DAUPHINE UNIVERSITÉ PARIS | CEREMADE UMR CNRS 7534

ÉCOLE DOCTORALE DE DAUPHINE

THÈSE DE DOCTORAT

pour obtenir le grade de

Docteur en sciences de l'université Paris-Dauphine

présentée par

Robin STEPHENSON

Divers aspects des arbres aléatoires : des arbres de fragmentation aux cartes planaires infinies

---

Soutenue le 27 Juin 2014 devant le jury composé de MM. :

Nathanaël BERESTYCKI	Examinateur
Nicolas CURIEN	Rapporteur
Massimiliano GUBINELLI	Examinateur
Bénédicte HAAS	Directrice de thèse
Jean-François LE GALL	Examinateur
Grégory MIERMONT	Directeur de thèse

Après avis de Svante JANSON (Rapporteur) et Nicolas CURIEN (Rapporteur).

Remerciements

Je souhaiterais en tout premier remercier mes directeurs de thèse Bénédicte Haas et Grégory Miermont pour m'avoir accompagné et guidé avec fidélité et amitié pendant ces quelques années, et ce malgré les nombreux kilomètres nous séparant parfois. Je leur suis en particulier reconnaissant pour avoir lu de manière précise et corrigé de manière implacable la quantité indénombrable de textes que je leur ai donnés.

Merci également à tous les professeurs que j'ai pu avoir pendant ma vie étudiante, qui m'ont fait découvrir le monde des mathématiques, puis les probabilités et la recherche. En ordre chronologique, Anne Lefèvre, Cyril Germain, Richard Antetomaso, Frédéric Massias, Jean Bertoin, Wendelin Werner, Jean-François Le Gall, et Nathanaël Berestycki.

Un grand merci à Svante Janson et Nicolas Curien pour avoir très rapidement rapporté ma thèse, et également à Massimiliano Gubinelli, dernier membre du jury, de s'intéresser à mon travail.

Merci, merci, *merci* à toutes les personnes avec qui j'ai pu jouer de la musique ces quelques dernières années. Sans vous, la vie serait trop triste. Merci à Mathieu (j'espère qu'on aura encore de nombreuses aventures), Ariane et Gabriel (une seule aventure... mais laquelle !), Chen (le meilleur cobaye), Judith (j'aurai toujours un faible pour ta voix), Grégory (l'unique personne que je mentionnerai deux fois), Jean-Marc (sans Musicami cette année aurait été bien plus pauvre), Auriane (comment préparer un concert en temps record). Merci également à Olivier Cazal, Gérard Bekerman et Hugues Leclère pour avoir fait de moi le pianiste que je suis et pour m'avoir guidé lors de mes périodes de doute.

Un grand merci aussi à tous les thésards et post-docs du CEREMADE, qui rendent la vie de tous les jours supportable et même agréable. En particulier le bureau C606 : Rémi, Jérémy, Isabelle, Adélaïde, Sofia, Viviana, Roxana.

Merci beaucoup à Marie, Isabelle et César, les secrétaires du CEREMADE, pour leur disponibilité, leur gentillesse et leur fiabilité.

Je voudrais également remercier mes amis (ceux qui n'ont pas déjà été mentionnés) pour leur soutien moral ou juste du bon temps passé. Mentions spéciales à Samuel pour ses remarques linguistiques et typographiques et pour m'avoir aidé à préparer ma soutenance (et puis son légendaire sens de l'humour), ainsi qu'à Ruocong pour tant de conversations sur les difficultés de la thèse (et de la vie en général).

Et enfin, je voudrais remercier toute ma famille, et surtout mes parents Adam et Géraldine, pour leur éternel soutien. Je ne le dis pas assez, mais savoir que vous êtes là derrière moi est un grand confort, et je ne sais pas ce que je ferais sans vous.

Table des matières

Introduction

Sommaire

1 Résumé

Nous nous intéressons à la notion d'arbre aléatoire, prise de différents points de vue. Ceux-ci se classent en deux grandes familles : les arbres discrets et les arbres continus. Les arbres discrets sont essentiellement les classiques graphes connexes sans cycles, éventuellement munis d'une structure supplémentaire comme un plongement précis dans le plan ou diverses étiquettes sur leurs sommets, et les arbres continus en sont une généralisation dans un cadre d'espaces métriques.

L'étude mathématique d'arbres discrets aléatoires commence implicitement au milieu du XIX^{e} siècle avec des travaux, hélas essentiellement disparus, du probabiliste français Irénée-Jules Bienaymé, qui traitent de la probabilité d'extinction de familles. Bienaymé utilise un modèle de généalogie simple appelé de nos jours modèle de Galton-Watson, d'après les deux mathématiciens britanniques qui l'ont popularisé. La notion d'arbre aléatoire continu, quant à elle, est bien plus récente : elle commence dans les années 1990 avec une série de trois articles fondateurs par David Aldous [3],[4],[5], où est étudié le célèbre arbre continu brownien, qui s'interprète comme la limite d'échelle de grands arbres généalogiques de Galton-Watson. Cet arbre brownien occupe une place centrale dans le monde des arbres continus car on l'obtient comme limite d'échelle de nombreux modèles d'arbres discrets (voir par exemple [42] et [24]), mais il existe une grande variété d'arbres continus aléatoires, notamment les arbres de Lévy (initiés par le travail de Le Gall et Le Jan [57] et proprement définis en tant qu'arbres par Duquesne et Le Gall dans [30]) et les arbres de fragmentation, ces derniers étant un des sujets principaux de notre travail.

Après un chapitre introductif sur la notion d'arbre continu, cette thèse se décompose en trois grands chapitres, chacun avec son propre thème.

- Le chapitre 2 a trait aux arbres de fragmentation généraux. Dans le prolongement de [41], nous nous intéressons à des arbres continus qui ont une propriété d'auto-similarité : tous les sous-arbres au-dessus d'une certaine hauteur ont la même loi que l'arbre original à un facteur de proportionnalité près, et ils sont indépendants conditionnellement à leurs tailles. Ces arbres peuvent s'interpréter comme arbres généalogiques des processus dits de fragmentation auto-similaire, premièrement étudiés par Jean Bertoin dans [9] et [10]. Après leur construction, nous en faisons un étude géométrique, ce qui aboutit au calcul de leur dimension fractale.

- Nous effectuons dans le chapitre 3 l'étude de la limite d'échelle d'une suite d'arbres k-aires construite algorithmiquement. En s'inspirant d'un algorithme introduit par Rémy dans [71] pour obtenir des arbres binaires plans uniformes, nous introduisons une suite d'arbres discrets où chaque nœud a exactement k enfants, k étant un entier fixé. La taille de l'arbre obtenu à l'instant $n \in \mathbb{Z}_+$ est alors de l'ordre de $n^{1/k}$ et nous montrons que, renormalisée, la suite converge vers un arbre de fragmentation. Nous donnons aussi une manière de coupler les arbres avec différentes valeurs de k de telle sorte qu'ils soient plongés les uns dans les autres.

- Nous établissons dans le chapitre 4 la convergence en loi de grands arbres de Galton-Watson critiques à plusieurs types vers un arbre infini. Il est bien connu qu'un arbre de Galton-Watson tel que chaque individu a en moyenne un enfant, conditionné à avoir un grand nombre de sommets, converge localement en distribution vers un arbre infini. Nous généralisons ceci aux cas de populations à plusieurs types, où la loi de l'ensemble des enfants de chaque individu dépend de son type. Nous appliquons ensuite ce résultat au domaine des cartes planaires aléatoires. Une carte planaire est la donnée d'un graphe connexe et d'un plongement de ce graphe dans la sphère, pris aux homéomorphismes directs de la

sphère près. Nous montrons que, pour de nombreuses distributions de cartes aléatoires, dont notamment les p-angulations uniformes pour p entier, la carte, conditionnée à avoir un nombre de sommets qui tend vers l'infini, converge en loi vers une carte infinie.

Plusieurs de nos résultats traitent de la convergence d'une certaine suite d'arbres. Avant de détailler plus, il convient d'expliciter ces notions, ainsi que d'autres préliminaires.

2 Notions de convergence d'arbres

Convergence locale d'arbres discrets. Tous les arbres (discrets ou continus) que nous considérerons seront enracinés, ce qui signifie qu'un point particulier sera sélectionné et sera appelé racine, typiquement noté ρ ou $\emptyset$. Ceci permet de donner une interprétation généalogique à tout arbre discret : si deux sommets x et y sont voisins, alors celui des deux qui est le plus proche de la racine est le parent de l'autre. La hauteur d'un sommet est alors sa distance à la racine au sens de la distance de graphe, et on peut aussi définir un arbre coupé à une certaine hauteur k : si T est un arbre, ρ sa racine et $k \in \mathbb{Z}_+$ un entier, on note $T_{\leqslant k}$ l'ensemble des sommets de T dont la hauteur est au plus k. Par exemple, $T_{\leqslant 0}$ est le singleton $\{\rho\}$ et $T_{\leqslant 1}$ est composé de ρ et de ses voisins.

Soit $(T_n)_{n\in\mathbb{N}}$ une suite d'arbres discrets. On dit alors qu'elle converge localement vers un arbre T si, pour tout entier k, la suite $\big((T_n)_{\leqslant k}\big)_{n\in\mathbb{N}}$ stationne à $T_{\leqslant k}$. Cette notion de convergence est équivalente à la stationarité si l'arbre limite T est fini, mais est particulièrement intéressante quand T est infini - on voit un arbre infini comme la limite projective de ses sous-arbres coupés. Notamment, la convergence en loi d'une suite d'arbres aléatoires $(T_n)_{n\in\mathbb{N}}$ vers un arbre aléatoire T est alors équivalente à la convergence en loi de $\big((T_n)_{\leqslant k}\big)_{n\in\mathbb{N}}$ vers $T_{\leqslant k}$ pour tout k.

Arbres réels, convergence au sens de Gromov-Hausdorff-Prokhorov. On appelle arbre réel tout espace métrique connexe (T, d) tel que, pour tout choix de deux points x et y de $\mathcal{T}$, il existe un unique chemin continu injectif reliant x à y, chemin qui s'avère de plus être isométrique à un segment. Dans la pratique, tous les arbres que nous considérerons seront compacts, enracinés, et équipés d'une mesure de probabilité sur leur tribu borélienne. De la présence d'une racine, on obtient de nouveau un point de vue généalogique sur l'arbre, en disant qu'un sommet x est un ancêtre d'un autre sommet y si x se trouve sur le chemin entre la racine et y. On obtient également une notion de hauteur, qui n'est rien d'autre que la distance à la racine.

La notion de convergence pour les arbres réels compacts enracinés et mesurés est celle de la métrique de Gromov-Hausdorff-Prokhorov. Donnons nous deux arbres $(\mathcal{T}, d, \rho, \mu)$ et $(\mathcal{T}', d', \rho', \mu')$ et posons

$$d_{\mathrm{GHP}}(\mathcal{T},\mathcal{T}') = \inf \Big[\max\big(d_{\mathcal{Z},\mathrm{H}}(\phi(\mathcal{T}),\phi'(\mathcal{T}')), d_{\mathcal{Z}}(\phi(\rho),\phi'(\rho')), d_{\mathcal{Z},\mathrm{P}}(\phi_*\mu,\phi'_*\mu')\big)\Big],$$

où :

- la borne inférieure est prise sur tous les plongements isométriques ϕ et ϕ' de $(\mathcal{T}, d)$ et $(\mathcal{T}', d')$ dans un même espace métrique $(\mathcal{Z}, d_{\mathcal{Z}})$.
- $d_{\mathcal{Z},\mathrm{H}}$ est la distance de Hausdorff entre deux parties compactes de $\mathcal{Z}$.
- $d_{\mathcal{Z},\mathrm{P}}$ est la distance de Prokhorov entre deux mesures de probabilité sur la tribu borélienne de $\mathcal{Z}$.

Cette notion permet de comparer deux arbres réels, ou même d'autres espaces métriques. On montre notamment que, si $d_{\mathrm{GHP}}(\mathcal{T},\mathcal{T}') = 0$ alors les deux arbres sont équivalents au sens où il existe un isométrie bijective entre les deux qui préserve égalements les mesures et les racines. La fonction d_{GHP} définit alors une métrique séparable et complète sur l'ensemble $\mathbb{T}_W$ des classes d'équivalence d'arbres réels compacts enracinés et mesurés (au sens ci-dessus), nous donnant ainsi une notion de convergence et justifiant aussi la notion d'arbre aléatoire au sens de variable aléatoire dans l'espace $\mathbb{T}_W$ muni de sa tribu borélienne. Les propriétés essentielles de la topologie sur $\mathbb{T}_W$ sont données notamment dans [32] (mais avec une métrique différente) ainsi que [2] (sans se restreindre à des arbres).

Deux exemples phares autour des arbres de Galton-Watson. Soit μ une mesure de probabilité sur l'ensemble des entiers positifs ou nuls $\mathbb{Z}_+$. On appelle arbre de Galton-Watson de loi de reproduction μ tout arbre aléatoire discret T tel que :

- le nombre d'enfants de la racine a pour loi μ.
- les sous-arbres issus des enfants de la racine sont indépendants entre eux et ont tous la même loi que T.

Cet arbre peut être interprété comme l'arbre généalogique d'une population simple où les individus n'interagissent pas entre eux et où le nombre d'enfants de chacun est aléatoire et a pour loi μ. On ajoute l'hypothèse $\mu(1) < 1$ pour retirer le cas dégénéré de l'arbre linéaire infini.

Il est bien connu que les arbres de Galton-Watson sont sujets à un phénomène de transition de phase : appelant m la valeur moyenne de μ, si $m \leqslant 1$ alors l'arbre T est fini presque sûrement, alors que si $m > 1$ il est infini avec probabilité non nulle. Quand $m = 1$ on dit que la loi μ (ou l'arbre T) est critique, et l'arbre est presque-sûrement fini bien que l'espérance de son nombre total de sommets soit infinie. Nous nous restreignons ici au cas critique.

Soit d le PGCD du support de μ. On remarque alors que le nombre total de sommets de T doit être de la forme $1 + dn$ avec $n \in \mathbb{Z}_+$. Par exemple, si chaque individu a nécessairement un nombre pair d'enfants, alors le nombre de sommets de T sera toujours impair. On note alors T_n une version de T conditionnée à avoir $1 + dn$ enfants, évènement qui a une probabilité non-nulle pour n suffisamment grand.

Depuis l'article de Kesten [51], on a cependant une notion de l'arbre T "conditionné à être infini". Cet arbre, noté $\widehat{T}$, est remarquable car sa loi peut être obtenue à partir de celle de T en biaisant par rapport à la taille : pour tout entier k, $\hat{T}_{\leqslant k}$ a une densité par rapport à T, cette densité étant le nombre Z_k de sommets de hauteur k de T. Précisément, on a

$$\mathbb{E}[f(\widehat{T}_{\leqslant k})] = \mathbb{E}[Z_k f(T_{\leqslant k})].$$

La description généalogique de $\widehat{T}$ est aussi intéressante : il est constitué d'une *épine dorsale* infinie, c'est-à-dire un chemin infini partant de la racine, épine dorsale sur laquelle les sommets ont une loi de reproduction particulière appelée $\widehat{\mu}$, qui est une version biaisée par la taille de μ, et qui vérifie, pour $i \in \mathbb{Z}_+$,

$$\widehat{\mu}(i) = i\mu(i).$$

En dehors de cette épine dorsale, les individus ont pour loi de reproduction μ.

Dans [49] ainsi que [1] sont données plusieurs manières de conditionner l'arbre T à être "grand" de sorte qu'il converge en distribution vers $\widehat{T}$, et on peut notamment conditionner par le nombre de sommets : T_n converge en loi vers $\widehat{T}$ quand n tend vers l'infini.

La suite T_n peut aussi servir à illustrer un exemple d'arbres continus. En plus de la criticalité, supposons que μ ait une variance finie, notée σ^2. On munit également T_n de la mesure uniforme sur ses sommets μ_n. On sait alors depuis les travaux d'Aldous (notamment [5]) que

$$\Big(\frac{\sigma}{2\sqrt{dn}} T_n, \mu_n\Big) \underset{n\to\infty}{\Longrightarrow} (\mathcal{T}_{\mathrm{Br}}, \mu_{\mathrm{Br}}),$$

où $(\mathcal{T}_{\mathrm{Br}}, \mu_{\mathrm{Br}})$ est un arbre continu appelé l'arbre brownien. Cette convergence a lieu en loi pour la topologie de Gromov-Hausdorff-Prokhorov (bien qu'Aldous n'utilisait pas encore ce formalisme). Pour être précis, la notation $\frac{\sigma}{2\sqrt{dn}}T_n$ signifie que l'on regarde T_n comme un espace métrique où toutes les arêtes sont en fait des segments de longueur $\frac{\sigma}{2\sqrt{dn}}$. Beaucoup de modèles naturels d'arbres aléatoires rentrent dans ce cadre d'arbres de Galton-Watson conditionnés, et c'est pourquoi l'arbre brownien a une place centrale dans ce monde. On trouve cependant d'autres limites d'échelle dans les cas où la variance est infinie, ce sont précisément les arbres de Lévy, voir notamment [30].

3 Arbres de fragmentation généraux

3.1 Définitions

Dans ce premier travail, on s'intéresse à une classe d'arbres aléatoires particuliers appelés *arbres de fragmentation auto-similaires*. Soit $(\mathcal{T}, d, \rho, \mu)$ un arbre aléatoire compact mesuré et α un réel strictement négatif. Pour $t \geqslant 0$, on note $\big(\mathcal{T}_i(t)\big)_{i\in\mathbb{N}}$ les composantes connexes de $\mathcal{T}_{>t} = \{x \in \mathcal{T} : \ d(x,\rho) > t\}$, ordonnées par ordre décroissant de leur μ-masses. Pour lever les ambiguïtés, si plusieurs composantes ont la même masse, on les ordonne uniformément, et s'il n'y a qu'un nombre fini de composantes alors on complète la suite par une répétition de l'ensemble vide. Pour chaque i, si $\mathcal{T}_i(t)$ est non vide, on note $x_i(t)$ l'unique point de hauteur t tel que $\mathcal{T}_i(t) \cup \{x_i(t)\}$ soit connexe, point qui sera la racine de cet arbre. De plus, si $\mu\big(\mathcal{T}_i(t)\big)$ est non-nul, on munit $\mathcal{T}_i(t)$ de la mesure $\mu_{\mathcal{T}_i(t)}$, qui est la mesure μ conditionnée au sous-ensemble $\mathcal{T}_i(t)$.

Définition 1. *On dit que $(\mathcal{T}, \mu)$ est un arbre de fragmentation auto-similaire d'indice α si, pour tout $t \geqslant 0$, on a p.s. $\mu(\mathcal{T}_i(t)) > 0$ pour chaque $i \in \mathbb{N}$ tel que $\mathcal{T}_i(t)$ soit non vide et, conditionnellement à $\big(\mu(\mathcal{T}_i(s)), i \in \mathbb{N}, s \leqslant t\big)$, les arbres $\big(\mathcal{T}_i(t) \cup \{x_i(t)\}, \mu_{\mathcal{T}_i(t)}\big)_{i\in\mathbb{N}}$ sont tous indépendants et chacun a la même loi que $\big(\mu(\mathcal{T}_i(t))^{-\alpha}\mathcal{T}', \mu'\big)$ où $(\mathcal{T}', \mu')$ est une copie indépendante de $(\mathcal{T}, \mu)$.*

La condition $\alpha < 0$ est nécessaire pour l'existence d'un tel arbre - avec $\alpha \geqslant 0$, on obtiendrait qu'il n'est pas borné, une contradiction.

Ces arbres ont été introduits dans [41], dans un contexte restreint : on y suppose que la mesure μ est non-atomique et donne masse pleine à l'ensemble des feuilles de $\mathcal{T}$. Notre travail est donc en partie une généralisation des résultats de [41]. La première étape consiste à construire les arbres de fragmentation et de caractériser leurs lois. Ceci passe par l'interprétation des arbres de fragmentation en tant qu'arbres généalogiques d'une famille de processus appelés processus de fragmentation, à valeurs dans l'ensemble des partitions de $\mathbb{N}$. Commençons par quelques rappels sur ces processus ainsi que les partitions aléatoires.

3.2 Processus de fragmentation auto-similaires

Soit $\mathbb{N}$ l'ensemble des entiers strictement positifs :

$$\mathbb{N} = \{1, 2, 3, \ldots\}.$$

Une partition π de $\mathbb{N}$ est représentée comme une suite infinie de blocs $(\pi_i)_{i\in\mathbb{N}}$ disjoints deux-à-deux dont l'union recouvre tout $\mathbb{N}$. Les blocs sont ordonnés de sorte que, pour $i < j$, le plus petit élément de π_i est inférieur à celui de π_j, ce qui garantit un unique ordre possible. On appelle $\mathcal{P}_{\mathbb{N}}$ l'ensemble de ces partitions, dont on fait un espace métrique compact en posant, pour deux partitions π et π', $d(\pi, \pi') = 2^{-n(\pi,\pi')}$ où $n(\pi, \pi')$ est l'entier maximal tel que π et π' sont égales sur les n premiers entiers, et l'infini si $\pi = \pi'$.

Une partition aléatoire π est dite échangeable si, pour toute permutation σ de $\mathbb{N}$, la partition $\sigma\pi$, dont les blocs sont les images réciproques par σ des blocs de π, a même loi que π. Une propriété des partitions aléatoires échangeables, démontrée initialement par Kingman dans [52], est que tous les blocs d'une telle partition π admettent des fréquences asymptotiques presque sûrement, la fréquence asymptotique $|A|$ de toute partie A de $\mathbb{N}$ étant définie par la limite

$$|A| = \lim_{n\to\infty} \frac{\#(A\cap[n])}{n}$$

quand cette limite existe, et où $[n] = \{1,\ldots,n\}$. La liste $|\pi|^{\downarrow}$ des fréquences asymptotiques des blocs, ordonnée par ordre décroissant, est alors un élément de l'ensemble des partitions de l'unité

$$\mathcal{S}^{\downarrow} = \Big\{\mathbf{s} = (s_i)_{i\in\mathbb{N}} : s_1 \geqslant s_2 \geqslant \ldots \geqslant 0, \sum s_i \leqslant 1\Big\},$$

que l'on métrise et rend compact en posant, pour deux partitions $\mathbf{s} = (s_i)_{i\in\mathbb{N}}$ et $\mathbf{s}' = (s'_i)_{i\in\mathbb{N}}$,

$$d(\mathbf{s},\mathbf{s}') = \sup_{i\in\mathbb{N}} |s_i - s'_i|.$$

Si π est une partition échangeable, alors, sachant que $|\pi|^{\downarrow} = \mathbf{s}$, on peut reconstruire la loi de π, par le procédé dit de la *boîte de peinture* de Kingman : soit $(X_i)_{i\in\mathbb{N}}$ une suite de variables indépendantes et uniformes sur $[0,1]$. On déclare alors que deux entiers i et j sont dans le même bloc si et seulement si X_i et X_j sont dans le même intervalle de la forme $[s_1+\ldots+s_k, s_1+\ldots+s_{k+1})$ avec $k \geqslant 0$. Si ν est la loi de $\mathbf{s} = (s_1, s_2, \ldots)$, on appelle alors κ_ν la loi de π. Les entiers i tels que $X_i \geqslant \sum_{k=1}^{\infty} s_k$ se retrouvent alors dans un singleton, on appelle ces singletons la *poussière* de π.

Soit α un réel. Un processus de fragmentation auto-similaire d'indice α est un processus $\Pi = (\Pi(t))_{t\geqslant 0}$ satisfaisant les trois propriétés suivantes.

- Pour toute permutation σ de $\mathbb{N}$, le processus $\sigma\Pi = \big(\sigma\Pi(t)\big)_{t\geqslant 0}$ a même loi que Π,
- Presque sûrement, pour tout temps t, $\Pi(t)$ admet des fréquences asymptotiques,
- Les trajectoires de Π sont càdlàg,
- Pour tout $t \geqslant 0$, conditionnellement à $\Pi(t) = \pi$, les processus $(\Pi(t+u)\cap\pi_i)_{u\geqslant 0}$ sont indépendants et chacun a la même loi que $(\Pi(|\pi_i|^{\alpha}u)\cap\pi_i)_{u\geqslant 0}$.

Notons que la notation $\pi\cap A$ où π est une partition de $\mathbb{N}$ et A un bloc représente la partition de A donc les blocs sont obtenus en intersectant les blocs de π avec A. On interprète ceci comme le fait que le bloc A est disloqué par la partition π.

Ces objets représentent le devenir d'une particule qui, au fur et à mesure du temps, se dégrade en des particules plus petites. Ils ont été introduits par Bertoin dans [9] pour $\alpha = 0$, puis pour tout α dans [10]. On sait notamment caractériser leur loi par trois paramètres : l'indice d'auto-similarité α, un *coefficient d'érosion* c qui est un réel positif ou nul, et une mesure σ-finie ν sur $\mathcal{S}^{\downarrow}$ qui intègre $1-s_1$, appelée *mesure de dislocation*. Le coefficient c code le fait que, si un entier i est dans un bloc A de $\Pi(t)$, alors pendant un intervalle de temps $\mathrm{d}t$, il a une probabilité de l'ordre de $c|A|^{\alpha}\mathrm{d}t$ d'être envoyé dans un singleton entre t et $t+\mathrm{d}t$. La mesure ν code les dislocations soudaines de fragments : de manière informelle, pour $\pi \in \mathcal{P}_{\mathbb{N}}$, alors un bloc A de Π va être remplacé par $A\cap\pi$ avec probabilité $|A|^{\alpha}\kappa_\nu(\mathrm{d}\pi)\mathrm{d}t$ entre t et $t+\mathrm{d}t$.

On dit que la fragmentation est *conservatrice* si ν ne charge que des suites de somme totale 1 et si $c = 0$. Intuitivement, cela signifie qu'à faible échelle, la fragmentation ne perd pas de masse et que la somme des fréquences asymptotiques des blocs reste égale à 1. Ceci n'est en fait pas vrai : si $\alpha < 0$, alors le processus accélère d'autant que les fragments sont petits, au point que l'on arrive en temps fini à la partition formée uniquement de singletons. Ce phénomène est précisément étudié dans [40].

Par la suite on supposera toujours que α est strictement négatif et que le couple (c, ν) est différent de $(0, a\delta_1)$ pour tout $a \geqslant 0$ pour éviter le cas dégénéré du processus constamment égal à la partition ayant pour seul bloc $\mathbb{N}$.

3.3 Les arbres de fragmentation en tant qu'arbres généalogiques de processus de fragmentation

Notre résultat, en deux parties, met en évidence une bijection entre les lois des arbres auto-similaires et celles processus de fragmentation auto-similaires.

Proposition 3.1. *Soit $(\mathcal{T}, \mu)$ un arbre de fragmentation auto-similaire d'indice α. Soit $(P_i)_{i\in\mathbb{N}}$ une suite de points de $\mathcal{T}$ telle que, conditionnellement à $(\mathcal{T}, \mu)$, les $(P_i)_{i\in\mathbb{N}}$ soient indépendants et de loi μ. Définissons alors un processus $\Pi_{\mathcal{T}} = \big(\Pi_{\mathcal{T}}(t)\big)_{t\geqslant 0}$ à valeurs dans $\mathcal{P}_{\mathbb{N}}$ de la manière suivante : pour tout temps $t \geqslant 0$, deux entiers i et j sont dans le même bloc de $\Pi_{\mathcal{T}}(t)$ si et seulement si P_i et P_j sont dans la même composante connexe de $\mathcal{T}_{>t}$. Alors $\Pi_{\mathcal{T}}$ est un processus de fragmentation auto-similaire d'indice α.*

La preuve de la Proposition 3.1 se fait de manière élémentaire à partir des définitions. Avec plus de travail, on démontre le théorème suivant, qui implique en particulier que la distribution de l'arbre de fragmentation $(\mathcal{T}, \mu)$ est intégralement caractérisée par les paramètres du processus de fragmentation $\Pi_{\mathcal{T}}$.

Théorème 3.1. *Pour tout triplet de paramètres (α, c, ν), il existe un arbre de fragmentation auto-similaire unique en loi $(\mathcal{T}, \mu)$ tel que le processus $\Pi_{\mathcal{T}}$ ait pour paramètres (α, c, ν).*

Ce théorème justifie alors le fait que l'on puisse parler d'un arbre de fragmentation avec pour paramètres (α, c, ν).

La preuve du Théorème 3.1 se fait en se donnant un processus de fragmentation Π avec paramètres (α, c, ν), et en construisant à la main son arbre généalogique. Si on note, pour $i \in \mathbb{N}$, D_i l'instant où i tombe dans un singleton de Π et, pour deux entiers distincts i et j, $D_{\{i,j\}}$ l'instant à partir duquel i et j sont dans deux blocs différents, alors $\mathcal{T}$ peut être défini comme l'arbre réel "minimal" tel que :

- pour chaque entier i, il y a un segment de longueur D_i partant de la racine.
- pour $i \neq j$, les chemins correspondant à i et j se séparent à la hauteur $D_{\{i,j\}}$.

L'extrémité du segment correspondant à un entier j est alors appelée le "point de mort" de j, et est notée Q_j. On notera également (j, t) le point de hauteur t de ce segment, pour $t \in [0, D_j]$. La Figure 1 ci-dessous donne un exemple dans le cas simplifié où on ne regarde que les 7 premiers entiers.

On montre notamment que $\mathcal{T}$, qui est plongé dans l'espace des suites sommables grâce à une variation de la méthode dite du "stick-breaking" d'Aldous, est une fonction mesurable de Π, ce qui implique que c'est effectivement une variable aléatoire.

Dans [41] était supposé le fait que la mesure sur l'arbre de fragmentation avait pour support l'ensemble des feuilles et n'avait pas d'atomes, et en conséquence on obtenait que le processus de fragmentation devait être conservatif. Ce n'est plus nécessairement le cas ici : on observe un atome de μ à chaque point de branchement correspondant à une dislocation où de la masse est perdue, et avec l'érosion on observe une composante à densité sur les branches de l'arbre.

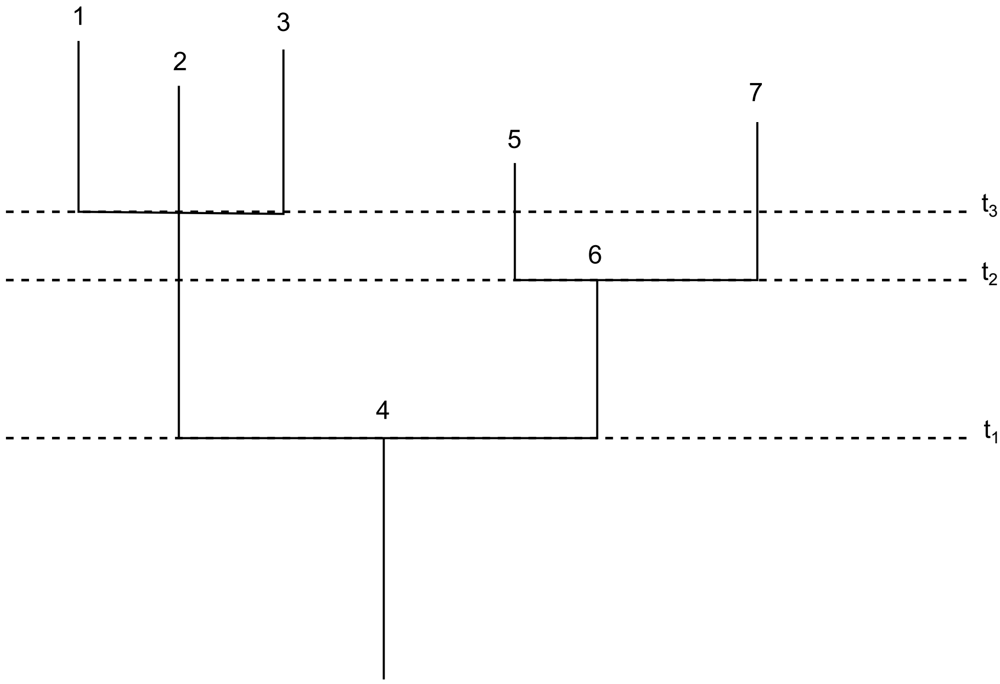

Figure 1 – Une coupe de $\mathcal{T}$ à sept entiers. On a par exemple ici $D_{\{1,4\}} = D_{\{2,4\}} = t_1$, $D_{\{5,6\}} = D_{\{5,7\}} = t_2$, $D_{\{1,2\}} = D_{\{2,3\}} = t_3$.

3.4 Calcul de la dimension fractale de Hausdorff des arbres de fragmentation malthusiens

Hypothèses malthusiennes. Soit $c \geqslant 0$ et ν une mesure de dislocation. On dit que le couple (c, ν) est malthusien avec pour exposant malthusien $p^* > 0$ si

$$cp^* + \int_{\mathcal{S}^\downarrow} \left(1 - \sum_{i=1}^{\infty} s_i^{p^*}\right) \mathrm{d}\nu(\mathbf{s}) = 0.$$

Informellement, cela signifie que, si l'on regarde un processus de fragmentation $(\Pi(t))_{t \geqslant 0}$ homogène (c'est-à-dire dont l'indice d'auto-similarité est 0), alors le processus $\big(M(t)\big)_{t \geqslant 0}$ défini par

$$M(t) = \sum_{i=1}^{\infty} |\Pi_i(t)|^{p^*}$$

est une martingale. Autrement dit, en mettant à la puissance p^* les masses des fragments, une forme de conservation apparait, bien que le processus ne soit pas conservatif. Notons que ceci est perdu quand on passe à un indice d'auto-similarité α strictement négatif, à cause du phénomène de singularité mentionné ci-dessus.

Dans la pratique, on a besoin d'une hypothèse malthusienne légèrement renforcée, ce que nous appelons hypothèse $(\mathbf{H})$:

> La fonction $p \mapsto cp + \int_{\mathcal{S}^\downarrow} (1 - \sum_i s_i^p) \nu(\mathrm{d}\mathbf{s}) \in [-\infty, +\infty)$ prend une valeur finie strictement négative sur l'intervalle $[0, 1]$.

Valeur de la dimension. Sous (**H**), on peut explicitement calculer la dimension de Hausdorff d'un arbre de fragmentation.

Théorème 3.2. *Supposons que (c,ν) vérifie (**H**), et soit $(\mathcal{T},\mu)$ un arbre de fragmentation d'indice α, de paramètre d'érosion c et de mesure de dislocation ν. Alors, presque sûrement, une seule des deux assertions suivantes est vraie :*

- *l'ensemble des feuilles de $\mathcal{T}$ est dénombrable.*
- *la dimension de Hausdorff de l'ensemble des feuilles de $\mathcal{T}$ est égale à $\frac{p^*}{|\alpha|}$.*

Le premier évènement est en quelque sorte un cas dégénéré, et est analogue au fait qu'un arbre de Galton-Watson surcritique a une probabilité non-nulle de s'éteindre. Il ne se produit que si la mesure de dislocation donne de la masse au singleton formé de la suite nulle $(0,0,\ldots)$.

Comme c'est le cas pour de nombreux calculs de dimensions de Hausdorff, la preuve du Théorème 3.2 se fait deux parties : la majoration et la minoration. La majoration de la dimension se fait de manière élémentaire, mais cependant sa minoration utilise des techniques assez élaborées qui méritent d'être expliquées.

Une mesure sur les feuilles de l'arbre. La preuve de la minoration de la dimension de Hausdorff de l'ensemble des feuilles de l'arbre de fragmentation commence avec la création d'une mesure borélienne sur cet ensemble. Le définition de cette mesure nécessite d'abord l'étude de certaines martingales associées au processus de fragmentation Π dont $\mathcal{T}$ est l'arbre généalogique. Pour cela, on doit d'abord revenir à un processus de fragmentation homogène. Posons, pour $i \in \mathbb{N}$ et $t \geqslant 0$,

$$\tau_i^{-1}(t) = \inf\Big\{u : \int_0^u |\Pi_{(i)}(r)|^\alpha \mathrm{d}r > t\Big\}.$$

La notation $\pi_{(i)}$, où π est une partition de $\mathbb{N}$, désigne le bloc de π qui contient i.

Grâce à cette famille de changements de temps, on peut se ramener à un processus homogène. Pour être précis, pour $t > 0$, on appelle $\Pi^0(t)$ la partition telle que deux entiers i et j sont dans le même bloc de $\Pi^0(t)$ si et seulement si $j \in \Pi_{(i)}(\tau_i^{-1}(t))$, et ceci définit un processus de fragmentation avec pour paramètres $(0,c,\nu)$. On pose alors, pour $i \in \mathbb{N}$, $t \geqslant 0$ et $s \geqslant 0$,

$$M_{i,t}(s) = \sum_{j=1}^{\infty} |\Pi_j^0(t+s) \cap \Pi_{(i)}^0(t)|^{p^*}.$$

L'hypothèse malthusienne a pour conséquence que le processus $\big(M_{i,t}(s)\big)_{s\geqslant 0}$ est une martingale. S'avérant de plus être positive et càdlàg, elle admet alors une limite à l'infini $W_{i,t}$ que l'on nomme $W_{i,t}$. Sous l'hypothèse que la martingale converge au sens L^1, on démontre alors que la famille $(W_{i,t}, i \in \mathbb{N}, t \geqslant 0)$ respecte la structure généalogique de Π^0, au sens où

$$W_{i,t} = \sum_{j \in \Pi^0_{(i)}(t) \cap \mathrm{rep}(\Pi^0(s))} W_{j,s},$$

où la notation $\mathrm{rep}(\Pi^0(s))$ signifie qu'on prend un unique entier par bloc de $\Pi^0(s)$. En conséquence de ceci, on peut alors montrer qu'il existe une unique mesure μ^* sur $\mathcal{T}$, supportée par l'ensemble de ses feuilles, telle que

$$\forall i \in \mathbb{N}, t \geqslant 0, \mu^*(\mathcal{T}_{(i,t+)}) = W_{i,\tau_i(t)},$$

où $\mathcal{T}_{(i,t+)}$ désigne, parmi les différentes composantes connexes de $\mathcal{T}_{>t}$, celle qui contient le point de mort Q_i.

L'étude de la mesure μ^* passe par une version de Π^0 biaisée par la taille du bloc contenant 1, qui fut introduite dans [13] : on peut définir un processus $\Pi^* = (\Pi^*(t))_{t\geqslant 0}$ à valeurs dans $\mathcal{P}_{\mathbb{N}}$ tel que

$$\mathbb{E}\big[F((\Pi^*(s))_{s\leqslant t})\big] = \mathbb{E}\Big[|\Pi_1^0(t)|^{p^*-1}F\big((\Pi^0(s))_{s\leqslant t}\big)\Big].$$

Le processus Π^* se comporte tout à fait comme le processus de fragmentation Π^0 mis à part le fait que le fragment contenant 1 se disloque de manière différente des autres : sur un intervalle $[t, t+dt]$, $\Pi_1^0(t)$ est remplacé par $\Pi_1^0(t)\cap\pi$, non pas avec probabilité $\kappa_\nu(\mathrm{d}\pi)\mathrm{d}t$ mais avec probabilité $\kappa_\nu^*(\mathrm{d}\pi)\mathrm{d}t$, quantité que l'on définit par

$$\kappa_\nu^*(\mathrm{d}\pi) = |\pi_1|^{p^*-1}\kappa_\nu(\mathrm{d}\pi).$$

Une conséquence de ceci est que, dans Π^*, l'entier 1 ne tombe pas dans un singleton en temps fini, ce qui est connecté au fait que la mesure μ^* charge uniquement les feuilles de $\mathcal{T}$.

Une fois l'étude de la mesure μ^* faite, la minoration de la dimension s'effectue en utilisant la méthode de l'énergie de Frostman. Ceci consiste à trouver $\gamma > 0$ tel que

$$\mathbb{E}\left[\int_{\mathcal{T}}\int_{\mathcal{T}}(d(L,L'))^{-\gamma}\mathrm{d}\mu^*(L)\mathrm{d}\mu^*(L')\right] < \infty,$$

impliquant que γ minore la dimension cherchée. De manière similaire à ce qui est obtenu dans [41], la borne inférieure obtenue par cette méthode s'avère souvent trop petite. Pour obtenir la bonne borne, on effectue un élagage de l'arbre de la manière suivante :

- on choisit $N \in \mathbb{N}$ et, à chaque dislocation, on ne regarde que les N fragments les plus grands.
- on choisit $\varepsilon > 0$ et, pour chaque dislocation dont le fragment le plus grand a une taille relative supérieure à $1-\varepsilon$, on ignore tous les autres fragments.

La méthode de Frostman s'applique bien à l'arbre tronqué et, en faisant tendre N vers l'infini ainsi que ε vers 0, on arrive alors à obtenir une minoration par $\frac{p^*}{|\alpha|}$ de la dimension de Hausdorff.

4 Limite d'échelle d'une suite d'arbres k-aires

Le modèle. On se donne un entier $k \geqslant 2$ et on s'intéresse à une suite d'arbres $(T_n(k))_{n\in\mathbb{Z}_+}$ définie par la récurrence aléatoire suivante :

- $T_0(k)$ est l'arbre formé d'une seule arête et deux sommets, l'un étant désigné comme racine.
- Pour $n \geqslant 1$, sachant $T_{n-1}(k)$, on choisit uniformément l'une de ses arêtes et on ajoute au milieu de cette arête un nouveau sommet, la séparant en deux. De ce nouveau sommet, on fait partir $k-1$ nouvelles arêtes, le reliant à $k-1$ nouvelles feuilles.

L'arbre $T_n(k)$ est également muni de la mesure uniforme sur ses feuilles, notée $\mu_n(k)$.

Cet algorithme est une généralisation de l'algorithme de Rémy utilisé dans [71] pour obtenir des arbres binaires uniformes, qui correspond au cas où $k=2$. Ces arbres binaires se trouvent en réalité être des arbres de Galton-Watson conditionnés, et les travaux d'Aldous impliquent alors que, renormalisé par $\sqrt{n}$, l'arbre $T_n(2)$, muni de la mesure uniforme sur ses feuilles, converge en loi vers un multiple scalaire de l'arbre brownien.

$$\left(\frac{T_n(2)}{n^{1/2}}, \mu_n(2)\right) \underset{n\to\infty}{\Longrightarrow} \left(2\sqrt{2}\mathcal{T}_{\mathrm{Br}}, \mu_{\mathrm{Br}}\right) \tag{1}$$

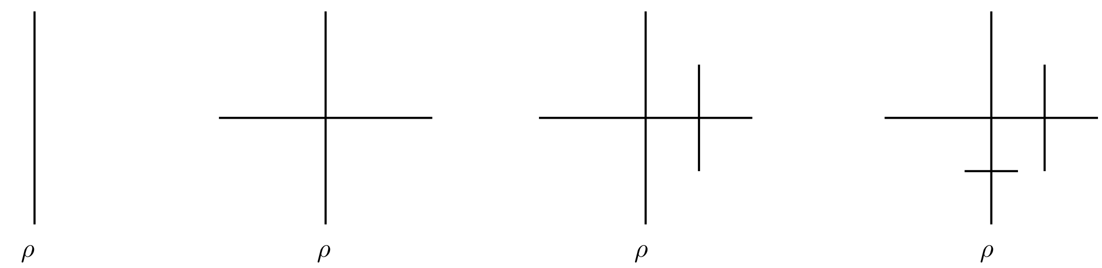

Figure 2 – Une représentation of $T_n(3)$ pour $n = 0, 1, 2, 3$

Il y a de nombreuses façons de définir l'arbre brownien, une manière adaptée ici serait de dire que c'est un arbre de fragmentation d'indice $-1/2$, sans érosion et dont les branchements sont tous binaires, si bien que sa mesure de dislocation ν_{Br} est caractérisée par

$$\nu_{\mathrm{Br}}(\mathrm{d}s_1) = \sqrt{\frac{2}{\pi}} s_1^{-3/2} s_2^{-3/2} \mathbf{1}_{\{s_1 \geqslant s_2\}} \mathrm{d}s_1 = \sqrt{\frac{2}{\pi}} s_1^{-1/2} s_2^{-1/2} \left(\frac{1}{1-s_1} + \frac{1}{1-s_2} \right) \mathbf{1}_{\{s_1 \geqslant s_2\}} \mathrm{d}s_1,$$

où s_2 est implicitement égal à $1 - s_1$. La convergence 1 se trouve en fait être une convergence presque sûre, ce qui a notamment été montré dans [24], où est traitée la convergence d'une autre suite d'arbres avec un autre algorithme généralisant également le cas binaire.

Convergence. Notre résultat principal est le suivant :

Théorème 4.1. *On a, lorsque n tend vers l'infini,*

$$\left(\frac{T_n(k)}{n^{1/k}}, \mu_n(k) \right) \underset{n \to \infty}{\overset{\mathbb{P}}{\Longrightarrow}} (\mathcal{T}_k, \mu_k)$$

où $(\mathcal{T}_k, \mu_k)$ est un arbre de fragmentation d'indice $-1/k$, sans érosion et dont la mesure de dislocation, appelée $\nu_k^{\downarrow}$ est donnée par

$$\nu_k^{\downarrow}(\mathrm{d}s_1 \mathrm{d}s_2 \ldots \mathrm{d}s_{k-1}) = \frac{(k-1)!}{k(\Gamma(\frac{1}{k}))^{k-1}} \prod_{i=1}^{k} s_i^{-(1-1/k)} \left(\sum_{i=1}^{k} \frac{1}{1-s_i} \right) \mathbf{1}_{\{s_1 \geqslant s_2 \geqslant \ldots \geqslant s_k\}} \mathrm{d}s_1 \mathrm{d}s_2 \ldots \mathrm{d}s_{k-1}.$$

La quantité s_k est ici implicitement définie par $s_k = 1 - \sum_{i=1}^{k-1} s_i$, se qui rend la mesure $\nu_k^{\downarrow}$ conservative et "k-aire" aux sens où seuls les k premiers termes des suites considérées sont non-nuls.

On perd en particulier la convergence presque sûre du cas binaire, pour obtenir à la place une convergence en probabilité légèrement plus faible.

La preuve du Théorème 4.1 se fait en deux parties. D'un côté on montre que la suite $(n^{-1/k} T_n(k), \mu_n(k))_{n \in \mathbb{N}}$ converge en probabilité, sans s'intéresser à la loi de l'arbre limite. Ceci se fait en numérotant les feuilles de $T_n(k)$ par ordre d'apparition dans l'algorithme. En appelant ainsi $(L_n^i)_{1 \leqslant i \leqslant (k-1)n+1}$ les feuilles de $T_n(k)$, on montre que la distance entre L_n^i et L_n^j, pour deux entiers distincs i et j, se comporte comme $n^{1/k}$ quand n tend vers l'infini. Ceci permet de définir l'arbre limite $\mathcal{T}_k$, et des propriétés de tension permettent de montrer la convergence. On utilise pour cela des résultats sur les processus dit de "restaurants chinois" étudiés notamment par Pitman dans [69].

L'autre moitié de la preuve du Théorème 4.1 consiste à montrer la convergence en loi de $(n^{-1/k} T_n(k), \mu_n(k))_{n \in \mathbb{N}}$ vers l'arbre de fragmentation considéré. On utilise alors la théorie des *arbres Markov branchants.* Informellement, on remarque que $T_n(k)$ est auto-similaire à l'échelle discrète. Plus précisément, soient $X_n^1, \dots, X_n^k$ des entiers aléatoires dont la somme vaut n et tels que

$$\mathbb{P}\big[(X_n^1, \dots, X_n^k) = (\lambda_1, \dots, \lambda_k)\big] = \\ \frac{1}{k(\Gamma(\frac{1}{k}))^{k-1}} \left(\prod_{i=1}^{k} \frac{\Gamma(\frac{1}{k} + \lambda_i)}{\lambda_i!} \right) \frac{n!}{\Gamma(\frac{1}{k} + n + 1)} \left(\sum_{j=1}^{\lambda_1+1} \frac{\lambda_1!}{(\lambda_1 - j + 1)!} \frac{(n-j+1)!}{n!} \right).$$

Conditionnellement à $(X_n^1, \dots, X_n^k)$, soient $T^1, \dots, T^k$ des arbres indépendants tels que, pour tout $i \in [k]$, T^i ait la même loi que $T_{X_n^i}(k)$. Soit T l'arbre obtenu de la manière suivante : sa racine est adjacente à une unique arête, à l'autre bout de laquelle on greffe $T^1, \dots, T^k$. L'arbre T a alors la même loi que $T_{n+1}(k)$. Les limites d'échelle d'arbres Markov branchants ont été étudiées en détail dans [42], et la convergence en loi s'obtient en étudiant les propriétés asymptotiques de la loi de $(X_n^1, \dots, X_n^k)$ quand n tend vers l'infini. Précisément, notons $(Y_n^1, \dots, Y_n^k)$ le réordonnement décroissant de $(X_n^1, \dots, X_n^k)$, et notons $\bar{q}_n^{\downarrow}$ la loi de $(\frac{Y_n^i}{n})_{1 \leqslant i \leqslant k}$. On démontre que

$$(1 - s_1)\bar{q}_n^{\downarrow}(\mathrm{d}\mathbf{s}) \underset{n\to\infty}{\Longrightarrow} (1 - s_1)\nu_k^{\downarrow}(\mathrm{d}\mathbf{s}),$$

et la convergence en loi de $\big(n^{-1/k} T_n(k), \mu_n(k)\big)$ est alors une conséquence du Théorème 5 de [42]

Couplage des arbres. Soit $k' < k$, il est possible d' "inclure" la suite $(T_n(k'))_{n\geqslant 0}$ à l'intérieur de $(T_n(k))_{n\geqslant 0}$ grâce à un couplage bien choisi. Précisément, on définit par récurrence une suite $(T_n(k, k'))_{n\geqslant 0}$ simultanément avec $(T_n(k))_{n\geqslant 0}$, de la manière suivante :

- $T_0(k, k')$ est égal à $T_0(k)$.
- Pour $n \geqslant 1$, sachant $T_{n-1}(k)$, ainsi que $T_{n-1}(k, k')$, deux cas se présentent. Si l'arête sélectionnée pour construire $T_n(k)$ n'est pas dans $T_{n-1}(k, k')$ alors on pose $T_n(k, k') = T_{n-1}(k, k')$. Si l'arête sélectionnée est dans $T_{n-1}(k, k')$, alors $T_n(k, k')$ est l'arbre obtenu en ajoutant $k' - 1$ des $k - 1$ nouvelles arêtes enracinées sur le nouveau nœud.

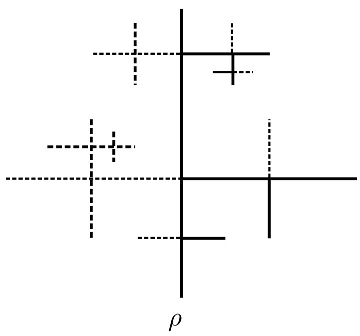

Figure 3 – Une représentation de $T_{10}(3)$ et $T_{10}(3, 2)$. Tous les "virages à gauche" de $T_{10}(3)$ sont retirés pour obtenir $T_{10}(3, 2)$.

De l'inclusion de $T_n(k, k')$ dans $T_n(k)$, on peut démontrer que

$$\frac{T_n(k, k')}{n^{1/k}} \underset{n\to\infty}{\overset{\mathbb{P}}{\Longrightarrow}} \mathcal{T}_{k,k'},$$

où $\mathcal{T}_{k,k'}$ est un sous-arbre de $\mathcal{T}_k$.

L'arbre $T_n(k,k')$ étant k'-aire, on s'attend à ce qu'il s'apparente à $T_m(k')$ pour un certain m. Pour être précis, on démontre que la suite $(T_n(k,k'))_{n\in\mathbb{Z}_+}$ a même loi que la suite $(T_{I_n}(k'))_{n\geqslant 0}$ où $(I_n)_{n\in\mathbb{Z}_+}$ est une chaîne de Markov simple indépendante de $(T_n(k'))_{n\in\mathbb{Z}_+}$. On peut démontrer que, à une constante aléatoire près, I_n est équivalent à $n^{k'/k}$ quand n tend vers l'infini, ce qui implique

$$\frac{T_n(k)}{n^{1/k}} \underset{n\to\infty}{\overset{\mathbb{P}}{\Longrightarrow}} M\tilde{\mathcal{T}}_{k'},$$

où $\tilde{\mathcal{T}}_{k'}$ est une copie de $\mathcal{T}_{k'}$ et M une variable aléatoire strictement positive indépendante de $\tilde{\mathcal{T}}_{k'}$. Ceci démontre que, à constante aléatoire près, l'arbre $\mathcal{T}_{k'}$ est inclus dans $\mathcal{T}_k$. Par ailleurs, sachant $(\mathcal{T}_k,\mu_k)$, il est même possible de reconstruire une version de $M\mathcal{T}_{k'}$ incluse dans $\mathcal{T}_k$, en choisissant à chaque branchement k' sous-arbres aléatoirement parmi les k présents.

5 Arbres de Galton-Watson multi-types critiques infinis

Nous généralisons la convergence locale en loi de grands arbres de Galton-Watson classiques critiques vers l'arbre infini biaisé par la taille au cas d'arbres avec plusieurs types. Commençons par définir ces arbres avec suffisamment de précision.

5.1 Arbres de Galton-Watson multi-types

Arbres plans et lois de Galton-Watson. Nous nous intéressons ici à des arbres *plans*, c'est-à-dire des arbres munis d'un plongement dans le plan orienté. Nous nous contenterons de les voir commes des arbres enracinés tels que l'ensemble des enfants de chaque sommet soit totalement ordonné, ce qui permet effectivement de dessiner l'arbre dans le plan, en ordonnant ces sommets de gauche à droite. Ces arbres seront de plus munis de types sur chaque sommet : on se donne un entier $K\in\mathbb{N}$ représentant le nombre de types, et à chaque sommet u d'un arbre T, on associe un type $\mathbf{e}(u)\in[K]$. Pour tout sommet u de T, les types ordonnés de ses enfants forment alors un élément de l'ensemble

$$\mathcal{W}_K=\bigcup_{n=0}^{\infty}[K]^n.$$

Les arbres de Galton-Watson multi-types se définissent à partir d'une famille de lois de reproduction : pour tout type $i\in[K]$, on se donne une mesure de probabilité $\zeta^{(i)}$ sur $\mathcal{W}_K$. De manière informelle, un arbre de Galton-Watson avec loi de reproduction $\zeta=(\zeta^{(i)})_{i\in\mathbb{N}}$ est un arbre où la descendance de chaque individu est indépendante des autres et, pour un sommet de type $i\in[K]$, la liste ordonnée des types de ses enfants a pour loi $\zeta^{(i)}$. Ceci modélise la généalogie d'une population d'individus à plusieurs types qui n'interagissent pas entre eux. Dans la pratique, on retire les cas dégénérés où chaque $\zeta^{(i)}$ a support dans $[K]$, pour éviter les généalogies linéaires infinies.

Tout comme pour les arbres de Galton-Watson monotypes, on a une notion de criticalité pour arbres de Galton-Watson multi-types, qui fait aussi intervenir la moyenne. Pour deux types i et j, on note $m_{i,j}$ la moyenne du nombre d'enfants de type j d'un individu de type i - il s'agit de l'espérance du nombre de fois qu'apparaît j dans une variable ayant pour loi $\zeta^{(i)}$. On définit alors la matrice des moyennes $M=(m_{i,j})_{i,j\in[K]}$, et on dit que ζ est critique si le rayon spectral de M est égal à 1. Si cette matrice est irréductible (au sens où, pour tout couple de types (i,j), il existe $n\in\mathbb{N}$ tel que le coefficient (i,j) de M^n est non nul) alors les espaces propres à gauche et à droite pour 1 sont tous deux de dimension 1. On y choisit respectivement deux éléments

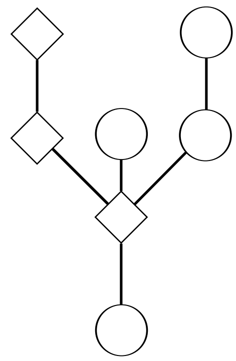

Figure 4 – Un exemple d'arbre à deux types. En admettant que les nœuds de type 1 soient représentés par des cercles ◯ et les nœuds de type 2 par des carrés ⋄, la racine a un enfant de type 2, qui a lui même trois enfants, dont la liste des types est $(1, 2, 2)$.

$\mathbf{a} = (a_1, \ldots, a_K)$ et $\mathbf{b} = (b_1, \ldots, b_K)$, qui vérifient $\sum_i a_i = \sum_i a_i b_i = 1$. Notons aussi que $a_i > 0$ et $b_i > 0$ pour tout i.

5.2 L'arbre infini.

On considère une famille de lois de reproduction critique ζ telle que la matrice M soit irréductible et on considère un arbre de Galton-Watson à K types T dont la racine a pour type $i \in [K]$. On définit pour $n \in \mathbb{N}$ la variable aléatoire

$$X_n = \sum_{u \in T_n} b_{\mathbf{e}(u)}$$

où T_n désigne l'ensemble des sommets avec hauteur n de T. La version infinie de T, notée $\widehat{T}$, a sa loi caractérisé par les formules

$$\mathbb{E}\left[f(\widehat{T}_{\leqslant n})\right] = \frac{1}{b_i}\mathbb{E}\left[X_n f(T_{\leqslant n})\right].$$

On vérifie que ceci définit bien la loi d'un arbre infini, notamment car le processus $(X_n)_{n\in\mathbb{Z}_+}$ est une martingale. On peut décrire $\widehat{T}$ de la manière suivante : il est composé d'une unique ligne infinie partant de la racine appelée *épine dorsale*, dont les sommets ont une loi de reproduction spéciale notée $\widehat{\zeta}$, définie par

$$\widehat{\zeta}^{(j)}(\mathbf{w}) = \frac{1}{b_j}\sum_{l=1}^{k} b_{w_l}\zeta^{(j)}(\mathbf{w}),$$

où j est un type et $\mathbf{w} = (w_1, \ldots, w_k)$ est une liste de types, et les autres sommets utilisent la loi de reproduction usuelle ζ. Sachant que u est dans l'épine dorsale et que $\mathbf{w}$ est la liste des types des enfants de u, l'élément suivant de l'épine dorsale sera le j-ème enfant avec probabilité proportionelle à b_{w_j}.

Convergence de l'arbre conditionné vers un arbre infini. Tout comme dans le cas à un type, nous obtenons que l'arbre T, conditionné à être "grand", converge vers $\widehat{T}$. Il faut cependant faire un bon choix de conditionnement. Il se trouve que le plus adapté est de conditionner par rapport au nombre d'éléments d'un unique type, et c'est pourquoi nous fixons un type $j \in [K]$.

Comme dans le cas monotype, des problèmes de périodicité apparaissent : le nombre d'éléments de type j dans un arbre T dont la racine a pour type i est nécessairement de la forme $\beta_i + dn$ avec $n \in \mathbb{Z}_+$. Nous n'expliciterons pas β_i et d dans cette introduction.

Théorème 5.1. *Soit T un arbre de Galton-Watson avec loi de reproduction ζ et dont la racine a pour type i. Pour $n \in \mathbb{Z}_+$ tel que $\mathbb{P}(\#_j T = \beta_i + dn) > 0$, soit T_n une version de T conditionnée à avoir $\beta_i + dn$ sommets de type j. Alors T_n converge en loi vers $\widehat{T}$ pour la convergence locale des arbres multi-types.*

Ce théorème contient en particulier le fait que T_n est bien défini pour n assez grand, autrement dit que $\mathbb{P}(\#_j T = \beta_i + dn) > 0$ à partir d'un certain rang.

5.3 Une application pour les cartes planaires

La motivation du Théorème 5.1 était en fait d'en déduire un résultat de convergence locale de grandes cartes planaires aléatoires. Donnons d'abord quelques rappels.

Cartes planaires et lois de Boltzmann. Une carte planaire m est un plongement d'un graphe connexe fini dans la sphère fait en sorte que les arêtes ne se croisent pas, et pris aux homéomorphismes directs de la sphère près. Ces objets connaissent un grand succès chez les probabilistes depuis le premier travail de Schaeffer [74] et le papier d'Angel et Schramm [7], notamment en ce qui concerne l'étude de la convergence d'échelle et de la convergence locale de grandes cartes aléatoires. On appelle *face* de la carte m toute composante connexe du complémentaire de m dans la sphère, et on appelle *degré* d'une face le nombre d'arêtes qui la touchent, en comptant deux fois toute arête rencontrée deux fois en faisant le tour de la face. L'ensemble des faces de m est noté $\mathcal{F}_m$.

Les cartes que nous étudierons seront enracinées, non pas en un sommet comme les arbres mais en une arête orientée e, partant d'un sommet e^- et pointant vers un sommet e^+. Elles seront également pointées, au sens ou un sommet r sera distingué. Ce pointage sert à l'étude de cartes aléatoires mais n'interviendra pas dans la convergence. On note $\mathcal{M}$ l'ensemble de toutes les cartes enracinées et pointées.

Soit $\mathbf{q} = (q_n)_{n\in\mathbb{N}}$ une suite de poids positifs ou nuls. On utilise $\mathbf{q}$ pour associer un poids à chaque carte enracinée pointée :

$$W_{\mathbf{q}}(m,e,r) = \prod_{f\in\mathcal{F}_m} q_{\deg(f)}$$

Si la suite $\mathbf{q}$ est admissible au sens où la somme

$$Z_{\mathbf{q}} = \sum_{(m,e,r)\in\mathcal{M}} W_{\mathbf{q}}(m,e,r)$$

est finie, alors on peut définir une mesure de probabilité dite de Boltzmann $B_{\mathbf{q}}$ par

$$B_{\mathbf{q}}(m,e,r) = \frac{W_{\mathbf{q}}(m,e,r)}{Z_{\mathbf{q}}}.$$

Autrement dit, quand la somme totale des poids est finie, alors la probabilité d'obtenir une carte particulière (m,e,r) est proportionelle à son poids.

L'admissibilité de la suite de poids $\mathbf{q}$ peut être caractérisée d'une autre manière. Pour deux réels positifs ou nuls x et y, on pose

$$f^\bullet(x,y) = \sum_{k,k'} \binom{2k+k'+1}{k+1}\binom{k+k'}{k} q_{2+2k+k'}\, x^k y^{k'}$$

ainsi que

$$f^\diamond(x,y) = \sum_{k,k'} \binom{2k+k'}{k}\binom{k+k'}{k} q_{1+2k+k'}\, x^k y^{k'}.$$

On sait d'après [65] que $\mathbf{q}$ est admissible si et seulement si le système

$$1 - \frac{1}{x} = f^\bullet(x,y) \tag{2}$$

$$y = f^\diamond(x,y) \tag{3}$$

admet une solution telle que le rayon spectral de la matrice

$$\begin{pmatrix} 0 & 0 & x-1 \\ \frac{x}{y}\partial_x f^\diamond(x,y) & \partial_y f^\diamond(x,y) & 0 \\ \frac{x^2}{x-1}\partial_x f^\bullet(x,y) & \frac{xy}{x-1}\partial_y f^\bullet(x,y) & 0 \end{pmatrix}$$

soit inférieur ou égal à 1. Cette solution est nécessairement unique, et on pose donc $Z^+_{\mathbf{q}} = x$ et $Z^\diamond_{\mathbf{q}} = y$. Pour des raisons qui apparaitront plus tard, on dira de plus que $\mathbf{q}$ est critique si le rayon spectral de la matrice est en fait 1.

Convergence vers une carte infinie. Posons

$$d = \mathrm{PGCD}\big(\{m \in \mathbb{N},\, q_{2m+2} > 0\} \cup \{m \in 2\mathbb{Z}_+ + 1,\, q_{m+2} > 0\}\big).$$

On démontre que le nombre de sommets d'une carte M de loi $B_{\mathbf{q}}$ est nécessairement de la forme $2 + dn$ avec avec $n \in \mathbb{Z}_+$ et que, pour n suffisamment grand, M a effectivement $2 + dn$ sommets avec probabilité non-nulle.

Théorème 5.2. *Pour $n \in \mathbb{N}$ suffisamment grand, soit M_n une carte de loi $B_{\mathbf{q}}$ conditionnée à avoir $2+dn$ sommets. Si $\mathbf{q}$ est critique, alors M_n converge en loi vers une carte infinie enracinée M_∞ quand n tend vers l'infini.*

Quelques remarques sur la notion de carte infinie : on appelle ici carte infinie enracinée toute suite $(M^i, e^i)_{i\in\mathbb{Z}_+}$ de cartes enracinées (non pointées) finies telle que, pour $i \leqslant j$, M^i soit égale à la boule de rayon i dans M^j : la carte formée par les sommets et arêtes de M^j à distance inférieure à i de $(e^j)^-$. Ces objets peuvent avoir une structure compliquée, et notamment il se peut que tout dessin dans le plan des $(M^i, e^i)_{i\in\mathbb{Z}_+}$ admette des points d'accumulation (au sens de la topologie du plan). On démontre que ça n'est pas le cas pour M_∞, qui reste ainsi une carte "planaire". La convergence en loi mentionée dans le Théorème 5.2 est alors une convergence locale, au sens où chaque boule de rayon i de M_n converge en loi vers celle de M_∞.

La preuve du Théorème 5.2, comme tant de preuves de résultats phares concernant des cartes, utilise une méthode bijective. Inaugurées par la bijection de Cori-Vauquelin-Schaeffer ([23] puis [74]), ces méthodes consistent à mettre en correspondance les cartes utilisées avec une certaine famille d'arbres et à déduire des propriétés asymptotiques des arbres un théorème limite sur les cartes. Nous utilisons ici la bijection introduite par Bouttier, Di Francesco et Guitter dans [19], qui permet de tranformer tout carte enracinée et pointée en un arbre de à quatre types avec des étiquettes sur ses sommets, et, de plus, envoie toute carte de loi Boltzmann critique vers un arbre de Galton-Watson critique. Cette opération s'avérant être continue, le Théorème 5.2 devient un corollaire du Théorème 5.1.

Cas des p-angulations. Soit $p \geqslant 3$ un entier. On appelle p-angulation toute carte planaire dont toutes les faces ont pour degré p. Les cas $p = 3$ et $p = 4$ ont déjà fait l'object de plusieurs travaux.

Pour $p = 3$, on a depuis [7] la convergence en loi de triangulations uniformes à n sommets, si on se restreint aux triangulations sans boucles (arêtes reliant un sommet à lui-même) ou bien sans boucles ni arêtes multiple, vers une carte infinie appelée *triangulation planaire infinie uniforme* ou UIPT (pour l'anglais Uniform Infinite Planar Triangulation). Du côté $p = 4$, Krikun a montré dans [53] que la quadrangulation uniforme à n face converge en distribution vers ce qu'on appelle la *quadrangulation planaire infinie* ou UIPQ. L'UIPQ fut également étudiée à l'aide de méthodes bijectives dans [21] et [26].

Le Théorème 5.2 s'adapte directement aux p-angulations, car on peut les représenter comme des cartes Boltzmann critiques, et il est équivalent de les conditionner par leur nombre de faces ou leur nombre de sommets. On obtient alors le résultat suivant :

Proposition 5.1. *Pour $n \geqslant 2$, soit M_n une p-angulation uniforme parmi les p angulations à n faces si n est pair et à $2n$ faces si n est impair. Alors M_n converge en loi vers une carte infinie appelée* p-angulation planaire infinie uniforme.

Récurrence de la carte infinie. Nous achevons ce travail par l'étude d'une première propriété de la carte infinie : nous montrons qu'elle forme un graphe récurrent, au sens où la marche aléatoire uniforme retourne presque sûrement à son point de départ en temps fini. Ceci est déjà connu depuis l'article [38] de Gurel-Gurevich et Nachmias pour l'UIPQ et l'UIPT, mais l'outil principal de cet article sert aussi à montrer la récurrence de M_∞ pour toute suite de poids critique $\mathbf{q}$.

Soit $(G_n, \rho_n)_{n\in\mathbb{N}}$ une suite de graphes planaires finis enracinés aléatoires qui converge localement en loi vers un graphe infini (G, ρ). On suppose que, conditionnellement à G_n, sa racine ρ_n est choisie selon la mesure invariante de la marche aléatoire simple (c'est-à-dire que ρ_n est égal à un sommet x avec probabilité proportionelle à son degré, qui est son nombre de voisins), et que le degré de ρ dans G est borné exponentiellement au sens où

$$\mathbb{P}[deg(\rho) > n] \leqslant \mathrm{e}^{-cn}$$

pour n assez grand et une certaine constante $c > 0$. Le Théorème 1.1 de [38] énonce alors que le graphe G est récurrent.

