

RANDOM ULTRAMETRIC TREES AND APPLICATIONS*

AMAURY LAMBERT¹

Abstract. Ultrametric trees are trees whose leaves lie at the same distance from the root. They are used to model the genealogy of a population of particles co-existing at the same point in time. We show how the boundary of an ultrametric tree, like any compact ultrametric space, can be represented in a simple way via the so-called comb metric. We display a variety of examples of random combs and explain how they can be used in applications. In particular, we review some old and recent results regarding the genetic structure of the population when throwing neutral mutations on the skeleton of the tree.

Résumé. Les arbres ultramétriques sont les arbres dont les feuilles se trouvent toutes à la même distance de la racine. Ces arbres sont utilisés pour modéliser la généalogie d'une population de particules qui co-existent à un temps donné. Nous montrons que la frontière d'un arbre ultramétrique, comme tout espace ultramétrique compact, peut être représentée simplement via ce que nous appelons la distance de peigne. Nous examinons plusieurs exemples de peignes aléatoires et nous expliquons comme ils peuvent être utilisés dans certaines applications. En particulier, nous évoquons quelques résultats anciens ou plus récents concernant la structure génétique de la population lorsque l'on jette des mutations neutres sur le squelette de l'arbre.

KEYWORDS AND PHRASES: random tree; real tree; reduced tree; coalescent point process; branching process; random point measure; allelic partition; regenerative set; coalescent; comb; phylogenetics; population dynamics; population genetics.

MSC 2000 subject classifications: primary 05C05, 60J80; secondary 54E45; 60G51; 60G55; 60G57; 60K15; 92D10.

CONTENTS

1. Introduction	71
2. Real Trees, Ultrametric Trees, Combs	72
2.1. The real tree	72
2.2. The comb metric	74
2.3. Sphere of a tree coded by a real function	76
2.4. The boundary of the infinite p -ary tree	77
3. Random Ultrametric Trees	78
3.1. The comb-based exchangeable coalescent	78
3.2. The coalescent point process	80

* The author thanks the Center for Interdisciplinary Research in Biology (Collège de France) for funding.

¹ Université Pierre et Marie Curie – UPMC Univ Paris 06, Laboratoire de Probabilités et Modèles Aléatoires, Case 188, 4 place Jussieu, 75252 PARIS Cedex 05; email: amaury.lambert@upmc.fr; url: <http://www.lpma-paris.fr/pageperso/amaury.lambert/>
© EDP Sciences, SMAI 2017

3.3. The boundary of a pure-birth process	81
3.4. Link between Kingman coalescent and CPP	82
4. Ultrametric Trees with Neutral Mutations	83
4.1. Throwing point mutations on a tree	83
4.2. Kingman coalescent with mutations	84
4.3. CPP with mutations	86
4.4. Supercritical processes with mutations	87
References	88

1. INTRODUCTION

In this paper, we review some mathematical properties of random tree models bearing in mind potential applications to evolutionary biology. Trees are used in *population genetics*, to trace the genealogy of a set of homologous genes, also of individuals sampled from asexual populations (in contrast to genealogies of individuals from sexual populations, which are called *pedigrees* and are not trees); in *phylogenetics*, to represent the ancestral relationships between species; in *epidemiology*, to model both the history of transmissions in epidemics and the genetic relationships between pathogenic strains. When these entities (genes, individuals, species, patients, pathogens) are sampled at the same point in time, simply called the present time, these trees are said *ultrametric*, which means that all the *leaves* of the tree lie at the same graph distance from the root. Mathematically, what actually is ultrametric w.r.t. the graph distance is the set of leaves of the tree, called its *boundary*.

From a theoretical point of view, a tree starts from one particle (coinciding with the root of the tree) and is generated by series of replication events (birth, speciation, transmission, division), which produce the so-called *branching points* of the tree (points whose complementary has at least three connected components) and of termination events (death, extinction, recovery, apoptosis), which produce the leaves, or *tips*, of the tree (points whose complement is connected). Probabilistic models for these processes abound [19,23]: branching processes, birth-death processes, Wright-Fisher and Moran model, lookdown process... The genealogy of particles present at time t is the subtree spanned by all points at the same distance t from the root. It is the ultrametric tree we have introduced in the previous paragraph, called the *reduced tree* in probability and the *coalescent tree* in population genetics. It is also the ball of radius t centered at the root.

Seen from an empirical point of view, reduced trees are also called *reconstructed trees*. Indeed, the tree itself is not available as data, and so has to be inferred (that is, reconstructed), typically from *multiple sequence alignments*. This can be done because there exist measures of *genetic distance* (i.e., dissimilarity between two aligned genetic sequences) which are good proxies of their *genealogical distance* (i.e., graph distance, or twice the time since most recent common ancestor). In most organisms, most mutations (said neutral) occur along the lineages of the tree at a more or less constant pace, hence the name of *molecular clock*. Modeling mutational processes by Poisson processes with possibly variable mutation rates in time and across genes then provides mathematical relations between genealogical distance and genetic distance, which can in turn be used to infer the tree from the sequences. Statistical methods for the inference of trees from sequences form a scientific field in its own right and will not be considered here any longer. Note that the empirical trees represented in Figures 1 and 2 are not ultrametric even though the sequences used to infer them do co-exist at present time. This is because evolutionary trees are often represented with genetic distances rather than genealogical distances.

The study of phylogenetic trees has fueled much mathematical research in graph theory, geometry and probability [7, 12, 35]. Here, we review some results, essentially those recently obtained by the author and his co-authors, around the study of ultrametric trees, initially motivated by two questions: the inference of the process most likely to have generated a given reduced tree (phylogenetics, epidemiology); the neutral genetic

composition expected to be observed in a branching population (population genetics).

We start with the general definition of real tree as a metric space, then we introduce the reduced tree (root-centered ball) and its boundary (root-centered sphere). We then show how the boundary of an ultrametric tree, like any compact ultrametric space, can be represented in a simple way via the so-called comb metric. We display a variety of examples of deterministic and random combs, the infinite p -ary tree, the Kingman comb and comb-based exchangeable combs in general, the boundary of a branching process and coalescent point processes in general. In the last section, we review some old and recent results regarding the genetic structure of the population when throwing neutral mutations on the skeleton of the tree.

2. REAL TREES, ULTRAMETRIC TREES, COMBS

2.1. The real tree

Definition 2.1. A real tree, or \mathbb{R} -tree, is a complete metric space (\mathfrak{t}, d) satisfying

(A) Uniqueness of geodesics. For any $x, y \in \mathfrak{t}$, there is a unique isometric map $\phi_{x,y} : [0, d(x, y)] \rightarrow \mathfrak{t}$ such that $\phi_{x,y}(0) = 0$ and $\phi_{x,y}(d(x, y)) = y$.

The geodesic $\phi_{x,y}([0, d(x, y)])$, also called *arc*, is denoted $\llbracket x, y \rrbracket$.

(B) No loop. For any continuous, injective map $\psi : [0, 1] \rightarrow \mathfrak{t}$, $\psi([0, 1]) = \llbracket \psi(0), \psi(1) \rrbracket$.

The *root* of an \mathbb{R} -tree \mathfrak{t} is a distinguished element of \mathfrak{t} denoted ρ .

Theorem 2.2 (Four-point condition). *The metric space (\mathfrak{t}, d) is a real tree if it is complete, path-connected and satisfies for any $x_1, x_2, x_3, x_4 \in \mathfrak{t}$*

$$d(x_1, x_2) + d(x_3, x_4) \leq \max\{d(x_1, x_3) + d(x_2, x_4), d(x_1, x_4) + d(x_2, x_3)\}$$

For major references on this topic, see [7, 12].

Definition 2.3. For any $x \in \mathfrak{t}$, the *multiplicity*, or *degree* of x denotes the number of connected components of $\mathfrak{t} \setminus \{x\}$. If $m(x) = 1$, then x is called a *leaf* or a *tip*, and if $m(x) \geq 3$, then x is called a *branching point*.

We will further need the following notation and terminology.

- **Mrca.** For any $x, y \in \mathfrak{t}$ the *most recent common ancestor* (in short mrca) of x and y , denoted $x \wedge y$, is the unique $z \in \mathfrak{t}$ such that $\llbracket \rho, x \rrbracket \cap \llbracket \rho, y \rrbracket = \llbracket \rho, z \rrbracket$.
- **Partial order.** For any $x, y \in \mathfrak{t}$, y is said to *descend* from x , and then x is called an *ancestor* of y if $x \in \llbracket \rho, y \rrbracket$, and this is denoted $x \preceq y$.
- **Length measure.** Whenever \mathfrak{t} is locally compact, there is a unique measure λ on the Borel σ -field of \mathfrak{t} , called *length measure*, such that for any $x, y \in \mathfrak{t}$, $\lambda(\llbracket x, y \rrbracket) = d(x, y)$ (see Section 4.3.5 in [12]).
- **Reduced tree.** For a real tree \mathfrak{t} and a fixed real number $T > 0$, the so-called *reduced tree* at height T is the tree spanned by points at distance T from the root, i.e.

$$\{y \in \mathfrak{t} : \exists x \in \mathfrak{t}, y \preceq x, d(\rho, x) = T\}.$$

The topology of the reduced tree can be understood from the topology of the *sphere* of \mathfrak{t} with center ρ and radius $T > 0$

$$\mathfrak{t}^{\{T\}} := \{x \in \mathfrak{t} : d(\rho, x) = T\},$$

Note that by the four-point condition, for any $x, y, z \in \mathfrak{t}^{\{T\}}$,

$$T + d(x, z) = d(\rho, y) + d(x, z) \leq \max\{d(\rho, x) + d(y, z), d(\rho, z) + d(y, x)\} = \max\{T + d(y, z), T + d(y, x)\},$$

which yields $d(x, z) \leq \max\{d(y, z), d(y, x)\}$, that is the metric induced by d on $\mathfrak{t}^{\{T\}}$ is *ultrametric*. From now on, we assume that (\mathfrak{t}, d) is **locally compact**, so that $(\mathfrak{t}^{\{T\}}, d)$ is a **compact ultrametric space** (by application

