-1} to the i -th unmarked leaf ℓ_i (excluding its final point). Here the number of appearances of a vertex in the coding sequence is the sum of its number of children and its number of marks.

By taking the $r \rightarrow \infty$ limit of this bijection and applying it to an i.i.d. sequence of elements of $[n]$, one recovers the construction of the so-called *p-trees* from [7].

Rooted forests. The bijection above also extends to forests. Given a set $S \subset [n]$, write \mathcal{F}_n^S for the set of forests F with vertex set $[n]$ and root set S . Setting $s = |S|$, we describe a bijection between \mathcal{F}_n^S and the set

$$\{v = (v_1, \dots, v_{n-s}) \in [n]^{n-s} : v_1 \in S\},$$

which has cardinality sn^{n-s-1} . The analogue of the "inverse bijection" is the easiest to describe, so we begin with that. To construct a sequence (v_1, \dots, v_{n-s}) from $F \in \mathcal{F}_n^S$, proceed just as in our proof of Cayley's formula, but start from the s -vertex forest F_0 containing only the root vertices S , and at each step append the labels along the path to the smallest labeled leaf not already in the current forest (excluding the leaf itself).

Conversely, here is how to construct a forest $F \in \mathcal{F}_n^S$ from a sequence $v = (v_1, \dots, v_{n-s})$ with $v_1 \in S$. Say that i is the location of a repeat if $i > 1$ and either $v_i \in S$ or there is

$1 \leq j < i$ such that $v_i = v_j$. Denote the locations of repeats of v by $j(1), \dots, j(k-1)$ in increasing order, and let $j(k) = n - s + 1$. List the integers from $[n] \setminus S$ which do not appear in v as ℓ_1, \dots, ℓ_k in increasing order.

Form a graph $F_S = F_S(v)$ with vertices $[n]$, root set S , and edge set

$$\{v_i v_{i+1} : i \in [n-s-1], i+1 \notin \{j(1), \dots, j(k)\}\} \cup \{v_{j(i)-1} \ell_i, 1 \leq i \leq k\}.$$

Essentially the same argument as in our proof of Cayley's formula shows that the connected components of F_S are trees and that there are s components of F_S , each containing exactly one vertex of S . Thus, rooting each component of F_S at its unique element of S turns F_S into an element of \mathcal{F}_n^S .

Since the construction is bijective, this yields the following extension of Cayley's formula (which is in fact stated by both Borchardt [6] and Cayley [8]):

$$|\mathcal{F}_n^S| = |\{v = (v_1, \dots, v_{n-s}) \in [n]^{n-s} : v_1 \in S\}| = sn^{n-s-1}.$$

This extension of the bijection can also be specialized to count forests with fixed types.

Other coding sequences of trees. A variant of Prüfer codes for rooted trees, first described by Neville [14], gives another bijection between \mathcal{T}_n and $[n]^{n-1}$. It coincides with the original Prüfer code if one sees vertex n as the root and then removes the resulting trailing n . Deo and Micikevicius [9] found that reading the rooted Prüfer code in reverse leads to a simpler description of the bijection, which shares some features of the bijection presented in this paper. This bijection was independently discovered by Seo and Shin [18] under the name *Reverse-Prüfer codes* to study leader vertices in trees, and also by Fleiner [10], who remarked that it can be used to describe the law of the height of a uniformly sampled vertex in a random labeled tree (see Proposition 3.2, below).

3. RANDOM TREES

The bijection above has numerous consequences for random trees. This section first explains how the bijection can be used to study typical and extreme distances in random trees, then discusses how it can be used to define growth procedures for random trees, somewhat in the spirit of Rémy's algorithm [17] for growing uniformly random leaf-labeled binary trees.

For a finite set \mathcal{S} , we will write $X \in_u \mathcal{S}$ to mean that X is chosen uniformly at random from the set \mathcal{S} .

Distances in random trees. For a tree $t = (V, E, \rho)$ and vertices $u, v \in V$, write $[u, v] = [u, v]_t$ for the unique path from u to v in t , and $\text{dist}_t(u, v)$ for the distance from u to v in t , which equals the number of edges of $[u, v]$. Also, write $|v|$ for the graph distance from ρ to v .

