

RESEARCH ARTICLE

WILEY

Combinatorial and stochastic properties of ranked tree-child networks

François Bienvenu^{1,2} | Amaury Lambert^{1,2} | Mike Steel³

¹Center for Interdisciplinary Research in Biology (CIRB), CNRS UMR 7241, Collège de France, PSL Research University, Paris, France

²Laboratoire de Probabilités, Statistique et Modélisation (LPSM), CNRS UMR 8001, Sorbonne Université, Paris, France

³Biomathematics Research Centre, University of Canterbury, Christchurch, New Zealand

Correspondence

François Bienvenu, Center for Interdisciplinary Research in Biology (CIRB), CNRS UMR 7241, Collège de France, PSL Research University, Paris, France.

Email: francois.bienvenu@normalesup.org

Funding information

Center for Interdisciplinary Research in Biology (CIRB); New Zealand Marsden Fund, Grant/Award Number: UOC1709.

Abstract

Tree-child networks are a class of directed acyclic graphs that have recently risen to prominence in phylogenetics. Although these networks have numerous, attractive mathematical properties, many combinatorial questions concerning them remain intractable. We show that endowing tree-child networks with a biologically relevant ranking structure yields mathematically tractable objects, which we term ranked tree-child networks (RTCNs). We derive explicit enumerative formulas and explain how to sample RTCNs uniformly at random. We study the properties of uniform RTCNs, including: lengths of random walks between root and leaves; distribution of number of cherries in the network; and sampling RTCNs conditional on displaying a given tree. We also formulate a conjecture regarding the scaling limit of the process counting the number of lineages in the ancestry of a leaf. The main idea in this paper, namely using ranking as a way to achieve combinatorial tractability, may also extend to other classes of networks.

KEYWORDS

phylogenetic networks, random graphs, reticulate evolution

1 | INTRODUCTION

Tree-child networks are a class of directed acyclic graphs (DAGs) introduced by [6] as a way to model reticulated phylogenies (i.e., phylogenies that take into account the possibility of hybridization or horizontal gene transfer). In addition to being biologically relevant, tree-child networks are mathematically interesting combinatorial structures and have thus gained attention recently to become one of the

most studied classes of phylogenetic networks. Although they are simpler and more structured than arbitrary phylogenetic networks, they are nevertheless notoriously hard to study. For instance, their enumeration is still an open problem [10, 14] and there is no known algorithm to sample them uniformly (although recursive procedures to enumerate them have recently been introduced [5, 7], as well as asymptotic methods [11]). As a result, relatively little is known about the properties of “typical” tree-child networks.

In this paper, we introduce a new class of phylogenetic networks that we term *ranked tree-child networks*, or RTCNs for short. These networks correspond to a subclass of tree-child networks that are endowed with an additional structure ensuring that they could have resulted from a time-embedded evolutionary process, something that is not required of tree-child networks.

Besides being arguably more biologically relevant than tree-child networks, one of the main advantages of RTCNs is that they are much easier to study. For instance, we show that there exist explicit formulas for the number of leaf-labeled RTCNs, as well as simple procedures to sample them uniformly at random (or even uniformly at random subject to some natural constraints such as containing a fixed number of reticulations, or displaying a given tree). These features make it possible to gain some insight into the structure of uniform RTCNs.

The structure of this paper is as follows. We begin by describing the class of tree-child networks, and then define the new class of RTCNs. We then collect together statements of the main results of this paper. In Section 2, we use forward-in-time and backward-in-time constructions of RTCNs to derive exact results for their enumeration and generation. In Section 3, we describe the probability distribution for the number of “cherries” and “tridents” in a RTCN that has been sampled uniformly at random. Section 5 provides an analysis of the lengths of random walks in a uniform RTCN, both from the root to a leaf, and from a leaf to the root. Finally, in Section 6 we investigate the number of lineages in the ancestry of a leaf. We characterize the Markov chain that counts these lineages as we move away from the leaf, and conjecture that when appropriately rescaled it converges, as the number of leaves of the RTCN goes to infinity, to a smooth function with a random shape parameter.

1.1 | Preliminaries

Let us start by recalling the definition of tree-child networks and introducing some vocabulary.

Definition 1. A binary phylogenetic network is a connected directed acyclic graph which has a single vertex of in-degree 0 and out-degree 2 (the *root*) and where every other vertex has either:

- in-degree 1 and out-degree 0 (the *leaves*)
- in-degree 1 and out-degree 2 (the *tree vertices*)
- in-degree 2 and out-degree 1 (the *reticulation vertices*.)

If V is the vertex set of a binary phylogenetic network, we write ∂V for the set of its leaves. The vertices that are not leaves are called *internal vertices* and we denote their set by $\overset{\circ}{V}$.

We refer to the elements of the set $\Gamma_{\text{in}}(v) = \{u : u \rightarrow v\}$ as the *parents* of v and to that of the set $\Gamma_{\text{out}}(v) = \{u : v \rightarrow u\}$ as the *children* of v .

Finally, an edge \vec{uv} is called a *reticulation edge* if v is a reticulation vertex and a *tree edge* if v is a tree vertex or a leaf.

Definition 2. A tree-child network is a binary phylogenetic network such that every internal vertex has at least one child that is a tree vertex or a leaf.




FIGURE 1 (A) A binary phylogenetic network that fails to be tree-child, since the only path from the vertex indicated by “*” to the leaves passes through a reticulation vertex; (B) a tree-child network that cannot be ranked (i.e., does not correspond to any RTCN); (C) a tree-child network that is not a normal network, as the arc indicated by “*” is redundant (see Definition 6)

Note that there are other simple characterizations of tree-child networks. For example, Lemma 2 in [6] gives the following alternative definition: a binary phylogenetic network is tree-child if and only if for every vertex v there exists a leaf such that every path going from the root to that leaf goes through v . See Figure 1 for examples of tree-child networks and of a binary phylogenetic network that is not a tree-child network.

1.2 | Ranked tree-child networks

First, note that every tree-child network is endowed with a partial order, which we refer to as the *genealogical order*, defined by

$$u \rightsquigarrow v \Leftrightarrow \begin{cases} \text{there exists a directed path from } u \text{ to } v \\ \text{that contains at least one tree-edge.} \end{cases}$$

Let us now introduce the notion of *events* of a tree-child network.

Definition 3. Let N be a tree-child network. Define a relation \mathcal{R} on the set \hat{V} of internal vertices of N by:

$$u \mathcal{R} v \Leftrightarrow u = v \text{ or } u \text{ and } v \text{ are linked by a reticulation edge or share a child.}$$

It is easy to see that, for tree-child networks, \mathcal{R} is an equivalence relation on \hat{V} (note that this is not the case for general networks). We call the equivalence classes of \mathcal{R} the *events* of N . Moreover, writing \bar{u} for the equivalence class of a vertex u ,

- either $\bar{u} = \{u\}$, in which case \bar{u} is called a *branching event*;
- or \bar{u} has three elements, and \bar{u} is called a *reticulation event*.

When there is no risk of confusion, we refer to these events as *branchings* and *reticulations*, for short.

Definition 4. A ranked tree-child network is an ordered pair $(N, <)$ where

- N is a tree-child network.

- The chronological order $<$ is a strict total order on the set of events of N that is compatible with the genealogical order; in other words, for every internal vertices u and v ,

$$u \rightsquigarrow v \Rightarrow \bar{u} < \bar{v}.$$

See Figure 2 for a graphical representation of a ranked tree-child network.

Note that in the case where N is a tree, Definition 4 agrees with the classical notion of ranked tree (see e.g., [13]). Indeed, in that case every internal vertex is its own equivalence class and thus the chronological order can be seen as a total strict order on \hat{V} .

The idea behind the notion of ranking is that the internal vertices of a phylogeny can be associated to evolutionary events. Under the assumption that these events are instantaneous and that no two events can happen at the same time, ranking a phylogeny is a straightforward way to endow it with the information of the order in which these events occurred, without having to specify the actual times at which they did.

The notion of ranking raises two questions:

1. Can all tree-child networks be ranked?
2. How many rankings are there for a given tree-child network?

It is not hard to see that the answer to the first question is no. In fact, we will prove in the next section that almost no tree-child network can be ranked. While this rules out the possibility of using RTCNs to gain insight into the structure of tree-child networks, this does not make RTCNs irrelevant, because the tree-child networks that we expect to see in nature should typically be rankable.¹ Thus, from a biological point of view, RTCNs might actually be more relevant than tree-child networks.

Regarding the second question, it is well-known (see e.g., Chapter 5, Exercice 20 of [12] or Lemma 1 in [18]) that the number of ways to rank a binary tree is

$$\frac{(\ell - 1)!}{\prod_{v \in \hat{V}} (\lambda(v) - 1)},$$

where ℓ is the number of leaves of the tree and $\lambda(v)$ the number of leaves subtended by v —that is, $\lambda(v) = \#\{u \in \partial V : v \rightsquigarrow u\}$. However, the proof of this result relies on the recursive structure of trees, which tree-child networks lack. Thus, counting the number of ways to rank a given tree-child network remains an open question.

1.3 | Relation to other types of phylogenetic networks

Let us start by recalling the definition of temporal networks.

Definition 5. A binary phylogenetic network $N = (V, \vec{E})$ is *temporal* if there exists a (time-stamp) function $t : V \rightarrow \mathbb{R}$ that satisfies the following two properties:

- (i) If \vec{uv} is a reticulation edge then $t(u) = t(v)$.
- (ii) If \vec{uv} is a tree edge then $t(u) < t(v)$.

The underlying tree-child network of a RTCN is always a temporal network, because to get a valid time-stamp function t one can assign to every internal vertex the rank of the corresponding event (i.e.,

¹This is not necessarily the case if some extinct lineages are not observed; but in that case we can always assume that there exists an underlying ranked tree-child network.




FIGURE 2 Graphical representation of a RTCN. Each “L” corresponds to a tree vertex (or the root), each “T” to a reticulation vertex, and the tip of each dangling vertical line to a leaf. The vertical lines represent tree edges, and the horizontal ones events. Note that $a < b$ if and only if the horizontal line representing b is below that of a

$t(u) = k$ where \bar{u} is the k th event, and $t(\rho) = 0$ for the root). Conversely, every temporal tree-child network can be ranked, because one can tweak the time-stamps to make sure that all reticulation vertices have different time-stamps, and then set

$$\bar{u} = \bar{v} \Leftrightarrow t(u) = t(v) \quad \text{and} \quad \bar{u} < \bar{v} \Leftrightarrow t(u) < t(v)$$

to get a valid partition into events and a valid ranking. This gives us the following proposition.

Proposition 1. *A tree-child network is temporal if and only if it is the underlying network of a RTCN.*

Although temporal tree-child networks and RTCNs are closely related, the two should not be mistaken: a RTCN contains more information than the corresponding temporal tree-child network, and a slightly different information than this network together with its time-stamp function (note that there is no requirement that the time-stamps of vertices that belong to different events be distinct in Definition 5). RTCNs should be thought of as being to temporal tree-child networks what ranked trees are to trees.

Let us now recall the notion of normal network, as introduced by [21].

Definition 6. An edge \vec{uv} is said to be *redundant* (or a *shortcut*) if there exists a directed path from u to v that does not contain \vec{uv} (see Figure 1). A normal network is a tree-child network that has no redundant edges.

It is well-known and not too hard to see that every temporal tree-child network is a normal network (see e.g., Proposition 10.12 in [17]). As a result, rankable tree-child networks are normal networks, in the sense that if $(N, <)$ is a RTCN then N is a normal network. Since by Theorem 1.4 in [14] the fraction of tree-child networks (leaf-labeled or vertex-labeled alike) that are also normal networks goes to zero as their number of vertices goes to infinity, this proves the next proposition.

Proposition 2. *Let \tilde{C}_n be the number of tree-child networks with n labeled vertices, and let \tilde{R}_n be the number of those tree-child networks that are rankable. Then, as n goes to infinity, $\tilde{R}_n/\tilde{C}_n \rightarrow 0$.*

Similarly, the fraction R_ℓ/C_ℓ of leaf-labeled tree-child networks that are rankable goes to 0 as ℓ goes to infinity.

1.4 | Main results

All the results presented in this paper concern leaf-labeled RTCNs—that is, the leaves carry labels that uniquely identify them, but the internal vertices do not. Note that the number of reticulations of a RTCN with ℓ leaves and b branchings is $r = \ell - b - 1$.

In Section 2, we give two constructions of RTCNs: one in backward time and one in forward time. Each of these constructions yields a proof of the following result.

Theorem 1. *The number of ranked tree-child networks with ℓ labeled leaves and b branchings is*

$$C_{\ell,b} = \left[\begin{matrix} \ell - 1 \\ b \end{matrix} \right] T_{\ell},$$

where $T_{\ell} = \ell! (\ell - 1)! / 2^{\ell-1}$ is the number of ranked trees with ℓ labeled leaves and $\left[\begin{matrix} \ell-1 \\ b \end{matrix} \right]$ is the number of permutations of $\{1, \dots, \ell - 1\}$ with b cycles (these quantities are known as the unsigned Stirling numbers of the first kind).

In particular, Theorem 1 implies that the number of branchings of a uniform RTCN with ℓ labeled leaves is asymptotically Poisson with mean $\sim \log(\ell)$, and thus satisfies a central limit theorem (see Corollary 3).