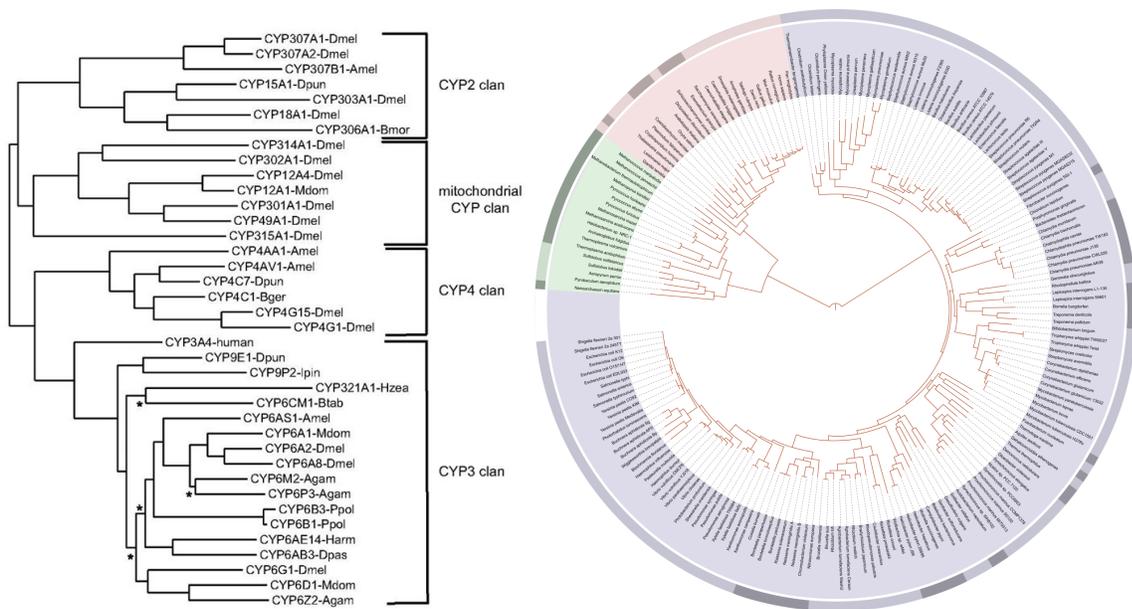


FIGURE 1. Left panel: A gene tree showing the evolution of the CYP gene family in insects. Right panel: A phylogenetic tree of contemporary species, with different colors for each of the three domains: Bacteria (grey), Archaea (green) and Eukarya (pink)¹.

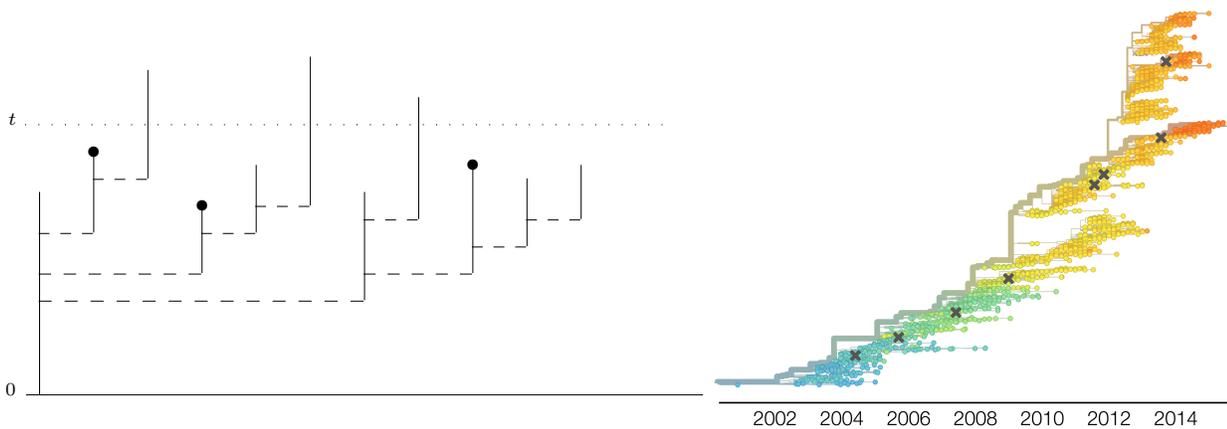


FIGURE 2. Left panel: A hypothetical transmission tree, with dots showing recovery events (exit from the epidemic). Right panel: A phylogenetic tree of H3N2 strains of influenza².

of the Hopf-Rinow theorem, since a real tree is a length-metric space). We will see in the next section that any compact ultrametric space can be represented by what we call a *comb*.

¹Sources: <http://what-when-how.com/insect-molecular-biology-and-biochemistry/insect-cyp-genes-and-p450-enzymes-part-2/> (left panel) and [https://en.wikipedia.org/wiki/Tree_of_life_\(biology\)](https://en.wikipedia.org/wiki/Tree_of_life_(biology)) (right panel).

²Source (right panel): <http://bedford.io/talks/real-time-tracking-vidd/>

2.2. The comb metric

Let I be a compact interval and $f : I \rightarrow [0, \infty)$ such that for any $\varepsilon > 0$, $\{f \geq \varepsilon\}$ is finite. For any $s, t \in I$, define \bar{d}_f by

$$\bar{d}_f(s, t) = 2 \max_{(s \wedge t, s \vee t)} f.$$

It is clear that \bar{d}_f is a pseudo-distance on $\{f = 0\}$ and that it is ultrametric, i.e.

$$\bar{d}_f(r, t) \leq \max\{\bar{d}_f(r, s), \bar{d}_f(s, t)\} \quad r, s, t \in I.$$

Let us assume additionally that $\{f \neq 0\}$ is dense in I for the usual topology, so that \bar{d}_f is a distance on $\{f = 0\}$.

Definition 2.4. We call f a *comb-like function* or *comb*, and \bar{d}_f the *comb metric* on $\{f = 0\}$.

The space $(\{f = 0\}, \bar{d}_f)$ is not complete in general. To make it complete, one has to distinguish for each point $t \in I$ between its **left face** (t, l) and its **right face** (t, r) . The distance \bar{d}_f is extended to the space $I \times \{l, r\}$ by the following definitions for $s < t \in I$

$$\begin{aligned} \bar{d}_f((s, r), (t, l)) &= 2 \max_{(s, t)} f, & \bar{d}_f((s, l), (t, l)) &= 2 \max_{[s, t]} f, \\ \bar{d}_f((s, r), (t, r)) &= 2 \max_{(s, t]} f, & \bar{d}_f((s, l), (t, r)) &= 2 \max_{[s, t]} f. \end{aligned}$$

If $f(s) = 0$, $\bar{d}_f((s, l), (s, r)) = 0$ so that (s, l) and (s, r) must be identified. It can be shown [30] that the associated **quotient space** (\bar{I}, \bar{d}_f) is a **compact, ultrametric space** called **comb metric space**. Actually the converse also holds, as we will see with Theorem 2.6.

Before stating this theorem, we wish to construct the ultrametric tree hidden behind the comb metric space, as illustrated on Figure 3b.

Definition 2.5. Let f be a comb on the interval $I = [0, a]$ and $T > \max f$. We define

$$\text{Sk} := \{0\} \times (0, T] \cup \{(t, y) \in I \times (0, T] : f(t) > y\},$$

endowed with the distance \bar{d}_f

$$\bar{d}_f((s, x), (t, y)) = \begin{cases} |\max_{(s, t]} f - x| + |\max_{(s, t]} f - y| & \text{if } s < t, \\ |x - y| & \text{if } s = t. \end{cases}$$

The tree $\tau_f(T)$ is defined as the **completion of** (Sk, \bar{d}_f) , so that in particular Sk is its skeleton. We will always take the root of $\tau_f(T)$ equal to $\rho = (0, T)$. For $t \in (0, T]$, we call the **lineage of** t the subset of the tree L_t defined as the closure of the set

$$\{(s, x) \in \text{Sk} : s \leq t, x \geq \max_{(s, t]} f\}.$$

It can be shown [8] that the boundary of (τ_f, \bar{d}_f) is indeed (\bar{I}, \bar{d}_f) , which explains why we keep the same notation for the two distances. Also note that for each $t \in I$, $L_t = \llbracket \rho, \alpha_t \rrbracket$, where $\alpha_t \in \bar{I}$ is equal to (t, r) .

Theorem 2.6 ([30]). *Any compact ultrametric space without isolated point is isometric to a comb metric space.*

Note that a comb metric space (\bar{I}, \bar{d}_f) is naturally endowed with the **finite measure** ℓ defined by

$$\ell(A \times B) := \text{Leb}(A),$$

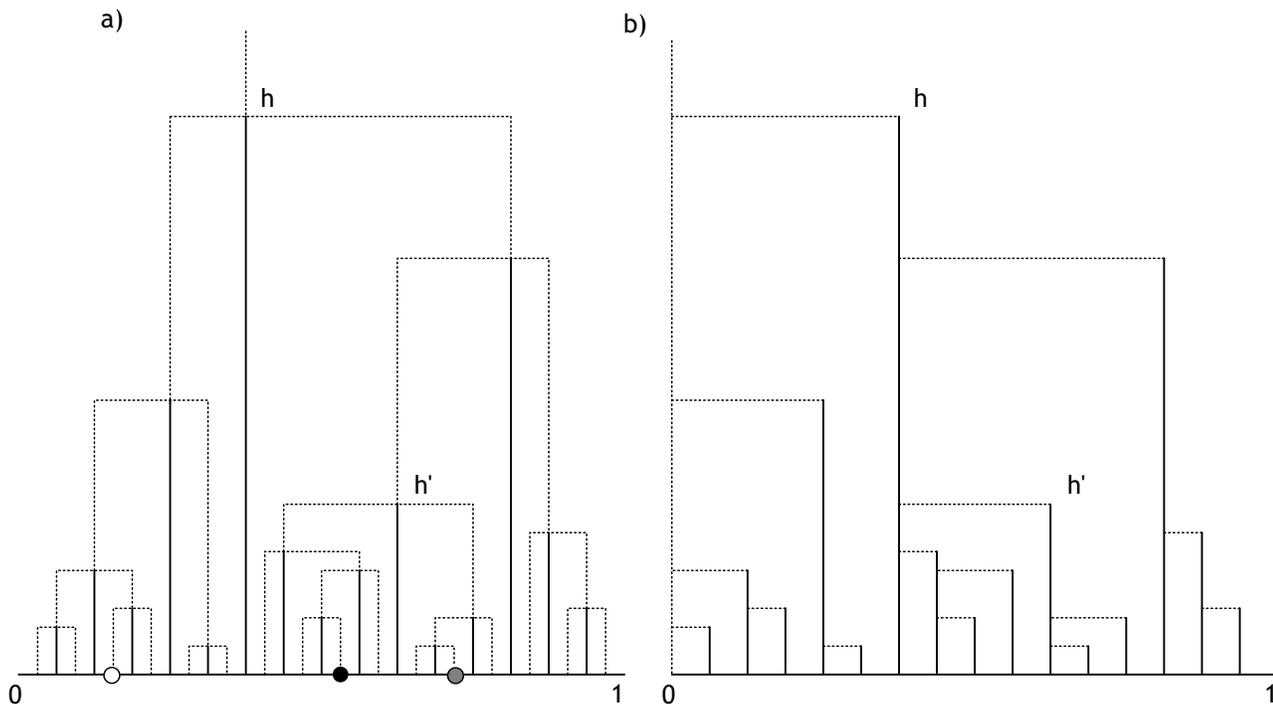


FIGURE 3. A comb-like function on $[0, 1]$ and two ways (a and b) of representing the associated ultrametric tree in dotted lines. a) The distance between the black dot and the grey dot is h' , whereas h is the distance from either of these dots to the white dot; b) Graphical representation of the ultrametric tree constructed from the comb, introduced in Definition 2.5.

