université PARIS-SACLAY | École doctorale de mathématiques Hadamard (EDMH)

UNIVERSITÉ PARIS SUD | FACULTÉ DES SCIENCES D'ORSAY — université PARIS-SACLAY

THÈSE DE DOCTORAT

de

L'UNIVERSITÉ PARIS-SACLAY

École doctorale de mathématiques Hadamard (EDMH, ED 574)

Établissement d'inscription : Université Paris-Sud

Laboratoire d'accueil : Laboratoire de mathématiques d'Orsay, UMR 8628 CNRS

Spécialité de doctorat : Mathématiques fondamentales

Thomas BUDZINSKI

Cartes aléatoires hyperboliques

Date de soutenance : 9 Novembre 2018

Après avis des rapporteurs : GRÉGORY MIERMONT (ÉNS Lyon)
OMER ANGEL (University of British Columbia)

Jury de soutenance :

MARIE ALBENQUE	(École Polytechnique)	Examinatrice
NICOLAS CURIEN	(Université Paris-Sud)	Directeur de thèse
JEAN-FRANÇOIS LE GALL	(Université Paris-Sud)	Examinateur
GRÉGORY MIERMONT	(ÉNS Lyon)	Rapporteur
ZHAN SHI	(UPMC)	Examinateur
ROMAIN TESSERA	(Université Paris-Diderot)	Examinateur
VINCENT VARGAS	(ÉNS Paris)	Examinateur

Fondation mathématique FMJH Jacques Hadamard — MATHEMATIQUES ORSAY — ENS ÉCOLE NORMALE SUPÉRIEURE

Remerciements

ou *Le bonus indispensable à toute bonne thèse.*

Je voudrais d'abord adresser ma plus profonde gratitude à Nicolas Curien pour le sujet de thèse riche et passionnant qu'il m'a proposé, et pour la liberté qu'il m'a laissée durant ces trois années. Sa disponibilité, sa bienveillance, son énergie inépuisable et son efficacité dans la re(re)lecture de tout ce que j'ai pu lui envoyer ont fait de lui un directeur de thèse idéal, auprès duquel j'espère encore beaucoup apprendre.

Je suis très reconnaissant envers Omer Angel et Grégory Miermont d'avoir accepté de rapporter cette thèse, et je les remercie pour l'intérêt qu'ils ont porté à mes travaux. Je remercie également Marie Albenque, Jean-François "Grand-chef" Le Gall, Zhan Shi, Romain Tessera et Vincent Vargas d'avoir accepté de faire partie de mon jury. Tous ont eu une influence considérable sur cette thèse, et je suis très honoré par leur présence.

J'ai eu la chance de passer ces trois dernières années entre deux laboratoires où j'ai bénéficié d'excellentes conditions de travail. Je remercie tous les chercheurs qui m'ont accueilli au DMA, grâce à qui le passage d'élève à caïman s'est fait tout naturellement. Merci à tous les occupants du bureau de l'ambiance et du bureau des plaisirs pour les déjeuners et les incontournables pauses baby-foot, et à Michel Pain avec qui c'est un plaisir de partager mon bureau depuis 2 ans. J'ai également beaucoup appris en enseignant à l'ENS, et je remercie tous les élèves avec qui j'ai pu échanger ces dernières années. Je n'oublie pas non plus le laboratoire de mathématiques d'Orsay, son cadre verdoyant et son merveilleux nouveau bâtiment, où j'ai pu profiter de l'excellente ambiance dans l'équipe de probabilités et entre doctorants. Je remercie également tous les collègues, carteux ou non, avec qui j'ai pu avoir des discussions mathématiques pendant ces années de thèse, dont Linxiao Chen, Olivier Hénard, Thomas Lehéricy, Baptiste Louf, Cyril Marzouk, Armand Riera, Delphin Sénizergues et Hugo Vanneuville. Enfin, je remercie profondément l'ensemble des personnels administratifs du DMA et de l'EDMH pour leur réactivité et leur efficacité, et en particulier Albane Trémeau, dont la bonne humeur contagieuse a illuminé le couloir vert, et Corentin Guénéron pour son aide plus que précieuse dans la paperasse de soutenance.