The backward-time and forward-time constructions also provide simple procedures to sample leaf-labeled RTCNs, be it:

- uniformly at random;
- uniformly at random conditional on their total number of reticulations;
- uniformly at random conditional on which events are reticulations.

The rest of our study focuses on the properties of the uniform distribution on the set of RTCNs with ℓ labeled leaves. One of the interesting characteristics of uniform leaf-labeled RTCNs is their close relationship with uniform leaf-labeled ranked trees. This is detailed in Section 3, where we also explain how to sample a uniform RTCN conditional on displaying a given tree (i.e., on “containing” the tree, in the sense of Definition 12).

Two of the most basic statistics of binary phylogenetic networks are their numbers of “cherries” and of “tridents” (we will define these precisely later, but briefly, a cherry is a vertex having two leaves as children; a trident is a reticulation event with three leaves as children). While almost nothing is known about the distribution of these two quantities in uniform tree-child networks, they prove very tractable in uniform RTCNs. Explicit expressions for their mean and variance are given in Section 4, where we also establish the following theorem.

Theorem 2. *Let κ_{ℓ} be the number of cherries of a uniform RTCN with ℓ labeled leaves and χ_{ℓ} be its number of tridents. Then, as $\ell \rightarrow \infty$,*

$$(i) \quad \kappa_{\ell} \xrightarrow{d} \text{Poisson} \left(\frac{1}{4} \right).$$

$$(ii) \quad \chi_{\ell} / \ell \xrightarrow{\mathbb{P}} \frac{1}{7}.$$

Sections 5 and 6 contain our most informative results about the structure of uniform RTCNs: in Section 5 we study the length of some typical paths joining the root to the leaf set of uniform RTCNs and prove the following theorem.




FIGURE 3 In red, the ancestry of a leaf and on the right the process counting the number of lineages in this ancestry

Theorem 3. Let ν be a uniform RTCN with ℓ labeled leaves and let

- γ^\downarrow be the path taken by a random walk going from the root of ν to its leaves, respecting the direction of the edges and choosing each of the two outgoing edges with equal probability when it reaches a tree-vertex.
- γ^\uparrow be the path taken by a random walk going from a uniformly chosen leaf of ν to its root, following the edges in reverse direction and choosing each of the two incoming edges with equal probability when it reaches a reticulation.

Then, letting $\text{length}(\cdot)$ denote the length of these paths, not counting reticulation edges, there exist two constants c^\downarrow and c^\uparrow such that, as $\ell \rightarrow \infty$,

- $\text{length}(\gamma^\downarrow) \approx \text{Poisson}(2 \log \ell + c^\downarrow)$
- $\text{length}(\gamma^\uparrow) \approx \text{Poisson}(3 \log \ell + c^\uparrow)$

in the sense that the total variation distance between these distributions goes to 0.

Remark 1. The proof of Theorem 3 can be adapted to the Yule model, which corresponds to uniform ranked trees. In that case, $\text{length}(\gamma^\downarrow) \approx \text{Poisson}(\log \ell + a)$ and $\text{length}(\gamma^\uparrow) \approx \text{Poisson}(2 \log \ell + b)$, for suitable constants a, b .

Finally, in Section 6 we study the number of lineages in the ancestry of a leaf, that is, in the subgraph consisting of all paths joining this leaf to the root, as illustrated in Figure 3 and defined formally in Definition 17. Starting from the leaves and going towards the root, one event of the RTCN after the other, this number of lineages will on average start by increasing but will eventually decrease as we get nearer the root and the total number of lineages of the RTCN itself decreases. A biological motivation for considering this concept is that an extant species (leaf) that has many lineages in its past is likely to have acquired more independent traits than a species with fewer ancestral lineages.

Let $X^{(\ell)} = (X_k^{(\ell)}, 0 \leq k \leq \ell - 2)$ denote the process counting these lineages. What can we say about the behavior of $X^{(\ell)}$ as ℓ goes to infinity? As it turns out, $X_k^{(\ell)}$ remains quite small as long as $k = \lfloor x\ell \rfloor$ with $0 < x < 1$. However, when we reach the regime $k = \lfloor \ell - x\sqrt{\ell} \rfloor$, $X_k^{(\ell)}$ starts to increase substantially, until a point where it will decrease suddenly. Moreover, even though the properly rescaled trajectories of $X^{(\ell)}$ seem to become smooth as $\ell \rightarrow \infty$, they do not become deterministic. In fact,




FIGURE 4 500 superimposed trajectories of the process $\tilde{X}^{(\ell)}$, rescaled so that their domain is $[0, 1]$ and that the total area under each curve is equal to 1

simulations presented in Section 6.2 suggest that they behave as a deterministic process with a random initial condition.

Another natural way to study the asymptotics of $X^{(\ell)}$, instead of restricting ourselves to an appropriate time-window, is to consider the embedded process $\tilde{X}^{(\ell)}$, that is, to ignore the steps in which $X^{(\ell)}$ does not change. As illustrated in Figure 4, simulations suggest that the scaling limit of $\tilde{X}^{(\ell)}$ is a remarkably regular deterministic process with a random shape parameter.

We were not able to complete the proof of the convergence of these rescaled processes, so our discussion relies in part on simulations and heuristics. Below are some of our results and our main conjecture.

Proposition 3. For all $\ell \geq 2$ and k such that $0 \leq k \leq \ell - 2$,

$$\frac{\ell}{\ell - k + 1} \left(1 - \frac{2k}{(\ell - k)(\ell - k - 1)} \right) \leq \mathbb{E} \left(X_k^{(\ell)} \right) \leq \frac{\ell}{\ell - k}.$$

As a result, for $M > 2$ and all ℓ large enough, for all $\epsilon > 0$,

$$(1 - \epsilon) \frac{1}{M} \left(1 - \frac{2}{M^2} \right) \leq \mathbb{E} \left(\frac{1}{\sqrt{\ell}} X_{\lfloor \ell - M\sqrt{\ell} \rfloor}^{(\ell)} \right) \leq \frac{1}{M},$$

so that the sequence of random variables $\frac{1}{\sqrt{\ell}} X_{\lfloor \ell - M\sqrt{\ell} \rfloor}^{(\ell)}$ is tight and bounded away from 0 in L^1 .

Conjecture 1. There exists a random variable $W_M > 0$ such that

$$\frac{1}{\sqrt{\ell}} X_{\lfloor \ell - M\sqrt{\ell} \rfloor}^{(\ell)} \xrightarrow[\ell \rightarrow \infty]{d} W_M.$$

Proposition 4. If Conjecture 1 holds, then for all ϵ such that $0 < \epsilon < 1$, as $\ell \rightarrow \infty$,

$$\left(\frac{1}{M\sqrt{\ell}} X_{\lfloor \ell - M\sqrt{\ell}(1-t) \rfloor}^{(\ell)}, t \in [0, 1 - \epsilon] \right) \Rightarrow (y(t, C_M), t \in [0, 1 - \epsilon])$$

where \Rightarrow denotes convergence in distribution in the Skorokhod space,

$$y(t, C_M) = \frac{1 - t}{C_M \cdot (1 - t)^2 + 1}$$

and $C_M = M/W_M - 1$.




FIGURE 5 A RTCN with its canonical labeling and the associated process $(P_k, 1 \leq k \leq \ell)$

2 | COUNTING AND GENERATING RTCNS

2.1 | Backward-time construction of RTCNs

Let us start by recalling that there is a simple way to label the internal vertices of any leaf-labeled tree-child network (or any leaf-labeled DAG), namely by labeling each internal vertex with the set of labels of its children.

Definition 7. Given a tree-child network with leaf set $\partial V = \{1, \dots, \ell\}$, the associated canonical labeling is the function ξ such that

- (i) $\forall v \in \partial V, \xi(v) = v$.
- (ii) $\forall v \in \mathring{V}, \xi(v) = \{\xi(u) : v \rightarrow u\}$.

While the canonical labeling of a tree-child network encodes it unambiguously (the whole network can be recovered from the label of the root), this is not the case for RTCNs because the information about the order of the events is missing. One way to retain this information is to encode RTCNs using the process $(P_k, 1 \leq k \leq \ell)$ defined as follows.

First, given a set $P = \{\xi_1, \dots, \xi_m\}$ of labels of vertices of $(N, <)$, define the two operations:

- $\text{coal}(P, \{\xi_i, \xi_j\}) = (P \setminus \{\xi_i, \xi_j\}) \cup \{\{\xi_i, \xi_j\}\}$
- $\text{ret}(P, \xi_i, \{\xi_j, \xi_k\}) = (P \setminus \{\xi_i, \xi_j, \xi_k\}) \cup \{\{\xi_i\}, \xi_j\}, \{\{\xi_i\}, \xi_k\}\}$

Now, let $U_1 < \dots < U_{\ell-1}$ denote the events of $(N, <)$. Then, starting from $P_1 = \{1, \dots, \ell\}$ and going backwards in time, for $k = 1$ to $\ell - 1$:

- If $U_{\ell-k}$ is the coalescence of v and w , that is, if there exists u such that $U_{\ell-k} = \{u\}$ and $\Gamma_{\text{out}}(u) = \{v, w\}$, then let $P_{k+1} = \text{coal}(P_k, \{\xi(v), \xi(w)\})$.
- If $U_{\ell-k}$ is the reticulation of u and v with the hybrid h , that is, if $U_{\ell-k} = \{u', h', v'\}$ with $\Gamma_{\text{out}}(u') = \{u, h'\}$, $\Gamma_{\text{out}}(v') = \{v, h'\}$ and $\Gamma_{\text{out}}(h') = \{h\}$, then let $P_{k+1} = \text{ret}(P_k, \xi(h), \{\xi(u), \xi(v)\})$.

The result of procedure is illustrated in Figure 5. Note that the RTCN that produced a process $(P_k, 1 \leq k \leq \ell)$ can unambiguously be recovered from that process.

In order to count RTCNs based on their number of reticulations (or, equivalently, branchings, since both are linked by $r + b = \ell - 1$), we need to introduce the notion of profile of a RTCN.

Definition 8. Let ν be a RTCN and let $U_1 < \dots < U_{\ell-1}$ denote its events. The *profile* of ν is the vector $\mathbf{q} = (q_1, \dots, q_{\ell-1})$ defined by

$$q_k = \begin{cases} 1, & \text{if } U_k \text{ is a branching;} \\ 0, & \text{otherwise.} \end{cases}$$

Note that we always have $q_1 = 1$.

We are now in position to prove our main result concerning the enumeration of RTCNs.

Theorem 1. *The number of RTCNs with profile \mathbf{q} is*

$$F(\mathbf{q}) = \prod_{k=1}^{\ell-1} (q_k + (k-1)(1-q_k)) T_{\ell},$$

where $T_{\ell} = \ell! (\ell-1)! / 2^{\ell-1}$. As a result, the number of RTCNs with ℓ labeled leaves and b branchings is

$$C_{\ell,b} = \left[\begin{matrix} \ell-1 \\ b \end{matrix} \right] T_{\ell},$$

where the bracket denotes the unsigned Stirling numbers of the first kind.

Proof. Let us count the processes $(P_k, 1 \leq k \leq \ell)$ with profile \mathbf{q} . When going from $k+1$ to k lineages, that is, when considering event U_k , there are:

- $\binom{k+1}{2}$ possible coalescence events;
- $(k-1) \binom{k+1}{2}$ possible reticulation events.

Therefore,

$$F(\mathbf{q}) = \prod_{k=1}^{\ell-1} \left(\binom{k+1}{2} q_k + (k-1) \binom{k+1}{2} (1-q_k) \right).$$

Factoring out the binomial coefficients, we get the first part of the theorem. Then, summing over all profiles with b branchings,

$$C_{\ell,b} = \left(\sum_{|\mathbf{q}|=b} \prod_{k=1}^{\ell-1} (q_k + (k-1)(1-q_k)) \right) \times T_{\ell}.$$

The first factor can be seen to be the coefficient of degree b of the polynomial $P(X) = (X+0)(X+1) \dots (X+\ell-2)$. Since by definition of the unsigned Stirling numbers of the first kind,

$$P(X) = \sum_{b=0}^{\ell-1} \left[\begin{matrix} \ell-1 \\ b \end{matrix} \right] X^b,$$

this concludes the proof. ■

Theorem 1 and its proof lead to several further results, which we now list.

Corollary 1. For $\ell \geq 2$, the number C_ℓ of RTCNs with ℓ labeled leaves is

$$C_\ell = \frac{\ell! (\ell - 1)!^2}{2^{\ell-1}}.$$

Proof. This follows from the expression for $C_{\ell,b}$ in Theorem 1 upon applying the identity $\sum_{b=0}^{\ell-1} \binom{\ell-1}{b} = (\ell - 1)!$ ■

Corollary 2. The following procedure yields a uniform RTCN with profile \mathbf{q} . Starting from ℓ labeled lineages, for $k = \ell - 1$ down to 1: If $q_k = 1$: let two lineages coalesce, uniformly at random among all $\binom{k+1}{2}$ possibilities. If $q_k = 0$: let three lineages reticulate, uniformly at random among all $(k - 1) \binom{k+1}{2}$ possibilities.

Proposition 5. Starting from ℓ labeled lineages, let pairs of lineages and triplets of lineages reticulate, choosing what to do uniformly among all possibilities at each step and stopping when there is only one lineage left. Then, the resulting RTCN has the uniform distribution on the set of RTCNs with ℓ labeled leaves.

