Compact metric measure spaces and Λ-coalescents coming down from infinity

Holger F. Biehler and Peter Pfaffelhuber

Albert-Ludwigs-Universität Freiburg
Abteilung Mathematische Stochastik
Eckerstr.1,
D-79104 Freiburg, Germany.
E-mail address: holger.biehler@web.de, p.p@stochastik.uni-freiburg.de
URL: http://www.stochastik.uni-freiburg.de/~pfaffelh

Abstract. We study topological properties of random metric spaces which arise by Λ-coalescents. These are stochastic processes, which start with an infinite number of lines and evolve through multiple mergers in an exchangeable setting. We show that the resulting Λ-coalescent measure tree is compact iff the Λ-coalescent comes down from infinity, i.e. only consists of finitely many lines at any positive time. If the Λ-coalescent stays infinite, the resulting metric measure space is not even locally compact. Our results are based on general notions of compact and locally compact (isometry classes of) metric measure spaces. In particular, we give characterizations for general (random) metric measure spaces to be (locally) compact using the Gromov-weak topology.

1. Introduction

Metric structures arise frequently in probability theory. Prominent examples are random trees (e.g. Aldous (1993); Evans and O'Connell (1994); Le Gall (1999); Berestycki (2009)), where the distance between two points is given by the length of the shortest path connecting the points. A class of random trees is given by coalescent processes, where a subset of an infinite number of *lines* can merge and the distance of two leaves is proportional to the coalescence time (Kingman (1982); Pitman (1999); Aldous (1999); Schweinsberg (2000b); Evans (2000)). The complexity of this class of processes is properly described by the concepts of Λ-coalescent, where any set of lines can merge to a single line (a multiple collision, Pitman (1999)) and Ξ-coalescents, where any set of lines can merge to several lines at the same time (a simultaneous multiple collision, Schweinsberg (2000a)). The resulting metric space

has so far mostly been studied in the simplest case, where only binary mergers are allowed, the Kingman-coalescent (Kingman, 1982; Evans, 2000).

Analyzing metric structures requires geometrical and topological foundations. In the context of Riemannian geometry, such foundations have already been laid by Gromov, summarized in his book (Gromov (1999), see also Vershik (1998); Burago et al. (2001)). These authors study convergence of (isometry classes of) compact metric spaces by the notion of Gromov-Hausdorff convergence. In addition, Gromov introduced a topology on the space of (isometry classes of) metric measure spaces (mm-spaces, for short), which are metric spaces equipped with a measure. We will call this the Gromov-weak topology in the sequel (see also Greven et al. (2009)).

In probability theory, results on weak convergence and stochastic process theory require that the underlying space is Polish. In addition, a characterization of the compact sets is required in order to show tightness. These concepts have been worked out based on Gromov's notions by Evans et al. (2006) and Greven et al. (2009).

The goal of the present paper is as follows: we concentrate on the spaces of *locally compact* and *compact* mm-spaces and give a characterization of these (see Theorems 2.10 and 2.15). In addition, we apply these general results to random mm-spaces (Λ-coalescent measure trees) which arise in connection to Λ-coalescents. Recall that Λ-coalescents fall into one of two categories, depending on Λ. Either a Λ-coalescent comes down from infinity, meaning that it can be started with an infinite number of lines and only finitely many are left at any positive time, or it stays infinite for all times (see Proposition 23 in Pitman (1999)). The proof of the following result is given in Section 4.

Theorem 1.1 (Coming down from infinity and compactness). *Let Λ be a finite measure on $[0,1]$ and $(\Pi_t)_{t \geq 0}$ the corresponding Λ-coalescent. Moreover, $\mathcal{L}$ is the associated Λ-coalescent measure tree, taking values in the space of mm-spaces. Then the following is equivalent.*

(1) *$(\Pi_t)_{t \geq 0}$ comes down from infinity, i.e. $\#\Pi_t < \infty$ almost surely, for all $t > 0$.*

(2) *$\mathcal{L}$ is compact, almost surely.*

If (1) (or 2) does not hold, $\mathcal{L}$ is not even locally compact.

We proceed as follows: In Section 2 we develop our general theory on compact and locally compact isometry classes of metric measure spaces. Section 3 contains a short introduction to Λ-coalescent measure trees. Finally, the proof of Theorem 1.1 is given in Section 4. We remark that the application of (locally) compact mm-spaces is not restricted to trees. For example, it is possible to study large random planar maps, as given in Le Gall (2007), or random Graphs (e.g. the Erdős-Renyi random graph, Addario-Berry et al. (2012)), by our notions.

2. Metric measure spaces

We start with some notation. Our main result, the characterization of compact and locally compact mm-spaces, is given in Theorems 2.10 and 2.15.