Ceci s'applique aux cartes de loi Boltzmann conditionnées. Notamment, si une carte enracinée et pointée (m, e, r) a pour loi $B_{\mathbf{q}}$, on voit que l'arête orientée e est choisie uniformément, et donc le sommet e^- est distribué selon la mesure invariante. Le point important est donc de montrer que le degré de ρ dans la carte M_∞ est borné exponentiellement, ce qui se démontre à l'aide de la construction de la carte par la bijection BDFG.

6 Conclusion et perspectives futures

Nous avons introduit dans cette thèse plusieurs nouveaux objets mathématiques. Naturellement, avec de nouveaux objets apparaissent aussi de nouvelles questions. Nous mentionnons ici quelques nouvelles avenues d'exploration.

- Le Chapitre 3 constitue une première étude de limites d'échelle d'arbres grandissant aléatoirement par une construction algorithmique. On s'attend à ce que nos résultats se généralisent et que l'on obtienne des limites similaires quand, au lieu d'ajouter une étoile déterministe au milieu de l'arête choisie uniformément, on ajoute une structure un peu plus riche. On peut notamment prendre le même algorithme qu'au Chapitre 3, mais choisir le degré k du nouveau sommet aléatoirement à chaque étape, de manière i.i.d. Si la loi commune admet une moyenne finie m, alors on s'attend à ce que la taille de l'arbre à l'étape n

soit de l'ordre de $n^{1/m}$. Il est possible de définir un arbre qui est le candidat pour l'arbre limite, dont on s'attend à ce qu'il soit un arbre de fragmentation. Si de plus les degrés sont bornés par un entier K, on s'attend aussi à ce que cet arbre se plonge naturellement dans $\mathcal{T}_K$.

- Soit T un arbre de Galton-Watson à un type, de loi de reproduction μ sous-critique, c'est-à-dire dont la moyenne est strictement inférieure à 1. Si l'on note T_n une version de T conditionnée à avoir $1 + dn$ sommets (en prenant toujours d le PGCD du support de la mesure μ), alors il est connu que T_n ne converge pas toujours vers un arbre biaisé par la taille. Dans certains cas dits *non-génériques*, on obtient un arbre limite $\mathcal{T}^*$ de nature très différente : il est constituée d'une épine dorsale *finie*, dont la longueur est une variable géométrique, et dont le sommet final a un nombre infini d'enfants. On appelle l'apparition d'un tel sommet un phénomène de condensation. On s'attend à ce que ceci s'applique aussi aux arbres multi-types, et que cela s'étende aux cartes. Quelques progrès ont déjà été faits dans cette direction dans le récent article [18] dans le cadre des cartes biparties, c'est-à-dire les cartes dont les faces ont toutes des degrés pairs.

- Notre Théorème 5.2 s'applique à toute suite de poids critique $\mathbf{q}$. Malheureusement, une grande partie de l'ensemble de ses suites est inconnue : on ne connait essentiellement que les suites à support fini (d'après la Proposition A.2 de [25], si $\mathbf{q}$ a est à support fini, alors on peut la rendre critique en la multipliant par une constante bien choisie) ainsi que certains cas bipartis étudiés dans [58]. Il serait intéressant de montrer que, si q_n est équivalent à $n^{-\alpha}$ avec $\alpha > 3/2$, alors la suite $\mathbf{q}$ peut être modifiée de sorte à être critique tout en donnant un poids assez élevé aux faces de haut degré.

Chapter 1

Preliminary results on measured $\mathbb{R}$-trees

Contents

We set up in this chapter the necessary background on continuum trees which will be used throughout the next two chapters. It consists mostly of definitions and classical results on the Gromov-Hausdorff-Prokhorov topology, but it also contains some new non-trivial work, most notably Proposition 2.1 which gives us a simple way of defining a measure on any compact $\mathbb{R}$-tree.

1 $\mathbb{R}$-trees and the GHP topology

1.1 $\mathbb{R}$-trees

Definition 1.1. *Let $(\mathcal{T}, d)$ be a metric space. We say that it is an* $\mathbb{R}$-tree *if it satisfies the following two conditions:*

- *for all $x, y \in \mathcal{T}$, there exists a unique distance-preserving map $\phi_{x,y}$ from $[0, d(x,y)]$ into $\mathcal{T}$ such $\phi_{x,y}(0) = x$ and $\phi_{x,y}(d(x,y)) = y$.*
- *for all continuous and one-to-one functions c: $[0,1] \to \mathcal{T}$, we have $c\big([0,1]\big) = \phi_{x,y}\big([0, d(x,y)]\big)$, where $x = c(0)$ and $y = c(1)$.*

For any x, y in a tree, we will denote by $[\![x, y]\!]$ the image of $\phi_{x,y}$, i.e. the path between x and y.

Here is a simple characterization of $\mathbb{R}$-trees which we will use in the future. It can be found in [31], Theorem 3.40.

Proposition 1.2. *A metric space $(\mathcal{T}, d)$ is an $\mathbb{R}$-tree if and only if it is connected and satisfies the following property, called the* four-point condition*:*

$$\forall x, y, u, v \in \mathcal{T}, d(x,y) + d(u,v) \leqslant \max\big(d(x,u) + d(y,v), d(x,v) + d(y,u)\big).$$

By permuting x, y, u, v, one gets a more explicit form of the four-point condition: out of the three numbers $d(x,y) + d(u,v)$, $d(x,u) + d(y,v)$ and $d(x,v) + d(y,u)$, at least two are equal, and the third one is smaller than or equal to the other two.

For commodity we will, for an $\mathbb{R}$-tree $(\mathcal{T}, d)$ and $a > 0$, call $a\mathcal{T}$ the $\mathbb{R}$-tree $(\mathcal{T}, ad)$ which is the same tree as $\mathcal{T}$, except that all distances have been rescaled by a.

1.2 Roots, partial orders and height functions

All the trees which we will consider will be *rooted*: we will fix a distinguished vertex ρ called the *root*. This provides $\mathcal{T}$ with a great amount of additional structure, the first one being the notion of *height*: for $x \in \mathcal{T}$ we call height of x its distance to the root, and write it as $ht(x) = d(\rho, x)$. We let also the total height of the tree $ht(\mathcal{T})$ be the supremum of the heights of all its points:

$$ht(\mathcal{T}) = \sup_{x \in \mathcal{T}} ht(x).$$

We use the height function to define, for $t \geqslant 0$, the subset $\mathcal{T}_{\leqslant t} = \{x \in \mathcal{T} : \ ht(x) \leqslant t\}$, as well as the similarly defined $\mathcal{T}_{<t}$, $\mathcal{T}_{\geqslant t}$ and $\mathcal{T}_{>t}$. Note that $\mathcal{T}_{\leqslant t}$ and $\mathcal{T}_{<t}$ are both $\mathbb{R}$-trees, as well as each of the connected components of $\mathcal{T}_{\geqslant t}$ and $\mathcal{T}_{>t}$, which we will call the *tree components* of $\mathcal{T}_{\geqslant t}$ and $\mathcal{T}_{>t}$.

Having fixed a root also lets us define a partial order on $\mathcal{T}$, by declaring that $x \leqslant y$ if $x \in [\![\rho, y]\!]$. We will often take a genealogical standpoint and say that x is an ancestor of y in this case, or simply that x is lower than y. We can then define for any x in $\mathcal{T}$ the *subtree of $\mathcal{T}$ rooted at x*, which we will call $\mathcal{T}_x$: it is the set $\{y \in \mathcal{T} : y \geqslant x\}$ of all the descendants of x. We will also say that two points x and y are *on the same branch* if they are comparable, i.e. if we have $x \leqslant y$ or $y \leqslant x$. For every subset S of $\mathcal{T}$ we can define the *greatest common ancestor* of S, which is the highest point to be lower than all the elements of S. The greatest common ancestor of two points x and y of $\mathcal{T}$ will be written $x \wedge y$.

One convenient property is that we can recover the metric from the order and the height function. Indeed, for any two points x and y, we have $d(x,y) = ht(x) + ht(y) - 2ht(x \wedge y)$.

Points of $\mathcal{T}$ can be classified into several categories. We call *degree* of $x \in \mathcal{T}$ the number of connected components of $\mathcal{T}_x \setminus \{x\}$. Points of degree 0 will be called *leaves*, and their set will be called $\mathcal{L}(\mathcal{T})$. We also call *leaf* of $\mathcal{T}$ any point L such that the $\mathcal{T}_L = \{L\}$. The set of leaves of $\mathcal{T}$ will be written $\mathcal{L}(\mathcal{T})$, and its complement is called the *skeleton* of $\mathcal{T}$. Amongst the point of the skeleton, those with degree greater than or equal to 2 are called *branch points*.

1.3 The Gromov-Hausdorff and Gromov-Hausdorff-Prokhorov metrics, spaces of trees

Recall that, if A and B are two compact nonempty subsets of a metric space $(\mathcal{Z}, d_{\mathcal{Z}})$, then the Hausdorff distance between A and B is defined by

$$d_{\mathcal{Z},\mathrm{H}}(A,B) = \inf\{\varepsilon > 0 : A \subset B^{\varepsilon} \text{ and } B \subset A^{\varepsilon}\},$$

where A^{ε} and B^{ε} are the closed ε-enlargements of A and B (that is, $A^{\varepsilon} = \{x \in E, \exists a \in A, d(x,a) \leqslant \varepsilon\}$ and the corresponding definition for B).

Now, if one considers two compact rooted $\mathbb{R}$-trees $(\mathcal{T}, \rho, d)$ and $(\mathcal{T}', \rho', d')$, define their *Gromov-Hausdorff distance*:

$$d_{\mathrm{GH}}(\mathcal{T}, \mathcal{T}') = \inf\left[\max\Big(d_{\mathcal{Z},\mathrm{H}}\big(\phi(\mathcal{T}), \phi'(\mathcal{T}')\big), d_{\mathcal{Z}}\big(\phi(\rho), \phi'(\rho')\big)\Big)\right],$$

where the infimum is taken over all pairs of isometric embeddings ϕ and ϕ' of $\mathcal{T}$ and $\mathcal{T}'$ in the same metric space $(\mathcal{Z}, d_{\mathcal{Z}})$.

We will also want to consider pairs $(\mathcal{T}, \mu)$, where $\mathcal{T}$ (d and ρ being implicit) is a compact rooted $\mathbb{R}$-tree and μ a Borel probability measure on $\mathcal{T}$. Between two such compact rooted measured trees $(\mathcal{T}, \mu)$ and $(\mathcal{T}', \mu')$, one can define the *Gromov-Hausdorff-Prokhorov* distance by

$$d_{\mathrm{GHP}}\big((\mathcal{T},\mu), (\mathcal{T}',\mu')\big) = \inf\left[\max\Big(d_{\mathcal{Z},\mathrm{H}}\big(\phi(\mathcal{T}), \phi'(\mathcal{T}')\big), d_{\mathcal{Z}}\big(\phi(\rho), \phi'(\rho')\big), d_{\mathcal{Z},\mathrm{P}}(\phi_*\mu, \phi'_*\mu')\Big)\right],$$

where the infimum is taken on the same space, and $d_{\mathcal{Z},\mathrm{P}}$ denotes the Prokhorov distance between two Borel probability measures on $\mathcal{Z}$, defined by

$$d_{\mathcal{Z},\mathrm{P}}(\nu, \nu') = \inf\big\{\varepsilon > 0 \ : \ \forall A \in \mathcal{B}(E), \nu(A) \leqslant \nu'(A^{\varepsilon}) + \varepsilon \text{ and } \nu'(A) \leqslant \nu(A^{\varepsilon}) + \varepsilon\big\}.$$

for two probability measures ν and ν'. It is well-known that convergence for $d_{\mathcal{Z},P}$ is equivalent to weak convergence of Borel probability measures on $\mathcal{Z}$ as soon as $\mathcal{Z}$ is separable, see for example [17], Section 6. /

The metrics d_{GH} and d_{GHP} allow us to study spaces of trees. We let $\mathbb{T}$ (respectively $\mathbb{T}_W$) be the set of equivalence classes of compact rooted $\mathbb{R}$-trees (respectively compact, rooted and measured $\mathbb{R}$-trees), where two trees are said to be equivalent if there is a root-preserving (respectively root-preserving and measure-preserving) bijective isometry between them. It can be shown (see in particular [32] and [2]) that these spaces are well-behaved.

Proposition 1.3. *The two functions d_{GH} and d_{GHP} are metrics on respectively $\mathbb{T}$ and $\mathbb{T}_W$, and the spaces $(\mathbb{T}, d_{\mathrm{GH}})$ and $(\mathbb{T}_W, d_{\mathrm{GHP}})$ are separable and complete.*

We note that the topology induced on $\mathbb{T}_W$ by d_{GHP} was first introduced in [36], and was also studied with a different metric in [33].

In practice, every $\mathbb{R}$-tree which we study will be embedded in the space ℓ^1 of summable real-valued sequences:

$$\ell^1 = \{x = (x_i)_{i\in\mathbb{N}} : \sum_{i=1}^{\infty} |x_i| < \infty\},$$

which is equipped with its usual metric d_{ℓ^1}. Proof of convergence for d_{GH} and d_{GHP} will thus always be done by proving the Hausdorff convergence of the sets and the Prokhorov convergence of the measures.

1.4 Identifying trees by using their finite-dimensional marginals

We recall here what is essentially the method used by Aldous in [3], [4] and [5] to study trees, which we will need in Chapter 3, Section 4.1. Let $(\mathcal{T}, d, \rho, \mu)$ be a random rooted measured compact $\mathbb{R}$-tree. Conditionally on $(\mathcal{T}, d, \rho, \mu)$, for $k \in \mathbb{N}$, let $X_1, \ldots, X_k$ be independant points with distribution μ. We call *k-dimensional marginal of $\mathcal{T}$* the random $\mathbb{R}$-tree

$$\Big(\bigcup_{i=1}^{k} [\![\rho, X_i]\!], d, \rho, \frac{\sum_{i=1}^{k} \delta_{X_i}}{k}\Big).$$

Two properties are worthy of note:

- if the finite-dimensional marginals of the trees $(\mathcal{T}, \mu)$ and $(\mathcal{T}', \mu')$ have the same distribution and if we know that μ and μ' are fully supported on $\mathcal{T}$ and $\mathcal{T}'$, then $\mathcal{T}$ and $\mathcal{T}'$ also have the same distribution. This is because, as k goes to infinity, the k-dimensional marginals converge almost surely to the tree in the GHP sense.

- if a sequence of trees $(\mathcal{T}_n, \mu_n)$ converges in distribution to a tree $(\mathcal{T}, \mu)$, then the finite-dimensional marginals of $(\mathcal{T}_n, \mu_n)$ converge in distribution to those of $(\mathcal{T}, \mu)$. This is because convergence of the marginals corresponds to convergence for a weaker topology called the Gromov-Prokhorov topology, see in particular [37].

2 Nonincreasing functions and measures on trees

It is well-known that probability distributions, and in fact all Borel measures on $\mathbb{R}$ are characterised by their cumulative distribution function. If μ is such a measure, we let, for $x \in \mathbb{R}$, $F_\mu(x) = \mu\big((\infty, x]\big)$, which is a right-continuous non-decreasing function which tends to 0 at $-\infty$. Conversely, for such a function F, there is a unique measure μ such that $F_\mu = F$. We establish here an analoguous theory for compact and rooted $\mathbb{R}$-trees.

Let $\mathcal{T}$ be a compact rooted tree. Let m be a non-increasing function on $\mathcal{T}$ (for the natural ordering on $\mathcal{T}$) taking values in $[0, \infty)$. As with monotone functions of a real variable, m satisfies some limiting properties everywhere. One can easily define the left-limit $m(x^-)$ of m at any point $x \in \mathcal{T}$, since $[\![\rho, x]\!]$ is isometric to a line segment, for example by setting $m(x^-) = \lim_{t\to ht(x)^-} m(\phi_{\rho,x}(t))$ (recall from Definition 1.1 that $\phi_{\rho,x}$ is the parametrization of the path from ρ to x). We say that m is *left-continuous* at a point x if $m(x^-) = m(x)$. Let us also define the *additive right-limit* $m(x^+)$: since $\mathcal{T}$ is compact, the set $\mathcal{T}_x \setminus \{x\}$ has countably many connected components, say $(\mathcal{T}_i)_{i\in S}$ for a finite or countable set S. Let, for all $i \in S$, $x_i \in \mathcal{T}_i$. We then set

$$m(x^+) = \sum_{i\in S} \lim_{t\to ht(x)^+} m(\phi_{\rho,x_i}(t)).$$

This is well-defined, because it does not depend on our choice of $x_i \in \mathcal{T}_i$ for all i. Note also that, by monotone convergence, we could swap the sum and the limit.

If μ is a finite Borel measure on $\mathcal{T}$, then the function m defined by $m(x) = \mu(\mathcal{T}_x)$ for $x \in \mathcal{T}$ is easily seen to be non-increasing and left-continuous. One can also check that $m(x) - m(x^+) = \mu(\{x\})$ for all $x \in \mathcal{T}$, and thus $m(x) \geqslant m(x^+)$. This function in fact completely characterizes the measure μ.

Proposition 2.1. *Let m be a decreasing, positive and left continuous function on $\mathcal{T}$ such that, for all $x \in \mathcal{T}$, $m(x) \geqslant m(x^+)$. Then there exists a unique Borel measure μ on $\mathcal{T}$ such that*

$$\forall x \in \mathcal{T}, \mu(\mathcal{T}_x) = m(x).$$

The proof of Proposition 2.1 is non-trivial and requires its own section.

2.1 Proof of Proposition 2.1

We will want to apply a variant of Caratheodory's extension theorem to a natural semi-ring of subsets of the tree $\mathcal{T}$ which generates the Borel topology. The reader is invited to look in [28] for definitions and its Theorem 3.2.4 which is the one we will use.

Definition 2.2. *Let $x \in \mathcal{T}$, and C be a finite subset of $\mathcal{T}_x$. We say that C is a* pre-cutset *of $\mathcal{T}_x$ if $x \leqslant y$ for all $y \in C$ and none of the elements of C are on the same branch as another. We then let $B(x, C) = \mathcal{T}_x \setminus \bigcup_{y\in C} \mathcal{T}_y$. Such a set is called a* pre-ball. *We let $\mathcal{B}$ be the set of all pre-balls of $\mathcal{T}$.*

We introduce the notation $[k] = \{1, 2, \dots, k\}$ for a positive integer k. Note that any set of the form $\mathcal{T}_x \setminus \bigcup_{i\in[k]} \mathcal{T}_{x_i}$ is a pre-ball, even if one does not specify that $\{x_i,\ i \in [k]\}$ is a pre-cutset of $\mathcal{T}_x$. Indeed, if x is not on the same branch as x_i for some i, then we can remove this one from the union, if we have $x_i \leqslant x$ for some i then we have just written the empty set, and, if for some $i \neq j$, we have $x_i \leqslant x_j$, we might as well remove x_j from the union. All these removals leave us with a pre-cutset of $\mathcal{T}_x$. Also note that, given a pre-ball B, there exists a unique $x \in \mathcal{T}$ and a finite pre-cutset C which is unique up to reordering such that $B = B(x, C)$.

Lemma 2.3. *$\mathcal{B}$ is a semi-ring which contains all the $\mathcal{T}_x$ for $x \in \mathcal{T}$, and it generates the Borel σ-field of $\mathcal{T}$.*

Proof. The fact that $\mathcal{D}$ contains all the sets of the form $\mathcal{T}_x$ for $x \in \mathcal{T}$, as well as the empty set, is in the definition. Stability by intersection is easily proven: let $B\big(x, (x_i)_{i\in[k]}\big)$ and $B\big(y, (y_i)_{i\in[l]}\big)$ be two pre-balls. If x and y are not on the same branch, then the intersection is the empty set, and otherwise, we can assume $y \geqslant x$, and we are left with $\mathcal{T}_x \setminus \big(\bigcup_{i\in[k]} \mathcal{T}_{x_i} \cup \bigcup_{j\in[l]} \mathcal{T}_{y_j}\big)$ which is indeed a pre-ball.

Now let $B\big(x, (x_i)_{i\in[k]}\big)$ and $B\big(y, (y_i)_{i\in[l]}\big)$ be two pre-balls, we want to check that $B\big(x, (x_i)_{i\in[k]}\big) \setminus B\big(y, (y_i)_{i\in[l]}\big)$ is a finite union of disjoint pre-balls. Exceptionally, we will write here for any subset A of $\mathcal{T}$, $\bar{A} = \mathcal{T} \setminus A$, for clarity's sake. We have:

$$\begin{aligned} B\big(x, (x_i)_{i\in[k]}\big) \cap \overline{B}\big(y, (y_i)_{i\in[l]}\big) &= \mathcal{T}_x \cap \bigcap_{i\in[k]} \bar{T}_{x_i} \cap (\bar{T}_y \cup \bigcup_{y\in[l]} \mathcal{T}_{y_i}) \\ &= (\mathcal{T}_x \cap \bar{\mathcal{T}}_y \cap \bigcap_{i\in[k]} \bar{T}_{x_i}) \cup \bigcup_{y\in[l]} (\mathcal{T}_x \cap \mathcal{T}_{y_i} \cap \bigcap_{i\in[l]} \bar{T}_{x_i}). \end{aligned}$$

Since for every i, $\mathcal{T}_x \cap \mathcal{T}_{y_i}$ is either equal to $\mathcal{T}_x$ or $\mathcal{T}_{y_i}$, we do have a finite union of pre-balls. This union is also disjoint, because $\bar{\mathcal{T}}_y, \mathcal{T}_{y_1}, \ldots, \mathcal{T}_{y_l}$ are all disjoint.

Finally, we want to check that $\mathcal{D}$ does indeed span the Borel σ-field of $\mathcal{T}$, which will be proven by showing that every open ball in $\mathcal{T}$ is the intersection of a countable amount of pre-balls. Let $x \in \mathcal{T}$ and $r \geqslant 0$, and let B the closed ball centered at x with radius r. Let y be the unique ancestor of x such that $ht(y) = (ht(x) - r) \vee 0$. Since $\mathcal{T}_y$ is compact and $B \in \mathcal{T}_y$ is open, we know that $\mathcal{T}_y \setminus B$ has a countable amount of closed tree components, which we will call $(\mathcal{T}_{x_i})_{i\in\mathbb{N}}$. Writing out $B = \big(\mathcal{T}_y \setminus \bigcup_{i\in\mathbb{N}} \mathcal{T}_{x_i}\big) \setminus \{y\}$ then shows that it is indeed a countable intersection of pre-balls. As a consequence, there exists at most one measure on $\mathcal{T}$ such that $\mu(\mathcal{T}_x) = m(x)$ for all $x \in \mathcal{T}$: uniqueness in Proposition 2.1 is proven.

□

Lemma 2.4. *For every $x \in \mathcal{T}$ and every finite pre-cutset C, we let*

$$\mu\big(B(x, C)\big) = m(x) - \sum_{y\in C} m(y).$$

This defines a nonnegative function on $\mathcal{D}$ which is σ-additive.

Proof. Let us first prove the non-negativity of μ. This can be done by induction on the number of elements k in the pre-cutset $C = \{x_i,\ i \in [k]\}$ of $\mathcal{T}_x$. If $k = 0$ then there is nothing to do, since $\mu(B(x, \emptyset)) = m(x) \geqslant 0$ by definition. Now assume $k \geqslant 1$ and that non-negativity has been proved up to $k - 1$. Let y be the greatest common ancestor of all the $(x_i)_{i\in[k]}$, we have $x \leqslant y$, and thus $m(x) \geqslant m(y)$, and it will suffice to prove $m(y) - \sum_{i=1}^{k} m(x_i) \geqslant 0$. The set $T_y \setminus \{y\}$ has a finite, but strictly greater than 1 number of connected components which contain the points $(x_i)_{i\in[k]}$, let us call them $C_1, \ldots, C_l$, with $1 \leqslant l \leqslant k$. Since every C_l contains at most $l-1 \leqslant k-1$ elements from the $(x_i)_{i\in[k]}$, one can use the induction hypothesis in every C_j: for all j, let $y_j \in C_j$ be such that, for all i such that $x_i \in C_j$, $y_j \leqslant x_i$, then we have $m(y_j) \geqslant \sum_{i:\, x_i\in C_j} m(x_i)$. Now, by letting every y_j converge to y, we end up with

$$m(y) \geqslant m(y^+) \geqslant \sum_j \lim_{y_j\to y^+} m(y_j) \geqslant \sum_i m(x_i)$$

which ends the proof of that μ is nonnegative.

The proof that μ is σ-additive on $\mathcal{D}$ will be done in three steps. First, we will prove that it is finitely additive, i.e. that, if a pre-ball can be written as a finite disjoint union of pre-balls, then the μ-masses add up properly. Next, we will prove that it is finitely subadditive, which means that if a pre-ball B can be written as a subset of the finite union of other pre-balls $B_1, \ldots, B_n$, we have $\mu(B) \leqslant \sum_i \mu(B_i)$. The σ-additivity itself will then be proved by proving both inequalities separately.

First, we want to show that μ is finitely additive, i.e. that if a pre-ball $B = B\big(x, (x_i)_{i\in[k]}\big)$ can be written as the disjoint union of pre-balls $B^j = B\big(x^j, (x_i^j)_{i\in[k^j]}\big)$ for $1 \leqslant j \leqslant n$, we have

$\mu(B) = \sum_j \mu(B^j)$. Note that since $\mathcal{D}$ is not stable under union, one cannot simply prove this for $n = 2$ and then do a simple induction. We will indeed do an induction on n, but it will be a bit more involved. The initial case, $n = 1$ is immediate. Now assume that $n \geqslant 2$ and that, for every pre-ball which can be written as the disjoint union of fewer than $n - 1$ pre-balls, the masses add up, and let $B = B\big(x, (x_i)_{i\in[k]}\big)$ be a pre-ball which is the union of $B^j = B\big(x^j, (x^j_i)_{i\in[k^j]}\big)$ for $1 \leqslant j \leqslant n$. We are first going to show that we can restrict ourselves to the case where $B = \mathcal{T}_x$. To do this, first notice that, since the union is disjoint, for every i with $1 \leqslant i \leqslant k$, there is only one j, which we will call $j(i)$, such that x_i is in the set $\{x^j_p,\ p \in [k^j]\}$. Thus, if we add T_{x_i} to the pre-ball $B^{j(i)}$ and do this for all i, the result is that T_x (which is none other than $B \cup \bigcup_{i\in[k]} T_{x_i}$) is written as the disjoint union of pre-balls $A^j = B^j \cup \bigcup_{i:j(i)=j} T_{x_i}$. Since $\mu(T_x) = \mu(B) + \sum_{i=1}^k m(x_i)$ and, for all j, $\mu(A^j) = \mu(B_j) + \sum_{i:j(i)=j} m(x_i)$, it suffices consider the case when $B = \mathcal{T}_x$. By reordering, one can also assume that $x^1 = x$. Now, for every i with $1 \leqslant i \leqslant k^1$, consider the pre-balls B^j with j such that $x^1_i \leqslant x_j$. These are disjoint, and their union is none other than $\mathcal{T}_{x^1_i}$, and they are strictly less than n in number. The induction hypothesis then tells us that $\mu(T_{x^1_l})$ is the sum of $\mu(B^j)$ for such j. Repeat this for all i, and we get $\sum_{j=2}^n \mu(B^j) = \sum_{i=1}^{k^1} \mu(\mathcal{T}_{x^1_i}) = \mu(\mathcal{T}_x) - \mu(B^1)$, which is what we wanted.

Now we go on to μ's finite subadditivity. This can actually be proven with pure measure theory. Let B be a pre-ball and $B_1, \dots, B_n$ be pre-balls such that $B \subset \underset{i\in[n]}{\cup} B_i$. Let us first start with the case where $n = 1$, in other words, let us show that μ is nondecreasing: since $\mathcal{D}$ is a semi-ring, $B_1 \setminus B$ can be rewritten as a finite disjoint of pre-balls $C_1, \dots, C_k$, and by finite additivity, we have $\mu(B_1) = \mu(B) + \sum_j \mu(C_j) \geqslant \mu(B)$. Now, going back to the general case, one can assume that for every i, we have $B_i \subset B$, because if it is not the case, one can replace B_i by $B_i \cap B$. Now, consider the sequence C_i defined by $C_1 = B_1$ and, for $i \geqslant 2$, $C_i = B_i \setminus (B_1 \cup B_2 \ldots \cup B_{i-1})$. Since $\mathcal{D}$ is a semi-ring, every B_i can be written as the disjoint union of a finiteamount of pre-balls: for every i, there exists disjoint pre-balls $D_1(i), \dots, D_{k(i)}(i)$ such that $C_i = \overset{k(i)}{\underset{j=1}{\cup}} D_j(i)$. By finite additivity, we then have $\mu(B) = \sum_{i=1}^n \sum_{j=1}^{k(i)} \mu(D_j(i))$. Now all that is left to do is show that, for all i, we have $\sum_{j=1}^{k(i)} \mu(D_j(i)) \leqslant \mu(B_i)$, which is immediate because $B_i \setminus (\overset{k(i)}{\underset{j=1}{\cup}} D_j(i))$ is a disjoint finite union of pre-balls.

Finally, we can move on to μ's σ-additivity . Assume that a pre-ball $B = B\big(x, (x_i)_{i\in[k]}\big)$ can be written as the disjoint union of pre-balls $B^j = B\big(x^j, (x^j_i)_{i\in[k^j]}\big)$ for $j \in \mathbb{N}$. Let us first prove the easy inequality $\mu(B) \geqslant \sum_i \mu(B_i)$. Fix $n \in \mathbb{N}$, since $\mathcal{B}$ is a semi-ring, the set $B \setminus (\underset{1\leqslant i\leqslant n}{\cup} B_i)$ is a finite disjoint union of pre-balls, which we will call $C_1, \dots, C_k$. By finite additivity, we have $\mu(B) = \sum_{i=1}^n \mu(B_i) + \sum_{j=1}^k \mu(C_j) \geqslant \sum_{i=1}^n \mu(B_i)$, and we just need to take the limit. To prove the reverse inequality, we will slightly modify our sets so that we can get a open cover of a compact set, and bring ourselves back to the finite case. Let $\varepsilon > 0$. For every j such that $x^j \neq \rho$ (and ε small enough), let $x^j(\varepsilon)$ be an ancestor of x^j such that $m(x^j(\varepsilon)) - m(x^j) \leqslant \varepsilon 2^{-j-1}$, and if $x_j = \rho$ we keep $x^j(\varepsilon) = \rho$. In the same vein, for $1 \leqslant i \leqslant k$, we choose an ancestor $x_i(\varepsilon)$ such that $m(x_i(\varepsilon)) - m(x_i) \leqslant \frac{1}{k}$, and such that $(x_i(\varepsilon))_{i\in[k]}$ is still a pre-cutset of $\mathcal{T}_x$. Now consider, for every j, the open set D^j which is equal to $B\big(x^j(\varepsilon), (x^j_i)_{i\in[k^j]}\big) \setminus \{x^j(\varepsilon)\}$ if $x^j \neq \rho$, and equal to B^j otherwise. These form a cover of $B\big(x, (x_i(\varepsilon))_{i\in[k]}\big)$ and therefore also cover its closure, $B\big(x, (x_i(\varepsilon))_{i\in[k]}\big) \cup \bigcup_{i\in[k]} \{x_i(\varepsilon)\}$. Since $\mathcal{T}$ is compact, $B\big(x, (x_i(\varepsilon))_{i\in[k]}\big)$ can be covered by a finite amount of the D^j, which we can assume are $D^1,\dots,D^n$. We can then use finite subadditivity:

$$\begin{aligned}
\mu(B) = m(x) - \sum_{i=1}^k m(x_i) & \leqslant m(x) - \sum_{i=1}^k m(x_i(\varepsilon)) + \varepsilon \\
\leqslant \mu(B(x,(x_i(\varepsilon))_{i\in[k]})) + \varepsilon & \leqslant \sum_{j=1}^n \mu(D^j) + \varepsilon \\
\leqslant \quad \sum_{j=1}^\infty \mu(D^j) + \varepsilon & \leqslant \sum_{j=1}^\infty \left(\mu(B^j) + \varepsilon 2^{-j-1}\right) + \varepsilon \\
\leqslant \quad \sum_{j=1}^\infty \mu(B^j) + 2\varepsilon. &
\end{aligned}$$

This gives us our final inequality.

□

Theorem 3.2.4 of [28] ends the proof of Proposition 2.1.

3 Gromov-Hausdorff-Prokhorov convergence of discrete trees with edge-lengths

Let T be a rooted finite graph-theoretical tree: we think of it as a set of vertices equipped with a set of edges E. For any non-negative function l on E, we let (T_l, d_l) be the abstract $\mathbb{R}$-tree obtained from T by considering every edge e as a line segment with length $l(e)$. Note that we allow edges with length 0, in which case the edge is identified to a single point. Our aim in this section is to give GHP results on T_l as l varies.

3.1 The stick-breaking embedding

We give here a practical way of constructing the $\mathbb{R}$-tree T_l in ℓ^1. This is essentially a slightly more general version of the one introduced by Aldous in [3]. This embedding has the convenient property that, for $x \in \mathcal{T}$, the path from the root ρ (which is embedded as the null vector) to x is such that "the coordinates increase one at a time", in the sense that

$$[\![\rho, x]\!] = \bigcup_{n=0}^{\infty} [p_n(x), p_{n+1}(x)]$$

where, for $n \in \mathbb{Z}_+$, p_n denotes the natural "orthogonal" projection on the first n coordinates (p_0 being the null map) and where, for two points a and b in ℓ^1, $[a,b]$ denotes the line segment from a to b when looking at ℓ^1 as a vector space.

We choose a finite family of marked points $(Q_i)_{i\in[p]}$ of T_l which includes all the leaves of T, so that $T_l = \cup_{i=1}^p [\![\rho, Q_i]\!]$. Let, for $i \in [p]$, $h_i = ht(Q_i)$ and, if we add another integer j, $h_{i,j} = ht(Q_i \wedge Q_j)$. Now, if $j \leqslant i$, let

$$Q_i^j = \max_{1\leqslant k\leqslant j} h_{k,i} - \max_{1\leqslant k\leqslant j-1} h_{k,i}.$$

We then embed Q_i in ℓ^1 as

$$Q_i = (Q_i^1, Q_i^2, \dots, Q_i^n, 0, 0, \dots)$$

The set $\cup_{i=1}^{p}[\![0,Q_i]\!]$ is then indeed a version of T_l. Without going into details, it is a tree because of the way we have built our paths and, by construction, the distances between the $(Q_i)_{i\in[p]}$ and the root are the correct ones.

3.2 GH and GHP convergence

We consider here a sequence $(l_n)_{n\in\mathbb{N}}$ of non-negative functions on the set of edges E, and state in this section results concerning the convergence of the sequence $(T_{l_n})_{n\in\mathbb{N}}$ in the GH or GHP sense.

Lemma 3.1. *Assume that, for all $e\in E$, $l_n(e)$ converges to a non-negative number $l(e)$ as n goes to infinity. We then have*

$$T_{l_n} \underset{n\to\infty}{\overset{\mathrm{GH}}{\longrightarrow}} T_l.$$

In fact, it is sufficient to know that, for all leaves L and L' of T, $d_{l_n}(L,L')$ and $d_{l_n}(\rho,L)$ converge as n goes to infinity.

Moreover, if we embed the trees in ℓ^1 with the stick-breaking method by using the leaves as marked points, we have Hausdorff convergence in ℓ^1.

Proof. We just need to prove the Hausdorff convergence in ℓ^1 of the embedded versions of the trees. For this, one only needs to notice that

$$d_{\ell^1,\mathrm{H}}(T_{l_n},T_l)\leqslant\sum_{e\in E}|l_n(e)-l(e)|,$$

which converges to 0.

The proof of the second point is merely a matter of noticing that, if we know the distances between the leaves (including the root), we can recover the complete metric structure by linear operations. □

Our next result mixes Proposition 2.1 and Lemma 3.1 to obtain a sufficient condition for Gromov-Hausdorff-Prokhorov convergence. We keep the assumptions of Lemma 3.1, and also assume that all the trees are embedded in ℓ^1, the marked points being the leaves of T.

Lemma 3.2. *For $n\in\mathbb{N}$, let μ_n be a probability measure on T_{l_n} with m_n the corresponding non-increasing function. Let S be any dense subset of T_l, and assume that, for all $x\in S$, there exists a sequence $(x_n)_{n\in\mathbb{N}}$, such that*

- *$x_n\in T_{l_n}$ for all n, x_n converges to x as n goes to infinity,*
- *$(T_{l_n})_{x_n}$ converges to $(T_l)_x$ in the Hausdorff sense,*
- *$m_n(x_n)$ converges to a number we call $f(x)$.*

We then have

$$(T_{l_n},\mu_n) \underset{n\to\infty}{\overset{\mathrm{GHP}}{\longrightarrow}} (T_l,\mu),$$

where μ is the unique probability measure on T_l such that, for all $x\in T_l$, $\mu((T_l)_x)=f(x^-)$, and $f(x^-)$ is defined as

$$f(x^-)=\lim_{\substack{y\to x\\ y\in S\cap[\![\rho,x[\![}} f(y), \tag{1.1}$$

and $f(\rho^-)=1$. More precisely, since we consider the versions of the trees embedded in ℓ^1, we have Hausdorff convergence of the sets and Prokhorov convergence of the measures.

Proof. Since T_l is compact, $(\cup T_{l_n}) \cup T_l$ also is and Prokhorov's theorem ensures us that a subsequence of $(\mu_n)_{n\in\mathbb{N}}$ converges weakly. Without loss of generality, we can therefore assume that $(\mu_n)_{n\in\mathbb{N}}$ converges to a measure μ on T_l, and we will show that μ must be as explicited in the statement of the lemma. This will be done by showing the following double inequality for all $x \in S$, which is inspired by the Portmanteau theorem,

$$\mu\big((T_l)_x \setminus \{x\}\big) \leqslant f(x) \leqslant \mu\big((T_l)_x\big). \tag{1.2}$$

We start by showing the right part of (1.2): $f(x) \leqslant \mu\big((T_l)_x\big)$. Let $\varepsilon > 0$, by Hausdorff convergence in ℓ^1, for n large enough, we have $(T_{l_n})_{x_n} \subset ((T_l)_x)^\varepsilon$ (we recall that A^ε is the closed ε-enlargement of the subset A). Since we also have $d_{\ell^1,\mathrm{P}}(\mu_n, \mu) \leqslant \varepsilon$ for n large enough, we obtain

$$\mu_n\big((T_{l_n})_{x_n}\big) \leqslant \mu_n\Big(\big((T_l)_x\big)^\varepsilon\Big) \leqslant \mu\Big(\big((T_l)_x\big)^{2\varepsilon}\Big) + \varepsilon,$$

and making n tend to infinity then gives us

$$f(x) \leqslant \mu\Big(\big((T_l)_x\big)^{2\varepsilon}\Big) + \varepsilon.$$

Letting ε tend to 0 and using the fact that $(T_l)_x$ is closed gives us $f(x) \leqslant \mu\big((T_l)_x\big)$.

A similar, slightly more involved argument will show that $\mu\big((T_l)_x \setminus \{x\}\big) \leqslant f(x)$ for $x \in S$. Let $x \in S$ and let $d+1$ be its degree (there is nothing to say if x is a leaf or the root). Let $T^1, \ldots, T^d$ be the tree components of $(T_l)_x \setminus \{x\}$ and let $y^1, \ldots, y^d$ be any points of $T^1, \ldots, T^d$ which also are in S. We give ourselves the corresponding sequences $(y^1_n)_{n\in\mathbb{N}}, \ldots, (y^d_n)_{n\in\mathbb{N}}$. Take $\varepsilon > 0$, we have, for n large enough,

$$\cup_{i=1}^d \big((T_l)_{y^i}\big) \subset \cup_{i=1}^d \big((T_{l_n})_{y^i_n}\big)^\varepsilon,$$

and therefore, using the Prokhorov convergence of measures, for possibly larger n,

$$\begin{aligned}\mu\Big(\cup_{i=1}^d \big((T_l)_{y^i}\big)\Big) &\leqslant \mu\Big(\cup_{i=1}^d \big((T_{l_n})_{y^i_n}\big)^\varepsilon\Big) \\ &\leqslant \mu_n\Big(\cup_{i=1}^d \big((\mathcal{T}_{l_n})_{y^i_n}\big)^{2\varepsilon}\Big) + \varepsilon.\end{aligned}$$

Since μ_n is supported on $\mathcal{T}_n$, if we take $2\varepsilon < \max\limits_{1\leqslant i\leqslant d} d(y^i, x)$, and n large enough, we obtain

$$\mu_n\Big(\cup_{i=1}^d \big((T_{l_n})_{y^i_n}\big)^{2\varepsilon}\Big) \leqslant \mu_n\big((T_{l_n})_{x_n}\big),$$

and thus also have

$$\mu\Big(\cup_{i=1}^d \big((T_l)_{y^i}\big)\Big) \leqslant \mu_n\big((\mathcal{T}_{l_n})_{x_n}\big) + \varepsilon.$$

Letting n tend to infinity and then ε tend to 0, we obtain

$$\mu\Big(\cup_{i=1}^d \big((T_l)_{y^i}\big)\Big) \leqslant f(x),$$

and finally we let all the y^i tend to x, which makes the left-hand side tend to $\mu\big((T_l)_x \setminus \{x\}\big)$.

Having proved (1.2), we only need to check that, calling m the decreasing function associated to μ, m is equal to the left-limit of f as defined in (1.1), which is immediate: let $x \in T_l \setminus \{\rho\}$ and evaluate (1.2) at a point $y \in [[\rho, x[[\cap S$. By left-continuity of m, if we let y tend to x, both the left and right members converge to $m(x)$, while the middle one converges to $f(x^-)$, which ends the proof. □

4 Subtrees and projections

Let $(\mathcal{T}, d, \rho)$ be a compact and rooted $\mathbb{R}$-tree and $\mathcal{T}'$ be a compact and connected subset of $\mathcal{T}$ containing ρ. The boundary $\partial\mathcal{T}'$ of $\mathcal{T}'$ in $\mathcal{T}$ is then finite or countable. We recall that $\mathcal{T}_x$ denotes the subtree of $\mathcal{T}$ rooted at x, $\forall x \in \mathcal{T}$, and similarly $\mathcal{T}'_x$ is the subtree of $\mathcal{T}'$ rooted at x, for $x \in \mathcal{T}'$. We then have

$$\mathcal{T} = \mathcal{T}' \cup \bigcup_{x \in \partial\mathcal{T}'} \mathcal{T}_x$$

with only the elements of $\partial\mathcal{T}'$ being counted multiple times in this union.

For $x \in \mathcal{T}$, there exists a highest ancestor of x which is in $\mathcal{T}'$. We call it $\pi(x)$, and call the map π the projection from $\mathcal{T}$ onto $\mathcal{T}'$. We consider it as a map from $\mathcal{T}$ to $\mathcal{T}$, so that, for any measure μ on $\mathcal{T}$, $\pi_*\mu$ defines a measure on $\mathcal{T}$ (that only charges $\mathcal{T}'$).

Lemma 4.1. *For any probability measure μ on $\mathcal{T}$, $\pi_*\mu$ is the unique probability measure ν on $\mathcal{T}'$ which satisfies*

$$\forall x \in \mathcal{T}', \quad \nu(\mathcal{T}'_x) = \mu(\mathcal{T}_x).$$

Proof. The fact that $\pi_*\mu$ satisfies the relation comes from the fact that, for all $x \in \mathcal{T}'$, we have $\mathcal{T}_x = \pi^{-1}(\mathcal{T}'_x)$. Uniqueness is a consequence of Proposition 2.1. □

Lemma 4.2. *The map π is 1-Lipschitz whether one considers points of $\mathcal{T}$, the Hausdorff distance between compact subsets of $\mathcal{T}$ or the Prokhorov distance between probability measures on $\mathcal{T}$:*

- $\forall x, y \in \mathcal{T}, d\big(\pi(x), \pi(y)\big) \leqslant d(x, y)$,
- *for A and B non-empty compact subsets of $\mathcal{T}$, $d_{\mathcal{T},\mathrm{H}}\big(\pi(A), \pi(B)\big) \leqslant d_{\mathcal{T},\mathrm{H}}(A, B)$,*
- *for any two probability measures μ and ν on $\mathcal{T}$, $d_{\mathcal{T},\mathrm{P}}(\pi_*\mu, \pi_*\nu) \leqslant d_{\mathcal{T},\mathrm{P}}(\mu, \nu)$.*

Proof. Let x and y be elements of $\mathcal{T}$. Assume first that x and y are on the same branch, and by symmetry, we can restrict that to $x \leqslant y$. If both of them are in $\mathcal{T}'$ then $\pi(x) = x$ and $\pi(y) = y$, while if they are both not in $\mathcal{T}'$, then $\pi(x) = \pi(y)$. If x is in $\mathcal{T}'$ but y is not, then $\pi(y) \in [[x, y]]$. In all these three cases, we have $d(\pi(x), \pi(y)) \leqslant d(x, y)$. If x and y are not on the same branch of $\mathcal{T}$, one just needs to consider $z = x \wedge y$, use the fact that $d(x, y) = d(x, z) + d(y, z)$ and use the previous argument twice.

Let A and B be compact subsets of $\mathcal{T}$ and let ε such that $A \subset B^\varepsilon = \{x \in \mathcal{T}, \exists b \in B, d(x, b) \leqslant \varepsilon\}$. Let $x \in \pi(A)$ and $a \in A$ such that $x = \pi(a)$ and then let $b \in B$ such that $d(a, b) \leqslant \varepsilon$. We then have $d(x, \pi(b)) \leqslant \varepsilon$ and thus $\pi(A) \subset \pi(B)^\varepsilon$. Reversing the roles of A and B then shows that $d_{\mathcal{T},\mathrm{H}}(\pi(A), \pi(B)) \leqslant d_{\mathcal{T},\mathrm{H}}(A, B)$.

Let μ and ν be two probability measures on $\mathcal{T}$ and let ε such that $d_P(\mu, \nu) \leqslant \varepsilon$. Let A be a measurable subset of $\mathcal{T}$, we then have $\pi_*\mu(A) = \mu\big(\pi^{-1}(A)\big) \leqslant \nu\big((\pi^{-1}(A))^\varepsilon\big) + \varepsilon$. We also have $(\pi^{-1}(A))^\varepsilon \subset \pi^{-1}(A^\varepsilon)$ and thus $\pi_*\mu(A) \leqslant \pi_*\nu(A^\varepsilon) + \varepsilon$. Reversing the roles of μ and ν yields $d_{\mathcal{T},\mathrm{P}}(\pi_*\mu, \pi_*\nu) \leqslant \varepsilon$. □

Let $Z_\pi = \sup_{x \in \mathcal{T}} d(x, \pi(x))$. This quantity controls all of the difference between $\mathcal{T}$ and $\mathcal{T}'$, even when measured:

Lemma 4.3. *We have*

$$Z_\pi = \sup_{x \in \partial\mathcal{T}'} ht(\mathcal{T}_x),$$

and

$$d_{\mathcal{T},\mathrm{H}}(\mathcal{T},\mathcal{T}') = Z_\pi$$

and, for any measure μ *on* $\mathcal{T}$*,*

$$d_{\mathcal{T},\mathrm{P}}(\mu,\pi_*\mu) \leqslant d_{\mathcal{T},\mathrm{H}}(\mathcal{T},\mathcal{T}').$$

Proof. The first point is a direct consequence from the fact that, if $x \in \mathcal{T}'$ then $\pi(x) = x$, while if $x \in \mathcal{T} \setminus \mathcal{T}'$, $x \in \mathcal{T}_{\pi(x)}$. The second point is also a fairly straightforward consequence of the definition of Z_π. The third point involves simple manipulations of the Prokhorov metric. Let A be a subset of $\mathcal{T}$. Since $A \subset \pi^{-1}(\pi(A))$ and $\pi(A) \subset A^{Z_\pi}$, we automatically have $\mu(A) \leqslant \pi_*\mu(\pi(A)) \leqslant \pi_*\mu(A^{Z_\pi})$. On the other hand, we have $\pi^{-1}(A) \subset A^{Z_\pi}$, which implies $\pi_*\mu(A) \leqslant \mu(A^{Z_\pi})$. □

Chapter 2

General fragmentation trees

Contents

We show that the genealogy of any self-similar fragmentation process can be encoded in a compact measured $\mathbb{R}$-tree. Under some Malthusian hypotheses, we compute the fractal Hausdorff dimension of this tree through the use of a natural measure on the set of its leaves. This generalizes previous work of Haas and Miermont which was restricted to conservative fragmentation processes.

1 Introduction

In this chapter, we study a family of trees derived from self-similar fragmentation processes. Such processes describe the evolution of an object which constantly breaks down into smaller fragments, each one then evolving independently from one another, just as the initial object would, but with a rescaling of time by the size of the fragment to a certain power called the index of self-similarity. This breaking down happens in two ways: erosion, a process by which part of the object is continuously being shaved off and thrown away, and actual splittings of fragments which are governed by a Poisson point process. Erosion is parametered by a nonnegative number c called the erosion rate, while the splitting Poisson point process depends on a dislocation measure ν on the space

$$\mathcal{S}^{\downarrow} = \{\mathbf{s} = (s_i)_{i\in\mathbb{N}} : s_1 \geqslant s_2 \geqslant \ldots \geqslant 0, \sum s_i \leqslant 1\}.$$

Precise definitions can be found in Section 2.

Our main inspiration is the 2004 article of Bénédicte Haas and Grégory Miermont [41]. Their work focused on *conservative* fragmentations, where there is no erosion and splittings of fragments do not change the total mass. They have shown that, when the index of self-similarity is negative, the genealogy of a conservative fragmentation process can be encoded in a continuum random tree, the genealogy tree of the fragmentation, which is compact and naturally equipped with a probability measure on the set of its leaves. Our main goal here will be to generalize the results they have obtained to the largest reasonable class of fragmentation processes: the conservation hypothesis will be discarded, though the index of self-similarity will be kept negative. We will show (Theorem 3.3) that we can still define some kind of fragmentation tree, but its natural measure will not be supported by the leaves, and we thus step out of the classical continuum random tree context set in [5].

That the measure of a general fragmentation tree gives mass to its skeleton will be a major issue for us here, and its study will therefore involve creating a new measure on the leaves of the tree. To do this we will restrict ourselves to *Malthusian* fragmentations. Informally, for a fragmentation process to be Malthusian means that there is a number $p^* \in (0,1]$ such that, infinitesimally, calling $(X_i(t))_{i\in\mathbb{N}}$ the sizes of the fragments of the process at time t, the expectation of $\sum_{i\in\mathbb{N}} X_i(t)^{p^*}$ is constant. Such conservation properties will let us define and study a family of martingales related to the tree and use them to define a Malthusian measure μ^* on the leaves of the tree. The use of this measure then lets us obtain the fractal Hausdorff dimension of the set of leaves of the fragmentation tree, under a light regularity condition, called "assumption (**H**)", which is a reinforcement of the Malthusian hypothesis:

> The function ψ defined on $\mathbb{R}$ by $\psi(p) = cp + \int_{\mathcal{S}^{\downarrow}} (1 - \sum_i s_i^p)\nu(\mathrm{d}\mathbf{s}) \in [-\infty, +\infty)$ takes at least one finite strictly negative value on the interval $[0,1]$.

Theorem 1.1. *Assume* (**H**) *and that $\alpha < 0$. Then, almost surely, if the set of leaves of the fragmentation tree derived from an α-self-similar fragmentation process with erosion rate c and dislocation measure ν is not countable, its Hausdorff dimension is equal to $\frac{p^*}{|\alpha|}$.*

In [41], a dimension of $\frac{1}{|\alpha|}$ was found for conservative fragmentation trees, also under a regularity condition. We can see that non-conservation of mass makes the tree smaller in the sense of dimension. Note as well that the event where the leaves of the tree are countable only has positive probability if $\nu(0,0,\ldots,0) > 0$, that is, if a fragment can suddenly disappear without giving any offspring.

The chapter is organized as follows: in Section 2 is presented the necessary background on fragmentation processes. In Section 3 we construct the tree associated to a fragmentation process with elementary methods, and give a few of its basic topological properties. The next three sections form the proof of Theorem 1.1: we build in Section 4 the random measure μ^* by combining martingale methods and Proposition 2.1, we then give in Section 5 an interpretation of this measure as a biased version of the distribution of the fragmentation tree, and in Section 6 we properly compute the Hausdorff dimension of the tree, using the results of Sections 4 and 5. Finally, Section 7 is dedicated various comments and applications, namely the effects of varying the parameters and the fact one can interpret continuous time Galton-Watson trees as fragmentation trees, giving us the Hausdorff dimension of their boundary.

Note: we use the convention that, when we take 0 to a nonpositive power, the result is 0. We therefore abuse notation slightly by omitting an indicator function such as $\mathbb{1}_{x\neq 0}$ most of the time. In particular, sums such as $\sum_{i\in\mathbb{N}} x_i^p$ are implicitly taken on the set of i such that $x_i \neq 0$ even when $p \leqslant 0$.

2 Background, preliminaries and some notation

2.1 Self-similar fragmentation processes

2.1.1 Partitions

We are going to look at two different kinds of partitions. The first ones are *mass partitions.* These are nonincreasing sequences $\mathbf{s} = (s_1, s_2, \ldots)$ with $s_i \geqslant 0$ for every i and such that $\sum_i s_i \leqslant 1$. These are to be considered as if a particle of mass 1 had split up into smaller particles, some of its mass having turned into *dust* which is represented by $s_0 = 1 - \sum_i s_i$. We call $\mathcal{S}^{\downarrow}$ the set of mass partitions, it can be metrized with the restriction of the uniform norm and is then compact.

The more important partitions we will consider here are the set-theoretic partitions of finite and countable sets. For such a set S, we let $\mathcal{P}_S$ be the set of partitions of S. The main examples are of course the cases of partitions of $\mathbb{N} = \{1, 2, 3, \ldots\}$ (for countable sets) and, for $n \in \mathbb{N}$, $[n] = \{1, 2, \ldots, n\}$. Let us focus here on $\mathcal{P}_{\mathbb{N}}$. A partition $\pi \in \mathcal{P}_{\mathbb{N}}$ will be written as a countable sequence of subsets of $\mathbb{N}$, called the blocks of the partition: $\pi = (\pi_1, \pi_2, \ldots)$ where every intersection between two different blocks is empty and the union of all the blocks is $\mathbb{N}$. The blocks are ordered by increasing smallest element: π_1 is the block containing 1, π_2 is the block containing the smallest integer not in π_1, and so on. If π has finitely many blocks, we complete the sequence with an infinite repeat of the empty set. (When not referring to a specific partition, the word "block" simply means "subset of $\mathbb{N}$".)

A partition can also be interpreted as an equivalence relation on $\mathbb{N}$: for a partition π and two integers i and j, we will write $i \sim_\pi j$ if i and j are in the same block of π. We will also call $\pi_{(i)}$ the block of π containing i.

We now have two ways to identify the blocks of a partition π: either with their rank in the partition's order or with their smallest element. Most of the time one will be more useful than the other, but sometimes we will want to mix both, which is why we will call $\mathrm{rep}(\pi)$ the set of smallest elements of blocks of π.

Let B be a block. For all $\pi \in \mathcal{P}_{\mathbb{N}}$, we let $\pi \cap B$ be the restriction of π to B, i.e. the partition of B whose blocks are, up to reordering, the $(\pi_i \cap B)_{i\in\mathbb{N}}$.

We say that a partition π is *finer* than another partition π' if every block of π is a subset of a block of π'. This defines a partial order on the set of partitions.

Intersection and union operators can be defined on partitions: let X be a set and, for $x \in X$, π_x be a partition. Then we let $\bigcap_{x\in X} \pi_x$ be the unique partition $\tilde{\pi}$ to satisfy $\forall i, j \in \mathbb{N}, i \sim_{\tilde{\pi}} j \Leftrightarrow \forall x \in X, i \sim_{\pi_x} j$. The blocks of $\bigcap_{x\in X} \pi_x$ are the intersections of blocks of the $(\pi_x)_{x\in X}$. Similarly, assuming that all the $(\pi_x)_{x\in X}$ are comparable, then we define $\bigcup_{x\in X} \pi_x$ to be the unique partition $\tilde{\pi}$ such that, $\forall i, j \in \mathbb{N}, i \sim_{\tilde{\pi}} j \Leftrightarrow \exists x \in X, i \sim_{\pi_x} j$.

We endow $\mathcal{P}_{\mathbb{N}}$ with a metric: for two partitions π and π', let $n(\pi, \pi')$ be the largest integer n such that $\pi \cap [n]$ and $\pi' \cap [n]$ are equal ($n(\pi, \pi') = \infty$ if $\pi = \pi'$) and let $d(\pi, \pi') = 2^{-n(\pi,\pi')}$. This defines a distance function on $\mathcal{P}_{\mathbb{N}}$, which in fact satisfies the ultra-metric triangle inequality. This metric provides a topology on $\mathcal{P}_{\mathbb{N}}$, for which convergence is simply characterized: a sequence $(\pi_n)_{n\in\mathbb{N}}$ of partitions converges to a partition π if, and only if, for every k, there exists n_k such that $\pi_n \cap [k] = \pi \cap [k]$ for n larger than n_k. The metric also provides $\mathcal{P}_{\mathbb{N}}$ with a Borel σ-field, which is easily checked to be the σ-field generated by the restriction maps, i.e. the functions which which map π to $\pi \cap [n]$ for all integers n.

Let S and S' be two sets with a bijection $f : S \to S'$. Then we can easily transform partitions of S' into partitions of S: let π be a partition of S', we let $f\pi$ be the partition defined by: $\forall i, j \in S, i \sim_{f\pi} j \Leftrightarrow f(i) \sim_{\pi} f(j)$. This can be used to generalize the metric d to $\mathcal{P}_S$ for infinite S (note that the notion of convergence does not depend on the chosen bijection), and then $\pi \mapsto f\pi$ is easily seen to be continuous.

Special attention is given to the case where f is a permutation: we call *permutation* of $\mathbb{N}$ any bijection σ of $\mathbb{N}$ onto itself. A $\mathcal{P}_{\mathbb{N}}$-valued random variable (or random partition) Π is said to be *exchangeable* if, for all permutations σ, $\sigma\Pi$ has the same law as Π.

Let B be a block. If the limit $\lim_{n\to\infty} \frac{1}{n}\#(B \cap [n])$ exists then we write it $|B|$ and call it the *asymptotic frequency* or more simply *mass* of B. If all the blocks of a partition π have asymptotic frequencies, then we call $|\pi|^{\downarrow}$ their sequence in decreasing order, which is an element of $\mathcal{S}^{\downarrow}$. This defines a measurable, but not continuous, map.

A well-known theorem of Kingman from the paper [52] links exchangeable random partitions of $\mathbb{N}$ and random mass partitions through the "paintbox construction". More precisely: let $\mathbf{s} \in \mathcal{S}^{\downarrow}$, and $(U_i)_{i\in\mathbb{N}}$ be independent uniform variables on $[0, 1]$, we define a random partition $\Pi_{\mathbf{s}}$ by

$$\forall i \neq j,\ i \sim_{\Pi_{\mathbf{s}}} j \Leftrightarrow \exists k,\ U_i, U_j \in [\sum_{p=1}^{k} s_p, \sum_{p=1}^{k+1} s_p).$$

This random partition is exchangeable, all its blocks have asymptotic frequencies, and $|\Pi_{\mathbf{s}}|^{\downarrow} = \mathbf{s}$. By calling $\kappa_{\mathbf{s}}$ the law of $\Pi_{\mathbf{s}}$, Kingman's theorem states that, for any exchangeable random partition Π, there exists a random mass partition S such that, conditionally on S, Π has law κ_S. A useful consequence of this theorem is found in [11], Corollary 2.4: for any integer k, conditionally on the variable S, the asymptotic frequency $|\Pi_{(k)}|$ of the block containing k exists almost surely and is a size-biased pick amongst the terms of S, which means that its distribution is $\sum_i S_i \delta_{S_i} + S_0 \delta_{S_0}$ (with $S_0 = 1 - \sum_{i\in\mathbb{N}} S_i$).

Let Π and Ψ be two independent exchangeable random partitions. Then, for any i and j, the block $\Pi_i \cap \Psi_j$ of $\Pi \cap \Psi$ almost surely has asymptotic frequency $|\Pi_i||\Psi_j|$. This stays true if we take countably many partitions, as is stated in [11], Corollary 2.5.

2.1.2 Definition of fragmentation processes

Partition-valued fragmentation processes were first introduced in [9] (homogeneous processes only) and [10] (the general self-similar kind).

Definition 2.1. *Let* $\Pi = \big(\Pi(t)\big)_{t\geqslant 0}$ *be a* $\mathcal{P}_{\mathbb{N}}$*-valued process with càdlàg paths, which satisfies* $\Pi(0) = (\mathbb{N}, \emptyset, \emptyset, \ldots)$*, which is exchangeable as a process (i.e. for all permutations* σ*, the process* $\big(\sigma\Pi(t)\big)_{t\geqslant 0}$ *has the same law as* Π *and such that, almost surely, for all* $t \geqslant 0$*, all the blocks of* $\Pi(t)$ *have asymptotic frequencies. Let* α *be any real number. We say that* Π *is a* self-similar fragmentation process with index α *if it also satisfies the following self-similar fragmentation property: for all* $t \geqslant 0$*, given* $\Pi(t) = \pi$*, the processes* $\big(\Pi(t+s) \cap \pi_i\big)_{s\geqslant 0}$ *(for all integers* i*) are mutually independent, and each one has the same distribution as* $\big(\Pi(|\pi_i|^\alpha(s)) \cap \pi_i\big)_{s\geqslant 0}$.

When $\alpha = 0$, we will say that Π is a *homogeneous* fragmentation process instead of 0-self-similar fragmentation process.

Remark 2.2. *One can give a Markov process structure to an* α*-self-similar fragmentation process* Π *by defining, for any partition* π*, the law of* Π *starting from* π*. Let* $(\Pi^i)_{i\in\mathbb{N}}$ *be independent copies of* Π *(each one starting at* $(\mathbb{N}, \emptyset, \ldots)$ *), then we let, for all* $t \geqslant 0$*,* $\Pi(t)$ *be the partition whose blocks are exactly those of* $\big(\Pi^i(|\pi_i|^\alpha t) \cap \pi_i\big)_{i\in\mathbb{N}}$*. In this case the process isn't exchangeable with respect to all permutations of* $\mathbb{N}$*, but only with respect to permutations which stabilize the blocks of the initial value* π*.*

We see fragmentation processes as random variables in the space $\mathcal{D} = \mathcal{D}([0, +\infty), \mathcal{P}_{\mathbb{N}})$, which is the set of càdlàg functions from $[0, +\infty)$ to $\mathcal{P}_{\mathbb{N}}$. An element of $\mathcal{D}$ will typically be written as $(\pi_t)_{t\geqslant 0}$. This space can be metrized with the *Skorokhod metric* and is then Polish. More importantly, the Borel σ-algebra on $\mathcal{D}$ is then the σ-algebra spanned by the *evaluation functions* $(\pi_t)_{t\geqslant 0} \mapsto \pi_s$ (for $s \geqslant 0$), implying that the law of a process is characterized by its finite-dimensional marginal distributions. The definition of the Skorokhod metric and generalities on the subject can be read in [48], Section VI.1.

Let us give a lemma which makes self-similarity easier to handle at times:

Lemma 2.3. *Let* $(\Pi(t))_{t\geqslant 0}$ *be any exchangeable* $\mathcal{P}_{\mathbb{N}}$*-valued process, and* A *any infinite block. Take any bijection* f *from* A *to* $\mathbb{N}$*, then the two* $\mathcal{P}_A$*-valued processes* $(\Pi(t) \cap A)_{t\geqslant 0}$ *and* $(f\Pi(t))_{t\geqslant 0}$ *have the same law.*

Proof. For all $n \in \mathbb{N}$, let $A_n = \{f^{-1}(1), f^{-1}(2), \ldots, f^{-1}(n)\}$. Recall then that, with the σ-algebra which we have on $\mathcal{P}_A$, we only need to check that, for all $n \in \mathbb{N}$, $(\Pi_{|A_n})$ has the same law as $f(\Pi \cap [n])$. If G is a nonnegative measurable function on $\mathcal{D}([0, +\infty), \mathcal{P}_{A_n})$, we have, by using the fact that the restriction of f from $[n]$ to A_n can be extended to a bijection of $\mathbb{N}$ onto itself

$$\mathbb{E}[G(\Pi \cap A_n)] = \mathbb{E}\Big[G\big((f\Pi) \cap A_n\big)\Big] = \mathbb{E}\Big[G(f(\Pi \cap [n]))\Big],$$

which is all we need. □

This lemma will make it easier to show the fragmentation property for some $\mathcal{D}$-valued processes we will build throughout the chapter.

2.1.3 Characterization and Poissonian construction

A famous result of Bertoin (detailed in [11], Chapter 3) states that the law of a self-similar fragmentation process is characterized by three parameters: the index of self-similarity α, an *erosion coefficient* $c \geqslant 0$ and a *dislocation measure* ν, which is a σ-finite measure on $\mathcal{S}^{\downarrow}$ such that

$$\nu(1,0,0,\ldots) = 0 \text{ and } \int_{\mathcal{S}^{\downarrow}} (1-s_1)\nu(\mathrm{d}\mathbf{s}) < \infty.$$

Bertoin's result can be formulated this way: for any fragmentation process, there exists a unique triple (α, c, ν) such that our process has the same distribution as the process which we are about to explicitly construct.

First let us describe how to build a fragmentation process with parameters $(0,0,\nu)$ which we will call $\Pi^{0,0}$. Let $\kappa_\nu(\mathrm{d}\pi) = \int_{\mathcal{S}^{\downarrow}} \kappa_{\mathbf{s}}(\mathrm{d}\pi)\nu(\mathrm{d}\mathbf{s})$ where $\kappa_s(\mathrm{d}\pi)$ denotes the paintbox measure on $\mathcal{P}_{\mathbb{N}}$ corresponding to $\mathbf{s} \in \mathcal{S}^{\downarrow}$. For every integer k, let $(\Delta_t^k)_{t\geqslant 0}$ be a Poisson point process with intensity κ_ν, such that these processes are all independent. Now let $\Pi^{0,0}(t)$ be the process defined by $\Pi^{0,0}(0) = (\mathbb{N}, \emptyset, \emptyset, \ldots)$ and which jumps when there is an atom (Δ_t^k): we replace the k-th block of $\Pi^{0,0}(t-)$ by its intersection with Δ_t^k. This might not seem well-defined since the Poisson point process can have infinitely many atoms. However, one can check (as we will do in Section 5.2 in a slightly different case) that this is well defined by restricting to the first N integers and taking the limit when N goes to infinity.

To get a $(0,c,\nu)$-fragmentation which we will call $\Pi^{0,c}$, take a sequence $(T_i)_{i\in\mathbb{N}}$ of exponential variables with parameter c which are independent from each other and independent from $\Pi^{0,0}$. Then, for all t, let $\Pi^{0,c}(t)$ be the same partition as $\Pi^{0,0}(t)$ except that we force all integers i such that $T_i \leqslant t$ to be in a singleton if they were not already.

Finally, an (α, c, ν)-fragmentation can then be obtained by applying a Lamperti-type time-change to all the blocks of $\Pi^{0,c}$: let, for all i and t,

$$\tau_i(t) = \inf\Big\{u, \int_0^u |\Pi_{(i)}^{0,c}(r)|^{-\alpha}\mathrm{d}r > t\Big\}.$$

Then, for all t, let $\Pi^{\alpha,c}(t)$ be the partition such that two integers i and j are in the same block of $\Pi^{\alpha,c}(t)$ if and only if $j \in \Pi_{(i)}^{0,c}(\tau_i(t))$. Note that if $t \geqslant \int_0^\infty |\Pi_{(i)}^{0,c}(r)|^{-\alpha}\mathrm{d}r$, then the value of $\tau_i(t)$ is infinite, and i is in a singleton of $\Pi^{\alpha,c}(t)$. Note also that the time transformation is easily inverted: for $s \in [0,\infty)$, we have

$$\tau_i^{-1}(s) = \inf\Big\{u, \int_0^u |\Pi_{(i)}^{\alpha,c}(r)|^{+\alpha}\mathrm{d}r > s\Big\}.$$

This time-change can in fact be done for any element π of $\mathcal{D}$: since, for all $i \in \mathbb{N}$ and $t \geqslant 0$, $\tau_i(t)$ is a measurable function of $\Pi^{0,c}$, there exists a measurable function G^α from $\mathcal{D}$ to $\mathcal{D}$ which maps $\Pi^{0,c}$ to $\Pi^{\alpha,c}$.