Proof. Proposition 5 follows by noting that the realization of the procedure correspond to those of the process $(P_k, 1 \leq k \leq \ell)$, which uniquely encodes every leaf-labeled RTCN. ■

Finally, let us point out the following fact about the profile and the number of branchings of uniform RTCNs.

Corollary 3. Let \mathbf{q} be the profile of a uniform RTCN with ℓ labeled leaves. Then,

$$\mathbf{q} \sim (X_1, \dots, X_{\ell-1}),$$

where $(X_1, \dots, X_{\ell-1})$ are independent Bernoulli variables such that

$$\mathbb{P}(X_k = 1) = \frac{1}{k}.$$

As a result, the number B_ℓ of branchings of a uniform RTCN with ℓ leaves, which is distributed as the number of cycles of a uniform permutation of $\{1, \dots, \ell - 1\}$, satisfies

- (i) $\mathbb{E}(B_\ell) = H_{\ell-1}$, where $H_{\ell-1} = \sum_{k=1}^{\ell-1} 1/k$ is the $(\ell - 1)$ th harmonic number.
- (ii) As $\ell \rightarrow \infty$, $d_{TV}(B_\ell, \text{Poisson}(H_{\ell-1})) \rightarrow 0$. In particular, $\frac{B_\ell - \log \ell}{\sqrt{\log \ell}} \xrightarrow{d} \mathcal{N}(0, 1)$.

Proof. The first part of the proposition follows from the fact that the steps of the algorithm described in Proposition 5 are independent and that, when going from $k + 1$ to k lineages there are $(k - 1) \binom{k+1}{2}$ possible reticulations and $\binom{k+1}{2}$ possible coalescences, so that choosing uniformly among those the probability of picking a coalescence is $1/k$.

Points (i) and (ii) for B_ℓ are classic properties of the distribution of the number of cycles in a uniform permutation—see for instance Section 3.1 of [15]—that follow easily from its representation as a sum of independent Bernoulli variables. Indeed, (i) is immediate and for (ii) we can use the

Stein–Chen bound on the total variation distance between a sum of independent Bernoulli variables and the corresponding Poisson distribution (recalled as Theorem A in Section 5) to get

$$d_{\text{TV}}(B_\ell, \text{Poisson}(H_{\ell-1})) \leq \min\{1, 1/H_{\ell-1}\} \sum_{k=1}^{\ell-1} \frac{1}{k^2} = O((\log \ell)^{-1}),$$

from which the central limit theorem follows readily (see e.g., [2, p. 17]). ■

Finally, let us close this section by pointing out an unexpected connection between RTCNs and a combinatorial structure known as “river-crossings.”

Remark 2. The number of RTCNs with ℓ labeled leaves is also the number of river-crossings using a two-person boat. It is recorded as sequence A167484 in the Online Encyclopedia of Integer Sequences [1], where it is described as follows:

For ℓ people on one side of a river, the number of ways they can all travel to the opposite side following the pattern of 2 sent, 1 returns, 2 sent, 1 returns, \dots , 2 sent.

We could not find a natural bijection between river-crossings and RTCNs, and thought that there would not be one: indeed, for $\ell = 3$ the $C_3 = 6$ river-crossings are completely equivalent up to permutation of the labels, while the 6 RTCNs are not (3 of them contain a reticulation while 3 of them do not). However, while this paper was under revision we were informed by Michael Fuchs (pers. comm.) that his group found one such bijection, which they will describe in a forthcoming paper [4].

2.2 | Forward-time construction of RTCNs

In this section, we give a forwards-in-time construction of RTCNs that will yield a second proof of Theorem 1. This proof is more combinatorial than the one we have already given, and will provide a different intuition as to why Stirling numbers arise in the enumeration of RTCNs.

Here we introduce a further definition: Two vertices are said to be *siblings* if they have a parent in common, and *step-siblings* if they have a sibling in common (thus, in a tree-child network two vertices are step-siblings if and only if they are the tree-vertex children of the two parents of a reticulation).

The following notion will be useful.

Definition 9. A decorated RTCN is a pair (ν, θ) , where

- (i) ν is a RTCN with vertex set V and root ρ .
- (ii) The *decoration* is a function $\theta : V \setminus \{\rho\} \rightarrow \{0, 1\}$ such that
 - If ν is a reticulation vertex or the child of a reticulation vertex, $\theta(\nu) = 0$.
 - If u and v are siblings or step-siblings and neither u nor v is a reticulation vertex, then $\theta(u) = 1 - \theta(v)$.

Note that this formal definition is just a way to say that:

- For every tree vertex ν , we distinguish one of the outgoing edges by assigning a “1” to one of the children of ν .

- For every reticulation vertex v , we distinguish one of the incoming edges by assigning a “1” to one of the siblings of v (this unambiguously determines their common parent u and therefore the incoming edge \vec{uv}).

The notion of a decorated RTCN is similar in spirit to that of an *ordered ranked tree*, where an ordering is specified for the children of each vertex—and indeed both notions are equivalent for trees. However, when working with RTCNs it is not useful to order the children of every vertex, in part because reticulated vertices already play a special role.

An important characteristic of decorated RTCNs is that, unlike their undecorated counterparts, they are intrinsically labeled. Indeed, it is possible to assign a unique label to every vertex of a decorated RTCN (ν, θ) as follows: first, remove every reticulation edge distinguished by the decoration. This yields a tree $\tilde{\nu}$. Then, assign to each vertex v the label $\theta(u) \dots \theta(v)$, where (ρ, u, \dots, v) is the unique path going from the root ρ to v in $\tilde{\nu}$. A consequence of this unique labeling of internal vertices is that there are $\ell!$ ways to label the leaves of a decorated RTCN with ℓ leaves.

Before describing the forward-time construction of decorated RTCNs, let us introduce one last combinatorial object.

Definition 10. A *subexcedant sequence of length n* is an integer-valued sequence $s = (s_1, \dots, s_n)$ such that, for all k , $1 \leq s_k \leq k$. For any two subexcedant sequences s and s' , the *number of encounters of s and s'* is defined as

$$\text{enc}(s, s') = \# \{k \geq 1 : s_k = s'_k\}$$

Lemma 1. For any subexcedant sequence s of length n , there are $\binom{n}{k}$ subexcedant sequences s' of length n such that $\text{enc}(s, s') = k$.

This classic result follows immediately by considering a uniform subexcedant sequence s' and noting that its number of encounters with s is distributed as the sum for $k = 1$ to n of independent Bernoulli variables with parameters $1/k$.

Let us now describe how to encode decorated RTCNs using subexcedant sequences. Let s° and s^\bullet be two subexcedant sequences of length $\ell - 1$. Start from a single lineage indexed “1” and, for $k = 1$ to $\ell - 1$:

- If $s_k^\circ = s_k^\bullet$, then let lineage s_k° branch to create lineage $k + 1$.
- If $s_k^\circ \neq s_k^\bullet$, then let lineages s_k° and s_k^\bullet hybridize to form lineage $k + 1$.

At each step of this procedure, decorate the lineage s_k° with a white dot. This construction is illustrated in Figure 7.

Observe that in this construction every pair of subexcedant sequences (s^\bullet, s°) will yield a different decorated RTCN ν° , and that given a decorated RTCN ν° it is possible to unequivocally recover the subexcedant sequences (s^\bullet, s°) that produced it. Thus, letting $\mathcal{S}_{\ell-1}$ denote the set of subexcedant of length $\ell - 1$ and \mathcal{C}_ℓ° that of decorated RTCNs with ℓ unlabeled leaves, this construction gives a bijection between $\mathcal{S}_{\ell-1} \times \mathcal{S}_{\ell-1}$ and \mathcal{C}_ℓ° .

Also note that if at every step of the procedure we only link lineage $k + 1$ to lineage s_k° and decorate the latter with a white dot, we get a bijection between $\mathcal{S}_{\ell-1}$ and the set \mathcal{T}_ℓ° of decorated ranked trees with ℓ unlabeled leaves.

From this forward-time encoding of decorated RTCNs and Lemma 1, we immediately get the following proposition.




FIGURE 6 An example of a decorated RTCN (left) with its simplified representation (right). In the simplified representation, decorated lineages are indicated as follows: for branchings, they are drawn to the right; for reticulations, we add a white dot to indicate which parent of the reticulation points to a 1




FIGURE 7 Graphical representation of the forward-time construction of RTCNs. On the left, the black and white dots represent the subexcedant sequences s^* and s° , respectively, and on the right the corresponding RTCN. At each step, the new lineage (marked by a cross) is linked to the black dot and the white dot. If these two dots fall on the same lineage, we get a branching

Proposition 6. *There are $\binom{\ell-1}{b} (\ell-1)!$ decorated RTCNs with ℓ unlabeled leaves and b branchings.*

To recover Theorem 1 from Proposition 6, it suffices to recall that there are $\ell!$ ways to label the leaves of a decorated RTCN and to note that there are $2^{\ell-1}$ ways to decorate a RTCN.

Finally, let us close this section by pointing out that the following simple stochastic process generates uniform RTCNs.

Proposition 7. *Starting from a single lineage representing the root, let every lineage branch at rate 1 and every ordered pair of lineages hybridize at rate 1, decorating a lineage when it branches and decorating the first vertex of an ordered pair of lineages when it hybridizes. The RTCN obtained by stopping upon reaching ℓ lineages is uniform on the set of decorated RTCN with ℓ leaves. Relabeling its leaves uniformly at random and discarding its decoration yields a uniform RTCN with ℓ labeled leaves.*




FIGURE 8 The operations of unreticulation and reticulation

3 | RTCNS AND RANKED TREES

The forward-time construction of the previous section gave us a way to encode a decorated RTCN using a decorated ranked tree and a subexcedant sequence. In fact, Theorem 1 shows that it is also possible to encode an *undecorated* leaf-labeled RTCN using a leaf-labeled ranked tree and, for example, a subexcedant sequence or a permutation—even though we were unable to find any meaningful such encoding.

In this section, we specialize the results of the previous section in order to explain how to obtain RTCNs from ranked trees using simple graph operations, without having to explicitly manipulate subexcedant sequences. In particular, we give a simple way to sample a uniform RTCN conditional on displaying a given ranked tree. Readers who are more interested in the structural properties of RTCNs than in how to sample them can jump to the next section.

3.1 | Reticulation, unreticulation, and base trees

Let us recall that the classic graph operation known as the *vertex identification* of u and v consists in replacing u and v by a new vertex w whose neighbors (in-neighbors and out-neighbors, respectively) are exactly the neighbors of u and those of v (without introducing a self-loop in the case where u and v were neighbors). Conversely, the *subdividing* of \vec{uv} consists in introducing an intermediate vertex w between u and v , that is: adding w to the graph, removing \vec{uv} and adding \vec{uw} and \vec{wv} .

Definition 11. The *unreticulation* of a reticulation edge \vec{uv} is the graph operation consisting in

1. Removing \vec{uv} .
2. Identifying u and its (now only) child.
3. Identifying v and its child.

Conversely, the *reticulation* of a tree edge e_1 into a tree edge e_2 consists in

1. Subdividing e_1 and e_2 .
2. Adding an edge from the vertex introduced in the middle of e_1 to that introduced in the middle of e_2 .

These operations are illustrated in Figure 8.

Unreticulating edges of a tree-child network N yields a tree N' with the same number of leaves as N as well as a partition into events and a genealogical order that are compatible with those of N . This justifies the following definition.

Definition 12. A ranked tree τ is called a *base tree* of a RTCN ν if it is possible to obtain it by unreticulating edges of ν . In that case, we write $\tau \sqsubset \nu$ and say that ν *displays* τ .

Note that since each reticulation vertex has two incoming reticulating edges and that unreticulating each of these yields different RTCNs, every leaf-labeled RTCN with r reticulations displays exactly 2^r ranked trees. This is also a consequence of the fact that an RTCN is a binary normal network, and each such network displays exactly 2^r trees (Corollary 3.4 of [22]).

3.2 | Uniform base trees of a uniform RTCN

Proposition 8. *The ranked tree obtained by unreticulating one of the two incoming edges of each reticulation vertex of a uniform RTCN with ℓ labeled leaves, each with probability $1/2$, has the uniform distribution on the set of ranked trees with ℓ labeled leaves.*

Proof. Let ν be a uniform leaf-labeled RTCN. The procedure described in the proposition amounts to:

1. Decorating ν uniformly at random to obtain a decorated leaf-labeled RTCN ν° .
2. Unreticulating each of the undecorated reticulation edges of ν° to produce a decorated leaf-labeled tree τ° .
3. Discarding the decoration of τ° .

In the forward-time encoding, ν° corresponds to a unique triplet $\nu^\circ \simeq (s^\bullet, s^\circ, \sigma)$, where the permutation σ represents the leaf labeling. Now, since ν° is uniform on the set of decorated leaf-labeled RTCNs, (s°, σ) is also uniform and as a result so is $\tau^\circ \simeq (s^\circ, \sigma)$. Finally, the tree obtained by forgetting the decoration of a uniform decorated leaf-labeled ranked tree is uniform on the set of leaf-labeled ranked trees, concluding the proof. ■

3.3 | Sampling RTCNs conditional on displaying a tree

To obtain a RTCN from a ranked tree, we need to pay attention to the constraints imposed by the chronological order. Indeed, it is not possible to reticulate any tree edge of a RTCN into any other tree edge and obtain a RTCN. To formulate this restriction, we need to introduce the notion of contemporary edges.