Proposition 3.1. *Let $T \in_u \mathcal{T}_n$ and let L be a uniformly random leaf of T . Also, let $V = (V_i, i \geq 1)$ be a sequence of independent uniformly random elements of $[n]$ and let $I = \min(i \geq 2 : V_i \in \{V_1, \dots, V_{i-1}\})$ be the index of the first repeated element of V . Then $|L| + 1 \stackrel{d}{=} \min(I, n)$.*

Proof. Write $V = (V_1, \dots, V_{n-1}) \in_u [n]^{n-1}$. Then $T = t(V) \in_u \mathcal{T}_n$. Moreover, recalling that the repeated entries of (V_1, \dots, V_{n-1}) are $j(1), \dots, j(k-1)$ and that $j(k) = n$, we have $\min(I, n) = j(1)$. The first leaf $\ell_1(T)$ is a child of $V_{j(1)-1}$, so

$$|\ell_1(T)| = |V_{j(1)-1}| + 1 = j(1) - 1 + 1 = \min(I, n).$$

But since T is a uniformly random tree, randomly permuting its leaf labels does not change its distribution, so $|\ell_1(T)|$ has the same distribution as $|L|$ for L a uniformly random leaf of T . \square

One can also use the version of the bijection for rooted trees with one marked vertex to show the following similar result, whose proof is left to the reader.

Proposition 3.2. *Let $T \in_u \mathcal{T}_n$ and let $U \in_u [n]$ be independent of T . Then $|U| \stackrel{d}{=} I - 2$, where I is as in Proposition 3.1.*

These propositions in particular imply that if D_n is the distance from the root to a uniformly random vertex (or leaf) in $T_n \in_u \mathcal{T}_n$, then $n^{-1/2}D_n \xrightarrow{d} R$, where R is Rayleigh distributed: $\mathbf{P}\{R \geq x\} = e^{-(x^2/2)\mathbf{1}_{[x \geq 0]}}$. It is not hard to build on these statements in order to precisely characterize the joint distribution of the lengths of the branches to the k smallest labeled leaves (using our bijection) or to k uniform vertices (using its extension to trees with k marked vertices). One thereby recovers the asymptotic *line-breaking construction* of uniformly random trees, proposed by Aldous [4].

These distributional identities also make it possible to compare the distributions of typical distances in different trees. Define a partial order \prec on type sequences $n = (n_c, c \geq 0)$ with $\sum_{c \geq 0} n_c = n$ and $\sum_{c \geq 0} cn_c = n - 1$ by the following covering¹ relation: $m = (m_c, c \geq 0)$ covers $\sum_{c \geq 0} n_c$ if there are positive integers a, b such that $m_a = n_a + 1, m_b = n_b + 1$ and $m_0 = n_0 - 1, m_{a+b} = n_{a+b} - 1$. (In words, to obtain m from n we replace one vertex with $a + b$ children by two vertices, one with a children and one with b children, and reduce the number of leaves accordingly; then $n \prec m$.)

Given type n, m related as in the previous paragraph, we may couple the constructions of random trees with types n and m as follows. Let $V_n = (V_1, \dots, V_{n-1})$ be a uniformly random sequence of elements of $[n]$ subject to the constraint that for each $c \geq 0$, the number of values from $[n]$ which occur precisely c times in the sequence is n_c . Conditionally on V_n , choose $X \in [n]$ uniformly at random from among those integers in $[n]$ which appear exactly d times in V_n , and independently choose Y uniformly at random from among those integers in $[n]$ which do not appear in V_n . Choose a of the instances where X appears in V_n , uniformly at random, and replace each them by the integer Y ; call the resulting sequence V_m . Then the trees T_n and T_m corresponding to V_n and V_m under our bijection are uniformly random elements of $\mathcal{T}(n)$ and $\mathcal{T}(m)$, respectively. Moreover, the index of the first repeated element of V_n is at most that of the first repeated element of V_m , so the distance from the root to the smallest labeled leaf in T_n is at most the corresponding distance in T_m .