Let us once and for all fix our notations for the processes: in the whole chapter, c and ν will be fixed (with $c \neq 0$ or ν not of the form $a\delta_{(1,0,0,\ldots)}$, $a \geqslant 0$, to remove the trivial case of constant fragmentations), however we will often jump between a homogeneous $(0,c,\nu)$-fragmentation and the associated self-similar (α,c,ν)-fragmentation, with $\alpha < 0$ fixed. This is why we will rename things and let $\Pi = \Pi^{0,c}$ as well as $\Pi^\alpha = \Pi^{\alpha,c}$. We then let $(\mathcal{F}_t)_{t\geqslant 0}$ be the canonical filtration associated to Π and $(\mathcal{G}_t)_{t\geqslant 0}$ the one associated to Π^α.

2.1.4 A few key results

One simple but important consequence of the Poissonian construction is that the notation $|\Pi_{(i)}^\alpha(t^-)|$ is well-defined for all i and t: it is equal to both the limit, as s increases to t, of $|\Pi_{(i)}^\alpha(s)|$, and the asymptotic frequency of the block of $\Pi^\alpha(t^-)$ containing i.

For every integer i, let $\mathcal{G}_i$ be the canonical filtration of the process $\big(\Pi^{\alpha}_{(i)}(t)\big)_{t\geqslant 0}$, and consider a family of random times $(L_i)_{i\in\mathbb{N}}$ such that L_i is a $\mathcal{G}_i$-stopping time for all i. We say that $(L_i)_{i\in\mathbb{N}}$ is a *stopping line* if, for all integers i and j, $j\in\Pi^{\alpha}_{(i)}(L_i)$ implies $L_i=L_j$. Under this condition, Π^{α} then satisfies an extended fragmentation property (proved in [11], Lemma 3.14): we can define for every t a partition $\Pi^{\alpha}(L+t)$ whose blocks are the $(\Pi^{\alpha}_{(i)}(L_i+t))_{i\in\mathbb{N}}$. Then conditionally on the sigma-field $\mathcal{G}_L$ generated by the $\mathcal{G}_i(L_i)$ $(i\in\mathbb{N})$, the process $(\Pi^{\alpha}(L+t))_{t\geqslant 0}$ has the same law as a version ofΠ starting at $\Pi^{\alpha}(L)$.

One of the main tools of the study of fragmentation processes is the *tagged fragment*: we specifically look at the block of Π^{α} containing the integer 1 (or any other fixed integer). Of particular interest, its mass can be written in terms of Lévy processes: one can write, for all t, $|\Pi^{\alpha}_{(1)}(t)|=\mathrm{e}^{-\xi_{\tau(t)}}$ where ξ is a killed subordinator with Laplace exponent ϕ defined for nonnegative q by

$$\phi(q)=c(q+1)+\int_{\mathcal{S}^{\downarrow}}(1-\sum_{n=1}^{\infty}s_n^{q+1})\nu(\mathrm{d}\mathbf{s}),$$

and $\tau(t)$ is defined for all t by $\tau(t)=\inf\Big\{u,\int_0^u \mathrm{e}^{\alpha\xi_r}\mathrm{d}r>t\Big\}$. Note that standard results on Poisson measures then imply that, if $q\in\mathbb{R}$ is such that $\int_{\mathcal{S}^{\downarrow}}(1-\sum_{n=1}^{\infty}s_n^{q+1})\nu(\mathrm{d}\mathbf{s})>-\infty$, then we still have $\mathbb{E}[e^{-q\xi_t}\mathbb{1}_{\{\xi_t<\infty\}}]=e^{-t\phi(q)}$.

In particular, the first time t such that the singleton $\{1\}$ is a block of $\Pi^{\alpha}(t)$ is equal to $\int_0^{\infty}\mathrm{e}^{\alpha\xi_s}\mathrm{d}s$, the exponential functional of the Lévy process $\alpha\xi$, which has been studied for example in [20]. In particular it is finite a.s. whenever α is strictly negative and Π is not constant.

3 The fragmentation tree

3.1 Main result

We are going to show a bijective correspondence between the laws of fragmentation processes with negative index and a certain class of random trees. We fix from now on an index $\alpha<0$. If $(\mathcal{T},\mu)$ is a measured tree and S is a measurable subset of $\mathcal{T}$ with $\mu(S)>0$, we let μ_S be the measure μ conditioned on S, which is a probability measure on S.

Definition 3.1. *Let $(\mathcal{T},\mu)$ be a random variable in $\mathbb{T}_W$. For all $t\geqslant 0$, let $\mathcal{T}_1(t),\mathcal{T}_2(t),\dots$ be the connected components of $\mathcal{T}_{>t}$, and let, for all i, $x_i(t)$ be the point of $\mathcal{T}$ with height t which makes $\mathcal{T}_i(t)\cup\{x_i(t)\}$ connected. We say that $\mathcal{T}$ is* self-similar *with index α if $\mu(\mathcal{T}_i(t))>0$ for all choices of $t\geqslant 0$ and i and if, for any $t\geqslant 0$, conditionally on $\big(\mu\big(\mathcal{T}_i(s)\big),i\in\mathbb{N},s\leqslant t\big)$, the trees $\big(\mathcal{T}_i(t)\cup\{x_i(t)\},\mu_{\mathcal{T}_i(t)}\big)_{i\in\mathbb{N}}$ are independent and, for any i, $(\mathcal{T}_i(t)\cup\{x_i(t)\},\mu_{\mathcal{T}_i(t)})$ has the same law as $(\mu(\mathcal{T}_i(t))^{-\alpha}\mathcal{T}',\mu')$ where $(\mathcal{T}',\mu')$ is an independent copy of $(\mathcal{T},\mu)$.*

The similarity with the definition of an α-self-similar fragmentation process must be pointed out: in both definitions, the main point is that each "component" of the process after a certain time is independent of all the others and has the same law as the initial process, up to rescaling. In fact, the following is an straightforward consequence of our definitions:

Proposition 3.2. *Assume that $(\mathcal{T},\mu)$ is a self-similar tree with index of similarity α. Let $(P_i)_{i\in\mathbb{N}}$ be an exchangeable sequence of variables directed by μ (i.e. conditionally on μ, they are independent and all have distribution μ). Define for every $t\geqslant 0$ a partition $\Pi_{\mathcal{T}}(t)$ by saying that i and j are in the same block of $\Pi_{\mathcal{T}}(t)$ if and only if P_i and P_j are in the same connected component of $\mathcal{T}_{>t}$ (in particular an integer i is in a singleton if $ht(P_i)\leqslant t$). Then $\Pi_{\mathcal{T}}$ is an α-self-similar fragmentation process.*

Proof. First of all, we need to check that, for all $t \geqslant 0$, $\Pi_{\mathcal{T}}(t)$ is a random variable. We therefore fix $t > 0$ and notice that the definition of $\Pi_{\mathcal{T}}(t)$ entails that, for all $i \in \mathbb{N}$ and $j \in \mathbb{N}$,

$$i \sim_{\Pi_{\mathcal{T}}(t)} j \Leftrightarrow ht(P_i \wedge P_j) > t,$$

which is a measurable event. Thus, for all integers n and all partitions π of $[n]$, the event $\{\Pi_{\mathcal{T}}(t) \cap [n] = \pi\}$ is also measurable. It then follows that $\Pi_{\mathcal{T}}(t) \cap [n]$ is measurable for all $n \in \mathbb{N}$, and therefore $\Pi_{\mathcal{T}}(t)$ itself is measurable.

Next we need to check that $\Pi_{\mathcal{T}}$ is càdlàg. It is immediate from the definition that $\Pi_{\mathcal{T}}$ is decreasing (in the sense that $\Pi_{\mathcal{T}}(s)$ is finer than $\Pi_T(t)$ for $s > t$), and then that, for any t, $\Pi_{\mathcal{T}}(t) = \underset{s>t}{\cup} \Pi_{\mathcal{T}}(s)$, and thus the process is right-continuous. Similarly, the process has a left-limit at t for all t, which is indentified as $\Pi_{\mathcal{T}}(t^-) = \underset{s<t}{\cap} \Pi_{\mathcal{T}}(s)$.

Exchangeability as a process of $\Pi_{\mathcal{T}}$ is an immediate consequence of the exchangeability of the sequence $(P_i)_{i\in\mathbb{N}}$.

The fact that, almost surely, all the blocks of $\Pi_{\mathcal{T}}(t)$ for $t \geqslant 0$ have asymptotic frequencies is a consequence of the Glivenko-Cantelli theorem (see [28], Theorem 11.4.2). For $i \geqslant 2$, let $Y_i = ht(P_1 \wedge P_i)$, then, for $t < Y_i$, 1 and i are in the same block of $\Pi_{\mathcal{T}}(t)$, and for $t \geqslant Y_i$, they are not. Then we have, for all $t \geqslant 0$,

$$\#(\Pi_{\mathcal{T}}(t) \cap [n])_{(1)} = 1 + \sum_{i=2}^{n} \mathbb{1}_{Y_i > t}.$$

It then follows from the Glivenko-Cantelli theorem (applied conditionally on $\mathcal{T}$, μ and P_1) that, with probability one, for all $t \geqslant 0$, $\frac{1}{n}\#(\Pi_{\mathcal{T}}(t) \cap [n])_{(1)}$ converges as n goes to infinity, the limit being the μ-mass of the tree component of $\mathcal{T}_{>t}$ containing P_1 (or 0 if $ht(P_1) < t$). By replacing 1 with any integer i, we get the almost sure existence of the asymptotic frequencies of $\Pi_{\mathcal{T}}$ at all times.

Let us now check that $\Pi_{\mathcal{T}}(0) = (\mathbb{N}, \emptyset, \ldots)$ almost surely, which amounts to saying that $\mathcal{T} \setminus \{\rho\}$ is connected. Apply the self-similar fragmentation property at time 0: the tree $\mathcal{T}_1(0) \cup \{\rho\}$ (as in Definition 3.1) has the same law as $\mathcal{T}$ up to a random multiplicative constant, and $\mathcal{T}_1$ is almost surely connected by definition. Thus $\mathcal{T} \setminus \{\rho\}$ is almost surely connected. A similar argument also shows that $\mu(\{\rho\})$ is almost surely equal to zero.

Finally, we need to check the α-self-similar fragmentation property for $\Pi_{\mathcal{T}}$. Let $t \geqslant 0$ and $\pi = \Pi_{\mathcal{T}}(t)$. For every integer k, we let $i(k)$ be the unique integer such that $k \in \pi_{i(k)}$ and, for every i, we let $\mathcal{T}_i(t)$ be the tree component of $\mathcal{T}_{>t}$ containing the points P_k with $k \in \mathbb{N}$ such that $i(k) = i$ (if π_i is a singleton, then $\mathcal{T}_i(t)$ is the empty set). We also add the natural rooting point x_i of $\mathcal{T}_i(t)$. Since, for all k, $i(k)$ is measurable knowing $\Pi_{\mathcal{T}}(t)$, we get that, conditionally on $(\mathcal{T}, \mu)$ and $\Pi_{\mathcal{T}}(t)$, P_k is distributed according to $\mu_{\mathcal{T}_{i(k)}}$. From the independence property in Definition 3.1 then follows that the $(\Pi_{\mathcal{T}}(t + .) \cap \pi_i)_{i\in\mathbb{N}}$ are independent. We now just need to identify their law. If $i \in \mathbb{N}$ is such that π_i is a singleton then there is nothing to do. Otherwise π_i is infinite: let f be any bijection $\mathbb{N} \to \pi_i$, and rename the points P_k with k such that $i(k) = i$ by letting $P'_k = P_{f(k)}$. By the self-similarity of the tree, the partition-valued process built from $\mathcal{T}_i \cup \{x_i\}$ and the P'_j (with $j \in \mathbb{N}$) has the same law as $\Pi_{\mathcal{T}}(|\pi_i|^{-\alpha} s)_{s\geqslant 0}$, and therefore $\Pi_{\mathcal{T}}(t + .) \cap \pi_i$ has the same law as $\big(f\Pi^i(|\pi_i|^{\alpha} s)\big)_{s\geqslant 0}$, which is what we wanted.

□

Our main result is a kind of converse of this proposition, in law.

Theorem 3.3. *Let Π^α be a non-constant fragmentation process with index of similarity $\alpha < 0$. Then there exists a random α-self-similar tree $(\mathcal{T}_{\Pi^\alpha}, \mu_{\Pi^\alpha})$ such that $\Pi_{\mathcal{T}_{\Pi^\alpha}}$ has the same law as Π^α.*

Remark 3.4. *This is analogous to a recent result obtained by Chris Haulk and Jim Pitman in [46], which concerns exchangeable hierarchies. An exchangeable hierarchy can be seen as a fragmentation of* $\mathbb{N}$ *where one has forgotten time. Haulk and Pitman show that, just as with self-similar fragmentations, in law, every exchangeable hierarchy can be sampled from a random measured tree.*

The rest of this section is dedicated to the proof of Theorem 3.3. We fix from now on a fragmentation process Π^α (defined on a certain probability space Ω) and will build the tree $\mathcal{T}$ and the measure μ (now omitting the index Π^α).

3.2 The genealogy tree of a fragmentation

We are here going to give an explicit description of $\mathcal{T}$ which has the caveat of not showing that $\mathcal{T}$ is a random variable, i.e. a d_{GH}-measurable function of Π^α (something we will do in the following section). Since this construction is completely deterministic, we will slightly change our assumptions and at first consider a single element π of $\mathcal{D}$ which is decreasing (the partitions get finer with time). For every integer i, let D_i be the smallest time at which i is in a singleton of π and for every block B with at least two elements, let D_B be the smallest time at which all the elements of B are not in the same block of π anymore. We will assume that π is such that all these are finite.

Proposition 3.5. *There is, up to bijective isometries which preserve roots, a unique complete rooted* $\mathbb{R}$*-tree* $\mathcal{T}$ *equipped with points* $(Q_i)_{i\in\mathbb{N}}$ *such that:*
(i) For all i, $ht(Q_i) = D_i$.
(ii) For all pairs of integers i *and* j, *we have* $ht(Q_i \wedge Q_j) = D_{\{i,j\}}$.
(iii) The set $\bigcup_{i\in\mathbb{N}} [\![\rho, Q_i]\!]$ *is dense in* $\mathcal{T}$.

$\mathcal{T}$ will then be called the *genealogy tree* of π and for all i, Q_i will be called the *death point* of i.

Proof. Let first prove the uniqueness of $\mathcal{T}$. We give ourselves another tree $\mathcal{T}'$ with root ρ' and points $(Q'_i)_{i\in\mathbb{N}}$ which also satisfy (i), (ii) and (iii). First note that, if i and j are two integers such that $Q_i = Q_j$, then $D_{\{i,j\}} = D_i = D_j$ and thus $Q'_i = Q'_j$. This allows us to define a bijection f between the two sets $\{\rho\} \cup \{Q_i,\ i \in \mathbb{N}\}$ and $\{\rho'\} \cup \{Q'_i,\ i \in \mathbb{N}\}$ by letting $f(\rho) = \rho'$ and, for all i, $f(Q_i) = Q'_i$. Now recall that we can recover the metric from the height function and the partial order: we have, for all i and j, $d(Q_i, Q_j) = D_i + D_j - 2D_{\{i,j\}}$, and the same is true in $\mathcal{T}'$. Thus f is isometric and we can (uniquely) extend it to a bijective isometry between $\bigcup_{i\in\mathbb{N}} [\![\rho, Q_i]\!]$ and $\bigcup_{i\in\mathbb{N}} [\![\rho', Q'_i]\!]$, by letting, for $i \in \mathbb{N}$ and $t \in [0, D_i]$, $f(\phi_{\rho,Q_i}(t)) = \phi_{\rho',Q'_i}(t)$. To check that this is well defined, we just need to note that, if i, j and t are such that $\phi_{\rho,Q_i}(t) = \phi_{\rho,Q_j}(t)$, then $t \leqslant D_{\{i,j\}}$ and thus we also have $\phi_{\rho',Q'_i}(t) = \phi_{\rho',Q'_j}(t)$. This extension is still an isometry because it preserves the height and the partial order and is surjective by definition, thus it is a bijection. By standard properties of metric completions, f then extends into a bijective isometry between $\mathcal{T}$ and $\mathcal{T}'$.

To prove the existence of $\mathcal{T}$, we are going to give an abstract construction of it. Let

$$\mathcal{A}_0 = \{(i,t),\ i \in \mathbb{N}, 0 \leqslant t \leqslant D_i\}.$$

A point (i,t) of $\mathcal{A}_0$ should be thought of as representing the block $\pi_{(i)}(t)$. We equip $\mathcal{A}_0$ with the pseudo-distance function d defined such: for all $x = (i,t)$ and $y = (j,s)$ in $\mathcal{A}_0$,

$$d(x,y) = t + s - 2\min(D_{\{i,j\}}, s, t).$$

(equivalently, $d(x,y) = t + s - 2D_{\{i,j\}}$ if $D_{\{i,j\}} \leqslant s,t$ and $d(x,y) = |t-s|$ otherwise.) Let us check that d verifies the four-point inequality from Chapter 1, Proposition 1.2(which in particular, implies the triangle inequality). Let $x = (i,t)$, $y = (j,s)$, $u = (k,a)$, $v = (l,b)$ be in $\mathcal{A}_0$, we want to check that, out of $\min(D_{\{i,j\}}, t, s) + \min(D_{\{k,l\}}, a, b)$, $\min(D_{\{i,k\}}, t, a) + \min(D_{\{j,l\}}, s, b)$ and $\min(D_{\{i,l\}}, t, b) + \min(D_{\{j,k\}}, s, a)$, two are equal and the third one is bigger. Now, there are, up to reordering, two possible cases: either i and j split from k and l at the same time or i splits from $\{j,k,l\}$ at time $t_1 \geqslant 0$, then splits j from $\{k,l\}$ at time $t_2 \geqslant t_1$ and then splits k from l at time $t_3 \geqslant t_2$. After distinguishing these two cases, the problem can be brute-forced through.

Now we want to get an actual metric space out of $\mathcal{A}_0$: this is done by identifying two points of $\mathcal{A}_0$ which represent the same block. More precisely, let us define an equivalence relation $\sim$ on $\mathcal{A}_0$ by saying that, for every pair of points (i,t) and (j,s), $(i,t) \sim (j,s)$ if and only if $d\big((i,t),(j,s)\big) = 0$ (which means that $s = t$ and that $i \sim_{\Pi(t^-)} j$). Then we let $\mathcal{A}$ we the quotient set of $\mathcal{A}_0$ by this relation:

$$\mathcal{A} = \mathcal{A}_0 / \sim \ .$$

The pseudo-metric d passes through the quotient and becomes an actual metric. Even better, the four-point condition also passes through the quotient, and $\mathcal{A}$ is trivially path-connected: every point (i,t) has a simple path connecting it to $(i,0) \sim (1,0)$, namely the path $(i,s)_{0 \leqslant s \leqslant t}$. Therefore, $\mathcal{A}$ is an $\mathbb{R}$-tree, and we will root it at $\rho = (1,0)$. Finally, we let $\mathcal{T}$ be the metric completion of $\mathcal{A}$. It is still a tree, since the four-point condition and connectedness easily pass over to completions.

It is simple to see that $\mathcal{T}$ does satisfy assumptions (i), (ii), (iii) by choosing $Q_i = (i, D_i)$ for all i: (i) and (iii) are immediate, and (ii) comes from the definition of d, which is such that for all i and j, $d\big((i,D_i),(j,D_j)\big) = D_i + D_j - 2D_{i,j}$. □

The natural order on $\mathcal{T}$ is simply described in terms of π:

Proposition 3.6. *Let (i,t) and (j,s) be in $\mathcal{A}$. We have $(i,t) \leqslant (j,s)$ if and only if $t \leqslant s$ and j and i are in the same block of $\pi(t^-)$.*

Proof. By definition, we have $(i,t) \leqslant (j,s)$ if and only if (i,t) is on the segment joining the root and (j,s). Since this segment is none other than $(j,u)_{u \leqslant s}$, this means that $(i,t) \leqslant (j,s)$ if and only if $t \leqslant s$ and $(i,t) \sim (j,t)$. Now, recall that $(i,t) \sim (j,t)$ if and only if $2t - 2\min(D_{i,j}, t) = 0$, i.e. if and only if $t \leqslant D_{i,j}$, and then notice that this last equation is equivalent to the fact that i and j are in the same block of $\pi(t^-)$. This ends the proof. □

The genealogy tree has a canonical measure to go with it, at least under a few conditions: assume that $\mathcal{T}$ is compact, that, for all times t, $\pi(t^-)$ has asymptotic frequencies, and that, for all i, the function $t \mapsto |\pi_{(i)}(t^-)|$ (the asymptotic frequency of the block of $\pi(t^-)$ containing i) is left-continuous (this is not necessarily true, but when it is true it implies that the notation is in fact not ambiguous). Then Proposition 2.1 from Chapter 1 tells us that there exists a unique measure μ on $\mathcal{T}$ such that, for all $(i,t) \in \mathcal{T}, \mu(T_{i,t}) = |\pi_{(i)}(t^-)|$.

3.3 A family of subtrees, embeddings in ℓ^1, and measurability

Proposition 3.7. *There exists a measurable function* TREE : $\mathcal{D} \to \mathbb{T}_W$ *such that, when* Π^α *is a self-similar fragmentation process,* TREE(Π^α) *is the genealogy tree* $\mathcal{T}$ *of* Π^α *equipped with its natural measure.*

This will be proven by providing an embedding of $\mathcal{T}$ in the space ℓ^1 of summable real-valued sequences:

$$\ell^1 = \{x = (x_i)_{i\in\mathbb{N}}; \sum_{i=1}^{\infty} |x_i| < \infty\}$$

and approximating $\mathcal{T}$ by a family of simpler subtrees. For any finite block B, let $\mathcal{T}_B$ be the tree obtained just as before but limiting ourselves to the integers which are in B:

$$\mathcal{T}_B = \{(i,t),\ i \in B, 0 \leqslant t \leqslant D_i\}/\sim\ .$$

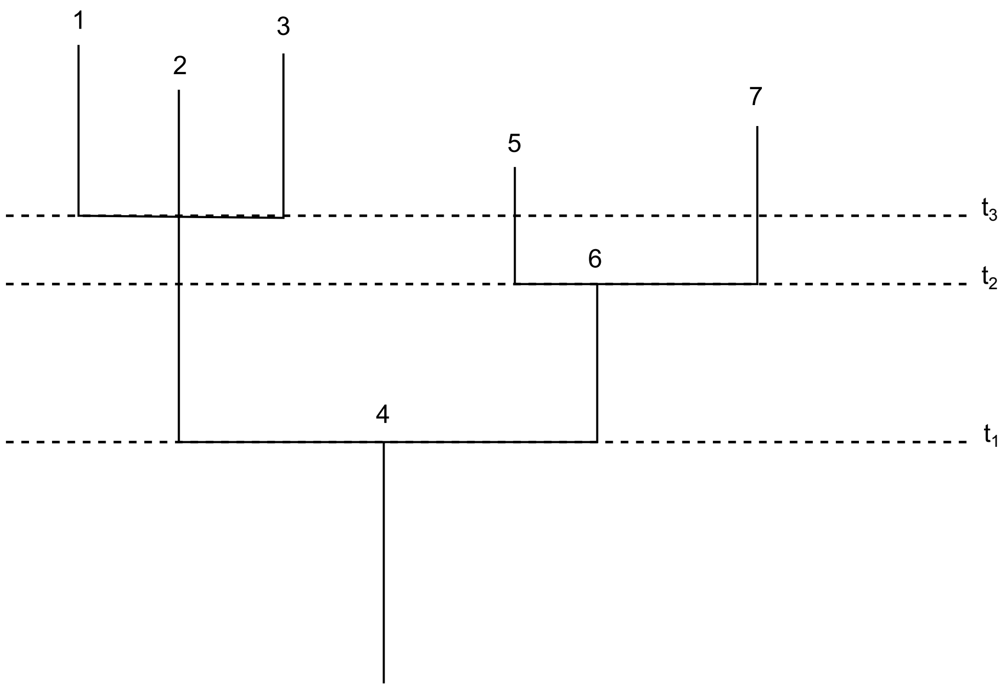

Figure 2.1: A representation of $\mathcal{T}_{[7]}$. Here, $D_{[7]} = t_1$, $D_{\{5,6,7\}} = t_2$ and $D_{\{1,2,3\}} = t_3$.

Do notice that we keep the times $(D_i)_i$ and that we do not change them to the time where i is in a singleton of $\pi \cap B$. Every $\mathcal{T}_B$ is easily seen to be an $\mathbb{R}$-tree since it is a path-connected subset of $\mathcal{T}$, and is also easily seen to be compact since it is just a finite union of segments. Also note that one can completely describe $\mathcal{T}_B$ by saying that it is the reunion of segments indexed by B, such that the segment indexed by integer i has length D_i and two segments indexed by integers i and j split at height $D_{\{i,j\}}$.

The tree $\mathcal{T}_B$ is also equipped with a measure called μ_B, which we define by

$$\mu_B = \frac{1}{\#B} \sum_{i\in B} \delta_{Q_i}.$$

The stick-breaking embedding of Chapter 1, Section 3.1 provides a simultaneous embedding of $\mathcal{T}_B$ in ℓ^1 for all B which is compatible with the natural inclusion of these trees, i.e. the fact that if $B \subset C$, $\mathcal{T}_B \subset \mathcal{T}_C$. The set of marked points used for the construction is the set of death points $(Q_i)_{i\in B}$ for each B.

The following lemma ensures that we are doing measurable operations.

Lemma 3.8. *For every finite block B, there exists a measurable function* $\mathrm{TREE}_B : \mathcal{D} \to \mathbb{T}_W$ *such that, when π is a decreasing element of $\mathcal{D}$ such that D_i is finite for all i,* $\mathrm{TREE}_B(\pi)$ *is the tree $\mathcal{T}_B$ defined above, equipped with the measure μ_B.*

Proof. Note that, since the set of decreasing functions in $\mathcal{D}$ is measurable and all the D_i all also measurable functions, we only need to define TREE_B in our case of interest, and can set it to be any measurable function otherwise.

We will now in fact prove that $\mathcal{T}_B$ is a measurable function of π as a compact subset of ℓ^1 with the Hausdorff metric. First notice that, for all i, Q_i is a measurable function of π (this is because all of its coordinates are themselves measurable). Note then that the map $x \to [\![0, x]\!] = \cup_{n=0}^{\infty}[p_n(x), p_{n+1}(x)]$ from ℓ^1 to the set of its compact subsets is a 1-Lipschitz continuous function of x. This follows from the fact that, for every $n \in \mathbb{N}$, and given two points $x = (x_i)_{i\in\mathbb{N}}$ and $y = (y_i)_{i\in\mathbb{N}}$,

$$\begin{aligned} d_H(\{p_n(x) + tx_{n+1}e_{n+1}, t \in [0,1]\}, \{p_n(y) + ty_{n+1}e_{n+1}, t \in [0,1]\}) &\leqslant ||p_{n+1}(x-y)|| \\ &\leqslant ||x-y||. \end{aligned}$$

Then finally notice that the union operator is continuous for the Hausdorff distance. Combining these three facts, one gets that $\mathcal{T}_B = \underset{i\in B}{\cup}[\![0, Q_i]\!]$ is indeed a measurable function of π.

The fact that μ_B is also a measurable function of π is immediate since all the Q_i are measurable.

□

Lemma 3.9. *For all $t > 0$ and $\varepsilon > 0$, let N_t^ε be the number of blocks of $\pi(t)$ which are not completely reduced to singletons by time $t+\varepsilon$. If, for any choice of t and ε, N_t^ε is finite, then the sequence $(\mathcal{T}_{[n]})_{n\in\mathbb{N}}$ is Cauchy for $d_{\ell^1,H}$, and the limit is isometric to $\mathcal{T}$. In particular, $\mathcal{T}$ is compact.*

Proof. We first want to show that the points $(Q_i)_{i\in\mathbb{N}}$ are *tight* in the sense that for every $\varepsilon > 0$, there exists an integer n such that any point Q_j is within distance ε of a certain Q_i with $i \leqslant n$. The proof of this is essentially the same as the second half of the proof of Lemma 5 in [41], so we will not burden ourselves with the details here. The main idea is that, for any integer l, all the points Q_i with i such that $ht(Q_i) \in (l\varepsilon, (l+1)\varepsilon]$ can be covered by a finite number of balls centered on points of height belonging to $((l-1)\varepsilon, l\varepsilon]$ because of our assumption.

From this, it is easy to see that the sequence $(\mathcal{T}_{[n]})_{n\in\mathbb{N}}$ is Cauchy. Let $\varepsilon > 0$, we take n just as in earlier. For $m \geqslant n$, we then have

$$d_{\ell^1,H}(\mathcal{T}_{[n]}, \mathcal{T}_{[m]}) \leqslant \max_{n+1\leqslant i\leqslant m}\Big(d(Q_i, \mathcal{T}_{[n]})\Big) \leqslant \varepsilon.$$

However, since our sequence is increasing, its limit has no choice but to be the completion of the union. By the uniqueness property of the genealogy tree, this limit is $\mathcal{T}$. □

Lemma 3.10. *The process Π^α almost surely satisfies the hypothesis of Lemma 3.9.*

Proof. Once again, we refer to [41], where this is proved in the first half of Lemma 5. The fact that we are restricted to conservative fragmentations in [41] does not change the details of the computations. □

Thus we have in particular proven that the genealogy tree of Π^α is compact. Let us now turn to the convergence of the measures μ_B to the measure on the genealogy tree.

Lemma 3.11. *Assume that $\mathcal{T}$ is compact, that, for all t, all the blocks of $\pi(t^-)$ and $\pi(t)$ have asymptotic frequencies, and that, for all i, the function $t \mapsto |\pi_{(i)}(t^-)|$ (the asymptotic frequency of the block of $\pi(t^-)$ containing i) is left-continuous. Then the sequence $(\mu_{[n]})_{n\in\mathbb{N}}$ of measures on $\mathcal{T}$ converges to μ.*

Proof. Since $\mathcal{T}$ is compact, Prokhorov's theorem assures us that a subsequence of $(\mu_{[n]})_{n\in\mathbb{N}}$ converges, and we will call its limit μ'. Use of the portmanteau theorem (see [17]) will show that $\mu'(\mathcal{T}_{(i,t)}) = |\pi_{(i)}(t^-)|$ for $(i,t) \in \mathcal{T}$, and the uniqueness part of Proposition 2.1 will imply that μ' and μ must be equal. Let us introduce the notation $\mathcal{T}_{(i,t^+)} = \cup_{s>t}\mathcal{T}_{(i,s)}$ (this is a sub-tree of $\mathcal{T}$, with its root removed). Notice that, for all n, by definition of $\mu_{[n]}$, we have $\mu_{[n]}(\mathcal{T}_{(i,t)}) = \frac{1}{n}\#\big(\pi_{(i)}(t^-)\cap[n]\big)$ and $\mu_{[n]}(\mathcal{T}_{(i,t^+)}) = \frac{1}{n}\#\big(\pi_{(i)}(t)\cap[n]\big)$ and, by definition of the asymptotic frequency of a block, these do indeed converge to $|\pi_{(i)}(t^-)|$ and $|\pi_{(i)}(t)|$. Since $\mathcal{T}_{(i,t)}$ is closed in $\mathcal{T}$ and $\mathcal{T}_{(i,t^+)}$ is open in $\mathcal{T}$, the portmanteau theorem tells us that $\mu'(\mathcal{T}_{(i,t^+)}) \geqslant |\pi_{(i)}(t)|$ and $\mu'(\mathcal{T}_{(i,t)}) \leqslant |\pi_{(i)}(t^-)|$. By writing out

$$\mathcal{T}_{(i,t)} = \bigcap_{n\in\mathbb{N}} \mathcal{T}_{(i,(t-\frac{1}{n})^+)},$$

we then get

$$\mu'(\mathcal{T}_{(i,t)}) \geqslant \lim_{s\to t^-}\mu'(\mathcal{T}_{(i,s^+)}) \geqslant \lim_{s\to t^-}|\pi_{(i)}(s)| \geqslant |\pi_{(i)}(t^-)|.$$

Thus $\mu'(\mathcal{T}_{(i,t)}) = |\pi_{(i)}(t^-)|$ for all choices of i and t, and Proposition 2.1 shows that $\mu' = \mu$. This ends the proof of the lemma. □

Note that, if we assume that $|\pi_{(i)}(t)|$ is right-continuous in t for all i, a similar argument would show that $\mu(\mathcal{T}_{(i,t^+)}) = |\pi_{(i)}(t)|$ for all i and t.

Combining everything we have done so far shows that, under a few conditions, $(\mathcal{T}_{[n]}, \mu_{[n]})$ converges as n goes to infinity to $(\mathcal{T}, \mu)$ in the d_{GHP} sense. We can now define the function TREE which was announced in Proposition 3.7. The set of decreasing elements π of $\mathcal{D}$ such that the sequence $(\mathcal{T}_{[n]}, \mu_{[n]})_{n\in\mathbb{N}}$ converges is measurable since every element of that sequence is measurable. Outside of this set, TREE can have any fixed value. Inside of this set, we let TREE be the aforementioned limit. Since, in the case of the fragmentation process Π^α, the conditions for convergence are met, TREE(Π^α) is indeed the genealogy tree of Π^α.

3.4 Proof of Theorem 3.3

We let $(\mathcal{T}, \mu) = \text{TREE}(\Pi^\alpha)$ and want to show that it is indeed an α-self-similar tree as defined earlier. Let $t \geqslant 0$, and let $\pi = \Pi^\alpha(t)$. For all $i \in \mathbb{N}$ such that π_i is not a singleton, let $\mathcal{T}_i(t)$ be the connected component of $\{x \in \mathcal{T}, ht(x) > t\}$ containing Q_j for all $j \in \pi_i$, and let $x_i = (j,t)$ for any such j. We let also f_i be any bijection: $\mathbb{N} \to \pi_i$ and Ψ_i be the process defined by $\Psi_i(s) = f_i\big(\Pi^\alpha(t+|\pi_i|^{-\alpha}s)\cap\pi_i\big)$ for $s \geqslant 0$. Let us show that, for all i, $(|\pi_i|^\alpha(\mathcal{T}_i(t)\cup\{x_i\}), \mu_{\mathcal{T}_i(t)}) = \text{TREE}(\Psi_i)$. First, $\mathcal{T}_i(t)\cup\{x_i\}$ is compact since it is a closed subset of $\mathcal{T}$. The death points of Ψ_i, which we will call $(Q'_j)_{j\in\mathbb{N}}$ are easily found: for all $j \in \mathbb{N}$, we let $Q'_j = Q_{f(j)}$, it is in $\mathcal{T}_i$ since $f(j)$ is in π_i. By the definition of Ψ, these points have the right distances between them. Similarly, the measure is the expected one: for $(j,s) \in \mathcal{T}_i$, we have $\mu(\mathcal{T}_{j,s}) = |\Pi^\alpha_{(j)}(s^-)| = |\pi_i||\Psi_{(j)}((s-t)^-)|$, which is what was expected.

From the equation $(|\pi_i|^\alpha(\mathcal{T}_i(t)\cup\{x_i\}), \mu_{\mathcal{T}_i(t)}) = \text{TREE}(\Psi_i)$ will come the α-self-similarity property. Recall that

$$\mathcal{G}_t = \sigma(\Pi^\alpha(s), s \leqslant t)$$

and let

$$\mathcal{C}_t = \sigma(|\Pi_i^\alpha(s)|, s \leqslant t, i \in \mathbb{N}) = \sigma(\mu(\mathcal{T}_i(s)), s \leqslant t, i \in \mathbb{N}).$$

We know that, conditionally on $\mathcal{F}_t$, the law of the sequence $(\Psi_i)_{i\in\mathbb{N}}$ is that of a sequence of independent copies of Π^α. Since this law is fixed and $\mathcal{C}_t \subset \mathcal{F}_t$, we deduce that this is also the law of the sequence conditionally on $\mathcal{C}_t$. Applying TREE then says that, conditionally on $\mathcal{C}_t$, the $(|\pi_i|^\alpha(\mathcal{T}_i(t) \cup \{x_i\}), \mu_{\mathcal{T}_i(t)})_{i\in\mathbb{N}}$ are mutually independent and have the same law as $(\mathcal{T}, \mu)$ for all choices of $i \in \mathbb{N}$.

Finally, we need to check that the fragmentation process derived from $(\mathcal{T}, \mu)$ has the same law as Π^α. Let $(P_i)_{i\in\mathbb{N}}$ be an exchangeable sequence of $\mathcal{T}$-valued variables directed by μ. The partition-valued process $\Pi_{\mathcal{T}}$ defined in Proposition 3.2 is an α-self-similar fragmentation process. To check that it has the same law as Π^α, one only needs to check that it has almost surely the same asymptotic frequencies as Π^α. Indeed, Bertoin's Poissonian construction shows that the distribution of the asymptotic frequencies of a fragmentation process determine α, c and ν. Let $t \geqslant 0$, take any non-singleton block B of $\Pi_{\mathcal{T}}(t)$, and let C be the connected component of $\{x \in \mathcal{T}, ht(x) > t\}$ containing P_i for all $i \in B$. By the law of large numbers, we have $|B| = \mu(C)$ almost surely. Thus the nonzero asymptotic frequencies of the blocks of $\Pi_{\mathcal{T}}(t)$ are the μ-masses of the connected components of $\mathcal{T}_{>t}$, which are of course the asymptotic frequencies of the blocks of $\Pi^\alpha(t)$. We then get this equality for all t almost surely by first looking only at rational times and then using right-continuity. □

3.5 Leaves of the fragmentation tree

Definition 3.12. *There are three kinds of points in $\mathcal{T} = \mathrm{TREE}(\Pi^\alpha)$:*

-skeleton points, which are of the form (i,t) with $t < D_i$.

-"dead" leaves, which come from the sudden total disappearance of a block: they are the points (i, D_i) such that $|\Pi^\alpha_{(i)}(D_i^-)| \neq 0$ but $\Pi^\alpha(D_i) \cap \Pi^\alpha_{(i)}(D_i^-)$ is only made of singletons. These only exist if ν gives some mass to $(0,0,\ldots)$, and are the leaves which are atoms of μ.

-"proper" leaves, which are either of the form (i, D_i) such that $|\Pi^\alpha_{(i)}(D_i^-)| = 0$ or which are limits of sequences of the form $(i_n, t_n)_{n\in\mathbb{N}}$ such that $(t_n)_{n\in\mathbb{N}}$ is strictly increasing and $|\Pi^\alpha_{(i_n)}(t_n)|$ tends to 0 as n goes to infinity.

Note that, if ν is conservative and the erosion coefficient is zero, then not only are there no dead leaves, but all the (i, D_i) are proper leaves: none of the processes $(|\Pi^\alpha_{(i)}(t)|)_{t<D_i}$ suddenly jump to 0. On the other hand, if ν is not conservative or if there is some erosion, then all the (i, D_i) are either skeleton points or dead leaves, and all the proper leaves can only be obtained by taking limits, which implies that μ does not charge the proper leaves at all.

Recall the construction of the α-self-similar fragmentation process through a homogeneous fragmentation process, which we call Π, and the time changes τ_i defined, for all i and t by $\tau_i(t) = \inf\{u, \int_0^u |\Pi_{(i)}(r)|^{-\alpha} \mathrm{d}r > t\}$. Notice also that if $t > D_i$, $\tau_i(t) = \infty$.

Proposition 3.13. *Let $(i_n, t_n)_{n\in\mathbb{N}}$ be a strictly increasing sequence of points of the skeleton of $\mathcal{T}$, which converges in $\mathcal{T}$. The following are equivalent:*

(i) $|\Pi^\alpha_{(i_n)}(t_n^-)|$ goes to 0 as n tends to infinity, making the limit of $(i_n, t_n)_{n\in\mathbb{N}}$ a proper leaf.

(ii) $\tau_{i_n}(t_n)$ goes to infinity as n tends to infinity.

Proof. To show that (ii) implies (i), first note that, for every pair (i,t) which is in $\mathcal{T}$, we have by definition $t \geqslant \tau_i(t)|\Pi^\alpha_{(i)}(t^-)|^{|\alpha|}$. Since $\mathcal{T}$ is bounded, the product $\tau_{i_n}(t_n)|\Pi^\alpha_{(i_n)}(t_n^-)|^{|\alpha|}$ must stay bounded. Thus, if one factor tends to infinity, the other one must tend to 0. For the converse, let us show that if (ii) does not hold, then (i) also does not. Assume that $\tau_{(i_n)}(t_n)$ converges to a finite number l. Now we know that, because of the Poissonian way that Π is constructed,

$\bigcap_{n\in\mathbb{N}} \Pi_{(i_n)}(\tau_{i_n}(t_n))$ is a block of $\Pi(l^-)$. Let i be in this block, we can now assume that $i_n = i$ for all n, and that t_n converges to D_i as n goes to infinity, with $\tau_{(i)}(D_i) = l$. The limit of $|\Pi^\alpha_{(i)}(t_n^-)|$ as n tends to infinity is then $|\Pi_{(i)}(l^-)|$, which is nonzero because the subordinator $-\log(|\Pi_{(i)}(t)|)_{t\geqslant 0}$ cannot continuously reach infinity in finite time. □

General leaves of $\mathcal{T}$ can also be described the following way: let L be a leaf. For all $t < ht(L)$, L has a unique ancestor with height t. This ancestor is a skeleton point of the form (j, t) with $j \in \mathbb{N}$. Letting $i_L(t)$ be the smallest element of $\Pi_{(j)}(t^-)$, then $(i_L(t), t)_{t<ht(L)}$ is a kind of canonical description of the path going to L and uniquely determines L.

4 Malthusian fragmentations, martingales, and applications

In order to study the fractal structure of $\mathcal{T}$ in detail, we will need some additional assumptions on c and ν: we turn to the *Malthusian* setting which was first introduced by Bertoin and Gnedin in [12], albeit in a very different environment, since they were interested in fragmentations with a nonnegative index of self-similarity.

4.1 Malthusian hypotheses and additive martingales

In this section, we will mostly be concerned with homogeneous fragmentations: $(\Pi(t))_{t\geqslant 0}$ is the $(\nu, 0, c)$-fragmentation process derived from a point process $(\Delta_t, k_t)_{t\geqslant 0}$, with dislocation measure ν and erosion coefficient c, and $(\mathcal{F}_t)_{t\geqslant 0}$ is the canonical filtration of the point process.

We first start with a few analytical preliminaries. For convenience's sake, we will do a translation of the variable p of the Laplace exponent ϕ defined in Section 2.1.4:

Lemma 4.1. *For all real p, let $\psi(p) = \phi(p-1) = cp + \int_{\mathcal{S}^\downarrow}(1 - \sum_i s_i^p)\mathrm{d}\nu(\mathbf{s})$. Then $\psi(p) \in [-\infty, +\infty)$, and this function is strictly increasing and concave on the set where it is finite.*

Proof. The only difficult point here is to prove for all real p that $\psi(p) \in [-\infty, +\infty)$. In other words, we want to give an upper bound to $1 - \sum_i s_i^p$ which is integrable with respect to ν. Such a bound is for example $1 - s_1^p$. Indeed, by letting $C_p = \sup_{x\in[0,1[} \frac{1-x^p}{1-x}$ (which is finite), we have $1 - s_1^p \leqslant C_p(1 - s_1)$, and $1 - s_1$ is integrable by assumption. □

Note that, even for negative p, as soon as $\phi(p) > -\infty$, we have, for all t,

$$\mathbb{E}\left[|\Pi_{(1)}(t)|^p \mathbb{1}_{\{|\Pi_{(1)}(t)|>0\}}\right] = e^{-t\phi(p)}.$$

This follows from the description of the Lévy measure of the subordinator $\xi_t = -\log|\Pi_{(1)}(t)|$ (see [9], Theorem 3).

Definition 4.2. *We say that the pair (c, ν) is* Malthusian *if there exists a strictly positive number p^* (which is necessarily unique), called the* Malthusian exponent *such that*

$$\phi(p^* - 1) = \psi(p^*) = cp^* + \int_{\mathcal{S}^\downarrow} \big(1 - \sum_{i=1}^{\infty} s_i^{p^*}\big)\mathrm{d}\nu(\mathbf{s}) = 0.$$

The typical example of pairs (c, ν) with a Malthusian exponent are *conservative* fragmentations, where $c = 0$ and $\sum_i s_i = 1$ ν-almost everywhere. In that case, the Malthusian exponent is simply 1. Note that assumption (**H**) defined in the introduction implies the existence of the Malthusian exponent, since $\psi(1) \geqslant 0$ for all choices of ν and c.

We assume from now on the existence of the Malthusian exponent p^*. This allows us to define, for $i \in \mathbb{N}$, $t \geqslant 0$ and $s \geqslant 0$,

$$M_{i,t}(s) = \sum_{j=1}^{\infty} |\Pi_j(t+s) \cap \Pi_{(i)}(t)|^{p^*}.$$

This is the sum of the sizes of the blocks of the part of the fragmentation which is issued from $\Pi_i(t)$, each one taken to the p^*-th power. In the case of $i = 1$, $t = 0$, we let $M(s) = M_{1,0}(s)$, the sum of the sizes of all the blocks of $\Pi(s)$ to the p^*-th power. These processes are interesting because the Malthusian hypothesis naturally makes them martingales.

Proposition 4.3. *For all $i \in \mathbb{N}$ and $t \geqslant 0$, the process $\big(M_{i,t}(s)\big)_{s\geqslant 0}$ is a càdlàg martingale with respect to the filtration $(\mathcal{F}_{t+s})_{s\geqslant 0}$.*

Proof. Let us first notice that, as a consequence of the fragmentation property, for every (i,t), the process $\big(M_{i,t}(s)\big)_{s\geqslant 0}$ has the same law as a copy of the process $(M(s))_{s\geqslant 0}$ which is independent of $\mathcal{F}_t$, multiplied by $|\Pi_{(i)}(t)|^{p^*}$ (which is an $\mathcal{F}_t$-measurable variable). Thus, we only need to prove the martingale property for $(M(s))_{s\geqslant 0}$. Recall that, given $\pi \in \mathcal{P}_{\mathbb{N}}$, $\mathrm{rep}(\pi)$ is the set of integers which are the smallest element of the block of π containing them and let $t \geqslant 0$ and $s \geqslant 0$, we have

$$\begin{aligned}
\mathbb{E}[M(t+s) \,|\, \mathcal{F}_s] &= \mathbb{E}\left[\sum_{i\in \mathrm{rep}(\Pi(s))} M_{i,s}(t) \,|\, \mathcal{F}_s\right] \\
&= \sum_{i\in\mathbb{N}} \mathbb{E}[\mathbb{1}_{\left\{i\in \mathrm{rep}(\Pi(s))\right\}} M_{i,s}(t) \,|\, \mathcal{F}_s] \\
&= \sum_{i\in\mathbb{N}} \mathbb{1}_{\left\{i\in \mathrm{rep}\left(\Pi(s)\right)\right\}} |\Pi_{(i)}(s)|^{p^*} \mathbb{E}[M(t)] \\
&= \sum_{i\in \mathrm{rep}\left(\Pi(s)\right)} |\Pi_{(i)}(s)|^{p^*} \mathbb{E}[M(t)] \\
&= M(s)\mathbb{E}[M(t)].
\end{aligned}$$

Thus we only need to show that $\mathbb{E}[M(t)] = 1$ for all t and our proof will be complete. To do this, one uses the fact that, since $\Pi(t)$ is an exchangeable partition, the asymptotic frequency of the block containing 1 is a size-biased pick from the asymptotic frequencies of all the blocks. This tells us that

$$\begin{aligned}
\mathbb{E}\left[\sum_i |\Pi_i(t)|^{p*}\right] &= \mathbb{E}\left[\sum_i |\Pi_i(t)||\Pi_i(t)|^{p*-1}\mathbb{1}_{|\Pi_{(i)}(t)|\neq 0}\right] \\
&= \mathbb{E}[|\Pi_{(1)}(t)|^{p^*-1}\mathbb{1}_{\{|\Pi_{(i)}(t)|\neq 0\}}] \\
&= \exp[-t\phi(p^*-1)] \\
&= 1.
\end{aligned}$$

We refer to [13] for the proof that $(M(t))_{t\geqslant 0}$ is càdlàg (it is assumed in [13] that $c = 0$ and that ν is conservative but these assumptions have no effect on the proof). ☐

Since these martingales are nonnegative, they all converge almost surely. For integer i and real t, we will call $W_{i,t}$ the limit of the martingale $M_{i,t}$ on the event where this martingale converges. We also write W instead of $W_{1,0}$ for simplicity. Our goal is now to investigate these limits. To this effect, let us introduce a family of integrability conditions indexed by a parameter $q > 1$: we let $(\mathbf{M_q})$ be the assumption that

$$\int_{\mathcal{S}^{\downarrow}} \left| 1 - \sum_{i=1}^{\infty} s_i^{p^*} \right|^q \nu(\mathrm{d}\mathbf{s}) < \infty.$$

We will assume through the rest of this section that there exists some $q > 1$ such that $(\mathbf{M_q})$ holds.

The following is a generalization of Theorem 1.1 and Proposition 1.5 of [11] which were restricted to the case where ν has finite total mass.

Proposition 4.4. *Assume* $(\mathbf{M_q})$ *for some* $q > 1$*. Then the martingale* $(M(t))_{t\geqslant 0}$ *converges to* W *in* L^q*.*

Proof. We will first show that the martingale $\big(M(t)\big)_{t\geqslant 0}$ is purely discontinuous in the sense of [27], which we will do by proving that it has finite variation on any bounded interval $[0,T]$ with $T > 0$. To this effect, write, for all t, $M(t) = \mathrm{e}^{-cp^* t} \sum_i \big(X_i(t)\big)^{p^*}$ where the $(X_i(t))_{i\in\mathbb{N}}$ are the sizes of the blocks of a homogeneous fragmentation with dislocation measure ν, but no erosion. Since the product of a bounded nonincreasing function with a bounded function of finite variation has finite variation, we only need to check that $t \mapsto \sum_i X_i(t)^{p^*}$ has finite variation on $[0,T]$. Since this function is just a sum of jumps, its total variation is equal to the sum of the absolute values of these jumps. Thus we want to show that $|\sum_{t\leqslant T}\sum_i \big(X_i(t)\big)^{p^*} - \big(X_i(t^-)\big)^{p^*}|$ is finite. This sum is equal to $\sum_{t\leqslant T} e^{cp^* t}|M(t) - M(t^-)|$, which is bounded above by $e^{cp^* T}\sum_{t\leqslant T}|M(t) - M(t^-)|$. We will not show the finiteness of this sum, because it can be done by computing its expectation similarly to our next computation.

Knowing that the martingale is purely discontinuous, according to [59] (at the bottom of page 299), to show that the martingale is bounded in L^q, one only needs to show that the sum of the $q-th$ powers of its jumps is bounded in L^1, i.e. that

$$\mathbb{E}\Big[\sum_t |M(t) - M(t^-)|^q\Big] < \infty.$$

This expected value can be computed with the Master formula for Poisson point processes (see [72], page 475). Recall from Section 2.1.2 the construction of Π through a family of Poisson point processes $((\Delta^k(t))_{t\geqslant 0})_{k\in\mathbb{N}}$: for t and k such that there is an atom $\Delta^k(t)$, the k-th block of $\Pi(t^-)$ is replaced by its intersection with $\Delta^k(t)$. We then have

$$\begin{aligned}
\mathbb{E}\left[\sum_{t\geqslant 0} |M(t) - M(t^-)|^q\right] &= \mathbb{E}\left[\sum_{k=1}^{\infty}\sum_{t\geqslant 0} |\Pi_k(t-)|^{qp^*}\left(\left|1 - \sum_{i=1}^{\infty}|\Delta_i^k(t)|^{p^*}\right|\right)^q\right] \\
&= \mathbb{E}\left[\int_0^{\infty}\sum_k |\Pi_k(t-)|^{qp^*} dt\right]\int_{\mathcal{S}^{\downarrow}} \big|1 - \sum_i s_i^{p^*}\big|^q \mathrm{d}\nu(\mathbf{s}) \\
&= \int_0^{\infty} \mathrm{e}^{-t\psi(qp^*)} dt \int_{\mathcal{S}^{\downarrow}} \big|1 - \sum_i s_i^{p^*}\big|^q \mathrm{d}\nu(\mathbf{s}).
\end{aligned}$$

Since $qp^* > p^*$, we have $\psi(qp^*) > 0$ and thus the expectation is finite. □

Proposition 4.5. *Assume that* $\mathbb{E}[W] = 1$ *(which is equivalent to assuming that the martingale* $\big(M(t)\big)_{t\geqslant 0}$ *converges in* L^1*). Then, almost surely, if* Π *does not die in finite time then* W *is strictly positive.*

Proof. We discretize the problem and only look at integer times: for $n \in \mathbb{N}$, let Z_n is the number of blocks of $\Pi(n)$ which have nonzero mass. The process $(Z_n)_{n\in\mathbb{N}}$ is a Galton-Watson process (possibly taking infinite values. See Appendix A to check that standard results stay true in this case). If it is critical or subcritical then there is nothing to say, and if it is supercritical, notice that the event $\{W = 0\}$ is hereditary (in the sense that $W = 0$ if and only if all the $W_{i,1}$ are also zero). This implies that the probability of the event $\{W = 0\}$ is either equal to 1 or to the probability of extinction. But since $\mathbb{E}[W] = 1$, W cannot be 0 almost surely and thus $\{W = 0\}$ and the event of extinction have the same probabilities. Since $\{W = 0\}$ is a subset of the event of extinction, W is nonzero almost surely on nonextinction.

□

The following proposition states the major properties of these martingale limits.

Proposition 4.6. *There exists an event of probability* 1 *on which the following are true:*
(i) For every i *and* t*, the martingale* $M_{i,t}$ *converges to* $W_{i,t}$*.*
(ii) For every integer i*, and any times* t *and* s *with* $s > t$*, we have*

$$W_{i,t} = \sum_{j\in\Pi_{(i)}(t)\cap\mathrm{rep}\big(\Pi(s)\big)} W_{j,s}.$$

(iii) For every i*, the function* $t \mapsto W_{i,t}$ *is nonincreasing and right-continuous. The left-limits can be described as follows: for every* t*, we have*

$$W_{i,t^-} = \sum_{j\in\Pi_{(i)}(t^-)\cap\mathrm{rep}(\Pi(t))} W_{j,t}.$$

To prove this we will need the help of several lemmas. The first is an intermediate version of point (ii)

Lemma 4.7. *For any integer* i *and any times* t *and* s *such that* $s > t$*, there exists an event of probability* 1 *on which the martingales* $M_{i,t}$ *and* $M_{j,s}$ *converge for all* j *and we have the relation*

$$W_{i,t} = \sum_{j\in\Pi_{(i)}(t)\cap\mathrm{rep}(\Pi(s))} W_{j,s}.$$

Proof. For clarity's sake, we are going to restrict ourselves to the case where $i = 1$ and $t = 0$, but the proof for the other cases is similar. We have, for all $r \geqslant s$,

$$M(r) = \sum_{j\in\mathrm{rep}(\Pi(s))} M_{j,s}(r - s).$$

We cannot immediately take the limits as r goes to ∞ because we do not have any kind of dominated convergence under the sum. However, Fatou's Lemma does give us the inequality

$$W \geqslant \sum_{j\in\mathrm{rep}(\Pi(s))} W_{j,s}.$$

To show that these are actually equal almost surely, we show that their expectations are equal. We know that $\mathbb{E}[W]=1$ and that, for all $j\in\mathbb{N}$ and $s\geqslant 0$, one can write $W_{j,s}=|\Pi_{(j)}(s)|^{p^*}W'_{j,s}$ where $W'_{j,s}$ is a copy of W which is independent of $\mathcal{F}_s$. We thus have

$$\begin{aligned}
\mathbb{E}\left[\sum_{j\in\mathrm{rep}(\Pi(s))} W_{j,s}\right] &= \mathbb{E}\left[\sum_{j\in\mathbb{N}} \mathbb{1}_{\left\{j\in\mathrm{rep}(\Pi(s))\right\}}|\Pi_{(j)}(s)|^{p^*}W'_{j,s}\right] \\
&= \sum_{j\in\mathbb{N}} \mathbb{E}[\mathbb{1}_{\left\{j\in\mathrm{rep}(\Pi(s))\right\}}|\Pi_{(j)}(s)|^{p^*}W'_{j,s}] \\
&= \sum_{j\in\mathbb{N}} \mathbb{E}[W'_{j,s}]\mathbb{E}[\mathbb{1}_{\left\{j\in\mathrm{rep}(\Pi(s))\right\}}|\Pi_{(j)}(s)|^{p^*}] \\
&= \sum_{j\in\mathbb{N}} \mathbb{E}[\mathbb{1}_{\left\{j\in\mathrm{rep}(\Pi(s))\right\}}|\Pi_{(j)}(s)|^{p^*}] \\
&= \mathbb{E}[M(s)] \\
&= 1.
\end{aligned}$$

□

Lemma 4.8. *For every pair of integers i and j, let $f_{i,j}$ be a nonnegative function defined on $[0,+\infty)$. For every i, we let f_i be the function $\sum_j f_{i,j}$, and we also let $f=\sum_i f_i$. We assume that, for every i and j, the function $f_{i,j}$ converges at infinity to a limit called $l_{i,j}$, and we also assume that f converges, its limit being $l=\sum_{i,j} l_{i,j}$. Then, for every i, the function f_i also converges at infinity and its limit is $l_i=\sum_j l_{i,j}$.*

Proof. We are going to prove that $\liminf f_i=\limsup f_i=l_i$ for all i. Let N be any integer, taking the upper limit in the relation $f\geqslant\sum_{i\leqslant N} f_i$ gives us $l\geqslant\sum_{i\leqslant N}\limsup f_i$, and by taking the limit as N goes to infinity, we have $l\geqslant\sum_i\limsup f_i$. Similarly, for every i, the relation $f_i=\sum_j f_{i,j}$ gives us $\liminf f_i\geqslant\sum_j l_{i,j}$. We thus have the following chain:

$$\sum_{i,j} l_{i,j}\leqslant\sum_i\liminf f_i\leqslant\sum_i\limsup f_i\leqslant\sum_{i,j} l_{i,j},$$

and this implies that, for every i, $\liminf f_i=\limsup f_i=l_i$. □

Proof of Proposition 4.6: let $t<s$ be two times and assume that the martingale $M_{j,s}$ converges for all j, and also assume the relation $W=\sum\limits_{j\in\mathrm{rep}(\Pi(s))} W_{j,s}$. Apply Lemma 4.8 with, for nonnegative r, $f(r)=M(s+r)$, $f_i(r)=\mathbb{1}_{\{i\in\mathrm{rep}(\Pi(t))\}}M_{i,t}(r+s-t)$ and $f_{i,j}(r)=\mathbb{1}_{\{i\in\mathrm{rep}(\Pi(t))\}}\mathbb{1}_{\{j\in\mathrm{rep}\Pi(s)\cap\Pi_{(i)}(t)\}}M_{j,s}(r)$. Then, for all i, the martingale $M_{i,t}$ does indeed converge, and point (ii) of the proposition is none other than the relation $l_i=\sum_j l_{i,j}$. We also get that $W=\sum_{i\in\mathrm{rep}\Pi(t)} W_{i,t}$ and thus can use the same reasoning to obtain $W_{i,r}=\sum\limits_{j\in\Pi_{(i)}(r)\cap\mathrm{rep}\Pi(t)} W_{j,t}$ for all $r<t<s$.

By Lemma 4.7, the assumption of the previous paragraph is true for any value of s with probability 1, we then obtain points (i) and (ii) by taking a sequence of values of s tending to infinity.

We can turn ourselves to point (iii). Fixing an integer i, it is clear that $t\mapsto W_{i,t}$ is non-increasing. Right-continuity is obtained by the monotone convergence theorem, noticing that $\Pi_{(i)}(t)\cap\mathrm{rep}(\Pi(s))$ is the increasing union, as u decreases to t, of sets $\Pi_{(i)}(u)\cap\mathrm{rep}(\Pi(s))$. Similarly, the fact that $W_{i,t^-}=\sum_{j\in\Pi_{(i)}(t^-)\cap\mathrm{rep}(\Pi(t))} W_{j,t}$ is only a matter of noticing that $\Pi_{(i)}(t^-)$

is the decreasing intersection, as u increases to t, of sets $\Pi_{(i)}(u)$ and taking the infimum on both sides of the relation $W_{i,u} = \sum_{j\in\Pi_{(i)}(u)\cap\text{rep}(\Pi(t))} W_{j,t}$. □

From now on we will restrict ourselves to the aforementioned almost-sure event: all the additive martingales are now assumed to converge, and the limits satisfy the natural additive properties.

4.2 A measure on the leaves of the fragmentation tree.

In this section we are going to assume that $\mathbb{E}[W] = 1$. We let $\mathcal{T}$ be the genealogy tree of the self-similar process Π^α and are going to use the martingale limits to define a new measure on $\mathcal{T}$.

Theorem 4.9. *On an event with probability one, there exists a unique measure μ^* on $\mathcal{T}$ which is fully supported by the proper leaves of $\mathcal{T}$ and which satisfies*

$$\forall i \in \mathbb{N}, t \geqslant 0, \mu^*(\mathcal{T}_{(i,t+)}) = W_{i,\tau_i(t)}.$$

where $\mathcal{T}_{i,t^+}$ is as defined in the proof of Lemma 3.11: $\mathcal{T}_{(i,t^+)} = \cup_{s>t}\mathcal{T}_{(i,s)}$.

Proof. This will be a natural consequence of Proposition 2.1 from Chapter 1, and our previous study of the convergence of additive martingales. Note that, since, for all $(i,t) \in \mathcal{T}$, we have

$$\mathcal{T}_{(i,t)} = \bigcup_{j\in\Pi_{(i)}(t^-)\cap\text{rep}\big(\Pi(t)\big)} \mathcal{T}_{(j,t^+)},$$

any candidate for μ^* would then have to satisfy, for every (i,t), the relation

$$\mu^*(\mathcal{T}_{(i,t)}) = \sum_{j\in\Pi_{(i)}(t^-)\cap\text{rep}(\Pi(t))} W_{j,\tau_i(t)} = W_{i,(\tau_i(t))^-}.$$

We thus know that we can apply Proposition 2.1 to the function m defined by $m(i,t) = W_{i,(\tau_i(t))^-}$. This function is indeed decreasing and left-continuous on $\mathcal{T}$, and we also have, for every point (i,t) of $\mathcal{T}$, $m\big((i,t)^+\big) = m(i,t)$ in the sense of Section 2.2.4 (this is point (iii) of Proposition 4.6). Thus μ^* exists and is unique, and we only now need to check that it is fully supported by the set of proper leaves of $\mathcal{T}$. To do this, notice first that, by Proposition 3.13, the complement of the set of proper leaves can be written as $\cup_{N\in\mathbb{N}}\{(i,s),\ i\in\mathbb{N}, \tau_i(s)\leqslant N\}$, and then that, for every integer N,

$$\mu^*\big(\{(i,s),\ i\in\mathbb{N},\tau_i(s)\leqslant N\}\big) = W - \sum_{i\in\text{rep}(\Pi(N))} W_{i,N} = 0.$$ □

The measure μ^* has total mass W, which is in general not 1. However, having assumed that $\mathbb{E}[W] = 1$, we will be able to create some probability measures involving μ^*. The following one can be interpreted as the "distribution" of the process of the size of the fragment associated to a leaf with "distribution" μ^*. Recall first that to every leaf L of $\mathcal{T}$ corresponds a family of integers $(i_L(t))_{t<ht(L)}$ such that, for all t, $i_L(t)$ is the smallest integer such that $(i_L(t),t)\leqslant L$ in $\mathcal{T}$.

Proposition 4.10. *Define a probability measure Q on the space $\mathcal{D}([0,+\infty))$ of càdlàg functions from $[0,+\infty)$ to $[0,+\infty)$ by setting, for all nonnegative measurable functionals $F:\mathcal{D}([0,+\infty))\to[0,+\infty)$,*

$$Q(F) = \mathbb{E}\left[\int_{\mathcal{T}} F\big((|\Pi^\alpha_{(i_L(t))}(t)|)_{t\geqslant 0}\big)\mu^*(\mathrm{d}L)\right].$$

Let $(x_t)_{t\geqslant 0}$ be the canonical process, and let ζ be the time-change defined for all $t \geqslant 0$ by:

$$\zeta(t) = \inf\Big\{u, \int_0^u x_t^\alpha dr > u\Big\}.$$

Under the law Q, the process $(\xi_t)_{t\geqslant 0}$ defined by $\xi_t = -\log(x_{\zeta(t)})$ for all $t \geqslant 0$ is a subordinator whose Laplace exponent ϕ^* satisfies, for p such that $\psi(p+p^*)$ is defined:

$$\phi^*(p) = cp + \int_{\mathcal{S}^\downarrow} \Big(\sum_i (1-s_i^p) s_i^{p^*}\Big)\nu(\mathrm{d}\mathbf{s}) = \psi(p+p^*).$$

As before, the function ϕ^* can be seen as defined on $\mathbb{R}$, in which case it takes values in $[-\infty, \infty)$.

Proof. Let us first show that, given a nonnegative and measurable function f on $[0,+\infty)$ and a time t, we have

$$Q(f(x_{\zeta(t)})) = \mathbb{E}\left[\sum_i |\Pi_i(t)|^{p^*} f(|\Pi_i(t)|)\right]. \tag{2.1}$$

To do this, notice first that we have $\Pi^\alpha_{(i_L(t))}(\tau^{-1}_{i_L(t)}(t)) = \Pi_{(i_L(t))}(t)$. Thus, using the definition of μ^*, one can change the integral with the respect to μ^* into a sum on the different blocks of $\Pi(t)$:

$$Q(f(x_{\zeta(t)})) = \mathbb{E}\left[\sum_{i\in \mathrm{rep}(\Pi(t))} W_{i,t} f(|\Pi_{(i)}(t)|)\right].$$

Finally, with the fragmentation property, one can write, for all t and i, $W_{i,t} = |\Pi_{(i)}(t)|^{p^*} W'_{i,t}$ where $W'_{i,t}$ is a copy of W which is independent of $|\Pi(t)|$. Since $\mathbb{E}[W] = 1$, we get formula (2.1).

Applying this to the function f defined by $f(x) = x^p$ gives us our moments formula:

$$Q(\mathrm{e}^{-p\xi_1}) = Q(x^p_{\zeta(1)}) = \mathbb{E}\left[\sum_i |\Pi_i(1)|^{p^*+p}\right] = \mathbb{E}[|\Pi_1(1)|^{p+p^*-1}] = \exp[-\big(\phi(p+p^*-1)\big)].$$

Independence and stationarity of the increments is proved the same way. Let $s < t$, f be any nonnegative measurable functions on $\mathbb{R}$ and G be any nonnegative measurable function on $\mathcal{D}([0,s])$. Let us apply the fragmentation property for Π at time s: for $i \in \mathrm{rep}\big(\Pi(s)\big)$, the partition of $\Pi_{(i)}(s)$ formed by the blocks of $\Pi(t)$ which are subsets of $\Pi_{(i)}(s)$ can be written as $\Pi_{(i)}(s) \cap \Pi^i(t-s)$ where $\big(\Pi^i(u)\big)_{u\geqslant 0}$ is an independent copy of Π. Thus one can write

$$\begin{aligned} &Q[f\big(\frac{x_{\zeta(t)}}{x_{\zeta(s)}}\big) G\big((x_{\zeta(u)})_{u\leqslant s}\big)] \\ &\qquad = \mathbb{E}\left[\sum_{i\in \mathrm{rep}\Pi(s)} |\Pi_{(i)}(s)|^{p^*} G\big((|\Pi_{(i)}(u)|)_{u\leqslant s}\big) \sum_{j\in\mathbb{N}} W_j^i |\Pi_j^i(t-s)|^{p^*} f\big(|\Pi^i_j(t-s)|\big)\right], \end{aligned}$$

where the W^i_j are copies of W independent of anything happening before time t, which all have expectation 1. We thus get

$$\begin{aligned} &Q\Big[f(\frac{x_{\zeta(t)}}{x_{\zeta(s)}}) G\big((x_{\zeta(u)})_{u\leqslant s}\big)\Big] \\ &\qquad = \mathbb{E}\left[\sum_{i\in \mathrm{rep}\Pi(s)} |\Pi_{(i)}(s)|^{p^*} G((|\Pi_{(i)}(u)|)_{u\leqslant s})\right] \mathbb{E}\left[\sum_j |\Pi_i(t-s)|^{p^*} f\big(|\Pi_j(t-s)|\big)\right], \end{aligned}$$

which is what we wanted.