Remark 2.1 (Notation). As usual, given a topological space $(X, \mathscr{O}_X)$, we denote by $\mathscr{M}_1(X)$ the space of all probability measures on the Borel-σ-algebra $\mathscr{B}(X)$. The *support* of $\mu \in \mathscr{M}_1(X)$, supp(μ), is the smallest closed set $X_0 \subseteq X$ such that

$\mu(X \setminus X_0) = 0$. The *push-forward* of μ under a measurable map φ from X into another topological space, $(Z, \mathcal{O}_Z)$, is the probability measure $\varphi_*\mu \in \mathcal{M}_1(Z)$ defined for all $A \in \mathcal{B}(Z)$ by $\varphi_*\mu(A) := \mu(\varphi^{-1}(A))$. We denote weak convergence in $\mathcal{M}_1(X)$ by $\Rightarrow$.

Definition 2.2 (Metric measure and mm-spaces)**.**

   (1) A *metric measure space* is a triple (X, r, μ) such that $X \subseteq \mathbb{R}$, (X, r) is a complete and separable metric space which is equipped with a probability measure μ on $\mathcal{B}(X)$. We say that (X, r, μ) and (X', r', μ') are *measure-preserving isometric* if there exists an isometry φ between $\mathrm{supp}(\mu) \subseteq X$ and $\mathrm{supp}(\mu') \subseteq X'$ such that $\mu'|_{\mathrm{supp}(\mu')} = \varphi_*(\mu|_{\mathrm{supp}(\mu)})$. It is clear that the property of being measure-preserving isometric is an equivalence relation.

   (2) The equivalence class of the metric measure space (X, r, μ) is called the mm-space of (X, r, μ) and is denoted $\overline{(X, r, \mu)}$. The set of mm-spaces is denoted $\mathbb{M}$ and generic elements are $\chi, y, ...$

   (3) An mm-space $\chi \in \mathbb{M}$ is *(locally) compact* if there is $(X, r, \mu) \in \chi$ such that (X, r) is (locally) compact. The space of (locally) compact mm-spaces is denoted $\mathbb{M}_c$ ($\mathbb{M}_{lc}$).

Following Greven et al. (2009), we equip $\mathbb{M}$ with the Gromov-weak topology as follows.

Definition 2.3 (Gromov-weak topology)**.** For a metric space (X, r) define

$$R^{(X,r)} : \begin{cases} X^{\mathbb{N}} & \to \mathbb{R}_+^{\binom{\mathbb{N}}{2}} \\ (x_i)_{i \in \mathbb{N}} & \mapsto (r(x_i, x_j))_{1 \le i < j} \end{cases}$$

the map which sends a sequence of points in X to its distance matrix and for an mm-space $\chi = \overline{(X, r, \mu)}$ we define the *distance matrix distribution* by

$$\nu^\chi := (R^{(X,r)})_* \mu^{\otimes \mathbb{N}} \in \mathcal{M}_1(\mathbb{R}_+^{\binom{\mathbb{N}}{2}}),$$

where $\mu^{\otimes \mathbb{N}}$ is the infinite product measure of μ, where $\mathbb{R}_+^{\binom{\mathbb{N}}{2}}$ is equipped with the product σ-field. We say that a sequence $\chi_1, \chi_2, ... \in \mathbb{M}$ converges *Gromov-weakly* to $\chi \in \mathbb{M}$ if

$$\nu^{\chi_n} \xRightarrow{n \to \infty} \nu^\chi.$$

Note that ν^χ does not depend on the representative $(X, r, \mu) \in \chi$, hence is well-defined.

Remark 2.4 (When is a random mm-space compact?). Recall from Theorem 1 of Greven et al. (2009) that the space $\mathbb{M}$, equipped with the Gromov-weak topology, is Polish. Hence, $\mathbb{M}$ allows to use standard tools from probability, e.g. from the theory of weak convergence.

    In order to show that a random variable taking values in $\mathbb{M}$ is supported by the space of locally compact or compact mm-spaces, there are two strategies, formulated here in the case of compact mm-spaces:

    Either, consider the Gromov-weak topology on $\mathbb{M}_c$. Defining an approximating sequence in $\mathbb{M}_c$ and showing that the sequence is tight in $\mathbb{M}_c$ ensures compactness of the limiting object. Note that any mm-space can be approximated by finite (hence compact) mm-spaces, so $\mathbb{M}_c$ is not closed in $\mathbb{M}$. So, this approach amounts

to knowing the compact sets in $\mathbb{M}_c$. See Proposition 6.2 of Greven et al. (2012) for an example.