for any Borel set $A \subseteq I$ and $B \subseteq \{l, r\}$ (where Leb denotes the Lebesgue measure, and it is important that $\{f \neq 0\}$ is dense), which suggests that any compact ultrametric space can be equipped with a finite measure (isolated points can be treated separately). Actually, any compact ultrametric space can be endowed with a finite measure charging every ball with non-zero radius. One example of such a measure is the so-called **visibility measure** [31], as shown by the following argument, which is actually also used in the proof of the theorem.

For any ultrametric space (U, d) and any $r > 0$, U can be partitioned into balls of radius r , because the relation \sim_r defined by $x \sim_r y \Leftrightarrow d(x, y) \leq r$ is an equivalence relation. If in addition U is assumed to be compact, the number M_r of blocks in this partition has to be finite, and it is nondecreasing in r . The visibility measure is constructed by putting mass 1 on U , and recursively at each jump time of M as r decreases, by dividing the mass of each fragmenting block equally between its new sub-blocks. The comb can be constructed simultaneously with this recursive construction of the visibility measure, by mapping each ball with measure m to an interval with length m , and putting ‘walls’ between such intervals (the graph of the comb). See [30] for the mathematical details.

Of course, if U is already endowed with a measure, the same construction can be done using the given measure. This is in particular the case when the ultrametric space is a sphere of a totally ordered, measured

tree. It can then be shown [29] that there is a càdlàg function $h : \mathbb{R} \rightarrow \mathbb{R}$ with no negative jumps which codes for the tree in a sense that we specify hereafter.

2.3. Sphere of a tree coded by a real function

Let $h : [0, \infty) \rightarrow [0, \infty)$ be càdlàg with no negative jumps and compact support. We are going to explain how h codes for a real tree. Set $\sigma_h := \sup\{t > 0 : h(t) \neq 0\}$ and

$$d_h(s, t) := h(s) + h(t) - 2 \inf_{[s \wedge t, s \vee t]} h.$$

It is clear that d_h is a pseudo-distance on $[0, \infty)$. Further let \sim_h denote the equivalence relation on $[0, \infty)$

$$s \sim_h t \Leftrightarrow d_h(s, t) = 0 \Leftrightarrow h(s) = h(t) = \inf_{[s \wedge t, s \vee t]} h.$$

Theorem 2.7. Denote by \mathfrak{t}_h the quotient space $[0, \sigma_h] / \sim_h$. Then (\mathfrak{t}_h, d_h) is a compact \mathbb{R} -tree.

Figure 4 shows how to uncover the tree coded by a function, for a very simple example. From now on, let

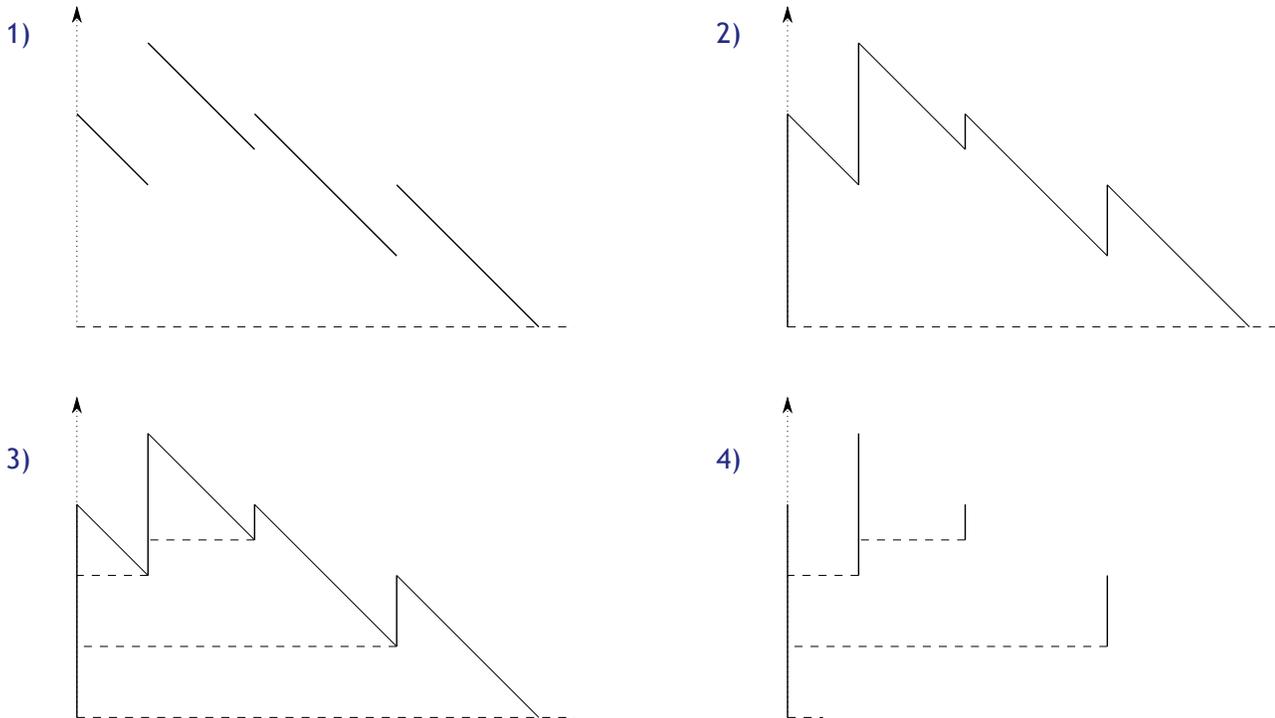


FIGURE 4. How to uncover the tree coded in a càdlàg map which has a.e. negative derivative and positive jumps; 1) Start from the graph of this map; 2) Draw vertical solid lines in the place of jumps; 3) Report horizontal dashed lines from each edge bottom left to the rightmost solid point; 4) erase diagonal lines.

$p_h : [0, \sigma_h] \rightarrow \mathfrak{t}_h$ map any element of $[0, \sigma_h]$ to its equivalence class relative to \sim_h . Note that the tree \mathfrak{t}_h is naturally endowed with a total order and a mass measure, as follows.

- **Total order.** We define \leq_h as the *order of first visits*, that is for any $x, y \in \mathfrak{t}_h$,

$$x \leq_h y \Leftrightarrow \inf p_h^{-1}(\{x\}) \leq \inf p_h^{-1}(\{y\}).$$

- **Mass measure.** The measure μ_h is defined as the *push forward of Lebesgue measure* by p_h .

Conversely, it can actually be shown [9, 29] that if (\mathfrak{t}, d) is a compact \mathbb{R} -tree endowed with a total order \leq and a finite mass measure μ satisfying some consistency conditions, there is a unique càdlàg map h called the **jumping contour process** of \mathfrak{t} such that the tree $(\mathfrak{t}_h, d_h, \leq_h, \mu_h)$ is isomorphic to $(\mathfrak{t}, d, \leq, \mu)$ (follow the panels of Figure 4 in reverse order).

Now let us consider a compact real tree \mathfrak{t} coded by a function h (its jumping contour process if the function is not given *a priori*) and let $T > 0$ be such that the sphere $\mathfrak{t}^{\{T\}}$ is not empty. We know that $\mathfrak{t}^{\{T\}}$ is an ultrametric space, but we have no guarantee that μ_h charges $\mathfrak{t}^{\{T\}}$, so we will directly construct an isometric comb metric space, following [30].