□

Lemma 4.11. *Assume that ν integrates the quantity $\sum_i \log(s_i)s_i^{p^*}$ and let $\underline{p} = \sup\{q \in \mathbb{R} : \phi^*(-q) > -\infty\}$. Then if $\gamma < 1 + \frac{\underline{p}}{|\alpha|}$, we have*

$$\mathbb{E}\left[\int_{\mathcal{T}} ht(L)^{-\gamma}\mu^*(\mathrm{d}L)\right] < \infty.$$

Proof. We know that the height of the leaf is equal to the death time of the fragment it marks: $ht(L) = \inf\{t : \tau_{i_L(t)}(t) = \infty\}$. Thus we can write, using the measure Q

$$\mathbb{E}\left[\int_{\mathcal{T}} ht(L)^{-\gamma}\mu^*(\mathrm{d}L)\right] = Q[I^{-\gamma}],$$

where $I = \int_0^\infty \mathrm{e}^{\alpha\xi_t}\mathrm{d}t$ is the exponential functional of the subordinator ξ with Laplace exponent ϕ^*. Following the proof of Proposition 2 in [15], one has, if $1 < \gamma < 1 + \frac{\underline{p}}{|\alpha|}$,

$$Q[I^{-\gamma}] = \frac{-\phi^*(-|\alpha|(\gamma-1))}{\gamma-1}Q[I^{-\gamma+1}].$$

By induction we then only need to show that $Q[I^{-\gamma}]$ is finite for $\gamma \in (0,1]$, and thus only need to show that $Q[I^{-1}]$ is finite. However, it is well known (see for example [14]) that $Q[I^{-1}] = (\phi^*)'(0^+) = c - \int_{\mathcal{S}^\downarrow}(\sum_i \log(s_i)s_i^{p^*})\nu(\mathrm{d}\mathbf{s})$, which is finite by assumption. □

The assumption that $\int_{\mathcal{S}^\downarrow}(\sum_i \log(s_i)s_i^{p^*})\nu(\mathrm{d}\mathbf{s})$ is finite is for example verified when ν has finite total mass, and (**H**) is satisfied: pick $\delta > 0$ such that $\psi(p^* - \delta) > -\infty$, then pick $K > 0$ such that $|\log(x)| \leqslant Kx^{-\delta}$ for all $x \in (0,1]$, then one can bound $\sum_i |\log(s_i)|s_i^{p^*}$ by $K - K(1 - \sum_i s_i^{p^*-\delta})$ which is indeed integrable.

5 Tilted probability measures and a tree with a marked leaf

Recall that $\mathcal{D}$ is the space of càdlàg $\mathcal{P}_\mathbb{N}$-valued functions on $[0,+\infty)$, and that it is endowed with the σ-field generated by all the evaluation functions. For all $t \geqslant 0$, let us introduce the space $\mathcal{D}_t$ of càdlàg functions from $[0,t]$ to $\mathcal{P}_\mathbb{N}$, which we endow with the product σ-field.

As was done in [13], we are going in this section to use the additive martingale to construct a new probability measure under which our fragmentation process has a special tagged fragment such that, heuristically, for all t, the tagged fragment is equal to a block $\Pi_i(t)$ of $\Pi(t)$ with "probability" $|\Pi_i(t)|^{p^*}$. Tagging a fragment will be done by forcing the integer 1 to be in it, and for this we need some additional notation. If π is a partition of $\mathbb{N}$, we let $R\pi$ be its restriction to $\mathbb{N}' = \mathbb{N} \setminus \{1\}$. Partitions of $\mathbb{N}'$ can still be denoted as sequences of blocks ordered with increasing least elements. Given a partition π of $\mathbb{N}'$ and any integer i, we let $H_i(\pi)$ be the partition of $\mathbb{N}$ obtained by inserting 1 in the i-th block of π. Similarly, let us also define a way to insert the integer 1 in a finite-time fragmentation process with state space the partitions of $\mathbb{N}'$. Let $i \in \mathbb{N}$, $t \geqslant 0$ and let $\big(\pi(s)\big)_{s \leqslant t}$ be a family of partitions of $\mathbb{N}'$. Now let j be any element of $\pi_i(t)$ (if this block is empty then the choice won't matter, one can just define $H_i^t(\pi)$ to be any fixed process) and, for all $0 \leqslant s \leqslant t$, let $H_i^t(\pi)(s)$ be the partition which is the same as $\pi(s)$, except that 1 is added to the block containing j. This defines a function H_i^t which maps a process taking values in $\mathcal{P}_{\mathbb{N}'}$ to processes taking value in $\mathcal{P}_{\mathbb{N}}$. What is important to note is that, if we now take $\big(\pi(s)\big)_{0\leqslant s\leqslant t} \in D_t$, then the process $\big(H_i^t R(\pi)(s)\big)_{0\leqslant s\leqslant t}$ is càdlàg (because the restrictions to finite subsets of $\mathbb{N}$ are pure-jump with finite numbers of jumps) and the map $H_i^t R$ from $\mathcal{D}_t$ to itself is also measurable (because, for all s, $H_i^t(\pi)(s)$ is a measurable function of $\pi(s)$ and $\pi(t)$).

5.1 Tilting the measure of a single partition

Here, we are going to work in a simple setting: we consider a random exchangeable partition of $\mathbb{N}$ called Π which has a positive Malthusian exponent p^*, in the sense that $\mathbb{E}\big[\sum_i |\Pi_i|^{p^*}\big] = 1$. Note that this implies that $\mathbb{E}\big[|\Pi_1|^{p^*-1}\mathbb{1}_{|\Pi_1|\neq 0}\big] = 1$ as well (we will omit the indicator function from now on).

Let us define two new random partitions Π^* and Π' through their distributions: we let, for nonnegative measurable functions f on $\mathcal{P}_{\mathbb{N}}$,

$$\mathbb{E}^*[f(\Pi^*)] = \mathbb{E}\left[\sum_i |R\Pi_i|^{p^*} f(H_i R\Pi)\right]$$

and

$$\mathbb{E}'[f(\Pi')] = \mathbb{E}\Big[|\Pi_1|^{p^*-1} f(\Pi)\Big].$$

These relations do define probability measures because p^* is the Malthusian exponent of Π, as can be checked by taking $f = 1$. We now state a few properties of these distributions.

Proposition 5.1. *(i) The two random partitions Π^* and Π' have the same distribution.*

(ii) If we call m the law of the asymptotic frequencies of the blocks of Π, and m' the law of the asymptotic frequencies of the blocks of Π', we have

$$m'(d\mathbf{s}) = \Big(\sum_i s_i^{p^*}\Big) m(d\mathbf{s}).$$

In particular, with probability 1, Π' is not the partition made uniquely of singletons.

(iii) Conditionally on the asymptotic frequencies of its blocks, the law of Π' (or Π^*) can be described as follows: the restriction of the partition to $\mathbb{N}'$ is built with a standard paintbox process from the law m'. Then, conditionally on $R\Pi'$, for every integer i, 1 is inserted in the block $R\Pi_i'$ with probability $\frac{|R\Pi_i'|^{p^*}}{\sum_j |R\Pi_j'|^{p^*}}$.

Proof. Item (i) is a simple consequence of the paintbox description of Π: we know that, conditionally on the restriction of Π to $\mathbb{N}'$, the integer 1 will be inserted in one of these blocks in a size-biased manner. Thus we get, for nonnegative measurable f,

$$\mathbb{E}'[f(\Pi')] = \mathbb{E}\big[|\Pi_1|^{p^*-1} f(\Pi)\big] = \mathbb{E}\Big[\sum_i |R\Pi_i||R\Pi_i|^{p^*-1} f(H_i R\Pi)\Big] = \mathbb{E}^*[f(\Pi^*)].$$

To prove (ii), we just need to use the definition of the law of Π': take any positive measurable function f on $\mathcal{S}^{\downarrow}$, we have

$$\mathbb{E}'[f(|\Pi'|^{\downarrow})] = \mathbb{E}[|\Pi_1|^{p^*-1} f(|\Pi|^{\downarrow})] = \int_{\mathcal{S}^{\downarrow}} (\sum_i s_i s_i^{p^*-1}) f(\mathbf{s}) m(d\mathbf{s}) = \int_{\mathcal{S}^{\downarrow}} (\sum_i s_i^{p^*}) f(\mathbf{s}) m(d\mathbf{s}),$$

which is all we need.

For (iii), first use the definition of Π^* to notice that its restriction to $\mathbb{N}'$ is exchangeable: if we take a measurable function f on $\mathcal{P}_{\mathbb{N}'}$ and a permutation σ of $\mathbb{N}'$, we have

$$\begin{aligned}\mathbb{E}^*[f(\sigma(R\Pi^*))] &= \mathbb{E}\Big[(\sum_i |R\Pi_i|^{p^*})f(\sigma(R\Pi))\Big] \\ &= \mathbb{E}\Big[(\sum_i |\sigma R\Pi_i|^{p^*})f(\sigma(R\Pi))\Big] \\ &= \mathbb{E}\Big[(\sum_i |R\Pi_i|^{p^*})f(R\Pi)\Big] \\ &= \mathbb{E}^*[f(R\Pi^*)].\end{aligned}$$

This exchangeability and Kingman's theorem then imply that the restriction of Π^* to $\mathbb{N}'$ can indeed be built with a paintbox process. Now we only need to identify which block contains 1, that is, find the distribution of Π^* conditionally of $R\Pi^*$. Thus, we take a nonegative measurable function f on $\mathcal{P}_{\mathbb{N}}$ and another one g on $\mathcal{P}_{\mathbb{N}'}$ and compute $\mathbb{E}^*[f(\Pi^*)g(R\Pi^*)]$:

$$\begin{aligned}\mathbb{E}^*[f(\Pi^*)g(R\Pi^*)] &= \mathbb{E}\Big[\sum_i |(R\Pi)_i|^{p^*} f(H_i R\Pi)g(R\Pi)\Big] \\ &= \mathbb{E}\left[\sum_j |(R\Pi)_j|^{p^*}\left(\sum_i \frac{|(R\Pi)_i|^{p^*}}{\sum_j |(R\Pi)_j|^{p^*}} f(H_i R\Pi)\right) g(R\Pi)\right] \\ &= \mathbb{E}^*\left[\left(\sum_i \frac{|(R\Pi^*)_i|^{p^*}}{\sum_j |(R\Pi^*)_j|^{p^*}} f(H_i R\Pi^*)\right) g(R\Pi^*)\right].\end{aligned}$$

This ends the proof. □

5.2 Tilting a fragmentation process

Here we aim to generalize the previous procedure to a homogeneous exchangeable fragmentation process. Let $t \geqslant 0$, we are going to define two random processes $\big(\Pi^*(s)\big)_{s\leqslant t}$ and $\big(\Pi'(s)\big)_{s\leqslant t}$, with corresponding expectation operators $\mathbb{E}^*_t$ and $\mathbb{E}'_t$, by letting, for measurable functions F on $\mathcal{D}_t$,

$$\mathbb{E}^*_t\big[F((\Pi^*(s))_{s\leqslant t})\big] = \mathbb{E}\Big[\sum_i |(R\Pi(t))_i|^{p^*} F\big((H_i^t R\Pi(s))_{s\leqslant t}\big)\Big]$$

and

$$\mathbb{E}'_t\big[F((\Pi'(s))_{s\leqslant t})\big] = \mathbb{E}\Big[|\Pi_1(t)|^{p^*-1} F\big((\Pi(s))_{s\leqslant t}\big)\Big].$$

For the same reason as before, these define probability measures. We then want to use Kolmogorov's consistency theorem to extend these two probability measures to $\mathcal{D}$. To do this we have to check that, if $u < t$ and $(\Pi^*(s))_{s\leqslant t}$ has law $\mathbb{P}^*_t$, then $(\Pi^*(s))_{s\leqslant u}$ has law $\mathbb{P}^*_u$, and the same for Π'. The argument is that the block of $\Pi^*(t)$ that 1 is inserted in only matters through its ancestor at time u: if i and j are such that $(R\Pi^*(t))_j \subset (R\Pi^*(u))_i$ then $(H_j^t\Pi(s))_{s\leqslant u} = (H_i^u\Pi(s))_{s\leqslant u}$.

Taking any nonnegative measurable function F on $\mathcal{D}$, we have

$$\begin{aligned}
\mathbb{E}_t^*\big[F\big((\Pi^*(s))_{s\leqslant u}\big)\big] &= \mathbb{E}\Big[\sum_j |(R\Pi(t))_j|^{p^*} F\big((H_j^t\Pi(s))_{s\leqslant u}\big)\Big] \\
&= \mathbb{E}\Big[\sum_i \sum_{j:(R\Pi^*(t))_j\subset(R\Pi^*(u))_i} |(R\Pi(t))_j|^{p^*} F\big((H_i^u\Pi(s))_{s\leqslant u}\big)\Big] \\
&= \mathbb{E}\Big[\sum_i F\big((H_i^u\Pi(s))_{s\leqslant u}\big) \sum_{j:(R\Pi^*(t))_j\subset(R\Pi^*(u))_i} |(R\Pi(t))_j|^{p^*}\Big] \\
&= \mathbb{E}\Big[\sum_i F\big((H_i^u\Pi(s))_{s\leqslant u}\big)|(R\Pi(u))_i|^{p^*}\Big].
\end{aligned}$$

The last equation comes from the martingale property of the additive martingale $M_{k,u}$ where k is any integer in $\big(R\Pi(u)\big)_i$. Consistency for Π' is a little bit simpler: it is once again a consequence of the fact that the process $(M'_t)_{t\geqslant 0}$ which we define by $M'_t = |\Pi_1(t)|^{p^*-1}\mathbb{1}_{\{|\Pi_1(t)|\neq 0\}}$ for all t is a martingale, which itself is an immediate consequence of the homogeneous fragmentation property.

Kolmogorov's consistency theorem then implies that there exist two random processes $\big(\Pi^*(t)\big)_{t\geqslant 0}$ and $\big(\Pi'(t)\big)_{t\geqslant 0}$ defined on probability spaces with probability measures $\mathbb{P}^*$ and $\mathbb{P}'$ and expectation operators $\mathbb{E}^*$ and $\mathbb{E}'$ such that, for any $t \geqslant 0$ and any nonnegative measurable function F on $\mathcal{D}_t$,

$$\mathbb{E}^*\big[F\big((\Pi^*(s))_{s\leqslant t}\big)\big] = \mathbb{E}\Big[\sum_i |(R\Pi(t))_i|^{p^*} F\big((H_i^t\Pi(s))_{s\leqslant t}\big)\Big] \tag{2.2}$$

and

$$\mathbb{E}'[F\big((\Pi'(s))_{s\leqslant t}\big)] = \mathbb{E}\big[|\Pi_1(t)|^{p^*-1}F\big((\Pi(s))_{s\leqslant t}\big)\big].$$

Just as in the previous section, these two definitions are in fact equivalent:

Proposition 5.2. *The two processes* $(\Pi^*(t))_{t\geqslant 0}$ *and* $(\Pi'(t))_{t\geqslant 0}$ *have the same law.*

To prove this, we only need to show that these two processes have the same finite-dimensional marginal distributions. The 1-dimensional marginals have already been proven to be the same and we will continue with an induction argument which uses the fact that the homogeneous fragmentation property generalizes to $\mathbb{P}^*$ and $\mathbb{P}'$.

Lemma 5.3. *Let* $t \geqslant 0$*, and* Ψ^* *and* Ψ' *be independent copies of respectively* Π^* *and* Π'*. Then, conditionally on* $(\Pi^*(s), s \leqslant t)$*, the process* $\big(\Pi^*(t+s)\big)_{s\geqslant 0}$ *has the same law as* $\big(\Pi^*(t)\cap\Psi^*(s)\big)_{s\geqslant 0}$ *and, conditionally on* $(\Pi'(s), s \leqslant t)$*, the process* $\big(\Pi'(t+s)\big)_{s\geqslant 0}$ *has the same law as* $(\Pi'(t)\cap\Psi'(s))_{s\geqslant 0}$*.*

Proof. Let $t \geqslant 0$ and $u \geqslant 0$, let F be a nonnegative measurable function on $\mathcal{D}_t$ and G be a nonnegative measurable function on $\mathcal{D}_u$. We have, by the fragmentation property,

$$\begin{aligned}
&\mathbb{E}^*\Big[F\big((\Pi^*(s))_{0\leqslant s\leqslant t}\big)G\big(\big(\Pi^*(t+s)\big)_{0\leqslant s\leqslant u}\big)\Big] \\
&\quad= \mathbb{E}\left[\sum_i |(R\Pi(t+u))_i|^{p^*} F\Big(\big(H_i^{t+u}R\Pi(s)\big)_{0\leqslant s\leqslant t}\Big)G\Big(\big(H_i^{t+u}R\Pi(t+s))_{0\leqslant s\leqslant u}\Big)\right] \\
&\quad= \mathbb{E}\left[\sum_i |R(\Pi(t)\cap\Psi(u))_i|^{p^*} F\Big(\big(H_i^{t+u}R\Psi\big)_{0\leqslant s\leqslant t}\Big)G\Big(H_i^{t+u}\big(R\Pi(t)\cap\Psi(s)\big)_{0\leqslant s\leqslant u}\Big)\right],
\end{aligned}$$

where Ψ is an independent copy of Π. The key now is to notice that a block of $\Pi(t) \cap \Psi(s)$ is the intersection of a block of $\Pi(t)$ and a block of $\Psi(s)$. Thus we replace our sum over integers (representing blocks of $\Pi(t) \cap \Psi(s)$) by two sums, one for the blocks of $\Pi(t)$ and another for those of $\Psi(s)$.

$$\begin{aligned}
&\mathbb{E}^*\Big[F\big((\Pi(s))_{0\leqslant s\leqslant t}\big)G\big((\Pi(t+s))_{0\leqslant s\leqslant u}\big)\Big] \\
&\quad= \mathbb{E}\left[\sum_i\sum_j |R\Pi_i(t)|^{p^*}|R\Psi_j(t)^{p^*}|F\big((H_i^t R\Pi(s))_{0\leqslant s\leqslant t}\big)G\Big(\big(H_i^t R\Pi(t)\cap H_j^u\Psi(s)\big)_{0\leqslant s\leqslant u}\Big)\right] \\
&\quad= \mathbb{E}^*\Big[F\big((\Pi(s))_{0\leqslant s\leqslant t}\big)G\big((\Pi(t)\cap\Psi(s))_{0\leqslant s\leqslant u}\big)\Big].
\end{aligned}$$

The proof for Π' again uses the same ideas but is simpler, so we will omit it. □

We can now complete the proof of Proposition 5.2: we show by induction that the finite-dimensional marginals of Π^* and Π' have the same distribution. Take an integer n, let $t_1 < t_2 < \ldots < t_{n+1}$ and assume that we have shown that $\big(\Pi^*(t_1), \ldots, \Pi^*(t_n)\big)$ and $\big(\Pi'(t_1), \ldots, \Pi'(t_n)\big)$ have the same law. Let Ψ be an independent copy of $\Pi^*(t_{n+1} - t_n)$ (which is then also an independent copy of $\Pi'(t_{n+1} - t_n)$), then

$$\begin{aligned}
\big(\Pi^*(t_1), \ldots, \Pi^*(t_{n+1})\big) &\overset{(d)}{=} \big(\Pi^*(t_1), \ldots, \Pi^*(t_n), \Pi^*(t_n)\cap\Psi\big) \\
&\overset{(d)}{=} \big(\Pi'(t_1), \ldots, \Pi'(t_{n+1})\big),
\end{aligned}$$

and the proof is complete. □

We can now proceed to the main part of this section, which is the description of Π^* with Poisson point processes. First, we let κ_ν^* be the measure on $\mathcal{P}_{\mathbb{N}}$ defined by

$$\kappa_\nu^*(\mathrm{d}\pi) = |\pi_1|^{p^*-1}\mathbb{1}_{\{|\pi_1|\neq 0\}}d\kappa_\nu(\mathrm{d}\pi),$$

where κ_ν is as in section 2.1.3 (a paintbox procedure where the asymptotic frequencies have distribution ν).

Let $\big(\Delta^1(t)\big)_{t\geqslant 0}$ be a P.p.p. with intensity k_ν^* and, for all $k \geqslant 2$, $\big(\Delta^k(t)\big)_{t\geqslant 0}$ a P.p.p. with intensity κ_ν. Let also $T_2, T_3, \ldots$ be exponential variables with parameter c (note that there is no T_1 in here). We assume that these variables are all independent. With these, we can create a $\mathcal{P}_{\mathbb{N}}$-valued process Π^*, just as is done in the case of classical fragmentation processes. We start with $\Pi^*(0) = (\mathbb{N}, \emptyset, \emptyset, \ldots)$. For every t such that there is an atom $\Delta^k(t)$, we let $\Pi^*(t)$ be equal to $\Pi^*(t^-)$, except that we replace the block $\Pi_k^*(t^-)$ by its intersection with all the blocks of $\Delta^k(t)$. Also, for every i, we let $\Pi^*(T_i)$ be equal to $\Pi^*(T_i^-)$, except that the integer i is removed from its block and placed into a singleton. Just as in the classical case, it might not be clear that this is well-defined. To make sure that it is the case, we are going to restrict this to finite subsets of $\mathbb{N}$. Let $n \in \mathbb{N}$, we now only need to look at integers $k \leqslant n$ and times t such that $\Delta^k(t)$ splits $[n]$ into at least two blocks. Conveniently enough, this set is in fact finite: indeed, we have

$$\kappa_\nu(\{[n] \text{ is split into two or more blocks}\}) = \int_{\mathcal{S}^\downarrow}(1 - \sum_{i=1}^{\infty} s_i^n)\nu(\mathrm{ds}) \leqslant \int_{\mathcal{S}^\downarrow}(1 - s_1^n)\nu(\mathrm{ds}) < \infty,$$

as well as

$$\begin{aligned}
\kappa_\nu^*\big(\{[n] \text{ is split into two or more blocks}\}\big) &= \int_{\mathcal{S}^\downarrow}(1-\sum_{i=1}^\infty s_i^n)\sum_{i=1}^\infty s_i^{p^*}\nu(\mathrm{d}\mathbf{s}) \\
&= cp^* + \int_{\mathcal{S}^\downarrow}(1-\sum_{i=1}^\infty s_i^n \sum_{i=1}^\infty s_i^{p^*})\nu(\mathrm{d}\mathbf{s}) \\
&\leqslant \int_{\mathcal{S}^\downarrow}(1-s_1^{p^*+n})\nu(\mathrm{d}\mathbf{s}) \\
&< \infty.
\end{aligned}$$

Since the set $(T_2,\dots,T_n)$ is also finite, the previous operations can be applied without ambiguity. From this, we get, for all t, a sequence $\big(\Pi^*(t)\cap[n]\big)_{n\in\mathbb{N}}$ of compatible partitions, which determine a unique partition $\Pi^*(t)$ of $\mathbb{N}$.

Theorem 5.4. *The process $\big(\Pi^*(t)\big)_{t\geqslant 0}$ constructed does have the distribution defined by 2.2.*

Proof. We start by extending the measure $\mathbb{P}'$, so that it contains not only the fragmentation process, but also the underlying Poisson point processes and exponential variables: for $t\leqslant 0$, and any nonegative measurable function F, let

$$\mathbb{E}'_t\Big[F\big((\Delta^i(t)_{s\leqslant t})_{i\in\mathbb{N}},(T_i)_{i\in\mathbb{N}'}\big)\Big] = \mathbb{E}\Big[|\Pi_1(t)|^{p^*-1}F\big((\Delta^i(t)_{s\leqslant t})_{i\in\mathbb{N}},(T_i)_{i\in\mathbb{N}'}\big)\Big]$$

(remember that, under $\mathbb{P}$, $(\Delta^1(t)_{s\leqslant t})$ is a P.p.p. with intensity k_ν, and not k_ν^*.) These probability measures are still compatible, and we can still use Kolmogorov's theorem to extend them to a single measure $\mathbb{P}'$. Note that under $\mathbb{P}'$, 1 never falls in a singleton, which is why we have ignored T_1. With this new law $\mathbb{P}'$, the partition-valued process $\big(\Pi'(t)\big)_{t\geqslant 0}$ is indeed built from the point processes $\big(\Delta^k(t)\big)_{t\geqslant 0}$ with $k\in\mathbb{N}$ and the T_i with $i\in\mathbb{N}$, and all we need to do is now find their joint distribution. We start with the harder part, which is finding the law of $\big(\Delta^1(t)\big)_{t\geqslant 0}$, and will use a Laplace transform method and the exponential formula for Poisson point processes. If $t\geqslant 0$ and f is a nonnegative measurable function on $\mathcal{P}_\mathbb{N}\times\mathbb{R}$, we have

$$\begin{aligned}
\mathbb{E}'[\mathrm{e}^{-\sum_{s\leqslant t} f(\Delta^1_s,s)}] &= \mathbb{E}\big[|\Pi_1(t)|^{p^*-1}\mathbb{1}_{\{|\Pi_1(t)|\neq 0\}}e^{-\sum_{s\leqslant t} f(\Delta^1_s,s)}\big] \\
&= \mathrm{e}^{-ct}\mathrm{e}^{-ct(p^*-1)}\mathbb{E}\Big[\prod_{s\leqslant t}|\Delta^1_1(s)|^{p^*-1}\mathbb{1}_{|\Delta^1_1(s)|\neq 0}e^{-\sum_{s\leqslant t} f(\Delta^1(s),s)}\Big] \\
&= \mathrm{e}^{-ctp^*}\mathbb{E}\left[\exp\left(-\sum_{s\leqslant t}(-(p^*-1)\log(|\Delta^1_1(s)|)+f(\Delta^1(s),s))\right)\right] \\
&= \mathrm{e}^{-ctp^*}\exp\left(-\int_0^t\int_{\mathcal{P}_\mathbb{N}}(1-\mathrm{e}^{-(-(p^*-1)\log(|\pi_1|)+f(\pi,s))})\kappa_\nu(\mathrm{d}\pi)\mathrm{d}s\right) \\
&= \mathrm{e}^{-ctp^*}\exp\left(-\int_0^t\int_{\mathcal{P}_\mathbb{N}}(1-|\pi_1|^{p^*-1}e^{-f(\pi,s)})\kappa_\nu(\mathrm{d}\pi)\mathrm{d}s\right)
\end{aligned}$$

Now we use the the Malthusian hypothesis: we have $cp^*+\int_{\mathcal{S}^\downarrow}(1-\sum s_i^{p^*})\mathrm{d}\nu(\mathbf{s})=0$. Translating this in terms of k_ν, we have

$$\begin{aligned}
\int_{\mathcal{P}_\mathbb{N}}(1-|\pi_1|^{p^*-1})\kappa_\nu(\mathrm{d}\pi) &= \int_{\mathcal{S}^\downarrow}\big(\sum_i s_i(1-s_i^{p^*-1})+s_0\big)\nu(\mathrm{d}\mathbf{s}) \\
&= -cp^*.
\end{aligned}$$

Thus, in the last integral with respect to k_ν, we can replace 1 by $|\pi_1|^{p^*-1}$, if we subtract cp^* outside of the integral:

$$\begin{aligned}\mathbb{E}'[\mathrm{e}^{-\sum_{s\leqslant t} f(\Delta^1_s,s)}] &= \mathrm{e}^{-ctp^*}\exp\left(-\int_0^t (-cp^* + \int_{\mathcal{P}_\mathbb{N}} (|\pi_1|^{p^*-1} - |\pi_1|^{p^*-1}e^{-f(\pi,s)})\kappa_\nu(\mathrm{d}\pi))\mathrm{d}s\right) \\ &= \exp\left(-\int_0^t\int_{\mathcal{P}_\mathbb{N}} (1-e^{-f(\pi,s)})|\pi_1|^{p^*-1}\kappa_\nu(\mathrm{d}\pi)\mathrm{d}s\right).\end{aligned}$$

This means that the point process $\big(\Delta^1(t)\big)_{t\geqslant 0}$ does indeed have the law of a Poisson point process with intensity $|\pi_1|^{p^*-1}d\kappa_\nu(\pi)$.

Let us now prove that the point processes and random variables are independent from each other and that, except for $\big(\Delta^1_t\big)_{t\geqslant 0}$, they have the same law as under $\mathbb{P}$. Take $n\in\mathbb{N}$ and $t\geqslant 0$, for every $i\in[n]$, F_i a nonnegative measurable function on the space of random measures on $\mathcal{P}_\mathbb{N}\times[0,t]$, and for $2\leqslant i\leqslant n$, a nonnegative measurable function g_i on $\mathbb{R}$. Using independence properties under $\mathbb{P}$, we have

$$\begin{aligned}&\mathbb{E}'\left[\prod_{i=1}^n F_i((\Delta^i(s))_{s\leqslant t})\prod_{i=2}^n g_i(T_i)\right] \\ &= \mathbb{E}\left[\prod_{s\leqslant t}|\Delta^1_1(s)|^{p^*-1}\mathbb{1}_{|\Delta^1_1(s)|\neq 0}F_1\left(\big(\Delta^1(s)\big)_{s\leqslant t}\right)\prod_{i=2}^n F_i((\Delta^i(s))_{s\leqslant t})g_i(T_i)\right] \\ &= \mathbb{E}\left[\prod_{s\leqslant t}|\Delta^1_1(s)|^{p^*-1}\mathbb{1}_{|\Delta^1_1(s)|\neq 0}F_1((\Delta^1(s))_{s\leqslant t})\right]\prod_{i=2}^n\mathbb{E}\big[F_i((\Delta^i(s))_{s\leqslant t})\big]\prod_{i=2}^n\mathbb{E}[g_i(T_i)] \\ &= \mathbb{E}'\big[F_1\big((\Delta^1(s)_{s\leqslant t})\big)\big]\prod_{i=2}^n\mathbb{E}\big[F_i((\Delta^i(s)_{s\leqslant t}))\big]\prod_{i=2}^n\mathbb{E}[g_i(T_i)],\end{aligned}$$

which is all we need. □

Remark 5.5. *Here is an alternative description of a Poisson point process $\big(\Delta^1(t)\big)_{t\geqslant 0}$ with intensity k^*_ν. Let $\big(s(t),i(t)\big)_{t\geqslant 0}$ be a $\mathcal{S}^\downarrow\times\mathbb{N}$-valued Poisson point process with intensity $s_i^{p^*}\nu(\mathrm{d}\mathbf{s})\#(\mathrm{d}i)$, where $\#$ is the counting measure on $\mathbb{N}$ (otherwise said, $\big(s(t)\big)_{t\geqslant 0}$ has intensity $\sum_i s_i^{p^*}\nu(\mathrm{d}\mathbf{s})$ and $i(t)$ is equal to an integer j with probability $\frac{s_j^{p^*}}{\sum_i s_i^{p^*}}$). When there is an atom, construct a partition of $\mathbb{N}'$ using the paintbox method (using for example a coupled process of uniform variables), and then add 1 to the $i(t)$-th block, where the blocks are ordered in decreasing order of their asymptotic frequencies.*

5.3 Link between μ^* and $\mathbb{P}^*$.

Let $\mathcal{T}$ be the fragmentation tree derived from Π^α, equipped with its list of death points $(Q_i)_{i\in\mathbb{N}}$, as well as the measure μ^* which has total mass W, and we keep the assumption that $\mathbb{E}[W]=1$. Given any leaf L, we can build a new partition process $(\Pi^\alpha_L(t))_{t\geqslant 0}$ from this, by declaring the "new death point" of 1 to be L. More precisely, for all $t\geqslant 0$, the restriction of $\Pi^\alpha_L(t)$ to $\mathbb{N}'$ is the same as that of $\Pi^\alpha(t)$, while 1 is put in the block containing all the integers j such that Q_j is in the same tree component of $\mathcal{T}_{>t}$ as L. As in the proof of Proposition 3.2, one can show that

Π_L^α is decreasing and in $\mathcal{D}$. Our main result here is that, if L is chosen with "distribution" μ^*, then Π_L^α has the same distribution as the $\Pi^{*,\alpha}$, where $\Pi^{*,\alpha}$ is the "α-self-similar" version of Π^*, obtained through the usual time-change.

Proposition 5.6. *Let F be any nonnegative measurable function of $\mathcal{D}$, then $\int_{\mathcal{T}} F(\Pi_L^\alpha)\mu^*(\mathrm{d}L)$ is a random variable and we have*

$$\mathbb{E}\left[\int_{\mathcal{T}} F(\Pi_L^\alpha)\mu^*(\mathrm{d}L)\right] = \mathbb{E}^*[F(\Pi^{*,\alpha})].$$

Proof. For any leaf L of $\mathcal{T}$, we let $\Pi_L = G^{-\alpha}(\Pi_L^\alpha)$, then $\Pi_L^\alpha = G^\alpha(\Pi_L)$ (recall from Section 2.1.3 that G^α and $G^{-\alpha}$ are the measurable functions which transform Π to Π^α and back). By renaming, we are reduced to proving that, for any nonnegative measurable function F on $\mathcal{D}$, $\int_{\mathcal{T}} F(\Pi_L)\mu^*(\mathrm{d}L)$ is a random variable and

$$\mathbb{E}\left[\int_{\mathcal{T}} F(\Pi_L)\mu^*(\mathrm{d}L)\right] = \mathbb{E}^*[F(\Pi^*)].$$

We let $M(F) = \int_{\mathcal{T}} F(\Pi_L)\mu^*(\mathrm{d}L)$. Assume first that F is of the form $F\big((\pi(s))_{s\geqslant 0}\big) = K\big((\pi(s))_{0\leqslant s\leqslant t}\big)$, for a certain $t \geqslant 0$ and a function K on $\mathcal{D}_t$. We then have, by definition of μ^*,

$$M(F) = \sum_i |R\Pi_i(t)|^{p^*} X_i K((H_i^t(R\Pi)(s))_{0\leqslant s\leqslant t}),$$

where X_i is defined for all i by $X_i = \frac{W_{j,t}}{|R\Pi_i(t)|^{p^*}}$ for any choice of $j \in \Pi_i(t)$, so X_i has the same law as W and is independent of $\big(\Pi(s)\big)_{s\leqslant t}$. We thus know that $M(F)$ is a random variable such that

$$\mathbb{E}[M(F)] = \mathbb{E}[W]\mathbb{E}\left[\sum_i |R\Pi_i(t)|^{p^*} K\Big((H_i^t(R\Pi)(s))_{0\leqslant s\leqslant t}\Big)\right] = \mathbb{E}^*[F(\Pi^*)].$$

A measure theory argument then extends this to any nonnegative measurable function F. Let $\mathcal{A}$ be the set of measurable subsets $A \in \mathcal{D}$ such that $M(\mathbb{1}_A)$ is a random variable and $\mathbb{E}[M(\mathbb{1}_A)] = \mathbb{P}^*[\Pi^* \in A]$. Standard properties of integrals show that $\mathcal{A}$ is a monotone class, and since it contains the generating π-system of sets of the form $A = \{\pi \in \mathcal{D},\ (\pi(s))_{0\leqslant s\leqslant t} \in B\}$ with $t \geqslant 0$ and $B \subset \mathcal{D}_t$, the monotone class theorem implies that $\mathcal{A}$ is $\mathcal{D}$'s Borel σ-field. We then conclude by approximating F by linear combinations of indicator functions.

□

5.4 Marking two points

We now want to go further and mark two points of $\mathcal{T}$ with distribution μ^*. However, in order to avoid having to manipulate partitions with both integers 1 and 2 being forced into certain blocks, we will instead work with the tree $\mathcal{T}^* = \mathrm{TREE}(\Pi^{*,\alpha})$. To make sure that this is properly defined, we need to check that $\Pi^{*,\alpha}$ satisfies the hypotheses of Lemmas 3.9 and 3.11. The first one is immediate because, for all $t \geqslant 0$, when restricted to the complement of $\Pi_1^{*,\alpha}(t)$, $\big(\Pi^{*,\alpha}(s)\big)_{s\geqslant t}$ is an α-self-similar fragmentation process, while the second one comes from the Poissonian construction.

Let us give an alternate description of $\mathcal{T}^*$ which we will use here. Let $\big(\Delta(t)\big)_{t\geqslant 0}$ be a Poisson point process with intensity measure κ_ν^*, and, for all $t \geqslant 0$, $\xi(t) = \mathrm{e}^{-ct}\prod_{s\leqslant t}|\Delta(s)|$. From this we define the usual time-change: for all $t \geqslant 0, \tau(t) = \inf\{u, \int_0^u \xi(t)^{-\alpha} dr > t\}$. The tree $\mathcal{T}^*$ is then made of a *spine* of length $T=\tau^{-1}(\infty)$ on which we have attached many small independent copies of $\mathcal{T}$. More precisely, for each t such that $(\Delta(s))_{s\geqslant 0}$ has an atom at time $\tau(t)$, we graft on the spine at height t a number of trees equal to the number of blocks of $\Delta(t)$ minus one (an infinite amount if Δ_t has infinitely many). These are indexed by $j\geqslant 2$ and, for every such j, we graft precisely a copy of $\Big((\xi(t^-)|\Delta_j(t)|)^{-\alpha}\mathcal{T},(\xi(t^-)|\Delta_j(t)|)\mu\Big)$, which will be called $(\mathcal{T}'_{j,t},\mu'_{j,t})$. All of these then naturally come with their copy of μ^* which we will call $\mu^*_{i,t}$. These can then all be added to obtain a measure μ^{**} on $\mathcal{T}^*$, which satisfies, for all $(i,t)\in\mathcal{T}^*$,

$$\mu^{**}(\mathcal{T}^*_{i,t+})=\lim_{s\to\infty}\sum_{j\in\Pi^*(\tau_i(t)+s)\cap\mathrm{rep}(\Pi^*(\tau_i(t)))}|\Pi^*_j(t+s)|^{p^*}.$$

The measure μ^{**} is the natural analogue of μ^* on the biased tree.

We will need a Gromov-Hausdorff-type metric for *trees with two extra marked points*: let $(\mathcal{T},\rho,d)$ and $(\mathcal{T}',\rho',d')$ be two compact rooted trees, and then let $(x,y)\in\mathcal{T}^2$ and $(x',y')\in(\mathcal{T}')^2$. We now let the 2-pointed Gromov-Hausdorff $d^2_{GH}\big((\mathcal{T},x,y),(\mathcal{T}',x',y')\big)$ be equal to

$$\inf\Big[\max\Big(d_{\mathcal{Z},H}(\phi(\mathcal{T}),\phi'(\mathcal{T}')),d_{\mathcal{Z}}(\phi(\rho),\phi'(\rho')),d_{\mathcal{Z}}(\phi(x),\phi'(x')),d_{\mathcal{Z}}(\phi(y),\phi'(y'))\Big)\Big],$$

where the infimum is once again taken on all possible isometric embeddings ϕ and ϕ' of $\mathcal{T}$ and $\mathcal{T}'$ in a common space $\mathcal{Z}$. Taking classes of such trees up to the relation d^2_{GH}, we then get a Polish space $\mathbb{T}^2$ which is the set of 2-pointed compact trees. For more details in a more general context (pointed metric spaces instead of trees), the reader can refer to [66], Section 6.4.

Proposition 5.7. *Let F be any nonnegative measurable function on $\mathbb{T}^2$. Then $\int_{\mathcal{T}}F(\mathcal{T},L,L')\mu^*(\mathrm{d}L)\mu^*(\mathrm{d}L')$ is a random variable, and we have*

$$\mathbb{E}\left[\int_{\mathcal{T}}\int_{\mathcal{T}}F(\mathcal{T},L,L')\mu^*(\mathrm{d}L)\mu^*(\mathrm{d}L')\right]=\mathbb{E}^*\left[\int_{\mathcal{T}^*}F(\mathcal{T}^*,L_1,L')\mu^{**}(\mathrm{d}L')\right].$$

Proof. As in the proof of Proposition 5.6, we let Π^α_L be the fragmentation-like process obtained by setting the leaf L as the new death point of the integer 1 in $\mathcal{T}$, and then we let Π_L be its homogeneous version. The other leaf L' will be represented by a sequence of integers $\big(j^\alpha_{L'}(t)\big)_{0\leqslant t<ht(L')}$ where, for all t with $0\leqslant t<ht(L')$, $j^\alpha_{L'}(t)$ is the smallest integer $j\neq 1$ such that $(j,t)\leqslant L'$ in $\mathcal{T}^*$. We then let $\big(j_{L'}(t)\big)_{t\geqslant 0}$ we the image of $(j^\alpha_{L'}(t))_{0\leqslant t\leqslant ht(L')}$ through the reverse Lamperti transformation.

Notice that $(\mathcal{T},L,L')$ is the image of $\big(\Pi_L(t),j_{L'}(t)\big)_{t\geqslant 0}$ by a measurable function. Indeed, going back to the representation in ℓ^1 of our trees, $\mathcal{T}$ is no more than TREE(Π^α_L), L_1 is Q_1, while L' is the limit as t goes to infinity of $Q_{j_{L'}(t)}$.

Thus, with some renaming, we now just need to check that, if F is a nonnegative measurable function on the space of $\mathcal{P}_{\mathbb{N}}\times\mathbb{N}$-valued càdlàg functions (equipped with the product σ-algebra generated by the evaluation functions), then $\int_{\mathcal{T}}F\big((\Pi_L(t),j_{L'}(t))_{t\geqslant 0}\big)\mu^*(\mathrm{d}L)\mu^*(\mathrm{d}L')$ is a random variable, and

$$\mathbb{E}\left[\int_{\mathcal{T}}\int_{\mathcal{T}}F\Big(\big(\Pi_L(t),j_{L'}(t)\big)_{t\geqslant 0}\Big)\mu^*(\mathrm{d}L)\mu^*(\mathrm{d}L')\right]=\mathbb{E}^*\left[\int_{\mathcal{T}^*}F\Big(\big(\Pi^*(t),j_{L'}(t)\big)_{t\geqslant 0}\Big)\mu^{**}_{i,t}(\mathrm{d}L')\right].$$

This will be done the same way as before: suppose that F is of the form $K\big((\pi(s),j(s))_{0\leqslant s\leqslant t}\big)$, then one can write

$$\int_{\mathcal{T}}\int_{\mathcal{T}}F\Big(\big(\Pi_L(t),j_{L'}(t)\big)_{t\geqslant 0}\Big)\mu^*(\mathrm{d}L)\mu^*(\mathrm{d}L')=\int_{\mathcal{T}}\sum_j W_{j(t),t}K\Big(\big(\Pi_L(s),j(s)\big)_{0\leqslant s\leqslant t}\Big)\mu^*(\mathrm{d}L).$$

(In the right-hand side, $j(s)$ denotes the smallest element of the block of $\Pi_L(s)$ which contains $(\Pi_L(t))_j$.) By Proposition 5.6, this is a random variable, and we know that its expectation is equal to

$$\mathbb{E}^*\left[\sum_j |\Pi_j^*(t)|^{p^*} K\Big(\big(\Pi^*(s), j(s)\big)_{0\leqslant s\leqslant t}\Big)\right] = \mathbb{E}^*\left[\int_{\mathcal{T}^*} F\Big(\big(\Pi^*(t), j_{L'}(t)\big)_{t\geqslant 0}\Big)\mu^{**}(\mathrm{d}L')\right].$$

A monotone class argument similar to the one at the end of Proposition 5.6 ends the proof.

□

6 The Hausdorff dimension of $\mathcal{T}$

The reader is invited to read [34] for the basics on the Hausdorff dimension $\dim_{\mathcal{H}}$ of a set, which we will not recall here.

6.1 The result

Theorem 6.1. *Assume* (**H**)*, that is that the function ψ takes at least one strictly negative value on $[0,1]$. Then there exists a Malthusian exponent p^* for (c,ν) and, almost surely, on the event that Π does not die in finite time, we have*

$$\dim_{\mathcal{H}}\big(\mathcal{L}(\mathcal{T})\big) = \frac{p^*}{|\alpha|}.$$

If Π does die in finite time, then the leaves of $\mathcal{T}$ form a countable set, which has dimension 0.

The last statement is a consequence of Proposition 3.13: if Π does die in finite time, then there are no proper leaves, which implies that every leaf of $\mathcal{T}$ is the death point of some integer.

6.2 The lower bound

An elaborate use of Frostman's lemma (Theorem 4.13 in [34]) with the measure μ^* combined with a truncation of the tree similar to what was done in [41] will show that $\dim_{\mathcal{H}}\big(\mathcal{L}(\mathcal{T})\big) \geqslant \frac{p^*}{|\alpha|}$ almost surely when Π does not die in finite time.

6.2.1 A first lower bound

Here we assume that $\mathbb{E}[W] = 1$, and thus Π dies in finite time if and only if μ^* is the zero measure. We also assume the integrability condition $\int_{\mathcal{S}^\downarrow}(\sum_i |\log(s_i)| s_i^{p^*})\nu(\mathrm{d}\mathbf{s}) < \infty$ of Lemma 4.11.

Lemma 6.2. *Recall that $\underline{p} = \sup\{q \in \mathbb{R} : \phi^*(-q) > -\infty\}$, and let also*

$$A = \sup\{a \leqslant p^* : \int_{\mathcal{S}^\downarrow} \sum_{i\neq j} s_i^{p^*-a} s_j^{p^*}\nu(\mathrm{d}\mathbf{s}) < \infty,\} \in [0, p^*].$$

On the event where Π does not die in finite time, we have the lower bound:

$$\dim_{\mathcal{H}}(\mathcal{L}(\mathcal{T})) \geqslant \frac{A \wedge (|\alpha| + \underline{p})}{|\alpha|}.$$

Proof. We want to apply Proposition 5.7 to the function F defined on the space $\mathbb{T}^2$ by $F(\mathcal{T},\rho,d,x,y) = d(x,y)^{-\gamma}\mathbb{1}_{x\neq y}$, where $\gamma > 0$. To do this we need to check that it is measurable, which can be done by showing that $d(x,y)$ is continuous. In fact, it is even Lipschitz-continuous: for all $(\mathcal{T},\rho,d,x,y)$ and $(\mathcal{T}',\rho',d',x',y')$ and any embeddings ϕ and ϕ' of $\mathcal{T}$ and $\mathcal{T}'$ in a common $\mathcal{Z}$, we have

$$|d(x,y)-d'(x',y')| = |d_{\mathcal{Z}}(\phi(x),\phi(y)) - d_{\mathcal{Z}}(\phi'(x'),\phi'(y'))| \leqslant d_{\mathcal{Z}}(\phi(x),\phi'(x')) + d_{\mathcal{Z}}(\phi(y),\phi'(y'))$$

and then taking the infimum, we obtain

$$|d(x,y)-d'(x',y')| \leqslant 2d^2_{GH}\big((\mathcal{T},x,y),(\mathcal{T}',x',y')\big).$$

Applying Proposition 5.7 to F, we then get

$$\mathbb{E}\left[\int_{\mathcal{T}}\int_{\mathcal{T}}(d(L,L'))^{-\gamma}\mu^*(\mathrm{d}L)\mu^*(\mathrm{d}L')\right] = \mathbb{E}^*\left[\int_{\mathcal{T}^*}(d(L_1,L'))^{-\gamma}\mu^{**}(\mathrm{d}L')\right].$$

Recall the Poisson description of $\mathcal{T}^*$ of Section 5.4. Let, for all relevant $j \geqslant 2$ and $t \geqslant 0$, $X_{j,t}$ be the root of $T'_{j,t}$ and $Z_{j,t} = \int_{\mathcal{T}'_{j,t}} d(L',X_{k,t})^{-\gamma}\mu^*(\mathrm{d}L')$. One can then write $Z_{j,t} = \big(\xi(t^-)|\Delta_j(t)|\big)^{p^*+\alpha\gamma}(I_{j,t})^{-\gamma}$ where $I_{i,t}$ is a copy of I (defined in the proof of Lemma 4.11) which is independent from the process $(\Delta)_{t\geqslant 0}$ and all the other $\mathcal{T}'_{k,s}$ for $(k,s)\neq(j,t)$. Thus, the process $(\Delta_t,(I_{j,t})_{j\geqslant 2})_{t\geqslant 0}$ is a Poisson point process whose intensity is the product of κ^*_ν and the law of an infinite sequence of i.i.d. copies of I. We then have

$$\begin{aligned}
\mathbb{E}^*\left[\int d(L_1,L')^{\gamma}\mu^*(\mathrm{d}L')\right] &= \mathbb{E}^*\left[\sum_{t\geqslant 0}\sum_{j\geqslant 2}\int_{\mathcal{T}'_{j,t}} d(L_1,L')^{-\gamma}\mu^{**}(\mathrm{d}L')\right] \\
&\leqslant \mathbb{E}^*\left[\sum_{t\geqslant 0}\sum_{j\geqslant 2}\int_{\mathcal{T}'_{j,t}} d(L',X_{i,t})^{-\gamma}\mu^{**}(\mathrm{d}L')\right] \\
&= \mathbb{E}^*\left[\sum_{t\geqslant 0}\sum_{j\geqslant 2}\big(\xi(t^-)|\Delta_j(t)|\big)^{p^*+\alpha\gamma}(I_{j,t})^{-\gamma}\right] \\
&= Q\big[I^{-\gamma}\big]\mathbb{E}^*\Big[\int \xi_{t^-}^{p^*+\alpha\gamma}dt\Big]\int_{\mathcal{S}^\downarrow}\sum_i s_i^{p^*}\sum_{j\neq i}s_j^{p^*+\alpha\gamma}\nu(\mathrm{d}\mathbf{s}).
\end{aligned}$$

The last equality directly comes from the Master Formula for Poisson point processes.

We have a product of three factors, and we want to know when they are finite. The case of the first factor has already been studied in Lemma 4.11, we know that it is finite when $\gamma < 1+\frac{p}{|\alpha|}$. For the second factor to be finite we simply need $\phi^*(p^*+\alpha\gamma) > 0$, which is true as soon as $p^*+\alpha\gamma > 0$ i.e. when $\gamma < \frac{p^*}{|\alpha|}$. Finally, by definition of A, the third factor is finite as soon as $\gamma < \frac{A}{|\alpha|}$. Since $A \leqslant p^*$ by definition, Frostman's lemma implies Lemma 6.2. □

6.2.2 A reduced fragmentation and the corresponding subtree

Let $N\in\mathbb{N}$ and $\varepsilon > 0$, we define a function $G_{N,\varepsilon}$ from $\mathcal{S}^\downarrow$ to $\mathcal{S}^\downarrow$ by

$$G_{N,\varepsilon}(\mathbf{s}) = \begin{cases}(s_1,\ldots,s_N,0,0,\ldots) & \text{if } s_1 \leqslant 1-\varepsilon \\ (s_1,0,0,\ldots) & \text{if } s_1 > 1-\varepsilon.\end{cases}$$

A similar function can be defined on partitions on $\mathcal{P}_{\mathbb{N}}$. If a partition π does not have asymptotic frequencies (a measurable event which doesn't concern us), we let $G_{N,\varepsilon}(\pi) = \pi$. If it does, we first reorder its blocks by decreasing order of their asymptotic frequencies by letting, for all i, $\pi_i^{\downarrow}$ be the block with i-th highest asymptotic frequency (if there is a tie, we just rank those blocks by increasing order of their first elements). Then we let

$$G_{N,\varepsilon}(\pi) = \begin{cases} (\pi_1^{\downarrow}, \ldots, \pi_N^{\downarrow}, \text{singletons}) & \text{if } |\pi_1^{\downarrow}| \leqslant 1-\varepsilon \\ (\pi_1^{\downarrow}, \text{singletons}) & \text{if } |\pi_1^{\downarrow}| > 1-\varepsilon. \end{cases}$$

We let $\nu_{N,\varepsilon}$ be the image of ν by $G_{N,\varepsilon}$. Then the image of k_ν by $G_{N,\varepsilon}$ on $\mathcal{P}_{\mathbb{N}}$ is $k_{\nu_{N,\varepsilon}}$. The following is immediate.

Proposition 6.3. *Let $\big((\Delta_t^k)_{t\geqslant 0}, k \in \mathbb{N}\big)$ be a family of independant Poisson point processes with intensity k_ν, then $\big((G_{N,\varepsilon}(\Delta_t^k))_{t\geqslant 0}, k \in \mathbb{N}\big)$ is a family of independant Poisson point processes with intensity $k_{\nu_{N,\varepsilon}}$. Using them, one gets two coupled fragmentation processes $(\Pi(t))_{t\geqslant 0}$ and $(\Pi^{N,\varepsilon}(t))_{t\geqslant 0}$ such that, for all t, $\Pi^{N,\varepsilon}(t)$ is finer than $\Pi(t)$. Also, $\mathcal{T}_{N,\varepsilon}$, the tree built from $(\Pi^{N,\varepsilon}(t))_{t\geqslant 0}$, is naturally a subset of $\mathcal{T}$.*

6.2.3 Using the reduced fragmentation

Recall the concave function ψ defined from $\mathbb{R}$ to $[-\infty, +\infty)$ by

$$\psi(p) = cp + \int_{\mathcal{S}^{\downarrow}} (1 - \sum_i s_i^p)\nu(\mathrm{d}\mathbf{s}).$$

We now assume (**H**): there exists $p > 0$ such that $-\infty < \psi(p) < 0$.

Proposition 6.4. *For $N \in \mathbb{N} \cup \{\infty\}$, $\varepsilon \in [0,1]$ and $p \in \mathbb{R}$, define the reduced Laplace exponent $\psi_{N,\varepsilon}(p) = cp + \int_{\mathcal{S}^{\downarrow}} (1 - \sum_i s_i^p)\nu_{N,\varepsilon}(\mathrm{d}\mathbf{s})$. One can then write*

$$\psi_{N,\varepsilon}(p) = cp + \int_{\mathcal{S}^{\downarrow}} \left(\Big(1 - \sum_{i=1}^N s_i^p\Big) \mathbb{1}_{\{s_1 \leqslant 1-\varepsilon\}} + (1 - s_1^p)\mathbb{1}_{\{s_1 > 1-\varepsilon\}} \right) \nu(\mathrm{d}\mathbf{s}).$$

(i) This is a nonincreasing function of N and a nondecreasing function of ε.

(ii) We have $\psi(p) = \inf\limits_{N,\varepsilon} \psi_{N,\varepsilon}(p)$.

*(iii) There exist N_0 and ε_0 such that, for $N > N_0$ and $\varepsilon < \varepsilon_0$, the pair $(c, \nu_{N,\varepsilon})$ satisfies (**H**) and has a Malthusian exponent $p^*_{N,\varepsilon}$.*

(iv) We have $p^* = \sup\limits_{N,\varepsilon} p^*_{N,\varepsilon}$.

Proof. The first point is immediate. The second one is a straightforward application of the monotone convergence theorem as N tends to infinity and ε tends to 0, which is valid because we have, for all $\mathbf{s}$, the upper bound

$$\left(1 - \sum_{i=1}^N s_i^p\right) \mathbb{1}_{\{s_1 \leqslant 1-\varepsilon\}} + (1 - s_1^p)\mathbb{1}_{\{s_1 > 1-\varepsilon\}} \leqslant (1 - s_1^p) \leqslant C_p(1 - s_1),$$

and $(1 - s_1)$ is ν-integrable.

The third point is a direct consequence of the second: let $p \in [0,1]$ such that $\psi(p) < 0$, there exist N_0 and ε_0 such that $\psi_{N_0,\varepsilon_0}(p) < 0$. Then by monotonicity, for all $N > N_0$ and $\varepsilon < \varepsilon_0$, $\psi_{N,\varepsilon}(p) < 0$ and thus $\nu_{N,\varepsilon}$ has a Malthusian exponent $p^*_{N,\varepsilon}$.

Now for the last point: first notice that, for all N and ε, we have $\phi_{N,\varepsilon}(p^*) \geqslant \phi(p^*) = 0$ and thus, if it exists, $p^*_{N,\varepsilon}$ is smaller than or equal to p^*. Then, for $p < p^*$, by taking N large enough and ε small enough, we have $\psi_{N,\varepsilon}(p) < 0$ and thus $p^*_{N,\varepsilon} \geqslant p$. This concludes the proof. □

Proposition 6.5. *For all N and ε such that $p^*_{N,\varepsilon}$ exists, and for all $q > 1$, the measure $\nu_{N,\varepsilon}$ satisfies assumption* $(\mathbf{M_q})$*: $\int_{\mathcal{S}^\downarrow} |1 - \sum_{i=1}^{\infty} s_i^{p^*_{N,\varepsilon}}|^q \, \nu_{N,\varepsilon}(\mathrm{d}\mathbf{s}) < \infty$.*

Proof. It is simply a matter of bounding $(1 - \sum_{i=1}^N s_i^{p^*_{N,\varepsilon}})\mathbb{1}_{s_1 \leqslant 1-\varepsilon} + (1 - s_1^{p^*_{N,\varepsilon}})\mathbb{1}_{s_1 > 1-\varepsilon}$ in such a way that both the upper and lower bound's absolute values have an integrable q-th power. For the upper bound, write

$$\left(1 - \sum_{i=1}^N s_i^{p^*_{N,\varepsilon}}\right)\mathbb{1}_{\{s_1 \leqslant 1-\varepsilon\}} + (1 - s_1^{p^*_{N,\varepsilon}})\mathbb{1}_{\{s_1 > 1-\varepsilon\}} \leqslant 1 - s_1^{p^*_{N,\varepsilon}} \leqslant C_{p^*_{N,\varepsilon}}(1 - s_1)$$

and since $q > 1$, we can bound $(1 - s_1)^q$ by $1 - s_1$ which is integrable. For the lower bound, write

$$\left(1 - \sum_{i=1}^N s_i^{p^*_{N,\varepsilon}}\right)\mathbb{1}_{\{s_1 \leqslant 1-\varepsilon\}} + (1 - s_1^{p^*_{N,\varepsilon}})\mathbb{1}_{\{s_1 > 1-\varepsilon\}} > (1 - N)\mathbb{1}_{\{s_1 \leqslant 1-\varepsilon\}}$$

and then note that, since ν integrates $1 - s_1$, the set $\{s_1 \leqslant 1 - \varepsilon\}$ has finite measure. □

Proposition 6.6. *Let N, ε be such that $p^*_{N,\varepsilon}$ exists. Let then $A_{N,\varepsilon}$ and $\underline{p}_{N,\varepsilon}$ corresponding quantities to A and $\underline{p}$ (see Lemma 6.2), replacing ν by $\nu_{N,\varepsilon}$. Then $A_{N,\varepsilon} = p^*_{N,\varepsilon}$ and $\underline{p}_{N,\varepsilon} \geqslant p^*_{N,\varepsilon}$.*

Proof. The important fact to note here is that, since $1 - s_1$ is integrable with respect to ν, we have $\nu(\{s_1 \leqslant 1 - \varepsilon\}) < \infty$. Now notice that, for all $p < p^*_{N,\varepsilon}$, we have

$$\begin{aligned}
\int_{\mathcal{S}^\downarrow} \sum_{i=1}^{\infty} (1 - s_i^{-p}) s_i^{p^*_{N,\varepsilon}} \nu_{N,\varepsilon}(\mathrm{d}\mathbf{s}) &= c p^*_{N,\varepsilon} + \int_{\mathcal{S}^\downarrow} \left(1 - \sum_{i=1}^{\infty} s_i^{p^*_{N,\varepsilon} - p}\right) \nu_{N,\varepsilon}(\mathrm{d}\mathbf{s}) \\
&= c p^*_{N,\varepsilon} + \int_{\mathcal{S}^\downarrow} \left(1 - \sum_{i=1}^{N} s_i^{p^*_{N,\varepsilon} - p}\mathbb{1}_{\{s_1 \leqslant 1-\varepsilon\}} + (1 - s_1^{p^*_{N,\varepsilon} - p})\mathbb{1}_{\{s_1 > 1-\varepsilon\}}\right) \nu(\mathrm{d}\mathbf{s}) \\
&\geqslant c p^*_{N,\varepsilon} - (N-1)\nu(\{s_1 \leqslant 1 - \varepsilon\}) \\
&> -\infty.
\end{aligned}$$

This shows that $\underline{p}_{N,\varepsilon} \geqslant p^*_{N,\varepsilon}$. Similarly, for $a < p^*_{N,\varepsilon}$, we have

$$\begin{aligned}
\int_{\mathcal{S}^\downarrow} \sum_{i \neq j} s_i^{p^*_{N,\varepsilon} - a} s_j^{p^*_{N,\varepsilon}} \nu_{N,\varepsilon}(\mathrm{d}\mathbf{s}) &= \int_{\mathcal{S}^\downarrow} \sum_{i \neq j \leqslant N} s_i^{p^*_{N,\varepsilon} - a} s_j^{p^*_{N,\varepsilon}} \mathbb{1}_{\{s_1 \leqslant 1-\varepsilon\}} \nu(\mathrm{d}\mathbf{s}) \\
&\leqslant N^2 \nu(\{s_1 \leqslant 1 - \varepsilon\}) \\
&< \infty.
\end{aligned}$$

Thus $A_{N,\varepsilon} = p^*_{N,\varepsilon}$ □

By applying Lemma 6.2 for N tending to infinity and ε going to 0 (recall that $(\mathbf{M_q})$ does imply $\mathbb{E}[W] = 1$), we immediately obtain the following:

Proposition 6.7. *Assume* (**H**)*. Then, on the event where at least one of the* $\Pi^{N,\varepsilon}$ *does not die in finite time, we almost surely have*

$$\dim_{\mathcal{H}}(\mathcal{T}) \geqslant \frac{\sup_{N,\varepsilon} p^*_{N,\varepsilon}}{|\alpha|} = \frac{p^*}{|\alpha|}.$$

Thus, to complete our proof, we want to check the following lemma:

Lemma 6.8. *Almost surely, if* Π *does not die in finite time, then for* N *large enough and* ε *small enough,* $\Pi^{N,\varepsilon}$ *also does not.*

Proof. We will argue using Galton-Watson processes. Let, for all integers n, $Z(n)$ be the number of non-singleton and nonempty blocks of $\Pi(n)$ and, for all N and ε, $Z_{N,\varepsilon}(n)$ be the number of non-singleton and nonempty blocks of $\Pi^{N,\varepsilon}(n)$. These are Galton-Watson processes, which might take infinite values. We want to show that, on the event that Z doesn't die, there exist N and ε such that $Z_{N,\varepsilon}$ also survives. Letting q be the extinction probability of Z and $q_{N,\varepsilon}$ be the extinction probability of $Z_{N,\varepsilon}$, this will be proved by showing that $q = \inf\limits_{N,\varepsilon} q_{N,\varepsilon}$. By monotonicity properties, this infimum is actually equal to $q' = \lim\limits_{N\to\infty} q_{N,\frac{1}{N}}$.

Assume that $q < 1$ (otherwise there is nothing to prove). This implies that $\mathbb{E}[Z(1)] > 1$, and by monotone convergence, there exists N such that $\mathbb{E}[Z_{N,\frac{1}{N}}(1)] > 1$, and thus $q_{N,\frac{1}{N}} < 1$. Let, for $x \in [0,1]$, $F(x) = \mathbb{E}[x^{Z(1)}]$ and, for all N and ε, $F_{N,\varepsilon}(x) = \mathbb{E}[x^{Z_{N,\varepsilon}(1)}]$. The sequence of nondecreasing functions $(F_{N,\frac{1}{N}})_{N\in\mathbb{N}}$ converges simply to F. Since F is continuous on the compact interval $[0, q_{N,\frac{1}{N}}]$, the convergence is in fact uniform on this interval. We can take the limit in the relation $F_{N,\frac{1}{N}}(q_{N,\frac{1}{N}}) = q_{N,\frac{1}{N}}$ and get $F(q') = q'$. Since $q' < 1$ and since F only has two fixed points on $[0,1]$ which are q and 1, we obtain that $q = q'$.

□

We have thus proved the lower bound of Theorem 6.1: assuming (**H**), almost surely, if Π does not die in finite time, then $\dim_{\mathcal{H}}\big(\mathcal{L}(\mathcal{T})\big) \geqslant \frac{p^*}{|\alpha|}$.

6.3 Upper bound

Here we will not need the existence of an exact Malthusian exponent, and we will simply let

$$p' = \inf\Big\{p \geqslant 0, \psi(p) \geqslant 0\Big\}.$$

Proposition 6.9. *We have almost surely*

$$\dim_{\mathcal{H}}\big(\mathcal{L}(\mathcal{T})\big) \leqslant \frac{p'}{|\alpha|}.$$

This statement is in fact slightly stronger than the upper bound of Theorem 6.1. In particular it states that, if there exists $p \leqslant 0$ such that $\psi(p) \geqslant 0$, then the Hausdorff dimension of the set of leaves of $\mathcal{T}$ is almost surely equal to zero.

Proof. We will find a good covering of the set of proper leaves, in the same spirit as in [41], but which takes account of the sudden death of whole fragments. Let $\varepsilon > 0$. For all $i \in \mathbb{N}$, let

$$t_i^\varepsilon = \inf\{t \geqslant 0 : |\Pi_{(i)}(t)| < \varepsilon\}.$$

Note that this is in fact a stopping line as defined in section 2.1.4. We next define an exchangeable partition Π^ε by saying that integers i and j are in the same block if $\Pi_{(i)}(t_i^\varepsilon) = \Pi_{(j)}(t_j^\varepsilon)$. This should be thought of as the partition formed by the blocks of Π the instant they get small enough. Now, for all integers i, consider

$$\tau^{\varepsilon}_{(i)} = \sup_{j\in\Pi_{(i)}(t^{\varepsilon}_i)} \inf\{t \geqslant t^{\varepsilon}_i : |\Pi_{(j)}(t)| = 0\} - t^{\varepsilon}_i,$$

the time this block has left before it is completely reduced to dust. This allows us to define our covering. For all integers i, we let b^{ε}_i be the vertex of $[0, Q_i]$ at distance t^{ε}_i from the root. We take a closed ball with center b^{ε}_i and radius $\tau^{\varepsilon}_{(i)}$. These balls are the same if we take two integers in the same block of Π^{ε}, so we will only need to consider one integer i representing each block of Π^{ε}.

Let us check that this covers all of the proper leaves of $\mathcal{T}$. Let L be a proper leaf and $\big(i(t)\big)_{0\leqslant t\leqslant ht(L)}$ be any sequence of integers such that, for all $0 \leqslant t \leqslant ht(L)$, $(i(t), t) \leqslant L$ in $\mathcal{T}$. By definition of a proper leaf, $|\Pi_{(i(t))}(t)|$ does not suddenly jump to zero, so there exists a $t < ht(L)$ such that $0 < |\Pi_{(i(t))}(t)| \leqslant \varepsilon$. This implies that L is in the closed ball centered at $b^{\varepsilon}_{i(t)}$ with radius $\tau^{\varepsilon}_{(i(t))}$.

The covering is also *fine* in the sense that $\sup_i \tau^{\varepsilon}_i$ goes to 0 as ε goes to 0; indeed, if that wasn't the case, one would have a sequence $(i_n)_{n\in\mathbb{N}}$ and a positive number η such that $\tau^{2^{-n}}_{i_n} \geqslant \eta$ for all n. By compactness, one could then take a limit point x or a sequence $(b^{2^{-n}}_{i_n})_{n\in\mathbb{N}}$, and we would have $\mu(\mathcal{T}_x) = 0$ despite x not being a leaf, a contradiction.