Our application to the Λ-coalescent measure tree in Section 4 relies on a different approach. It is possible to give handy characterizations of compact mm-spaces; see Theorem 2.10. Hence, if we are given a random variable taking values in $\mathbb{M}$ through a sequence of mm-spaces, it is possible to check directly if the limiting object is compact.

Definition 2.5 (Distance distribution, Moduli of mass distribution). Let $\chi \in \mathbb{M}$. We set $\underline{r} := (r_{ij})_{1 \leq i < j} \in \mathbb{R}_+^{\binom{\mathbb{N}}{2}}$.

(a) Let $r : \mathbb{R}_+^{\binom{\mathbb{N}}{2}} \to \mathbb{R}_+$ be given by $r(\underline{r}) := r_{12}$. Then, the *distance distribution* is given by $w_\chi := r_* \nu^\chi$, i.e.,

$$w_\chi(\cdot) := \nu^\chi\{\underline{r} : r_{12} \in \cdot\}.$$

(b) For $\varepsilon > 0$, define $s_\varepsilon : \mathbb{R}_+^{\binom{\mathbb{N}}{2}} \to \mathbb{R}_+$ by

$$s_\varepsilon(\underline{r}) := \lim_{n \to \infty} \frac{1}{n} \sum_{j=1}^n 1_{\{r_{1j} \leq \varepsilon\}}$$

if the limit exists (and zero otherwise). Note that $s_\varepsilon(\underline{r})$ exists for ν^χ-almost all $\underline{r}$ by exchangeability and de Finetti's Theorem. For $\delta > 0$, the *moduli of mass distribution* are

$$v_\delta(\chi) := \inf\{\varepsilon > 0 : \ \nu^\chi\{\underline{r} : s_\varepsilon(\underline{r}) \leq \delta\} \leq \varepsilon\}$$

and

$$\widetilde{v}_\delta(\chi) := \inf\{\varepsilon > 0 : \ \nu^\chi\{\underline{r} : s_\varepsilon(\underline{r}) \leq \delta\} = 0\}.$$

Example 2.6 (Representatives of χ). Let $\chi = \overline{(X, r, \mu)}$. Without loss of generality we assume that $\mathrm{supp}(\mu) = X$. Since $\nu^\chi = (R^{(X,r)})_* \mu^{\otimes \mathbb{N}}$, we have that

$$w_\chi(\cdot) = \mu^{\otimes 2}\{(x, y) : r(x, y) \in \cdot\}.$$

Moreover,

$$\nu^\chi\{\underline{r} : s_\varepsilon(\underline{r}) \in \cdot\} = \mu\{x : \mu(B_\varepsilon(x)) \in \cdot\} \tag{2.1}$$

by construction, where $B_\varepsilon(x)$ is the closed ball of radius ε around x. This implies that

$$v_\delta(\chi) \leq \varepsilon \qquad \Longleftrightarrow \qquad \mu\{x : \mu(B_\varepsilon(x)) \leq \delta\} \leq \varepsilon.$$

In particular, $v_\delta(\chi) \leq \varepsilon$ means, that thin points (in the sense that $\mu(B_\varepsilon(x)) \leq \delta$) are rare (i.e. carry mass at most ε). Moreover,

$$\widetilde{v}_\delta(\chi) \leq \varepsilon \qquad \Longleftrightarrow \qquad \mu\{x : \mu(B_\varepsilon(x)) \leq \delta\} = 0.$$

This means that there are μ-almost surely no points which are too thin (in the sense that $\mu(B_\varepsilon(x)) \leq \delta$).

Definition 2.7 (Size of ε-separated set). Let $\underline{r} \in \mathbb{R}_+^{\binom{\mathbb{N}}{2}}$. For $\varepsilon > 0$, define the *maximal size of an ε-separated set* by

$$\xi_\varepsilon(\underline{r}) := \sup\left\{N \in \mathbb{N} : \ \exists k_1 < \ldots < k_N : (r_{k_i, k_j})_{1 \leq i < j \leq N} \in (\varepsilon, \infty)^{\binom{N}{2}}\right\}.$$

Lemma 2.8 (ξ_ε is constant, ν^χ-almost surely)**.** *Let* $\chi \in \mathbb{N}$ *and* $\varepsilon > 0$*. Then,* ξ_ε *is constant,* ν^χ-*almost surely and equals*

$$\xi_\varepsilon(\chi) := \inf \left\{ N \in \mathbb{N} : \nu^\chi \left(\rho_N^{-1} \left((\varepsilon, \infty)^{\binom{N}{2}} \right) \right) > 0 \right\},$$

where $\rho_N : \mathbb{R}^{\binom{\mathbb{N}}{2}} \to \mathbb{R}^{\binom{N}{2}}$ *is the projection on the first* $\binom{N}{2}$ *coordinates.*