Assume that the sphere $\mathfrak{t}^{\{T\}}$ has no isolated point in itself. Then it can be shown that $\{h = T\}$ has no isolated point and has empty interior. So we can construct a *local time* at level T for h , that is a nondecreasing, continuous map $L : [0, \infty) \rightarrow [0, \infty)$ such that $L(0) = 0$ and for any $s < t$

$$L_t > L_s \Leftrightarrow (s, t) \cap \{h = T\} \neq \emptyset.$$

Let $I = [0, L_{\sigma_h}]$. Let J denote the right inverse of the local time

$$J_t := \inf\{s > 0 : L_s > t\} \quad t \in I.$$

Every jump time s of J corresponds to an interval $(g_s, d_s) := (J_{s-}, J_s)$ where L is constant (to s) such that $h \neq T$ on (g_s, d_s) and $h(g_s-) = h(g_s) = T = h(d_s)$ (because h has no negative jump). In particular, $\mathfrak{t}^{\{T\}}$ is the closure of the range of $p_h \circ J$. Now for any $s < t$ such that $h(s) = h(t) = T$,

$$d_h(s, t) = h(s) + h(t) - 2 \inf_{[s, t]} h = 2(T - \inf_{[s, t]} h).$$

So for any $s < t$ in I , the distance between the two points $p_h(J_s)$ and $p_h(J_t)$ of $\mathfrak{t}^{\{T\}}$ is $2 \max_{[s, t]} f$, where

$$f(s) := \begin{cases} T - \inf_{[g_s, d_s]} h & \text{if } \Delta J_s \neq 0 \\ 0 & \text{otherwise.} \end{cases}$$

This indicates that $\mathfrak{t}^{\{T\}}$ should be represented by the comb metric space associated to the comb f just defined. Note that f is a comb with values equal to the **depths of the excursions of h away from T** . Further define

$$\begin{aligned} \bar{\theta} : (\bar{I}, \bar{d}_f) &\longrightarrow (\mathfrak{t}^{\{T\}}, d) \\ s' &\longmapsto \begin{cases} p_h(J_{s'-}) & \text{if } s' = (s, l) \text{ and } f(s) \neq 0 \\ p_h(J_s) & \text{otherwise.} \end{cases} \end{aligned}$$

Theorem 2.8 ([30]). *The map $\bar{\theta} : (\bar{I}, \bar{d}_f) \rightarrow (\mathfrak{t}^{\{T\}}, d_h)$ is a global isometry preserving the order and mapping the Lebesgue measure to the push forward of the measure dL by p_h .*

2.4. The boundary of the infinite p -ary tree

The fundamental example of compact ultrametric space is the boundary of the infinite p -ary tree. This is the set U_p of sequences $x = (x_n)_{n \geq 1}$ with values in $\{0, \dots, p-1\}$ endowed with the distance d_u defined for any pair $x = (x_n)$ and $y = (y_n)$ of elements of U_p by

$$d_u(x, y) = p^{-v_u(x-y)},$$

where

$$v_u(x) := \min\{n \geq 1 : x_n \neq 0\},$$

with the convention $p^{-v_u(0)} = 0$. It is known that (U_p, d_u) is a compact ultrametric space. The distance d_u is actually the graph distance associated to the tree where each edge between generation $n - 1$ and generation n has length equal to p^{-n} . We show how to construct explicitly the isometry of the theorem between (U_p, d_u) and a comb metric space which can intuitively be guessed from the previous remark, as shown on Figure 5.

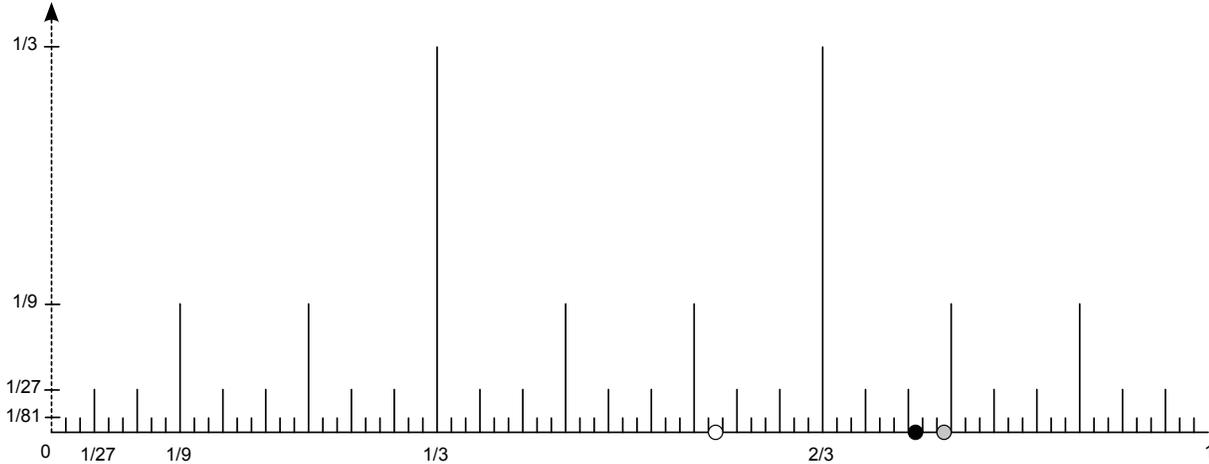


FIGURE 5. Triadic comb. The white dot represents a sequence starting with $(1, 2, 0, 1, \dots)$, the black dot represents a sequence starting with $(2, 0, 2, 0, \dots)$ and the grey dot a sequence starting with $(2, 0, 2, 2, \dots)$. The distances white/black and white/grey are both equal to $(1/3)^1$, and the distance black/grey equals $(1/3)^4$.

Recall that $t \in [0, 1]$ is called a *p-adic number* if it has two distinct *p-adic* decompositions. We denote by $\phi_l(t)$ its *p-adic* decomposition stationary at $p - 1$ and by $\phi_r(t)$ its *p-adic* decomposition stationary at 0. When t is not *p-adic*, its *p-adic* decomposition is simply denoted $\phi(t)$. Now for any $x \in U_p$, set $w(x) := \max\{n \geq 1 : x_n \neq 0\}$, and for any $t \in [0, 1]$, define

$$F_p(t) := \begin{cases} p^{-w(\phi_r(t))} & \text{if } t \text{ is } p\text{-adic} \\ 0 & \text{otherwise.} \end{cases}$$

The function F_p is a comb called the ***p-adic comb*** and it is not difficult to see that \bar{d}_{F_p} defines a comb metric space (\bar{I}, \bar{d}_{F_p}) . Then we define $\bar{\phi} : (\bar{I}, \bar{d}_{F_p}) \rightarrow (U_p, d_u)$ that maps every $t' \in [0, 1] \times \{l, r\}$ to the *p-adic* decomposition of t , except when t is *p-adic* which then has two faces, each mapped to a distinct *p-adic* decomposition as follows.

$$\bar{\phi}(t') := \begin{cases} \phi(t) & \text{if } t' = t \in \{F_p = 0\} \\ \phi_l(t) & \text{if } t' = (t, l) \\ \phi_r(t) & \text{if } t' = (t, r). \end{cases}$$

It can be seen that $\bar{\phi}$ is a global isometry between (\bar{I}, \bar{d}_{F_p}) and (U_p, d_u) conserving the measures, where U_p is endowed with the visibility measure μ_p that puts mass p^{-n} on any ball with radius p^{-n} . The measure μ_p is also the law of a sequence of i.i.d. random variables uniformly distributed in $\{0, \dots, p - 1\}$. See Figure 6 for an illustration of the ***dyadic comb*** showing the left and right faces of the same dyadic number.

3. RANDOM ULTRAMETRIC TREES

3.1. The comb-based exchangeable coalescent

A first example of random ultrametric tree is the following. Start with a comb f on $[0, 1]$ and an independent sequence (V_i) of independent and identically distributed (i.i.d.) random variables uniform in $(0, 1)$. Note that

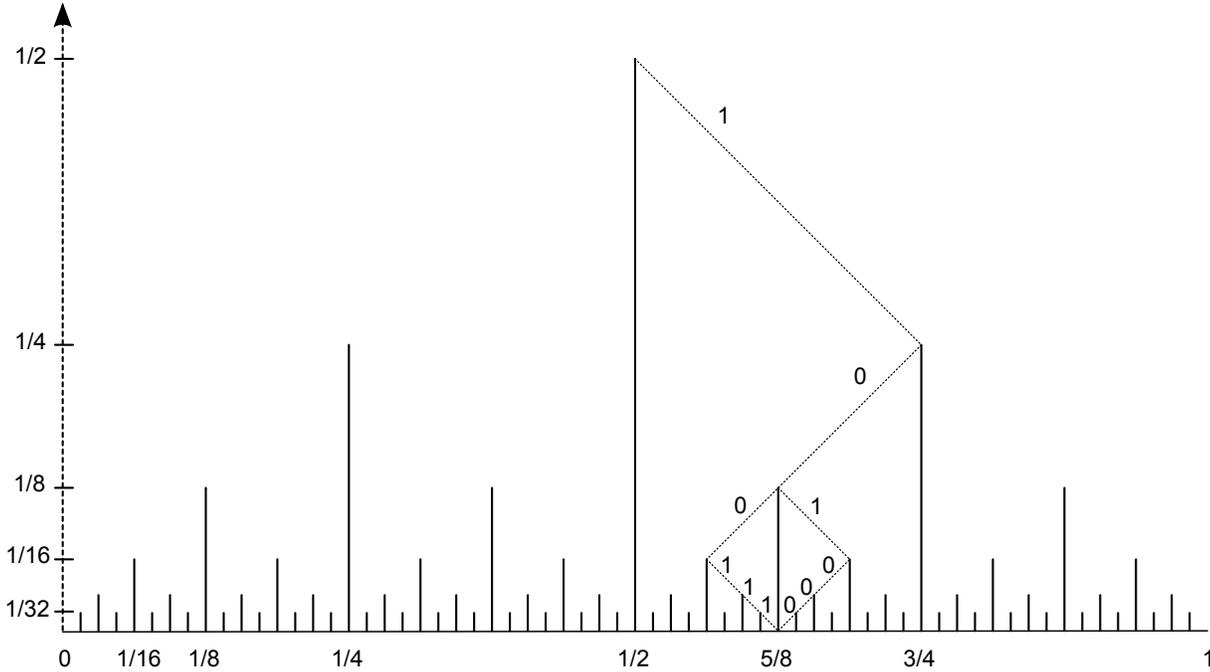


FIGURE 6. Dyadic comb. The dashed lines represent two infinite sequences of 0's and 1's: on the right, the sequence $x = (1, 0, 1, 0, 0, 0, \dots)$ and on the left the sequence $\hat{x} = (1, 0, 0, 1, 1, 1, \dots)$. They correspond to the right and left faces of $5/8$, i.e., the dyadic decompositions of $5/8$ stationary at 0 and at 1, respectively. They lie at distance $(1/2)^3$ from each other.

a.s. $f(V_i) = 0$ for all i , so we can define the ultrametric distance δ_f on \mathbb{N} by

$$\delta_f(i, j) := \bar{d}_f(V_i, V_j).$$

Now for any $t > 0$, define the partition $R_f(t)$ on \mathbb{N} induced by the equivalence relation \sim_t

$$i \sim_t j \Leftrightarrow \delta_f(i, j) \leq t.$$

Note that $R_f(0+)$ is the finest partition (all singletons) because $\{f \neq 0\}$ is assumed dense in $[0, 1]$.