Now, for $0 < \gamma \leqslant 1$, we have, summing one integer i per block of Π^{ε}, and using the extended fragmentation property with the stopping line $(t^{\varepsilon}_i)_{i\in\mathbb{N}}$,

$$\begin{aligned}\mathbb{E}\left[\sum_{i\in\mathrm{rep}(\Pi^{\varepsilon})} (\tau^{\varepsilon}_{(i)})^{\frac{\gamma}{|\alpha|}}\right] &\leqslant \mathbb{E}\left[\sum_{i\in\mathrm{rep}(\Pi^{\varepsilon})} \mathbb{E}\big[\tau^{\gamma/|\alpha|}\big]|\Pi^{\varepsilon}_{(i)}|^{\gamma}\right]\\ &\leqslant \mathbb{E}\big[\tau^{\gamma/|\alpha|}\big]\mathbb{E}\left[\sum_{i\in\mathrm{rep}(\Pi^{\varepsilon})} |\Pi^{\varepsilon}_{(i)}|^{\gamma}\right].\end{aligned}$$

Since τ has exponential moments (see [40], Proposition 14), the first expectation is finite and we only need to check when the second one is finite. Since Π^{ε} is an exchangeable partition, we know that, given its asymptotic frequencies, the asymptotic frequency of the block containing 1 is a size-biased pick among them and we therefore have

$$\begin{aligned}\mathbb{E}\left[\sum_i |\Pi^{\varepsilon}_i|^{\gamma}\right] &= \mathbb{E}\big[|\Pi^{\varepsilon}_1|^{\gamma-1}\mathbb{1}_{\{|\Pi^{\varepsilon}_1|\neq 0\}}\big]\\ &= \mathbb{E}\big[|\Pi_1(T_{\varepsilon})|^{\gamma-1}\mathbb{1}_{\{|\Pi_1(T_{\varepsilon})|\neq 0\}}\big]\\ &\leqslant \mathbb{E}\big[|\Pi_1(T_0^-)|^{\gamma-1}\big],\end{aligned}$$

where $T_{\varepsilon} = \inf\{t, |\Pi_1(t)| \leqslant \varepsilon\}$ and $T_0 = \inf\{t, |\Pi_1(t)| = 0\}$. Now recall that, up to a time-change which does not concern us here, the process $(|\Pi_1(t)|_{t\geqslant 0})$ is the exponential of the opposite of a killed subordinator $(\xi(t))_{t\geqslant 0}$ with Laplace exponent ϕ. This last expectation can be easily computed: let k be the killing rate of ξ and $\phi_0 = \phi - k$, ϕ_0 is then the Laplace exponent of a subordinator ξ' which evolves as ξ, but is not killed. By considering an independent exponential random time T with parameter k and killing ξ' at time T, one obtains a process with the same distribution as ξ. We thus have

$$\mathbb{E}[\mathrm{e}^{-(\gamma-1)\xi_{T^-}}] = \mathbb{E}[\mathrm{e}^{-(\gamma-1)\xi'_{T^-}}] = \int_0^{\infty} k\mathrm{e}^{-kt}\mathrm{e}^{-t\big(\phi_0(\gamma-1)\big)}dt = \int_0^{\infty} k\mathrm{e}^{-\phi(\gamma-1)t}dt.$$

Thus, if $\psi(\gamma) > 0$, then $\frac{\gamma}{|\alpha|}$ is greater than the Hausdorff dimension of the leaves of $\mathcal{T}$. $\square$

7 Some comments and applications

7.1 Comparison with previous results

In [41], the dimension of some conservative fragmentation trees was computed. The result was, as expected, $\frac{1}{|\alpha|}$, but this was obtained with very different assumptions on the dislocation measure:

Proposition 7.1. *Let ν be a conservative dislocation measure, $\alpha < 0$, and let $\mathcal{T}$ be a fragmentation tree with parameters $(\alpha, 0, \nu)$. Assume that ν satisfies the assumption $(\mathbf{H'})$ which we define by*

$$\int_{\mathcal{S}^{\downarrow}} (s_1^{-1} - 1)\nu(\mathrm{d}\mathbf{s}) < \infty.$$

Then, almost surely, we have

$$dim_{\mathcal{H}}(\mathcal{L}(\mathcal{T})) = \frac{1}{|\alpha|}.$$

This result complements ours - neither $(\mathbf{H})$ nor $(\mathbf{H'})$ is stronger than the other, which we are going to show by producing two corresponding examples.

For all $n \geqslant 2$, let $s_1^n = 1 - \frac{1}{n}$ and, for $i \geqslant 2$, $s_i^n = \frac{S}{n}\frac{1}{i(\log(i))^2}$, where $S = \Big(\sum_{i=2}^{\infty} \frac{1}{(i(\log(i))^2)}\Big)^{-1}$ (this ensures that $\sum_i s_i^n = 1$). Let then $\mathbf{s}^n = (s_i^n)_{i\in\mathbb{N}} \in \mathcal{S}^{\downarrow}$ and

$$\nu_1 = \sum_{n\geqslant 2} \frac{1}{n}\delta_{\mathbf{s}^n}.$$

We will show that this σ-finite measure on $\mathcal{S}^{\downarrow}$ is a dislocation measure which satisfies $(\mathbf{H'})$ but not $(\mathbf{H})$. First, $\int_{\mathcal{S}^{\downarrow}}(1-s_1)\nu_1(\mathrm{d}\mathbf{s}) = \sum_{n\geqslant 2}\frac{1}{n^2} < \infty$ so we do have a dislocation measure. Next, let us check $(\mathbf{H'})$:

$$\int_{\mathcal{S}^{\downarrow}} (s_1^{-1} - 1)\nu_1(\mathrm{d}\mathbf{s}) = \sum_{n\geqslant 2}\frac{1}{n}(\frac{n}{n-1} - 1) = \sum_{n\geqslant 2}\frac{1}{n(n-1)} < \infty.$$

Finally, $(\mathbf{H})$ is not verified: indeed, for any $p < 1$, $n \geqslant 2$ and $i \geqslant 2$, $(s_i^n)^p = \frac{S^p}{n^p}\big(i(\log(i))^2\big)^{-p}$ which is the general term of a divergent series.

Now we are going to do the same on the other side. For all $n \in \mathbb{N}$, let $t_1^n = \frac{1}{n}$ and, for $i \geqslant 2$, let $t_i^n = T(1-\frac{1}{n})\frac{1}{i^2}$, where $T = \big(\sum_{i=2}^{\infty}(\frac{1}{i^2})\big)^{-1}$. Since $t_2^n > t_1^n$ for large n, the sequence $\mathbf{t}^n = (t_i^n)_{i\in\mathbb{N}}$ is not a mass partition (despite its sum being equal to 1), and we will solve this problem by splitting its terms. Let $N(n) = \left\lceil \frac{t_2^n}{t_1^n} \right\rceil$, and then let $\mathbf{u}^n = (u_i^n)_{n\in\mathbb{N}} \in \mathcal{S}^{\downarrow}$ such that $u_1^n = t_1^n$ and, for $i \geqslant 2$, $u_i^n = \frac{t_k^n}{N(n)}$ where k is such that $i \in \{(k-2)N(n)+2, \ldots, (k-1)N(n)+1\}$. In other words, $\mathbf{u}^n$ starts with t_1^n, and then every term of $\mathbf{t}^n$ is divided by $N(n)$ and repeated $N(n)$ times. Now let us define

$$\nu_2 = \sum_{n\in\mathbb{N}} \frac{1}{n^2}\delta_{\mathbf{u}^n}.$$

The measure ν_2 integrates $1-s_1$ since it is finite, but $\sum_{n\in\mathbb{N}} \frac{1}{n^2}(\frac{1}{t_1^n}-1) = \sum_{n\in\mathbb{N}} \frac{1}{n} - \frac{1}{n^2} = \infty$, so $(\mathbf{H}')$ is not verified. On the other hand, for any $p<1$, we have

$$\int_{\mathcal{S}^\downarrow} \sum_i s_i^{p^*} \nu_2(\mathrm{d}\mathbf{s}) = \sum_{n\in\mathbb{N}} \frac{1}{n^2} \left(\frac{1}{n^p} + N(n) \Big(\frac{T(1-\frac{1}{n})}{N(n)}\Big)^p \big(\sum_{i\geqslant 2} \frac{1}{i^{2p}}\big) \right),$$

which is finite as soon as $p > \frac{1}{2}$, since $N(n)$ is asymptotically equivalent to $\frac{Tn}{4}$ as n goes to infinity. Thus ν_2 satisfies $(\mathbf{H})$.

7.2 Influence of parameters on the Malthusian exponent

We will here investigate what happens when we change some parameters of the fragmentation process. We start with a "basic" function ψ to which we will add either a constant (which amounts to increasing $\nu(\{(0,0,\ldots)\})$) or a linear part (which amounts to adding some erosion). We let $p_0 = \inf\{p \geqslant 0, \psi(p) > -\infty\}$. We also exclude the trivial case where $\nu(s_2>0)=0$, where the tree is always a line segment.

7.2.1 Influence of the killing rate

We assume here that $\nu((0,0,\ldots))=0$, which implies that $\psi(0)<0$, while we do not make any assumptions on the erosion parameter $c \geqslant 0$. We will quickly study how the Malthusian exponent changes when we add to ν a component of the form $k\delta_{(0,0,\ldots)}$ with $k\geqslant 0$. Let therefore, for $k \geqslant 0$, $\nu_k = \nu + k\delta_{(0,0,\ldots)}$ and, for $p\in\mathbb{R}$, $\psi_k(p) = cp + \int_{\mathcal{S}^\downarrow}(1-\sum_i s_i^p)\nu_k(\mathrm{d}\mathbf{s}) = \psi(p)+k$ and, if it exists, $p^*(k)$ the only number in $(0,1]$ which nulls the function ψ_k.

Proposition 7.2. *Assume* $(\mathbf{H})$ *for* (c,ν)*, that is* $\psi(p_0^+)<0$*, and let* $k_{max} = |\psi(p_0^+)|$*. Then, for* $k\in[0,k_{max})$*, the pair* (c,ν_k) *also satisfies* $(\mathbf{H})$*. Letting* $p^*(k_{max}) = p_0$ *(though it is not a Malthusian exponent in our sense when* $p_0=0$*), the function* $p^*(k)$ *on* $[0,k_{max}]$ *is the inverse function of* $-\psi$*. It is thus strictly decreasing and is differentiable as many times as* ψ*. For* $k\geqslant k_{max}$*,* $(\mathbf{H})$ *is no longer satisfied (in fact there is no Malthusian exponent if* $k>k_{max}$*), however we have in this case* $p_0 = \inf\{p\geqslant 0, \psi_k(p)\geqslant 0\}$ *which is the equivalent of* p' *in Section 6.3.*

7.2.2 Influence of erosion

Here we do not make any assumptions of ν, and let, for nonnegative c and any p, $\psi_c(p) = cp + \int_{\mathcal{S}^\downarrow}(1-\sum_i s_i^p)\nu(\mathrm{d}\mathbf{s})$. Note that, unlike in the previous section, the standard coupling between (α,c,ν)-fragmentations of Section 2.1.3 for all $c\geqslant 0$ is such that, almost surely, if for one c, $\Pi^{0,c}$ dies in finite time, then $\Pi = \Pi^{0,c}$ dies in finite time for all c. Thus, placing ourselves on the event where they do not die in a finite time, and calling $\mathcal{T}_c = \mathrm{TREE}(\Pi^{\alpha,c})$, we have $dim_{\mathcal{H}}(\mathcal{L}(\mathcal{T}_c)) = \frac{p^*(c)}{|\alpha|}$, $p^*(c)$ being the corresponding Malthusian exponent.

Proposition 7.3. *Assume* $(\mathbf{H})$ *for* $(0,\nu)$*, that is* $\psi(p_0^+)<0$*. If* $p_0=0$ *then the couple* (c,ν) *satisfies* $(\mathbf{H})$ *for all* c*, and its Malthusian exponent* $p^*(c)$ *tends to zero as* c *tends to infinity with the following asymptotics:*

$$p^*(c) \underset{c\to\infty}{\sim} \frac{|\psi(0)|}{c}.$$

If $p_0>0$*, then* (c,ν) *satisfies* $(\mathbf{H})$ *for* $c<c_{max}$ *with* $c_{max} = \frac{|\psi(p_0^+)|}{p_0}$*. By setting* $p^*(c_{max}) = p_0$*, the function* $c\to p^*(c)$ *is decreasing and is differentiable as many times as* ψ *is. For* $c\geqslant c_{max}$ $(\mathbf{H})$ *is no longer satisfied, however we do have* $p_0 = \inf\{p\geqslant 0, \psi_k(p)\geqslant 0\}$*.*

7.3 An application to the boundary of Galton-Watson trees

In this part we generalize some simple well-known results on the boundary of discrete Galton-Watson trees (see for example [47]) to trees where the lifetime of an individual is exponentially distributed. Unsurprisingly, the Hausdorff dimension of this boundary is the same in both cases.

Let $\xi = \sum p_i \delta_i$ be a probability measure on $\mathbb{N} \cup \{0\}$ which is supercritical in the sense that $m = \sum_i i p_i > 1$. Let $\mathcal{T}$ be a Galton-Watson tree with offspring distribution ξ and such that the individuals have exponential lifetimes with parameter 1. Seeing $\mathcal{T}$ as an $\mathbb{R}$-tree, we define a new metric on it by changing the length of every edge: let $a \in (1, \infty)$ and e be an edge of $\mathcal{T}$ connecting a parent and the child, we define the new length of e to be the old length of e times a^{-n}, where the parent is in the n-th generation of the Galton-Watson process. We let d' be this new metric.

The metric completion of $(\mathcal{T}, d')$ can then be seen as $\mathcal{T} \cup \partial\mathcal{T}$ where $\partial\mathcal{T}$ are points at the end of the infinite rays of $\mathcal{T}$.

Proposition 7.4. *On the event where $\mathcal{T}$ is infinite, we have*

$$\dim_{\mathcal{H}}(\partial\mathcal{T}) = \frac{\log m}{\log a}.$$

Proof. We start with the case where there exists $N \in \mathbb{N}$ such that, for $i \geqslant N+1$, $p_i = 0$. We aim to identify $(\mathcal{T}, d')$ as a fragmentation tree and apply Theorem 1.1. To do this, we first have to build a measure μ on it, as usual with Proposition 2.1. Let $x \in \mathcal{T}$, and let n be its generation, we then let $m(x) = \frac{1}{N^n}$. What this means is that the mass of the whole tree is 1, then each of the subtrees spawned by the death of the initial ancestor have mass $\frac{1}{N}$, then the death of each of these spawns trees with mass $\frac{1}{N^2}$, and so on.

We leave to the reader the details of the proof that $(\mathcal{T}, d', \mu)$ is a fragmentation tree, the corresponding parameters being $c = 0$, $\alpha = -\frac{\log a}{\log N}$ and $\nu = \sum p_i \delta_{\mathbf{s}^i}$, with $\mathbf{s}^i = (s_1^i, s_2^i, \ldots)$ such that $s_j^i = \frac{1}{N}$ if $j \leqslant i$ and $s_j^i = 0$ otherwise. One method of proof would be to couple $\mathcal{T}$ with an actual $(\alpha, 0, \nu)$-fragmentation process which would be obtained by constructing the death points one by one, following the tree and choosing a branch uniformly at each branching point, which is possible since the branching points of $\mathcal{T}$ form a countable set.

We then just need to compute the Malthusian exponent and check condition (**H**). We are looking for a number p^* such that $\int_{\mathcal{S}^{\downarrow}} (1 - \sum_{i=1}^N s_i^{p^*}) \nu(\mathrm{d}s) = 0$. This can be rewritten:

$$\begin{aligned}\int_{\mathcal{S}^{\downarrow}} (1 - \sum_{j=1}^N s_j^{p^*}) \nu(\mathrm{d}s) &= \sum_i p_i (1 - i \frac{1}{N^{p^*}}) \\ &= 1 - \frac{m}{N^{p^*}}.\end{aligned}$$

Thus we have $p^* = \frac{\log m}{\log N}$. Condition (**H**) is also easily checked, since $\psi(0) = 1 - m < 0$ and we thus get

$$\dim_{\mathcal{H}}(\partial\mathcal{T}) = \frac{p^*}{|\alpha|} = \frac{\log m}{\log a}.$$

The proof in the general case is once again done with a truncation argument, as in Section 6.2.3: once again leaving the details, we let, for all $N \in \mathbb{N}$, ξ_N be the law of $X \wedge N$ where X has law ξ. The monotone convergence theorem shows that the average of ξ_N converges to that of ξ, and the tree $\mathcal{T}$ with offspring distribution ξ can be simultaneously coupled with trees $(\mathcal{T}_N)_{N \in \mathbb{N}}$ with offspring distributions $(\xi_N)_{N \in \mathbb{N}}$, such that $\mathcal{T}$ has finite height (for its original metric) if and only if all the $(\mathcal{T}_N)_{N \in \mathbb{N}}$ also do. □

A Possibly infinite Galton-Watson processes

The purpose of this section is to extend the most basic results from the theory of discrete time Galton-Watson processes to the case where one parent may have an infinite amount of children. We refer to [45] for the classical results. Let Z be a random variable taking values in $\mathbb{Z}_+ \cup \{\infty\}$ with $\mathbb{P}[Z \geqslant 1] \neq 1$, and $(Z_n^i)_{i,n\in\mathbb{N}}$ be independent copies of Z. Let also, for $x \geqslant 0$, $F(x) = \mathbb{E}[x^Z]$. We define the process $(X_n)_{n\in\mathbb{N}}$ by $X_1 = 1$ and, for all n, $X_{n+1} = \sum_{i=1}^{X_n} Z_n^i$.

Proposition A.1. *The following are all true:*

(i) Almost surely, X either hits 0 in finite time or tends to infinity.

(ii) If X hits the infinite value once, then it stays there almost surely.

(iii) If $\mathbb{E}[Z] > 1$ then the function F has two fixed points on $[0,1]$: one is the probability of extinction q, and the other is 1. If $\mathbb{E}[Z] \leqslant 1$ then $q = 1$ and F only has one fixed point.

Proof. The proof of (i) is the same proof as in the classical case. For (ii), it is only a matter of seeing that, if we have $X_n = \infty$ for some n, then $\mathbb{P}[Z = 0] \neq 1$ and $\mathbb{E}[Z] > 0$, thus X_{n+1} is infinite by the law of large numbers. For (iii), in the case where $\mathbb{P}[Z = \infty] \neq 0$, we first show that $q \neq 1$ by taking an integer k such that $\mathbb{E}[\min(Z,k)] > 1$, and noticing that X dominates the classical Galton-Watson process where we have replaced, for all n and i, Z_n^i by $\min(Z_n^i, k)$, which is supercritical and thus has an extinction probability which is different from 1. Then, the fact that q is a fixed point of F and that F has at most two fixed points on $[0,1]$ are proved the same way as in the classical case. □

Chapter 3

Scaling limits of k-ary growing trees

Contents

For each integer $k \geqslant 2$, we introduce a sequence of k-ary discrete trees constructed recursively by choosing at each step an edge uniformly among the present edges and grafting on "its middle" $k-1$ new edges. When $k=2$, this corresponds to a well-known algorithm which was first introduced by Rémy. Our main result concerns the asymptotic behavior of these trees as the number of steps n of the algorithm becomes large: for all k, the sequence of k-ary trees grows at speed $n^{1/k}$ towards a k-ary random real tree that belongs to the family of self-similar fragmentation trees. This convergence is proved with respect to the Gromov-Hausdorff-Prokhorov topology. We also study embeddings of the limiting trees when k varies.

1 Introduction

The model. Let $k \geqslant 2$ be an integer. We introduce a growing sequence $\big(T_n(k)\big)_{n \in \mathbb{Z}_+}$ of k-ary discrete trees, constructed recursively as follows:

- STEP 0: $T_0(k)$ is the tree with one edge and two vertices: one root, one leaf.
- STEP n: given $T_{n-1}(k)$, choose uniformly at random one of its edges and graft on "its middle" $(k-1)$ new edges, that is split the selected edge into two so as to obtain two edges separated by a new vertex, and then add $k-1$ new edges to the new vertex.

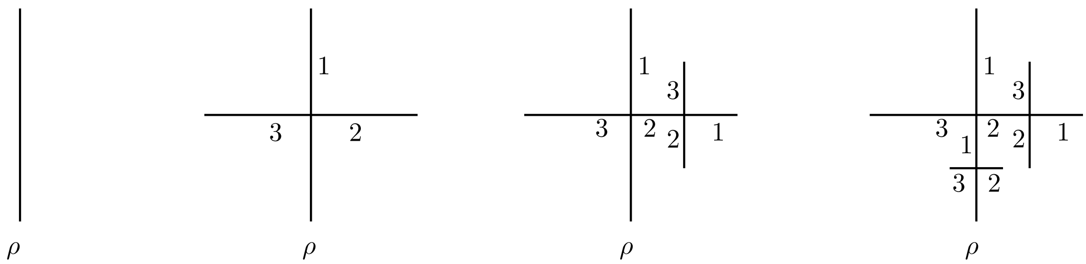

Figure 3.1: A representation of $T_n(3)$ for $n = 0, 1, 2, 3$. The edges are also labelled as explained in the paragraph just above Theorem 1.3.

For all n, this gives a tree $T_n(k)$ with $(k-1)n+1$ leaves, n internal nodes and $kn+1$ edges. In the case where $k=2$, edges are added one by one and our model corresponds to an algorithm introduced by Rémy [71] to generate trees uniformly distributed among the set of binary trees with $n+1$ labelled leaves. Many other dynamical models of trees growing by adding edges one by one exist in the literature, see for example [75, 16, 73, 35, 22].

Scaling limits. We are interested in the description of the metric structure of the growing tree $T_n(k)$ as n becomes large. For $k=2$, it is easy to explicitly compute the distribution of $T_n(2)$ (see for example [62]), which turns out to be that of a binary critical Galton-Watson tree conditioned to have $2n+2$ nodes (after forgetting the order). According to the original work of Aldous on scaling limits of Galton-Watson trees [5], the tree $T_n(2)$ then grows at speed $n^{1/2}$ towards a multiple of the Brownian continuum random tree (Brownian CRT). To be precise, viewing $T_n(2)$ for $n \in \mathbb{Z}_+$ as an $\mathbb{R}$-trees by considering that its edges all have length 1, and endowing it with the uniform probability measure on its leaves, which we denote by $\mu_n(2)$, we then have

$$\left(\frac{T_n(2)}{n^{1/2}}, \mu_n(2)\right) \underset{n \to \infty}{\overset{\text{a.s.}}{\longrightarrow}} \left(2\sqrt{2}\mathcal{T}_{\text{Br}}, \mu_{\text{Br}}\right) \tag{3.1}$$

for the GHP-topology, where $(\mathcal{T}_{\mathrm{Br}}, \mu_{\mathrm{Br}})$ is a random compact $\mathbb{R}$-tree called the Brownian continuum random tree (Brownian CRT for short). We point out that the almost sure convergence was not proved initially in [5], which states, in a more general setting, convergence in distribution of rescaled Galton-Watson trees. However, it is implicit in [69] and [61]. See also [24, Theorem 5] for an explicit statement.

Many classes of random trees are known to converge after rescaling towards the Brownian CRT. However, other limits are also possible, amongst which two important classes of random $\mathbb{R}$-trees: the class of Lévy trees introduced by Duquesne, Le Gall and Le Jan [57, 29, 30] (which is the class of all possible limits in distribution of rescaled sequences of Galton-Watson trees [29]) and the class of self-similar fragmentation trees which were studied in Chapter 1 (which is the class of scaling limits of the so-called Markov branching trees [42, 43]). We will see in this chapter that the sequence $(T_n(k), n \geqslant 0)$ has a scaling limit which to this second category.

The Brownian CRT belongs to the family of fragmentation trees, as shown in [10]. Its index of self-similarity is $-1/2$, its erosion coefficient is 0 and its dislocation measure $\nu_2^{\downarrow}$ is binary, conservative and such that

$$\nu_2^{\downarrow}(\mathrm{d}s_1) = \sqrt{\frac{2}{\pi}} s_1^{-3/2} s_2^{-3/2} \mathbb{1}_{\{s_1 \geqslant s_2\}} \mathrm{d}s_1 = \sqrt{\frac{2}{\pi}} s_1^{-1/2} s_2^{-1/2} \left(\frac{1}{1-s_1} + \frac{1}{1-s_2} \right) \mathbb{1}_{\{s_1 \geqslant s_2\}} \mathrm{d}s_1,$$

with $s_2 = 1 - s_1$. Of course the constraint $s_1 \geqslant s_2$ is here equivalent to $s_1 \geqslant 1/2$, but we keep the first notation in view of generalizations.

Our main goal is to generalize the convergence (3.1) to the sequences of trees $\big(T_n(k)\big)_{n \in \mathbb{Z}_+}$ for all integers $k \geqslant 2$. For reasons which will be apparent later on, we will want to consider mass partitions which are not necessarily ordered with the decreasing order, and will only reorder them when specifically looking at dislocation measures. We thus work with $\mathcal{S}_k$, the closed $(k-1)$-dimensional simplex and its variant $\mathcal{S}_{k,\leqslant}$ of dimension k obtained by allowing the sum to be less than 1,

$$\mathcal{S}_k = \left\{ \mathbf{s} = (s_1, s_2, ..., s_k) \in [0,1]^k : \sum_{i=1}^k s_i = 1 \right\}; \quad \mathcal{S}_{k,\leqslant} = \left\{ \mathbf{s} = (s_1, s_2, ..., s_k) \in [0,1]^k : \sum_{i=1}^k s_i \leqslant 1 \right\}.$$

Both spaces are endowed with the distance

$$d_k(\mathbf{s}, \mathbf{s}') = \sum_{i=1}^k |s_i - s'_i|,$$

which makes them compact. The Lebesgue measure on $\mathcal{S}_k$ can be written as $\mathrm{d}\mathbf{s} = \prod_{i=1}^{k-1} \mathrm{d}s_i$, with s_k being implicitly defined by $1 - \sum_{i=1}^{k-1} s_i$, whereas that on $\mathcal{S}_{k,\leqslant}$ should be understood as $\mathrm{d}\mathbf{s} = \prod_{i=1}^{k} \mathrm{d}s_i$.

Theorem 1.1. *Let $\mu_n(k)$ be the uniform measure on the leaves of $T_n(k)$. There exists a k-ary $\mathbb{R}$-tree $\mathcal{T}_k$, endowed with a probability measure on its leaves μ_k, such that*

$$\left(\frac{T_n(k)}{n^{1/k}}, \mu_n(k) \right) \xrightarrow{\mathbb{P}} (\mathcal{T}_k, \mu_k)$$

for the GHP-topology. The measured tree $(\mathcal{T}_k, \mu_k)$ is a fragmentation tree with index of self-similarity $-1/k$ and erosion coefficient 0. Its dislocation measure $\nu_k^{\downarrow}$ is supported on $\mathcal{S}_k$ and is defined by

$$\nu_k^{\downarrow}(\mathrm{d}\mathbf{s}) = \frac{(k-1)!}{k(\Gamma(\frac{1}{k}))^{k-1}} \prod_{i=1}^k s_i^{-(1-1/k)} \left(\sum_{i=1}^k \frac{1}{1-s_i} \right) \mathbb{1}_{\{s_1 \geqslant s_2 \geqslant \ldots \geqslant s_k\}} \mathrm{d}\mathbf{s},$$

where Γ *stands for Euler's Gamma function.*

Note that the convergence is a little weaker than (3.1) since it is only a convergence in probability. However some smaller objects akin to finite dimensional marginals of $T_n(k)$ converge almost surely as we will see later in Proposition 3.1. Note also that $\nu_k^{\downarrow}$ is a σ-finite measure on $\mathcal{S}_k$ such that

$$\int_{\mathcal{S}_k} (1 - s_1)\nu_k^{\downarrow}(\mathrm{d}\mathbf{s}) < \infty$$

but with infinite total mass. This fact implies in particular that the leaves of the tree $\mathcal{T}_k$ are dense in $\mathcal{T}_k$ (see [41, Theorem 1]).

Since the limiting tree is a fragmentation tree, we immediately have its Hausdorff dimension. Indeed, we know from [41, Theorem 2] that the Hausdorff dimension of a conservative fragmentation tree with index of self-similarity $\alpha < 0$ and dislocation measure ν (and no erosion) is equal to $\max(|\alpha|^{-1}, 1)$ provided that the measure ν integrates $(s_1^{-1} - 1)$. Here,

$$\int_{\mathcal{S}_k} (s_1^{-1} - 1)\nu_k^{\downarrow}(\mathrm{d}\mathbf{s}) \leqslant \int_{\mathcal{S}_k} k^{-1}(1 - s_1)\nu_k^{\downarrow}(\mathrm{d}\mathbf{s}) < \infty$$

since $s_1 \geqslant s_2 \geqslant ... \geqslant s_k$ together with $\sum_{i=1}^k s_i = 1$ implies that $s_1 \geqslant 1/k$.

Corollary 1.2. *The Hausdorff dimension of tree* $\mathcal{T}_k$ *is almost surely* k*.*

Remark. From the recursive construction of the sequence $(T_n(k))$ one could believe at first sight that the trees $T_n(k), n \geqslant 0$, as well as their continuous counterparts $\mathcal{T}_k$, are invariant under uniform re-rooting (which means that the law of the tree re-rooted at a leaf chosen uniformly at random is the same as the initial tree). However, this is only true for $k = 2$. For $k = 2$, this is a well-known property of the Brownian CRT (see [4]). For $k \geqslant 3$, it is easy to check for small values of n that this property is not satisfied for $T_n(k)$. In the continuous setting, it is known that a fragmentation tree which has the property of invariance under re-rooting necessarily belongs to the family of stable Lévy trees ([44]). It is also well-known that, up to a multiplicative scaling, the unique stable Lévy tree without vertices of infinite degree is the Brownian CRT ([30]). Hence $\mathcal{T}_k$ is not invariant under uniform re-rooting for $k \geqslant 3$.

Labels on edges and subtrees. Partly for technical reasons, we want to label all the edges of $T_n(k)$, with the exception of the edge adjacent to the root, with integers from 1 to k (see Figure 3.1 for an illustration). We do this recursively. The unique edge of $T_0(k)$ has no label since it is adjacent to the root. Given $T_n(k)$ and its labels, focus on the new vertex added in the middle of the selected edge. This edge was split into two: one new edge going towards the root, the other going away from it. Have the edge going towards the root keep the original label of the selected edge (no label if it is adjacent to the root), and have the other one be labelled 1. The $k-1$ new edges added after that will be labelled $2, \ldots, k$, say uniformly at random (actually, the way these $k-1$ additional edges are labelled is not important for our purpose, but the index 1 is important).

Now fix $2 \leqslant k' < k$. We consider, for all n, the k'-ary subtree of $T_k(n)$ obtained by discarding all edges with label larger than or equal to $k' + 1$ as well as their descendants. This subtree is denoted by $T_{k,k'}(n)$. We are interested in the sequence of subtrees $(T_{k,k'}(n), n \geqslant 0)$, because up to a (discrete) time-change in n, it is distributed as the sequence $(T_{k'}(n), n \geqslant 0)$ (see Lemma 4.5). As a consequence, we will see that a rescaled version of $\mathcal{T}_{k'}$ is nested in $\mathcal{T}_k$. Moreover this version can be identified as a non-conservative fragmentation tree. All this is precisely stated in the following theorem.

Theorem 1.3. *For each $n \in \mathbb{Z}_+$, endow $T_n(k)$ with the uniform probability on its leaves $\mu_n(k)$ and $T_n(k,k')$ with the image of this probability by the projection on $T_n(k,k')$. This image measure is denoted by $\mu_n(k,k')$. Then*

$$\left(\left(\frac{T_n(k)}{n^{1/k}}, \mu_n(k)\right), \left(\frac{T_n(k,k')}{n^{1/k}}, \mu_n(k,k')\right)\right) \underset{n\to\infty}{\overset{\mathbb{P}}{\longrightarrow}} \left((\mathcal{T}_k, \mu_k), (\mathcal{T}_{k,k'}, \mu_{k,k'})\right)$$

for the GHP-topology, where $\mathcal{T}_{k,k'}$ is a closed subtree of $\mathcal{T}_k$ and $(\mathcal{T}_{k,k'}, \mu_{k,k'})$ has the distribution of a non-conservative fragmentation tree with index $-1/k$ and no erosion. Its dislocation measure $\nu^{\downarrow}_{k,k'}$ is supported on $\mathcal{S}_{k',\leqslant}$ and defined by

$$\nu^{\downarrow}_{k,k'}(\mathrm{d}\mathbf{s}) = \frac{(k'-1)!}{k(\Gamma(\frac{1}{k}))^{k'-1}\Gamma(1-\frac{k'}{k})} \times \frac{1}{(1-\sum_{i=1}^{k'} s_i)^{k'/k}} \prod_{i=1}^{k'} s_i^{-(1-1/k)} \left(\sum_{i=1}^{k'} \frac{1}{1-s_i}\right) \mathbb{1}_{\{s_1 \geqslant \ldots \geqslant s_{k'}\}} \mathrm{d}\mathbf{s}.$$

Moreover,

$$\mathcal{T}_{k,k'} \overset{(\mathrm{d})}{=} M^{1/k'}_{k'/k,1/k} \cdot \mathcal{T}_{k'} \tag{3.2}$$

where in the right side, $M_{k'/k,1/k}$ has a generalized Mittag-Leffler distribution with parameters $(k'/k, 1/k)$ and is independent of $\mathcal{T}_{k'}$.

The identity (3.2) is similar to results established in [24] on the embedding of stable Lévy trees. The precise definition of generalized Mittag-Leffler distribution will be recalled in Section 4. In that section we will also see how to extract a random rescaled version of $\mathcal{T}_{k'}$ directly from the limiting fragmentation tree $\mathcal{T}_k$ by adequately pruning subtrees on each of its branch point (Proposition 4.6).

Organization of the chapter. We will use two approaches to prove our results. The first one, developed in Section 2, consists in checking that the sequence $(T_n(k), n \geqslant 0)$ possesses the so-called Markov branching property and then use results of Haas and Miermont [42] on scaling limits of Markov branching trees to obtain the convergence in distribution of the rescaled trees $(T_n(k))$ towards a fragmentation tree. Our second approach, in Section 3, is based on urn schemes and the Chinese restaurant process of Pitman [69]. It provides us the convergence in probability of the rescaled trees $(T_n(k))$ towards a compact $\mathbb{R}$-tree, but does not allow us to identify the limiting tree as a fragmentation tree. Combination of these two approaches then fully proves Theorem 1.1. In Section 3, we also treat the convergence in probability of the rescaled subtrees $(T_n(k,k'))$. The distribution of the limit will be identified in Section 4, hence giving the convergence results of Theorem 1.3. Lastly, still in Section 4, we study the embedding of the limiting trees $\mathcal{T}_k$ as k varies: for all $k' < k$, we show how to extract directly from $\mathcal{T}_k$ a tree with the distribution of $\mathcal{T}_{k,k'}$ and prove the relation (3.2).

From now on, k and k' are fixed, with $2 \leqslant k' < k$. To lighten notation, we will use, until Section 3.5, T_n instead of $T_n(k)$ and T'_n instead of $T_n(k,k')$.

2 Convergence in distribution and identification of the limit

In this section, we use Theorem 5 of [42] on the scaling limits of Markov branching trees to obtain the convergence in distribution of the rescaled trees $n^{-1/k}T_n$ and identify the limit distribution. While the method used in Section 3 will yield a stronger convergence, convergence in probability, that approach does not allow us to identify the distribution of the limit. We will also set up here some material needed to identify the distribution of the limit of the subtrees $n^{-1/k}T'_n$, which will be done in Section 4.

Let $n \in \mathbb{Z}_+$ and consider the tree T_{n+1}. Its root is connected to only one edge, after which there are k subtrees. These subtrees can be identified by the label given to their first edge, and we call them $(T_n^i)_{i\in[k]}$, where $i \in [k]$ refers to the edge number i. For all $i \in [k]$, we let X_n^i be the number of internal nodes of T_n^i and we let q_n be the distribution of $(X_n^i)_{i\in[k]}$ seen as an element of

$$\mathcal{C}_n^k = \left\{ \lambda = (\lambda_1, ..., \lambda_k) \in \mathbb{Z}_+^k : \sum_{i=1}^{k} \lambda_i = n \right\}.$$

To use the results of [42], we have to check

(i) that the sequence (T_n) is Markov branching, which roughly means that conditionally on their sizes, the trees $(T_n^i)_{i\in[k]}$ are mutually independent and have, respectively on $i \in [k]$, the same distribution as $T_{X_n^i}$;

(ii) that appropriately rescaled, the distribution q_n converges.

We start by studying this probability q_n in Section 2.1 and then prove the Markov branching property and get the limit distribution in Section 2.2.

2.1 Description and asymptotics of the measure q_n

Let $\bar{q}_n$ be the distribution of $(X_n^i/n)_{i\in[k]}$, it is a probability measure on $\mathcal{S}_k$, $\forall n \geqslant 1$. As we will see below in Proposition 2.1, the continuous scaling limit of these distributions is the measure ν_k on $\mathcal{S}_k$ defined by

$$\nu_k(\mathrm{d}\mathbf{s}) = \frac{1}{k(\Gamma(\frac{1}{k}))^{k-1}} \frac{1}{1-s_1} \prod_{i=1}^{k} s_i^{-(1-1/k)} \mathrm{d}\mathbf{s}.$$

Note the dissymmetry between the index 1 and the others. This is due the fact that the subtree T_n^1 is often much larger than the other ones, since, in the n-th step of the recursive construction, in the case where the new $k-1$ edges are added on the edge adjacent to the root, the subtree T_n^1 has n internal nodes whereas the $k-1$ other ones have none.

Since we are also interested in describing the asymptotic behaviour of the subtrees T'_n, we will also need to consider, for $n \geqslant 1$, the probability measures $\overline{q}'_n$ on $S_{k',\leqslant}$ obtained by considering the first k' elements of $(X_n^i/n)_{i\in[k]}$. Their continuous scaling limit (see Corollary 2.2) is denoted by $\nu_{k,k'}$ and defined on $\mathcal{S}_{k',\leqslant}$ by

$$\nu_{k,k'}(\mathrm{d}\mathbf{s}) = \frac{1}{k(\Gamma(\frac{1}{k}))^{k'-1}\Gamma(1-k'/k)} \times \frac{1}{(1-s_1)\,(1-\sum_{i=1}^{k'} s_i)^{k'/k}} \prod_{i=1}^{k'} s_i^{-(1-1/k)} \mathrm{d}\mathbf{s}.$$

For $\mathbf{s} \in \mathcal{S}_k$, we let $\mathbf{s}^{\downarrow}$ be the sequence obtained by reordering the elements of $\mathbf{s}$ in the decreasing order. This map is continuous from $\mathcal{S}_k$ to $\mathcal{S}_k$. For any measure μ on $\mathcal{S}_k$, let $\mu^{\downarrow}$ be the image of μ by it.

Examples. For instance, one can check that the measure $\nu_k^{\downarrow}$ associated to ν_k indeed coincides with the definition of $\nu_k^{\downarrow}$ in Theorem 1.1. And similarly for the measure $\nu_{k,k'}^{\downarrow}$ and its expression in Theorem 1.3.

The main goal of this section is to prove the following result.

Proposition 2.1. *We have the following weak convergence of measures on $\mathcal{S}_k$:*

$$n^{1/k}(1-s_1)\bar{q}_n(\mathrm{d}\mathbf{s}) \underset{n\to\infty}{\Rightarrow} (1-s_1)\nu_k(\mathrm{d}\mathbf{s}).$$

As a consequence,

$$n^{1/k}(1-s_1)\bar{q}_n^{\downarrow}(\mathrm{d}\mathbf{s}) \underset{n\to\infty}{\Rightarrow} (1-s_1)\nu_k^{\downarrow}(\mathrm{d}\mathbf{s}).$$

The *symmetric Dirichlet measure* on $\mathcal{S}_k$ with parameter k^{-1} is $\Gamma(1/k)^{-k}(\prod_{i=1}^k s_i)^{-(1-1/k)}\mathrm{d}\mathbf{s}$. It is well-known and easy to check that this defines a probability measure on $\mathcal{S}_k$. As a direct consequence, we see that

$$\int_{\mathcal{S}_k}(1-s_1)\nu_k(\mathrm{d}\mathbf{s}) = \frac{\Gamma(1/k)}{k}.$$

More generally, we will need several times the well-known fact that for any integer $K \geqslant 2$ and all K-uplets $\alpha_1,...,\alpha_K>0$,

$$\int_{\mathcal{S}_K}\prod_{i=1}^{K} s_i^{\alpha_i-1}\mathrm{d}\mathbf{s} = \frac{\prod_{i=1}^K\Gamma(\alpha_i)}{\Gamma\big(\sum_{i=1}^K\alpha_i\big)}, \tag{3.3}$$

where $x_K = 1-\sum_{i=1}^{K-1}x_i$.

The results of Proposition 2.1 can easily be transferred to $\bar{q}'_n$:

Corollary 2.2. *We have the following weak convergences of measures on $\mathcal{S}_{k',\leqslant}$:*

$$n^{1/k}(1-s_1)\bar{q}'_n(\mathrm{d}\mathbf{s}) \underset{n\to\infty}{\Rightarrow} (1-s_1)\nu_{k,k'}(\mathrm{d}\mathbf{s})$$

and,

$$n^{1/k}(1-s_1)(\bar{q}'_n)^{\downarrow}(\mathrm{d}\mathbf{s}) \underset{n\to\infty}{\Rightarrow} (1-s_1)\nu_{k,k'}^{\downarrow}(\mathrm{d}\mathbf{s}).$$

In order to prove Proposition 2.1, we start by explicitly computing the measure q_n in Section 2.1.1. We then set up preliminary lemmas in Section 2.1.2 and lastly turn to the proofs of Proposition 2.1 and Corollary 2.2 in Section 2.1.3.

2.1.1 The measure q_n

Proposition 2.3. *For all $\lambda \in \mathcal{C}_n^k$,*

$$q_n(\lambda) = \frac{1}{k(\Gamma(\frac{1}{k}))^{k-1}}\left(\prod_{i=1}^{k}\frac{\Gamma(\frac{1}{k}+\lambda_i)}{\lambda_i!}\right)\frac{n!}{\Gamma(\frac{1}{k}+n+1)}\left(\sum_{j=1}^{\lambda_1+1}\frac{\lambda_1!}{(\lambda_1-j+1)!}\frac{(n-j+1)!}{n!}\right).$$

Proof. Let $N_1,\dots,N_{n+1}$ be the $n+1$ internal nodes of T_{n+1}, listed in order of apparition, and let J be the random variable such that N_J is the first node encountered after the root of T_n. Recall that $T_n^1,\dots,T_n^k$ denote the ordered subtrees rooted at N_J. For $\lambda\in\mathcal{C}_n^k$ and $j\in\mathbb{N}$, we first compute the probability $p_j(\lambda)$ that $J=j$, T_n^1 contains the nodes $N_1,\dots,N_{j-1},N_{j+1},N_{\lambda_1+1}$, T_n^2 contains the nodes $N_{\lambda_1+2},\dots,N_{\lambda_1+\lambda_2+1}$ and so on until T_n^k, which contains the nodes $N_{\lambda_1+\ldots+\lambda_{k-1}+2},\dots,N_{n+1}$. This probability is null for $j>\lambda_1+1$. For $1\leqslant j\leqslant\lambda_1+1$, since each edge is chosen with probability $1/(1+kp)$ when constructing T_{p+1} from T_p, $p\geqslant 1$, we get

$$\begin{aligned} p_j(\lambda) &= \frac{1}{1+k(j-1)}\prod_{p=j+1}^{\lambda_1+1}\frac{1+k(p-2)}{1+k(p-1)}\prod_{i=2}^{k}\prod_{p=1}^{\lambda_i}\frac{1+k(p-1)}{1+k(\lambda_1+\ldots+\lambda_{i-1}+p)} \\ &= \frac{\prod_{i=1}^k\prod_{p=1}^{\lambda_i-1}(1+kp)}{\prod_{p=j-1}^{n}(1+kp)} \end{aligned}$$

(by convention, a product indexed by the empty set is equal to 1). Note that $p_j(\lambda) = p_1(\lambda)$ for all $j \leqslant \lambda_1 + 1$. Note also that this probability does not change if we permute the indices of nodes $N_{j+1}, \dots, N_n$ (both the numerator and the denominator have the same factors, just in different orders). We thus have

$$\begin{aligned} q_n(\lambda) &= \sum_{j=1}^{\lambda_1+1} \frac{(n-j+1)!}{(\lambda_1 - j+1)! \prod_{i=2}^k \lambda_i!} p_1(\lambda) \\ &= \frac{n!}{\prod_{p=1}^n (1+pk)} \prod_{i=1}^k \frac{\prod_{p=1}^{\lambda_i - 1}(1+pk)}{\lambda_i!} \sum_{j=1}^{\lambda_1+1} \frac{\lambda_1!}{(\lambda_1 - j+1)!} \frac{(n-j+1)!}{n!}. \end{aligned}$$

The proof is then ended by using the fact that $\Gamma(\frac{1}{k} + q) = \Gamma(\frac{1}{k}) k^{-q} \prod_{p=0}^{q-1}(1+kp)$ for any $q \in \mathbb{Z}_+$. $\square$

2.1.2 Preliminary lemmas

The proof of Proposition 2.1 relies on the convergence of some Riemann sums. To set up these convergences, we first rewrite $q_n(\lambda), n \geqslant 1$, in the form

$$q_n(\lambda) = \frac{1}{k\Gamma(\frac{1}{k})^{k-1}} \frac{\prod_{i=1}^k \gamma_k(\lambda_i)}{(n+1)\gamma_k(n+1)} \beta_n\left(\frac{\lambda_1}{n}\right), \tag{3.4}$$

where, for all $x \geqslant 0$,

$$\gamma_k(x) = \frac{\Gamma(\frac{1}{k} + x)}{\Gamma(1+x)}$$

and, for all $x \in [0,1]$ and $n \in \mathbb{N}$,

$$\beta_n(x) = 1 + \sum_{j=1}^{\lfloor nx \rfloor} \frac{nx(nx-1)\dots(nx-j+1)}{n(n-1)\dots(n-j+1)}.$$

Lemma 2.4. *The following convergence of functions*

$$x \mapsto n^{1-1/k}\gamma_k(nx) \underset{n\to\infty}{\longrightarrow} x \mapsto x^{-(1-1/k)}$$

holds uniformly on all compact subsets of $(0,1]$. *Moreover, there exists a finite constant A such that* $\gamma_k(x) \leqslant Ax^{-(1-1/k)}$ *for all* $x \geqslant 0$.

Proof. Pointwise convergence comes from a direct application of Stirling's formula. The uniformity of this convergence on all compact subsets of $(0,1]$ can be proved by a standard monotonicity argument (sometimes known as Dini's Theorem): we only need to notice that γ_k is a nonincreasing function of $x \geqslant 0$. This can be done through differentiating; indeed, γ_k is differentiable and we have, for all $x \geqslant 0$,

$$\gamma_k'(x) = \frac{\Gamma'(\frac{1}{k}+x)\Gamma(x+1) - \Gamma(\frac{1}{k}+x)\Gamma'(x+1)}{(\Gamma(x+1))^2}.$$

Notice that the function $x \mapsto \Gamma'(x)/\Gamma(x)$ is nondecreasing on $(0,+\infty)$, since the Gamma function is logarithmically convex (see for example [8]). Therefore, the derivative of γ_k is indeed nonpositive. Lastly, the domination of γ_k by a constant times the power function $x^{-(1-1/k)}$ for all $x \geqslant 0$ follows immediately from Stirling's formula and the fact that γ_k is continuous on $[0,+\infty)$. $\square$

Lemma 2.5. *The function β_n converges uniformly to the function $x \mapsto (1-x)^{-1}$ on all compact subsets of $[0,1)$. Moreover $(1-x)\beta_n(x) \leqslant 1$ for all $x \in [0,1]$ and all $n \in \mathbb{N}$.*

Proof. The proof works on the same principle as the previous one: since β_n is obviously a nondecreasing function, we only need to show that the sequence converges pointwise. This is immediate for $x = 0$, and will be done with the help of the dominated convergence theorem in the other cases. Note that, for all $x \in [0,1]$ and $j \in \mathbb{N}$

$$\frac{nx(nx-1)\dots(nx-j+1)}{n(n-1)\dots(n-j+1)} \underset{n\to\infty}{\longrightarrow} x^j \quad \text{and} \quad \frac{nx(nx-1)\dots(nx-j+1)}{n(n-1)\dots(n-j+1)} \leqslant x^j,\ \forall n \in \mathbb{N}, j \leqslant \lfloor nx \rfloor,$$

which is summable for $x \in [0,1)$. The dominated convergence theorem then ensures us that $\beta_n(x)$ converges to $\sum_{j=0}^{\infty} x^j = (1-x)^{-1}$ uniformly on all compact subsets of $[0,1)$. □

Lemma 2.6. *The sequence of measures $n^{\frac{1}{k}}(1-s_1)\bar{q}_n(\mathrm{d}\mathbf{s})$ satisfies:*

$$\forall \varepsilon > 0, \exists \eta > 0, \forall n \in \mathbb{N}, n^{\frac{1}{k}} \sum_{\lambda \in \mathcal{C}_n^k} \left(1 - \frac{\lambda_1}{n}\right) q_n(\lambda) \mathbb{1}_{\{\exists i, \lambda_i < \eta n\}} < \varepsilon.$$

Proof. We use (3.4). By individually bounding all the instances of $\gamma_k(x)$ by $Ax^{-(1-1/k)}$ and $(1-x)\beta_n(x)$ by 1 we are reduced to showing

$$\forall \varepsilon > 0, \exists \eta > 0, \forall n \in \mathbb{N}, \sum_{\lambda \in \mathcal{C}_n^k} \prod_{i=1}^{k} \lambda_i^{-(1-1/k)} \mathbb{1}_{\{\exists i, \lambda_i < \eta n\}} < \varepsilon.$$

By virtue of symmetry, we can restrict ourselves to the case where λ is nonincreasing. The condition $\exists i, \lambda_i < \eta n$ then boils down to $\lambda_k < \eta n$. Summation over λ nonincreasing and in $\in \mathcal{C}_n^k$ is done by choosing first λ_k then λ_{k-1}, going on until λ_2, the first term λ_1 being then implicitly defined as $n - \lambda_2 - \ldots - \lambda_k$. Let $\varepsilon > 0$, it is enough to find $\eta > 0$ such that, for any $n \in \mathbb{N}$,

$$\sum_{\lambda_k=1}^{\lfloor \eta n \rfloor} \sum_{\lambda_{k-1}=\lambda_k}^{\lfloor n/(k-1) \rfloor} \cdots \sum_{\lambda_2=\lambda_3}^{\lfloor n/2 \rfloor} \mathbb{1}_{\{\lambda_1 \geqslant \lambda_2\}} \prod_{i=1}^{k} \lambda_i^{-(1-1/k)} < \varepsilon.$$

By using $\lambda_1 \geqslant n/k$, we obtain

$$\sum_{\lambda_k=1}^{\lfloor \eta n \rfloor} \sum_{\lambda_{k-1}=\lambda_k}^{\lfloor n/(k-1) \rfloor} \cdots \sum_{\lambda_2=\lambda_3}^{\lfloor n/2 \rfloor} \mathbb{1}_{\{\lambda_1 \geqslant \lambda_2\}} \prod_{i=1}^{k} \lambda_i^{-(1-1/k)} \leqslant \left(\frac{n}{k}\right)^{-(1-1/k)} \sum_{\lambda_k=1}^{\lfloor \eta n \rfloor} \sum_{\lambda_{k-1}=1}^{\lfloor n/(k-1) \rfloor} \cdots \sum_{\lambda_2=1}^{\lfloor n/2 \rfloor} \prod_{i=2}^{k} \lambda_i^{-(1-1/k)}.$$

Standard comparison results between series and integrals imply that, since the function $t \mapsto t^{-(1-1/k)}$ is nonincreasing and has an infinite integral on $[1,\infty)$, there exists a finite constant B such that, for all $n \geqslant 1$, $\sum_{j=1}^{n} j^{-(1-1/k)} \leqslant Bn^{\frac{1}{k}}$. We thus get

$$\sum_{\lambda_k=1}^{\lfloor \eta n \rfloor} \sum_{\lambda_{k-1}=\lambda_k}^{\lfloor n/(k-1) \rfloor} \cdots \sum_{\lambda_2=\lambda_3}^{\lfloor n/2 \rfloor} \mathbb{1}_{\{\lambda_1 \geqslant \lambda_2\}} \prod_{i=1}^{k} \lambda_i^{-(1-1/k)} \leqslant B' \eta^{\frac{1}{k}} n^{-(1-1/k)} \eta^{\frac{1}{k}} (n^{\frac{1}{k}})^{k-1} \leqslant B'' \eta^{\frac{1}{k}}$$

where B' and B'' are finite constants. Choosing $\eta \leqslant (B'')^{-k}\varepsilon$ makes our sum smaller than ε for all choices of n. □

2.1.3 Proof of Proposition 2.1 and Corollary 2.2

Proof of Proposition 2.1. First note that since $(1-s_1)\nu_k(\mathrm{d}\mathbf{s})$ is a finite measure on $\mathcal{S}_k$ and since $\nu_k(\exists i : s_i = 0) = 0$, for all $\varepsilon > 0$ there exists a $\eta > 0$ such that $\int_{\mathcal{S}_k}(1-s_1)\mathbb{1}_{\{\exists i: s_i<\eta\}}\nu_k(\mathrm{d}\mathbf{s}) < \varepsilon$. Together with Lemma 2.6, this implies that Proposition 2.1 will be proved once we have checked that

$$n^{\frac{1}{k}}\int_{\mathcal{S}_k}(1-s_1)f(\mathbf{s})\prod_{i=1}^{k}\mathbb{1}_{\{s_i\geqslant\eta\}}\bar{q}_n(\mathrm{d}\mathbf{s}) \underset{n\to\infty}{\longrightarrow} \int_{\mathcal{S}_k}(1-s_1)f(\mathbf{s})\prod_{i=1}^{k}\mathbb{1}_{\{s_i\geqslant\eta\}}\nu_k(\mathrm{d}\mathbf{s})$$

for all $\eta > 0$ and all continuous functions f on $\mathcal{S}_k$. In the following, we fix such a real number $\eta > 0$ and a function f. Using the expression (3.4) and Lemmas 2.4 and 2.5, we see that

$$n^{\frac{1}{k}}\int_{\mathcal{S}^k}(1-s_1)f(\mathbf{s})\prod_{i=1}^{k}\mathbb{1}_{\{s_i\geqslant\eta\}}\bar{q}_n(\mathrm{d}\mathbf{s}) \underset{n\to\infty}{\sim} \frac{n^{1-k}}{k\Gamma(\frac{1}{k})^{k-1}}\sum_{\lambda\in\mathcal{C}_n^k} f\left(\frac{\lambda}{n}\right)\prod_{i=1}^{k}\left(\frac{\lambda_i}{n}\right)^{-(1-1/k)}\mathbb{1}_{\{\lambda_i\geqslant\eta n\}}.$$

We conclude by noticing that this last term is in fact a Riemann sum of a (Riemann) integrable function on $[0,1]^{k-1}$: to sum over $\lambda\in\mathcal{C}_n^k$, we only need to choose $\lambda_1,\ldots,\lambda_{n-1}$ in $\{0,...,n\}$ such that $n-(\lambda_1+\ldots+\lambda_{n-1})\geqslant 0$. Standard results on Riemann sums then imply that it converges towards the integral

$$\int_{\mathcal{S}_k}(1-s_1)f(\mathbf{s})\prod_{i=1}^{k}\mathbb{1}_{\{s_i\geqslant\eta\}}\nu_k(\mathrm{d}\mathbf{s}).$$

The convergence of the decreasing versions of the measures follows immediately. A continuous function f on $\mathcal{S}_k$ being fixed, we let g_f be the function defined on $\mathcal{S}_k$ by $g_f(\mathbf{s}) = (1-s_1^{\downarrow})f(\mathbf{s}^{\downarrow})/(1-s_1)$. The function g_f is then continuous and bounded on $\mathcal{S}_k$ (there is no singularity when $s_1 = 1$ since $s_1^{\downarrow} = s_1$ as soon as $s_1 \geqslant 1/2$). By the first part of this proof, we then have

$$\begin{aligned} n^{\frac{1}{k}}\int_{\mathcal{S}_k}(1-s_1)f(\mathbf{s})\bar{q}_n^{\downarrow}(\mathrm{d}\mathbf{s}) &= n^{\frac{1}{k}}\int_{\mathcal{S}_k}(1-s_1^{\downarrow})f(\mathbf{s}^{\downarrow})\bar{q}_n(\mathrm{d}\mathbf{s}) = n^{\frac{1}{k}}\int_{\mathcal{S}_k}(1-s_1)g_f(\mathbf{s})d\bar{q}_n(\mathbf{s}) \\ &\underset{n\to\infty}{\rightarrow} \int_{\mathcal{S}_k}(1-s_1)g_f(\mathbf{s})\nu_k(\mathrm{d}\mathbf{s}) = \int_{\mathcal{S}_k}(1-s_1)f(\mathbf{s})\nu_k^{\downarrow}(\mathrm{d}\mathbf{s}). \end{aligned}$$

□

Proof of Corollary 2.2. Let f be a continuous function on $\mathcal{S}_{k',\leqslant}$ and assume for the moment that $k'\leqslant k-2$. Applying first Proposition 2.1 and then the identity (3.3), we get

$$\begin{aligned} &n^{1/k}\int_{\mathcal{S}_{k',\leqslant}} f(\mathbf{s})(1-s_1)\bar{q}'_n(\mathrm{d}\mathbf{s}) \underset{n\to\infty}{\longrightarrow} \frac{1}{k(\Gamma(\frac{1}{k}))^{k-1}}\int_{\mathcal{S}_k} f(s_1,...,s'_k)\prod_{i=1}^{k}s_i^{-(1-1/k)}\mathrm{d}\mathbf{s} \\ &= \frac{1}{k(\Gamma(\frac{1}{k}))^{k-1}}\int_{S_{k',\leqslant}} f(\mathbf{s})\prod_{i=1}^{k'}s_i^{-(1-1/k)}\left(\int_{\left(0,1-\sum_{i=1}^{k'}s_i\right]^{k-1-k'}}\prod_{i=k'+1}^{k}s_i^{-(1-1/k)}\mathrm{d}s_{k'+1}\ldots\mathrm{d}s_{k-1}\right)\mathrm{d}\mathbf{s} \\ &= \frac{1}{k(\Gamma(\frac{1}{k}))^{k-1}}\int_{S_{k',\leqslant}} f(\mathbf{s})\prod_{i=1}^{k'}s_i^{-(1-1/k)}\left(\frac{\Gamma(\frac{1}{k})^{k-k'}}{\Gamma(1-k'/k)}\left(1-\sum_{i=1}^{k'}s_i\right)^{-k'/k}\right)\mathrm{d}\mathbf{s}, \end{aligned}$$

which gives the result for $k'\leqslant k-2$. For $k' = k-1$ the calculation is more direct since we do not need (3.3). Finally, the convergence of decreasing measures follows immediately by mimicking the end of the proof of Proposition 2.1. □

2.2 Markov branching property and identification of the limit

Proposition 2.7 (Markov branching property)**.** *Let $n \in \mathbb{Z}_+$. Conditionally on $(X_n^i)_{i\in[k]}$, the $(T_n^i)_{i\in[k]}$ are mutually independent and, for $i \in [k]$, T_n^i has the same law as $T_{X_n^i}$.*

Proof. We prove this statement by induction on $n \in \mathbb{Z}_+$. Starting with $n = 0$, we have $X_0^i = 0$ and $T_0^i = T_0$ for all i, everything is deterministic.

Assume now that the Markov branching property has been proven until some integer $n-1$, and let us prove it for n. Let e be the random selected edge of T_{n-1} used to build T_n and let J be the random variable defined by: $J = j$ if e is an edge of T_{n-1}^j with $j \in [k]$, and $J = 0$ if e is the edge adjacent to the root of T_n. Note that J and T_n are independent conditionally on $(X_{n-1}^i)_{i\in[k]}$. Let us then determine the law of $(T_n^i, i \in [k])$ conditionally on J and $(X_n^i)_{i\in[k]}$.

If $J = j \neq 0$ then $(T_n^i)_{i\in[k]\setminus\{j\}}$ is the same sequence as $(T_{n-1}^i)_{i\in[k]\setminus\{j\}}$ and we have added an extra edge to T_{n-1}^j. Hence, for all $j \leqslant k$ and $(x_1, ..., x_k) \in \mathcal{C}_n^k$, with $x_j \geqslant 1$, we have for all k-uplet of rooted k-ary trees $(t_1, ..., t_k)$ with respectively $x_1, ..., x_k$ internal nodes,

$$\begin{aligned}
& \mathbb{P}\left(\forall i \in [k], T_n^i = t_i \mid \forall i \in [k], X_n^i = x_i, J = j\right) \\
= \; & \mathbb{P}\left(\forall i \in [k] \setminus \{j\}, T_{n-1}^i = t_i, T_n^j = t_j \mid \forall i \in [k] \setminus \{j\}, X_{n-1}^i = x_i, J = j\right) \\
= \; & \mathbb{P}\left(\forall i \in [k] \setminus \{j\}, T_{n-1}^i = t_i \mid \forall i \in [k] \setminus \{j\}, X_{n-1}^i = x_i, J = j\right) \mathbb{P}\left(T_{x_j} = t_j\right) \\
= \; & \textstyle\prod_{i=1}^k \mathbb{P}\left(T_{x_i}^i = t_i\right)
\end{aligned}$$

where we have used that a conditioned uniform variable is uniform in the set of conditioning to get the second equality and then that J and T_n are independent conditionally on $(X_{n-1}^i)_{i\in[k]}$, together with the Markov branching property at $n-1$, to get the third equality.

When $J = 0$, $(T_n^i)_{i\leqslant k} = (T_n, T_0, \ldots, T_0)$. Since T_0 is deterministic and the event $\{J = 0\}$ is independent of T_n, the distribution of $(T_n^i)_{i\in[k]}$ conditional on $J = 0$ and $(X_n^i)_{i\in[k]} = (n, 0, ..., 0)$ is indeed the same as that of sequence of independent random variables $(T_n, T_0, \ldots, T_0)$.

Finally, since the distribution of $(T_n^i)_{i\in[k]}$ conditionally on J and $(X_n^i)_{i\in[k]}$ is independent of J, one can remove J in the conditioning, which ends the proof. □

We now have the material to prove the convergence in distribution of $n^{-1/k}T_n$ and identify its limit as a fragmentation tree.

Proof of Theorem 1.1 (convergence in distribution part). Theorem 5 of [42] concerns sequences of Markov branching trees indexed by their number of leaves, however our sequence (T_n) is indexed by the number of internal nodes of the tree. This is not a real problem since T_n has $1 + (k-1)n$ leaves for all n, the sequence $(T_n)_{n\in\mathbb{N}}$ can be seen as a sequence of Markov branching trees $(T_p^\circ)_{p\in 1+(k-1)\mathbb{N}}$ indexed by their number of leaves. For all $p \in 1 + (k-1)\mathbb{N}$, we let $\bar{q}_p^\circ$ denote its associated splitting distribution, that is, if $p = 1+(k-1)n$, $\bar{q}_p^\circ$ is the distribution on $\mathcal{S}_k$ of the sequence

$$\left(\frac{1 + (k-1)X_{n-1}^i}{1 + (k-1)n}\right)_{i\in[k]}.$$

As an immediate consequence of Proposition 2.1, we have that

$$(k-1)^{-1/k} p^{1/k} (1 - s_1) \bar{q}_p^{\circ,\downarrow}(\mathrm{d}\mathbf{s}) \underset{\substack{p\to\infty \\ p\in 1+(k-1)\mathbb{N}}}{\Rightarrow} (1 - s_1)\nu_k^\downarrow(\mathrm{d}\mathbf{s}). \tag{3.5}$$

Indeed, for any bounded Lipschitz function $f : \mathcal{S}_k \to \mathbb{R}$, let $g_f : \mathcal{S}_k \to \mathbb{R}$ be defined by $g_f(\mathbf{s}) = (1-s_1)f(\mathbf{s})$. Then g_f is also Lipschitz, say with Lipschitz constant c_g. It is then easy to see that

$$n^{1/k}\left|\mathbb{E}\left[g\left(\left(\frac{1+(k-1)X^i_{n-1}}{1+(k-1)n}\right)^{\downarrow}\right)\right] - \mathbb{E}\left[g\left(\left(\frac{X^i_{n-1}}{n-1}\right)^{\downarrow}\right)\right]\right| \leqslant \frac{n^{1/k}2kc_g}{1+(k-1)n} \underset{n\to\infty}{\to} 0.$$

Together with Proposition 2.1 this immediately leads to (3.5).

Hence the sequence $(T^\circ_p)_{p\in 1+(k-1)\mathbb{N}}$ is Markov branching with a splitting distribution sequence $(\bar{q}^\circ_p)$ satisfying (3.5). This is exactly the hypotheses we need to apply Theorem 5 of [42], except that this theorem is stated for sequences of Markov branching trees indexed by the full set $\mathbb{N}$, not by one of its subsets. However, without any modifications, it could easily be adapted to that setting. Hence we obtain from this theorem that

$$\left((k-1)^{1/k}p^{-1/k}T^\circ_p, \mu^\circ_p\right) \underset{\substack{p\to\infty \\ p\in 1+(k-1)\mathbb{N}}}{\longrightarrow} (\mathcal{T}_k, \mu_k)$$

where μ°_p is the uniform probability on the leaves of (T°_p) and $(\mathcal{T}_k,\mu_k)$ the fragmentation tree of Theorem 1.1. This convergence holds in distribution, for the GHP topology. Otherwise said, $(n^{-1/k}T_n)$ endowed with the uniform probability on its leaves converges in distribution towards $(\mathcal{T}_k,\mu_k)$. □

3 Convergence in probability and joint convergence

This section is dedicated to improving the convergence in distribution we have just obtained. We will construct the limiting tree in the space ℓ^1 of summable real-valued sequences (equipped with its usual metric d_{ℓ^1}), and convergence will be proved by using subtrees akin to finite-dimensional marginals. The almost sure convergence of these marginals can be proved using urn schemes and results concerning Chinese restaurant processes, as studied by Pitman in [69, Chapter 3]. Tightness properties will extend this to the convergence of $(n^{-1/k}T_n, \mu_n)$. Unfortunately, almost sure convergence is lost by this method and we are left with convergence in probability. Also, due to some technical issues, we first have to study the Gromov-Hausdorff convergence of the non-measured trees before adding the measures.

3.1 Finite-dimensional marginals and the limiting tree

In this section we will need to define an ordering of the leaves of T_n for $n\in\mathbb{Z}_+$, calling them $(L^i_n)_{i\in[(k-1)n+1]}$. They are labelled by order of apparition: the single leaf of T_0 is called L^1_0, while, given T_n and its leaves, the leaves $L^1_{n+1},\ldots,L^{(k-1)n+1}_{n+1}$ of T_{n+1} are those inherited from T_n, and the leaves $L^{(k-1)n+2}_{n+1},\ldots,L^{(k-1)n+k}_{n+1}$ are the leaves at the ends of the new edges labelled $2,3,\ldots,k$ respectively.

Let $p\in\mathbb{Z}_+$. For all $n\geqslant p$, consider the subtree T^p_n of T_n spanned by the root and all the leaves L^i_n with $i\in[(k-1)p+1]$:

$$T^p_n = \bigcup_{i=1}^{(k-1)p+1} [\![\rho, L^i_n]\!].$$

The tree T^p_n has the same graph structure as T_p, however the metric structure isn't the same: the distance between two vertices of T^p_n is the same as the distance between the corresponding

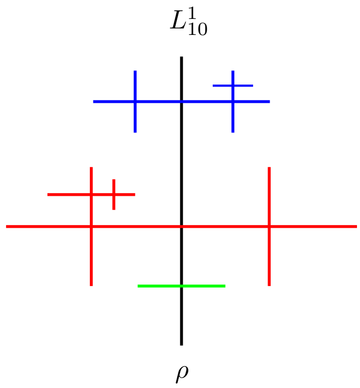

Figure 3.2: Colour-coding of the tables of T_{10} (here $k = 3$). The green table has one client, the red table has five and the blue table has four.

vertices of T_n. The study of the sequence $(T_n^p)_{n\geqslant p}$ for all p will give us much information on the sequence $(T_n, \mu_n)_{n\in\mathbb{Z}_+}$.