Definition 13. We now introduce a relation on the set V of vertices of an RTCN. Let us write $\bar{u} \preceq \bar{v}$ to indicate that $\bar{u} < \bar{v}$ or $\bar{u} = \bar{v}$. In addition, if u is a leaf then set $\bar{u} = \partial V$ and $\bar{v} < \bar{u}$ for any internal vertex v .

The edge \vec{uv} is said to be *alive between two events* $U < U'$ if it meets the two following conditions:

- (i) It is a tree edge.
- (ii) $\bar{u} \preceq U$ and $U' \preceq \bar{v}$.

Two edges are said to be *contemporary* if there exist two events such that they are both alive between these events.

As for events, these definitions become very intuitive when using the graphical representation of RTCNs. Recall that, in this representation, we think of the vertical axis as time, and that tree edges correspond to vertical lines while events correspond to horizontal ones. An edge is alive between two events if a portion of it is located in the horizontal strip of the plane delimited by the two events.




FIGURE 9 A RTCN with its events numbered in increasing order (and with the convention that $U_7 = \partial V$). The red edges are the four edges that are alive between events U_3 and U_4

Two edges are contemporary if they overlap when projected on the vertical axis. This is illustrated in Figure 9.

Note that if we let $U_1 < \dots < U_{\ell-1}$ denote the events of a RTCN with ℓ leaves, with the convention that $U_\ell = \partial V$, then there are exactly $k + 1$ edges alive between U_k and U_{k+1} .

Recall the Definition 8 of a profile.

Proposition 9. Let τ be a ranked tree with ℓ labeled leaves and let \mathbf{q} be a profile. Letting $U_1 < \dots < U_{\ell-1}$ denote the branchings of τ , for $k = 1$ to $\ell - 1$:

- If $q_k = 1$, do nothing.
- If $q_k = 0$, letting u denote the vertex such that $U_k = \{u\}$:
 1. Pick an edge e uniformly at random among the two outgoing edges of u .
 2. Pick an edge e' uniformly at random among the $k - 1$ edges that are contemporary with e and do not contain u .
 3. Reticulate e' into e .

Then, the resulting RTCN ν has the uniform distribution on the set of RTCNs with profile \mathbf{q} displaying τ . Moreover, if τ is uniform then ν is uniform on the set of RTCNs with profile \mathbf{q} . If in addition \mathbf{q} is distributed as the profile of a uniform RTCN (see Corollary 3), then the resulting RTCN is uniform on the set of RTCNs with ℓ labeled leaves.

Proof. The first part of the proposition follows from the fact (1) that for every fixed ranked tree θ and for every RTCN μ displaying θ there is exactly one sequence of modifications of θ that yields μ , and (2) that each possible realization of the procedure described in the proposition has the same probability, namely

$$\alpha(\mathbf{q}) = \prod_{k=1}^{\ell-1} (q_k + 2(k-1)(1-q_k))^{-1},$$

although the exact value of this probability does not matter here. Thus, letting $Q(\mu)$ denote the profile of μ ,

$$\mathbb{P}(\nu = \mu | \tau = \theta) = \alpha(\mathbf{q}) \mathbb{1}_{\{Q(\mu)=\mathbf{q}, \theta \sqsubset \mu\}}.$$

Now let us assume that τ is uniform. We have to show that for any RTCN μ ,

$$\mathbb{P}(v = \mu) = \frac{1}{F(\mathbf{q})} \mathbb{1}_{\{Q(\mu)=\mathbf{q}\}},$$

where $F(\mathbf{q})$ is the number of RTCNs with profile \mathbf{q} , whose exact value does not matter here. Since $\mathbb{P}(\tau = \theta) = 1/T_\ell$ for all fixed ranked tree θ ,

$$\mathbb{P}(v = \mu) = \sum_{\theta \sqsubset \mu} \alpha(\mathbf{q}) \cdot \frac{1}{T_\ell} \cdot \mathbb{1}_{\{Q(\mu)=\mathbf{q}\}}.$$

Now let $r = \sum_{k=1}^{\ell-1} (1 - q_k)$ denote the number of reticulations of the profile \mathbf{q} . Since every RTCN with profile \mathbf{q} displays exactly 2^r ranked trees, we have

$$\mathbb{P}(v = \mu) = \frac{2^r \alpha(\mathbf{q})}{T_\ell} \cdot \mathbb{1}_{\{Q(\mu)=\mathbf{q}\}},$$

which concludes the proof since the factor $2^r \alpha(\mathbf{q})/T_\ell$ does not depend on μ . ■

To close this section, note that to generate all RTCNs displaying a ranked tree it suffices to apply the procedure above without the restrictions of the profile, and that there are then $2(k - 1) + 1$ possible actions at step k of the procedure. This gives us the following proposition counting the number of RTCNs displaying a ranked tree.

Proposition 10. *The number of RTCNs displaying any ranked tree τ is*

$$\#\{v : \tau \sqsubset v\} = \prod_{k=1}^{\ell-1} (2k - 1) = (2\ell - 3)!!$$

Remark 3. Surprisingly, the right-hand side of the expression in Proposition 10 is also the number of rooted unranked binary trees with ℓ labeled leaves—see e.g., [13] or [17]. While it is easy to give a bijective proof of this, we have not found a “graphical” bijection that would make this intuitive.

4 | CHERRIES AND TRIDENTS

Cherries and tridents are among the most basic statistics of tree-child networks. In biological terms, a cherry is a pair of nonhybrid sibling species and a trident is a group of three extant species such that one of these species was produced by the hybridization of the two others. These notions can be formalized as follows.

Definition 14. An event is said to be *external* when

$$\bigcup_{u \in U} \Gamma_{\text{out}}(u) \subset \partial V,$$

where $\Gamma_{\text{out}}(u)$ denotes the set of children of vertex u and ∂V is the leaf set of N .

- A *cherry* is an external branching event.

- A *trident* is an external reticulation event.

These notions are illustrated in Figure 11.

4.1 | Number of cherries

In this section, we prove the first part of what was announced in the introduction as Theorem 2.

Theorem 2 (Part (i)). *Let κ_ℓ be the number of cherries of a uniform RTCN with ℓ labeled leaves. Then,*

$$\kappa_\ell \xrightarrow[\ell \rightarrow \infty]{d} \text{Poisson}\left(\frac{1}{4}\right).$$

The proof consists in coupling κ_ℓ with a Markov chain to show that the factorial moments of κ_ℓ (i.e., $m_\ell^n := \mathbb{E}(\kappa_\ell(\kappa_\ell - 1) \cdots (\kappa_\ell - n + 1))$) satisfy the recursion:

$$m_{\ell+1}^n = \left(\frac{\ell - 2n}{\ell}\right)^2 m_\ell^n + \frac{n(\ell - 2n + 2)}{\ell^2} m_\ell^{n-1},$$




FIGURE 10 Example of the random construction of a RTCN from a ranked tree described in Proposition 9. Here, nothing happens for U_1 and U_3 . For U_2 and U_4 , the black dot represents the edge e and the white one the edge e' . Note that the modifications of branching events can be performed in any order, but that they have to be performed sequentially, so that contemporary edges remain well-defined at every step of the procedure




FIGURE 11 A RTCN with its only cherry (in blue) and its two tridents (in red)

and then using this recursion to prove that $m_\ell^n \rightarrow 1/4^n$ as $\ell \rightarrow \infty$ for all n . We separate each of these steps into different propositions.

Proposition 11. *Let $(X_\ell)_{\ell \geq 2}$ be the Markov chain defined by:*

- (i) $X_2 = 1$.
- (ii) For $\ell \geq 2$, conditional on $X_\ell = k$,

$$X_{\ell+1} = \begin{cases} k-2, & \text{with probability } \frac{2k(2k-2)}{\ell^2}; \\ k-1, & // \frac{2k+4k(\ell-2k)}{\ell^2}; \\ k, & // \frac{(\ell-2k)(\ell-2k-1)+2k}{\ell^2}; \\ k+1, & // \frac{\ell-2k}{\ell^2}. \end{cases}$$

Then, for all $\ell \geq 2$, X_ℓ is distributed as the number κ_ℓ of cherries of a uniform RTCN with ℓ labeled leaves.

Proof. Let us use the forward-time construction to build a nested sequence $(v_\ell^\circ)_{\ell \geq 2}$ of uniform decorated RTCNs. Recall from Section 2.2 that to go from v_ℓ° to $v_{\ell+1}^\circ$, we choose an ordered pair (u, v) of lineages v_ℓ° uniformly at random. If $u = v$, we let the lineage branch; otherwise we let u and v hybridize.

Now, assume that v_ℓ° has k cherries, so that there are $2k$ lineages that belong to cherries (which we refer to as C leaves) and $\ell - 2k$ lineages that do not (non-C leaves). Let us list all possible choices of (u, v) and see their effect on the number of cherries.

1. If the next event is a branching:

- (i) If a C leaf is chosen, the number cherries does not change. This happens with probability $2k/\ell^2$.
- (ii) If a non-C leaf is chosen, one cherry is created (probability: $(\ell - 2k)/\ell^2$).

2. If the next event is a hybridization:

- (i) If the two leaves of a same cherry are chosen, that cherry is destroyed (probability: $2k/\ell^2$).
- (ii) If two leaves of two different cherries are chosen, these two cherries are destroyed (probability: $2k(2k - 2)/\ell^2$).
- (iii) If a C leaf and a non-C leaf are chosen, one cherry is destroyed (probability: $4k(\ell - 2k)/\ell^2$).
- (iv) If two non-C leaves are chosen, the number of cherries does not change (probability: $(\ell - 2k)(\ell - 2k - 1)/\ell^2$).

Doing the book-keeping and observing that a RTCN with two leaves has one cherry concludes the proof. ■

Notation. We denote by $x^n = \prod_{k=0}^{n-1} (x - k)$ the n th falling factorial of x .

Proposition 12. *Let $(X_\ell)_{\ell \geq 2}$ be the Markov chain defined in Proposition 11 and let $m_\ell^n = \mathbb{E}(X_\ell^n)$ denote the n th factorial moment of X_ℓ . Then,*

$$m_{\ell+1}^n = \left(\frac{\ell - 2n}{\ell} \right)^2 m_\ell^n + \frac{n(\ell - 2n + 2)}{\ell^2} m_\ell^{n-1}.$$

Proof. Let $p_{(-2)}, p_{(-1)}, p_{(+0)}$ and $p_{(+1)}$ denote the transition probabilities of X_ℓ , conditional on $X_\ell = k$. Then,

$$\begin{aligned} \mathbb{E} \left(X_{\ell+1}^n | X_\ell = k \right) &= (k-2)^n p_{(-2)} + (k-1)^n p_{(-1)} + k^n p_{(+0)} + (k+1)^n p_{(+1)} \\ &= k^n \left(\frac{(k-n)(k-n-1)}{k(k-1)} p_{(-2)} + \frac{k-n}{k} p_{(-1)} + p_{(+0)} + \frac{k+1}{k-n+1} p_{(+1)} \right) \\ &= \frac{k^n}{\ell^2} (4(k-n)(k-n-1) + 2(k-n)(1+2(\ell-2k)) \\ &\quad + (\ell-2k)(\ell-2k-1) + 2k + \frac{(k+1)(\ell-2k)}{k-n+1}). \end{aligned}$$

After a little algebra to rearrange this last expression, we get

$$\mathbb{E} \left(X_{\ell+1}^n | X_\ell = k \right) = k^n \left(\frac{\ell-2n}{\ell} \right)^2 + k^{n-1} \frac{n(\ell+2-2n)}{\ell^2},$$

and the proposition follows by integrating in k . ■

Proposition 13.

(i) For all $\ell \geq 2$, $\mathbb{E}(\kappa_\ell) = \frac{(3\ell-5)\ell}{12(\ell-1)(\ell-2)}$.

(ii) For all $n \geq 1$, $\mathbb{E}(\kappa_\ell^n) \rightarrow 1/4^n$ as $\ell \rightarrow \infty$. As a result, $\kappa_\ell \xrightarrow[\ell \rightarrow \infty]{d} \text{Poisson}(1/4)$.