It follows from this coupling that for any types n and m with $n \preceq m$, if $T_n \in_u \mathcal{T}(n)$ and $T_m \in_u \mathcal{T}(m)$, and L_n and L_m are uniformly random leaves of T_n and T_m , respectively, then $|L_n|$ is stochastically dominated by $|L_m|$, by which we mean that $\mathbf{P}\{|L_n| \leq t\} \geq \mathbf{P}\{|L_m| \leq t\}$ for all $t \in \mathbb{R}$; we denote this relation by $|L_n| \preceq_{\text{st}} |L_m|$. We conjecture that this stochastic relation also holds for the heights of the trees: writing $\text{ht}(t) := \max\{|v|, v \text{ is a vertex of } t\}$, then

$$\text{ht}(T_n) \preceq_{\text{st}} \text{ht}(T_m),$$

whenever $T_n \in_u \mathcal{T}(n)$ and $T_m \in_u \mathcal{T}(m)$ and $n \preceq m$.

In forthcoming work, Blanc-Renaudie [5] uses the bijection in order to prove convergence toward inhomogeneous continuum random trees for random combinatorial trees with given types, in great generality; he also proves exponential upper tail bounds for the heights of such random trees, which are tight up to constant factors in many cases. In other forthcoming work, Addario-Berry and Donderwinkel [2] use the bijection in order to prove several conjectures from [3] and [11], regarding the asymptotic behaviour of the height of random trees; the latter work [2] in particular shows that *all* random combinatorial trees with n nodes have typical height $O(n^{1/2})$, unless they are extremely “path-like”, possessing

¹ For a partially ordered set (\mathcal{P}, \prec) , $y \in \mathcal{P}$ covers $x \in \mathcal{P}$ if $x \prec y$ and for all $z \in \mathcal{P}$, if $x \preceq z \preceq y$ then $z \in \{x, y\}$.

$n - o(n)$ nodes with exactly one child. Our bijection can also be adapted to study the structure of random graphs which are not trees (see [1, Chapter 1]), and we expect that it will prove useful in proving convergence results in such settings as well.

Growth procedures for rooted trees with a given degree sequence. In this section, we present a growth procedure for rooted trees with a given degree sequence — and in particular for d -ary trees — which is built from a version of the inverse bijection presented in the proof of Theorem 1.1.

Fix an abstract set $\{\ell_i, i \geq 1\}$ of leaf labels. Let $\mathbf{d} = (d_1, \dots, d_m)$ be a sequence of positive integers, and let $L_{\mathbf{d}} = 1 + \sum_{i=1}^m (d_i - 1) \geq 1$. Let $\mathcal{T}_{\mathbf{d}}$ denote the set of rooted trees t with vertex set $[m] \cup \{\ell_1, \dots, \ell_{L_{\mathbf{d}}}\}$ such that for each $i \in [m]$, i has d_i children in t (and $\ell_1, \dots, \ell_{L_{\mathbf{d}}}$ are leaves of t). Also, write

$$\mathcal{S}_{\mathbf{d}} := \{(v_1, \dots, v_{m+L_{\mathbf{d}}-1}) : |\{k : v_k = i\}| = d_i \text{ for all } i \in [m]\}$$

Then the following modification of the inverse bijection from Theorem 1.1 gives a bijection between $\mathcal{S}_{\mathbf{d}}$ and $\mathcal{T}_{\mathbf{d}}$. Fix $\mathbf{v} = (v_1, \dots, v_{m+L_{\mathbf{d}}-1}) \in \mathcal{S}_{\mathbf{d}}$.

Bijection between $\mathcal{S}_{\mathbf{d}}$ and $\mathcal{T}_{\mathbf{d}}$.

- Let $j(0) = 1$, let $j(1) < j(2) < \dots < j(L_{\mathbf{d}} - 1)$ be the locations of the repeats of the sequence \mathbf{v} , and let $j(L_{\mathbf{d}}) = m + L_{\mathbf{d}}$.
- For $i = 1, \dots, L_{\mathbf{d}}$, let P_i be the path $(v_{j(i-1)}, \dots, v_{j(i)-1}, \ell_i)$.
- Let $t(\mathbf{v}) \in \mathcal{T}_{\mathbf{d}}$ have root v_1 and edge set given by the union of the edges of the paths $P_1, P_2, \dots, P_{L_{\mathbf{d}}}$.