Proposition 3.1. *Let $p \in \mathbb{Z}_+$. We have, in the Gromov-Hausdorff sense, as n goes to infinity:*

$$\frac{T_n^p}{n^{1/k}} \xrightarrow{\text{a.s.}} \mathcal{T}^p, \tag{3.6}$$

where $\mathcal{T}^p$ is a rooted compact $\mathbb{R}$-tree with $(k-1)p+1$ leaves which we will call $(L^i)_{i\in[(k-1)p+1]}$. Under a suitable embedding in ℓ^1, for $p' < p$, $\mathcal{T}^{p'}$ is none other than the subtree of $\mathcal{T}^p$ spanned by the root and the leaves L^i for $i \in [(k-1)p'+1]$, making this notation unambiguous.

Proof. The proof hinges on our earlier description of T_n^p for $n \geqslant p$: it is the graph T_p, but with distances inherited from T_n. As explained in Lemma 3.1, we only need to show that, for i and j smaller than $(k-1)p+1$, both $n^{-1/k}d(L_n^i, L_n^j)$ and $n^{-1/k}d(\rho, L_n^i)$ have finite limits as n goes to infinity. We first concentrate on the case of $n^{-1/k}d(\rho, L_n^1)$. This could be done by noticing that $(d(\rho, L_n^1))_{n\geqslant 0}$ is a Markov chain and using martingale methods, however, in view of what will follow, we will use the theory of Chinese restaurant processes.

For $n \in \mathbb{N}$, we consider a set of tables indexed by the vertices of T_n which are strictly between ρ and L_n^1. We then let the number of clients on the table indexed by a vertex v be the number of internal nodes u of T_n such that v is the branch point of u and L_1^n (including the case $u = v$).

Let us check that this process is part of the two-parameter family introduced by Pitman in [69], Chapter 3, with parameters $(1/k, 1/k)$. Indeed, assume that, at time $n \in \mathbb{N}$, we have $l \in \mathbb{N}$ tables with respectively $n_1, \ldots, n_l$ clients (the tables can be ordered by their order of apparition in the construction). For any $i \leqslant l$, table i corresponds to a subset of T_n with $kn_i - 1$ edges, thus there is a probability of $(kn_i - 1)/(kn+1)$ that the next client comes to this table. This next client will sit at a new table if the selected edge is between ρ and L_n^1, an event with probability $(l+1)/(kn+1)$.

Since, for all n, $d(\rho, L_n^1)$ is equal to the number of tables plus one, Theorem 3.8 of [69] tells us that $n^{-1/k}d(\rho, L_n^1)$ converges almost surely towards a $(1/k, 1/k)$-generalized Mittag-Leffler random variable (the definition of generalized Mittag-Leffler distributions is recalled in Section 4.2, however we will not need here the exact distribution of this limit). The cases of $d(\rho, L_n^i)$ and $d(L_n^i, L_n^j)$ for $i \neq j$ can be treated very much the same way: the main difference is that the tables of the restaurant process are now indexed by the nodes between L_n^i and L_n^j, and they have a non-trivial initial configuration. Lemma 3.1 finally implies that $n^{-1/k}T_n^p$ does converge a.s. to a tree with $(k-1)p+1$ leaves in the Gromov-Hausdorff sense.

The trees $(\mathcal{T}^p, p\in\mathbb{Z}_+)$ and $(T_n^p, p\in\mathbb{Z}_+, n\in\mathbb{Z}_+)$ can be embedded in ℓ^1 using the stick-breaking method of Chapter 1, Section 3.1. By taking as marked points the leaves, it is indeed the case that this embedding respects the natural inclusion $\mathcal{T}^p\subset\mathcal{T}^{p+1}$. ☐

Under this embedding in ℓ^1 we let

$$\mathcal{T}=\overline{\cup_{p=0}^{\infty}\mathcal{T}^p},$$

which is also an $\mathbb{R}$-tree. We will see in Lemma 3.4 and Proposition 3.5 that this tree is compact and is the limiting tree for $n^{-1/k}T_n$, the tree which was called $\mathcal{T}_k$ in the introduction.

3.2 A tightness property

To move from the convergence of T_n^p for all $p\in\mathbb{N}$ to the convergence of T_n, we need some kind of compactness to not be bothered by the choice of p, which the following proposition gives.

Proposition 3.2. *For all $\varepsilon>0$ and $\eta>0$, there exists an integer p such that, for n large enough,*

$$\mathbb{P}\big(d_{\ell^1,\mathrm{H}}(T_n^p,T_n)>n^{\frac{1}{k}}\eta\big)<\varepsilon.$$

The same is then true if we replace p by any greater integer p'.

Before proving this proposition, we need an intermediate result. Fix $p\in\mathbb{Z}_+$. All the variables in the following lemma depend on a variable $n\geqslant p$, however we omit mentioning n for the sake of readability.

Lemma 3.3. *Let $v_1,\dots,v_N$ be the internal nodes of T_n which are part of T_n^p but are not branch points of T_n^p, listed in order of apparition. At each of these vertices are rooted $k-1$ subtrees of T_n which we call $(S_j^i\,;\,j\leqslant N, i\leqslant k-1)$, S_j^i being the tree rooted at v_j with a unique edge adjacent to v_j, this edge having label $i+1$. Letting, for $j\leqslant N$ and $i\leqslant k-1$, Y_j^i be the number of internal nodes of S_j^i then, conditionally on $(Y_q^l; q\leqslant N, l\leqslant k-1)$, the tree S_j^i has the same distribution as $T_{Y_j^i}$.*

Furthermore, these subtrees allow us to define some restaurant processes by letting n vary: for $j\leqslant N$, let $S_j=\cup_{i=1}^{k-1}S_j^i$, and Y_j be the number of vertices of S_j, including v_j but excluding all leaves. Considering S_j as a table with Y_j clients for all j, we have defined a restaurant process whose initial configuration is zero tables at time $n=p$ and has parameters $(1/k,p+1/k)$.

The subtrees (S_j^i) are also conditionally independent, however this will not be useful to us.

Proof. The proof that S_j^i is, conditionally on $(Y_q^l; q\leqslant N, l\leqslant k-1)$, distributed as $T_{Y_j^i}$ is a straightforward induction on n. We will not give details since this is very similar to the Markov branching property (Proposition 2.7) but the main point is that, conditionally on the event that the selected edge at a step of the algorithm is an edge of S_j^i, this edge is then uniform amongst the edges of S_j^i.

The restaurant process nature of these subtrees is proved just as in Proposition 3.1: if table S_j has Y_j clients at time $n\geqslant p$, then the subtree S_j has kY_j-1 edges, and a new client will therefore be added to this table with probability $(kY_j-1)/(kn+1)$, while a new table is formed with probability $(kp+N+1)/(kn+1)$. These are indeed the transition probabilities of a restaurant process with parameters $(1/k,p+1/k)$ taken at time $n-p$. ☐

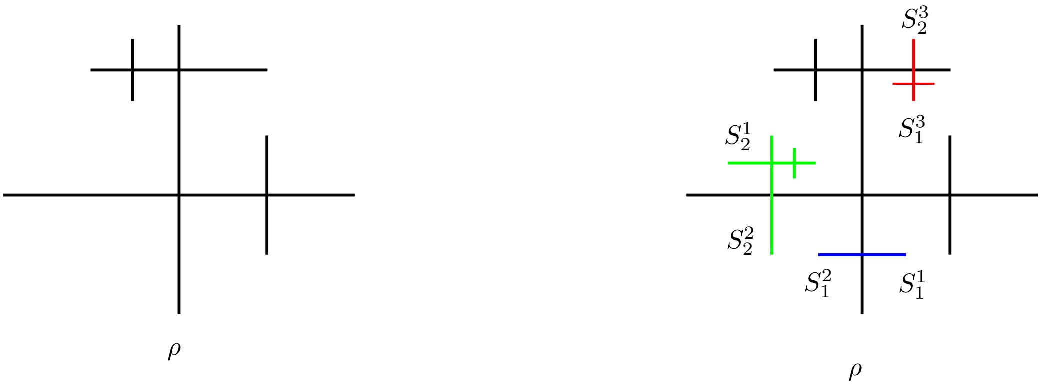

Figure 3.3: The tree T_{10} seen as an extension of T_4 ($k = 3$). The colored sections correspond to the tables of the Chinese restaurant, each table corresponding to two subtrees.

Proof of Proposition 3.2. We will need Lemma 33 of [42]: since the sequence $(T_n)_{n\in\mathbb{Z}_+}$ is Markov branching and we have the convergence of measures of Proposition 2.1, we obtain that, for any $q > 0$, there exists a finite constant C_q such that, for any $\varepsilon > 0$ and $n \in \mathbb{N}$,

$$\mathbb{P}\Big(ht(T_n) \geqslant \varepsilon n^{1/k}\Big) \leqslant \frac{C_q}{\varepsilon^q}.$$

Choosing $q > k$, applying this to all of the S_j^i conditionally on Y_j^i ($j \leqslant N, i \leqslant k-1$), and using the simple fact that

$$d_{\ell^1,\mathrm{H}}(T_n^p, T_n) \leqslant \max_{i,j} ht(S_i^j),$$

we obtain

$$\begin{aligned}
\mathbb{P}\Big(d_{\ell^1,\mathrm{H}}(T_n^p, T_n) > \eta n^{1/k} \mid (Y_j^i)_{i,j}\Big) &\leqslant \sum_{i,j} \mathbb{P}\Big(ht(S_j^i) > \eta n^{\frac{1}{k}} \mid (Y_j^i)_{i,j}\Big) \\
&\leqslant \sum_{i,j} \mathbb{P}\left(ht(S_j^i) > \eta \Big(\frac{n}{Y_j^i}\Big)^{\frac{1}{k}} (Y_j^i)^{\frac{1}{k}} \mid (Y_j^i)_{i,j}\right) \\
&\leqslant \frac{C_q}{\eta^q} \sum_{i,j} \left(\frac{Y_j^i}{n}\right)^{\frac{q}{k}} \\
&\leqslant \frac{C_q}{\eta^q} \sum_{j} \left(\frac{Y_j}{n}\right)^{\frac{q}{k}}.
\end{aligned}$$

Let us now reorder the (Y_j) in decreasing order. Theorem 3.2 of [69] states the following convergence for all j as n goes to infinity:

$$\frac{Y_j}{n} \xrightarrow{\text{a.s.}} V_j$$

where $(V_j)_{j\in\mathbb{N}}$ is a Poisson-Dirichlet random variable with parameters $(1/k, p+1/k)$. By writing out, for each j, $(Y_j)^{\frac{q}{k}} \leqslant (Y_1)^{\frac{q}{k}-1} Y_j$, we then get

$$\limsup_{n\to\infty} \mathbb{P}\Big(d_{\ell^1,\mathrm{H}}(T_n^p, T_n) > \eta n^{1/k}\Big) \leqslant \frac{C_q}{\eta^q} \mathbb{E}[(V_1)^{\frac{q}{k}-1}].$$

We then use an estimation of the density of V_1 found in Proposition 19 of [70] to obtain

$$\begin{aligned}\mathbb{E}[(V_1)^{\frac{q}{k}-1}] &\leqslant \frac{\Gamma(p+1+\frac{1}{k})}{\Gamma(p+\frac{2}{k})\Gamma(1-\frac{1}{k})}\int_0^1 u^{\frac{q-1}{k}-2}(1-u)^{\frac{2}{k}+p-1}du \\ &\leqslant \frac{\Gamma(p+1+\frac{1}{k})\Gamma(\frac{q-1}{k}-1)\Gamma(p+\frac{2}{k})}{\Gamma(p+\frac{2}{k})\Gamma(1-\frac{1}{k})\Gamma(p-1+\frac{q+1}{k})}.\end{aligned}$$

As p goes to infinity, this is, up to a constant, equivalent to $p^{2-\frac{q}{k}}$, which tends to 0 if we take $q > 2k$, thus ending the proof. □

3.3 Gromov-Hausdorff convergence

Lemma 3.4. *As p tends to infinity, we have the following convergence, in the sense of Hausdorff convergence for compact subsets of ℓ^1:*

$$\mathcal{T}^p \xrightarrow{a.s.} \mathcal{T}.$$

In particular, the tree $\mathcal{T}$ is in fact compact and $\mathcal{T}^p$ converges a.s. to $\mathcal{T}$ in the Gromov-Hausdorff sense.

Proof. Let us first prove that the sequence $(\mathcal{T}^p)_{p\in\mathbb{N}}$ is Cauchy in probability for the Hausdorff distance in ℓ^1, in the sense of [50], Chapter 3: we want to show that, for any $\varepsilon > 0$ and $\eta > 0$, if p and q are large enough, $\mathbb{P}\big(d_{\ell^1,\mathrm{H}}(\mathcal{T}^p,\mathcal{T}^q) > \eta\big) < \varepsilon$. Let therefore $\eta > 0$ and $\varepsilon > 0$. We have, for integers p and q,

$$\begin{aligned}\mathbb{P}\left(d_{\ell^1,\mathrm{H}}(\mathcal{T}^p,\mathcal{T}^q) > \eta\right) &= \mathbb{P}\left(\lim_{n\to\infty} n^{-1/k} d_{\ell^1,\mathrm{H}}(T_n^p,T_n^q) > \eta\right) \\ &\leqslant \liminf_{n\to\infty} \mathbb{P}\left(n^{-1/k} d_{\ell^1,\mathrm{H}}(T_n^p,T_n^q) > \eta\right) \\ &\leqslant \liminf_{n\to\infty} \mathbb{P}\left(d_{\ell^1,\mathrm{H}}(T_n^p,T_n) + d_{\ell^1,\mathrm{H}}(T_n^q,T_n) > n^{1/k}\eta\right) \\ &\leqslant \limsup_{n\to\infty} \mathbb{P}\left(d_{\ell^1,\mathrm{H}}(T_n^p,T_n) > n^{1/k}\frac{\eta}{2}\right) + \limsup_{n\to\infty} \mathbb{P}\left(d_{\ell^1,\mathrm{H}}(T_n^q,T_n) > n^{1/k}\frac{\eta}{2}\right).\end{aligned}$$

Thus, by Proposition 3.2, choosing p and q large enough yields

$$\mathbb{P}\Big(d_{\ell^1,\mathrm{H}}(\mathcal{T}^p,\mathcal{T}^q) > \eta\Big) \leqslant \varepsilon.$$

Since the Hausdorff metric on the set of nonempty compact subsets of ℓ^1 is complete, the sequence $(\mathcal{T}^p)_{p\in\mathbb{N}}$ does converge in probability, and thus has an a.s. converging subsequence. Since it is also monotonous (in the sense of inclusion of subsets), it in fact does converge to a limit we call $\mathcal{L}$, and we only need to show that $\mathcal{L} = \mathcal{T}$. Since $\mathcal{L}$ is a compact subset of ℓ^1 and contains $\mathcal{T}^p$ for all p, we have $\mathcal{T} \subset \mathcal{L}$. On the other hand, assuming the existence of a point $x \in \mathcal{L}\setminus\mathcal{T}$ would yield $\varepsilon > 0$ such that $d_{\ell^1}(x,\mathcal{T}) \geqslant \varepsilon$ and also $d_{\ell^1}(x,\mathcal{T}^p) \geqslant \varepsilon$ for all p, negating the Hausdorff convergence of $\mathcal{T}^p$ to $\mathcal{L}$. □

Proposition 3.5. *We have*

$$\frac{T_n}{n^{1/k}} \xrightarrow{\mathbb{P}} \mathcal{T}$$

as n goes to infinity, in the Gromov-Hausdorff sense.

Proof. All the work has already been done, we only need to stick the pieces together. Let $n \in \mathbb{N}$ and $p \leqslant n$, we use the triangle inequality:

$$d_{\mathrm{GH}}\left(\frac{T_n}{n^{1/k}}, \mathcal{T}\right) \leqslant d_{\mathrm{GH}}\left(\frac{T_n}{n^{1/k}}, \frac{T_n^p}{n^{1/k}}\right) + d_{\mathrm{GH}}\left(\frac{T_n^p}{n^{1/k}}, \mathcal{T}^p\right) + d_{\mathrm{GH}}(\mathcal{T}^p, \mathcal{T}).$$

For $\eta > 0$, we then have

$$\begin{aligned} &\mathbb{P}\left(d_{\mathrm{GH}}\left(\frac{T_n}{n^{1/k}}, \mathcal{T}\right) > \eta\right) \\ \leqslant \ &\mathbb{P}\left(d_{\mathrm{GH}}\left(\frac{T_n}{n^{1/k}}, \frac{T_n^p}{n^{1/k}}\right) > \frac{\eta}{3}\right) + \mathbb{P}\left(d_{\mathrm{GH}}\left(\frac{T_n^p}{n^{1/k}}, \mathcal{T}^p\right) > \frac{\eta}{3}\right) + \mathbb{P}\left(d_{\mathrm{GH}}(\mathcal{T}^p, \mathcal{T}) > \frac{\eta}{3}\right). \end{aligned}$$

Let $\varepsilon > 0$. By Lemma 3.4 and Proposition 3.2, there exists p such that the third term of the sum is smaller than ε, and the first term also is for all n large enough. Apply then Proposition 3.1 with this fixed p, to make the second term smaller than ε for large n, and the proof is over. □

3.4 Adding in the measures

We now know that $\mathcal{T}$ is compact. This compactness will enable us to properly obtain a measure on $\mathcal{T}$ and the desired GHP convergence. For all n and $p \leqslant n$, let μ_n^p be the image of μ_n by the projection from T_n to T_n^p (see Chapter 1, Section 4 for the definition of the projection). Let also, for all p, π^p be the projection from $\mathcal{T}$ to $\mathcal{T}^p$. We start by proving an extension of Proposition 3.1 to the measured case.

Proposition 3.6. *There exists a probability measure μ^p on $\mathcal{T}^p$ such that, in the GHP sense,*

$$\left(\frac{T_n^p}{n^{1/k}}, \mu_n^p\right) \overset{a.s.}{\to} (\mathcal{T}^p, \mu^p).$$

What's more, we have, for $p' \geqslant p$, $\mu^p = (\pi^p)_* \mu^{p'}$.

Proof. We aim to apply Lemma 3.2 from Chapter 1. The trees T_n^p and $\mathcal{T}^p$ in ℓ^1 with the stick-breaking method, by sequentially adding the leaves according to their indices, as recalled at the end of the proof of Proposition 3.1.

The first step to apply Chapter 1, Lemma 3.2 is then to find an appropriate dense subset of $\mathcal{T}^p$. Since we know from Section 2 that the distribution of the metric space $\mathcal{T}$ is that of a fragmentation tree and that the dislocation measure ν_k has infinite total mass, Theorem 1 from [41] tells us that it is leaf-dense. As a consequence, its branch points are also dense. Let S_p be the set of points of $\mathcal{T}^p$ which are also branch points of $\mathcal{T}$, we then know that S_p is a dense subset of $\mathcal{T}^p$. In fact S_p can be simply explicited:

$$S_p = \{L^i \wedge L^j;\ i \leqslant (k-1)p+1 \text{ or } j \leqslant (k-1)p+1\}$$

(recall that $\{L^i, i \geqslant 1\}$ is the set of leaves of $\mathcal{T}$ that belong to $\cup_{p=0}^{\infty} \mathcal{T}^p$). Let i and j be integers such that either i or j is smaller than or equal to $(k-1)p+1$, and let $x = L^i \wedge L^j$. For n such that $i \leqslant (k-1)n+1$ and $j \leqslant (k-1)n+1$, define x_n as the branch point in T_n of L_n^i and L_n^j. It is immediate that x_n converges to x, and moreover, calling $(T_n^p)_{x_n}$ the subtree of descendants of x_n in (T_n^p), that $(T_n^p)_{x_n}$ converges to $\mathcal{T}_x^p$ (the subtree of descendants of x in $\mathcal{T}^p$) in the Hausdorff sense in ℓ^1. What is left for us to do is to prove that $\mu_n^p\big((T_n^p)_{x_n}\big) = \mu_n\big((T_n)_{x_n}\big)$ converges a.s. as

n goes to infinity. To this effect, we let Z_n be the number of internal nodes of $(T_n)_{x_n}$, including x_n itself. Since we have

$$\mu_n\big((T_n)_{x_n}\big) = \frac{(k-1)Z_n+1}{(k-1)n+1},$$

convergence of $\mu_n\big((T_n)_{x_n}\big)$ as n goes to infinity is equivalent to convergence of $n^{-1}Z_n$. However the distribution of Z_n is governed by a simple recursion: for all n, given Z_n, $Z_{n+1} = Z_n + 1$ with probability $(kZ_n)/(1+kn)$, while $Z_{n+1} = Z_n$ with the complementary probability. It is then easy to check that the rescaled process $(Z_n/(kn+1))_{n\in\mathbb{N}}$ is a nonnegative martingale, hence converges a.s.. Then, so does $\mu_n\big((T_n)_{x_n}\big) = n^{-1}Z_n$. Hence we can apply Chapter 1, Lemma 3.2 to conclude.

The fact that $\mu^p = (\pi^p)_*\mu^{p+1}$ is then a direct consequence of the fact that $\mu_n^p = (\pi_n^p)_*\mu_n^{p+1}$ for all n: for any x in S_p, we have $\mu_n^p\big((T_n^p)_{x_n}\big) = \mu_n^{p+1}\big((T_n^{p+1})_{x_n}\big)$ and, letting n tend to infinity (and taking left-continuous versions in x as stated in Chapter 1, Lemma 3.2), we obtain $\mu^p\big((\mathcal{T}^p)_x\big) = \mu^{p+1}\big((\mathcal{T}^{p+1})_x\big)$, and Lemma 4.1 ends the proof. □

Lemma 3.7. *As p tends to infinity, μ^p converges a.s. to a probability measure μ on $\mathcal{T}$ which satisfies, for all p, $\mu^p = (\pi^p)_*\mu$*

Proof. Since $\mathcal{T}$ is compact, Lemma 2.1 shows that we can define a unique measure μ on $\mathcal{T}$ such that, for all p and $x \in \mathcal{T}^p$, $\mu(\mathcal{T}_x) = \mu^p(\mathcal{T}_x^p)$ (Proposition 3.6 assures us that this is well-defined since it does not depend on the choice of p). By definition, we then have $\mu^p = (\pi^p)_*\mu$ for all p, and Lemma 4.3 ends the proof. □

Proof of Theorem 1.1 (convergence in probability part). We want to prove that

$$\Big(\frac{T_n}{n^{1/k}}, \mu_n\Big) \xrightarrow{\mathbb{P}} (\mathcal{T}, \mu). \tag{3.7}$$

Once this will be done, the distribution of $(\mathcal{T}, \mu)$ will be that of the fragmentation tree mentioned in Theorem 1.1, since we have already proved the convergence in distribution to that measured tree in Section 2. To get (3.7), notice that Lemma 4.3 directly improves Proposition 3.2, since we can replace the GH distance by the GHP distance, adding the measures μ_n and μ_n^p respectively to the trees T_n and T_n^p. Once we know this, as well as Proposition 3.6 and Lemma 3.7, the same proof as that of Proposition 3.5 works. □

3.5 Joint convergence

For the sake of clarity, we return to the notations of the introduction: for $n \in \mathbb{Z}_+$, $T_n(k)$ is the tree at the n-th step of the algorithm, its scaling limit is $\mathcal{T}_k$. For $p \leqslant n$, we let $T_n^p(k)$ and $\mathcal{T}_k^p$ be the respective finite-dimensional marginals we have studied, endowed, respectively, with the probability measures $\mu_n^p(k)$ and μ_k^p. Let k' be an integer with $2 \leqslant k' < k$. Recall now that $T_n(k,k')$ is the subtree of $T_n(k)$ obtained by discarding all edges with labels greater than or equal to $k'+1$, as well as their descendants. The objective of this section is to prove the convergence in probability of $n^{-1/k}T_n(k,k')$ by using what we know of the convergence of $n^{-1/k}T_n(k)$. This method once again fails to give the distribution of the limiting tree, which will be obtained in Section 4.1.

For all n, the tree $T_n(k,k')$ comes with a measure $\mu_n(k,k')$ which is the image of $\mu_n(k)$ by the projection from $T_n(k)$ onto $T_n(k,k')$. Similarly, for $p \leqslant n$, define

$$T_n^p(k,k') = T_n(k,k') \cap T_n^p(k),$$

and the image measure $\mu_n^p(k,k')$. For fixed p, the almost sure convergence of $n^{-1/k}T_n^p(k)$ to $\mathcal{T}_k^p$ as n goes to infinity allows us to extend the edge labellings to $\mathcal{T}_k^p$, and thus define $\mathcal{T}_{k,k'}^p$ and $\mu_{k,k'}^p$ in analogous fashion. Note that the sequence $\big(n^{-1/k}T_n^p(k,k'),\mu_n^p(k,k')\big)$ converges almost surely to $(\mathcal{T}_{k,k'}^p,\mu_{k,k'}^p)$ as n goes to infinity, by using Lemmas 3.1 and 3.2 and imitating the proofs of Propositions 3.1 and 3.6. Finally, considering again versions of all these trees embedded in ℓ^1 via the stick-breaking construction, we let

$$\mathcal{T}_{k,k'} = \overline{\cup_{p=0}^{\infty}\mathcal{T}_{k,k'}^p}.$$

Clearly, $\mathcal{T}_{k,k'} \subset \mathcal{T}_k$ and we let $\mu_{k,k'}$ be the image of μ_k under the projection from $\mathcal{T}_k$ onto $\mathcal{T}_{k,k'}$.

Proof of Theorem 1.3 (convergence in probability part). What we want to show is that the sequence of measured trees $\big(n^{-1/k}T_n(k,k'),\mu_n(k,k')\big)$ converges in probability to $(\mathcal{T}_{k,k'},\mu_{k,k'})$ as n goes to infinity, and it is in fact a simple consequence of Lemma 4.2. Indeed, this lemma directly gives us the fact that, for $p\leqslant n$,

$$d_{\mathrm{GHP}}\left(\Big(\frac{T_n^p(k,k')}{n^{1/k}},\mu_n^p(k,k')\Big),\Big(\frac{T_n(k,k')}{n^{1/k}},\mu_n(k,k')\Big)\right)\leqslant d_{\mathrm{GHP}}\left(\Big(\frac{T_n^p}{n^{1/k}},\mu_n^p\Big),\Big(\frac{T_n}{n^{1/k}},\mu_n\Big)\right),$$

as well as, for any p,

$$d_{\mathrm{GHP}}\big((\mathcal{T}_{k,k'}^p,\mu_{k,k'}^p),(\mathcal{T}_{k,k'},\mu_{k,k'})\big)\leqslant d_{\mathrm{GHP}}\big((\mathcal{T}_k^p,\mu_k^p),(\mathcal{T}_k,\mu_k)\big).$$

Since we know that $(\mathcal{T}_k^p,\mu_k^p)\to(\mathcal{T}_k,\mu_k)$ a.s. as $p\to\infty$, that $\big(n^{-1/k}T_n^p(k,k'),\mu_n^p(k,k')\big)\to(\mathcal{T}_{k,k'}^p,\mu_{k,k'}^p)$ a.s. for all p as $n\to\infty$, and that there exists a GHP version of Proposition 3.2 (see the convergence in probability part of the proof of Theorem 1.1), the proof can then be ended just as that of Proposition 3.5. $\square$

4 Stacking the limiting trees

This section is devoted to the study of $\mathcal{T}_{k,k'}$, seen as a subtree of $\mathcal{T}_k$. We start by giving the distribution of the measured tree $(\mathcal{T}_{k,k'},\mu_{k,k'})$, then move on to prove (3.2), which is the last part of Theorem 1.3, and then finally show that, even without the construction algorithm, one can extract from $\mathcal{T}_k$ a tree distributed as $\mathcal{T}_{k,k'}$.

4.1 The distribution of $(\mathcal{T}_{k,k'},\mu_{k,k'})$

In Section 2, the distribution of $\mathcal{T}_k$ was obtained by using the main theorem of [42]. We would like to do the same with $\mathcal{T}_{k,k'}$, but the issue is that the results of [42] are restricted to conservative fragmentations. The aim of this section is therefore to concisely show that the arguments used in [42] still apply in our context and prove the last part of Theorem 1.3: that $(\mathcal{T}_{k,k'},d,\rho,\mu_{k,k'})$ has the distribution of a fragmentation tree with index $-1/k$ and dislocation measure $\nu_{k,k'}^{\downarrow}$ (which was defined in Theorem 1.3). For reference, we let $(\mathcal{T}^0,d^0,\rho^0,\mu^0)$ be such a fragmentation tree.

We are going to use the method of finite-dimensional marginals introduced in Chapter 1, Section 1.4. In our cases, for both $(\mathcal{T}_{k,k'},\mu_{k,k'})$ and $(\mathcal{T}^0,\mu^0)$, we know that the measure is fully supported on the tree. For $\mathcal{T}^0$, this is because it is a self-similar fragmentation tree with infinite dislocation measure, by using [41, Theorem 1]. The same theorem shows that μ_k is fully supported on $\mathcal{T}_k$, and by projection, $\mu_{k,k'}$ is fully supported on $\mathcal{T}_{k,k'}$. What this entails is that, since $(n^{-1/k}T_n(k,k'),\mu_n(k,k'))$ converges in probability to $(\mathcal{T}_{k,k'},\mu_{k,k'})$, we can prove that $(\mathcal{T}_{k,k'},\mu_{k,k'})$ and $(\mathcal{T}^0,\mu^0)$ have the same distribution by showing that the finite-dimensional marginals of $(n^{-1/k}T_n(k,k'),\mu_n(k,k'))$ converge to those of $(\mathcal{T}^0,\mu^0)$.

Our method of proof will use interpretations of those trees using partition-valued processes.

4.1.1 Partitions of finite sets with holes

In order to interpret the sequence $\big(T_n(k,k')\big)_{n\in\mathbb{Z}_+}$ as a sequence of Markov branching trees, we are going to need an alternative notion of partitions of finite sets. Let A be a finite subset of $\mathbb{N}$. We call partition with holes of A any partition π of A in the usual sense except that some of the singletons will be marked as holes. Holes are marked by saying that they have "size" zero: $\#\pi_i = 0$, even though the singleton block π_i is not actually empty. We then let $\lambda(\pi)$ be the list of the sizes of the blocks of π, in decreasing order, without writing the zeroes. For example, if $A=[5]$, then the partition $\pi = \big(\{1,4,5\},\{2\},\{3\}\big)$ is a partition with one hole when we specify $\#\{2\}=0$ and $\#\{3\}=1$, and then $\lambda(\pi)=(3,1)$. We let $\mathcal{P}'_A$ be the set of partitions with holes of A.

The main reason we introduce these new partitions is to clearly identify the nature of singletons of partitions of $\mathbb{N}$ when restricted to finite sets. If π is a partition of $\mathbb{N}$, then the partition $\pi\cap A$ of A really has two kinds of singletons: those who are singletons of π, and which will be singletons for any choice of A, and those who just happen to be singletons because A is not large enough. This is why we now identify $\pi\cap A$ as a partition with holes by saying that, if an integer $i\in A$ is in a singleton of π then it is in a hole of $\pi\cap A$.

4.1.2 Seeing $\big(T_n(k,k')\big)_{n\in\mathbb{Z}_+}$ as non-conservative Markov branching trees

As in Section 2.2, we will want to match the approach of [42] and index the trees not by their number of internal nodes, but by the number of leaves of $T_n(k)$. This is why we let, for $n\in\mathbb{N}$, $r_n=(k-1)n+1$, and then $(T^\circ_{r_n},\mu^\circ_{r_n}) = \big(T_n(k),\mu_n(k)\big)$ and $(T^\bullet_{r_n},\mu^\bullet_{r_n}) = \big(T_n(k,k'),\mu_n(k,k')\big)$. Similarly, we let $(q^\circ_{r_n})^\downarrow$ and $(q^\bullet_{r_n})^\downarrow$ be the associated splitting distributions: $(q^\circ_{r_n})^\downarrow$ is the distribution of $\big((k-1)X_i+1\big)_{i\in[k]}$ where $(X_i)_{i\in[k]}$ has distribution $q^\downarrow_{n-1}$, while $(q^\bullet_{r_n})^\downarrow$ is the distribution of $\big((k-1)X_i+1\big)_{i\in[k'']}$ where $(X_i)_{i\in[k]}$ has distribution $(q'_{n-1})^\downarrow$. For coherence, we also let $(q^\bullet_1)^\downarrow = (\bar{q}^\bullet_1)^\downarrow = \delta_{(0)}$, the Dirac mass on the sequence with only term 0.

For $r\in(k-1)\mathbb{N}+1$, and A a subset of $\mathbb{N}$ such that $\#A=r$, let p°_A be the probability measure on $\mathcal{P}'_A$ defined the following way: take a partition λ of r with distribution $(q^\circ_r)^\downarrow$ and then choose a random partition π of A such that $\lambda(\pi)=\lambda$, uniformly amongst such partitions. Similarly, we let $p^\bullet_A$ be the distribution of the partition obtained if we take λ with distribution $(q^\bullet_r)^\downarrow$ instead.

We will want to represent $T^\bullet_r$ as some discrete variant of a "fragmentation tree". To do so, first let $P_1(r),\dots,P_r(r)$ be a uniform ordering of the leaves of T°_r. Next, let $Q_1(r),\dots,Q_r(r)$ be their respective projections on $T^\bullet_r$. Now, for $n\in\mathbb{Z}_+$, let $\Pi^{(r)}(n)$ be the partition of $[r]$ defined the following way: an integer i is in a hole if $ht\big(Q_i(r)\big)\leqslant n$ and, for two integers i and j which are not in a hole, they are in the same block if $ht\big(Q_i(r)\wedge Q_j(r)\big) > n$. Since the sequence $\big(T_n(k)\big)_{n\in\mathbb{Z}_+}$ is Markov branching, the distribution of $(\Pi^{(r)}(n))_{n\in\mathbb{Z}_+}$ can be described by the following induction:

- $\Pi^{(r)}(0)=[r]$ almost surely.
- Conditionally on $\Pi^{(r)}(n)=(\pi_1,\dots,\pi_p)$, let $\big(\pi'(l)\big)_{l\in[p]}$ be independent variables, each with distribution $p^\bullet_{\pi_l}$, then $\Pi^{(r)}(n+1)$ has the same distribution as the partition obtained by fragmenting every block π_l into $\pi'(l)$.

This makes $T^\bullet_r$ a discrete analogue of a fragmentation tree, also equipped with a set of death points. Moreover, the measure $\mu^\bullet_r$ is the empirical measure associated to the death points $Q_1(r),\dots,Q_r(r)$.

We also mention the natural coupling between $p^\circ_{[r]}$ and $p^\bullet_{[r]}$ which comes from this construction. Of course, $\Pi^{(r)}(1)$ has distribution $p^\bullet_{[r]}$. Consider the partition without holes π' of $[r]$ obtained by saying that two integers i and j are in the same block if $ht\big(P_i(r)\wedge P_j(r)\big)>1$, where $P_1(r),\dots,P_r(r)$ are the earlier mentioned leaves in uniform order. This partition has distribution $p^\circ_{[r]}$, and since the holes of $\Pi^{(r)}(1)$ correspond to the integers i such that $P_i(r)\neq Q_i(r)$ and $ht(Q_i)=1$, the non-hole blocks of $\Pi^{(r)}(1)$ are all indeed blocks of π'.

4.1.3 Convergence of 1-dimensional marginals

For $r\in(k-1)\mathbb{Z}_++1$, let $X(r)\in T^\bullet_r$ have distribution $\mu^\bullet_r$. We will show that, as n tends to infinity, $n^{-1/k}d\big(\rho,X(r_n)\big)$ converges in distribution to the height of a point of $\mathcal{T}^0$ with distribution μ^0. Our proof will essentially be the same as that of Lemma 28 in [42], and will use the same main ingredient which is Theorem 2 of [39].

By exchangeability, we can assume that $X(r)=Q_1(r)$. This enables us to use a Markov chain: let $M_m=\#\Pi_1(m)$ for $m\in\mathbb{Z}_+$. This is a decreasing Markov chain on $\mathbb{Z}_+$, with starting value r and for which 0 is an absorbing state. Moreover, its transition probabilites do not depend on r, and, calling them $(p_{a,b})_{a,b\in\mathbb{Z}_+}$, are simply characterized by

$$\sum_{b=0}^{r} f(b)p_{r,b}=\sum_{\lambda=(\lambda_1,\dots,\lambda_{k'})\in\mathbb{Z}_+^{k'}:\sum_{i=1}^{k'}\lambda_i\leqslant r}(q^\bullet_r)^{\downarrow}(\lambda)\Big(\sum_{i=1}^{k'}f(\lambda_i)\frac{\lambda_i}{r}+f(0)\frac{r-\sum_{i=1}^{k'}\lambda_i}{r}\Big)$$

for a measurable function f.

With this and Corollary 2.2, it then follows that the measure $n^{1/k}(1-x)\sum_{b=0}^{r_n}p_{r_n,b}\delta_{b/r_n}(\mathrm{d}x)$ converges weakly to

$$\int_{\mathcal{S}_{k',\leqslant}}\Big(\sum_{i=1}^{k'}(1-s_i)s_i\delta_{s_i}+\Big(1-\sum_{i=1}^{k'}s_i\Big)\delta_0\Big)\nu^{\downarrow}_{k,k'}(\mathrm{d}\mathbf{s}).$$

Theorem 2 of [39] is then applicable and shows that, when renormalized by $n^{-1/k}$, the height of $X(r_n)$ does converge in distribution to the height of a point of $\mathcal{T}^0$ with distribution μ^0, which, as explained in Chapter 2, Section 2.1.4 can be written $\int_0^\infty \mathrm{e}^{-\xi_t/k}\mathrm{d}t$, where $(\xi_s)_{s\geqslant 0}$ is a subordinator with Laplace exponent defined for $q\geqslant 0$ by $\int_{\mathcal{S}_{k',\leqslant}}(1-\sum_{i=1}^{k'}s_i^{q+1})\nu^{\downarrow}_{k,k'}(\mathrm{d}\mathbf{s})$.

4.1.4 Convergence of l-dimensional marginals for $l\geqslant 2$

The general proof of convergence of the finite-dimensional marginals requires an induction, which we have already initialized at $l=1$. The main ingredient for the l-th step will be investigating the structure or $T^\bullet_r$ (respectively $\mathcal{T}^0$) around the branch point of $X_1(r),\dots,X_l(r)$ (respectively $X_1^0,\dots,X_l^0$), which are independent with distribution $\mu^\bullet_r$ (respectively μ^0) and this will be done by studying the respective discrete and continuous fragmentation processes at the time where l different integers split. More precisely, our aim is to prove the following proposition:

Proposition 4.1. *Let $(\Pi(t))_{t\geqslant 0}$ be a partition-valued fragmentation process with parameters $(-1/k,0,\nu_{k,k'})$. Recall from Chapter 2, Section 3.2 that $D_{[l]}$ is the time from which integers $1,\dots,l$ are not all in the same block of $(\Pi(t))_{t\geqslant 0}$. We let in analoguous fashion $D^{r_n}_{[l]}$ be the time from which $\Pi^{(r_n)}$ splits the block $[l]$. We then have the following convergence in distribution as* *n goes to infinity:*

$$\left(\frac{D_{[l]}^{r_n}}{n^{1/k}}, [l]\cap \Pi^{(r_n)}(D_{[l]}^{r_n}), \Big(\frac{\#\Pi_{(i)}^{(r_n)}(D_{[l]}^{r_n})}{r_n}, 1 \leqslant i \leqslant l\Big)\right) \qquad (4.8)$$
$$\underset{n\to\infty}{\Longrightarrow} \left(D_{[l]}, [l]\cap \Pi(D_{[l]}), \Big(|\Pi_{(i)}(D_{[l]}), 1 \leqslant i \leqslant l\Big)\right)$$

The proof of this proposition will be split in several steps. The following two lemmas state the distributions of the right-hand side and left-hand side of (4.8), while the one after that gives in some fashion the scaling limit of the measure $p^{\bullet}_{r_n}$. In all three lemmas, in order to manipulate information on $\pi \cap [l]$ (where π is a random partition studied in the current lemma), we look at an event where it is equal to a particular partition of $[l]$ called π'. This partition will have some holes, and we call $\pi'_{i_1}, \dots, \pi'_{i_b}$ the blocks of π' which are not holes. Once we are on this event, the information on the sizes of the blocks containing the first l integers is entirely contained in $\Big(|\pi_{i_j}|, 1 \leqslant j \leqslant b\Big)$ (replacing the asymptotic frequency by cardinality if π is a partition of a finite set).

Lemma 4.2. *If f and g are functions from $(0,\infty)$ to $\mathbb{R}$ and h from $(0,\infty)^b$ to $\mathbb{R}$, all nonnegative and measurable, we have*

$$\mathbb{E}\left[f\big(D_{[l]}\big)g\big(|\Pi_1(D^-_{[l]}|)\big)h\Big(\big(\frac{|\Pi_{i_j}(D_{[l]})|}{|\Pi_1(D^-_{[l]}|)}\big), 1 \leqslant j \leqslant b\Big)\mathbb{1}_{\{\Pi(D_{[l]})\cap[l]=\pi'\}}\right]$$
$$= \int_0^\infty f(u)\mathrm{d}u\mathbb{E}\Big[|\Pi_1(u)|^{l-1+1/k}g(|\Pi_1(u)|)\Big]\int_{\mathcal{P}_{\mathbb{N}}}\kappa^{\downarrow}_{\nu_{k,k'}}(\mathrm{d}\pi)h(|\pi_{i_j}|, 1 \leqslant j \leqslant b)\mathbb{1}_{\{\pi\cap[l]=\pi'\}}$$

Proof. This is an elementary computation involving the Poissonian construction and the Lamperti-type time-change, which were both explained in Chapter 2, Section 2.1.3. We do not reproduce the computation which was already done in [42], Proposition 18. □

For the next lemma, we introduce the notation $(a)_b = a(a-1)\dots(a-b+1)$ for nonnegative integers a and b, and also add $(-1)_b = 0$ for any b.

Lemma 4.3. *If f and g are functions from $(0,\infty)$ to $\mathbb{R}$ and h from $(0,\infty)^b$ to $\mathbb{R}$, all nonnegative and measurable, we have*

$$\mathbb{E}\left[f\big(D_{[l]}^{(r)}\big)g\big(\#\Pi_1^{(r)}(D_{[l]}^{(r)}-1)\big)h\Big(\#\Pi_{i_j}^{(r)}(D_{[l]}^{(r)}), 1 \leqslant j \leqslant b\Big)\mathbb{1}_{\{[l]\cap\Pi^{(r)}(D_{[l]}^{(r)})=\pi'\}}\right]$$
$$= \sum_{t\in\mathbb{N}}\mathbb{E}\left[\frac{(\#\Pi_1^{(r)}(t-1)-1)_{l-1}}{(r-1)_{l-1}}f(t)g\big(\#\Pi_1^{(r)}(t-1)\big)p^{\bullet}_{\Pi^{(r)}(t-1)}\Big(h(|\pi_{i_j}|, 1 \leqslant j \leqslant b)\mathbb{1}_{\{[l]\cap\pi=\pi'\}}\Big)\right]$$

Proof. We refer to the proof of Lemma 27 in [42]. The general idea is that $t \in \mathbb{N}$ represents the possible values of $D_{[l]}^{(r)}$ and that the fraction $\frac{(\#\Pi_1^{(r)}(t-1)-1)_{l-1}}{(r-1)_{l-1}}$ is the probability that $[l]$ is not yet split at time $t-1$, conditionally on $(\#\Pi_1^{(r)}(p), p \leqslant t-1)$. □

Lemma 4.4. *Let $g : (0,\infty)^b \to \mathbb{R}$ be a continuous function with compact support, then*

$$(r_n)^{1/k}p^{\bullet}_{[r_n]}\Big(g\big(\frac{\#\pi_{i_j}}{r_n}, 1 \leqslant j \leqslant b\big)\mathbb{1}_{\{\pi\cap[l]=\pi'\}}\Big) \underset{n\to\infty}{\to} \int_{\mathcal{P}_{\mathbb{N}}}\kappa_{\nu_{k,k'}}(\mathrm{d}\pi)g(|\pi_{i_j}|, 1 \leqslant j \leqslant b)\mathbb{1}_{\{\pi\cap[l]=\pi'\}}$$

Proof. This is our equivalent version of Lemma 26 in [39]. We could prove it by using similar methods (though the presence of holes would require several modifications), however we can here use the coupling between $p^{\bullet}_{[r_n]}$ and $p^{\circ}_{[r_n]}$ explained at the end of Section 4.1.2 and not do any computations.

Consider all partitions of $[l]$ with holes π'' such that any non-hole block of π' is a block of π''. We call these blocks $\pi''_{\tilde{i}_1},\dots,\pi''_{\tilde{i}_b}$. The coupling between $p^{\bullet}_{[r_n]}$ and $p^{\circ}_{[r_n]}$ then shows that

$$p^{\bullet}_{[r_n]}\Big(g\big(\frac{\#\pi_{i_j}}{r_n},1\leqslant j\leqslant b\big)\mathbb{1}_{\{\pi\cap[l]=\pi'\}}\Big)=\sum_{\pi''}p^{\circ}_{[r_n]}\Big(g\big(\frac{\#\pi_{\tilde{i}_j}}{r_n},1\leqslant j\leqslant b\big)\mathbb{1}_{\{\pi\cap[l]=\pi''\}}\Big),$$

where we sum over all such possible π''. There also exists a coupling between the measures $\kappa_{\nu_{k,k'}}$ and κ_{ν_k} such that all non-hole blocks of one partition are blocks of the other one. It can be obtained by slightly specifying how we do the paintbox construction. For $\mathbf{s}\in\mathcal{S}_k$, let $(U_i)_{i\in\mathbb{N}}$ be independent uniform variables in $[0,1]$, and let two partitions π^1 and π^2 be defined such: two integers i and j are in the same block of π^1 and π^2 if U_i and U_j are in the same interval of the form $[\sum_{p=1}^m s_p,\sum_{p=1}^{m+1}s_p)$, with the exception that, for i such that $U_i\geqslant\sum_{p=1}^{k'}s_p$, i is in a singleton of π^2. Taking a "random" sequence $\mathbf{s}$ with "distribution" ν_k yields our coupling, and we thus have

$$\int_{\mathcal{P}_{\mathbb{N}}}\kappa_{\nu_{k,k'}}(\mathrm{d}\pi)g(|\pi_{i_j}|,1\leqslant j\leqslant b)\mathbb{1}_{\{\pi\cap[l]=\pi'\}}=\sum_{\pi''}\int_{\mathcal{P}_{\mathbb{N}}}\kappa_{\nu_k}(\mathrm{d}\pi)g(|\pi_{\tilde{i}_j}|,1\leqslant j\leqslant b)\mathbb{1}_{\{\pi\cap[l]=\pi''\}}.$$

Now, since the measures $q^{\circ}_{r_n}$ and ν_k are conservative, we can use Lemma 26 from [39], and we have, for any fixed π'',

$$(r_n)^{1/k}p^{\circ}_{[r_n]}\Big(g\big(\frac{\#\pi_{\tilde{i}_j}}{r_n},1\leqslant j\leqslant b\big)\mathbb{1}_{\{\pi\cap[l]=\pi''\}}\Big)\underset{n\to\infty}{\to}\int_{\mathcal{P}_{\mathbb{N}}}\kappa_{\nu_k}(\mathrm{d}\pi)g(|\pi_{\tilde{i}_j}|,1\leqslant j\leqslant b)\mathbb{1}_{\{\pi\cap[l]=\pi''\}}.$$

The proof is then ended by summing over all possible π''. □

Proof of Proposition 4.1. There is no significant difference between the proof of Proposition 4.1 and that of Lemma 29 in [42], so we only give the general idea here. We take a partition with holes π' of $[l]$, and compactly supported functions f, g and h, and look at the quantity

$$\mathbb{E}\left[f\Big(\frac{D^{(r_n)}_{[l]}}{n^{1/k}}\Big)g\Big(\frac{\#\Pi^{(r_n)}_1(D^{(r_n)}_{[l]}-1)}{r_n}\Big)h\Big(\frac{\#\Pi^{(r)}_{i_j}(D^{(r_n)}_{[l]})}{\#\Pi^{(r_n)}_1(D^{(r_n)}_{[l]}-1)},1\leqslant j\leqslant b\Big)\mathbb{1}_{\{[l]\cap\Pi^{(r)}(D^{(r_n)}_{[l]})=\pi'\}}\right].$$

We apply Lemma 4.3 and normalize $t\in\mathbb{N}$ by $n^{1/k}$ to obtain that this expecation is equal to

$$\int_{n^{-1/k}}^{\infty}f\Big(\frac{\lfloor un^{1/k}\rfloor}{n^{1/k}}\Big)du\mathbb{E}\left[\frac{(\Pi^{(r_n)}_1\lfloor n^{1/k}u-1\rfloor)_{l-1}}{(r_n-1)_{l-1}}g\Big(\frac{\Pi^{(r_n)}_1(\lfloor n^{1/k}u-1\rfloor)}{r_n}\Big)\right.$$
$$\left.n^{1/k}p^{\bullet}_{\Pi^{(r_n)}(\lfloor n^{1/k}u-1\rfloor)-1)}h\Big(\frac{\#\pi_{i_j}}{\Pi^{(r_n)}_1(\lfloor n^{1/k}u-1\rfloor)},1\leqslant j\leqslant b\Big)\right].$$

Applying Lemma 4.4 and using the dominated convergence theorem finally leaves us with all the terms of Lemma 4.2. □

Proof by induction of convergence of l-dimensional marginals. Let $X_1(r_n),\dots,X_l(r_n)$ be l independent points of $T_n(k,k')$ distributed as $\mu_n(k,k')$ conditionally on $\big(T_n(k,k'),\mu_n(k,k')\big)$.

We set out to prove that the tree

$$\Big(\cup_{i=1}^l [\![\rho, X_i(r_n)]\!], \frac{1}{l}\sum_{i=1}^l \delta_{X_i(r_n)}\Big).$$

converges to

$$\Big(\cup_{i=1}^l [\![\rho, X_i^0]\!], \frac{1}{l}\sum_{i=1}^l \delta_{X_i^0}\Big).$$

where $X_1^0, \ldots, X_l$ are independent points of $\mathcal{T}^0$ with distribution μ^0. We know of course that X_i^0 can be seen as the death point of i in a fragmentation process, and first show that the same is true for the $X_i(r_n)$.

We first want to show that we can replace $X_i(r_n)$ by $Q_i^{r_n}$ for all $i \in [l]$ and not change the limit. To do this, first let $Y_1(r_n), \ldots, Y_l(r_n)$ be l independent points of $T_n(k)$ with distribution $\mu_n(k)$. Conditionally on the event where they are all distinct, they are distributed as $P_1(r_n), \ldots, P_l(r_n)$, and therefore, taking $X_i(r_n)$ as the projection on $T_n(k,k')$ of $Y_i(r_n)$ for all i, we get that $(X_1(r_n), \ldots, X_l(r_n))$ are distributed as $Q_1^{r_n}, \ldots, Q_l^{r_n}$. Moreover, the probability of the event of conditioning tends to 1 as n goes to infinity, since this probability is equal to $r_n^{-l+1}(r_n-1)_{l-1}$. This means that we can now focus on proving the convergence in distribution of

$$\Big(\bigcup_{i=1}^l [\![\rho, Q_i^{(r_n)}]\!], \frac{1}{l}\sum_{i=1}^l \delta_{Q_l^{(r_n)}}\Big)$$

to

$$\Big(\bigcup_{i=1}^l [\![\rho, Q_i^0]\!], \frac{1}{l}\sum_{i=1}^l \delta_{Q_i^0}\Big),$$

where $Q_1^0, \ldots, Q_l^0$ are the death points of $1, \ldots, l$ in a $(-1/k, 0, \nu_{k,k'}^{\downarrow})$-fragmentation process. The distribution of $\big(\cup_{i=1}^l [\![\rho, Q_i^{(r_n)}]\!], l^{-1}\sum_{i=1}^l \delta_{Q_l^{(r_n)}}\big)$ can be characterized by saying that it is the tree formed by having an initial segment of length $D_{[l]}^{r_n}$ ending at a branch point, to which we append independent trees, the distribution of which are determined by $\Pi^{(r_n)}(D_{[l]}^{r_n}) \cap [l]$ and $(\#\Pi_{(i)}^{(r_n)})_{i\in[l]}$. Precisely, let $\pi' = \Pi^{(r_n)}(D_{[l]}^{r_n}) \cap [l]$ and let $i_1, \ldots, i_b$ be the indices of the non-hole blocks of π'. Then, for all $j \in [b]$, let $n_j = \#\Pi_{i_j}^{(r_n)}(D_{[l]}^{r_n})$ and $m_j = \#\pi'_{i_j}$ and graft a copy of the m_j-dimensional marginal of $T_{n_j}^\bullet$ where the measure has been renormalized to have total mass m_j/l. Moreover, we give to the branchpoint mass $1 - \sum_{j=1}^b m_j/l$. Applying Proposition 4.1 as well as the induction hypothesis for each subtree then ends the proof. □

4.2 Proof of (3.2)

For $n \geqslant 0$, let I_n denote the number of internal nodes of $T_n(k)$ which are in $T_n(k,k')$.

Lemma 4.5. *One has*

$$(T_n(k,k'), n \geqslant 0) = \big(\tilde{T}_{I_n}(k'), n \geqslant 0\big),$$

where $(\tilde{T}_i(k'), i \geqslant 0)$ is a sequence distributed as $(T_i(k'), i \geqslant 0)$ and independent of $(I_n, n \geqslant 0)$. Moreover, $(I_n, n \geqslant 0)$ is a Markov chain with transition probabilities

$$\mathbb{P}\left(I_{n+1} = i+1 \mid I_n = i\right) = 1 - \mathbb{P}\left(I_{n+1} = i \mid I_n = i\right) = \frac{k'i+1}{kn+1},$$

and as a consequence,

$$\frac{I_n}{n^{k'/k}} \xrightarrow{\text{a.s.}} M_{k'/k,1/k},$$

where the limit is a $(k'/k, 1/k)$-generalized Mittag-Leffler random variable.

We recall that a generalized Mittag-Leffler random variable $M_{\alpha,\theta}$ with parameters $\alpha \in (0,1)$ and $\theta > -\alpha$ has its distribution characterized by its positive moments, given by

$$\mathbb{E}\left[M_{\alpha,\theta}^p\right] = \frac{\Gamma(\theta+1)\Gamma(\theta/\alpha+p+1)}{\Gamma(\theta/\alpha+1)\Gamma(\theta+p\alpha+1)}, \quad p \geqslant 0.$$

Proof. This proof is very similar to those of Lemma 8 and Lemma 9 of [24]. Given $T_i(k)$ and $T_i(k,k')$ for $0 \leqslant i \leqslant n$, the new node added to get $T_{n+1}(k)$ from $T_n(k)$ will belong to $T_{n+1}(k,k')$ if and only if the selected edge is in $T_n(k,k')$, which occurs with probability $(k'I_n+1)/(kn+1)$ since $k'I_n+1$ is the number of edges of $T_n(k,k')$ and $kn+1$ that of $T_n(k)$. Moreover, conditionally to the fact that this new node belongs to $T_{n+1}(k,k')$, it is located uniformly at random on one of the edges of $T_n(k,k')$, independently of the whole process $(I_n, n \geqslant 0)$ and of $T_i(k,k')$ for $0 \leqslant i \leqslant I_{I_n}^{-1} - 1$ where $I_m^{-1} := \inf\{n \geqslant 0 : I_n = m\}$, $m \geqslant 0$. From this, it should be clear that the process defined for all $i \geqslant 0$ by

$$\tilde{T}_i(k') = T_{I_i^{-1}}(k,k')$$

is distributed as $(T_i(k'), i \geqslant 0)$ and independent of $(I_n, n \geqslant 0)$. Moreover, we have that $T_n(k,k') = \tilde{T}_i(k')$ if $I_n = i$, hence $T_n(k,k') = \tilde{T}_{I_n}(k')$.

Lastly, the few lines above show that $(I_n, n \geqslant 0)$ is a Markov chain with the expected transition probabilities. It turns out that these probabilities are identical to those of the number of tables in a $(k'/k, 1/k)$ Chinese restaurant process. Therefore, using again Theorem 3.8 in [69], $n^{-k'/k}I_n$ converges almost surely towards a $(k'/k, 1/k)$-generalized Mittag-Leffler random variable. □

Proof of (3.2). This is a straightforward consequence of the joint convergence in probability settled in Theorem 1.3 and of Lemma 4.5. Indeed, we know that

$$\left(\frac{T_n(k)}{n^{1/k}}, \frac{T_n(k,k')}{n^{1/k}}\right) \underset{n\to\infty}{\overset{\mathbb{P}}{\to}} (\mathcal{T}_k, \mathcal{T}_{k,k'})$$

in the GH sense. Then, for $n \geqslant 1$,

$$\frac{T_n(k,k')}{n^{1/k}} = \frac{T_n(k,k')}{I_n^{1/k'}} \times \left(\frac{I_n}{n^{k'/k}}\right)^{1/k'}.$$

On the one hand, the left hand side converges in probability towards $\mathcal{T}_{k,k'}$. On the other hand, by Lemma 4.5 and since I_n converges a.s. to $+\infty$,

$$\frac{T_n(k,k')}{I_n^{1/k'}} \underset{n\to\infty}{\overset{\mathbb{P}}{\to}} \tilde{\mathcal{T}}_{k'},$$

where $\tilde{\mathcal{T}}_{k'}$ is distributed as $\mathcal{T}_{k'}$. Moreover this holds independently of the a.s. convergence of $I_n/n^{k'/k}$ towards the generalized Mittag-Leffler r.v. $M_{k'/k,1/k}$. The result follows by identification of the limits. □

4.3 Extracting a tree with distribution $\mathcal{T}_{k'}$ from $\mathcal{T}_k$

We know from the discrete approximation that there is a subtree of $\mathcal{T}_k$ which is distributed as $M^{1/k'}_{k'/k,1/k} \cdot \mathcal{T}_{k'}$ (or, equivalently, as a fragmentation tree with index $-1/k$ and dislocation measure $\nu^{\downarrow}_{k,k'}$). Our goal is now to explain how to extract such a tree directly from $\mathcal{T}_k$. Our approach strongly relies on the fact that $(\mathcal{T}_k, \mu_k)$ is a fragmentation tree.

As a fragmentation tree, $\mathcal{T}_k$ has a countable number of branch points, almost surely. We denote this set of branch points $\{b(n), n \in \mathbb{N}\}$. For each $n \in \mathbb{N}$, we recall that

$$\mathcal{T}_{b(n)} = \big\{v \in \mathcal{T}_k : b(n) \in [[\rho, v]]\big\}$$

is the subtree of descendants of $b(n)$ (ρ denotes the root of $\mathcal{T}_k$). Since $\mathcal{T}_k$ is k-ary, the set $\mathcal{T}_{b(n)}\backslash\{b(n)\}$ has exactly k connected components. We label them as follows: $\mathcal{T}_{b(n),1}$ is the connected component with the largest μ_k-mass, $\mathcal{T}_{b(n),2}$ is the connected component with the second largest μ_k-mass, and so on (if two or more trees have the same mass, we label them randomly).

For $n \in \mathbb{N}$ and $i = 1, ..., k$, let

$$s_i(n) = \frac{\mu_k(\mathcal{T}_{b(n),i})}{\mu_k(\mathcal{T}_{b(n)})}.$$

Almost surely, for all $n \in \mathbb{N}$, these quotients are well-defined, strictly positive and sum to 1. We then mark the sequences $\mathbf{s}(n)$, *independently* for all $n \in \mathbb{N}$, by associating to each sequence $\mathbf{s} \in \mathcal{S}_k$ an element $\mathbf{s}^* \in \mathcal{S}_{k',\leqslant}$ by deciding that for all $1 \leqslant i_1 < ... < i_{k'} \leqslant k$

$$(s^*_1, ..., s^*_{k'}) = (s_{i_1}, ..., s_{i_{k'}}) \text{ with probability } \frac{(k'-1)!(k-k')!}{(k-1)!} \frac{\sum_{j\in\{i_1,...,i'_k\}} \prod_{1\leqslant i\neq j\leqslant k}(1-s_i)}{\sum_{j=1}^k \prod_{1\leqslant i\neq j\leqslant k}(1-s_i)}. \tag{3.8}$$

This means that we attribute a weight $\prod_{i\neq j}(1-s_i)$ to the jth term of the sequence $\mathbf{s}$, for all $1 \leqslant j \leqslant k$, and then choose at random a k'-uplet of terms (with strictly increasing indices) with a probability proportional to the sum of their weights. One can easily check that, for any sequence $\mathbf{s}$, the quotient in (3.8) indeed defines a probability distribution since $(k-1)!/((k'-1)!(k-k')!)$ is the number of k'-uplets $(i_1, ..., i_{k'})$, with $1 \leqslant i_1 < ... < i_{k'} \leqslant k$, containing a given integer $j \in \{1, ..., k\}$. For $n \in \mathbb{N}$, if $(s^*_1(n), ..., s^*_{k'}(n)) = (s_{i_1}(n), ..., s_{i_{k'}}(n))$, we then let

$$\mathcal{T}^*_{b(n)} = \bigcup_{j\in\{1,...,k\}\backslash\{i_1,...,i_{k'}\}} \mathcal{T}_{b(n),j}.$$

Finally we set

$$\mathcal{T}^*_{k,k'} = \mathcal{T}_k \backslash \bigcup_{n\in\mathbb{N}} \mathcal{T}^*_{b(n)}. \tag{3.9}$$

In words, $\mathcal{T}^*_{k,k'}$ is obtained from $\mathcal{T}_k$ by removing all groups of trees $\mathcal{T}^*_{b(n)}$ for $n \in \mathbb{N}$. This tree (which is well-defined almost surely) has the required distribution:

Proposition 4.6. *The tree $\mathcal{T}^*_{k,k'}$ is a non-conservative fragmentation tree, with index of self-similarity $-1/k$ and dislocation measure $\nu^{\downarrow}_{k,k'}$.*

Proof. Let $(A_i)_{i\in\mathbb{N}}$ be an exchangeable sequence of leaves of $\mathcal{T}_k$ directed by μ_k. We know from Chapter 2, Proposition 3.2 that the partition-valued process $(\Pi(t))_{t\geqslant 0}$ obtained by declaring, for $t \geqslant 0$, that two different integers i and j are in the same block of $\Pi(t)$ if A_i and A_j are in the same connected component of $\{x \in \mathcal{T}_k, ht(x) > t\}$ is a partition-valued fragmentation process with dislocation measure $\nu_k^\downarrow$ and self-similarity index $-1/k$ (and no erosion). As explained in Chapter 2, Section 2.1.3, the process Π can be constructed from a Poisson point process $\big((\Delta(s), i(s)), s \geqslant 0\big)$ on $\mathcal{P}_{\mathbb{N}} \times \mathbb{N}$, with intensity measure $\kappa_{\nu_k^\downarrow} \otimes \#$, where $\#$ denotes the counting measure on $\mathbb{N}$ and $\kappa_{\nu_k^\downarrow}$ is the measure on $\mathcal{P}_{\mathbb{N}}$ associated to ν_k by the paintbox method.

We then mark the Poisson point process as follows just as we marked elements of $\mathcal{S}_k$ earlier: for each atom $(\Delta(s), i(s))$, we extract randomly k' blocks of $\Delta(s)$ by setting

$$(\Delta_1^*(s), ..., \Delta_{k'}^*(s)) = (\Delta_{i_1}(s), ..., \Delta_{i_{k'}}(s))$$

with probability

$$\frac{(k'-1)!(k-k')!}{(k-1)!} \frac{\sum_{j\in\{i_1,...,i'_k\}} \prod_{1\leqslant i\neq j\leqslant k}(1-|\Delta(s)_i|)}{\sum_{j=1}^k \prod_{1\leqslant i\neq j\leqslant k}(1-|\Delta(s)_i)|}.$$

Then, we make $\Delta^*(s)$ into a partition of $\mathbb{N}$ with dust by putting every integer which is not originally in a block $\Delta_1^*(s), ..., \Delta_{k'}^*(s)$ into a singleton. The process $\big((\Delta^*(s), i(s)), s \geqslant 0\big)$ is then a marked Poisson point process with intensity $\kappa_{\nu_k^{\downarrow,*}} \otimes \#$, where

$$\kappa_{\nu_k^{\downarrow,*}}(\mathrm{d}\pi) = \int_{\mathcal{S}_{k',\leqslant}} \kappa_{\mathbf{s}}(\mathrm{d}\pi)\nu_k^{\downarrow,*}(\mathrm{d}\mathbf{s}) \quad \text{and} \quad \int_{\mathcal{S}_{k',\leqslant}} f(\mathbf{s})\nu_k^{\downarrow,*}(\mathrm{d}\mathbf{s}) = \int_{\mathcal{S}_k} \mathbb{E}[f(\mathbf{s}^*)]\nu_k^\downarrow(\mathrm{d}\mathbf{s}),$$

for all suitable test functions f. Now, the key-point is that

$$\nu_k^{\downarrow,*} = \nu_{k,k'}^\downarrow.$$

This is easy to check by using the definitions of $\nu_k^\downarrow$, $\nu_{k,k'}^\downarrow$ and of the marking procedure (3.8), together with the identity (3.3). The details of this calculation are left to the reader.

To finish, let Π^* be the $(-1/k, 0, \nu_{k,k'}^\downarrow)$-fragmentation process derived from the Poisson point process $\big((\Delta^*(s), i(s)), s \geqslant 0\big)$. For all $i \in \mathbb{N}$, let $D_i^* = \inf\{t \geqslant 0, \{i\} \in \Pi^*(t)\}$ and note that $D_i^* \leqslant D_i$, where $D_i := \inf\{t \geqslant 0, \{i\} \in \Pi(t)\}$ is the height of A_i in $\mathcal{T}_k$. Let then A_i^* be the unique point of $\mathcal{T}_k$ belonging to the geodesic $[[\rho, A_i]]$ which has height D_i^*. It is not hard to see that $\mathcal{T}_{k,k'}^*$, defined by (3.9), is the closure of the subtree $\cup_{i\geqslant 1}[[\rho, A_i^*]]$ of $\mathcal{T}_k$ spanned by the root and all the vertices A_i^* (almost surely). But by definition this closure is the genealogy tree of Π^*, in the sense of Chapter 2, Proposition 3.5. Thus $\mathcal{T}_{k,k'}^*$ has the distribution of a $(-1/k, \nu_{k,k'}^\downarrow)$-fragmentation tree. $\square$

Chapter 4

Infinite multi-type Galton-Watson trees and infinite Boltzmann maps

Contents

We show that large critical multi-type Galton-Watson trees, when conditioned to have a large amount of vertices of one fixed type, converge locally in distribution to an infinite tree which is analoguous to Kesten's infinite monotype Galton-Watson tree. Use of the well-known Bouttier-Di Francesco-Guitter bijection then allows us to apply this result to the theory of random maps, showing that critical large Boltzmann-distributed random planar maps converge in distribution to an infinite variant, which is in fact a recurrent planar graph.

1 Introduction

A planar map is a proper embedding of a finite connected planar graph in the sphere, taken up to orientation-preserving homeomorphisms. These objects were first studied from a combinatorial point of view in the works of Tutte in the 1960s (see for example [77]), and have since been of use in different domains of mathematics, such as algebraic geometry (see for example [55]) and theoretical physics (as in [6]). There has been great progress in their probabilistic study ever since the work of Schaeffer [74], which has amongst other things led to finding the scaling limit of many large random maps (we mention [67] and [56]).

Our subject of interest here is the local convergence of large random maps, which means that we are not interested in scaling limits but in the combinatorial structure of a map around a chosen root. Such problems were first studied by Angel and Schramm ([7]) and Krikun ([53]), who showed that the distributions of uniform triangulations and quadrangulations with n vertices converge weakly as n goes to infinity. Each limit is the distribution of an infinite random map, respectively the uniform infinite planar triangulation (UIPT) and the uniform infinite planar quadrangulation (UIPQ). Of particular interest to us is the paper [26] where the convergence to the UIPQ is shown by a method involving the well-known Cori-Vauquelin-Schaeffer bijection ([74]).

We will generalize this to a large family of random maps called the class of Boltzmann-distributed random maps. Let $\mathbf{q} = (q_n)_{n\in\mathbb{N}}$ be a sequence of nonnegative numbers. We assign to every finite planar map a weight which is equal to the product of the weights of its faces, the weight of a face being q_d where d is the number of edges adjacent to said face, counted with multiplicity. If the sum of all the weights of all the maps is finite, then one can normalize this into a probability distribution.