Proof. The expression for $\mathbb{E}(\kappa_\ell)$ given in (i) follows from Lemma A.1 and routine calculations. To prove (ii), we proceed by induction on n . First, we see from the expression in (i) that $m_\ell^1 \rightarrow 1/4$ as $\ell \rightarrow \infty$. Now suppose that $n \geq 2$ and that $m_\ell^{n-1} \rightarrow 1/4^{n-1}$ as $\ell \rightarrow \infty$. By Proposition 12,

$$m_{\ell+1}^n = a_\ell m_\ell^n + b_\ell,$$

where

$$a_\ell = \left(\frac{\ell-2n}{\ell} \right)^2 \quad \text{and} \quad b_\ell = \frac{n(\ell-2n+2)}{\ell^2} m_\ell^{n-1}.$$

Since $a_\ell \neq 0$ for all $\ell \geq 2n+1$, application of Lemma A.1 together with the identity

$$\prod_{j=2n+1}^k \left(\frac{j-2n}{j} \right)^2 = \frac{1}{\binom{k}{2n}^2}$$

gives

$$m_\ell^n = \left(m_{2n+1}^n + \sum_{k=2n+1}^{\ell-1} \frac{n(k-2n+2)}{k^2} \binom{k}{2n}^2 m_k^{n-1} \right) / \binom{\ell-1}{2n}^2.$$

Now, as $k \rightarrow \infty$, $\binom{k}{2n} \sim k^{2n}/(2n)!$ and, by the induction hypothesis, $m_k^{n-1} \sim 1/4^{n-1}$. As a result,

$$\frac{n(k-2n+2)}{k^2} \binom{k}{2n}^2 m_k^{n-1} \sim \frac{n/4^n}{(2n)!^2} k^{4n-1}.$$

Using Lemma A.2 to get an asymptotic equivalent of the sum of these terms,

$$m_\ell^n \sim \frac{n/4^{n-1}}{(2n)!^2} \cdot \frac{\ell^{4n}}{4n} \cdot \left(\frac{(2n)!}{\ell^{2n}}\right)^2 = 1/4^n.$$

The convergence in distribution of X_ℓ to a Poisson distribution with mean $1/4$ is then a classic result (see e.g., Theorem 2.4 in [20]). ■

4.2 | Number of tridents

In this section, we prove the second part of Theorem 2.

Theorem 2 (Part (ii)). *Let χ_ℓ be the number of tridents of a uniform RTCN with ℓ labeled leaves. Then,*

$$\frac{\chi_\ell}{\ell} \xrightarrow[\ell \rightarrow \infty]{\mathbb{P}} \frac{1}{7}.$$

The proof is quite similar to that used in the previous section to study the number κ_ℓ of cherries – namely, we couple χ_ℓ with a Markov chain in order to compute its moments. The difference is that the moments of χ_ℓ are not as tractable as those of κ_ℓ . As a result, we only compute the first two moments explicitly and then use Chebyshev’s inequality to prove the convergence in probability.

Proposition 14. *Let $(X_\ell)_{\ell \geq 2}$ be the Markov chain defined by*

- (i) $X_2 = 0$
- (ii) For $\ell \geq 2$, conditional on $X_\ell = k$,

$$X_{\ell+1} = \begin{cases} k-1, & \text{with probability } \frac{3k(3k-2)}{\ell^2}; \\ k, & // \quad 1 - \frac{3k(3k-2) + (\ell-3k)(\ell-3k-1)}{\ell^2}; \\ k+1, & // \quad \frac{(\ell-3k)(\ell-3k-1)}{\ell^2}. \end{cases}$$

Then, for all $\ell \geq 1$, X_ℓ has the distribution of the number of tridents of a uniform RTCN with ℓ leaves.

Proof. As in the proof of Proposition 11, let us consider a nested sequence of uniform decorated RTCNs $(v_\ell^\circ)_{\ell \geq 2}$ produced by the forward-time construction, and see how the number of tridents is affected when we go from v_ℓ° to $v_{\ell+1}^\circ$.

Assuming that there are k tridents in χ_ℓ , there are $3k$ leaves associated with tridents (RC leaves) and $\ell - 3k$ other leaves (non-RC leaves). The possible cases are the following:

1. The next event is a branching:
 - (i) If a non-RC leaf is chosen, the number of tridents does not change. This happens with probability $(\ell - 3k)/\ell^2$.
 - (ii) If a RC leaf is chosen, the corresponding trident is destroyed (probability: $3k/\ell^2$).
2. The next event is a reticulation:
 - (i) If two non-RC leaves are chosen, one trident is created (probability: $(\ell - 3k)(\ell - 3k - 1)/\ell^2$).

- (ii) If a non-RC leaf and a RC leaf are chosen, one trident is destroyed and one is created (probability: $6k(\ell - 3k)$).
- (iii) If two RC leaves are chosen:
 - a. If they do not belong to two different tridents, these are destroyed and a new trident is created (probability: $3k(3k - 3)/\ell^2$).
 - b. If they belong to the same trident, this trident is destroyed and another one is created (probability: $6k/\ell^2$).

Doing the book-keeping and observing that a RTCN with two leaves has zero tridents yields the proposition. ■

Proposition 15. *The expected number $\mu_\ell = \mathbb{E}(\chi_\ell)$ of tridents of a uniform RTCN with ℓ leaves μ_ℓ satisfies the recursion*

$$\mu_{\ell+1} = \left(\frac{\ell-3}{\ell}\right)^2 \mu_\ell + \frac{\ell-1}{\ell}.$$

As a result, we have $\mu_2 = 0$, $\mu_3 = 1/2$, $\mu_4 = 2/3$ and, for $\ell \geq 4$,

$$\mu_\ell = \frac{(15\ell^3 - 85\ell^2 + 144\ell - 71)\ell}{105(\ell-1)(\ell-2)(\ell-3)}.$$

Proof. Using the Markov chain of Proposition 14, we have

$$\begin{aligned} \mathbb{E}(X_{\ell+1}|X_\ell = k) &= k + \mathbb{P}(X_{\ell+1} = k + 1|X_\ell = k) - \mathbb{P}(X_{\ell+1} = k - 1|X_\ell = k) \\ &= k \left(\frac{\ell-3}{\ell}\right)^2 + \frac{\ell-1}{\ell}, \end{aligned}$$

and the recursion follows by integrating against the law of X_ℓ .

The expression of μ_ℓ then follows from Lemma A.1 and calculations that are better performed by a symbolic computation software such as [19]. ■

Proposition 16. *The variance of the number of tridents of a uniform RTCN with ℓ leaves is*

$$\text{Var}(\chi_\ell) = \frac{24}{637} \ell + \frac{1}{21} + o(1).$$

The proof of this proposition is exactly the same as that of Proposition 15 but involves more complex expressions that can be found in Section B of the Appendix.

Finally, the convergence in probability of χ_ℓ/ℓ to $1/7$ follows readily from Chebyshev's inequality and the fact that $\mathbb{E}(\chi_\ell) \sim \ell/7$ and $\text{Var}(\chi_\ell) = O(\ell)$.

5 | RANDOM PATHS BETWEEN THE ROOT AND THE LEAVES

In this section, we study the length of two random paths going from the root to the leaf set:

1. A path obtained by starting from the root and going “down” towards the leaves, choosing each outgoing edge with equal probability whenever we reach a tree vertex.

2. A path obtained by starting from a uniformly chosen leaf and going “up” towards the root, choosing each incoming edge with equal probability whenever we reach a reticulation vertex.

Definition 15. The *length* of a path γ is its number of tree edges.

The reason why we do not count reticulation edges when calculating the length of a path is that, from a biological point of view, reticulation edges are supposed to correspond to “instantaneous” hybridization events. As will become apparent from their proof, our results can be adapted to the case where all edges are counted. However, this yields a compound Poisson distribution as the approximation distribution for the lengths of the paths (see e.g., Lemma 8 in [3]).

Before starting with the proofs, let us introduce some notation.

Notation. We denote by

$$H_n^{(m)} = \sum_{k=1}^n \frac{1}{k^m}$$

the n th generalized harmonic number of order m . We also use the notation $\gamma = \lim_n H_n^{(1)} - \log n$ for the Euler–Mascheroni constant, where as in the rest of this text “log” denotes the natural logarithm.

Finally, let us recall a classic bound on the total variation distance between a sum of independent Bernoulli variables and the corresponding Poisson distribution.

Theorem A. Let X_1, \dots, X_n be independent Bernoulli variables with parameters $\mathbb{P}(X_i = 1) = p_i$, and let $\lambda_n = \sum_{i=1}^n p_i$. Then,

$$d_{\text{TV}} \left(\sum_{i=1}^n X_i, \text{Poisson}(\lambda_n) \right) \leq \min(1, 1/\lambda_n) \sum_{i=1}^n p_i^2,$$

where d_{TV} denotes the total variation distance.

This inequality is a consequence of the Stein–Chen method and can be found, for example, as Theorem 4.6 in [16].

5.1 | Length of a random walk from the root to a leaf

In this section, we prove the first part of what was announced as Theorem 3 in the introduction.

Theorem 3 (Point (i)). Let ν be a uniform RTCN with ℓ leaves, and let γ^\downarrow be a random path obtained by starting from the root and following the edges of ν , choosing each of the two out-going edges of a tree vertex with equal probability and stopping when we reach a leaf. Then,

$$\text{length}(\gamma^\downarrow) = \sum_{k=1}^{\ell-1} I_k,$$

where $I_1, \dots, I_{\ell-1}$ are independent Bernoulli variables with parameter

$$\mathbb{P}(I_k = 1) = \frac{2k-1}{k^2}.$$

In particular, letting $c^\downarrow = 2g - \pi^2/6$, where g is the Euler–Mascheroni constant,

- (i) $\mathbb{E}(\text{length}(\gamma^\downarrow)) = 2 \log \ell + c^\downarrow + o(1)$.
- (ii) $\text{Var}(\text{length}(\gamma^\downarrow)) = 2 \log \ell + O(1)$.
- (iii) $d_{\text{TV}}(\text{length}(\gamma^\downarrow), \text{Poisson}(2 \log \ell + c^\downarrow)) \rightarrow 0$ as $\ell \rightarrow \infty$.

Proof. The idea of the proof is to use the forward-time construction to build jointly a nested sequence $(\nu_\ell^\circ)_{\ell \geq 2}$ of uniform decorated RTCNs and the random path γ^\downarrow . With the convention that ν_1° consists of a single lineage, for $k \geq 2$ let (u_k, v_k) denote the pair of lineages that was chosen to turn ν_{k-1}° into ν_k° (recall that if $u_k = v_k$ then the next event is a branching) and let x_k record the position of the random walk among the leaves of ν_{k-1}° . With this notation, the length of γ^\downarrow in ν_ℓ° is

$$\text{length}(\gamma^\downarrow) = \sum_{k=1}^{\ell-1} \mathbb{1}_{\{x_k \in \{u_k, v_k\}\}},$$

where the variables $\mathbb{1}_{\{x_k \in \{u_k, v_k\}\}}$ are independent because (x_{k-1}, x_k) is independent of (u_k, v_k) . Moreover, since (u_k, v_k) is chosen uniformly among the pairs of lineages of ν_{k-1}° and independently of x_k ,

$$\mathbb{P}(x_k \in \{u_k, v_k\}) = \frac{2k-1}{k^2},$$

which proves the first part of the proposition.

The rest of the proposition follows immediately from Theorem A since, letting $p_k = (2k-1)/k^2$,

- $\mathbb{E}(\text{length}(\gamma^\downarrow)) = \sum_{k=1}^{\ell-1} p_k = 2H_{\ell-1}^{(1)} - H_{\ell-1}^{(2)}$.
- $\text{Var}(\text{length}(\gamma^\downarrow)) = \sum_{k=1}^{\ell-1} p_k(1-p_k) = 2H_{\ell-1}^{(1)} - 5H_{\ell-1}^{(2)} + 4H_{\ell-1}^{(3)} - H_{\ell-1}^{(4)}$.
- $\sum_{k=1}^{\ell-1} p_k^2 = 4H_{\ell-1}^{(2)} - 4H_{\ell-1}^{(3)} + H_{\ell-1}^{(4)} = O(1)$. ■

5.2 | Length of a random walk from a leaf to the root

In this section, we prove the second part of Theorem 3.

Theorem 3 (Point (ii)). *Let ν be a uniform RTCN with ℓ leaves, and let γ^\uparrow be a random path obtained by starting from a uniformly chosen leaf and following the edges of ν in reverse direction, choosing each of the two incoming edges of a reticulation vertex with equal probability and stopping when we reach the root. Then,*

$$\text{length}(\gamma^\uparrow) = \sum_{k=2}^{\ell} J_k,$$

where J_2, \dots, J_ℓ are independent Bernoulli variables with parameter

$$\mathbb{P}(J_k = 1) = \frac{3k-4}{k(k-1)}.$$

In particular, letting $c^\uparrow = 3g - 4$, where g is the Euler–Mascheroni constant,

- (i) $\mathbb{E}(\text{length}(\gamma^\uparrow)) = 3 \log \ell + c^\uparrow + o(1)$.
- (ii) $\text{Var}(\text{length}(\gamma^\uparrow)) = 3 \log \ell + O(1)$.
- (iii) $d_{\text{TV}}(\text{length}(\gamma^\uparrow), \text{Poisson}(3 \log \ell + c^\uparrow)) \rightarrow 0$.

Remark 4. Note that the random path γ^\uparrow is not uniformly chosen among all the paths going from the focal leaf to the root.