In the above bijection, the path P_i ends with the leaf ℓ_i . However, any other fixed ordering of $\ell_1, \dots, \ell_{L_{\mathbf{d}}}$ would also yield a bijective correspondence. In fact, this is even true if the choice of leaf ordering depends on the sequence \mathbf{v} , provided that the leaf-ordering rule still has the property that different sequences result in different trees. We exploit this flexibility below, in order to design a simple growth procedure.

Write \preceq for the total ordering of the vertices of $t(\mathbf{v})$ which is the order the vertices first appear along the paths $P_1, \dots, P_{L_{\mathbf{d}}}$.

For $d \in \mathbb{N}$ and $\mathbf{d}' := (d_1, \dots, d_m, d)$, note that we can construct an element of $\mathcal{S}_{\mathbf{d}'}$ starting from an element $S_{\mathbf{d}} \in \mathcal{S}_{\mathbf{d}}$ by inserting d repeats of integer $m + 1$. (Later, we also write $\mathbf{d}' = (d_1, \dots, d_m, d_{m+1})$, so $d_{m+1} = d$.) This corresponds to choosing $S_{\mathbf{d}} \in \mathcal{S}_{\mathbf{d}}$ and a multiset of size d with elements in $[m + L_{\mathbf{d}}]$, or equivalently, $T_{\mathbf{d}} \in \mathcal{T}_{\mathbf{d}}$ and a multiset of d vertices of $T_{\mathbf{d}}$, where we use the convention that $i \in [m + L_{\mathbf{d}}]$ corresponds to the i^{th} vertex in the total ordering \preceq of the vertices of $T_{\mathbf{d}}$ defined above. The following growth procedure provides a way to sample a uniform random element of $\mathcal{T}_{\mathbf{d}'}$ given $T_{\mathbf{d}} \in_u \mathcal{T}_{\mathbf{d}}$ and an independent uniform random multiset of size d of vertices of $T_{\mathbf{d}}$. It yields the same shape of tree that one would obtain starting from $T_{\mathbf{d}}$ by first transforming the sequence $S_{\mathbf{d}}$ corresponding to $T_{\mathbf{d}}$ under the above bijection as just described, then applying the bijection to the resulting sequence. However, it yields a different labelling of the leaves.

Constructing a tree $T_{\mathbf{d}'}$ from a tree $T_{\mathbf{d}} \in \mathcal{T}_{\mathbf{d}}$ and a multiset of d vertices of $T_{\mathbf{d}}$

Denote the multiset as $\{w_1, \dots, w_d\}$, where $w_1 \preceq w_2 \preceq \dots \preceq w_d$. Then $T_{\mathbf{d}'}$ is obtained from $T_{\mathbf{d}}$ by adding vertices $m+1, \ell_{L_{\mathbf{d}+1}}, \dots, \ell_{L_{\mathbf{d}+d-1}}$ and modifying the edge-set as follows.

- (1) If w_1 was the root of $T_{\mathbf{d}}$, connect $m+1$ to w_1 and re-root the tree at $m+1$. Otherwise, replace the edge vw_1 that connects w_1 to its parent v with two edges $v(m+1)$ and $(m+1)w_1$.
- (2) For $j = 2, \dots, d$,
 - (a) if $w_j = w_{j-1}$, add an edge from $m+1$ to $\ell_{L_{\mathbf{d}+j-1}}$,
 - (b) if $w_j \neq w_{j-1}$, remove the edge from w_j to its parent v , then add edges from v to $\ell_{L_{\mathbf{d}+j-1}}$ and from $m+1$ to w_j .




FIGURE 2. The growth procedure is used to add one internal vertex to the quaternary tree $T_{\mathbf{d}}$. Here, $(w_1, w_2, w_3, w_4) = (3, 3, 1, \ell_4)$.