The use of the so-called Bouttier-Di Francesco-Guitter bijection (see [19], or Section 4.3) allows us to obtain the convergence to infinite maps for a fairly large class of weight sequences $\mathbf{q}$. For $\mathbf{q}$ in this class, let (M_n, E_n) be a $\mathbf{q}$-Boltzmann rooted map conditioned to have n vertices, our main Theorem 5.1 states this sequence converges in distribution to a random map (M_∞, E_∞), which we call the *infinite* $\mathbf{q}$-*Boltzmann map*. Due to combinatorial reasons, we have to restrict n to a lattice of the form $2 + d\mathbb{Z}_+$ where d is an integer depending on the sequence $\mathbf{q}$.

The class of weight sequences for which this is true is the class of *critical* (as defined in Section 4.2) sequences, which contains all sequences with finite support (up to multiplicative constants). Taking $q_n = \mathbb{1}_{\{n=p\}}$ with $p \geqslant 3$ gives us the case of the uniform p-angulation, making our results an extension of what was known about the UIPT and UIPQ.

Local limits of Boltzmann random maps have notably been studied recently in [18]. A central difference with our work here is the fact that the maps are supposed to be bipartite in [18] (the weight sequence $\mathbf{q}$ is supported on the even integers). In this context, it is more natural to condition maps by their number of edges instead of their number of vertices, leading to a result which complements ours.

The proof of convergence to an infinite map hinges on a similar result for critical multi-type Galton-Watson trees and forests, Theorem 3.3. This theorem itself generalizes the well-known fact that critical monotype Galton-Watson trees, when conditioned to be large, converge to an infinite tree formed by a unique infinite spine to which many finite trees are grafted. This infinite tree was first indirectly mentioned in [51], Lemma 1.14, and many details about the convergence are given in [1] and [49]. One of its properties is that one can obtain its distribution from the distribution of the finite tree by a size-biasing process, as is explained in [60].

As will be apparent when we discuss the Bouttier-Di Francesco-Guitter bijection in Section 4.3, we will want to condition multi-type Galton-Watson trees on the number of vertices of one fixed type only, leaving the other types free. As it happens, such a conditioning is for now the only one under which we are able to prove the local convergence of the tree to an infinite version. The distribution of the infinite tree can once again be described by a size-biasing process from the original tree, as explained in Proposition 3.1, something which was anticipated in [54].

A fairly important issue in Theorem 3.3 is the problem of periodicity: as with maps, a multi-type Galton-Watson tree cannot have any number of vertices. To be precise, the number of vertices of a fixed type in the tree is always in $\alpha + d\mathbb{Z}_+$, where d only depends on the offspring distribution and α also depends on the type of the root vertex. Particular care must thus be taken when counting the vertices of forests or specific subtrees.

The chapter is split into two halves: we start by working on trees, and later on apply the results to maps. To be precise, after recalling facts about multi-type Galton-Watson trees in Section 2, we prove in Section 3 the convergence of large critical multi-type Galton-Watson forests to their infinite counterpart. Section 4 then states the basic background on planar maps, and we state and prove Theorem 5.1, our main theorem of convergence of maps, in Section 5. The final section is then dedicated to an application, namely showing that the infinite Boltzmann map is almost surely a recurrent graph.

2 Background on multi-type Galton-Watson trees

2.1 Basic definitions

Multi-type plane trees. We recall the standard formalism for family trees, first introduced by Neveu in [68]. Let

$$\mathcal{U} = \bigcup_{k=0}^{\infty} \mathbb{N}^k$$

be the set of finite words on $\mathbb{N}$, also known as the Ulam-Harris tree. Elements of $\mathcal{U}$ are written as sequences $u = u^1u^2\dots u^k$, and we call $|u| = k$ the height of u. We also let $u^- = u^1u^2\dots u^{k-1}$ be the father of u when $k > 0$. In the case of the empty word $\emptyset$, we let $|\emptyset| = 0$ and we do not give it a father. If $u = u^1\dots u^k$ and $v = v^1\dots v^l$ are two words, we define their concatenation $uv = u^1\dots u^k v^1\dots v^l$.

A plane tree is a subset $\mathbf{t}$ of $\mathcal{U}$ which satisfies the following conditions:

- $\emptyset \in \mathbf{t}$,
- $u \in \mathbf{t} \setminus \{\emptyset\} \Rightarrow u^- \in \mathbf{t}$,
- $\forall u \in \mathbf{t}, \exists k_u(\mathbf{t}) \in \mathbb{Z}_+, \forall j \in \mathbb{N}, uj \in \mathbf{t} \Leftrightarrow j \leqslant k_u(\mathbf{t})$.

Given a tree $\mathbf{t}$ and an integer $n \in \mathbb{Z}_+$, we let $\mathbf{t}_n = \{u \in \mathbf{t}, |u| = n\}$ and $\mathbf{t}_{\leqslant n} = \{u \in \mathbf{t}, |u| \leqslant n\}$. We call *height* of $\mathbf{t}$ the supremum $ht(t)$ of the heights of all its elements. If $u \in \mathbf{t}$, we let $\mathbf{t}_u = \{v \in \mathcal{U}, uv \in \mathbf{t}\}$ be the subtree of $\mathbf{t}$ rooted at u.

Note that the finiteness of $k_u(\mathbf{t})$ for any vertex u implies that all the trees which we consider are locally finite: a vertex can only have a finite number of neighbours. We do however allow infinite trees.

Let now $K \in \mathbb{N}$ be an integer. A K-type tree is a pair $(\mathbf{t}, \mathbf{e})$ where $\mathbf{t}$ is a plane tree and $\mathbf{e}$ is a function: $\mathbf{t} \to [K]$, which gives a type $\mathbf{e}(u)$ to every vertex $u \in \mathbf{t}$. For a vertex $u \in \mathbf{t}$, we also let $\mathbf{w}_{\mathbf{t}}(u) = \big(\mathbf{e}(u1), \dots, \mathbf{e}(uk_u(\mathbf{t}))\big)$ be the list of types of the ordered offspring of u. Note of course that the knowledge of $\mathbf{e}(\emptyset)$ and of all the $\mathbf{w}_{\mathbf{t}}(u)$, $u \in \mathbf{t}$ gives us the complete type function $\mathbf{e}$.

We let

$$\mathcal{W}_K = \bigcup_{n=0}^{\infty} [K]^n.$$

be the set of finite type-lists. Given such a list $\mathbf{w} \in \mathcal{W}_K$ and a type $i \in [K]$, we let $p_i(\mathbf{w}) = \#\{j, w_j = i\}$ and $p(\mathbf{w}) = (p_i(\mathbf{w}))_{i \in [K]}$. This defines a natural projection from $\mathcal{W}_K$ onto $(\mathbb{Z}_+)^K$. We also let $|\mathbf{w}| = \sum_i p_i(\mathbf{w})$ be the length of $\mathbf{w}$. Elements of $\mathcal{W}_K$ should be seen as orderings of types, such that the type i appears $p_i(\mathbf{w})$ times in the order $\mathbf{w}$.

Offspring distributions. We call *ordered offspring distribution* any sequence $\zeta = (\zeta^{(i)})_{i \in [K]}$ where, for all $i \in [K]$, $\zeta^{(i)}$ is a probability distribution on $\mathcal{W}_K$. Letting $\mu^{(i)} = (p_i)_* \zeta^{(i)}$ for all i, we then call $\mu = (\mu^{(i)})_{i \in [K]}$ the associated *unordered* offspring distribution.

We will always assume the condition

$$\exists i \in [K], \mu^{(i)} \left(\left\{ \mathbf{z} \in (\mathbb{Z}_+)^k, \sum_{j=1}^{K} z_j \neq 1 \right\} \right) > 0$$

to avoid degenerate cases which will lead to infinite linear trees.

Uniform orderings. Let us give details about a particular case of ordered offspring distribution. For $\mathbf{n} = (n_i)_{i \in [K]} \in (\mathbb{Z}_+)^K$, we call uniform ordering of $\mathbf{n}$ any uniformly distributed random variable on the set of words $\mathbf{w} \in \mathcal{W}_K$ satisfying $p(\mathbf{w}) = \mathbf{n}$. Such a random variable can be obtained by taking the word $(1, 1, \dots, 1, 2, \dots, 2, 3, \dots, K, \dots, K)$ (where each i is repeated n_i times) and applying a uniform permutation to it. Now let $\mu = (\mu^{(i)})_{i \in [K]}$ be a family of distributions on $(\mathbb{Z}_+)^K$, we call *uniform ordering of* μ the ordered offspring distribution $\zeta = (\zeta^{(i)})_{i \in [K]}$ where, for each i, $\zeta^{(i)}$ is the distribution of a uniform ordering of a random variable with distribution $\mu^{(i)}$.

Galton-Watson distributions. We can now define the distribution of a K-type Galton-Watson tree rooted at a vertex of type $i \in [K]$ and with ordered offspring distribution ζ, which we call $\mathbb{P}_\zeta^{(i)}$, by

$$\mathbb{P}_\zeta^{(i)}(\mathbf{t}, \mathbf{e}) = \mathbb{1}_{\{\mathbf{e}(\emptyset) = i\}} \prod_{u \in \mathbf{t}} \zeta^{(\mathbf{e}(u))}(\mathbf{w}_{\mathbf{t}}(u)) \tag{4.1}$$

for any finite tree $(\mathbf{t}, \mathbf{e})$. This formula only defines a sub-probability measure in general, however in the cases which interest us (namely, critical offspring distributions, see the next section) we will indeed have a probability distribution. In practice we are not interested in this formula as much as in the *branching property*, which also characterizes these distributions: the types of the children of the root of a tree $(\mathbf{T}, \mathbf{E})$ with law $\mathbb{P}_\zeta^{(i)}$ are determined by a random variable with law $\zeta^{(i)}$ and, conditionally on the offspring of the root being equal to a word $\mathbf{w}$, the subtrees rooted at points j with $j \in [|\mathbf{w}|]$ are independent, each one with distribution $\mathbb{P}_\zeta^{(j)}$.

Criticality. Let $M = (m_{i,j})_{i,j \in [K]}$ be the $K \times K$ matrix defined by

$$m_{i,j} = \sum_{\mathbf{z} \in (\mathbb{Z}_+)^K} z_j \mu^{(i)}(\mathbf{z}), \qquad \forall i, j \in [K].$$

We assume that M is *irreducible*, which means that, for all i and j in $[K]$, there exists some power p such that the (i,j)-th entry of M^p is nonzero. In this case, we know by the Perron-Frobenius theorem that the spectral radius ρ of M is in fact an eigenvalue of M. We say that ζ (or μ, or M) is *subcritical* if $\rho < 1$ and *critical* if $\rho = 1$, which both in particular impliy that equation 4.1 does define a probability distribution and that Galton-Watson trees with ordered offspring distribution ζ are almost surely finite. We will always assume criticality in the rest of the chapter. The Perron-Frobenius theorem also tells us that, up to multiplicative constants, the left and right eigenvectors of M for ρ are unique. We call them $\mathbf{a} = (a_1, \dots, a_K)$ and $\mathbf{b} = (b_1, \dots, b_K)$ and normalize them such that $\sum_i a_i = \sum_i a_i b_i = 1$, in which case their components are all strictly positive.

The fact that $\mathbf{b}$ is a right-eigenvector of M translates as

$$b_i = \sum_{\mathbf{z} \in (\mathbb{Z}_+)^K} \mu^{(i)}(\mathbf{z}) \mathbf{b} \cdot \mathbf{z},$$

where $\cdot$ is the usual dot product. One can deduce from this the existence of a martingale naturally associated to the Galton-Watson tree. Let $(\mathbf{T}, \mathbf{E})$ have the distribution $\mathbb{P}_\zeta^{(i)}$ for some $i \in [K]$ and, for all $n \in \mathbb{N}$ and $j \in [K]$, let $Z_n^{(j)}$ be the number of vertices of $\mathbf{T}$ which have height n and type j, and set $\mathbf{Z}_n = (Z_n^{(j)})_{j \in [K]}$. Define then, for $n \in \mathbb{N}$,

$$X_n = \mathbf{b} \cdot \mathbf{Z}_n = \sum_{j=1}^{K} b_j Z_n^{(j)}. \tag{4.2}$$

The process $(X_n)_{n \in \mathbb{Z}_+}$ is then a martingale.

Spatial trees. Later on in this paper we will be looking at *spatial* K-type trees, that is trees coupled with labels on their vertices. We define a K-type spatial tree to be a triple $(\mathbf{t}, \mathbf{e}, \mathbf{l})$ where $(\mathbf{t}, \mathbf{e})$ is a K-type tree and $\mathbf{l}$ is any real-valued function on $\mathbf{t}$. Note that, given $\mathbf{t}$, $\mathbf{e}$ and $\mathbf{l}(\emptyset)$, the rest of $\mathbf{l}$ is completely determined by the differences $\mathbf{l}(u) - \mathbf{l}(u^-)$ for $u \in \mathbf{t} \setminus \{\emptyset\}$. This is why we let, for $u \in \mathbf{t}$, $\mathbf{y}_u = \Big(\mathbf{l}(u_1) - \mathbf{l}(u), \mathbf{l}(u2) - \mathbf{l}(u), \dots, \mathbf{l}\big(uk_u(\mathbf{t})\big) - \mathbf{l}(u)\Big) \in \mathbb{R}^{|\mathbf{w}_\mathbf{t}(u)|}$ be the list of ordered label displacements of the offspring of u.

Consider, for all types $i \in [K]$ and words $\mathbf{w} \in \mathcal{W}_K$, a probability distribution $\nu_\mathbf{w}^{(i)}$ on $\mathbb{R}^{|\mathbf{w}|}$, as well as a number ε. We let $\mathbb{P}_{\zeta,\nu}^{(i,\varepsilon)}$ be the distribution of a triple $(\mathbf{T}, \mathbf{E}, \mathbf{L})$ where $(\mathbf{T}, \mathbf{E})$ is a K-type tree with distribution $\mathbb{P}_\zeta^{(i)}$, the root $\emptyset$ has label ε and the label displacements $\big(\mathbf{L}(u_1) - \mathbf{L}(u),$ $\mathbf{L}(u2) - \mathbf{L}(u), \dots, \mathbf{L}(uk_u(\mathbf{T})) - \mathbf{L}(u)\big)$ (with $u \in \mathbf{T}$) are all independent, each one having distribution $\nu_{\mathbf{w}_\mathbf{T}(u)}^{(\mathbf{E}(u))}$ conditionally on $\mathbf{E}(u)$ and $\mathbf{w}_\mathbf{T}(u)$.

Forests. We will not only look at trees but also at multi-type (and, when needed, labelled) *forests*, a forest being defined as a ordered finite collection of trees: elements of the form $(\mathbf{f}, \mathbf{e}, \mathbf{l}) = \big((\mathbf{t}^1, \mathbf{e}^1, \mathbf{l}^1), \dots, (\mathbf{t}^p, \mathbf{e}^p, \mathbf{l}^p)\big)$.

A Galton-Watson random forest will be a forest where the trees are mutually independent and each one has a Galton-Watson distribution with the same ordered offspring distribution (and label increment distribution, in the labelled case). We can thus let, for $\mathbf{w} \in \mathcal{W}_K$, $\mathbb{P}_\zeta^{(\mathbf{w})}$ be the distribution of $(\mathbf{T}^i, \mathbf{E}^i)_{i \in [|\mathbf{w}|]}$ where the $(\mathbf{T}^i, \mathbf{E}^i)$ are independent, and each $(\mathbf{T}^i, \mathbf{E}^i)$ has distribution $\mathbb{P}_\zeta^{(w_i)}$ and, given also a list of initial labels $\varepsilon = (\varepsilon_1, \dots, \varepsilon_{|\mathbf{w}|})$, $\mathbb{P}_{\zeta,\nu}^{(\mathbf{w}),(\epsilon)}$ be the distribution of $(\mathbf{T}^i, \mathbf{E}^i, \mathbf{L}^i)_{i \in [|\mathbf{w}|]}$ where the terms of the sequence are independent and, for a given i, $(\mathbf{T}^i, \mathbf{E}^i, \mathbf{L}^i)$ has distribution $\mathbb{P}_{\zeta,\nu}^{(w_i,\varepsilon_i)}$.

All previous notation will be adapted to forests, for example, the height of a forest $\mathbf{f}$ is the maximum of the heights of its elements, $\mathbf{f}_{\leqslant n}$ is the forest where each tree has been cut at height n, and so on.

Canonical variable notation. For readability, we will throughout the chapter use the canonical variable $(\mathbf{T}, \mathbf{E})$, which is simply the identity function of the space of K-type trees, as well as $(\mathbf{T}, \mathbf{E}, \mathbf{L})$, $(\mathbf{F}, \mathbf{E})$ $(\mathbf{F}, \mathbf{E}, \mathbf{L})$ when looking at labelled trees or forests. Thus we will, for instance, write $\mathbb{P}_\zeta^{(i)}\big((\mathbf{T}, \mathbf{E}) = (\mathbf{t}, \mathbf{e})\big)$ instead of $\mathbb{P}_\zeta^{(i)}(\mathbf{t}, \mathbf{e})$, for a given type i and a given K-type tree $(\mathbf{t}, \mathbf{e})$.

Local convergence of multi-type trees and forests. Take a sequence of K-type forests $(\mathbf{f}^{(n)}, \mathbf{e}^{(n)})_{n\in\mathbb{N}}$. We say that this sequence converges locally to a K-type forest $(\mathbf{f}, \mathbf{e})$ if, for all $k \in \mathbb{N}$, and $n \in \mathbb{N}$ large enough (depending on k), we have $(\mathbf{f}^{(n)}_{\leqslant k}, \mathbf{e}^{(n)}_{\leqslant k}) = (\mathbf{f}_{\leqslant k}, \mathbf{e}_{\leqslant k})$. This convergence can be metrized: we can for example set, for two K-type forests $(\mathbf{f}, \mathbf{e})$ and $(\mathbf{f}', \mathbf{e}')$, $d\big((\mathbf{f}, \mathbf{e}), (\mathbf{f}', \mathbf{e}')\big) = \frac{1}{1+p}$ where p is the supremum of all integers k such that $(\mathbf{f}_{\leqslant k}, \mathbf{e}_{\leqslant k}) = (\mathbf{f}'_{\leqslant k}, \mathbf{e}'_{\leqslant k})$.

Convergence in distribution of random forests for this metric is simply characterized: if $(\mathbf{F}^{(n)}, \mathbf{E}^{(n)})_{n\in\mathbb{N}}$ is a sequence of random K-type forests, it converges in distribution to a certain random forest $(\mathbf{F}, \mathbf{E})$ if and only if, for all $k \in \mathbb{N}$ and finite K-type forests $(\mathbf{f}, \mathbf{e})$, the quantity $\mathbb{P}\big((\mathbf{F}^{(n)}_{\leqslant k}, \mathbf{E}^{(n)}_{\leqslant k}) = (\mathbf{f}, \mathbf{e})\big)$ converges to $\mathbb{P}\big((\mathbf{F}_{\leqslant k}, \mathbf{E}_{\leqslant k}) = (\mathbf{f}, \mathbf{e})\big)$.

All these definitions can directly be adapted to the case of spatial forests: when asking for equality between the forests below height k, we also ask equality of the labels below this height.

2.2 The first generation of fixed type

In this section, the next and also later on in the chapter, we fix a reference type $j \in [K]$, and are interested in counting the number of vertices of type j in a Galton-Watson tree with ordered offspring distribution ζ. A useful tool for this is the *first generation of type* j, that is, in a K-type tree $(\mathbf{t}, \mathbf{e})$, the set of vertices of $\mathbf{t}$ with type j which have no ancestors of type j, except maybe for the root. If $(\mathbf{T}, \mathbf{E})$ has distribution $\mathbb{P}_\nu^{(i)}$ for some type i, we let $\mu_{i,j}$ be the distribution of the number of vertices in the first generation of type j of $(\mathbf{T}, \mathbf{E})$.

Lemma 2.1. *For all types i, the average of the probability distribution $\mu_{i,j}$ is equal to*

$$\sum_{k=0}^{\infty} k\mu_{i,j}(k) = \frac{b_i}{b_j}$$

Proof. For all $i \in [K]$, let $c_i = \sum_{k=0}^{\infty} k\mu_{i,j}(k)$. The proof that $c_i = \frac{b_i}{b_j}$ for all i is done in two steps: first, show that $c_j = 1$ and then that the vector $\mathbf{c} = (c_i)_{i\in[K]}$ is a right eigenvector of M for the eigenvalue 1.

The fact that $c_j = 1$ is proven in [65], Proposition 4. It is obtained by removing the types different from j one by one, and noticing that criticality is conserved at every step until we are left with a critical monotype Galton-Watson tree.

To prove that $\mathbf{c}$ is a right eigenvector of M, consider a type $i \in [K]$ and apply the branching property at height 1 in a tree with distribution $\mathbb{P}_\nu^{(i)}$, we get

$$\begin{aligned} c_i &= \sum_{\mathbf{z}\in(\mathbb{Z}_+)^K} \mu^{(i)}(\mathbf{z})\Big(\sum_{l\in[K]\setminus\{j\}} z_l c_l + z_j\Big) \\ &= \sum_{\mathbf{z}\in(\mathbb{Z}_+)^K} \mu^{(i)}(\mathbf{z})\Big(\sum_{l=1}^{K} z_l c_l\Big) \\ &= \sum_{l=1}^{K} m_{i,l} c_l. \end{aligned}$$

Since $\sum_{l=1}^k m_{i,l}c_l$ is the i-th component of $(M\mathbf{c})$, the proof is complete. □

2.3 Periodicity

We keep a fixed type $j \in [K]$, and want to find out for which n it is possible for a tree with distribution $\mathbb{P}_\zeta^{(i)}$ (for all types $i \in [K]$) to have n vertices of type j. We introduce the notation $\#_j\mathbf{t}$ for the number of vertices of type j in any K-type tree $\mathbf{t}$. Let $d = \gcd\{n, \mu_{j,j}(n) > 0\}$. It is straightforward that, if $\mathbf{T}$ has distribution $\mathbb{P}_\zeta^{(j)}$, $\#_j\mathbf{T}$ is always of the form $1 + dn$ with $n \in \mathbb{Z}_+$. A similar statement is also true if the root has a different type:

Lemma 2.2. *(i) For all $i \in [k]$, there exists an integer $\beta_i \in \{0, 1, \dots, d-1\}$ such that the measure $\mu_{i,j}$ is supported on $\beta_i + d\mathbb{Z}_+$.*

(ii) Setting $\alpha_i = \beta_i + \mathbb{1}_{\{i=j\}}$ for all $i \in [K]$, the number of vertices of type j in a tree with distribution $\mathbb{P}_\zeta^{(i)}$ is then always of the form $\alpha_i + dn$ for some integer n.

(iii) For $i \in [K]$ and $\mathbf{w} \in \mathcal{W}_K$ such that $\zeta^{(i)}(\mathbf{w}) > 0$, we have

$$\beta_i \equiv \sum_{k=1}^{|\mathbf{w}|} \alpha_{w_k} \pmod d.$$

Proof. The main point is to prove the existence of β_i. We therefore choose two integers m and n such that $\mu_{i,j}(n) > 0$ and $\mu_{i,j}(m) > 0$, and set out to prove $n \equiv m \pmod d$. Note that, under $\mathbb{P}_\zeta^{(j)}$, the probability that there is a vertex of type i before the first generation of type j is nonzero, because irreducibility could not be satisfied otherwise. Since such a vertex can have n or m children in the first generation of type j, we obtain the existence of an integer p such that $\mu_{j,j}(p) > 0$ and $\mu_{j,j}(p + n - m) > 0$, which implies that d divides $n - m$.

Points (*ii*) and (*iii*) are both immediate consequence of the fact that the number of descendants of type j of a vertex of type j is always a multiple of d. □

Remark 2.3. *We will later on indirectly prove that, for any type $i \in [K]$, if n is large enough, then $\mathbb{P}^{(i)}(\#_j\mathbf{T} = \alpha_i + dn) > 0$.*

3 Infinite multi-type Galton-Watson trees and forests

In this section we will consider unlabelled trees and forests with a critical ordered offspring distribution ζ, and will omit mentioning ζ for readability purposes. We could in fact work with spatial trees, however, since the labellings are done conditionally on the tree and in independent fashion for each vertex, the reader can check that the proofs do not change at all if we add the labellings in.

Just as in the case of critical monotype Galton-Watson trees, multi-type trees have an infinite variant which is obtained through a size-biasing method which was first introduced in [54].

3.1 Existence of the infinite forest

Proposition 3.1. *There exists a unique probability measure $\widehat{\mathbb{P}}^{(\mathbf{w})}$ on the space of infinite K-type forests such that, for any $n \in \mathbb{N}$ and for any finite K-type forest $(\mathbf{f}, \mathbf{e})$ with height n,*

$$\widehat{\mathbb{P}}^{(\mathbf{w})}\big((\mathbf{F}_{\leqslant n}, \mathbf{E}_{\leqslant n}) = (\mathbf{f}, \mathbf{e})\big) = \frac{1}{Z_{\mathbf{w}}}\left(\sum_{u \in \mathbf{f}_n} b_{\mathbf{e}(u)}\right)\mathbb{P}^{(\mathbf{w})}\big((\mathbf{F}_{\leqslant n}, \mathbf{E}_{\leqslant n}) = (\mathbf{f}, \mathbf{e})\big). \tag{4.3}$$

The normalizing constant $Z_{\mathbf{w}}$ is equal to $\sum_{i=1}^{|\mathbf{w}|} b_{w_i} = p(\mathbf{w}) \cdot \mathbf{b}$.

Proof. Our proof is structured as the one given in [60] for monotype trees. Let $n \in \mathbb{N}$, we will first define a probability distribution $\widehat{\mathbb{P}}_n^{(\mathbf{w})}$ on the space of K-type forests with height exactly n paired with a point of height n. Let $(\mathbf{f}, \mathbf{e})$ be such a forest and $u \in \mathbf{f}_n$, and set

$$\widehat{\mathbb{P}}_n^{(\mathbf{w})}\big(\mathbf{f}, \mathbf{e}, u\big) = \frac{b_{\mathbf{e}(u)}}{Z_{\mathbf{w}}}\mathbb{P}^{(\mathbf{w})}\big((\mathbf{F}, \mathbf{E}) = (\mathbf{f}, \mathbf{e})\big).$$

The martingale property of the process $(X_n)_{n\in\mathbb{N}}$ defined by (4.2) under $\mathbb{P}^{(\mathbf{w})}$ ensures us that we do have probability measures: the total mass of $\widehat{\mathbb{P}}_n^{(\mathbf{w})}$ is $\frac{1}{Z_{\mathbf{w}}}\mathbb{E}_{\zeta}^{(\mathbf{w})}(X_n) = \frac{p(\mathbf{w})\cdot\mathbf{b}}{Z_{\mathbf{w}}} = 1$.

We will check that these are compatible in the sense that, for $n \in \mathbb{N}$, if $(\mathbf{F}, \mathbf{E}, U)$ has distribution $\widehat{\mathbb{P}}_{n+1}^{(\mathbf{w})}$ then $(\mathbf{F}_{\leqslant n}, \mathbf{E}_{\leqslant n}, U^-)$ has distribution $\widehat{\mathbb{P}}_n^{(\mathbf{w})}$. Fix therefore $(\mathbf{f}, \mathbf{e})$ a K-type forest of height n and u a vertex of t at height n. We have

$$\begin{aligned}
\widehat{\mathbb{P}}_{n+1}^{(\mathbf{w})}\big((\mathbf{F}_{\leqslant n}, \mathbf{E}_{\leqslant n}, U^-) = (\mathbf{f}, \mathbf{e}, u)\big) &= \frac{1}{Z_{\mathbf{w}}}\mathbb{P}^{(\mathbf{w})}\big((\mathbf{F}_{\leqslant n}, \mathbf{E}_{\leqslant n}) = (\mathbf{f}, \mathbf{e})\big) \sum_{\mathbf{x} \in \mathcal{W}_K} \zeta^{(\mathbf{e}(u))}(\mathbf{x}) \sum_{j=1}^{|\mathbf{x}|} b_{x_j} \\
&= \frac{1}{Z_{\mathbf{w}}}\mathbb{P}^{(\mathbf{w})}\big((\mathbf{F}_{\leqslant n}, \mathbf{E}_{\leqslant n}) = (\mathbf{f}, \mathbf{e})\big) \sum_{\mathbf{z} \in (\mathbb{Z}_+)^K} \mu^{(\mathbf{e}(u))}(\mathbf{z})\, \mathbf{z} \cdot \mathbf{b} \\
&= \frac{1}{Z_{\mathbf{w}}}\mathbb{P}^{(\mathbf{w})}\big((\mathbf{F}_{\leqslant n}, \mathbf{E}_{\leqslant n}) = (\mathbf{f}, \mathbf{e})\big)\, b_{\mathbf{e}(u)}.
\end{aligned}$$

Kolmogorov's consistency theorem then allows us to define a distribution $\widehat{\mathbb{P}}_{\infty}^{(\mathbf{w})}$ on the set of forests where one of the trees has a distinguished infinite path. Forgetting the infinite path then gives us the distribution $\widehat{\mathbb{P}}^{(\mathbf{w})}$ which we were looking for.

□

For $n \in \mathbb{Z}_+$, $(\mathbf{f}, \mathbf{e})$ a forest of height $n+1$ and $u \in \mathbf{f}_{n+1}$, we have

$$\widehat{\mathbb{P}}_{n+1}^{(\mathbf{w})}\Big((\mathbf{F}, \mathbf{E}, U) = (\mathbf{f}, \mathbf{e}, u) \mid (\mathbf{F}_{\leqslant n}, \mathbf{E}_{\leqslant n}, U^-) = (\mathbf{f}_{\leqslant n}, \mathbf{e}_{\leqslant n}, u^-)\Big) =$$
$$\frac{b_{\mathbf{e}(\mathbf{u})}}{b_{\mathbf{e}(\mathbf{u}^-)}}\mathbb{P}_{n+1}^{(\mathbf{w})}\Big((\mathbf{F}, \mathbf{E}) = (\mathbf{f}, \mathbf{e}) \mid (\mathbf{F}_{\leqslant n}, \mathbf{E}_{\leqslant n}) = (\mathbf{f}_{\leqslant n}, \mathbf{e}_{\leqslant n})\Big).$$

From this formula follows a simple description of these infinite forests.

Given a type $i \in [K]$, a random tree with distribution $\widehat{\mathbb{P}}^{(i)}$ can be described the following way: it is made of a *spine*, that is an infinite ascending chain starting at the root, on which we have grafted independent trees with offspring distribution ζ. Elements of the spine have a different offspring distribution, called $\widehat{\zeta}$, which is a size-biased version of ζ. It is defined by

$$\widehat{\zeta}^{(j)}(\mathbf{x}) = \frac{1}{b_j} \sum_{l=1}^{|\mathbf{x}|} b_{x_l} \zeta^{(j)}(\mathbf{x}), \tag{4.4}$$

with $j \in [K]$ and $\mathbf{x} \in \mathcal{W}_K$. Given an element of the spine $u \in \mathcal{U}$ and its offspring $\mathbf{x} \in \mathcal{W}_K$, the probability that the next element of the spine is uj for $j \in [|\mathbf{x}|]$ is proportional to b_{x_j}, and therefore equal to

$$\frac{b_{x_j}}{\sum_{l=1}^{|\mathbf{x}|} b_{x_l}}.$$

To get a forest with distribution $\widehat{\mathbb{P}}^{(\mathbf{w})}$, let first J be a random variable taking values in $[|\mathbf{w}|]$ such that $J = j$ with probability proportional to b_{w_j}. Conditionally on J, let $(\mathcal{T}_J, \mathbf{E}_J)$ be a tree with distribution $\widehat{\mathbb{P}}^{(J)}$, and let $(\mathbf{T}_i, \mathbf{E}_i)$, for $i \in [|\mathbf{w}|]$, $i \neq J$ be a tree with distribution $\mathbb{P}^{(i)}$, all these trees being mutually independent. Then the forest $(\mathbf{T}_i, \mathbf{E}_i)_{i \in [|\mathbf{w}|]}$ has distribution $\widehat{\mathbb{P}}^{(\mathbf{w})}$.

Remark 3.2. *Recall that a tree with law $\mathbb{P}^{(i)}$ is finite for any $i \in [K]$. Therefore, a forest with distribution $\widehat{\mathbb{P}}^{(\mathbf{w})}$ can only have one infinite path, and thus we do not lose any information by going from $\widehat{\mathbb{P}}^{(\mathbf{w})}_\infty$ to $\widehat{\mathbb{P}}^{(\mathbf{w})}$.*

3.2 Convergence to the infinite forest

Let $j \in [K]$ be any type and $\mathbf{w} \in \mathcal{W}_K$ be any word. Recall from Section 2.3 that d is the gcd of the support of the measure $\mu_{j,j}$, and let also $\alpha_\mathbf{w} = \sum_i \alpha_{w_i}$, such that the number of vertices of type j in a forest with distribution $\mathbb{P}^{(\mathbf{w})}$ is of the form $\alpha_\mathbf{w} + dn$, with integer n.

Theorem 3.3. *As n tends to infinity, a forest with distribution $\mathbb{P}^{(\mathbf{w})}$, conditioned on having $\alpha_\mathbf{w} + dn$ vertices of type j, converges in distribution to a forest with distribution $\widehat{\mathbb{P}}^{(\mathbf{w})}$. In other words, given a forest $(\mathbf{f}, \mathbf{e})$ of height k, we have*

$$\mathbb{P}^{(\mathbf{w})}\big((\mathbf{F}_{\leqslant k}, \mathbf{E}_{\leqslant k}) = (\mathbf{f}, \mathbf{e}) \mid \#_j \mathbf{F} = \alpha_\mathbf{w} + dn\big) \underset{n\to\infty}{\longrightarrow} \widehat{\mathbb{P}}^{(\mathbf{w})}\big((\mathbf{F}_{\leqslant k}, \mathbf{E}_{\leqslant k}) = (\mathbf{f}, \mathbf{e})\big)$$

The proof of this theorem will rely on the following asymptotics, indexed by any word $\mathbf{w} \in \mathcal{W}_K$

$$\mathbb{P}^{(\mathbf{w})}(\#_j \mathbf{F} = \alpha_\mathbf{w} + dn) \underset{n\to\infty}{\sim} \frac{Z_\mathbf{w}}{b_j} \mathbb{P}^{(j)}\big(\#_j \mathbf{T} = 1 + d(n+p)\big), \qquad \forall p \in \mathbb{Z} \tag{$H_\mathbf{w}$}$$

We will prove that $(H_\mathbf{w})$ holds for any word $\mathbf{w}$ and that both terms are nonzero for large enough n (implying that the conditioning of the theorem is well-defined). Let us first prove that the theorem is indeed a consequence of this. Take a forest $(\mathbf{f}, \mathbf{e})$ with height k, and let $\mathbf{x}$ be the word obtained by taking the types of the vertices of $\mathbf{f}$ with height k (the order of the elements $\mathbf{x}$ actually has no influence). For n large enough, we have

$$\begin{aligned}\mathbb{P}^{(\mathbf{w})}\big((\mathbf{F}_{\leqslant k}, \mathbf{E}_{\leqslant k}) = (\mathbf{f}, \mathbf{e}) \mid \#_j \mathbf{F} = \alpha_\mathbf{w} + dn\big) &= \frac{\mathbb{P}^{(\mathbf{w})}\big((\mathbf{F}_{\leqslant k}, \mathbf{E}_{\leqslant k}) = (\mathbf{f}, \mathbf{e}), \#_j \mathbf{F} = \alpha_\mathbf{w} + dn\big)}{\mathbb{P}^{(\mathbf{w})}(\#_j \mathbf{F} = \alpha_\mathbf{w} + dn)} \\ &= \mathbb{P}^{(\mathbf{w})}\big((\mathbf{F}_{\leqslant k}, \mathbf{E}_{\leqslant k}) = (\mathbf{f}, \mathbf{e})\big) \frac{\mathbb{P}^{(\mathbf{x})}(\#_j \mathbf{F} = \alpha_\mathbf{w} + dn - q)}{\mathbb{P}^{(\mathbf{w})}(\#_j \mathbf{F} = \alpha_\mathbf{w} + dn)}\end{aligned}$$

where q is the number of vertices of $\mathbf{f}$ of type j which do not have height n. Repeatedly using point (iii) of Lemma 2.2 shows that, if $\mathbb{P}^{(\mathbf{w})}\big((\mathbf{F}_{\leqslant k}), (\mathbf{E}_{\leqslant k}) = (\mathbf{f}, \mathbf{e})\big) > 0$ then $\alpha_{\mathbf{w}} - q$ must be congruent to $\alpha_{\mathbf{x}}$ modulo d, giving us

$$\begin{aligned}\mathbb{P}^{(\mathbf{w})}&\big((\mathbf{F}_{\leqslant k}, \mathbf{E}_{\leqslant k}) = (\mathbf{f},\mathbf{e}) \mid \#_j\mathbf{F} = \alpha_{\mathbf{w}} + dn\big)\\ &= \mathbb{P}^{(\mathbf{w})}\big((\mathbf{F}_{\leqslant k}), (\mathbf{E}_{\leqslant k}) = (\mathbf{f}, \mathbf{e})\big)\frac{\mathbb{P}^{(\mathbf{x})}\big(\#_j\mathbf{F} = \alpha_{\mathbf{x}} + d(n+p)\big)}{\mathbb{P}^{(\mathbf{w})}(\#_j\mathbf{F} = \alpha_{\mathbf{w}} + dn)}\end{aligned}$$

for a certain integer p. Now if we let n tend to infinity, using both $(H_{\mathbf{w}})$ and $(H_{\mathbf{x}})$, we obtain

$$\frac{\mathbb{P}^{(\mathbf{x})}\big(\#_j\mathbf{F} = \alpha_{\mathbf{x}} + d(n+p)\big)}{\mathbb{P}^{(\mathbf{w})}(\#_j\mathbf{F} = \alpha_{\mathbf{w}} + dn)} \underset{n\to\infty}{\longrightarrow} \frac{Z_{\mathbf{x}}}{Z_{\mathbf{w}}} = \frac{1}{Z_{\mathbf{w}}}\left(\sum_{u\in\mathbf{f}_k} b_{\mathbf{e}(u)}\right),$$

which concludes the proof of Theorem 3.3.

3.3 Proof of $(H_{\mathbf{w}})$

Obtaining $(H_{\mathbf{w}})$ for every word $\mathbf{w}$ will be done in several small steps. We will first prove it for some fairly simple words and gradually enlarge the class of $\mathbf{w}$ for which it holds, until we have every element of $\mathcal{W}_K$.

3.3.1 Ratio limit theorems for a random walk

Let $(S_n)_{n\in\mathbb{N}}$ be a random walk which starts at 0 and whose jumps are all greater than or equal to -1, their distribution being given by $\mathbb{P}(S_1 = k) = \mu_{j,j}(k+1)$ for $k \geqslant -1$.

Lemma 3.4. *For all $\alpha \in \{0, \dots, d-1\}$ we have*

$$\mathbb{P}(S_{\alpha+dn} = -\alpha) \underset{n\to\infty}{\sim} \mathbb{P}(S_{dn} = 0) \underset{n\to\infty}{\sim} \mathbb{P}(S_{d(n+1)} = 0)$$

Proof. The first thing to notice is that the random walk $(\frac{S_{dn}}{d})_{n\in\mathbb{N}}$ is irreducible, recurrent and aperiodic on $\mathbb{Z}$. First, it is indeed integer-valued because, by definition, for every n, $S_{n+1} \equiv S_n - 1 \pmod d$, and thus we stay in the same class modulo d if we take d steps at a time. Irreducibility comes from the fact that steps of $(S_n)_{n\in\mathbb{N}}$ has a nonzero probability of being equal to -1 because $\mu_{j,j}(0) > 0$, and thus $(\frac{S_{dn}}{d})_{n\in\mathbb{N}}$ can have positive jumps or jumps equal to -1. Since the jumps of $(S_n)_{n\in\mathbb{N}}$ are centered by Lemma 2.1, this makes $(\frac{S_{dn}}{d})_{n\in\mathbb{N}}$ an irreducible and centered random walk on $\mathbb{Z}$, so that it is recurrent (see for example Theorem 8.2 in [50]). Finally, aperiodicity is obtained from the fact that, if $\mu_{j,j}(n) > 0$, then $\mathbb{P}(S_n = 0) > 0$ by jumping straight to $n-1$ and going down to 0 one step at a time.

As a consequence of this, we can apply Spitzer's strong ratio theorem (see [76], p.49) to the random walk $(\frac{S_{dn}}{d})_{n\in\mathbb{N}}$. We obtain that, for any $k \in \mathbb{Z}$,

$$\mathbb{P}(S_{dn} = 0) \underset{n\to\infty}{\sim} \mathbb{P}(S_{d(n+1)} = 0) \underset{n\to\infty}{\sim} \mathbb{P}(S_{dn} = dk).$$

This proves the second half of Lemma 3.4, and can also be used to prove the first half. Let $\mu_{j,j}^{*\alpha}$ be the distribution of the sum of α independent variables with distribution $\mu_{j,j}$. For $n \in \mathbb{N}$, we then have

$$\mathbb{P}(S_{\alpha+dn} = -\alpha) = \sum_{p\in\mathbb{Z}} \mathbb{P}(S_{dn} = -\alpha - p)\mu_{j,j}^{*\alpha}(p+\alpha).$$

Fatou's lemma then gives us

$$\liminf_{n\to\infty} \frac{\mathbb{P}(S_{\alpha+dn}=-\alpha)}{\mathbb{P}(S_{dn}=0)} \geqslant \sum_{p\in\mathbb{Z}} \mu_{j,j}^{*\alpha}(p+\alpha) = 1$$

A similar argument also shows that

$$\liminf_{n\to\infty} \frac{\mathbb{P}(S_{d(n+1)}=0)}{\mathbb{P}(S_{\alpha+dn}=-\alpha)} \geqslant 1,$$

and this ends the proof. □

3.3.2 The case where $\mathbf{w} = (j, j \ldots, j)$

Consider a tree $(\mathbf{T}, \mathbf{E})$ with distribution $\mathbb{P}^{(j)}$. Consider then the reduced tree $\Pi^{(j)}(\mathbf{T})$ where all the vertices with types different from j have been erased but ancestral lines are kept (such that the father of a vertex of $\Pi^{(j)}(\mathbf{T})$ is its closest ancestor of type j in $\mathbf{T}$). This tree is precisely studied in [65], where it is shown that it is a monotype Galton-Watson tree, its offspring distribution naturally being $\mu_{j,j}$. As a result, the well-known *cyclic lemma* (see [69], Sections 6.1 and 6.2) tells us that

$$\mathbb{P}^{(j)}(\#_j\mathbf{T} = 1+dn) = \frac{1}{1+dn}\mathbb{P}(S_{1+dn}=-1).$$

where $(S_n)_{n\in\mathbb{N}}$ is the random walk defined in Section 3.3.1. This has two particular consequences. The first one is that $\mathbb{P}^{(j)}(\#_j\mathbf{T} = 1+dn)$ is nonzero for large n, as was first announced in Section 2.3, and the second is the fact that, thanks to Lemma 3.4, in order to prove $(H_{\mathbf{w}})$ for a certain word $\mathbf{w}$, we can restrict ourselves to proving the asymptotic equivalence for a single value of p, which will we take to be 0.

Consider now a word $\mathbf{w} = (j, j, \ldots, j)$ of length k, where k is any integer. The cyclic lemma can be adapted to forests (see [69] again), and we have

$$\mathbb{P}^{(\mathbf{w})}(\#_j\mathbf{F} = k+dn) = \frac{k}{k+dn}\mathbb{P}(S_{k+dn}=-k),$$

Lemma 3.4 then implies $(H_{\mathbf{w}})$ in this case since $Z_{\mathbf{w}} = kb_j$ and $\alpha_{\mathbf{w}} = k$.

The cases where $\mathbf{w}$ contains types different from j will be much less simple, and we first start with an inequality.

3.3.3 A lower bound

Let $\mathbf{w} \in \mathcal{W}_K$. In order to count the number of vertices of type j of a forest with distribution $\mathbb{P}^{(\mathbf{w})}$, we cut it at its first generation of type j.

$$\mathbb{P}^{(\mathbf{w})}(\#_j\mathbf{F} = \alpha_{\mathbf{w}}+dn) = \sum_{i=1}^{|\mathbf{w}|}\sum_{k_i=0}^{\infty} \mu_{w_i,j}(k_i)\mathbb{P}^{(j,\ldots,j)}(\#_j\mathbf{F} = \alpha_{\mathbf{w}}-q+dn)$$

where q is the number of times j appears in $\mathbf{w}$ and j is repeated $k_1 + k_2, \ldots + k_{|\mathbf{w}|}$ times in $\mathbb{P}^{(j,\ldots,j)}$. By Lemma 2.2, part (iii), whenever $\mu_{w_i,j}(k_i) > 0$, we have $\beta_{w_i} \equiv k_i \pmod d$, and thus the use of $H_{(j,\ldots,j)}$, combined with Fatou's lemma, gives us the following lower bound:

$$\liminf_{n\to\infty} \frac{\mathbb{P}^{(\mathbf{w})}(\#_j\mathbf{F} = \alpha_{\mathbf{w}}+dn)}{\mathbb{P}^{(j)}(\#_j\mathbf{T} = 1+dn)} \geqslant \sum_{i=1}^{|\mathbf{w}|}\sum_{k_i} k_i\mu_{w_{k_i},j}(k_i).$$

We can then use Lemma 2.1 to identify the right-hand side and obtain

$$\liminf_{n\to\infty} \frac{\mathbb{P}^{(\mathbf{w})}(\#_j\mathbf{F} = \alpha_{\mathbf{w}} + dn)}{\mathbb{P}^{(j)}(\#_j\mathbf{T} = 1 + dn)} \geqslant \frac{Z_{\mathbf{w}}}{b_j}. \tag{4.5}$$

To prove the reverse inequality for the limsup, we will try to fit a forest with distribution $\mathbb{P}^{(\mathbf{w})}$ "inside" a tree with distribution $\mathbb{P}^{(j)}$. We first need some additional notions.

3.3.4 The extension relation

Let $(\mathbf{t}, \mathbf{e})$ and $(\mathbf{t}', \mathbf{e}')$ be two K-type trees. We say that $\mathbf{t}'$ *extends* $\mathbf{t}$, which we write $\mathbf{t}' \vdash \mathbf{t}$ (omitting the type functions for clarity) if $\mathbf{t}'$ can be obtained from $\mathbf{t}$ by grafting trees on the leaves of $\mathbf{t}'$. More precisely, $\mathbf{t}' \vdash \mathbf{t}$ if:

- $\mathbf{t} \subset \mathbf{t}'$.
- $\forall u \in \mathbf{t}$, $\mathbf{e}(u) = \mathbf{e}'(u)$.
- $\forall u \in \mathbf{t}' \setminus \mathbf{t}, \exists v \in \partial\mathbf{t}, w \in \mathcal{U} : \; u = vw$.

Here, $\partial\mathbf{t}$ is the set of leaves of $\mathbf{t}$, that is the set of vertices v of $\mathbf{t}$ such that $k_v(\mathbf{t}) = 0$.

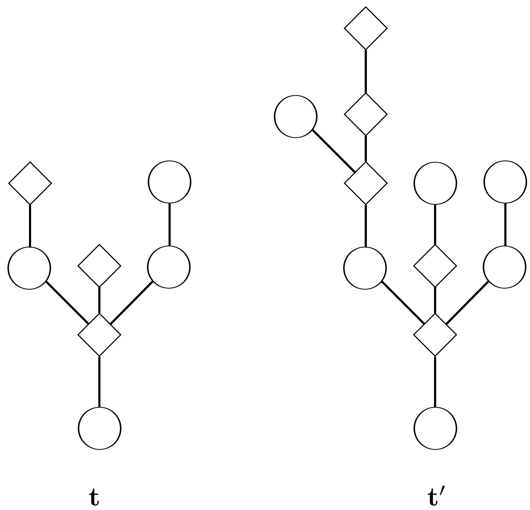

Figure 4.1: An example of a 2-type tree extending another. Here, $\mathbf{t}' \vdash \mathbf{t}$.

This is once again adaptable to forests: if $(\mathbf{f}, \mathbf{e})$ and $(\mathbf{f}', \mathbf{e}')$ are two k-type forests, then we say that $\mathbf{f}' \vdash \mathbf{f}$ if they have the same number of tree components and each tree of $\mathbf{f}'$ extends the corresponding tree of $\mathbf{f}$.

This relation behaves well with Galton-Watson random forests. For example, the following is immediate from the branching property:

Lemma 3.5. *If* $(\mathbf{f}, \mathbf{e})$ *is a finite forest and* $\mathbf{w}$ *the list of types of the roots of its components, then*

$$\mathbb{P}^{(\mathbf{w})}(\mathbf{F} \vdash \mathbf{f}) = \prod_{u \in \mathbf{f}\setminus\partial\mathbf{f}} \zeta^{(\mathbf{e}(u))}(\mathbf{w}_{\mathbf{f}}(u))$$

Moreover, we have a generalization of the branching property: conditionally on $\mathbf{F} \vdash \mathbf{f}$, $\mathbf{F}$ *is obtained by appending independent trees at the leaves of* $\mathbf{f}$*, and for every such leaf* v*, the tree grafted at* v *has distribution* $\mathbb{P}^{(\mathbf{e}(v))}$.

For infinite trees, we get a generalization of (4.3):

Lemma 3.6. *If* $(\mathbf{f},\mathbf{e})$ *is a finite forest, let* $\mathbf{x}$ *be the word formed by the types of the leaves of* $\mathbf{f}$ *in lexicographical order. We have*

$$\widehat{\mathbb{P}}^{(\mathbf{w})}(\mathbf{F}\vdash\mathbf{f})=\frac{Z_{\mathbf{x}}}{Z_{\mathbf{w}}}\mathbb{P}^{(\mathbf{w})}(\mathbf{F}\vdash\mathbf{f}).$$

Proof. Let n be the height of $\mathbf{f}$. Any forest of height n which extends $\mathbf{f}$ can be obtained by adding after each leaf u of $\mathbf{f}$ a tree with height smaller than $n-|u|$. Let $u_1,\dots,u_p$ be the leaves of $\mathbf{f}$, and $e_1,\dots,e_p$ be their types, we will append for all i a tree $(\mathbf{t}^i,\mathbf{e}^i)$ to the leaf u_i and call the resulting forest $(\tilde{\mathbf{f}},\tilde{\mathbf{e}})$, implicitly a function of $\mathbf{f}$ and $\mathbf{t}^1,\dots,\mathbf{t}^p$. Thus, recalling the notation X for the martingale defined in equation (4.2),

$$\begin{aligned}
\widehat{\mathbb{P}}^{(\mathbf{w})}(\mathbf{F}\vdash\mathbf{f})&=\sum_{\mathbf{t}_1,\dots,\mathbf{t}_p}\sum_{v\in\tilde{\mathbf{f}}_n}\frac{b_{\tilde{\mathbf{e}}(v)}}{Z_{\mathbf{w}}}\mathbb{P}^{(\mathbf{w})}\big((\mathbf{F}_{\leqslant n},\mathbf{E}_{\leqslant n})=(\tilde{\mathbf{f}},\tilde{\mathbf{e}})\big)\\
&=\sum_{\mathbf{t}_1,\dots,\mathbf{t}_p}\sum_{i=1}^{p}\sum_{v\in\mathbf{t}^i_{n-|u_i|}}\frac{b_{\tilde{\mathbf{e}}(v)}}{Z_{\mathbf{w}}}\mathbb{P}^{(\mathbf{w})}(\mathbf{F}\vdash\mathbf{f})\prod_{i=1}^{p}\mathbb{P}^{(e_i)}(\mathbf{T}_{\leqslant n-|u_i|}=\mathbf{t}_i)\\
&=\frac{\mathbb{P}^{(\mathbf{w})}(\mathbf{F}\vdash\mathbf{f})}{Z_{\mathbf{w}}}\sum_{i=1}^{p}\sum_{\mathbf{t}_1,\dots,\mathbf{t}_p}\sum_{v\in\mathbf{t}^i_{n-|u_i|}}b_{\mathbf{e}_i(v)}\prod_{i=1}^{p}\mathbb{P}^{(e_i)}(\mathbf{T}_{\leqslant n-|u_i|}=\mathbf{t}_i)\\
&=\frac{\mathbb{P}^{(\mathbf{w})}(\mathbf{F}\vdash\mathbf{f})}{Z_{\mathbf{w}}}\sum_{i=1}^{p}\mathbb{E}^{(e_i)}[X_{n-|u_i|}]\\
&=\frac{\mathbb{P}^{(\mathbf{w})}(\mathbf{F}\vdash\mathbf{f})}{Z_{\mathbf{w}}}\sum_{i=1}^{p}b_{e_i}\\
&=\frac{Z_{\mathbf{x}}}{Z_{\mathbf{w}}}\mathbb{P}^{(\mathbf{w})}(\mathbf{F}\vdash\mathbf{f}).
\end{aligned}$$

□

3.3.5 The case where there is a tree $\mathbf{t}$ such that $\mathbb{P}^{(j)}(\mathbf{T}\vdash\mathbf{t})>0$ and $\mathbf{w}$ is the word formed by the leaves of $\mathbf{t}$

Let $(\mathbf{t},\mathbf{e})$ be a tree with root of type j such that $\mathbb{P}^{(j)}(\mathbf{T}\vdash\mathbf{t})>0$. Let $\mathbf{w}$ be the word formed by the types of the leaves of $\mathbf{t}$, we will prove $(H_{\mathbf{w}})$. We first need an intermediate lemma.

Lemma 3.7. *There exists a a countable family of trees* $(\mathbf{t}^{(2)},\mathbf{e}^{(2)}),(\mathbf{t}^{(3)},\mathbf{e}^{(3)})\dots$ *such that, for any* K*-type tree* $(\mathbf{t}',\mathbf{e}')$ *with root of type* j*:*

- *either* $\mathbf{t}\vdash\mathbf{t}'$.
- *or* $\mathbf{t}'\vdash\mathbf{t}$.
- *or there is a unique* i *such that* $\mathbf{t}'\vdash\mathbf{t}^{(i)}$.

Proof. For all $k\in\{2,3,\dots,ht(\mathbf{t})\}$, take all the trees $(\mathbf{t}',\mathbf{e}')$ which have height k and which satisfy both $(\mathbf{t}'_{\leqslant k-1},\mathbf{e}'_{\leqslant k-1})=(\mathbf{t}_{\leqslant k-1},\mathbf{e}_{\leqslant k-1})$ and $(\mathbf{t}'_k,\mathbf{e}'_k)\neq(\mathbf{t}_k,\mathbf{e}_k)$. These are in countable amount and we can therefore call them $(\mathbf{t}^{(i)},\mathbf{e}^{(i)})_{i\geqslant 2}$ in any order. Now for any tree $(\mathbf{t}',\mathbf{e}')$ with root of type j, by considering the highest integer k such that $(\mathbf{t}'_{\leqslant k-1},\mathbf{e}'_{\leqslant k-1})=(\mathbf{t}_{\leqslant k-1},\mathbf{e}_{\leqslant k-1})$, we directly obtain that, if none of $\mathbf{t}$ and $\mathbf{t}'$ extend the other, then $\mathbf{t}$ extends one of the $\mathbf{t}^{(i)}$. □

Now let $(\mathbf{t}^{(1)}, \mathbf{e}^{(1)}) = (\mathbf{t}, \mathbf{e})$, and, for all $i \in \mathbb{N}$, let also $\mathbf{w}^i$ be the word formed by the types of the leaves of $\mathbf{t}^{(i)}$. Write

$$\begin{aligned}\mathbb{P}^{(j)}(\#_j\mathbf{T} = 1 + dn) &= \sum_{i=1}^{\infty} \mathbb{P}^{(j)}(\mathbf{T} \vdash \mathbf{t}^{(i)}, \#_j\mathbf{T} = 1 + dn) + \mathbb{P}^{(j)}(\mathbf{t} \vdash \mathbf{T}, \mathbf{t} \neq \mathbf{T}, \#_j\mathbf{T} = 1 + dn) \\ &= \sum_{i=1}^{\infty} \mathbb{P}^{(j)}(\mathbf{T} \vdash \mathbf{t^i})\mathbb{P}^{(\mathbf{w}^i)}(\#_j\mathbf{F} = 1 - q^{(i)} + dn) + \mathbb{P}^{(j)}(\mathbf{t} \vdash \mathbf{T}, \mathbf{t} \neq \mathbf{T}, \#_j\mathbf{T} = 1 + dn)\end{aligned}$$

where $q^{(i)}$ is the number of vertices of type j of $\mathbf{t}^{(i)}$ which are not leaves. Divide by $\mathbb{P}^{(j)}(\#_j\mathbf{T} = 1 + dn)$ on both sides of the equation to obtain

$$\sum_{i=1}^{\infty} \mathbb{P}^{(j)}(\mathbf{T} \vdash \mathbf{t}^{(i)}) \frac{\mathbb{P}^{(\mathbf{w}^i)}(\#_j\mathbf{F} = 1 - q^{(i)} + dn)}{\mathbb{P}^{(j)}(\#_j\mathbf{T} = 1 + dn)} + \mathbb{P}^{(j)}(\mathbf{t} \vdash \mathbf{T} \mid \#_j\mathbf{T} = 1 + dn) = 1 \tag{4.6}$$

Note that

$$\mathbb{P}^{(j)}(\mathbf{t} \vdash \mathbf{T} \mid \#_j\mathbf{T} = 1 + dn)$$

tends to 0 as n tends to infinity. This is because we can bound it by $\mathbb{P}(ht(\mathbf{T}) \leqslant ht(\mathbf{t}) \mid \#\mathbf{T} = 1+dn)$ where here $\mathbf{T}$ is a monotype Galton-Watson tree with offspring distribution $\mu_{j,j}$, and this tends to 0 by the monotype case of Theorem 3.3 (proved for example in [49]), since the limiting tree has infinite height.

By Lemma 2.2, we have $1 - q^{(i)} \equiv \alpha_{\mathbf{w}^{(i)}} \pmod d$ for all $i \in \mathbb{N}$, and thus, using the lower bound 4.5, we have

$$\liminf_{n\to\infty} \mathbb{P}^{(j)}(\mathbf{T} \vdash \mathbf{t}^{(i)}) \frac{\mathbb{P}^{(\mathbf{w}^i)}(\#_j\mathbf{F} = 1 - q^{(i)} + dn)}{\mathbb{P}^{(j)}(\#_j\mathbf{T} = 1 + dn)} \geqslant \mathbb{P}^{(j)}(\mathbf{T} \vdash \mathbf{t}^{(i)}) \frac{Z_{\mathbf{w}^i}}{b_j}$$

for all $i \in \mathbb{N}$. However, by Lemma 3.6 and Lemma 3.7, we have

$$\sum_{i=1}^{\infty} \mathbb{P}^{(j)}(\mathbf{T} \vdash \mathbf{t}^{(i)}) \frac{Z_{\mathbf{w}^i}}{b_j} = \sum_{i=1}^{\infty} \widehat{\mathbb{P}}^{(i)}(\mathbf{T} \vdash \mathbf{t}^{(i)}) = 1,$$

and thus, whenever $\mathbb{P}^{(j)}(\mathbf{T} \vdash \mathbf{t}^{(i)})$ is nonzero, we must have

$$\limsup_{n\to\infty} \frac{\mathbb{P}^{(\mathbf{w}^i)}(\#_j\mathbf{F} = 1 - q^{(i)} + dn)}{\mathbb{P}^{(j)}(\#_j\mathbf{T} = 1 + dn)} \leqslant \frac{Z_{\mathbf{w}^i}}{b_j},$$

which ends the proof of $(H_{\mathbf{w}})$.

3.3.6 Removing one element from w

Lemma 3.8. *Let $\mathbf{w} \in \mathcal{W}_K$ be such that $(H_{\mathbf{w}})$ holds. Let m be any integer in $[|\mathbf{w}|]$ and let $\tilde{\mathbf{w}}$ be $\mathbf{w}$, except that we remove w_m from the list. Then $(H_{\tilde{\mathbf{w}}})$ also holds.*

Proof. For $n \in \mathbb{N}$, we split the event $\{\#_j\mathbf{F} = \alpha_{\mathbf{w}} + dn\}$ according to the first and second generations of type j in the m-th tree of the forest. By calling k the number of vertices in the first generation of type j issued from the m-th tree, and then $i_1, \ldots, i_k$ the numbers of vertices in the first generation of type j of each corresponding subtree, we have

$$\mathbb{P}^{(\mathbf{w})}(\#_j\mathbf{F} = \alpha_{\mathbf{w}}+dn) = \sum_k \mu_{w_m,j}(k) \sum_{i_1,\ldots,i_k} \prod_{r=1}^{k} \mu_{j,j}(i_r) \mathbb{P}^{(\tilde{\mathbf{w}}^{i_1+\ldots+i_r})}(\#_j\mathbf{F} = \alpha_{\mathbf{w}}-k-\mathbb{1}_{\{w_m=j\}}+dn)$$

where $\tilde{\mathbf{w}}^{i_1+\ldots+i_r}$ is the word $\mathbf{w}$ where w_m has been replaced by j, repeated $i_1+\ldots+i_r$ times. Note that the term of the sum where $k=0$ is to be interpreted as $\mathbb{P}^{(\tilde{\mathbf{w}})}(\#_j\mathbf{F} = \alpha_{\mathbf{w}} - \mathbb{1}_{\{w_m=j\}} + dn)$.

We now use the same argument as in the end of the previous section: we first divide by $\mathbb{P}^{(\mathbf{w})}(\#_j\mathbf{F} = \alpha_{\mathbf{w}} + dn)$ to get

$$\sum_k \mu_{w_m,j}(k) \sum_{i_1,\ldots,i_k} \prod_{r=1}^k \mu_{j,j}(i_r) \frac{\mathbb{P}^{(\tilde{\mathbf{w}}^r)}(\#_j\mathbf{F} = \alpha_{\mathbf{w}} - k - \mathbb{1}_{\{w_m=j\}} + dn)}{\mathbb{P}^{(\mathbf{w})}(\#_j\mathbf{F} = \alpha_{\mathbf{w}} + dn)} = 1.$$

For each choice of k and $i_1,\ldots,i_k$, using lower bound 4.5 as well as $(H_{\mathbf{w}})$, we have

$$\liminf_{n\to\infty} \mu_{w_m,j}(k) \prod_{r=1}^k \mu_{j,j}(i_r) \frac{\mathbb{P}^{(\tilde{\mathbf{w}}^r)}(\#_j\mathbf{F} = \alpha_{\mathbf{w}} - k - \mathbb{1}_{\{w_m=j\}} + dn)}{\mathbb{P}^{(\mathbf{w})}(\#_j\mathbf{F} = \alpha_{\mathbf{w}} + dn)} \geqslant \mu_{w_m,j}(k) \prod_{r=1}^k \mu_{j,j}(i_r) \frac{Z_{\tilde{\mathbf{w}}} + \sum_r i_r}{Z_{\mathbf{w}}}.$$

A repeated use of point (iii) of Lemma 2.1 shows that these add up to 1, and thus, for k and $i_1,\ldots,i_k$ such that $\mu_{w_m,j}(k)\prod_{r=1}^k \mu_{j,j}(i_r) \neq 0$, we do have

$$\lim_{n\to\infty} \frac{\mathbb{P}^{(\tilde{\mathbf{w}}^r)}(\#_j\mathbf{F} = \alpha_{\mathbf{w}} - k - \mathbb{1}_{\{w_m=j\}} + dn)}{\mathbb{P}^{(\tilde{\mathbf{w}})}(\#_j\mathbf{F} = \alpha_{\mathbf{w}} + dn)} = \frac{Z_{\tilde{\mathbf{w}}} + \sum_{r=1}^k i_r}{Z_{\mathbf{w}}}.$$

By irreducibility, one can find k such that $\mu_{w_m,j}(k) \neq 0$, and by criticality one has $\mu_{j,j}(0) \neq 0$, meaning that we can take $i_1,\ldots,i_k$ all equal to zero, and this ends the proof. ☐

3.3.7 End of the proof

By applying Lemma 3.8 repeatedly and using the fact that $(H_{\mathbf{w}})$ stays true if we permute the terms of $\mathbf{w}$, we obtain that, if $\mathbf{w}$ and $\mathbf{w}'$ are two words such that any type features fewer times in $\mathbf{w}'$ than in $\mathbf{w}$, then $(H_{\mathbf{w}})$ implies $(H_{\mathbf{w}'})$. Thus, by Section 3.3.5, we now only need to show the following lemma.

Lemma 3.9. *For all nonnegative integers $n_1,\ldots,n_K$, there exists a K-type tree $(\mathbf{t},\mathbf{e})$ which has more than n_i leaves of type i for all $i\in[K]$, and such that $\mathbb{P}^{(j)}(\mathbf{T}\vdash\mathbf{t})>0$.*

Proof. The first step is showing that, for p large enough, the p-th generation of type j of $\mathbf{T}$ has positive probability of having more than $n_1+\ldots+n_K$ vertices, where the p-th generation of type j is the set of vertices of type j which have exactly p ancestors of type j including the root. This is immediate because the average of $\mu_{j,j}$ is 1 and we are not in a degenerate tree, and thus the size of each generation of type j has positive probability of being strictly larger than the previous generation.

Irreducibility then tells us that, after each vertex of the p-th generation of type j, there is a positive probability of finding a vertex of type i for any i. ☐

4 Background on random planar maps

4.1 Planar maps

As stated in the introduction, a planar map is a proper embedding m of a finite connected planar graph in the sphere, in the sense that edges do not intersect. These are taken up to orientation-preserving homeomorphisms of the sphere, thus making them combinatorial objects. We call *faces* of a map m the connected components of its complement in the sphere, and let $\mathcal{F}_m$ be their set. The *degree* of a face f, denoted by $\deg(f)$, is the number of edges it is adjacent to, counting multiplicity: we count every edge as many times as we encounter it when circling around f.

We are going to look at maps which are both *rooted* and *pointed*. These are triplets (m, e, r), where m is a planar map, e is an oriented edge of m called the root edge, starting at a vertex e^- and pointing to a vertex e^+, and r is a vertex of m. We call $\mathcal{M}$ the set of all such maps and $\mathcal{M}_n$ the set of such maps with n vertices for $n \in \mathbb{N}$. A map (m, e, r) will be called positive (resp. null, negative) if $d(r, e^+) = d(r, e^-) + 1$ (resp. $d(r, e^-)$, $d(r, e^-) - 1$). We call $\mathcal{M}^+$, $\mathcal{M}^0$ and $\mathcal{M}^-$ the corresponding sets of maps and, for $n \in \mathbb{N}$, $\mathcal{M}_n^+$, $\mathcal{M}_n^0$ and M_n^- the corresponding sets of maps which have n vertices. Since there is a trivial bijection between positive and negative maps, we will mostly restrict ourselves to $\mathcal{M}^+$ and $\mathcal{M}^0$. By convention, we add to $\mathcal{M}^+$ the vertex map $\dagger$, which consists of one vertex, no edges and one face.