Proof. The proof is similar to that of the previous section, but this time the idea is to use the backward-time construction to jointly build the RTCN ν and the random path γ^\uparrow . Recall that, in the backward-time construction, for $k = \ell$ down to 2, we go from k to $k - 1$ lineages by choosing an event uniformly at random among the $k(k - 1)/2$ possible coalescences and $k(k - 1)(k - 2)/2$ possible reticulations. Out of these, $k - 1$ coalescences and $3(k - 1)(k - 2)/2$ reticulations involve the lineage through which γ^\uparrow goes, and the choice is independent of the position of γ^\uparrow . As a result, the probability that the lineage containing γ^\uparrow is involved in the event that is chosen is

$$\frac{k - 1 + 3(k - 1)(k - 2)/2}{k(k - 1)/2 + k(k - 1)(k - 2)/2} = \frac{3k - 4}{k(k - 1)},$$

proving the first part of the proposition. The rest of the proposition then follows from Theorem A and from the fact that, letting $p_k = \frac{3k-4}{k(k-1)}$,

- $\mathbb{E}(\text{length}(\gamma^\uparrow)) = \sum_{k=2}^{\ell} p_k = 3H_{\ell}^{(1)} - 4 + 3/\ell.$
- $\text{Var}(\text{length}(\gamma^\uparrow)) = \sum_{k=2}^{\ell} p_k(1 - p_k) = 3H_{\ell}^{(1)} + 20 - 17H_{\ell}^{(2)} - 7/\ell + 1/\ell^2.$
- $\sum_{k=2}^{\ell} p_k^2 = -17H_{\ell}^{(2)} + 24 - 10/\ell + 1/\ell^2 = O(1).$

■

5.3 | An alternative proof of Theorem 3

In this section, we give another proof of Theorem 3. This proof is less direct than the previous one, but it provides an alternative intuition as to where the Poisson distribution, the $\log \ell$ order of magnitude and the factors 2 and 3 come from.

Because writing down this proof formally would require introducing additional notation, and because we already have a formal proof, we allow ourselves to present it as a heuristic. In what follows, the symbol “ \approx ” will be used loosely to indicate that two quantities are approximately equal.

Let us start with γ^\downarrow . Slowing-down time in Proposition 7, consider the uniform decorated RTCN with ℓ leaves ν° obtained by:

1. Starting from one lineage.
2. Conditional on there being k lineages, letting:
 - each lineage branch at rate $1/k$;
 - each ordered pair of lineages hybridize at rate $1/k$.
3. Stopping upon reaching ℓ lineages.

Note that in this construction a branching event is viewed as the production of a new particle by another, rather than as the splitting of a particle into two new particles. Thus, we can consider the path $\tilde{\gamma}^\downarrow$ obtained by always following the edge corresponding to the lineage of the first particle, as illustrated in Figure 12.

From the forward-time joint construction of ν° and γ^\downarrow , we see that the distribution of γ^\downarrow does not depend on which lineage it chooses to follow, as long as this choice only depends on the past of the process. As a result,

$$\text{length}(\gamma^\downarrow) \stackrel{d}{=} \text{length}(\tilde{\gamma}^\downarrow).$$




FIGURE 12 On the left, the path γ^\downarrow and on the right the path $\tilde{\gamma}^\downarrow$ corresponding to the lineage of the first particle

Now, $\text{length}(\tilde{\gamma}^\downarrow)$ is simply the number of events affecting the lineage of the first particle, and the rate at which each given lineage is affected by events is $(k - 1)\frac{2}{k} + \frac{1}{k} \approx 2$. Therefore, conditional on the time T it takes for the process to reach ℓ lineages,

$$\text{length}(\tilde{\gamma}^\downarrow) \approx \text{Poisson}(2T).$$

Finally, the total number of lineages increases by 1 at rate $k \cdot \frac{1}{k} + k(k - 1) \cdot \frac{1}{k} = k$ and therefore follows a Yule process $(Y(t), t \geq 0)$. Since as $t \rightarrow \infty, Y(t)e^{-t} \rightarrow W$ almost surely, where W is a random variable (namely, an exponential variable with parameter 1), we see that the random time T it takes for the process to reach ℓ lineages is asymptotically

$$T \approx \log \ell - \log W.$$

Putting the pieces together, we see that $\text{length}(\gamma^\downarrow) \approx \text{Poisson}(2 \log \ell)$.

Let us now give a similar, forward-in-time construction of γ^\uparrow where it can be identified with the lineage of the first particle. For this, we need to “straighten” γ^\uparrow thanks to a set of deterministic rules telling us how to fix each bend, as illustrated in Figure 13. This yields a $\tilde{\nu}^\circ = f(\nu^\circ, \gamma^\uparrow)$ in which $\tilde{\gamma}^\uparrow$, the image of γ^\uparrow , is the lineage of the first particle.

Having done this, $\text{length}(\gamma^\uparrow) = \text{length}(\tilde{\gamma}^\uparrow)$ and, conditional on $\tilde{\nu}^\circ$, the length of $\tilde{\gamma}^\uparrow$ is the number of events affecting the lineage of the first particle.

Now, $\tilde{\nu}^\circ$ has the same distribution as the decorated RTCN $\hat{\nu}^\circ$ generated by the following continuous-time Markov chain:

- Each lineage branches at rate $1/k$, except for that of the first particle, which branches at rate $2/k$.
- Each ordered pair of lineages hybridizes at rate $1/k$, except for pairs involving the lineage of the first particle, which hybridize at rate $3/(2k)$.

This equality in distribution is proved by writing down explicitly the laws $\tilde{\nu}^\circ$ and $\hat{\nu}^\circ$, something that requires introducing notation to give a formal description of $\tilde{\nu}^\circ$. However, to see where the factors 2 and $3/2$ for the lineage of the first particle come from in the construction of $\hat{\nu}^\circ$, it suffices to note that when going from $k - 1$ to k lineages the probability that the next event involves the lineage of the first particle is




FIGURE 13 Illustration of the deterministic procedure used to “straighten” the path γ^\uparrow in order to make it coincide with the lineage of the first particle. Left: the local modifications that are made each time a branching or a reticulation is encountered; these essentially consist of swapping lineages. Right: an example of application of the procedure to a RTCN

$$\frac{3(k-2)+2}{3(k-2)+2+(k-2)(k-3)+(k-2)} = \frac{3k-4}{k(k-1)},$$

and so is indeed the same as the probability the lineage containing γ^\uparrow is involved when going from k to $k-1$ lineages in the joint backward-time construction of ν° and γ^\uparrow given in Section 5.2.

In the construction of $\hat{\nu}^\circ$, conditional on there being k lineages, events affect the lineage of the first particle at rate $2(k-1) \cdot \frac{3}{2k} + \frac{2}{k} \approx 3$ so that letting \hat{T} denote the random time it takes for the process to reach ℓ lineages, $\text{length}(\gamma^\uparrow) \approx \text{Poisson}(3\hat{T})$. Finally, since the total number of lineages increases by 1 at rate $k+1$ when there are k lineages, the total number of lineages is distributed as $\hat{Y}(t) - 1$, where $(\hat{Y}(t), t \geq 0)$ is a Yule process started from 2, so that $\hat{Y}(t)e^{-t} \rightarrow W + W'$, where W and W' are independent exponential variables with parameter 1. Therefore, $\hat{T} \approx \log \ell$ and we recover $\text{length}(\gamma^\uparrow) \approx \text{Poisson}(3 \log \ell)$.

6 | NUMBER OF LINEAGES IN THE ANCESTRY OF A LEAF

Let us start by giving a formal definition of the process counting the number of lineages in the ancestry of a leaf that was described in Section 1.4.

6.1 | Definition and characterization of $X^{(\ell)}$

Definition 16. The *ancestry* of a vertex u of a RTCN ν is the subgraph $\nu|_u$ consisting of all paths going from the root of ν to u .

Definition 17. The number of lineages of a subgraph μ of a RTCN ν is the process $(X_k, 0 \leq k \leq \ell - 2)$ defined by

$$X_k = \#\{e \in \mu : \text{the edge } e \text{ is alive between } U_k \text{ and } U_{k+1}\},$$

where $U_1 > \dots > U_{\ell-1}$ are the events of ν , taken in inverse chronological order, with the same convention as in Definition 13 of alive edges that $U_0 = \partial V$.




FIGURE 14 List of all the possible types of events and their effect on the number of marked lineages

In the rest of this section, we study the number of lineages in the ancestry of a uniformly chosen leaf of a uniform RTCN with ℓ labeled leaves, and denote it by $X^{(\ell)}$. See Figure 3 in Section 1.4 for an illustration. We also study the embedded process $\tilde{X}^{(\ell)}$ defined by $\tilde{X}_i = X_{k_i}$, where $k_0 = 0$ and, for $i \geq 1$, $k_i = \inf \{k > k_{i-1} : X_k \neq X_{k_{i-1}}\}$.

Let us start by characterizing the law of $X^{(\ell)}$.

Proposition 17. *The process $X_k^{(\ell)}$ is the Markov chain characterized by $X_0^{(\ell)} = 1$ and the transition probabilities:*

- $\mathbb{P}\left(X_{k+1}^{(\ell)} = x + 1 \mid X_k^{(\ell)} = x\right) = \frac{x(\ell - k - x)(\ell - k - x - 1)}{(\ell - k)(\ell - k - 1)^2}$
- $\mathbb{P}\left(X_{k+1}^{(\ell)} = x - 1 \mid X_k^{(\ell)} = x\right) = \frac{x(x-1)^2}{(\ell - k)(\ell - k - 1)^2}$
- $\mathbb{P}\left(X_{k+1}^{(\ell)} = x \mid X_k^{(\ell)} = x\right) = 1 - \mathbb{P}\left(X_{k+1}^{(\ell)} = x \pm 1 \mid X_k^{(\ell)} = x\right)$

Proof. The proof relies on the backward construction of a uniform RTCN and a bit of book-keeping to see how the $(k + 1)$ th event, which takes us from $\ell - k$ lineages to $\ell - k - 1$ lineages, affects the number of lineages in the ancestry of a leaf. Recall that in the backward construction there are $(\ell - k)(\ell - k - 1)^2/2$ possibilities for the $(k + 1)$ th event. Let us refer to the lineages in the ancestry of the focal leaf as marked lineages. Conditional on $X_k^{(\ell)} = x$, there are x marked lineages and $\ell - k - x$ unmarked lineages and thus there are:

- $x(x - 1)/2$ possible coalescences between marked lineages. These decrease the number of lineages by 1.
- $x(x - 1)(x - 2)/2$ reticulations involving only marked lineages. These also decrease the number of lineages by 1.
- $x(\ell - k - x)(\ell - k - x - 1)/2$ possible reticulations where the hybrid is a marked lineage and the other two lineages are unmarked. These increase the number of marked lineages by 1.

Other types of events (coalescences between two unmarked lineages, coalescences between a marked and an unmarked lineage, etc.) leave the number of marked lineages unchanged, as illustrated in Figure 14. Since the event is chosen uniformly among all possibilities, this concludes the proof. ■

6.2 | Simulations

In this section, we present simulations supporting the conjectures about $X^{(\ell)}$ and $\tilde{X}^{(\ell)}$ made in Section 1.4 and outline some ideas to approach these conjectures. Let us start by looking at some individual trajectories of these processes, for increasing values of ℓ . As can be seen in Figure 15A, most of the interesting behavior of $X^{(\ell)}$ seems to happen very close to the root so that to obtain a

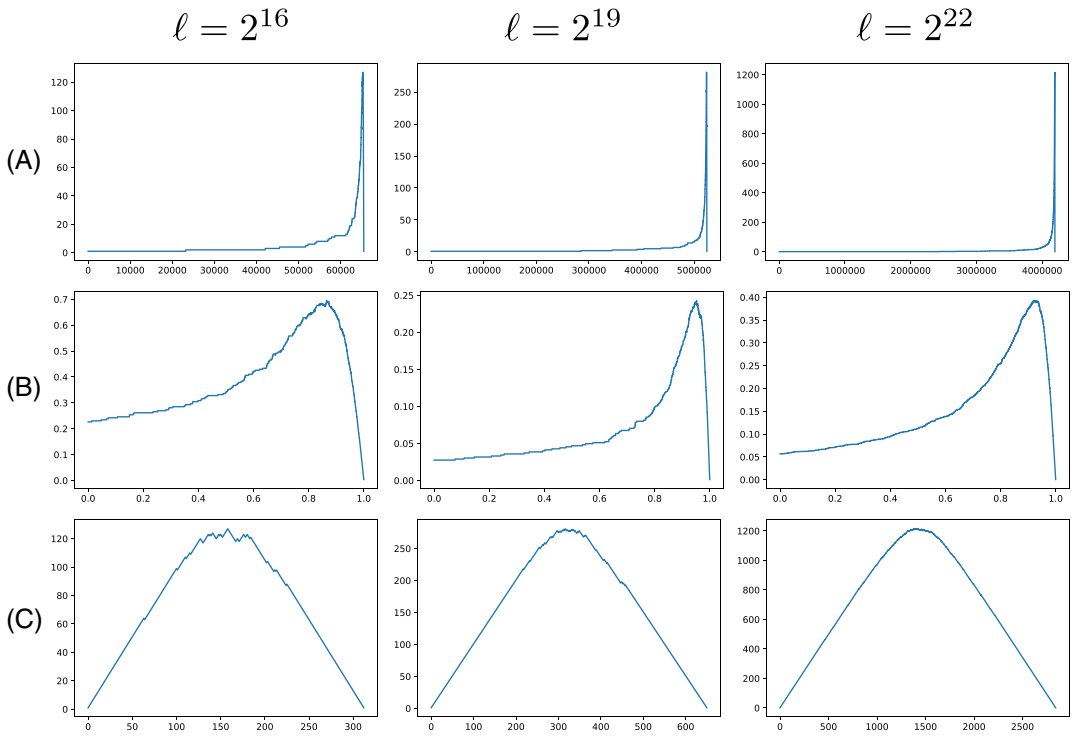


FIGURE 15 Individual trajectories of the processes described in the main text, for different values of ℓ . A, the process $X^{(\ell)}$, B the process $t \mapsto \frac{1}{\sqrt{\ell}} X_{\lfloor \ell - M\sqrt{\ell}(1-t) \rfloor}^{(\ell)}$ and C, the process $\tilde{X}^{(\ell)}$

nondegenerate scaling-limit, we need to focus on a small window of time, for instance by considering $X_k^{(\ell)}$ for $k = \lfloor \ell - M\sqrt{\ell}(1-t) \rfloor$ and $t \in [0, 1]$.