Definition 3.1. The *comb-based coalescent process* $(R_f(t); t > 0)$ is an *exchangeable coalescent process*, in the sense that its law is invariant under permutations of \mathbb{N} .

The converse statement is given in the next proposition.

Proposition 3.2. Let $(R(t); t > 0)$ be a random exchangeable coalescent process such that for each $t > 0$, $R(t)$ has a finite number of blocks and no singleton. Then there is a random comb f such that the comb-based coalescent process R_f is equal in distribution to R .

Proof. We can define the distance δ on \mathbb{N} by

$$\delta(i, j) := \inf\{t \geq 0 : i \text{ and } j \text{ are in the same block of } R(t)\}.$$

Then it is straightforward that δ is ultrametric. Indeed, for any integers i, j, k , and for any $t \geq 0$, if $\delta(i, j) \leq t$ and $\delta(j, k) \leq t$, then i and j on the one hand, and j and k on the other hand, are in the same block of $R(t)$, so

that i and k are in the same block of $R(t)$. This shows that $\delta(i, k) \leq t$, so that $\delta(i, j) \leq \max\{\delta(i, j), \delta(j, k)\}$. Now by de Finetti's theorem, each ball A of the ultrametric space (\mathbb{N}, δ) has an asymptotic frequency $m(A)$ defined by

$$m(A) := \lim_{n \uparrow \infty} \frac{\#A \cap \{1, \dots, n\}}{n}.$$

Unfortunately, m is not a measure (it does not satisfy Caratheodory's property) and (\mathbb{N}, δ) is not compact. But because we assume that for all $t > 0$, $R(t)$ has a finite number of blocks and no singleton, we can use the same procedure as outlined after the statement of Theorem 2.6 and construct a comb f on $I = [0, 1]$ and a map ϕ between balls of (\mathbb{N}, δ) and balls of (I, f) , which conserves radii, measures and partial order (inclusion). Now let (V_i) be i.i.d. uniform random variables independent of the comb (I, f) and of the map ϕ and define δ_f by $\delta_f(i, j) = \bar{d}_f(V_i, V_j)$. By construction, the processes of ranked frequencies of blocks of R and of R_f are a.s. equal. Kingman's representation theorem of exchangeable partitions ensures that the law of $R(\varepsilon)$ can be obtained from its frequencies by a paintbox process, so that $R_f(\varepsilon)$ has the same law as $R(\varepsilon)$. By monotonicity, this implies that $(R_f(t); t \geq \varepsilon)$ and $(R(t); t \geq \varepsilon)$ are equally distributed. Since ε is arbitrary, the result is proved. \square

The archetypal example of comb-based coalescent is the following.

Definition 3.3. Define the random comb f on $I = [0, 1]$ by

$$f = \sum_{j \geq 1} \tau_j \mathbb{1}_{\{U_j\}},$$

where the families of r.v. (U_j) and (τ_j) are independent, the (U_j) are i.i.d. uniform on $(0, 1)$ and $\tau_j = \sum_{k \geq j+1} e_k$, where e_k are independent exponential r.v. with parameter $k(k-1)/2$. Then the comb-based coalescent R_f has the same distribution as the **Kingman coalescent** [18]. We call f the **Kingman comb**.

In the rest of this section, we assume that \mathfrak{t} denotes a **binary** \mathbb{R} -tree, that is, $m(x) \leq 3$ for all $x \in \mathfrak{t}$. See [26] for extensions to random real trees with arbitrarily large degree.

3.2. The coalescent point process

Recall from Theorem 2.7 that any càdlàg map h codes for a compact \mathbb{R} -tree denoted \mathfrak{t}_h . Let us call **Brownian tree** the tree coded by a Brownian excursion conditioned to have height larger than T . The Brownian excursion has a local time at level T , which allows one to construct as in Theorem 2.8 the comb giving the metric of the reduced tree at level T . This comb is a 'list', in the plane order, of the depths of excursions of the contour away from T . These depths form a Poisson point process, which motivates the following definition.

Definition 3.4. Let ν be a σ -finite measure on $(0, \infty)$ such that $\nu([\varepsilon, \infty)) < \infty$ for all $\varepsilon > 0$. Let \mathcal{M} be a Poisson point process on $(0, \infty)^2$ with intensity $\text{Leb} \otimes \nu$ and denote by $(S_i, H_i)_i$ its atoms. Finally, let (D, H) denote the first (i.e., with smallest first component) atom such that $H > T$. We will say that the random comb metric space associated with the comb $\sum_{i: S_i < D} H_i \mathbb{1}_{\{S_i\}}$ is a **coalescent point process** with height T and intensity measure ν .

Proposition 3.5 ([1, 34]). *The sphere $\mathfrak{t}^{\{T\}}$ of the Brownian tree \mathfrak{t} is isometric to a coalescent point process with height T and intensity measure ν_0 , where ν_0 is the push forward of the Brownian excursion measure by the function which maps an excursion to its depth, i.e.*

$$\nu_0(dx) = \frac{dx}{2x^2}$$

We will call **Brownian CPP** a CPP with intensity measure ν_0 (or a multiple of it).

More generally, a (root-centered) **sphere of any real tree whose contour process is strongly Markovian is isometric to a coalescent point process (CPP)**. This is the case of **splitting trees**, which are the trees generated by binary branching processes where particles give birth at constant rate during a lifetime which follows an arbitrary distribution [14]. A splitting tree is actually isometric to a tree coded by a Lévy process with finite variation [21].

More generally, we say that a random binary \mathbb{R} -tree satisfies the **splitting property** if for any $x \in \mathfrak{t}$, the subtrees rooted at the branching points of the segment $\llbracket \rho, x \rrbracket$ form a Poisson point process on $\llbracket \rho, x \rrbracket \times \mathcal{E}$, where \mathcal{E} denotes the space of locally compact \mathbb{R} -trees. For example, the tree coded by a Brownian excursion satisfies the splitting property. We have proved that a tree satisfying the splitting property, again called splitting tree, is isometric to a tree coded by a Lévy process with possibly infinite variation [29], so that its sphere is again isometric to a CPP.

From a practical point of view, an interesting question is to characterize the intensity measure ν of the CPP. In the case of strongly Markovian contour processes, ν is **simply the push forward of the excursion measure away from T of the contour process, by the function which maps an excursion to its depth**, and a lot is known on this measure in the Lévy case (see [21] for details). There is also a whole class of random trees whose spheres are CPPs with finite intensity measure (derogating to our definition of combs having dense support) and which have applications in evolutionary biology.

Consider a population where all individuals live and reproduce independently, and each individual is endowed with a trait (some random character living in \mathbb{R} for simplicity) that evolves through time according to independent copies of the same, possibly time-inhomogeneous, Markov process. Further assume what follows.

- This trait is non-heritable, in the sense that any individual born at time t draws the value of her trait at birth from the same distribution, independently of her ancestors' histories;
- All individuals give birth during their lifetime, according to a Poisson point process with intensity β , where β is a diffuse Radon measure on $[0, \infty)$;
- An individual holding trait x at time t dies at rate $d(t, x)$.

Theorem 3.6 ([28]). *For a tree generated by the previously defined population model, starting with one individual at time 0 and conditional on having at least one alive individual at time T , the sphere of radius T is isometric to a CPP with intensity measure ν given by*

$$W(t) := \frac{1}{\nu([t, \infty))} = \exp \left(\int_{T-t}^T (1 - q(s)) \beta(ds) \right) \quad t \in [0, T],$$

where $q(t)$ denotes the probability that an individual born at time t has no descendants alive by time T . Note that the knowledge of $W(t)$ for $t > T$ is not needed. In addition, if $g(t, \cdot)$ denotes the density of the death time of an individual born at time t , then W is solution to the following integro-differential equation

$$W'(t) = b(T - t) \left(W(t) - \int_0^t W(s) g(T - t, T - s) ds \right) \quad t \geq 0,$$

with initial condition $W(0) = 1$.

The last statement can be used to compute the likelihood of a given reconstructed tree (cf. Introduction) to infer the model that most likely has generated this tree. Two beautiful biology papers applying this procedure are [32, 36]. More recent examples of applications can be found in Section 3.3 of the lecture notes [23], see in particular [2, 11, 24, 25].

3.3. The boundary of a pure-birth process

Here, we consider a pure-birth process on $[0, \infty)$ with birth rate β , assumed to be a diffuse Radon measure on $[0, \infty)$ such that $\beta([0, \infty)) = \infty$. We will write $\underline{\beta}(t) := \beta([0, t])$.

The genealogy of such a process can be encoded by the infinite binary tree \mathcal{T} (finite sequences of 0's and 1's) endowed with the birth dates $\alpha(v)$ of each finite sequence v . If u is a prefix of v , we write $u \preceq v$ and we say that v is a *descendant* of u . We will denote $\omega(v)$ the date of death of v and assume that $\omega(v) = \alpha(v0) = \alpha(v1)$. The tree \mathcal{T} may then be equipped with a **measure \mathcal{L} on its boundary** $\partial\mathcal{T} = \{0, 1\}^{\mathbb{N}}$ defined by

$$\mathcal{L}(B_u) := \lim_{t \uparrow \infty} \frac{N_u(t)}{e^{\beta(t)}} \quad u \in \mathcal{T},$$

where B_u denotes the set of infinite sequences v with prefix u and $N_u(t)$ is the number of descendants of u at time t

$$N_u(t) := \#\{v \in \mathcal{T}, u \preceq v, \alpha(v) < t \leq \omega(v)\}.$$

The mass of \mathcal{L} is an exponential r.v. with parameter 1.

Theorem 3.7. *Let $\varphi : [0, \infty) \rightarrow (0, T]$ be a decreasing bijection and let $\varphi(\mathfrak{t})$ denote the tree obtained from the previous pure-birth tree \mathfrak{t} by mapping all distances to the root by $T - \varphi$ and letting the measure \mathcal{L} unchanged. Then $\varphi(\mathfrak{t})$ is isometric to the tree τ_f introduced in Definition 2.5 constructed from a CPP f with intensity ν and height T , naturally measured by ℓ (defined page 74), where*

$$\nu([t, \infty)) = \exp(\underline{\beta} \circ \varphi^{-1}(t)) \quad t \in [0, T].$$

Note that $\nu([T, \infty)) = 1$ due to the normalization of \mathcal{L} .