4.2 Boltzmann distributions

Let $\mathbf{q} = (q_n)_{n \in \mathbb{N}}$ be a sequence of nonnegative numbers such that there exists $i \geqslant 3$ such that $q_i > 0$. For any map m, let

$$W_{\mathbf{q}}(m) = \prod_{f \in \mathcal{F}_m} q_{\deg(f)}.$$

We say that the sequence $\mathbf{q}$ is *admissible* if the sum

$$Z_{\mathbf{q}} = \sum_{(m,e,r) \in \mathcal{M}} W_{\mathbf{q}}(m)$$

is finite. When $\mathbf{q}$ is admissible, we can define the Boltzmann probability distribution $B_{\mathbf{q}}$ by setting, for a pointed rooted map (m, e, r),

$$B_{\mathbf{q}}(m, e, r) = \frac{W_{\mathbf{q}}(m)}{Z_{\mathbf{q}}}.$$

We also introduce the versions of $B_{\mathbf{q}}$ conditioned to be positive or null: let $Z_{\mathbf{q}}^+ = \sum_{(m,e,r) \in \mathcal{M}^+} W_{\mathbf{q}}(m)$ and $Z_{\mathbf{q}}^0 = \sum_{(m,e,r) \in \mathcal{M}^0} W_{\mathbf{q}}(m)$ and, for any map (m, e, r), $B_{\mathbf{q}}^+(m, e, r) = \frac{W_{\mathbf{q}}(m)}{Z_{\mathbf{q}}^+}$ if it is positive and $B_{\mathbf{q}}^0(m, e, r) = \frac{W_{\mathbf{q}}(m)}{Z_{\mathbf{q}}^0}$ if it is null. Lastly, given an integer n, we introduce versions of $B_{\mathbf{q}}$, $B_{\mathbf{q}}^+$ and $B_{\mathbf{q}}^0$ where we also condition the map to have n vertices (for n such that this event has positive probability), which we call $B_{\mathbf{q}}^n$, $B_{\mathbf{q}}^{n,+}$ and $B_{\mathbf{q}}^{n,0}$

For nonnegative numbers x and y, let

$$f^\bullet(x, y) = \sum_{k,k'} \binom{2k + k' + 1}{k + 1} \binom{k + k'}{k} q_{2+2k+k'} \, x^k y^{k'}$$

and

$$f^\diamond(x, y) = \sum_{k,k'} \binom{2k + k'}{k} \binom{k + k'}{k} q_{1+2k+k'} \, x^k y^{k'}.$$

It was shown in [64], Proposition 1, that $\mathbf{q}$ is admissible if and only if the system

$$1 - \frac{1}{x} = f^\bullet(x, y) \tag{4.7}$$

$$y = f^\diamond(x, y) \tag{4.8}$$

has a solution with $x > 1$, such that the spectral radius of the matrix

$$\begin{pmatrix} 0 & 0 & x-1 \\ \frac{x}{y}\partial_x f^{\diamond}(x,y) & \partial_y f^{\diamond}(x,y) & 0 \\ \frac{x^2}{x-1}\partial_x f^{\bullet}(x,y) & \frac{xy}{x-1}\partial_y f^{\bullet}(x,y) & 0 \end{pmatrix}$$

is smaller than or equal to 1. The existence of such a solution implies its uniqueness, with $x = Z^+_{\mathbf{q}}$ and $y = \sqrt{Z^0_{\mathbf{q}}}$. We let $Z^{\diamond}_{\mathbf{q}} = \sqrt{Z^0_{\mathbf{q}}}$.

We then say that $\mathbf{q}$ is *critical* if the spectral radius of the aforementioned matrix is exactly 1.

4.3 The Bouttier-Di Francesco-Guitter bijection

In [19] was exposed a bijection between rooted and pointed maps and a certain class of 4-type labelled trees called *mobiles*. Let us quickly recall the facts here, with a few variations to make the bijection more adapted to our study.

4.3.1 Mobiles

A finite spatial 4-type tree $(\mathbf{t}, \mathbf{e}, \mathbf{l})$ is called a *mobile* if the types satisfy the following conditions:

- The root has type 1 or 2,
- The children of a vertex of type 1 all have type 3,
- If a vertex has type 2, then it has only one child, which has type 4, except if it is the root, if $\emptyset$ has type 2 then it has exactly two children, both of type 4,
- Vertices of type 3 and 4 can only have children of types 1 and 2,

and the labels satisfy the following conditions:

- Vertices of type 1 and 3 have integer labels, vertices of type 2 and 4 have labels in $\mathbb{Z} + \frac{1}{2}$,
- The root has label 0 if it is of type 1, $\frac{1}{2}$ if it is of type 2,
- Vertices of type 3 or 4 have the same label as their father.
- If $u \in \mathbf{t}$ has type 3 or 4, let by convention $u0 = u\underline{k_u(\mathbf{t}) + 1} = u^-$. Then, for all $i \in \{1, \ldots, k_u(\mathbf{t}+1)\}$, $\mathbf{l}\big(u\underline{i+1}\big) - \mathbf{l}(ui) \geqslant -\frac{1}{2}(\mathbb{1}_{\{\mathbf{e}(ui)=1\}} + \mathbb{1}_{\{\mathbf{e}(u\underline{i+1})=1\}})$.

The notation $u\underline{i+1}$ means that we are looking at $i+1$ as a letter, the word $u\underline{i+1}$ being the concatenation of u and $i+1$.

Traditionally, vertices of type 1 are represented as white circles $\bigcirc$, vertices of type 2 are "flags" $\diamond$ while the other two types are dots $\bullet$. Notice also that we do not need to mention the labels of vertices with type 3 and 4 since the label of such a vertex is the same as that of its father. We let $\mathbb{T}_M$ be the set of finite mobiles, $\mathbb{T}^+_M$ be the set of finite mobiles such that $\mathbf{e}(\emptyset) = 1$ and $\mathbb{T}^0_M$ be the set of finite mobiles such that $\mathbf{e}(\emptyset) = 2$

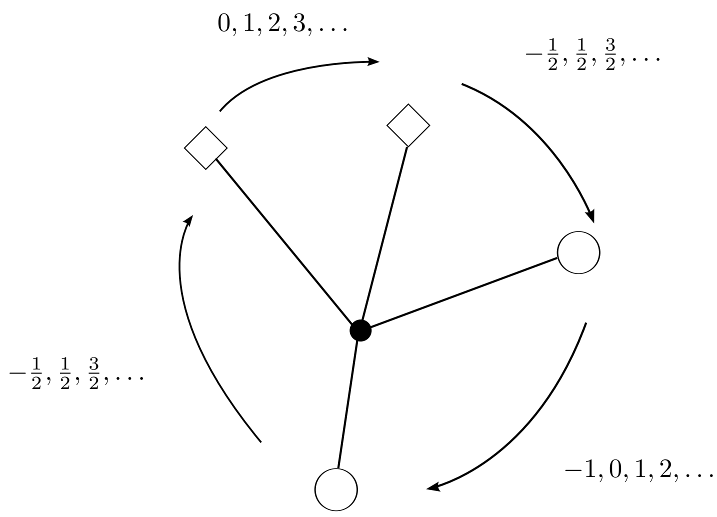

Figure 4.2: The authorized labelling differences when circling around a vertex of type 3 or 4.

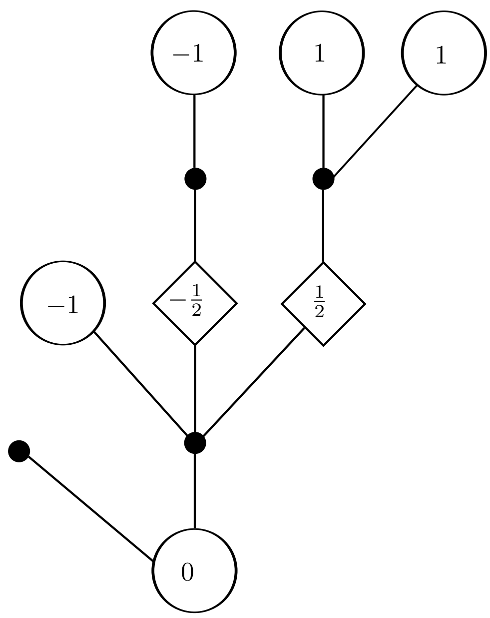

Figure 4.3: An example of a mobile, with root of type 1.

4.3.2 The bijection

Let $(\mathbf{t}, \mathbf{e}, \mathbf{l})$ be a mobile and let us describe how to transform it into a map. Let $v_1, v_2, \dots, v_p$ be, in order, the vertices of type 1 or 2 of $\mathbf{t}$ appearing in the standard contour process and $e_1, e_2, \dots, e_p$ and $l_1, l_2, \dots, l_p$ be the corresponding types and labels. We refer to $v_1, \dots, v_p$ as the *corners* of the tree because a vertex will be visited a number of times equal to the number of angular sectors around it delimited by the tree. Draw $\mathbf{t}$ in the plane and add an extra type 1 vertex r outside of $\mathbf{t}$, giving it label $\min_{\mathbf{e}(u)=1} \mathbf{l}(u) - 1$. Now, for every $i \in [p]$, define the successor of the i-th corner as the next corner of type 1 with label $l_i - 1$ if $e_i = 1$ and $l_i - \frac{1}{2}$ is $e_i = 2$. If there is no such vertex, then let its successor be r. In both cases, draw an arc between v_i and the successor. This construction can be done without having any of the arcs intersect. Now erase all the original edges of the tree, as well as vertices of types 3 and 4. Erase as well all the vertices of type 2, merging the corresponding pairs of arcs. We are left with a planar map, with a distinguished vertex r. The root edge depends on the type of the root of the tree: if $\mathbf{e}(\emptyset) = 1$ then we let the root edge be the first arc which was drawn (have it point to $\emptyset$ for a positive map, and away from $\emptyset$ for a negative map). If $\mathbf{e}(\emptyset) = 2$ then we let the root edge be the result of the merging of the two edges adjacent to $\emptyset$, pointing to the successor of the first corner encountered in the contour process.

This construction gives us two bijections: one between $\mathbb{T}_M^+$ and $\mathcal{M}^+$ and one between $\mathbb{T}_M^0$ and $\mathcal{M}^0$, which we both call Ψ.

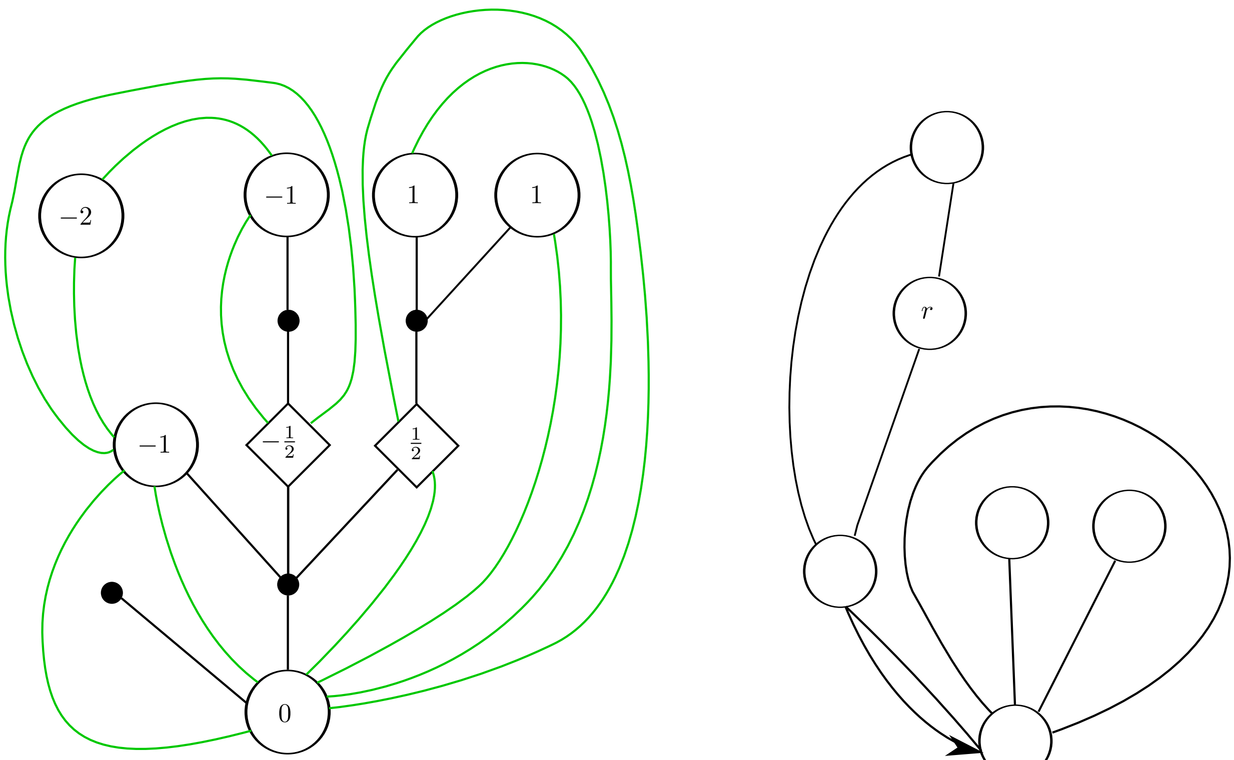

Figure 4.4: Having added a vertex with label -2 to the mobile of Figure 4.3, we transform it into a map.

It was shown in [64] that the BDFG bijection serves as a link between Galton-Watson mobiles and Boltzmann maps.

Proposition 4.1. *Consider and admissible weight sequence* $\mathbf{q}$ *and define an unordered* 4*-type* *offspring distribution* μ *by*

$$\mu^{(1)}(0,0,k,0) = \frac{1}{Z_{\mathbf{q}}^{+}}(1-\frac{1}{Z_{\mathbf{q}}^{+}})^{k}$$
$$\mu^{(2)}(0,0,0,1) = 1$$
$$\mu^{(3)}(k,k',0,0) = \frac{(Z_{\mathbf{q}}^{+})^{k}(Z_{\mathbf{q}}^{\diamond})^{k'}\binom{2k+k'+1}{k+1}\binom{k+k'}{k}q_{2+2k+k'}}{f^{\bullet}(Z_{\mathbf{q}}^{+},Z_{\mathbf{q}}^{\diamond})}$$
$$\mu^{(4)}(k,k',0,0) = \frac{(Z_{\mathbf{q}}^{+})^{k}(Z_{\mathbf{q}}^{\diamond})^{k'}\binom{2k+k'}{k}\binom{k+k'}{k}q_{1+2k+k'}}{f^{\diamond}(Z_{\mathbf{q}}^{+},Z_{\mathbf{q}}^{\diamond})}.$$

Let then ζ *be the ordered offspring distribution which is uniform ordering of* μ*, as explained in Section 2.1. This offspring distribution is irreducible, and it is critical if the weight sequence* $\mathbf{q}$ *is critical, while it is subcritical if* $\mathbf{q}$ *is admissible but not critical. Define also, for all ordered offspring type-list* $\mathbf{w}$*,* $\nu_{\mathbf{w}}^{(i)}$ *as the uniform measure on the set* $D_{\mathbf{w}}^{(i)}$ *of allowed displacements to have a mobile, which is precisely* $D_{\mathbf{w}}^{(i)} = \{0\}^{|\mathbf{w}|}$ *if* $i=1$ *or* $i=2$ *and*

$$D_{\mathbf{w}}^{(i)} = \left\{\mathbf{y}=(y_i)_{i\in[|\mathbf{w}|]}:\ \ \forall i\in\{0,1,\dots,|\mathbf{w}|\},\ y_{i+1}-y_i+\frac{1}{2}(\mathbb{1}_{\{w_i=1\}}+\mathbb{1}_{\{w_{i+1}=1\}})\in\mathbf{Z}_+\right\},$$

if $i=3$ *or* $i=4$*, in which case we set by convention* $w_0 = w_{|\mathbf{w}|+1} = i-2$ *and* $y_0 = y_{|\mathbf{w}|+1} = 0$.

Then:

- *if* $(\mathbf{T},\mathbf{E},\mathbf{L})$ *has distribution* $\mathbb{P}_{\zeta,\nu}^{(1),(0)}$*, then the random map* $\Psi(\mathbf{T},\mathbf{E},\mathbf{L})$ *has distribution* $\mathbb{B}_{\mathbf{q}}^{+}$.
- *if* $(\mathbf{F},\mathbf{E},\mathbf{L})$ *is a forest with distribution* $\mathbb{P}_{\zeta,\nu}^{(2,2),(\frac{1}{2},\frac{1}{2})}$*, consider the mobile formed by merging both tree components at their roots. The image of this mobile by* Ψ *has law* $\mathbb{B}_{\mathbf{q}}^{0}$.

Remark 4.2. *The operation of merging two trees at their roots can be formalized the following way. Consider two trees* $(\mathbf{t}_1,\mathbf{e}_1)$ $(\mathbf{t}_2,\mathbf{e}_2)$ *which are such that, in both trees, the root has type* 2 *and has a unique child, with type* 4*. For* $u\in\mathbf{t}_2\setminus\{\emptyset\}$*, we can write* $u = 1u^2\dots u^k$*. Let then* $u' = 2u^2\dots u^k$*, and let* $\mathbf{t}_2' = \{u',\ u\in\mathbf{t}_2\setminus\{\emptyset\}\}$*. We can now define* $\mathbf{t} = \mathbf{t}_1\cup\mathbf{t}_2'$*, which is easily checked to be a tree. Types can then simply be assigned by setting, for* $u\in\mathbf{t}_1$*,* $\mathbf{e}(u)=\mathbf{e}_1(u)$ *and, for* $u\in\mathbf{t}_2\setminus\{\emptyset\}$*,* $\mathbf{e}(u')=\mathbf{e}_2(u)$.

This operation is of course continuous for the local convergence topology since, for any $k\in\mathbb{Z}_+$*, the* k*-th generation of* $\mathbf{t}$ *is completely determined by the* k*-th generations of* $\mathbf{t}_1$ *and* $\mathbf{t}_2$.

Remark 4.3. *If the weight sequence* $\mathbf{q}$ *is such that* $q_{2n+1}=0$ *for all* $n\in\mathbb{Z}_+$*, then a* $\mathbf{q}$*-Boltzmann map is a.s. bipartite, which implies that there will be no vertices of type* 2 *or* 4 *in the corresponding tree, in which case we consider the mobile as a tree with two types, and it stays irreducible. Moreover, we then have* $Z_{\mathbf{q}}^{\diamond}=0$.

4.4 Infinite maps and local convergence

If (m,e) is a rooted map and $k\in\mathbb{N}$, we let $B_{m,e}(k)$ be the map formed by all vertices whose graph distance to e^+ is less than or equal to k, and all edges connecting such vertices, except if the distance between each vertex and the e^+ is exactly k. The map $B_{m,e}(k)$ is still rooted at the same oriented edge e. For two rooted maps (m,e) and (m',e'), let $d\big((m,e),(m',e')\big) = \frac{1}{1+p}$ where p is the supremum of all integers k such that $B_{m,e}(k)$ is equivalent to $B_{m',e'}(k)$. This defines a metric on the set of rooted maps. Call then $\overline{\mathcal{M}}$ the completion of this set. Elements of $\overline{\mathcal{M}}$ which are not finite maps are then called *infinite maps*, which we mostly consider as a sequence of compatible finite maps: $(m,e) = (m_i, e_i)_{i\in\mathbb{N}}$ with $(m_i, e_i) = B_{m_{i+1},e_{i+1}}(i)$ for all i. Note in particular that infinite maps are not pointed.

As with trees and forests, convergence in distribution is simply characterized: if $(M_n, E_n)_{n\in\mathbb{N}}$ is a sequence of random rooted maps, one can check that it converges in distribution to a certain random map (M, E) if and only if, for all finite deterministic maps (m', e') and all $k \in \mathbb{N}$, $\mathbb{P}\big(B_{(M_n,E_n)}(k) = (m', e')\big)$ converges to $\mathbb{P}\big(B_{(M,E)}(k) = (m', e')\big)$

5 Convergence to infinite Boltzmann maps

We now take a critical a critical weight sequence $\mathbf{q}$, and take μ, ζ and ν as defined in Proposition 4.1. Since, in the BDFG bijection, the vertices of the map correspond to the vertices of type 1 of the tree (and one extra vertex), we expect Theorem 3.3 to tell us that Boltzmann maps with large amounts of vertices converge locally. This section is dedicated to establishing the fact that this is indeed the case.

Let d be the greatest common divisor of all integers m which are either such that $q_{2m+2} > 0$ or which are odd and such that $q_{m+2} > 0$:

$$d = gcd\Big(\{m \in \mathbb{N} : q_{2m+2} > 0\} \cup \{m \in 2\mathbb{Z}_+ + 1 : q_{m+2} > 0\}\Big)$$

Theorem 5.1. *For $n \in \mathbb{N}$ such that $B_{\mathbf{q}}$ gives mass to maps with n vertices, let (M_n, E_n, R_n) be a variable with distribution $B^n_{\mathbf{q}}$. We then have*

$$(M_n, E_n) \underset{\substack{n\to\infty \\ n\in 2+d\mathbb{Z}_+}}{\Longrightarrow} (M_\infty, E_\infty)$$

in distribution for the local convergence, where (M_∞, E_∞) is an infinite rooted map which we call the infinie $\mathbf{q}$-Boltzmann map.

The choice of the subsequence $(2 + dn)_{n\in\mathbb{Z}_+}$ is explained by the fact that the number of vertices of a map with distribution $B_{\mathbf{q}}$ can only be of the form $2 + dn$ for integer n. This will be explained in Section 5.2.1, as will be the fact that, for n large enough, $B_{\mathbf{q}}$ does give mass to maps with $2 + dn$ vertices.

The infinite map (M_∞, E_∞) is moreover *planar*, in the sense that it is possible to embed it in the plane in such a way that bounded subsets of the plane only encounter a finite number of edges.

5.1 The example of uniform p-angulations

Here we take an integer $p \geqslant 3$ and consider maps which only have faces of degree p, which we call p-angulations. The well-known Euler's formula will show us that the number of faces of such a map is directly connected to its number of vertices. Let m be a finite p-angulation, and let V be its number of vertices, E be its number of edges and F be its number of faces. Since each edge is adjacent to two faces, we have $pF = 2E$. Euler's formula, on the other hand, states that $V - E + F = 2$. Combining the two shows that

$$V = 2 + (\frac{p}{2} - 1)F.$$

At this point, we must split the discussion according to the parity of p.

Uniform infinite 2p-angulation Let $p \geqslant 2$. It has been shown in [63] that the weight sequence $\mathbf{q}$ defined by

$$q_n = \frac{(p-1)^{p-1}}{p^p \binom{2p-2}{p-1}} \mathbb{1}_{\{n=2p\}}$$

is critical. Since the weight of a map here only depends on its number of faces (or vertices), it is immediate that conditioning the distribution $B_{\mathbf{q}}$ to the set of maps with $2+(p-1)n$ vertices yields the distribution of the uniform $2p$-angulation with n faces. We thus obtain the following.

Proposition 5.2 (Uniform $2p$-angulation)**.** *Let $p \geqslant 2$ and, for $n \in \mathbb{N}$, let (M_n, E_n) be a uniform rooted map amongst the set of rooted $2p$-angulation with n faces. Then (M_n, E_n) converges locally in distribution as n goes to infinity, the limit being a random rooted map which we call the* uniform infinite $2p$-angulation*.*

In the case where $2p = 4$, we obtain the local convergence in distribution of large uniform quadrangulation to the UIPQ which was first obtained by Krikun in [53]. In fact our method here ends up being essentially the same as that of [26], where we have used the BDFG bijection in a situation where the simpler Cori-Vauquelin-Schaeffer bijection would have sufficed.

Uniform infinite $2p+1$-angulation Let $p \in \mathbb{N}$ and consider $2p+1$-angulations. It follows the relation $V = (p-1/2)F$ that a $2p+1$-angulation must have an even number of faces, and, for $n \in \mathbb{N}$, a $2p+1$-angulation with $2n$ faces has $2+(2p-1)n$ vertices. As in the even case, a uniform $2p+1$-angulation with a prescribed vertex can be seen as a conditioned Boltzmann-distributed random map for the weight sequence $\mathbf{q}$ defined by

$$q_n = \alpha \mathbb{1}_{\{n=2p+1\}},$$

for any positive number α. It has been shown in [25], Proposition A.2 that there is one value of α which makes this sequence critical. Theorem 5.1 then gives us the following.

Proposition 5.3 (Uniform infinite $2p+1$-angulation)**.** *Let $p \in \mathbb{N}$ and, for $n \in \mathbb{N}$, let (M_n, E_n) be a uniform rooted map amongst the set of rooted $2p+1$-angulation with $2n$ faces. Then (M_n, E_n) converges locally in distribution as n goes to infinity, to a random rooted map called the* uniform infinite $(2p+1)$-angulation*.*

5.2 Proof of Theorem 5.1

The proof of Theorem 5.1 involves first showing the convergence for maps conditioned to be null or positive by using the BDFG bijection and identifying the limiting map as the image of an infinite tree by the bijection, and then removing the conditionings.

5.2.1 On the trees associated to $B_{\mathbf{q}}$

We want to investigate the periodic structure of Galton-Watson trees with ordered offspring distribution ν. We thus adopt the notations of Section 2.2 and 2.3, our reference type being 1: for i in $\{1,2,3,4\}$, $\mu_{i,1}$ is the distribution of the size of the first generation of type 1 in a tree with distribution $\mathbb{P}_\zeta^{(i)}$, d is the greatest common divisor of the support of $\mu_{1,1}$ and, for $i \in \{1,2,3,4\}$, β_i is the common value of all elements of the support of $\mu_{i,1}$ modulo d. To remove any confusion with the previous section, we let $d' = gcd\big(\{m \in \mathbb{N},\, q_{2m+2} > 0\} \cup \{m \in 2\mathbb{Z}_+ + 1,\, q_{m+2} > 0\}\big)$, which was called d in the previous section.

Lemma 5.4. *We have*

- $d = d'$.
- $\beta_3 = 0$.

Moreover, if the weight sequence $\mathbf{q}$ *is not bipartite, then we also have*

- $2\beta_2 \equiv 1 \pmod{d}$.
- $\beta_4 = \beta_2$.

Proof. We first treat the bipartite case separately. In this case, types 1 and 3 alternate in tree, and it is straightforward that $d = gcd\big(\{m \in \mathbb{N},\, q_{2m+2} > 0\}\big)$ and that $\beta_3 = 0$. We now assume not to be in this case.

It is immediate that $\beta_3 = 0$ and $\beta_4 = \beta_2$ because a vertex of type 1 or 2 can give birth to a single vertex of type 2 or 4, respectively.

To prove the other two more interesting equations, first take $m \in \mathbb{N}$ such that $q_{2m+2} > 0$. Using the fact that a vertex of type 3 can give birth to m vertices of type 1, one obtains $m \equiv 0 \pmod{d}$.

Now take an odd integer $m = 2n + 1$ such that $q_{m+2} = q_{2n+3} > 0$. A vertex of type 3 can then give birth to n vertices of type 1 and one vertex of type 2, and a vertex of type 4 can give birth to $n + 1$ vertices of type 1. We thus obtain $n + \beta_2 \equiv 0 \pmod{d}$ and $n + 1 \equiv \beta_2 \pmod{d}$. Combining these nets us $m = 2n + 1 \equiv 0 \pmod{d}$ and $2\beta_2 \equiv m + 1 \equiv 1 \pmod{d}$.

We have thus shown $2\beta_i \equiv 1 \pmod{d}$ and that d divides d'. To show that they are equal we require some more refined analysis.

Notice that for words $\mathbf{w} = (k, k', 0, 0)$ such that $\mu^{(3)}(\mathbf{w}) > 0$, we have $2k + k' \equiv 0 \pmod{d'}$. Indeed, if k' is even, letting $n = k + \frac{k'}{2}$, we then have $q_{2n+2} > 0$, implying that d' divides n, while if k' is odd, we let $n = 2k + k'$, and then $q_{n+2} > 0$ and therefore d' divides n. Similarly, if $\mu^{(4)}(\mathbf{w}) > 0$, then $2k + k' \equiv 1 \pmod{d'}$. Applying this repeatedly to a tree $(\mathbf{t}, \mathbf{e})$ such that $\mathbb{P}^{(1)}_{\zeta}(\mathbf{T} \vdash \mathbf{t})$ and such that all its leaves are of type 1 or 2, one obtains $2k + k' = 0 \pmod{d'}$ where k and k' are respectively the number of leaves of type 1 and 2 in $\mathbf{t}$. Taking $(\mathbf{t}, \mathbf{e})$ which has only one generation of type 1, and we do obtain that d' divides every member of the support of $\mu_{1,1}$. □

5.2.2 Infinite mobiles and the BDFG bijection

We call an infinite mobile any infinite 4-type labelled tree $(\mathbf{t}, \mathbf{e}, \mathbf{l})$ which satisfies the conditions of Section 4.3.2, which has a unique infinite spine and such that the labels of vertices of type 1 of the spine do not have a lower bound. We let $\overline{\mathbb{T}}_M$ be the set of all finite and infinite mobiles, and split it in $\overline{\mathbb{T}}_M = \overline{\mathbb{T}}^+_M \cup \overline{\mathbb{T}}^0_M$ as before.

The BDFG bijection Ψ can be naturally extended to $\overline{\mathbb{T}}_M$. Let $(\mathbf{t}, \mathbf{e}, \mathbf{l}) \in \overline{\mathbb{T}}_M$, we let $(u_n)_{n\in\mathbb{N}}$ be the sequence of the elements of the spine. This sequence splits the tree in two: the part which is on the left-hand side of the spine, and the part which is on the right-hand side. To be precise, we say that $v \in \mathbf{t}$ is on the left-hand side of the spine if there exists three integers n, k and l such that $v = u_n k$, $u_{n+1} = u_n l$ and $k \leqslant l$, and v is on the right-hand side if we have the same, but with $k \geqslant l$.

This splitting allows us to define a contour process, but it has to be indexed by $\mathbb{Z}$: since every subtree branching out of the spine is finite, we can let $\big(v(n)\big)_{n\in\mathbb{N}}$ be the contour process of the left-hand side and $\big(v(-n)\big)_{n\in\mathbb{N}}$ be the other half. This determines a unique sequence $\big(v(n)\big)_{n\in\mathbb{Z}}$. Since we have assumed that the labels of the vertices of type 1 do not have a lower bound, the notion of successor we used for finite trees is still valid, and in fact, unlike in the case of a finite tree, we do not need to add an extra vertex. As in the finite case, we connect every vertex of type 1 or 2 to its successor, erase all the original edges of the tree, erase vertices of types 2, 3 and 4, merging the two edges adjacent to every vertex of type 2. This leaves us with an infinite map (by construction, the arcs do not intersect, while the following Lemma 5.5 implies that it is locally finite). We give this map a root edge which is determined with the same rules as in the finite case, however it is not pointed. We call this rooted map $\Psi(\mathbf{t}, \mathbf{e}, \mathbf{l})$.

Lemma 5.5. *The extended BDFG function* Ψ *is continuous on* $\overline{\mathbb{T}}_M$.

Proof. Let $(\mathbf{t}, \mathbf{e}, \mathbf{l})$ be an infinite mobile. We assume $\mathbf{e}(\emptyset) = 1$, the other case can be treated the same way. For $n \in \mathbb{N}$, we need to find $p \in \mathbb{N}$ such that, for another mobile $(\mathbf{t}', \mathbf{e}', \mathbf{l}')$, if $(\mathbf{t}_{\leqslant p}, \mathbf{e}_{\leqslant p}, \mathbf{l}_{\leqslant p}) = (\mathbf{t}_{\leqslant p}, \mathbf{e}_{\leqslant p}, \mathbf{l}_{\leqslant p})$ then $B_{m,e}(n) = B_{m',e'}(n)$, where $(m, e, r) = \Psi(\mathbf{t}, \mathbf{e}, \mathbf{l})$ and $(m', e', r') = \Psi(\mathbf{t}', \mathbf{e}', \mathbf{l}')$. Let $s \in \mathbb{N}$ be large enough such that all the arcs in $B_m(n)$ connect vertices of $\mathbf{t}_{\leqslant s}$, let $x = \inf_{v \in \mathbf{t}_{\leqslant s}} \mathbf{l}(v)$ and let u be any type 1 vertex of the spine such that $\mathbf{l}(u) < x-1$. Notice now that there are no arcs connecting $\mathbf{t}_{\leqslant s}$ and the subtree above u. Indeed, the successor of any vertex of $\mathbf{t}_{\leqslant s}$ will be encountered below u while, if v is above u, $\mathbf{l}(v) \geqslant x$ would imply that its successor is also above u, while $\mathbf{l}(v) \leqslant x - 1$ would make it impossible for its successor to be in $B_{(t,l)}(s)$. Taking p to be the height of u then ends the proof. □

Lemma 5.6. *For any infinite mobile* $(\mathbf{t}, \mathbf{e}, \mathbf{l})$*, the infinite map* $\Psi(\mathbf{t}, \mathbf{e}, \mathbf{l})$ *is planar, in the sense that it can be embedded in the plane in such a way that bounded subsets of the plane only encounter a finite number of edges.*

Proof. We first start by embedding the mobile in the plane in a convenient way. We draw its infinite spine as the subset $\{0\} \times \mathbb{Z}_+$, where the child of $(0, n)$ is $(0, n+1)$ for $n \in \mathbb{Z}_+$. Starting from this, we can then embed the tree in $\mathbb{Z} \times \mathbb{Z}_+$ such that the second coordinate of a vertex is always its graph distance to the root, and also such that the children of any vertex u always form a set of the type $\{(n, m), (n+1, m), \dots, (n + k_u(\mathbf{t}) - 1, m)\}$, with $n \in \mathbb{Z}$ and $m \in \mathbb{N}$ and their first coordinates are in the correct order. With such an embedding, it is also apparent that there exists a continuous function $f : \mathbb{R} \to \mathbb{R}$ with is decreasing on $(-\infty, 0]$ and increasing on $[0, +\infty)$, which has limit $+\infty$ at both $-\infty$ and $+\infty$ such that $\mathbf{t}$ is strictly above the graph of f.

We point out an important fact of the bijection: let u be any corner of type 1 or 2 and let v be its successor. Then, for any corner w of type 1 or 2 which is encountered between u and v in the contour process of the mobile, the successor of w, which we call x, is then encountered between w and v, and the arc between w and x is then enclosed between $\mathbf{t}$ and the arc connecting u and v. From this fact, we obtain that all the arcs which connect two points on the left-hand side of the tree can be embedded without any issues: first draw the arcs connecting the line of successors starting at the root, and enclose in each of them the other necessary arcs.

The arcs which originate from the right-hand side of the tree are a more complex issue, because some of them might start very high on the right-hand side, go around a large part of the tree and end up high on the left-hand side. To make sure that these are well separated, we introduce for $n \in \mathbb{Z}_+$ the "strip"

$$S_n = \Big\{ (x, y) \in \mathbb{R}^2 : \ f(x) - n + 1 \leqslant y < f(x) - n \Big\}$$

We now explore the right-hand side of the tree in counter-clockwise order and, when we encounter the n-th corner of type 1 or 2, we join it to S_n. We point out that it is possible to do this in such a way that second coordinate along the path is nondecreasing. We then do the same thing for the corner's successor, and then join both halves by a path which stays in S_n.

The paths we have drawn this way still do not intersect because of the "enclosure" property as before, and this embedding is indeed such that bounded subsets of $\mathbb{R}^2$ only encounter a finite number of edges. This is because we have split these edges in parts which are in S_n, of which there is only one for every n, and parts which originate from vertices of the tree and have nondecreasing second coordinate, of which there are a finite amount in bounded subsets because there is a finite amount of vertices of $\mathbf{t}$ with bounded second coordinate. □

5.2.3 Behaviour of the labels on the spine of the infinite tree

Let $(\mathbf{T}, \mathbf{E}, \mathbf{L})$ be a 4-tree with law $\widehat{\mathbb{P}}^{(1),(0)}_{\zeta,\nu}$ or the tree obtained from merging both components of a forest with distribution $\widehat{\mathbb{P}}^{(2,2),(\frac{1}{2},\frac{1}{2})}_{\zeta,\nu}$ at their roots. The aim of this section is to show that it is an infinite mobile, that is, that the labels on the spine do not have a lower bound. Let us first describe it quickly.

The root of the $(\mathbf{T}, \mathbf{E}, \mathbf{L})$ has either type 1 and label 0, or type 2 and label 1/2, in which case it has (exceptionally) two children of type 4, one of them (uniformly selected) being on the spine. The vertices which are not on the spine have offspring distribution ζ, which was defined in Proposition 4.1 as the uniform ordering of μ, while vertices which are on the spine have offspring $\widehat{\zeta}$, defined by (4.4). The distribution $\widehat{\zeta}$ is itself the uniform ordering of a distribution $\widehat{\mu}$ on $(\mathbb{Z}_+)^4$ which we defined by

$$\widehat{\mu}^{(i)}(k_1, k_2, k_3, k_4) = \frac{k_1 b_1 + k_2 b_2 + k_3 b_3 + k_4 b_4}{b_i} \mu^{(i)}(k_1, k_2, k_3, k_4)$$

for $i \in [4]$ and $k_1, k_2, k_3, k_4 \in \mathbb{Z}_+$ and where b_1, b_2, b_3, b_4 are some positive numbers which depend on $\mathbf{q}$. The label displacement distribution $\nu^{(i)}_{\mathbf{w}}$ for a type $i \in [K]$ and a word $\mathbf{w}$ is then the uniform distribution on the set $D^{(i)}_{\mathbf{w}}$ which was defined in Proposition 4.1.

Lemma 5.7. *Let $i \in \{1, 2, 3, 4\}$ and $\mathbf{w} \in \mathcal{W}_4$ such that $\zeta^{(i)}_{\mathbf{w}} > 0$. Define the reversed word $\overleftarrow{\mathbf{w}} = (w_{|\mathbf{w}|}, \dots, w_1)$, and, for a label sequence $\mathbf{y} = (y_i)_{i \in [|\mathbf{w}|]}$, let $\overleftarrow{\mathbf{y}} = (-y_{|\mathbf{w}|}, -y_{|\mathbf{w}|-1}, \dots, -y_1)$. The function which maps $\mathbf{y}$ to $\overleftarrow{\mathbf{y}}$ is a bijection between $D^{(i)}_{\mathbf{w}}$ and $D^{(i)}_{\overleftarrow{\mathbf{w}}}$, sets which are defined in Proposition 4.1.*

As a corollary, we get that, if $\mathbf{W}$ has distribution $\widehat{\zeta}^{(i)}$ for some i and $\mathbf{Y}$ has distribution $\nu^{(i)}_{\mathbf{W}}$ conditionally on $\mathbf{W}$, then the pair $(\overleftarrow{\mathbf{W}}, \overleftarrow{\mathbf{Y}})$ has the same distribution as $(\mathbf{W}, \mathbf{Y})$.

Proof. If $i = 1$ or $i = 2$ then the result is immediate, since $\overleftarrow{\mathbf{w}} = \mathbf{w}$ and $D^{(i)}_{\mathbf{w}}$ only has one element.

If $i = 3$ or $i = 4$, bijectivity of the map comes from the fact that reversing a sequence (and eventually changing the signs of its elements) is an involutive operation, and thus we only need to check that $\overleftarrow{\mathbf{y}} \in D^{(i)}_{\overleftarrow{\mathbf{w}}}$ for any displacement list $\mathbf{y}$, which is straightforward given the definitions, since $(-y_{|\mathbf{w}|+1-(i+1)}) - (-y_{|\mathbf{w}|+1-i}) = y_{|\mathbf{w}|-i+1} - y_{|\mathbf{w}|-i}$ for $i \in \{0, \dots, |\mathbf{w}|\}$. □

Lemma 5.8. *Let, for $n \in \mathbb{Z}_+$, U_n be the $(n+1)$-th vertex of type 1 of the spine of $\mathbf{T}$. We then have*

$$\inf_{n \in \mathbb{N}} \mathbf{L}(U_n) = -\infty.$$

Proof. Note that U_n is well-defined for all $n \in \mathbb{Z}_+$, because the number of vertices of type 1 on the spine of $\mathbf{T}$ is a.s. infinite. Indeed, if it were not the case then all the vertices on the spine after a certain height would have type 2 or 4, but since a vertex of type 4 has positive probability of having at least one child of type 1, having an infinite sequence of vertices 2 and 4 has probability 0.

Notice then that $\big(\mathbf{L}(U_n)\big)_{n\in\mathbb{Z}_+}$ is in fact a centered random walk in $\mathbf{Z}$. It is a random walk because of the construction - the set of descendants of a vertex of type 1 of the spine will have distribution $\mathbb{P}_\zeta^{(1)}$. We can see that it is centered thanks to Lemma 5.7. Define the mirrored tree $(\overleftarrow{\mathbf{T}},\overleftarrow{\mathbf{E}},\overleftarrow{\mathbf{L}})$ by reversing the order of all the offspring of $\mathbf{T}$. To precise, if $u = u_1u_2\dots u_n \in \mathbf{T}$, then let, for $i \in [n]$, $v_i = k_{u_1\dots u_{i-1}} - i + 1$ and let then $\overleftarrow{u} = v_1\dots v_n$. Let then $\overleftarrow{\mathbf{E}}(\overleftarrow{u}) = \mathbf{E}(u)$ and, define the labels $\overleftarrow{\mathbf{L}}$ on $\overleftarrow{\mathbf{T}}$ by $\overleftarrow{\mathbf{L}}(\emptyset) = \mathbf{L}(\emptyset)$ and, for all u, $\mathbf{y}_{\overleftarrow{u}} = \overleftarrow{\mathbf{y}}_u$ (as defined in Lemma 5.7). Since, for $i \in [4]$ and $\mathbf{w} \in \mathcal{W}_4$ the distribution $\zeta^{(i)}$ is the uniform ordering of $\mu^{(i)}$ and $\nu_{\mathbf{w}}^{(i)}$ is uniform on $D_{\mathbf{w}}^{(i)}$, we obtain from Lemma 5.7 that $(\overleftarrow{\mathbf{T}},\overleftarrow{\mathbf{E}},\overleftarrow{\mathbf{L}})$ has the same distribution as $(\mathbf{T},\mathbf{E},\mathbf{L})$. In particular, $\mathbf{L}(U_1) - \mathbf{L}(U_0)$ has the same distribution as $\mathbf{L}(U_0) - \mathbf{L}(U_1)$, making its distribution centered. In particular, the centered random walk $\big(\mathbf{L}(U_n)\big)_{n\in\mathbb{Z}_+}$ then has no upper or lower bounds, for example by [50], Theorem 8.2.

□

5.2.4 Removing conditionings

The work done in the previous sections shows that maps with distribution $B_{\mathbf{q}}^{n,+}$ and $B_{\mathbf{q}}^{n,0}$ converge in distribution along the subsequence $(2+dn)_{n\in\mathbb{N}}$ (considering $B_{\mathbf{q}}^{n,0}$ only in the non-bipartite case). To show that maps with distribution $B_{\mathbf{q}}^{n}$ converge, all that is left for us to do is to show that the two quantities $B_{\mathbf{q}}^{n}(\mathcal{M}_n^+)$ and $B_{\mathbf{q}}^{n}(\mathcal{M}_n^0)$ converge (along the same subsequence). Since $2B_{\mathbf{q}}^{n}(\mathcal{M}_n^+) + B_{\mathbf{q}}^{n}(\mathcal{M}_n^0) = 1$, we can in fact restrict ourselves to showing that the quotient $\frac{B_{\mathbf{q}}^{n}(\mathcal{M}_n^+)}{B_{\mathbf{q}}^{n}(\mathcal{M}_n^0)}$ converges. Elementary calculations on conditionings give us

$$\begin{aligned}\frac{B_{\mathbf{q}}^{n}(\mathcal{M}_n^+)}{B_{\mathbf{q}}^{n}(\mathcal{M}_n^0)} &= \frac{B_{\mathbf{q}}\Big((M,E,R)\in\mathcal{M}^+ \mid (M,E,R)\in\mathcal{M}_n\Big)}{B_{\mathbf{q}}\Big((M,E,R)\in\mathcal{M}^0 \mid (M,E,R)\in\mathcal{M}_n\Big)} \\ &= \frac{B_{\mathbf{q}}^{+}\Big((M,E,R)\in\mathcal{M}_n\Big)}{B_{\mathbf{q}}^{0}\Big((M,E,R)\in\mathcal{M}_n\Big)}\frac{B_{\mathbf{q}}(\mathcal{M}^+)}{B_{\mathbf{q}}(\mathcal{M}^0)}.\end{aligned}$$

Recall that, in the BDFG bijection, the number of vertices of the map is exactly one more than the number of vertices of type 1 in the mobile. As a consequence, we have

$$\frac{B_{\mathbf{q}}^{+}\Big((M,E,R)\in\mathcal{M}_n\Big)}{B_{\mathbf{q}}^{0}\Big((M,E,R)\in\mathcal{M}_n\Big)} = \frac{\mathbb{P}_\zeta^{(1)}\big(\#_1\mathbf{T} = n-1\big)}{\mathbb{P}^{(2,2)}\big(\#_1\mathbf{F} = n-1\big)}.$$

We then deduce from $(H_{\mathbf{w}})$ and Lemma 5.4 that this quotient indeed converges as n converges to infinity, along the $(2+dn)_{n\in\mathbb{N}}$ subsequence.

6 Recurrence of the infinite map

The aim of this section if to prove the following:

Theorem 6.1. *The random rooted graph (M_∞, E_∞^+) is almost surely recurrent.*

Our principal tool for the proof will be the main result of [38]: since (M_∞, E^+_∞) is the limit in distribution of $\big((M_n, E^+_n), n \in \mathbb{N}\big)$, and since E^+_n is chosen according to the stationary distribution on M_n (that is, a vertex is chosen with probability proportional to its degree, i.e. its number of adjacent edges), then Theorem 1.1 of [38] states that if we can find positive constants λ and C such that, for all $n \in \mathbb{N}$,

$$\mathbb{P}(\deg(E^+) \geqslant n) \leqslant C\mathrm{e}^{-\lambda n},$$

then Theorem 6.1 will be proven.

Before showing this bound, we need a few simple results concerning variables which such exponential tails.

6.1 Around exponentially integrable variables

Let X be a nonnegative random variable. We say that X is *exponentially integrable with parameter* $\lambda > 0$ (which we will shorten as $EI(\lambda)$ from now on, and simply EI if we are not interested in the value of λ) if we have

$$\mathbb{E}[\mathrm{e}^{\lambda X}] < \infty.$$

The use of Markov's inequality shows that this implies that the tail of X is bounded by an exponential with parameter λ:

$$\forall n \in \mathbb{N}, \mathbb{P}(X \geqslant n) \leqslant \mathbb{E}[\mathrm{e}^{\lambda X}]\mathrm{e}^{-\lambda n}.$$

The converse is no quite true, but almost is: if the tail of X is bounded by an exponential with parameter λ, then X is $EI(\lambda')$ for $\lambda' < \lambda$.

If X and Y are two $EI(\lambda)$ variables then $X+Y$ will be $EI(\lambda')$ for $\lambda' < \frac{\lambda}{2}$, simply by bounding $\mathbb{P}(X+Y > n)$ by $\mathbb{P}(X > \frac{n}{2}) + \mathbb{P}(Y > \frac{n}{2})$. With an extra independence assumption, one can also do random sums:

Lemma 6.2. *Let $(X_i)_{i\in\mathbb{N}}$ be i.i.d nonnegative variables which are $EI(\lambda)$ for some $\lambda > 0$. Let N be an independent integer-valued variable which is $EI(\mu)$ for some $\mu > 0$. If $\mathbb{E}[\mathrm{e}^{\lambda X_1}] \leqslant \mathrm{e}^{\mu}$ (which is always possible by taking λ small enough), the the variable*

$$Y = \sum_{i=1}^{N} X_i$$

is also $EI(\lambda)$.

Proof. Conditioning on N and integrating with respect to all of the X_i, one immediately obtains

$$\mathbb{E}[\mathrm{e}^{\lambda Y}] = \mathbb{E}\Big[\mathbb{E}[e^{\lambda X_1}]^N\Big],$$

and this is enough. □

This could of course be generalized to the case where the (X_i) do not have the same distribution, but uniformally bounded exponential moments - we will not need such a generalization.

6.2 The case of positive maps

Picture a mobile $(\mathbf{T}, \mathbf{E}, \mathbf{L})$ with distribution $\widehat{\mathbb{P}}^{(1),(0)}_{\zeta,\nu}$: it has an infinite spine, and on its right and left sides are grafted some finite trees. Since the BDFG bijection makes $\emptyset$ into e^+, we will want to show that $\emptyset$ has an exponentially integrable number of successors and is the successor of an exponentially integrable number of vertices. We start with a simplified case.

Proposition 6.3. *Let* $\mathbf{A}$ *be a mobile with distribution* $\mathbb{P}^{(1),(0)}_{\zeta,\nu}$ *conditioned to the event where* $\emptyset$ *has exactly one child. Let* X *be the number of corners of* $\mathbf{A}$ *for which* $\emptyset$ *is the successor. Then* X *is an* $SE(\lambda)$ *variable for a certain* $\lambda > 0$.

Proof. Recall that X is the number of corners labelled 1 or $\frac{1}{2}$ met before encountering a vertex labelled 0 while circling counterclockwise around the tree $\mathbf{A}$. We will separately treat corners of types 1 and 2.

Let X_1 be the number of corners of type 1 encountered. We claim that, for all n,

$$\mathbb{P}(X_1 = n \mid X_1 \geqslant n) \geqslant \alpha(1 - \frac{1}{Z^+}), \tag{4.9}$$

where $\alpha > 0$ is the probability that, given a vertex of type 3 labelled 1, its rightmost offspring is of type 1 and has label 0. The fact that α is strictly positive comes from the fact that there exists $i \geqslant 3$ such that $q_i > 0$. In the case where such an i is different from 3, the vertex of type 3 can have offspring with at least one child of type 1, the uniform ordering of the offspring means that this child can be the rightmost one, and the distribution of the label displacements shows that it can have label 0. For the case where $q_3 > 0$ and $q_i = 0$ for $i \geqslant 4$, the type 3 vertex can have a unique child of type 2 with label $\frac{1}{2}$, which can have a unique child of type 4 which can have a unique child of type 1 with label 0.

Inequality (4.9) is obtained by recalling from Proposition 4.1 that the offspring of vertices of type 1 is only made of vertices of type 3, and that their number follows a geometric distribution with parameter $1 - \frac{1}{Z^+}$. Thus, whenever we visit a corner of a type 1 vertex with label 1, there is a $1 - \frac{1}{Z^+}$ chance that this vertex has another child. This immediately gives us (4.9), and a simple induction shows that X_1 is indeed a SE variable.

Let now X_2 be the number of *vertices* of type 2 with label $\frac{1}{2}$ encountered before the first vertex labelled. We insist that we count each vertex exactly once, when we meet them for the first time on the counter-clockwise exploration path. Then the same argument as for vertices of label 1 shows that $\mathbb{P}(X_2 = n \mid X_2 \geqslant n) \geqslant \alpha'$ for some strictly positive α', and X_2 is a SE variable.

Since $X \leqslant X_1 + 2X_2$, we now have our conclusion. □

The following lemma provides some additional on the structure $\mathbf{T}$.

Lemma 6.4. *Let* $n \in \mathbb{Z}_+$*, and let* V *be the* n*-th vertex of the spine of* $\mathbf{T}$ *to have type* 1*. Let also* N_r *and* N_l *be the numbers of subtrees rooted at* v *on the right and left sides of the spine. These variables are i.i.d. and their common distribution is geometric with parameter* $1 - \frac{1}{Z^+}$.

Proof. By combining Proposition 4.1 and Proposition 3.1, we obtain that the total offspring N of V follows a size-biased geometric distribution: we have

$$\mathbb{P}(N = k) = \frac{k}{(Z^+)^2}(1 - \frac{1}{Z^+})^{k-1}$$

for $k \geqslant 1$. Recall also that the child of V which is on the spine is chosen uniformly amongst the offspring of V. We thus have

$$\mathbb{P}(N_l = k, N_r = k') = \frac{\mathbb{P}(N = 1 + k + k')}{1 + k + k'} = (\frac{1}{Z^+})(1 - \frac{1}{Z^+})^{k+k'},$$

ending the proof. □

Proof of Theorem 6.1 for positive maps: First off, by Lemma 6.4, we know that $\emptyset$ has an EI amount of children, since geometric variables are EI, and therefore has an EI amount of successors. Next, look at all the subtrees of $\mathbf{T}$ which are rooted at $\emptyset$, excluding the subtree containing the spine. These are in EI amount, all independent, and, by Proposition 6.3, the root $\emptyset$ is connected to an EI amount of vertices in each of them. Lemma 6.2 allows to combine all of this: outside of the subtree containing the spine, $\emptyset$ is connected to an EI amount of vertices. Thus we now only need to prove a variation of Proposition 6.3 for this very subtree. This is done in the same way since, when doing the counterclockwise exploration process, the number of children of a vertex of type 1 on the spine is still geometric by Lemma 6.4, while vertices of type 2 only correspond to one corner. □

6.3 The case of null maps

The situation for null maps is slightly different, because the vertex E^+ is no longer the root of the mobile. Consider a mobile $(\mathbf{T}, \mathbf{E}, \mathbf{L})$ obtained by merging at their roots the two components of a forest with distribution $\widehat{\mathbb{P}}_{\zeta,\nu}^{(2,2),(\frac{1}{2},\frac{1}{2})}$, and let (M, E, R) be the map obtained after applying the BDFG bijection. Recall that E^+ is the first vertex of type 1 and label 0 encountered when running the clockwise countour process of $\mathbf{T}$. Note that it is either on the spine or on its left side.

An adaptation of the reasoning used in the previous section will work and give us that the number of vertices E^+ is connected to is indeed EI. First, for the number such vertices which are descendants of E^+ in $\mathbf{T}$, we find ourselves exactly back to the positive case: if E^+ is not on the spine then we apply Proposition 6.3 to an EI number of subtrees rooted at E^+, and if E^+ is on the spine, we separate the subtrees on the left side of the spine, on the right side of the spine and the subtree containing the spine. Secondly, we look for points of which E^+ is the successor, but which are not descendants of E^+. These can be obtained by running both the clockwise and counter-clockwise contour processes, starting at the root, and stopping them the first time we reach a 0 label. The same arguments as in the proof of Proposition 6.3 show that we encounter an EI number of vertices of labels 1 and $\frac{1}{2}$ on the way, thus ending the complete proof of Theorem 6.1. □

Bibliography

[1] Romain Abraham and Jean-François Delmas. Local limits of conditioned Galton-Watson trees: the infinite spine case. *Electron. J. Probab.*, 19:no. 2, 1–19, 2014.

[2] Romain Abraham, Jean-François Delmas, and Patrick Hoscheit. A note on the Gromov-Hausdorff-Prokhorov distance between (locally) compact metric measure spaces. *Electron. J. Probab.*, 18:no. 14, 1–21, 2013.

[3] David Aldous. The continuum random tree. I. *Ann. Probab.*, 19(1):1–28, 1991.

[4] David Aldous. The continuum random tree. II. An overview. In *Stochastic analysis (Durham, 1990)*, volume 167 of *London Math. Soc. Lecture Note Ser.*, pages 23–70. Cambridge Univ. Press, Cambridge, 1991.

[5] David Aldous. The continuum random tree III. *Ann. Probab*, 21(1):248–289, 1993.

[6] Jan Ambjørn, Thórdhur Jónsson, and Bergfinnur Durhuus. *Quantum geometry: a statistical field theory approach.* Cambridge University Press, 1997.

[7] Omer Angel and Oded Schramm. Uniform infinite planar triangulations. *Comm. Math. Phys*, 241(2-3):191–213, 2003.

[8] Emil Artin. *The gamma function.* Translated by Michael Butler. Athena Series: Selected Topics in Mathematics. Holt, Rinehart and Winston, New York, 1964.

[9] Jean Bertoin. Homogeneous fragmentation processes. *Probab. Theory Relat. Field.*, 121:301–318, 2001.

[10] Jean Bertoin. Self-similar fragmentations. *Ann. Inst. H. Poincaré Probab. Statist.*, 38(3):319–340, 2002.

[11] Jean Bertoin. *Random fragmentation and coagulation processes*, volume 102. Cambridge University Press, 2006.

[12] Jean Bertoin and Alexander Gnedin. Asymptotic laws for nonconservative self-similar fragmentations. *Electron. J. Probab*, 9:no.19, 575–593, 2004.

[13] Jean Bertoin and Alain Rouault. Discretization methods for homogeneous fragmentations. *J. London Math. Soc.*, 72(1):91–109, 2005.

[14] Jean Bertoin and Marc Yor. On subordinators, self-similar Markov processes and some factorizations of the exponential variable. *Electron. Commun. Probab.*, 6(10):95–106, 2001.

[15] Jean Bertoin and Marc Yor. On the entire moments of self-similar Markov processes and exponential functionals of Lévy processes. *Ann. Fac. Sci. Toulouse VI. Ser. Math.*, 11(1):33–45, 2002.

[16] Shankar Bhamidi. Universal techniques to analyze preferential attachment trees: Global and Local analysis, 2007. Prepublication.

[17] Patrick Billinglsey. *Convergence of probability measures.* Wiley Series in Probability and Statistics. Wiley, 2009.

[18] Jakob Björnberg and Sigurdur Stefánsson. Recurrence of bipartite planar maps. *Electron. J. Probab.*, 19:no. 31, 1–40, 2014.

[19] J. Bouttier, P. Di Francesco, and E. Guitter. Planar maps as labeled mobiles. *Electron. J. Combin.*, page 69.

[20] Philippe Carmona, Frédéric Petit, and Marc Yor. On the distribution and asymptotic results for exponential functionals of Lévy processes. *Rev. Mat. Iberoamericana*, pages 73–130, 1997.

[21] Philippe Chassaing and Bergfinnur Durhuus. Local limit of labeled trees and expected volume growth in a random quadrangulation. *Ann. Probab.*, 34(3):879–917, 2006.

[22] Bo Chen, Daniel Ford, and Matthias Winkel. A new family of Markov branching trees: the alpha-gamma model. *Electron. J. Probab.*, 14:no. 15, 400–430, 2009.

[23] Robert Cori and Bernard Vauquelin. Planar maps are well-labeled trees. *Canad. J. Math.*, (33):1023–1042, 1981.

[24] Nicolas Curien and Bénédicte Haas. The stable trees are nested. *Probab. Theory Related Fields*, 157(1):847–883, 2013.

[25] Nicolas Curien, Jean-François Le Gall, and Grégory Miermont. The brownian cactus i. scaling limits of discrete cactuses. *Ann. Inst. H. Poincaré Probab. Statist.*, 49(2):340–373, 2013.

[26] Nicolas Curien, Laurent Ménard, and Grégory Miermont. A view from infinity of the uniform planar quadrangulation. *ALEA, Lat. Am. J. Probab. Math. Stat.*, 10(1):45–88, 2013.

[27] C. Dellacherie and P.A. Meyer. *Probabilités et potentiel: Chapitres V à VIII: Théorie des martingales.* Actualités scientifiques et industrielles. Hermann, 1980.

[28] Richard Mansfield Dudley. *Real Analysis and Probability.* Cambridge University Press, 2002.

[29] Thomas Duquesne and Jean-François Le Gall. Random trees, Lévy processes and spatial branching processes. *Astérisque*, (281), 2002.

[30] Thomas Duquesne and Jean-François Le Gall. Probabilistic and fractal aspects of Lévy trees. *Probab. Theory Related Fields*, 131(4):553–603, 2005.

[31] Steven Evans. *Probability and real trees: École D'Été de Probabilités de Saint-Flour XXXV - 2005*, volume 1920 of *Lecture notes in mathematics.* Springer, 2008.

[32] Steven Evans, Jacob Pitman, and Amanda Winter. Rayleigh processes, real trees, and root growth with re-grafting. *Probab. Theory Related Fields*, 134(1):918–961, 2006.

[33] Steven Evans and Anita Winter. Subtree prune and regraft: a reversible tree-valued Markov process. *Ann. Probab.*, 34(3):81–126, 2006.

[34] Kenneth Falconer. *Fractal Geometry.* John Wiley & Sons, 1990.

[35] Daniell Ford. Probabilities on cladograms: introduction to the alpha model. 2005. `arXiv:math/0511246`.

[36] Kenji Fukaya. Collapsing of Riemannian manifolds and eigenvalues of Laplace operator. *Inventiones mathematicae*, 87:517–548, 1987.

[37] Andreas Greven, Peter Pfaffelhuber, and Anita Winter. Convergence in distribution of random metric measure spaces (lambda-coalescent measure trees). *Probab. Theory Related Fields*, 145(1-2):285–322, 2009.

[38] Ori Gurel-Gurevich and Asaf Nachmias. Recurrence of planar graph limits. *Annals of Mathematics*, 177(2):761–781, 2013.

[39] Bénédict Haas and Grégory Miermont. Self-similar scaling limits of non-increasing Markov chains. *Bernoulli Journal*, 17(4):1217–1247, 2011.

[40] Bénédicte Haas. Loss of mass in deterministic and random fragmentations. *Stoch. Proc. App.*, 106:411–438, 2003.

[41] Bénédicte Haas and Grégory Miermont. The genealogy of self-similar fragmentations with negative index as a continuum random tree. *Electron. J. Probab.*, 9:no.4, 57–97, 2004.

[42] Bénédicte Haas and Grégory Miermont. Scaling limits of Markov branching trees with applications to Galton-Watson and random unordered trees. *Ann. Probab.*, 40(6):2589–2666, 2012.

[43] Bénédicte Haas, Grégory Miermont, Jim Pitman, and Matthias Winkel. Continuum tree asymptotics of discrete fragmentations and applications to phylogenetic models. *Ann. Probab.*, 36(5):1790–1837, 2008.

[44] Bénédicte Haas, Jim Pitman, and Matthias Winkel. Spinal partitions and invariance under re-rooting of continuum random trees. *Ann. Probab.*, 37(4):1381–1411, 2009.

[45] Theodore Edward Harris. *The Theory of Branching Processes.* Springer, 1963.

[46] Chris Haulk and Jim Pitman. A representation of exchangeable hierarchies by sampling from random real trees. 2011. `arXiv:1101.5619`.

[47] John Hawkes. Trees generated by a simple branching process. *J. London Math. Soc.*, s2-24(2):373–384, 1981.

[48] J. Jacod and A.N. Shiryaev. *Limit Theorems for Stochastic Processes.* Grundlehren der Mathematischen Wissenschaften. Springer, second edition, 1987.

[49] Svante Janson. Simply generated trees, conditioned Galton-Watson trees, random allocations and condensation. *Probability Surveys*, 9:103–252, 2012.

[50] O. Kallenberg. *Foundations of Modern Probability.* Applied probability. Springer, 2002.

[51] Harry Kesten. Subdiffusive behavior of random walk on a random cluster. *Ann. Inst. H. Poincaré Probab. Statist.*, 22(4):425–487, 1986.

[52] John Frank Charles Kingman. The coalescent. *Stochastic Process. Appl*, 13:235–248, 1982.

[53] Maxim Krikun. Local structure of random quadrangulations. 2005. `arXiv:math/0512304`.

[54] Thomas Kurtz, Russell Lyons, Robin Pemantle, and Yuval Peres. A conceptual proof of the Kesten-Stigum theorem for multi-type branching processes. Athreya, Krishna B. (ed.) et al., Classical and modern branching processes. Proceedings of the IMA workshop, Minneapolis, MN, USA, June 13–17, 1994. New York, NY: Springer. IMA Vol. Math. Appl. 84, 181-185 (1997)., 1997.

[55] S.K. Lando, A.K. Zvonkin, and D. Zagier. *Graphs on Surfaces and Their Applications*. Encyclopaedia of Mathematical Sciences. Springer, 2004.

[56] Jean-François Le Gall. Uniqueness and universality of the Brownian map. *Ann. Probab.*, 41(4):2880–2960, 07 2013.

[57] Jean-Francois Le Gall and Yves Le Jan. Branching processes in Lévy processes: the exploration process. *Ann. Probab.*, 26(1):213–252, 1998.

[58] Jean-François Le Gall and Grégory Miermont. Scaling limits of random planar maps with large faces. *Ann. Probab.*, 39(1):1–69, 2011.

[59] Dominique Lépingle. La variation d'ordre p des semimartingales. *Z. Wahrscheinlichkeitstheorie verw. Gebiete*, 36:295–316, 1976.

[60] Russell Lyons and Yuval Peres. *Probability on trees and networks*. Cambridge University Press, in preparation, 2014.

[61] Philippe Marchal. Constructing a sequence of random walks strongly converging to Brownian motion. In *Discrete random walks (Paris, 2003)*, Discrete Math. Theor. Comput. Sci. Proc., AC, pages 181–190 (electronic). Assoc. Discrete Math. Theor. Comput. Sci., Nancy, 2003.

[62] Philippe Marchal. A note on the fragmentation of a stable tree. In *Fifth Colloquium on Mathematics and Computer Science*, Discrete Math. Theor. Comput. Sci. Proc., AI, pages 489–499. Assoc. Discrete Math. Theor. Comput. Sci., Nancy, 2008.

[63] Jean-François Marckert and Grégory Miermont. Invariance principles for random bipartite planar maps. *Ann. Probab.*, 35(5):1642–1705, 2007.

[64] Grégory Miermont. An invariance principles for random maps. *DMTCS Proceedings*, pages 39–58, 2006.

[65] Grégory Miermont. Invariance principles for spatial multitype Galton-Watson trees. *Ann. Inst. H. Poincaré Probab. Statist.*, 44(6):1128–1161, 2008.

[66] Grégory Miermont. Tessellations of random maps of arbitrary genus. *Ann. Sci. Éc. Norm. Supér.*, 42(5):725–781, 2009.

[67] Grégory Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Mathematica*, 210(2):319–401, 2013.

[68] Jacques Neveu. Arbres et processus de Galton-Watson. *Ann. Inst. H. Poincaré Probab. Statist.*, 22(2):199–207, 1986.

[69] J. Pitman. *Combinatorial stochastic processes*, volume 1875 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 2006. Lectures from the 32nd Summer School on Probability Theory held in Saint-Flour, July 7–24, 2002, With a foreword by Jean Picard.

[70] Jim Pitman and Marc Yor. The two-parameter Poisson-Dirichlet distribution derived from a stable subordinator. *Ann. Probab.*, 25(2):855–900, 1997.

[71] Jean-Luc Rémy. Un procédé itératif de dénombrement d'arbres binaires et son application à leur génération aléatoire. *RAIRO Inform. Théor.*, 19(2):179–195, 1985.

[72] D. Revuz and M. Yor. *Continuous Martingales and Brownian Motion.* Grundlehren der Mathematischen Wissenschaften. Springer, third edition, 2004.

[73] Anna Rudas, Bálint Tóth, and Benedek Valkó. Random trees and general branching processes. *Random Structures Algorithms*, 31(2):186–202, 2007.

[74] Gilles Schaeffer. *Conjugaison d'arbres et de cartes combinatoires aléatoires.* PhD thesis, Université Bordeaux I, 1998.

[75] Robert T. Smythe and Hosam M. Mahmoud. A survey of recursive trees. *Teor. Ĭmovīr. Mat. Stat.*, (51):1–29, 1994.

[76] F. Spitzer. *Principles of random walk.* Graduate texts in mathematics. Springer-Verlag, 2001.

[77] William Thomas Tutte. A census of planar maps. *Canad. J. Math*, 15:249–271, 1963.

Résumé

Nous nous intéressons à trois problèmes issus du monde des arbres aléatoires discrets et continus. Dans un premier lieu, nous faisons une étude générale des arbres de fragmentation auto-similaires, étendant certains résultats de Haas et Miermont en 2006, notamment en calculant leur dimension de Hausdorff sous des hypothèses malthusiennes.

Nous nous intéressons ensuite à une suite particulière d'arbres discrets k-aires, construite de manière récursive avec un algorithme similaire à celui de Rémy de 1985. La taille de l'arbre obtenu à la n-ième étape est de l'ordre de $n^{1/k}$, et après renormalisation, on trouve que la suite converge en probabilité vers un arbre de fragmentation. Nous étudions également des manières de plonger ces arbres les uns dans les autres quand k varie.

Dans une dernière partie, nous démontrons la convergence locale en loi d'arbres de Galton-Watson multi-types critiques quand on les conditionne à avoir un grand nombre de sommets d'un certain type fixé. Nous appliquons ensuite ce résultat aux cartes planaires aléatoire pour obtenir la convergence locale en loi de grandes cartes de loi de Boltzmann critique vers une carte planaire infinie.

Abstract

We study three problems related to discrete and continuous random trees. First, we do a general study of self-similar fragmentation trees, extending some results obtained by Haas and Miermont in 2006, in particular by computing the Hausdorff dimension of these trees under some Malthusian hypotheses.

We then work on a particular sequence of k-ary growing trees, defined recursively with a similar method to Rémy's algorithm from 1985. We show that the size of the tree obtained at the n-th step if of order $n^{1/k}$, and, after renormalization, we prove that the sequence convergences to a fragmentation tree. We also study embeddings of the limiting trees as k varies.

In the last chapter, we show the local convergence in distribution of critical multi-type Galton-Watson trees conditioned to have a large number of vertices of a fixed type. We then apply this result to the world of random planar maps, obtaining that large critical Boltzmann-distributed maps converge locally in distribution to an infinite planar map.