Je remercie également tous les enseignants qui m'ont transmis le goût des mathématiques. Je suis particulièrement reconnaissant envers Pascale Gillmann pour son soutien en Première. La préparation aux olympiades internationales fut une véritable révélation, et je remercie du fond du coeur tous ceux qui y ont contribué, dont Pierre Bornsztein, Xavier Caruso, Claude Deschamps, Theresia Eisenkoelbl, Bodo Lass et Johan Yebbou. J'ai eu la chance de profiter de cours de grande qualité en classes préparatoires grâce à Jérôme Dégot et Bernard Randé. Je remercie également tous les enseignants que j'ai eu en cours à l'ENS puis en M2, et dont je garde un excellent souvenir. Un merci particulier à Raphaël Cerf, qui a été mon tuteur pendant 3 ans à l'ENS, ainsi qu'à Zhan Shi, dont les deux cours limpides en L3 puis en M2 ont largement contribué à m'orienter vers les probabilités.

Au cours des dernières années, j'ai été impliqué dans un certain nombre d'activités mathématiques pour lycéens au sein de l'association Animath. Je voudrais saluer tous ceux avec qui

j'ai pu participer à un stage olympique, à un jury de TFJM ou à une séance de correction de copies pour l'énergie formidable avec laquelle ils organisent ces activités. Sans que la liste soit exhaustive, merci à Martin, Mathieu, Margaret, Guillaume, Colin, Raphaël, Nicolas, Cécile, Vincent, Igor, Matthieu, François, Giancarlo, Antoine, Eva et Jean-Louis.

Je n'oublie pas tous les amis, matheux ou non, avec qui j'ai pu passer de bons moments pendant ces trois ans. Merci à tous les anciens du C6 2013-2014 pour les soirées jeux de société et les voyages au ski ou en rando, à tous les pongistes de Ping Paris 14 (allez l'équipe 3 !), à mes copains de prépa pour les soirées à la Bodega et les LAN. Merci surtout à mes colocs Aymeric, François, Jacko et Delphin, sans qui ces années de thèse auraient été bien monotones.

Pour finir, je remercie du fond du coeur ma famille pour son soutien de tous les instants. Merci à Dominique pour m'avoir accueilli à Paris il y a bien des années, et merci à Nicolas et Julia, que j'ai toujours grand plaisir à retrouver. Enfin, merci à mes parents qui m'ont toujours encouragé.

Table des matières

Chapitre 1

Introduction

ou *Ils pèlent sur leurs graphes, et vont même explorer des hectomètres.*

Dans cette introduction, nous commençons par présenter la théorie des cartes planaires aléatoires, en mettant l'accent sur les différentes limites de grandes triangulations uniformes et sur leurs propriétés. Nous décrivons ensuite les outils les plus couramment utilisés pour étudier ces limites, et présentons une application à un modèle de triangulations aléatoires dynamiques [44]. Nous abordons ensuite les triangulations aléatoires hyperboliques, qui constituent l'objet d'étude central de cette thèse. Les articles [45] et [46] leur sont consacrés, ainsi que les Appendices A, B et C. Enfin, nous introduisons un autre modèle naturel de cartes aléatoires hyperboliques, les cartes causales surcritiques, qui fait l'objet de [47]. Les contributions originales de cette thèse sont les résultats encadrés.

Table des matières

Index des notations

La liste qui suit récapitule les notations les plus utilisées dans l'introduction de cette thèse. Ces notations sont pour la plupart cohérentes avec celles des articles qui suivent.