Proof: Assume $\chi = \overline{(X, r, \mu)}$. Let $x_1, x_2, ... \in X$ be such that $\xi_\varepsilon((r(x_i, x_j))_{1 \le i < j}) = N$. Then, N is the maximal size of an ε-separated set in (X, r), $\mu^{\otimes \mathbb{N}}$-almost surely. All results follow, since $\nu^\chi = (R^{(X,r)})_* \mu^{\otimes \mathbb{N}}$ and since ν^χ is exchangeable. $\quad\square$

Remark 2.9 (Tightness in $\mathbb{M}$). Recall from Theorem 2 in Greven et al. (2009) that for any $\chi \in \mathbb{M}$, it holds that $v_\delta(\chi) \xrightarrow{\delta \to 0} 0$. Moreover, a set $\Gamma \subseteq \mathbb{M}$ is pre-compact iff $\{ w_\chi : \chi \in \Gamma \}$ is tight (as a family in $\mathcal{M}_1(\mathbb{R}_+)$) and $\sup_{\chi \in \Gamma} v_\delta(\chi) \xrightarrow{\delta \to 0} 0$.

This leads to a characterization of tightness for a family of random mm-spaces, see Greven et al. (2009), Theorem 3: Here, (the distributions of) a family $\{ \mathcal{X} : \mathcal{X} \in \Gamma \}$ of $\mathbb{M}$-valued random variables is tight iff $\{ \langle w_\mathcal{X} \rangle : \mathcal{X} \in \Gamma \}$ is tight (where $\langle w_\mathcal{X} \rangle$ is the first moment measure of $(w_\mathcal{X})_* \mathbf{P} \in \mathcal{M}_1(\mathcal{M}_1(\mathbb{R}_+))$ and $\sup_{\mathcal{X} \in \Gamma} \mathbf{E}[v_\delta(\mathcal{X})] \xrightarrow{\delta \to 0} 0$. Given a sequence of random mm-spaces, we can use these results in order to obtain limiting objects, at least along subsequences.

Now we come to a characterization of compact mm-spaces.

Theorem 2.10 (Compact mm-spaces)**.** *Let* $\chi \in \mathbb{M}$*. The following conditions are equivalent.*

(1) *The mm-space* χ *is compact, i.e.* $\chi \in \mathbb{M}_c$.
(2) *For all* $\varepsilon > 0$*, it holds that* $\xi_\varepsilon(\chi) < \infty$*.*
(3) *For all* $\varepsilon > 0$*, there is* $\delta > 0$ *such that* $\widetilde{v}_\delta(\chi) \le \varepsilon$*.*

The following characterization of random, almost surely compact mm-spaces is immediate.

Corollary 2.11 (Random compact mm-spaces)**.** *Let* $\mathcal{X}$ *be a random variable taking values in* $\mathbb{M}$*. The following conditions are equivalent.*

(1) *The mm-space* $\mathcal{X}$ *is compact, almost surely, i.e.* $\mathbf{P}(\mathcal{X} \in \mathbb{M}_c) = 1$*.*
(2) *For all* $\varepsilon > 0$*, it holds that* $\mathbf{P}(\xi_\varepsilon(\mathcal{X}) < \infty) = 1$*.*
(3) *For all* $\varepsilon > 0$*, there is a random variable* $\Delta > 0$ *with* $\mathbf{P}(\widetilde{v}_\Delta(\mathcal{X}) \le \varepsilon) = 1$*.*

Remark 2.12 (Size of ε-separated set and size of ε-covering). The following observation will be used in the proof of Theorem 2.10: Let (X, r) be a metric space and $\varepsilon > 0$, let ξ_ε be the maximal size of an ε-separated set and N_ε be the minimal number of ε-balls needed to cover (X, r). Then

$$N_\varepsilon \le \xi_\varepsilon \le N_{\varepsilon/2}.$$

In order to see this, let $x_1, ..., x_{\xi_\varepsilon}$ be a maximal ε-separated set. Then, $X = \bigcup_{i=1}^{\xi_\varepsilon} B_\varepsilon(x_i)$, since otherwise, we find $x \in X \setminus \left(\bigcup_{i=1}^{\xi_\varepsilon} B_\varepsilon(x_i) \right)$ and hence, the set is not maximal. This shows $N_\varepsilon \le \xi_\varepsilon$. For the second inequality, it is clear that $B_{\varepsilon/2}(x_1), ..., B_{\varepsilon/2}(x_{\xi_\varepsilon})$ are disjoint. Hence, any set of centers of $\varepsilon/2$-balls which cover (X, r) must hit each $B_{\varepsilon/2}(x_i)$ at least once. As a consequence, $\xi_\varepsilon \le N_{\varepsilon/2}$.