An example is provided in Figure 2. In the example, the coding sequence of $T_{\mathbf{d}}$ under the bijection presented above is $(2, 2, 3, 2, 4, 4, 1, 1, 2, 1, 3, 4, 3, 4, 1, 3)$. The total ordering \preceq of vertices of $T_{\mathbf{d}}$ is

$$2 \prec \ell_1 \prec 3 \prec \ell_2 \prec 4 \prec \ell_3 \prec 1 \prec \ell_4 \prec \ell_5 \prec \ell_6 \prec \dots \prec \ell_{13}.$$

The chosen multiset $(3, 3, 1, \ell_4)$ of vertices of $T_{\mathbf{d}}$ corresponds to the multiset $(3, 3, 7, 8)$, of elements of $[m + L_{\mathbf{d}}]$, and thence to transforming the coding sequence into

$$(2, 2, \mathbf{5}, \mathbf{5}, 3, 2, 4, 4, \mathbf{5}, 1, \mathbf{5}, 1, 2, 1, 3, 4, 3, 4, 1, 3).$$

The new sequence gives $T_{\mathbf{d}'}$ if we modify the bijection so that the new leaves $\ell_{14}, \ell_{15}, \ell_{16}$ are inserted when the repeats of 5 occur, and the other leaves are inserted according to their original ordering. In general, the growth procedure yields a modified bijection which reads as follows. With \mathbf{d}' as above, fix $v' = (v'_1, \dots, v'_{m+L_{\mathbf{d}'}}) \in \mathcal{S}_{\mathbf{d}'}$, and note that $L_{\mathbf{d}'} = L_{\mathbf{d}} + d - 1$.

Modified bijection between $\mathcal{S}_{\mathbf{d}'}$ and $\mathcal{T}_{\mathbf{d}'}$.

- Let $j(0) = 1$, let $j(1) < j(2) < \dots < j(L_{\mathbf{d}'} - 1)$ be the locations of the repeats of the sequence \mathbf{v}' , and let $j(L_{\mathbf{d}'}) = m + 1 + L_{\mathbf{d}'}$.
- Reorder $\ell_1, \dots, \ell_{L_{\mathbf{d}'}}$ as $\hat{\ell}_1, \dots, \hat{\ell}_{L_{\mathbf{d}'}}$ as follows. For $1 \leq i \leq L_{\mathbf{d}'} - 1$:
 - If $v'_{j(i)}$ is the $(k + 1)$ 'st appearance of $(m + 1)$ then let $\hat{\ell}_i = \ell_{L_{\mathbf{d}'} + k}$.
 - If $v'_{j(i)}$ is the k 'th repeated entry which is not equal to $(m + 1)$ then let $\hat{\ell}_i = \ell_k$.
 Let $\hat{\ell}_{L_{\mathbf{d}'}} = \ell_{L_{\mathbf{d}'}}$.
- For $i = 1, \dots, L_{\mathbf{d}'}$, let P'_i be the path $(v_{j(i-1)}, \dots, v_{j(i)-1}, \hat{\ell}_i)$.
- Let $t'(\mathbf{v}') \in \mathcal{T}_{\mathbf{d}'}$ have root v'_1 and edge set given by the union of the edges of the paths $P'_1, P'_2, \dots, P'_{L_{\mathbf{d}'}}$.

We now describe the inverse of the modified bijection, which takes as input a tree $t' \in \mathcal{T}_{\mathbf{d}'}$ — i.e., a tree t' with vertex set $[m + 1] \cup \{\ell_1, \dots, \ell_{L_{\mathbf{d}'}}\}$ such that for each $i \in [m + 1]$, i has d_i children, and $\ell_1, \dots, \ell_{L_{\mathbf{d}'}}$ are leaves — and outputs a sequence $\mathbf{v}' = (v'_1, \dots, v'_{m+L_{\mathbf{d}'}}) \in \mathcal{S}_{\mathbf{d}'}$. We leave it to the reader to verify that the two procedures are indeed inverses. This justifies the fact that the above growth procedure takes the uniform distribution on $\mathcal{T}_{\mathbf{d}}$ to the uniform distribution on $\mathcal{T}_{\mathbf{d}'}$.