ρ : racine d'un graphe,
d_G : distance de graphe dans le graphe G,
γ : désigne une géodésique,
B_r : boule de rayon r,
$B_r^\bullet$: enveloppe de rayon r, vue comme un graphe ou une carte,
$\overline{B_r}$: enveloppe de rayon r, munie de la distance induite,
$\mathscr{T}_n$: ensemble des triangulations de la sphère à n sommets,
$T_n, T_n(\infty)$: triangulation aléatoire uniforme de la sphère à n sommets,
$T_\infty, \mathbb{T}_{\lambda_c}$: la triangulation uniforme infinie du plan, ou UIPT,
$\mathbb{H}_\alpha$: les triangulations markoviennes du demi-plan,
$\mathbb{T}_\lambda$: les triangulations markoviennes hyperboliques du plan, ou PSHIT,
λ_c : vaut $\frac{1}{12\sqrt{3}}$,
$\mathbf{m}$: la carte brownienne,
$\mathcal{P}$: le plan brownien,
$\mathcal{P}^h$: le plan brownien hyperbolique,
τ : désigne un arbre,
$\boldsymbol{\tau}$: désigne un arbre de Galton–Watson,
μ : loi de reproduction (une mesure sur $\mathbb{N}$),
GW_μ : loi d'un arbre de Galton–Watson de loi de reproduction μ,
$\boldsymbol{\tau}_\lambda^g$: arbre des géodésiques infinies les plus à gauche de $\mathbb{T}_\lambda$,
$\partial\tau$: ensemble des branches infinies d'un arbre τ,
$\widehat{\partial}\tau$: $\partial\tau$, où on a identifié les paires de branches voisines,
$\mathcal{C}(\tau)$: carte causale associée à l'arbre τ,
$\mathcal{M}(\tau, (s_i))$: carte obtenue en remplissant les faces de l'arbre τ par les bandes s_i.

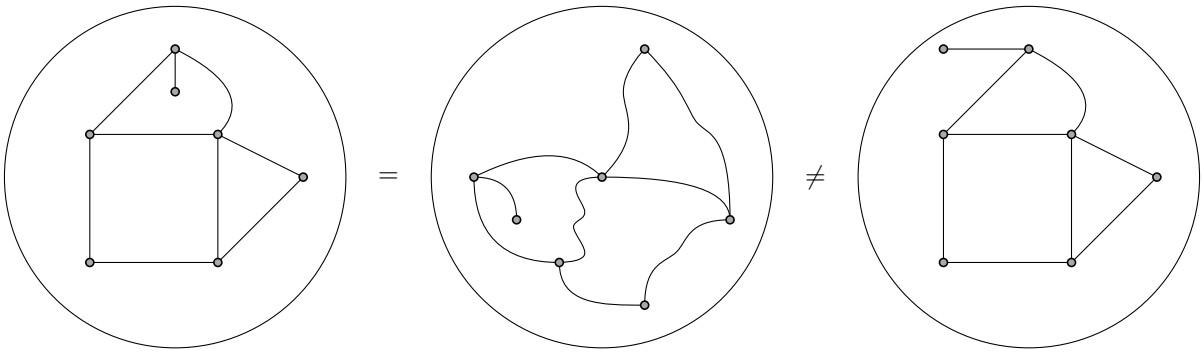

FIGURE 1.1 – Un exemple de carte planaire. La carte à gauche et celle au centre sont les mêmes. La carte à droite a le même graphe sous-jacent, mais c'est une carte différente. Par exemple, elle a une face de degré 8, tandis que les deux autres cartes n'ont que des faces de degré au plus 6.

1.1 Limites de cartes aléatoires

1.1.1 Définitions et combinatoire

Cartes planaires. On commence par définir les cartes planaires, qui sont l'objet d'étude principal de cette thèse. Rappelons qu'un graphe est une paire (V, E), où V est un ensemble de *sommets*, et E un ensemble d'*arêtes*, chaque arête ayant deux extrémités dans V. On autorise les arêtes multiples (plusieurs arêtes reliant la même paire de sommets) et les boucles (une arête reliant un sommet à lui-même).

Définition 1.1. Une *carte planaire finie* est un plongement propre d'un graphe fini et connexe dans la sphère $\mathbb{S}^2$, considéré à homéomorphisme préservant l'orientation près.

Par *plongement propre*, on entend que les images des arêtes ne s'intersectent pas, sauf éventuellement en une extrémité commune. De plus, on considère que deux cartes sont les mêmes s'il existe un homéomorphisme de $\mathbb{S}^2$ dans $\mathbb{S}^2$ préservant l'orientation qui envoie l'une sur l'autre (voir Figure 1.1). Les composantes connexes du complémentaire du plongement sont appelées *faces* de la carte, et le *degré* d'une face est le nombre d'arêtes qu'on doit longer pour faire le tour de la face (si on doit longer deux fois la même arête, elle est comptée deux fois).