Proof of Theorem 2.10: Let $\chi = \overline{(X, r, \mu)}$. We use the notation laid out in Remark 2.6. In particular, recall (2.1).

(1) $\Rightarrow$ (2): Let χ be compact and $\varepsilon > 0$. Then (X, r) is totally bounded and there is $N_{\varepsilon/2} \in \mathbb{N}$ such that (X, r) can be covered by $N_{\varepsilon/2}$ balls of radius $\varepsilon/2$. Then we find $\xi_\varepsilon(\chi) \leq N_{\varepsilon/2} < \infty$ by the last remark.

(2) $\Rightarrow$ (3): Let $\varepsilon > 0$. The space (X, r) can be covered by $\xi_{\varepsilon/2}(\chi) < \infty$ balls of radius $\varepsilon/2$, again by the last remark. Let $x_1, \ldots, x_{\xi_{\varepsilon/2}}$ be centers of such balls and $\delta := \min\{\mu(B_{\varepsilon/2}(x_i)) : \mu(B_{\varepsilon/2}(x_i)) > 0\}$. Then $\delta > 0$. Now take any $x \in X$ and choose $i \in \{1, \ldots, \xi_{\varepsilon/2}\}$ such that $x \in B_{\varepsilon/2}(x_i)$. Then we have

$$\mu(B_\varepsilon(x)) \geq \mu(B_{\varepsilon/2}(x_i)) \geq \delta.$$

Hence,

$$\nu^\chi\{\underline{r} : s_\varepsilon(\underline{r}) \leq \delta\} = \mu\{x \in X : \mu(B_\varepsilon(x)) \leq \delta\} = 0.$$

(3) $\Rightarrow$ (1): It suffices to show that (X, r) is totally bounded. Let $\varepsilon > 0$. By assumption, there is $\delta > 0$ such that

$$\nu^\chi\{\underline{r} : s_\varepsilon(\underline{r}) \leq \delta\} = \mu\{x \in X : \mu(B_\varepsilon(x)) \leq \delta\} = 0.$$

We show that there is a finite maximal 2ε-separated set in X. For this, take a maximal 2ε-separated set $S \subseteq X$ (and without loss of generality assume that $\mathrm{supp}(\mu) = X$). Then, using the last remark,

$$1 = \mu(X) = \mu\Big(\bigcup_{x \in S} B_{2\varepsilon}(x)\Big) \geq \mu\Big(\bigcup_{x \in S} B_\varepsilon(x)\Big) = \sum_{x \in S} \mu(B_\varepsilon(x)) \geq |S| \cdot \delta,$$

since $\mu(B_\varepsilon(x)) > \delta$ holds μ-almost surely by assumption. Now, $|S| \leq 1/\delta < \infty$ and $\varepsilon > 0$ was arbitrary, so (X, r) is totally bounded. $\qquad\square$

Next, we come to a characterization of locally compact mm-spaces. Again some notation is needed.

Definition 2.13 (δ-restriction)**.** Let $\underline{r} := (r_{ij})_{1 \leq i < j} \in \mathbb{R}_+^{\binom{\mathbb{N}}{2}}$. Set $\widehat{\tau}_\delta(0) := 1$ and

$$\widehat{\tau}_\delta(i + 1) := \inf\{j > \widehat{\tau}_\delta(i) : r_{1j} \leq \delta\}.$$

Then,

$$\tau_\delta(\underline{r}) := (r_{\widehat{\tau}_\delta(i), \widehat{\tau}_\delta(j)})_{1 \leq i < j}$$

is called the *δ-restriction* of $\underline{r}$.

Remark 2.14 (δ-restriction for distance matrices.)**.** Let $\chi = \overline{(X, r, \mu)} \in \mathbb{M}$ and $x_1, x_2, \ldots \in X$. We note that $x_k \in \widehat{\tau}_\delta(\mathbb{N})$ iff $r(x_1, x_k) \leq \delta$. Hence, $\tau_\delta((r(x_i, x_j))_{1 \leq i \leq j})$ is the distance matrix distribution for points among $x_2, x_3, \ldots$ which have distance at most δ to x_1. So,

$$(\tau_\delta)_* \nu^\chi(\cdot) = \nu^\chi\{\tau_\delta(\underline{r}) \in \cdot\} = \nu^\chi\{\underline{r} \in \cdot | r_{12}, r_{13}, \ldots \leq \delta\}$$
$$= \mu^{\otimes \mathbb{N}}\{(r(x_i, x_j))_{1 \leq i < j} \in \cdot | r(x_1, x_j) \leq \delta \text{ for all } j = 2, 3, \ldots\}$$

Clearly, $(\tau_\delta)_* \nu^\chi$ is exchangeable, since ν^χ is exchangeable.