Inverse of the modified bijection

- Let $v'_{m+L_{\mathbf{d}'}}$ be equal to the parent of $\ell_{L_{\mathbf{d}'}}$ in t' . Say that $\ell_{L_{\mathbf{d}'}}$ is used, and that $\ell_1, \dots, \ell_{L_{\mathbf{d}'}-1}$ and $\ell_{L_{\mathbf{d}'}+1}, \dots, \ell_{L_{\mathbf{d}'}}$ are unused.
- For $j = m + L_{\mathbf{d}'} - 1, \dots, 2, 1$:
 - If the number of occurrences of v'_{j+1} in $(v'_{j+1}, \dots, v'_{m+L_{\mathbf{d}'}})$ is equal to the number of children of v'_{j+1} in t' , then let v'_j be equal to the parent of v'_{j+1} in t' .
 - Otherwise, define v'_j as follows.
 - * If $v'_{j+1} = m + 1$, let i^* be the maximum $L_{\mathbf{d}'} + 1 \leq i \leq L_{\mathbf{d}'}$ such that ℓ_i is unused.
 - * If $v'_{j+1} \neq m + 1$, let i^* be the maximum $1 \leq i \leq L_{\mathbf{d}'} - 1$ such that ℓ_i is unused.
 Let v'_j be equal to the parent of ℓ_{i^*} in t' . Say that ℓ_{i^*} is used.

We conclude the paper by justifying the assertion from Section 1 that the growth procedure yields an algorithm for sampling a sequence of random d -ary leaf-labeled trees. For this, let $\mathbf{d}^{(m)} = (d, d, \dots, d)$ have length m . Starting from the unique tree $T^{(1)} \in_u \mathcal{T}_{\mathbf{d}^{(1)}}$, for each $m \geq 1$ let $T^{(m+1)}$ be constructed from $T^{(m)}$ according to the above procedure. Now let T_m be obtained from $T^{(m)}$ by unlabeled the non-leaf vertices; then T_m is a uniformly random leaf-labeled d -ary tree with m internal nodes, and $(T_m, m \geq 1)$ is the random sequence alluded to in the introduction. (In fact, given a multiset of size d of vertices of $T^{(m)}$, its ordering, as defined in the construction of $T^{(m+1)}$ from $T^{(m)}$, does not depend on the labels of the internal vertices. Therefore, it is straightforward to define the pushforward of the growth procedure directly on leaf-labeled trees.) Then, given T_{m+1} , the non-leaf vertex that was added to obtain T_{m+1} from T_m is a uniformly random non-leaf vertex of T_{m+1} . In the binary case, this distinguishes the procedure from Rémy's algorithm, for which, given the shape of T_{m+1} , the non-leaf vertex that was added is instead the parent of a uniformly random leaf of T_{m+1} .

4. ACKNOWLEDGEMENTS

We thank Arthur Blanc-Rénaudie and Nicolas Broutin for useful discussions.