In particular, if we take $T = 1$ and φ given by $\varphi(t) = e^{-\beta(t)}$, then $\varphi(\mathfrak{t})$ has the distribution of τ_f , where f is a CPP with height 1 and intensity $\frac{dx}{x^2}$.

Proof. We know from [8] that $\varphi(\mathfrak{t})$ is the genealogy of a reversed (i.e. time flows from T to 0) pure-birth tree with birth rate $\beta' = \beta \circ \varphi^{-1}$. From the same paper, we know that a reversed pure-birth tree with birth rate β' and measure \mathcal{L} (invariant by the time-change) on its boundary is the tree τ_f associated with a CPP f with intensity measure ν satisfying that the Laplace-Stieltjes measure associated with the increasing function $\ln(\bar{\nu})$ is equal to β' , where we wrote $\bar{\nu}(t) = \nu([t, \infty))$. The proof follows from a simple calculation. \square

Notice that we could have applied the same reasoning to a supercritical birth-death process conditioned to non-extinction, by considering the subtree spanned by its boundary, i.e. the tree of indefinitely surviving lineages, which has the law of the tree generated by the pure-birth process with birth rate $\tilde{\beta}(dt) = \beta(dt)(1 - q(t))$, where $q(t)$ is the probability of survival of a single particle alive at time t . Then everything works as in Theorem 3.7 provided that the counter $N_u(t)$ is restricted to the particles alive at t with indefinitely surviving descendance, and that β is replaced with $\tilde{\beta}$ in the displayed formula.

Furthermore, we can recover the first part of Theorem 3.6 by replacing in the last paragraph ‘infinite survival’ with ‘survival up to T ’ and by taking $\varphi(t) = T - t$ in Theorem 3.7.

3.4. Link between Kingman coalescent and CPP

Recall the **Kingman comb** and the **Brownian CPP** defined respectively in Definition 3.3 and Proposition 3.5. Both objects code for the genealogy of a large exchangeable population, but the Kingman coalescent is based on the assumption of a stationary population with constant size (total size constraint), whereas in the CPP the size of the population is fluctuating like a branching process, and its foundation time is fixed (time constraint). Following [27], we show that one of the two is embedded in the other. In an exchangeable population with large constant size, the descendance of a small subpopulation is blind to the total size constraint and it is constant in expectation (see for example Theorem 1 in [3]). Our goal is to state a backward-in-time version of the last informal observation, namely ‘in a large stationary population with constant size, the genealogy of a subpopulation with recent MRCA is given by a CPP’, and to derive some consequences of this fact.

Let $K : [0, 1] \rightarrow [0, \infty)$ denote the Kingman comb and let $C : [0, \infty) \rightarrow [0, \infty)$ denote the CPP with intensity measure $\nu(dx) = 2x^{-2}dx$. Now for each $\varepsilon > 0$, we define S_ε the linear operator mapping each real function f to the function $S_\varepsilon(f) : t \mapsto \varepsilon^{-1}f(\varepsilon^{-1}t)$. The next statement ensures that if we zoom out on the Kingman comb restricted to a small interval, we end up with the Brownian CPP.

Proposition 3.8 ([27]). *As $\varepsilon \rightarrow 0$, the convergence $S_\varepsilon(K) \implies C$ holds weakly in law on every compact set.*

The genealogy generated by a comb restricted to a small interval focuses with high probability on parts of the ultrametric tree which are closely related (small genealogical distances). Now conditional on K , let us sample n points conditioned to have a time to mrca smaller than ε , and consider the genealogy of the whole subtree spanned by this sample (quenched conditional sampling). In other words, we let $0 = x_0^\varepsilon < x_1^\varepsilon < \dots < x_{N_\varepsilon}^\varepsilon < 1 = x_{N_\varepsilon+1}^\varepsilon$ denote the ranked enumeration of the finite set $\{K > \varepsilon\}$ and we let B_i^ε denote the comb restricted to the i -th interval of the subdivision ($1 \leq i \leq N_\varepsilon + 1$)

$$B_i^\varepsilon(x) = K(x_{i-1}^\varepsilon + x) \quad 0 < x < l_i^\varepsilon,$$

where $l_i^\varepsilon := x_i^\varepsilon - x_{i-1}^\varepsilon$ is the length of the i -th interval. Then for any measurable $F : \Omega \times (0, \infty) \rightarrow [0, \infty)$ (where Ω denotes the set of combs endowed with the vague topology), we define conditional on K

$$\mu_K^\varepsilon(F) := \frac{\sum_{i=1}^{N_\varepsilon} (l_i^\varepsilon)^n F(B_i^\varepsilon, l_i^\varepsilon)}{\sum_{i=1}^{N_\varepsilon} (l_i^\varepsilon)^n}$$

We naturally extend the definition of S_ε to bivariate F , by $S'_\varepsilon : (f, l) \mapsto (S_\varepsilon(f), \varepsilon^{-1}l)$.

Theorem 3.9 ([27]). *For any continuous $F : \Omega \times (0, \infty) \rightarrow [0, \infty)$,*

$$\lim_{\varepsilon \downarrow 0} \mathbb{E}(\mu_K^\varepsilon(F \circ S'_\varepsilon)) = \mathbb{E}(F(M, L)),$$

where L is a Gamma $(n+1, 2)$ r.v. and M is an independent CPP with intensity measure $\nu(dx) \mathbb{1}_{x < 1}$ restricted to the interval $[0, L]$.

In words, the previous statement ensures that after proper rescaling and averaging over the Kingman coalescent, the subtree spanned by a quenched conditional sample can be described in terms of a n -size-biased (i.e., biased by the n -th power of its size) Brownian CPP with height 1.

4. ULTRAMETRIC TREES WITH NEUTRAL MUTATIONS

4.1. Throwing point mutations on a tree

As explained in the Introduction, phylogenetic trees are inferred from genetic distances, due to the existence of a so-called molecular clock which regulates the pace at which new mutations appear on the lineages of the tree. This provides statistical relationships between genealogical distances and genetic distances, that allow biologists to infer the former from the latter.

So we consider a point measure M on the skeleton of a real tree (\mathfrak{t}, d) , that we will call mutation point measure, whose atoms are viewed as mutation events. Assume that each point x of the tree is further given a type, or **allele**, inherited from the most recent atom of M on $[[\rho, x]]$, that is, the point

$$\sigma(x) := \arg \max\{d(\rho, y) : y \text{ atom of } M, y \preceq x\},$$

which is set equal to ρ if the above set is empty. This assumption is known as the **infinitely-many allele model**. A point which carries the same allele as the root will be said **clonal**. The partition of the boundary

into distinct alleles is the so-called **allelic partition**.

There are usually two ways of studying the allelic partition in the framework of combs.

One possibility is to consider the allelic partition of a sample of size n . This can readily be done as in Subsection 3.1, by considering n r.v. V_1, \dots, V_n i.i.d. uniform in the interval I of definition of the comb, and associate each i with the most recent atom of the lineage $\llbracket \rho, V_i \rrbracket$. This induces a partition of $\{1, \dots, n\}$ which can be described by the so-called **allele frequency spectrum** $(A(1), \dots, A(n))$, where $A(k)$ is the number of blocks of the partition with cardinality k .

A second possibility is to consider the allelic partition of the whole population. If the comb has finite support (finite number of teeth), then one can proceed as previously. Otherwise, the allele frequency spectrum has to be expressed thanks to the measure ℓ defined on the boundary, by defining the point measure on $(0, \infty)$

$$A(dx) := \sum_m \delta_{\ell(R_m)}(dx),$$

where the sum is taken over all atoms m of M and R_m denotes the set of points x in the boundary such that $\sigma(x) = m$.

In the context of the molecular clock, the most natural way of modeling mutation events is to use Poisson point processes. We have seen that a locally compact \mathbb{R} -tree (\mathfrak{t}, d) has a length measure λ , so for any non-negative Borel function $\mu(t)$ (mutation rate at time t), we could define the mutation point measure as the Poisson point measure with intensity $\mu(d(\rho, x)) \lambda(dx)$. But if we are only interested in the genetic composition of the population of individuals/species co-existing at the same time, we can restrict our attention to ultrametric trees and in virtue of Theorem 2.6, focus on the tree $\tau_f(T)$ associated with a comb f defined on the interval $I = [0, a]$. Let μ denote a diffuse Radon measure on $(0, T]$. From now on, we define the **mutation point measure** M on Sk as the Poisson point measure with intensity measure

$$\sum_{t \in [0, a]: f(t) \neq 0} \delta_t(dx) \otimes \mu(dy) \mathbb{1}_{y < f(x)},$$

where we assume that $f(0) = T$ so that there can also be mutations on the **origin branch** L_0 . We will refer to the measure μ as the **mutation rate**. Usually, $\mu(dx)$ is taken equal to θdx , except when the comb studied is the image of an infinite tree by some map φ , as in Subsection 3.3.

4.2. Kingman coalescent with mutations

Here, we consider that the comb f is the Kingman comb and that $\mu(dx) = \theta dx$. The well-known results reviewed here are wonderfully exposed in [10].

Let us focus first on the allelic partition of a sample of n individuals as defined in the previous subsection. Let $A_n(k)$ denote the number of blocks of the allelic partition containing k elements. Observe that we must have $\sum_{k=1}^n k A_n(k) = n$.