Une définition combinatoire. Un inconvénient de la définition "topologique" donnée ci-dessus est qu'il n'est pas immédiat qu'il n'y a qu'un nombre fini de cartes avec, par exemple, un nombre fixé d'arêtes. Pour remédier à ce problème, on donne une autre définition, plus combinatoire, d'une carte planaire. Cette définition est équivalente à la première (voir Figure 1.2).

Définition 1.2. Une carte planaire est un ensemble de polygones dont les arêtes ont été collées deux à deux, de manière à obtenir topologiquement une sphère.

En particulier, étant donné un ensemble de polygones, il n'y a qu'un nombre fini de manières de recoller leurs arêtes deux à deux, donc il n'y a qu'un nombre fini de cartes dont les faces ont des degrés fixés.

Triangulations enracinées. On se restreindra la plupart du temps à certaines classes particulières de cartes planaires.

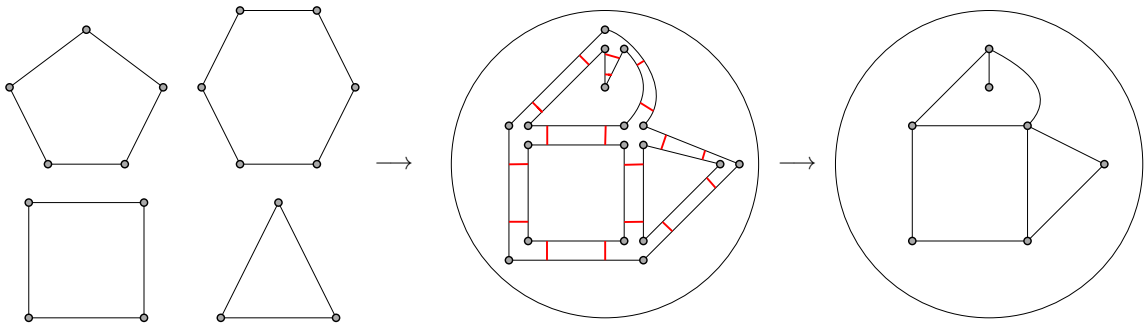

FIGURE 1.2 – Une carte planaire vue comme un recollement de polygones.

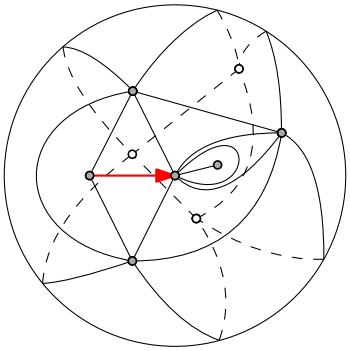

FIGURE 1.3 – Une triangulation de la sphère enracinée à 14 faces. L'arête racine est en rouge.

Définition 1.3. Soit $p \geq 3$. Une *p-angulation de la sphère* est une carte planaire dont toutes les faces sont de degré p.

En particulier, une carte dont toutes les faces sont de degré 3 est appelée *triangulation*, et une carte dont toutes les faces sont de degré 4 est appelée *quadrangulation*. Plus précisément, une *triangulation de type I* est une triangulation pouvant contenir des arêtes multiples et des boucles. Une *triangulation de type II* est une triangulation pouvant contenir des arêtes multiples, mais pas de boucle. Finalement, les cartes qu'on considère seront généralement *enracinées*, c'est-à-dire munies d'une arête orientée distinguée. Cette convention permet d'éliminer d'éventuels problèmes de symétrie.

Triangulations du p-gone et triangulations avec un trou. Une partie d'une triangulation de la sphère n'est pas une triangulation de la sphère. Pour cette raison, il est souvent nécessaire, pour décomposer des triangulations en parties plus petites, d'introduire une notion plus générale de triangulations à bord. Plus précisément, soit $p \geq 1$. Une *triangulation du p-gone* est une carte planaire enracinée où la face située à droite de l'arête racine est de degré p et a un bord simple (c'est-à-dire que son bord ne passe pas deux fois par le même sommet), et où toutes les autres faces sont de degré 3. Une *triangulation avec un trou de périmètre p* est une triangulation où toutes les faces sont de degré 3 sauf une, qui a un bord simple et est de degré p. Dans les deux cas, la face de degré p est appelée *face externe*.