REFERENCES

- [1] Louigi Addario-Berry. Lecture notes for the 2021 CRM-PIMS summer school in probability, 2021. URL <http://problab.ca/louigi/notes/ssprob2021.pdf>. 6
- [2] Louigi Addario-Berry and Serte Donderwinkel. Large degrees yield short trees. in preparation, 2021. 5
- [3] Louigi Addario-Berry, Luc Devroye, and Svante Janson. Sub-Gaussian tail bounds for the width and height of conditioned Galton-Watson trees. *Ann. Probab.*, 41(2): 1072–1087, 2013. ISSN 0091-1798. doi: 10.1214/12-AOP758. URL <https://doi.org/10.1214/12-AOP758>. 5
- [4] David Aldous. The continuum random tree. I. *The Annals of Probability*, 19(1):1–28, 1991. URL <https://doi.org/10.1214/aop/1176990534>. 5
- [5] Arthur Blanc-Rénaudie. Limit of trees with fixed degree sequence. in preparation, 2021. 5
- [6] C.W. Borchardt. Ueber eine der Interpolation entsprechende Darstellung der Eliminations-Resultante. *J. Reine Angew. Math.*, 57:111–121, 1860. ISSN 0075-4102. doi: 10.1515/crll.1860.57.111. URL <https://doi.org/10.1515/crll.1860.57.111>. 1, 4
- [7] Michael Camarri and Jim Pitman. Limit distributions and random trees derived from the birthday problem with unequal probabilities. *Electron. J. Probab.*, 5:no. 2, 18, 2000. ISSN 1083-6489. doi: 10.1214/EJP.v5-58. URL <https://doi.org/10.1214/EJP.v5-58>. 3
- [8] Arthur Cayley. A theorem on trees. *The quarterly journal of pure and applied mathematics*, 23:376–378, 1889. URL http://resolver.sub.uni-goettingen.de/purl?PPN600494829_0023. 4
- [9] Narsingh Deo and Paulius Micikevicius. Prüfer-like codes for labeled trees. In *Proceedings of the Thirty-second Southeastern International Conference on Combinatorics, Graph Theory and Computing (Baton Rouge, LA, 2001)*, volume 151, pages 65–73, 2001. 4
- [10] Tamás Fleiner. On Prüfer codes. Technical report of Egerváry Research Group, TR-2005-16, 2005. URL <http://www.cs.elte.hu/egres/www/tr-05-16.html>. 4
- [11] Svante Janson. Simply generated trees, conditioned Galton-Watson trees, random allocations and condensation. *Probab. Surv.*, 9:103–252, 2012. doi: 10.1214/11-PS188. URL <https://doi.org/10.1214/11-PS188>. 5
- [12] André Joyal. Une théorie combinatoire des séries formelles. *Adv. in Math.*, 42(1):1–82, 1981. ISSN 0001-8708. URL [https://doi.org/10.1016/0001-8708\(81\)90052-9](https://doi.org/10.1016/0001-8708(81)90052-9). 1
- [13] J. W. Moon. Various proofs of Cayley's formula for counting trees. In *A seminar on Graph Theory*, pages 70–78. Holt, Rinehart and Winston, New York, 1967. 1
- [14] E. H. Neville. The codifying of tree-structure. *Proc. Cambridge Philos. Soc.*, 49: 381–385, 1953. ISSN 0008-1981. doi: 10.1017/s030500410002853x. URL <https://doi.org/10.1017/s030500410002853x>. 4
- [15] Jim Pitman. Enumerations of trees and forests related to branching processes and random walks. In *Microsurveys in discrete probability (Princeton, NJ, 1997)*, volume 41 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 163–180. Amer. Math. Soc., Providence, RI, 1998. 3
- [16] Jim Pitman. Coalescent random forests. *J. Combin. Theory Ser. A*, 85(2):165–193, 1999. ISSN 0097-3165. doi: 10.1006/jcta.1998.2919. URL <https://doi.org/10.1006/jcta.1998.2919>. 1

- [17] Jean-Luc Rémy. Un procédé itératif de dénombrement d'arbres binaires et son application à leur génération aléatoire. *RAIRO - Theoretical Informatics and Applications - Informatique Théorique et Applications*, 19(2):179–195, 1985. URL http://www.numdam.org/item/ITA_1985__19_2_179_0/. 1, 4
- [18] Seunghyun Seo and Heesung Shin. A generalized enumeration of labeled trees and reverse Prüfer algorithm. *J. Combin. Theory Ser. A*, 114(7):1357–1361, 2007. ISSN 0097-3165. doi: 10.1016/j.jcta.2007.01.010. URL <https://doi.org/10.1016/j.jcta.2007.01.010>. 4
- [19] Richard P. Stanley. *Enumerative combinatorics. Vol. 2*, volume 62 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1999. ISBN 0-521-56069-1; 0-521-78987-7. doi: 10.1017/CBO9780511609589. URL <https://doi.org/10.1017/CBO9780511609589>. With a foreword by Gian-Carlo Rota and appendix 1 by Sergey Fomin. 3
- [20] Remco van der Hofstad. *Random graphs and complex networks. Vol. 1*. Cambridge Series in Statistical and Probabilistic Mathematics, [43]. Cambridge University Press, Cambridge, 2017. ISBN 978-1-107-17287-6. doi: 10.1017/9781316779422. URL <https://doi.org/10.1017/9781316779422>. 1

DEPARTMENT OF MATHEMATICS AND STATISTICS, MCGILL UNIVERSITY, MONTRÉAL, CANADA

DEPARTMENT OF STATISTICS, OXFORD, UK

Email address: louigi.addario@mcgill.ca

Email address: serte.donderwinkel@st-hughs.ox.ac.uk

Email address: mickael.maazoun@stats.ox.ac.uk

Email address: martin@stats.ox.ac.uk