Theorem 2.15 (Locally compact mm-spaces)**.** *Let $\chi \in \mathbb{M}$. The following conditions are equivalent.*

(1) *The mm-space χ is locally compact, $\chi \in \mathbb{M}_{lc}$.*

(2) *It holds that*

$$\nu^{\chi}\Big(\bigcap_{0<\eta<\delta} \big\{\underline{\underline{r}} : \xi_{\eta}(\tau_{\delta}(\underline{\underline{r}})) < \infty\big\}\Big) \xrightarrow{\delta\to 0} 1.$$

Proof: Let $\chi = \overline{(X, r, \mu)}$. Then, χ is locally compact iff for μ-almost all $x \in X$ there is $\delta > 0$, such that the ball $B_{\delta}(x)$ can be covered by a finite number of balls with radius η, for all $0 < \eta < \delta$. Hence,

$$1 = \mu\Big(\bigcup_{\delta>0}\bigcap_{0<\eta<\delta} \big\{x : B_{\varepsilon}(x) \text{ can be covered by finitely many balls of radius } \eta\big\}\Big)$$

$$= \lim_{\delta\to 0} \mu\Big(\bigcap_{0<\eta<\delta} \big\{x : \text{ the maximal } \eta\text{-separated set in } B_{\delta}(x) \text{ is finite}\big\}\Big)$$

$$= \lim_{\delta\to 0} \mu^{\otimes\mathbb{N}}\Big(\bigcap_{0<\eta<\delta} \big\{(x_1, x_2, ...) : \xi_{\eta}((r_{x_i,x_j})_{2\leq i<j}) < \infty | r(x_1, x_2), r(x_1, x_3), ... < \delta\big\}\Big)$$

$$= \lim_{\delta\to 0} \mu^{\otimes\mathbb{N}}\Big(\bigcap_{0<\eta<\delta} \big\{(x_1, x_2, ...) : \xi_{\eta}(\tau_{\delta}((r_{x_i,x_j})_{1\leq i<j})) < \infty\big\}\Big)$$

$$= \lim_{\delta\to 0} \nu^{\chi}\Big(\bigcap_{0<\eta<\delta} \big\{\underline{\underline{r}} : \xi_{\eta}(\tau_{\delta}(\underline{\underline{r}})) < \infty\big\}\Big).$$

$\square$

3. Λ-coalescents

We come to the application of the general results from the last section to metric spaces which arise in the context of coalescents which allow for multiple mergers. The proof of Theorem 1.1 is given in the next section. Introduced by Pitman (1999), Λ-coalescents are usually described by Markov processes taking values in partitions of $\mathbb{N}$, which become coarser as time evolves, almost surely, and are exchangeable. More exactly, we define $(\Pi_t)_{t\geq 0} = (\Pi_t^{\Lambda})_{t\geq 0}$, starting in the trivial partition of $\mathbb{N}$. For a finite measure Λ on $[0, 1]$, set

$$\lambda_{b,k} = \int_0^1 x^{k-2}(1-x)^{b-k}\Lambda(dx). \tag{3.1}$$

Among any set of b partition elements in Π_t, each subset of size k merges to one partition element at rate $\lambda_{b,k}$. It is easy to check that such a process is well-defined (i.e. the $\lambda_{b,k}$'s are consistent) and leads to an exchangeable partition of $\mathbb{N}$ for all $t \geq 0$. In our analysis we restrict ourselves to measures Λ which do not have an atom at 1; see Example 20 in Pitman (1999) for a discussion of this case.

One intuitive way to construct a Λ-coalescent (given Λ has no atom at 0) is as follows: consider a Poisson-process with intensity measure $\frac{\Lambda(dx)}{x^2} \cdot dt$ on $[0, 1] \times \mathbb{R}_+$. At any Poisson point (x, t), mark all partition elements, which are available by time t with probability x and merge all marked partition elements.

The set of Λ-coalescents falls into (at least) three classes. The class of Λ-coalescents coming down from infinity (see Property 1 in Theorem 1.1), the larger class of processes having the *dust-free*-property, i.e. $\{f(\Pi_t^1) > 0 \text{ for all } t > 0\}$, almost surely, where $f(\Pi_t^j)$ is the frequency of the partition element containing j at

time t, $j \in \mathbb{N}$). All other Λ-coalescents contain *dust*, which is a positive frequency of natural numbers forming their own partition element.