Even though the trajectories represented in Figure 15 seem to become smooth as $\ell \rightarrow \infty$, they do not become deterministic, as made apparent by Figure 16, where the distributions of some statistics of $X^{(\ell)}$ are given. In particular, these simulations suggest that the relevant scaling limit for X^ℓ is indeed $\frac{1}{\sqrt{\ell}} X_{\lfloor \ell - M\sqrt{\ell}(1-t) \rfloor}^{(\ell)}$, and support Conjecture 1.

Our idea to approach the study of the process $X^{(\ell)}$ is to separate it into two phases:

1. A slow, stochastic phase, where up to time $k = \lfloor \ell - M\sqrt{\ell} \rfloor$ the process $X_k^{(\ell)}$ remains relatively small and highly stochastic
2. A fast, deterministic phase, during which the internal dynamics of the rescaled process become deterministic, but retain a trace of the stochasticity of the first phase in the form of random initial conditions.

6.3 | The stochastic phase

In this section, to avoid clutter we will sometimes drop the superscript in $X^{(\ell)}$.

Let us start by considering the process $Z = (Z_k, 0 \leq k \leq \ell - 2)$ characterized by $Z_0 = 1, Z_{k+1} - Z_k \in \{0, 1\}$ and

$$\mathbb{P}(Z_{k+1} = z + 1 | Z_k = z) = \frac{z}{\ell - k - 1}.$$

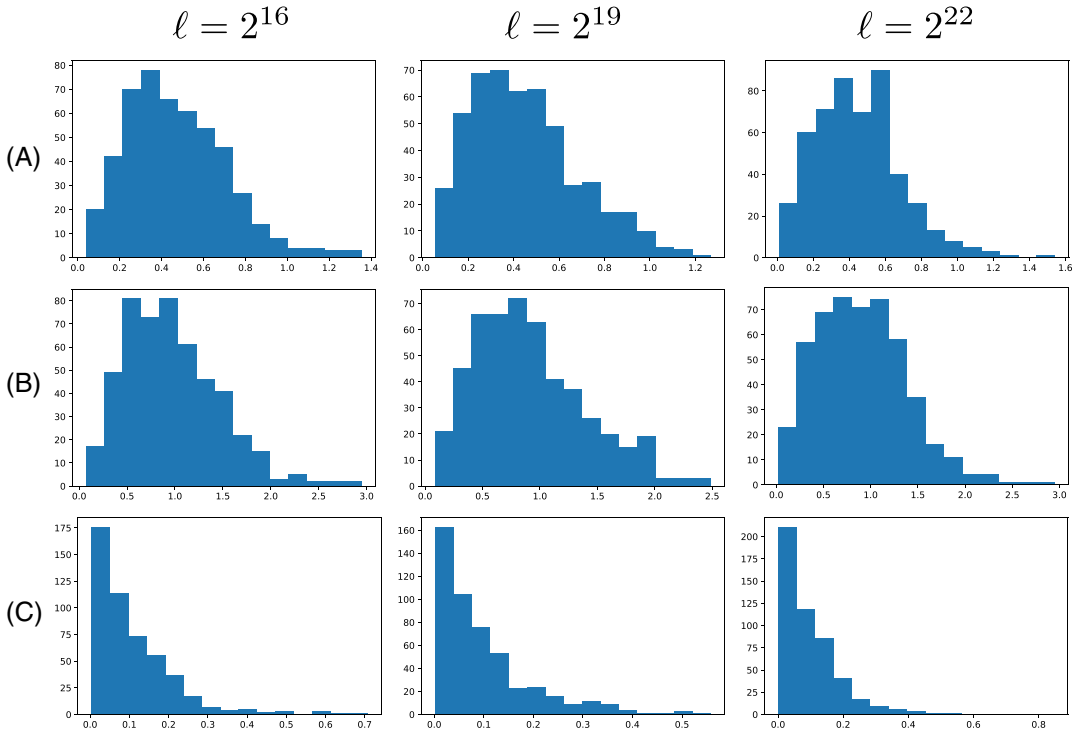


FIGURE 16 Distribution of some statistics of $X^{(\ell)}$ for $M = 10$ and 500 trajectories. (A), $\frac{1}{\sqrt{\ell}} \max \{X_k^{(\ell)} : 0 \leq k \leq \ell - 2\}$; B, $\frac{1}{\sqrt{\ell}} (\ell - \operatorname{argmax} \{X_k^{(\ell)} : 0 \leq k \leq \ell - 2\})$, and C, $\frac{1}{\sqrt{\ell}} X_{\lfloor \ell - M\sqrt{\ell} \rfloor}^{(\ell)}$

Comparing the transition probabilities of Z with those of X , we see that we can couple these two processes in such a way that $X_k \leq Z_k$ for all k . Let us now compute the first moments of Z .

Proposition 18.

- (i) $\mathbb{E}(Z_k) = \frac{\ell}{\ell - k}$
- (ii) $\mathbb{E}(Z_k^2) = \frac{\ell(\ell + k + 1)}{(\ell - k)(\ell - k + 1)}$

Proof. We have

$$\mathbb{E}(Z_{k+1}) = \left(1 + \frac{1}{\ell - k - 1}\right) \mathbb{E}(Z_k).$$

Using that $\mathbb{E}(Z_0) = 1$, we get

$$\mathbb{E}(Z_k) = \prod_{i=0}^{k-1} \left(1 + \frac{1}{\ell - i - 1}\right) = \frac{\ell}{\ell - k},$$

proving (i). For (ii), we note that

$$\mathbb{E}(Z_{k+1}^2 | Z_k) = \frac{Z_k}{\ell - k - 1} + \left(\frac{2}{\ell - k - 1} + 1\right) Z_k^2.$$

Taking expectation and substituting $\mathbb{E}(Z_k)$ by $\ell/(\ell - k)$, we get

$$\mathbb{E}(Z_{k+1}^2) = \frac{\ell}{(\ell - k)(\ell - k - 1)} + \frac{\ell - k + 1}{\ell - k - 1} \mathbb{E}(Z_k^2).$$

Multiplying both sides by $(\ell - k)(\ell - k - 1)$,

$$(\ell - k)(\ell - k - 1) \mathbb{E}(Z_{k+1}^2) = \ell + (\ell - k - 1)(\ell - k) \mathbb{E}(Z_k^2)$$

and therefore

$$\mathbb{E}(Z_k^2) = \frac{\ell(\ell + k + 1)}{(\ell - k)(\ell - k + 1)},$$

concluding the proof. ■

Proposition 19. For all $\ell \geq 2$ and k such that $0 \leq k \leq \ell - 2$,

$$\frac{\ell}{\ell - k + 1} \left(1 - \frac{2k}{(\ell - k)(\ell - k - 1)}\right) \leq \mathbb{E}(X_k) \leq \frac{\ell}{\ell - k}.$$

As a result, for $M > 2$ and all ℓ large enough, for all $\epsilon > 0$,

$$(1 - \epsilon) \frac{1}{M} \left(1 - \frac{2}{M^2}\right) \leq \mathbb{E} \left(\frac{1}{\sqrt{\ell}} X_{\lfloor \ell - M\sqrt{\ell} \rfloor}^{(\ell)} \right) \leq \frac{1}{M},$$

so that the sequence of random variables $\frac{1}{\sqrt{\ell}} X_{\lfloor \ell - M\sqrt{\ell} \rfloor}^{(\ell)}$ is tight and bounded away from 0 in L^1 .

Proof. The upper bound on $\mathbb{E}(X_k)$ follows immediately from $X_k \leq Z_k$ and Proposition 18. For the lower bound, let us denote by $\mathcal{F}_k = \sigma(X_0, \dots, X_k; Z_0, \dots, Z_k)$ the filtration generated by Z and X . Then,

$$\begin{aligned} \mathbb{E}(X_{k+1} | \mathcal{F}_k) &= X_k + \frac{X_k(\ell - k - X_k)(\ell - k - X_k - 1)}{(\ell - k)(\ell - k - 1)^2} - \frac{X_k(X_k - 1)^2}{(\ell - k)(\ell - k - 1)^2} \\ &= X_k \left(1 + \frac{(\ell - k)(\ell - k - 1) - 1}{(\ell - k)(\ell - k - 1)^2} - \frac{2(\ell - k) - 3}{(\ell - k)(\ell - k - 1)^2} X_k \right) \\ &\geq X_k \left(1 + \frac{1}{\ell - k} \right) - \frac{2Z_k^2}{(\ell - k - 1)^2} \end{aligned}$$

Taking expectations,

$$\mathbb{E}(X_{k+1}) \geq \left(1 + \frac{1}{\ell - k}\right) \mathbb{E}(X_k) - \frac{2}{(\ell - k - 1)^2} \mathbb{E}(Z_k^2),$$

and using Proposition 18 we get

$$\mathbb{E}(X_{k+1}) \geq \left(1 + \frac{1}{\ell - k}\right) \mathbb{E}(X_k) - \frac{2\ell(\ell + k + 1)}{(\ell - k + 1)(\ell - k)(\ell - k - 1)^2}.$$

Multiplying both side of the inequality by $(\ell - k)$, we get

$$(\ell - k) \mathbb{E}(X_{k+1}) \geq (\ell - (k - 1)) \mathbb{E}(X_k) - \frac{2\ell(\ell + k + 1)}{(\ell - k + 1)(\ell - k - 1)^2}.$$

and as a result,

$$(\ell - (k - 1)) \mathbb{E}(X_k) \geq \ell + 1 - 2\ell \sum_{i=0}^{k-1} \frac{\ell + i + 1}{(\ell - i + 1)(\ell - i - 1)^2}.$$

Using the fact that

$$\frac{\ell + i + 1}{(\ell - i + 1)(\ell - i - 1)^2} \leq \frac{\ell + i}{(\ell - i)(\ell - i - 1)(\ell - i - 2)}$$

and that

$$\sum_{i=0}^{k-1} \frac{\ell + i}{(\ell - i)(\ell - i - 1)(\ell - i - 2)} = \frac{k}{(\ell - k)(\ell - k - 1)},$$

we get

$$\mathbb{E}(X_k) \geq \frac{\ell}{\ell - k + 1} \left(1 - \frac{2k}{(\ell - k)(\ell - k - 1)} \right).$$

Finally, by taking $k = \lfloor \ell - M\sqrt{\ell} \rfloor$ in these inequalities we get

$$\frac{\ell}{\sqrt{\ell}(M\sqrt{\ell} + 2)} \left(1 - \frac{2(\ell - M\sqrt{\ell})}{M\sqrt{\ell}(M\sqrt{\ell} - 1)} \right) \leq \frac{1}{\sqrt{\ell}} \mathbb{E}(X_{\lfloor \ell - M\sqrt{\ell} \rfloor}) \leq \frac{1}{M}$$

where the term on the left-hand side of the inequality goes to $\frac{1}{M}(1 - \frac{2}{M^2})$ as $\ell \rightarrow \infty$. Finally, the tightness follows from Markov's inequality. ■

Proposition 19 suggests the following conjecture.

Conjecture 1. *The sequence of random variables $\frac{1}{\sqrt{\ell}} X_{\lfloor \ell - M\sqrt{\ell} \rfloor}^{(\ell)}$ converges in distribution to a positive random variable W_M .*

One problem that arises in attempting to prove Conjecture 1 it is that as soon as we get in the regime $k = \lfloor \ell - M\sqrt{\ell}(1 - t) \rfloor$, the random variables $Z^{(\ell)}$ and $X^{(\ell)}$ start to differ significantly and therefore the coupling is not so useful.

A very natural idea would be to use the backward-time construction to couple $X^{(\ell)}$ and $X_k^{(\ell+1)}$, but a difficulty with this approach is that this coupling lacks continuity in the sense that, with probability $\Theta(1/\ell)$, $X_k^{(\ell)}$ and $X^{(\ell+1)}$ will differ by a factor $\Theta(\ell)$ for $k = \lfloor \ell - M\sqrt{\ell} \rfloor$.

6.4 | The deterministic phase

Proposition 20. *If Conjecture 1 holds, that is, if*

$$\frac{1}{\sqrt{\ell}} X_{\lfloor \ell - M\sqrt{\ell} \rfloor}^{(\ell)} \xrightarrow{\ell \rightarrow \infty} W_M > 0$$

then, for all ϵ such that $0 < \epsilon < 1$, as $\ell \rightarrow \infty$,

$$\left(\frac{1}{M\sqrt{\ell}} X_{\lfloor \ell - M\sqrt{\ell}(1-t) \rfloor}^{(\ell)}, t \in [0, 1 - \epsilon] \right) \Rightarrow (y(t, C_M), t \in [0, 1 - \epsilon])$$

where \Rightarrow denotes convergence in distribution in the Skorokhod space,

$$y(t, C_M) = \frac{1-t}{C_M \cdot (1-t)^2 + 1}$$

and $C_M = M/W_M - 1$.