Theorem 4.1 (Ewens sampling formula [13]). *The random vector $(A_n(1), \dots, A_n(n))$ has the same law as the random vector (Y_1, \dots, Y_n) conditional on $\sum_{k=1}^n k Y_k = n$, where the Y_k 's are independent, and Y_k is a Poisson r.v. with parameter θ/k . For any vector (a_1, \dots, a_n) such that $\sum_{k=1}^n k a_k = n$,*

$$\mathbb{P}(A_n(1) = a_1, \dots, A_n(n) = a_n) = c_{\theta, n} \prod_{k=1}^n \frac{\left(\frac{\theta}{k}\right)^{a_k}}{a_k!}$$

where

$$c_{\theta, n} := \frac{n!}{\theta(\theta+1)\cdots(\theta+n-1)}$$

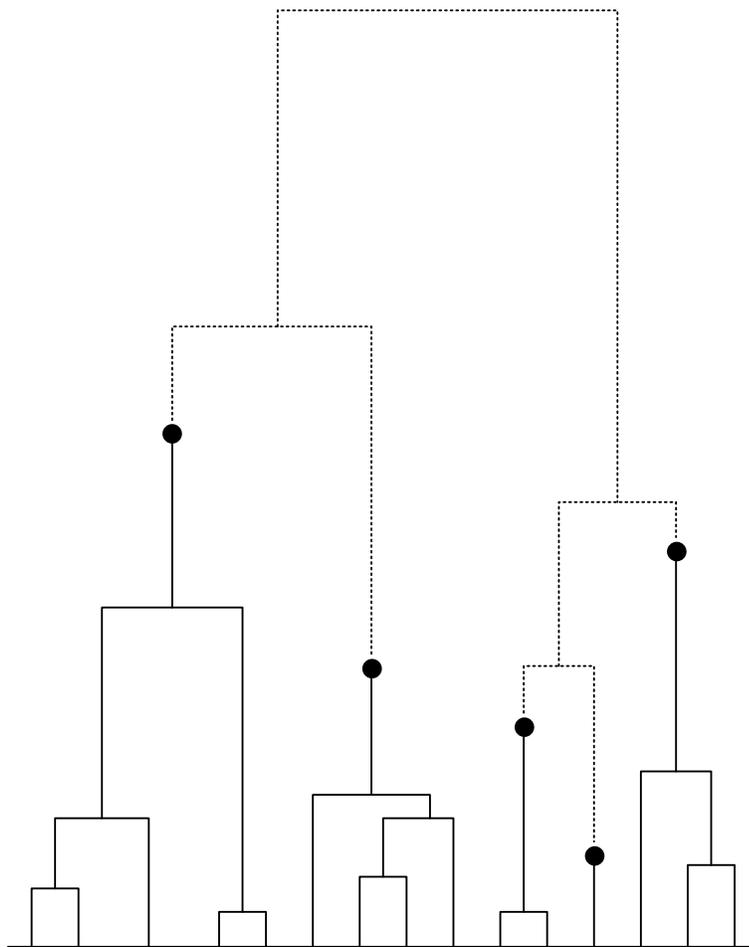


FIGURE 7. Coalescent with mutations. Backward time flows upwards, black dots symbolize mutation events. Lineages further in the past of mutations are dotted to recall they are independent of the frequency spectrum, see proof of Theorem 4.1. In the present example, the resulting allele frequency spectrum is $(1, 1, 1, 1, 1, 0, 0, \dots)$ (one block of each cardinality up to 5).

Proof. The trick is to follow the lineages of the n individuals backwards in time. It is known (but not obvious from Definition 3.3) that regardless of mutations, each pair of lineages coalesces at unit rate independently. Now each lineage is hit independently by a mutation at constant rate θ . Once it is hit, mutations occurring further in the past can be ignored because they have no consequence on the allelic partition at present time, thanks to the infinite-allele assumption. In addition, by the sampling consistency of the Kingman coalescent, the lineage itself can be frozen and put aside because its presence has no consequence on the process of coalescences of other lineages. So we are left with a Markov process in backward time where pairs of lineages coalesce at unit rate and single lineages are frozen at rate θ . Once all lineages have been frozen by a mutation, each of these mutations corresponds to a distinct allele and the maximum number of lineages downstream from this mutation is the size of the block corresponding to this allele. See Figure 7 for an illustration.

Now it is not difficult to see that changing the arrow of time once again and changing time, this process becomes a pure-birth process with immigration started at 0 and stopped when it hits n . It has unit birth rate and immigration rate θ . So the allele frequency spectrum has the same law as the vector of sizes of immigrant families after the process is stopped. For the calculations ending the proof see [22] or the references therein. \square

Corollary 4.2. *As $n \rightarrow \infty$, the random vector $(A_n(1), \dots, A_n(n))$ converges in the sense of finite-dimensional distributions to a sequence $(Y_k)_{k \geq 1}$ of independent r.v., where Y_k is a Poisson r.v. with parameter θ/k .*

This limiting spectrum of **small families** (i.e., blocks with size $O(1)$) is sometimes called the *harmonic frequency spectrum*. Besides the dependence in $1/k$, it is important to remember that the number of small families is $O(1)$, which is in deep contrast with what happens for coalescent point processes (see next subsection).

Now to understand the behavior of the **large families** (the other end of the spectrum), we set $X_n(i)$ the size of the i -th largest block of the allelic partition. The following result follows from the description of the allelic partition in terms of a pure-birth process with immigration, as in the proof of Theorem 4.1.

Proposition 4.3 ([6]). *The random vector $(n^{-1}X_n(k))_{k \geq 1}$ converges in the sense of finite-dimensional distributions to the sequence $(P_k)_{k \geq 1}$ defined as*

$$P_k = Z_k \prod_{i=1}^{k-1} (1 - Z_i),$$

where the (Z_i) are i.i.d. with density $\theta(1-z)^{\theta-1}$ (Beta $(1, \theta)$).

In words, the sizes of the largest blocks are of the order of the sample size n , and they represent fractions of the sample size which converge to the GEM distribution with parameter θ , distribution that can be obtained by a recursive stick-breaking procedure of the unit interval.

One could also tackle directly the problem on the Kingman comb, by considering the measures, say $X_k := \ell(R_k)$, of the largest blocks $(R_k)_{k \geq 1}$ of the allelic partition of the whole population. One would theoretically expect (X_k) to follow the GEM distribution. Another interesting question is to study the geometry of the subsets (R_k) of the unit interval analogously to what is done below in Theorem 4.4.

4.3. CPP with mutations

CPPs are less prone than the Kingman comb to calculations based on a sample of a fixed size. On the other hand, it can be shown that the genealogy of a sample of points thrown on the boundary according to a Poisson point process with constant intensity is equal in distribution to a new CPP, with finite intensity measure [20]. Modulo a modification of ν , we can therefore focus on the allelic partition of the whole population without loss of generality.

From now on, we assume that the comb f is a CPP killed at its first atom with second component larger than T , and with intensity ν , a diffuse measure on $(0, \infty)$ such that $\bar{\nu}(x) := \nu([x, \infty)) < \infty$ for all $x > 0$. We additionally assume that the mutation intensity measure μ is a diffuse Radon measure on $[0, \infty)$. In particular, we can define $\underline{\mu}(t) := \mu([0, t])$. Then it can be seen [8] that if $\int_0 \underline{\mu}(x)\nu(dx) < \infty$ the total number of mutations is finite a.s. whereas if $\int_0 \underline{\mu}(x)\nu(dx) = \infty$ then the number of mutations in any clade (set of all descendants of a point) is infinite a.s.

Recall that a point of I is clonal if it carries the same allele as the root, that is, there is no mutation on its lineage. The next statement is based on the idea that for any clonal t , the right-hand side of t only ‘sees’ the mutation-free lineage of t .

Theorem 4.4 ([8]). *Let $T = \infty$ so that $I = [0, \infty)$ and let R denote the closure of the clonal set. Conditional on the absence of mutation on the origin branch L_0 , R is a regenerative set that can be described as the range of a subordinator whose Laplace exponent ϕ is given by*

$$\frac{1}{\phi(\lambda)} = \int_{(0, \infty)} \frac{e^{-\mu(x)}}{\lambda + \bar{\nu}(x)} \mu(dx).$$

In addition, there exist semi-explicit formulae (that we will not provide explicitly here) for the **expectation of the allele frequency spectrum**, namely for

$$\Lambda_T(dx) := \mathbb{E}(A_T(dx)),$$

where T is here to recall the dependence on the height of the CPP, and the measure Λ_T is defined on \mathbb{N} when ν is finite [20] and on $(0, \infty)$ when ν is infinite [8]. In addition, the following convergence holds weakly on $[0, \infty)$

$$\Lambda(dx) := \lim_{T \rightarrow \infty} \frac{\Lambda_T(dx)}{\mathbb{E}(a(T))},$$

where $a(T)$ is the width of the CPP killed at its first atom with second component larger than T (in particular, $\mathbb{E}(a(T)) = 1/\nu([T, \infty))$). In the case of a unit rate critical birth-death process with mutations at rate θ , we get [20]

$$\Lambda(\{k\}) = \frac{\theta}{k} (1 + \theta)^{-k} \quad k \geq 1.$$

Analogously, in the case of the Brownian CPP with intensity measure $\frac{dx}{x^2}$ and mutation measure θdx , elementary calculations show that

$$\Lambda(dx) = \frac{\theta}{x} e^{-\theta x} dx \quad x > 0.$$

The last two formulae are reminiscent of the harmonic frequency spectrum displayed in Corollary 4.2.

Next, a LLN type of argument entails the following result.

Theorem 4.5 ([4, 8]). *Let $\bar{A}_T(q)$ denote the number of alleles whose carriers on the boundary form a set of measure larger than q (counting measure in the finite case, ℓ otherwise). Then the following limit holds a.s.*

$$\lim_{T \rightarrow \infty} \frac{\bar{A}_T(q)}{a(T)} = \Lambda([q, \infty)).$$

A similar convergence will be mentioned in the next subsection for the allele frequency spectrum at time t of the supercritical branching process as $t \rightarrow \infty$.

4.4. Supercritical processes with mutations

Here, we wish to explain how the allelic partition at the boundary of a supercritical branching process can be understood from the allelic partition of a CPP, in the vein of Theorem 3.7. We also say a word on the limiting allele frequency spectrum for large times.

Let \mathfrak{t} be the infinite tree generated by a supercritical birth-death process with birth intensity measure β conditioned on non extinction and let $q(t)$ denote the probability of extinction of a particle born at time t (which depends on some unspecified, possibly 0, possibly age-dependent death rate). Also assume that conditional on \mathfrak{t} , mutations occur on the lineages of \mathfrak{t} according to some intensity measure μ on $[0, \infty)$. Also recall that there is a measure \mathcal{L} defined on the boundary of \mathfrak{t} by

$$\mathcal{L}(B_u) := \lim_{t \uparrow \infty} \frac{\tilde{N}_u(t)}{e^{\tilde{\beta}(t)}},$$

where B_u denotes the set of infinite sequences v with prefix u , $\tilde{N}_u(t)$ is the number of descendants of u at time t which have infinite descendance and $\tilde{\beta}(t) = \int_0^t (1 - q(s)) \beta(ds)$.