Starting with Schweinsberg (2000b), sharp conditions for a Λ-coalescents coming down from infinity have been given. Precisely, it was stated that a Λ-coalescent comes down from infinity iff

$$\sum_{b=2}^{\infty} \Big(\sum_{k=2}^{b} k \binom{b}{k} \lambda_{b,k} \Big)^{-1} < \infty. \tag{3.2}$$

It has been shown by Bertoin and Le Gall (2006) that this is equivalent to

$$\int_{t}^{\infty} \psi(q)^{-1} dq < \infty.$$

for some $t > 0$ where

$$\psi(q) = \int_{0}^{1} (e^{-qx} - 1 + qx) x^{-2} \Lambda(dx).$$

The larger class of coalescents having the dust-free property is characterized by the requirement that

$$\int_{0}^{1} x^{-1} \Lambda(dx) = \infty, \tag{3.3}$$

see Theorem 8 in Pitman (1999).

Let $\Pi^{\Lambda} := \Pi = (\Pi_t : t \geq 0)$ be the Λ-coalescent. Then for almost all sample paths of Π^{Λ}, there is a metric r^{Π} on $\mathbb{N}$, associated to Π, defined by

$$r^{\Pi}(i,j) := \inf\{t \geq 0 : i, j \text{ in the same partition element of } \Pi_t\},$$

that is the time needed for i and j to coalesce. We denote by (L^{Π}, r^{Π}) the completion of $(\mathbb{N}, r^{\Pi})$. In order to equip (L^{Π}, r^{Π}) with a probability measure, we use a limit procedure. Set

$$H^n(\Pi) := \overline{\Big(L^{\Pi}, r^{\Pi}, \tfrac{1}{n} \sum_{i=1}^{n} \delta_i \Big)}$$

Then, the family of $\mathbb{M}$-valued random variables $(H^n(\Pi))_{n=1,2,\dots}$ converges in distribution with respect to the Gromov-weak topology iff Π^{Λ} is dust-free, i.e. (3.3) holds (see Theorem 5 in Greven et al. (2009)). Since coalescent processes are associated with tree-like structures, we call the limiting mm-space $\mathcal{L} = \overline{(L^{\Pi}, r^{\Pi}, \mu^{\Pi})}$ the Λ-*coalescent measure tree*.

4. Proof of Theorem 1.1

Let $N(t) := \#\Pi_t$ denote the number of blocks in the partition Π_t and note that $\xi_\varepsilon(\mathcal{L}) \leq N(\varepsilon)$ where $\xi_\varepsilon(\mathcal{L}) < N(\varepsilon)$ is only possible if there are partition elements in Π_ε which carry no mass in $\mathcal{L}$.

(1) $\Rightarrow$ (2): Using Corollary 2.11, we must show that for all $\varepsilon > 0$, we have $\xi_\varepsilon(\mathcal{L}) < \infty$ almost surely. This follows directly from the fact that $\xi_\varepsilon(\mathcal{L}) \leq N(\varepsilon)$ and the assumption that Π comes down from infinity.

(2) $\Rightarrow$ (1): The proof is by contradiction. Assume $\mathcal{L}$ is compact and Π stays infinite for some time $\varepsilon > 0$. Since Π_ε contains no dust, we have that $f((\Pi_\varepsilon^j)) > 0$ for all $j = 1, 2, \dots$, almost surely. Since there are infinitely many lines up to time

ε, we find partition elements of arbitrarily small mass. This implies that $\nu^{\mathcal{L}}\{\underline{r} : s_\varepsilon(\underline{r}) \leq \delta\} > 0$ almost surely, for all $\delta > 0$. On the other hand, since $\mathcal{L}$ is compact, there is a random variable $\Delta > 0$ such that $\nu^{\mathcal{L}}\{\underline{r} : s_\varepsilon(\underline{r}) \leq \Delta\} = 0$, almost surely by Corollary 2.11. In particular, there is $\delta > 0$ such that

$$\nu^{\mathcal{L}}\{\underline{r} : s_\varepsilon(\underline{r}) \leq \delta\} = 0$$

with positive probability, which gives a contradiction.

Last, assume that $\mathcal{L}$ does not come down from infinity and recall that Λ cannot have an atom at 0 in this case. It has been shown in Proposition 23 of Pitman (1999) that the total coalescence rate of all lines is infinite for all times, almost surely. This is easy to see from the construction of Λ-coalescence using the Poisson process with intensity $\Lambda(dx)/x^2$, since the total coalescence rate of the partition element containing 1, given that there are infinitely many lines, is

$$\int_0^1 x \frac{\Lambda(dx)}{x^2} = \int_0^1 x^{-1} \Lambda(dx) = \infty,$$

since the dust-free property, (3.3), holds by assumption.