Proof. Let us write for convenience $M_\ell := M\sqrt{\ell}$ and

$$\mathcal{T}_\ell = \frac{1}{M_\ell} \{0, \dots, \ell - 2 - \lfloor \ell - M_\ell \rfloor\}.$$

Define the Markov chain $(Y_t^{(\ell)}, t \in \mathcal{T}_\ell)$ taking values in $\frac{1}{M_\ell}\mathbb{N}$ by

$$Y_t^{(\ell)} = \frac{1}{M_\ell} X_{\lfloor \ell - M_\ell \rfloor + tM_\ell}^{(\ell)}.$$

The Markov chain $Y^{(\ell)}$ has infinitesimal mean

$$\begin{aligned} b^{(\ell)}(y, t) &= M_\ell \mathbb{E} \left(Y_{t+1/M_\ell}^{(\ell)} - y \mid Y_t^{(\ell)} = y \right) \\ &= \mathbb{E} \left(X_{\lfloor \ell - M_\ell \rfloor + tM_\ell + 1}^{(\ell)} - yM_\ell \mid X_{\lfloor \ell - M_\ell \rfloor + tM_\ell}^{(\ell)} = yM_\ell \right) \\ &= \frac{yM_\ell(\ell - k - yM_\ell)(\ell - k - yM_\ell - 1)}{(\ell - k)(\ell - k - 1)^2} - \frac{yM_\ell(yM_\ell - 1)^2}{(\ell - k)(\ell - k - 1)^2} \end{aligned}$$

where $k = \lfloor \ell - M_\ell \rfloor + tM_\ell$. Let us show that, for any $R > 0$ and any $\epsilon > 0$,

$$b^{(\ell)}(y, t) \xrightarrow{\ell \rightarrow \infty} \frac{y(1-t-y)^2}{(1-t)^3} - \frac{y^3}{(1-t)^3},$$

uniformly in $(y, t) \in [0, R] \times [0, 1 - \epsilon]$. Let us write

$$\begin{cases} b_+^{(\ell)}(y, t) = \frac{yM_\ell(\ell - k - yM_\ell)(\ell - k - yM_\ell - 1)}{(\ell - k)(\ell - k - 1)^2} \\ b_-^{(\ell)}(y, t) = \frac{yM_\ell(yM_\ell - 1)^2}{(\ell - k)(\ell - k - 1)^2} \end{cases}$$

Using that $(1-t)M_\ell - 1 \leq \ell - k \leq (1-t)M_\ell$, we get

$$\frac{y(1-t-y-2/M_\ell)^2}{(1-t)^3} \leq b_+^{(\ell)}(y, t) \leq \frac{y(1-t-y)^2}{(1-t-2/M_\ell)^3}$$

As a result,

$$\begin{aligned} b_+^{(\ell)}(y, t) - \frac{y(1-t-y)^2}{(1-t)^3} &\geq \frac{y}{(1-t)^3} \left(-\frac{4}{M_\ell}(1-t-y) + \frac{4}{M_\ell^2} \right) \\ &\geq -\frac{4R}{\epsilon^3 M_\ell} + O(1/M_\ell^2) \end{aligned}$$

Similarly,

$$\begin{aligned} b_+^{(\ell)}(y, t) - \frac{y(1-t-y)^2}{(1-t)^3} &\leq \frac{y(1-t-y)^2}{(1-t)^3(1-t-2/M_\ell)^3} \left((1-t)^3 - (1-t-2/M_\ell)^3 \right) \\ &\leq \frac{R(1+R)^2}{\epsilon^3(\epsilon-2/M_\ell)^3} \left(\frac{6}{M_\ell} + O(1/M_\ell^2) \right). \end{aligned}$$

This proves the uniform convergence of $b_+^{(\ell)}(y, t)$. The uniform convergence of $b_-^{(\ell)}(y, t)$ is treated similarly.

Now, since $X^{(\ell)}$ has jumps of size one, $Y^{(\ell)}$ has infinitesimal variance

$$\begin{aligned} a^{(\ell)}(y, t) &= M_\ell \mathbb{E} \left(\left(Y_{t+1/M_\ell}^{(\ell)} - y \right)^2 \mid Y_t^{(\ell)} = y \right) \\ &= \frac{1}{M_\ell} \left(b_+^{(\ell)}(y, t) + b_-^{(\ell)}(y, t) \right), \end{aligned}$$

which goes to zero uniformly in $(y, t) \in [0, R] \times [0, 1 - \epsilon]$. Assuming that

$$\frac{1}{M_\ell} X_{\lfloor \ell - M_\ell \rfloor}^{(\ell)} \xrightarrow[\ell \rightarrow \infty]{d} W_M / M,$$

the convergence of the piecewise constant interpolation of $Y^{(\ell)}$ to the solution of the Cauchy problem

$$\begin{cases} \frac{dy}{dt} = \frac{y(1-t-y)^2}{(1-t)^3} - \frac{y^3}{(1-t)^3} \\ y(0) = W_M / M \end{cases}$$

follows from Corollary 4.2 of [9] (see for instance Chapter 8.7 of [8] for a more practical introduction). Note that in these references, the results are stated for time-homogeneous Markov chains. However, they are easily adapted to time-inhomogeneous ones by extending the state space with time in order to obtain a time-homogeneous process.

Finally, the function given in the Proposition is then readily checked to be the unique solution of that Cauchy problem, concluding the proof. ■

To close this section, let us mention briefly that an idea to study the embedded process $\tilde{X}^{(\ell)}$ is to introduce the process $S^{(\ell)}$ that counts the jumps of $X^{(\ell)}$. Indeed, with this process,

$$\tilde{X}_i^{(\ell)} = X_{(S^{(\ell)})^{-1}(i)}^{(\ell)}.$$

As a result, proving the convergence

$$\frac{1}{M_\ell} \left(X_{\lfloor \ell - M_\ell \rfloor + i M_\ell}^{(\ell)}, S_{\lfloor \ell - M_\ell \rfloor + i M_\ell}^{(\ell)} \right) \Rightarrow (y(t), s(t))$$

would show that

$$\frac{1}{M_\ell} \tilde{X}_{\lfloor \ell - M_\ell(1-t) \rfloor}^{(\ell)} \Rightarrow y(s^{-1}(t)),$$

and we might be able to take $M \rightarrow \infty$ to study the convergence of $\frac{1}{\sqrt{\ell}} \tilde{X}_{\lfloor \ell t \rfloor}^{(\ell)}$.

ACKNOWLEDGMENTS

FB and AL thank the Center for Interdisciplinary Research in Biology (CIRB) for funding. MS thanks the New Zealand Marsden Fund (UOC1709) for funding. The authors thank Régine Marchand, Michael Fuchs and two anonymous referees for helpful corrections and suggestions. FB thanks Jean-Jil Duchamps and Peter Czuppon for discussions about Conjecture 1.

REFERENCES

1. *The on-line encyclopedia of integer sequences*, <http://oeis.org>, 2019.
2. A. D. Barbour, L. Holst, and S. Janson, *Poisson approximation*, Clarendon Press, Oxford, 1992.
3. M. V. Boutsikas and M. V. Koutras, *A bound for the distribution of the sum of discrete associated or negatively associated random variables*, *Ann. Appl. Probab.* **10** (2000), no. 4, 1137–1150.
4. A. Caraceni, M. Fuchs, and G.-R. Yu, *Bijections for ranked tree-child networks*. arXiv preprint arXiv:2105.10137, 2021.
5. G. Cardona, J. C. Pons, and C. Scornavacca, *Generation of tree-child phylogenetic networks*, *PLOS Comput. Biol.* **15** (2019), e1007347.
6. G. Cardona, F. Rossello, and G. Valiente, *Comparison of tree-child phylogenetic networks*, *IEEE/ACM Trans. Comput. Biol. Bioinform.* **6** (2009), no. 4, 552–569.
7. G. Cardona and L. Zhang, *Counting and enumerating tree-child networks and their subclasses*, *J. Comput. Syst. Sci.* **114** (2020), 84–104.
8. R. Durrett, *Stochastic calculus: A practical introduction*, CRC Press, Boca Raton, FL, 1996.
9. S. N. Ethier and T. G. Kurtz, *Markov processes: Characterization and convergence*, 2nd ed., John Wiley & Sons, New York, NY, 2005.
10. M. Fuchs, B. Gittenberger, and M. Mansouri, *On the asymptotic growth of the number of tree-child networks*, *Australas. J. Comb.* **73** (2019), no. 2, 385–423.
11. M. Fuchs, G.-R. Yu, and L. Zhang, *On the asymptotic growth of the number of tree-child networks*, *Eur. J. Comb.* **93** (2021), 103278.
12. D. E. Knuth, *The art of computer programming: Sorting and searching*, Vol 3, Addison-Wesley, Redwood City, CA, 1997.
13. A. Lambert, *Probabilistic models for the (sub)tree(s) of life*, *Braz. J. Probab. Stat.* **31** (2017), no. 3, 415–475.
14. C. McDiarmid, C. Semple, and D. Welsh, *Counting phylogenetic networks*, *Ann. Comb.* **19** (2015), no. 1, 205–224.
15. J. Pitman, *Combinatorial stochastic processes: École d'été de probabilités de Saint-Flour XXXII-2002*, Springer-Verlag, Berlin, Heidelberg/Germany, 2006.
16. N. Ross, *Fundamentals of Stein's method*, *Probab. Surv.* **8** (2011), 201–293.
17. M. Steel, *Phylogeny: Discrete and random processes in evolution*, SIAM, Philadelphia, PA, 2016.
18. M. Steel and A. McKenzie, *Properties of phylogenetic trees generated by yule-type speciation models*, *Math. Biosci.* **170** (2001), no. 1, 91–112.
19. W. A. Stein, *Sage mathematics software (version 8.2)*, The Sage Development Team, 2019. <http://www.sagemath.org>
20. R. Van Der Hofstad, *Random graphs and complex networks*, Cambridge University Press, Cambridge, UK, 2016.
21. S. J. Willson, *Unique determination of some homoplasies at hybridization events*, *Bull. Math. Biol.* **69** (2007), no. 5, 1709–1725.
22. S. J. Willson, *Tree-average distances on certain phylogenetic networks have their weights uniquely determined*, *Algorithms Mol. Biol.* **7** (2012), 13.

How to cite this article: F. Bienvenu, A. Lambert, and M. Steel, *Combinatorial and stochastic properties of ranked tree-child networks*, *Random Struct. Algorithms*. (2021), 1–37. <https://doi.org/10.1002/rsa.21048>

APPENDIX A: LEMMAS USED IN SECTION 4

In this section, we recall without proof two elementary lemmas that were used to study the number of cherries and of tridents in Section 4.

Lemma A.1. Let (u_ℓ) be a sequence satisfying the recursion

$$u_{\ell+1} = a_\ell u_\ell + b_\ell$$

and let i be such that $\forall \ell \geq i, a_\ell \neq 0$. Then,

$$\forall \ell \geq i, \quad u_\ell = \left(u_i + \sum_{k=i}^{\ell-1} \frac{b_k}{\prod_{j=i}^k a_j} \right) \prod_{k=i}^{\ell-1} a_k$$

Lemma A.2. Let (v_ℓ) be such that $v_\ell \sim \alpha \ell^p$, where $p \geq 0$ and $\alpha \neq 0$. Then,

$$\sum_{k=i}^{\ell-1} v_k \sim \frac{\alpha}{p+1} \ell^{p+1}.$$

APPENDIX B: VARIANCE OF χ_ℓ

In this section, we prove Proposition 16 by obtaining an explicit expression for the variance of the number χ_ℓ or tridents of a uniform RTCN with ℓ labeled leaves.

Consider the Markov chain $(X_\ell)_{\ell \geq 2}$ defined in Proposition 14, and let $s_\ell = \mathbb{E}(X_\ell^2)$. From the transition probabilities of X_ℓ , we get

$$\mathbb{E}(X_{\ell+1}^2 | X_\ell = k) = \left(\frac{\ell-6}{\ell} \right)^2 k^2 + \frac{2\ell^2 - \ell - 3}{\ell^2} k + \frac{\ell-1}{\ell}.$$

Integrating in k , this yields

$$s_{\ell+1} = \left(\frac{\ell-6}{\ell} \right)^2 s_\ell + \frac{2\ell^2 - 8\ell - 3}{\ell^2} \mu_\ell + \frac{\ell-1}{\ell}$$

where, for $\ell \geq 4$, we can substitute the expression of μ_ℓ given in Proposition 15. Rearranging a bit get that for all $\ell \geq 4$,

$$s_{\ell+1} = \left(\frac{\ell-6}{\ell} \right)^2 s_\ell + \frac{30 \ell^5 - 185 \ell^4 + 188 \ell^3 + 746 \ell^2 - 1649 \ell + 843}{105 (\ell-1)(\ell-2)(\ell-3)\ell}$$

Using Lemma A.1 and a symbolic computation software, we get an explicit expression for s_ℓ , and, from there,

$$\text{Var}(\chi_\ell) = \frac{(59400 \ell^9 - 1618650 \ell^8 + O(\ell^7)) \ell}{1576575 (\ell-1)^2 (\ell-2)^2 (\ell-3)^2 (\ell-4) (\ell-5) (\ell-6)}$$

from which Proposition 16 follows.