Now from Theorem 3.7 and the remarks following it, if $\varphi : [0, \infty) \rightarrow (0, T]$ is a decreasing bijection, then $\varphi(\mathfrak{t})$ is isometric to the tree τ_f constructed from a CPP f with intensity ν and height T , naturally measured by ℓ , where

$$\nu([t, \infty)) = \exp(\tilde{\beta} \circ \varphi^{-1}(t)) \quad t \in [0, T].$$

Now it is straightforward to see that the mutations on τ_f occur according to a mutation point measure with mutation rate $\mu_\varphi := \mu \circ \varphi^{-1}$.

If we assume that μ has finite mass θ and we take φ given by $\varphi(t) = \theta^{-1}\mu([t, \infty))$, then the allelic partition at the boundary is equal in distribution to that of a CPP (with intensity ν previously displayed) with height 1 and mutations at constant rate θ .

On the contrary, if we take $\mu(dx) = \theta dx$ and φ given by $\varphi(t) = e^{-\tilde{\beta}(t)}$, then $\varphi(\mathfrak{t})$ has the distribution of τ_f where f is a Brownian CPP with height 1, intensity $\frac{dx}{x^2}$ and mutation rate μ_φ satisfying $\mu_\varphi[(x, 1]) = \theta\varphi^{-1}(x)$. In particular, if the birth-death process is time-homogeneous with birth rate b and if we set $a := b(1 - q)$, then $\varphi(t) = e^{-at}$, so that $\varphi^{-1}(t) = -\ln(t)/a$. As a conclusion, the allelic partition at the boundary of a time-homogeneous supercritical birth-death process with mutations at constant rate θ is equal in distribution to that of a **Brownian CPP** with height 1, intensity $\frac{dx}{x^2}$ and mutations at inhomogeneous rate $\theta/(at)$.

Similarly to Theorem 4.5, the allele frequency spectrum at time t of the supercritical birth-death process properly rescaled converges a.s. as $t \rightarrow \infty$. The LLN type of argument invoked here is known as the theory of branching processes **counted with random characteristics** [15–17, 33]. In our setting, the random characteristic of individual i , say, can be for example the number $\chi_i^k(t)$ of mutations that i has experienced during her lifetime and which are carried by k alive individuals, t units of time after her birth ($\chi_i(t) = 0$ if $t < 0$). Then the total number $A_t(k)$ of alleles carried by k individuals at time t (except possibly the ancestral type) is the sum $\sum_i \chi_i(t - \alpha_i)$ over all individuals i (dead or alive), where α_i is the birth time of individual i . The theory of branching processes counted with random characteristics ensures that these sums rescaled by e^{at} converge a.s. on the survival event.

This method has been used extensively in [37]. Further refinements have been obtained in [5] thanks to the use of coalescent point processes. For example, we have shown the convergence of the (properly rescaled) largest blocks of the partition at time t , as $t \rightarrow \infty$. In particular, we showed, letting $N \equiv N_t$ denote the total population size at time t , that the largest blocks are roughly of size $N^{\theta/a}$ when $\theta < a$, $\log(N)^2$ when $a = \theta$ and $\log(N)$ when $\theta > a$. Similarly, we showed that the oldest mutations with alive carriers at time t appeared roughly at time $O(1)$ when $\theta < a$, at time $\ln(t)/a$ when $\theta = a$ and at time $(1 - a/\theta)t$ when $a < \theta$. Notice that in contrast to the Kingman case, no block has a size of the order of the total population size when the mutation rate is constant. It is an open question to investigate whether the compactification mentioned earlier in this subsection of the tree into a CPP can shed extra light on these results. In particular, the case of large families is not straightforward at all since the compactification focuses precisely on subsets of the boundary with measure of the order of the total population size.

REFERENCES

- [1] David Aldous and Lea Popovic. A critical branching process model for biodiversity. *Advances in Applied Probability*, 37(4):1094–1115, 2005.
- [2] Helen K. Alexander, Amaury Lambert, and Tanja Stadler. Quantifying age-dependent extinction from species phylogenies. *Systematic Biology*, 65(1):35, 2015.
- [3] Jean Bertoin and Jean-François Le Gall. Stochastic flows associated to coalescent processes. iii. limit theorems. *Illinois Journal of Mathematics*, 50(1-4):147–181, 2006.
- [4] Nicolas Champagnat and Amaury Lambert. Splitting trees with neutral Poissonian mutations I: Small families. *Stochastic Processes and their Applications*, 122(3):1003–1033, 2012.
- [5] Nicolas Champagnat and Amaury Lambert. Splitting trees with neutral Poissonian mutations II: Largest and Oldest families. *Stochastic Processes and their Applications*, 123(4):1368–1414, 2013.

- [6] Peter Donnelly and Simon Tavaré. The ages of alleles and a coalescent. *Adv. in Appl. Probab.*, 18(1):1–19, 1986.
- [7] Andreas Dress, Vincent Moulton, and Werner Terhalle. T-theory: An overview. *European Journal of Combinatorics*, 17(2–3):161–175, February 1996.
- [8] Jean-Jil Duchamps and Amaury Lambert. Mutations on a random binary tree with measured boundary. *Eprint arXiv:1701.07698*, 2017.
- [9] Thomas Duquesne. The coding of compact real trees by real valued functions. *arXiv:math/0604106*, April 2006. arXiv: math/0604106.
- [10] R. Durrett. *Probability models for DNA sequence evolution*. Springer Verlag, 2008.
- [11] Rampal S. Etienne, Hélène Morlon, and Amaury Lambert. Estimating the duration of speciation from phylogenies. *Evolution*, 68(8):2430–2440, April 2014.
- [12] Steven Neil Evans. *Probability and Real Trees: École d’été de Probabilités de Saint-Flour XXXV-2005*. Springer, 2008.
- [13] Warren J. Ewens. The sampling theory of selectively neutral alleles. *Theoretical Population Biology*, 3:87–112; erratum, *ibid.* 3 (1972), 240; erratum, *ibid.* 3 (1972), 376, 1972.
- [14] Jochen Geiger and Götz Kersting. Depth-first search of random trees, and Poisson point processes. In *Classical and modern branching processes (Minneapolis, MN, 1994)*, volume 84 of *IMA Vol. Math. Appl.*, pages 111–126. Springer, New York, 1997.
- [15] Peter Jagers and Olle Nerman. The growth and composition of branching populations. *Adv. in Appl. Probab.*, 16(2):221–259, 1984.
- [16] Peter Jagers and Olle Nerman. Limit theorems for sums determined by branching and other exponentially growing processes. *Stochastic Process. Appl.*, 17(1):47–71, 1984.
- [17] Peter Jagers and Olle Nerman. The asymptotic composition of supercritical multi-type branching populations. In *Séminaire de Probabilités, XXX*, volume 1626 of *Lecture Notes in Math.*, pages 40–54. Springer, Berlin, 1996.
- [18] J.F.C. Kingman. The coalescent. *Stochastic processes and their applications*, 13(3):235–248, 1982.
- [19] Amaury Lambert. Population dynamics and random genealogies. *Stochastic Models*, 24(suppl. 1):45–163, 2008.
- [20] Amaury Lambert. The allelic partition for coalescent point processes. *Markov Processes and Related Fields*, 15(3):359–386, 2009.
- [21] Amaury Lambert. The contour of splitting trees is a Lévy process. *The Annals of Probability*, 38(1):348–395, 2010.
- [22] Amaury Lambert. Species abundance distributions in neutral models with immigration or mutation and general lifetimes. *Journal of mathematical biology*, 63(1):57–72, 2011.
- [23] Amaury Lambert. Probabilistic models for the subtrees of life. *Brazilian Journal of Probability and Statistics* (in press) Eprint arXiv:1603.03705, 2017.
- [24] Amaury Lambert, Helen K. Alexander, and Tanja Stadler. Phylogenetic analysis accounting for age-dependent death and sampling with applications to epidemics. *Journal of Theoretical Biology*, 352:60–70, July 2014.
- [25] Amaury Lambert, Hélène Morlon, and Rampal S. Etienne. The reconstructed tree in the lineage-based model of protracted speciation. *Journal of Mathematical Biology*, 70(1):367–397, 2015.
- [26] Amaury Lambert and Lea Popovic. The coalescent point process of branching trees. *Annals of Applied Probability*, 23(1):99–144, 2013.
- [27] Amaury Lambert and Emmanuel Schertzer. Recovering the Brownian coalescent point process from the Kingman coalescent by conditional sampling. *Bernoulli* (in press) Eprint arXiv:1611.01323, 2017.
- [28] Amaury Lambert and Tanja Stadler. Birth–death models and coalescent point processes: The shape and probability of reconstructed phylogenies. *Theoretical Population Biology*, 90:113–128, December 2013.
- [29] Amaury Lambert and Geronimo Uribe Bravo. Totally ordered, measured trees and splitting trees with infinite variation. *Eprint arXiv:1607.02114*, 2016.
- [30] Amaury Lambert and Geronimo Uribe Bravo. The comb representation of compact ultrametric spaces. *p-Adic Numb. Ultrametric Anal. Appl.*, 9(1):22–38, 2017.
- [31] Russell Lyons. Equivalence of boundary measures on covering trees of finite graphs. *Ergodic Theory and Dynamical Systems*, 14(03):575–597, September 1994.
- [32] Hélène Morlon, Todd L. Parsons, and Joshua B. Plotkin. Reconciling molecular phylogenies with the fossil record. *Proceedings of the National Academy of Sciences*, 108(39):16327–16332, 2011.
- [33] Olle Nerman. On the convergence of supercritical general (C-M-J) branching processes. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 57(3):365–395, 1981.
- [34] Lea Popovic. Asymptotic genealogy of a critical branching process. *Annals of Applied Probability*, pages 2120–2148, 2004.
- [35] Charles Semple and Mike A Steel. *Phylogenetics*, volume 24. Oxford University Press, 2003.
- [36] Tanja Stadler. Mammalian phylogeny reveals recent diversification rate shifts. *Proceedings of the National Academy of Sciences*, 108(15):6187–6192, December 2011.
- [37] Ziad Taib. *Branching processes and neutral evolution*. Lecture Notes in Biomathematics. 93. Berlin: Springer-Verlag. viii, 112 p., 1992.