Let $0 < \eta < \delta$ and consider the δ-ball around 1 in L^Π. Since the coalescence rate is infinite and an infinite number of lines coalesce to the line containing 1 between times η and δ, there is an infinite η-separated set in $B_\delta(\{1\})$. Hence,

$$\nu^{\mathcal{L}}\{\underline{r} : \xi_\eta(\tau_\delta(\underline{r})) < \infty\} = 0,$$

almost surely. Hence, for any sequences $0 < \eta_n < \delta_n$ with $\delta_n \xrightarrow{n \to \infty} 0$, we find that

$$\nu^{\mathcal{L}}\Big(\bigcap_{0 < \eta < \delta_n} \{\underline{r} : \xi_\eta(\tau_{\delta_n}(\underline{r})) < \infty\} \Big) \leq \nu^{\mathcal{L}}(\{\underline{r} : \xi_{\eta_n}(\tau_{\delta_n}(\underline{r})) < \infty\}) = 0,$$

almost surely. By Theorem 2.15, $\mathcal{L}$ cannot be locally compact.

References

L. Addario-Berry, N. Broutin and C. Goldschmidt. The continuum limit of critical random graphs. *Probab. Theory Related Fields* **152** (3-4), 367–406 (2012). MR2892951.

D. Aldous. The continuum random tree. III. *Ann. Probab.* **21** (1), 248–289 (1993). MR1207226.

D. Aldous. Deterministic and stochastic models for coalescence (aggregation and coagulation): a review of the mean-field theory for probabilists. *Bernoulli* **5** (1), 3–48 (1999). MR1673235.

N. Berestycki. *Recent progress in coalescent theory*, volume 16 of *Ensaios Matemáticos [Mathematical Surveys]*. Sociedade Brasileira de Matemática, Rio de Janeiro (2009). ISBN 978-85-85818-40-1. MR2574323.

J. Bertoin and J.-F. Le Gall. Stochastic flows associated to coalescent processes. III. Limit theorems. *Illinois J. Math.* **50** (1-4), 147–181 (electronic) (2006). MR2247827.

D. Burago, Y. Burago and S. Ivanov. *A course in metric geometry*, volume 33 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI (2001). ISBN 0-8218-2129-6. MR1835418.

S. Evans. Kingman's coalescent as a random metric space. In *Stochastic models: Proceedings of the International Conference (ICSM) held in honour of Professor Donald A. Dawson at Carleton University, Ottawa, ON, June 10–13, 1998*, volume 26 of *CMS Conference Proceedings*, pages 105–114. Published by the American Mathematical Society, Providence, RI (2000). ISBN 0-8218-1063-4. Edited by Luis G. Gorostiza and B. Gail Ivanoff, MR1775475.

S. Evans and N. O'Connell. Weighted occupation time for branching particle systems and a representation for the supercritical superprocess. *Canad. Math. Bull.* **37** (2), 187–196 (1994). MR1275703.

S. Evans, J. Pitman and A. Winter. Rayleigh processes, real trees, and root growth with re-grafting. *Probab. Theory Related Fields* **134** (1), 81–126 (2006). MR2221786.

A. Greven, P. Pfaffelhuber and A. Winter. Convergence in distribution of random metric measure spaces (Λ-coalescent measure trees). *Probab. Theory Related Fields* **145** (1-2), 285–322 (2009). MR2520129.

A. Greven, P. Pfaffelhuber and A. Winter. Tree-valued resampling dynamics: Martingale problems and applications. *Probab. Theory Related Fields* (2012). In press.

M. Gromov. *Metric structures for Riemannian and non-Riemannian spaces*, volume 152 of *Progress in Mathematics*. Birkhäuser Boston Inc., Boston, MA (1999). ISBN 0-8176-3898-9. Based on the 1981 French original, With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by Sean Michael Bates; MR1699320.

J. F. C. Kingman. The coalescent. *Stochastic Process. Appl.* **13** (3), 235–248 (1982). MR671034.

J.-F. Le Gall. *Spatial branching processes, random snakes and partial differential equations.* Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel (1999). ISBN 3-7643-6126-3. MR1714707.

J.-F. Le Gall. The topological structure of scaling limits of large planar maps. *Invent. Math.* **169** (3), 621–670 (2007). MR2336042.

J. Pitman. Coalescents with multiple collisions. *Ann. Probab.* **27** (4), 1870–1902 (1999). MR1742892.

J. Schweinsberg. Coalescents with simultaneous multiple collisions. *Electron. J. Probab.* **5**, Paper no. 12, 50 pp. (electronic) (2000a). MR1781024.

J. Schweinsberg. A necessary and sufficient condition for the Λ-coalescent to come down from infinity. *Electron. Comm. Probab.* **5**, 1–11 (electronic) (2000b). MR1736720.

A. M. Vershik. The universal Uryson space, Gromov's metric triples, and random metrics on the series of natural numbers. *Uspekhi Mat. Nauk* **53** (5(323)), 57–64 (1998). MR1691182.