La seule différence entre ces deux définitions réside dans la position de l'arête racine : dans la première, elle doit être adjacente à la face externe, tandis que dans la seconde, elle peut se trouver n'importe où. Les cartes avec un trou seront en général utilisées pour décrire un voisinage

de la racine dans une triangulation de la sphère, et les triangulations du p-gone pour décrire une partie ne contenant pas la racine.

Quelques résultats généraux. L'un des résultats les plus simples et les plus anciens concernant les cartes est la célèbre formule d'Euler : soit m une carte planaire, et soit $V(m)$ (resp. $E(m)$, $F(m)$) l'ensemble de ses sommets (resp. arêtes, faces). Alors on a

$$|V(m)| + |F(m)| - |E(m)| = 2.$$

En particulier, si m est une triangulation à n sommets, il est facile de montrer qu'elle a $2n-4$ faces et $3n-6$ arêtes. Si m est une triangulation d'un p-gone avec n sommets internes, elle contient $2n+p-2$ faces triangulaires (sans compter la face externe), et $3n+2p-3$ arêtes au total.

Par ailleurs, un résultat tout aussi simple à énoncer, mais beaucoup (beaucoup) plus difficile à prouver est le non moins célèbre théorème des quatre couleurs : il est toujours possible de colorier les sommets d'une carte planaire (sans boucle) avec quatre couleurs, de telle manière que deux sommets reliés par une arête sont toujours de couleurs différentes. Ce résultat a longtemps été la principale motivation pour l'étude des cartes planaires, et a finalement été démontré par Kenneth Appel et Wolfgang Haken en 1976.

Énumération. Avant d'intéresser les probabilistes, les cartes ont d'abord fait l'objet d'études combinatoires dans les années 60. Motivé par le théorème des quatre couleurs, William Tutte a obtenu de nombreux résultats d'énumération [132, 133, 134, 135, 136]. Par exemple, le nombre de quadrangulations de la sphère enracinées à n faces vaut

$$\frac{2\times 3^n}{(n+1)(n+2)}\binom{2n}{n} \underset{n\to+\infty}{\sim} \frac{2}{\sqrt{\pi}} 12^n n^{-5/2}, \tag{1.1}$$

tandis que le nombre de triangulations (de type I) de la sphère enracinées à n sommets vaut

$$2\times 4^{n-2}\frac{(3n-6)!!}{(n-1)!\,n!!} \underset{n\to+\infty}{\sim} \frac{1}{72\sqrt{6\pi}}\left(12\sqrt{3}\right)^n n^{-5/2}. \tag{1.2}$$

La seconde formule, sous cette forme, est due à Krikun [95].

Notons que la motivation initiale de Tutte était de démontrer le théorème des quatre couleurs en dénombrant d'une part les cartes planaires, et d'autre part les cartes planaires 4-coloriables. De plus, ces résultats ont été démontrés par des méthodes calculatoires : en utilisant une décomposition récursive des cartes, Tutte obtient une équation sur une fonction génératrice à deux variables, puis développe une nouvelle méthode, la *méthode quadratique*, pour résoudre cette équation.

D'autres méthodes ont depuis été développées pour dénombrer des cartes : des liens avec des intégrales de matrices ont été mis en évidence par des physiciens théoriciens dans les années 70 [131]. Enfin, dans les années 80, Robert Cori et Bernard Vauquelin on établi une bijection entre des arbres étiquetés et des cartes, qui donne une interprétation combinatoire de (1.1) (l'apparition des nombres de Catalan dans (1.1) rend naturel de chercher une telle interprétation). On reviendra sur cette bijection et ses multiples généralisations dans la Section 1.2.3.

Genre supérieur. Notons finalement que la notion de carte planaire peut se généraliser, en remplaçant la sphère par des surfaces de genre supérieur (tores à g trous). Il est alors important de demander dans la définition que chaque face soit homéomorphe à un disque. Les résultats combinatoires connus sont alors moins précis : des relations de récurrence à deux indices (la

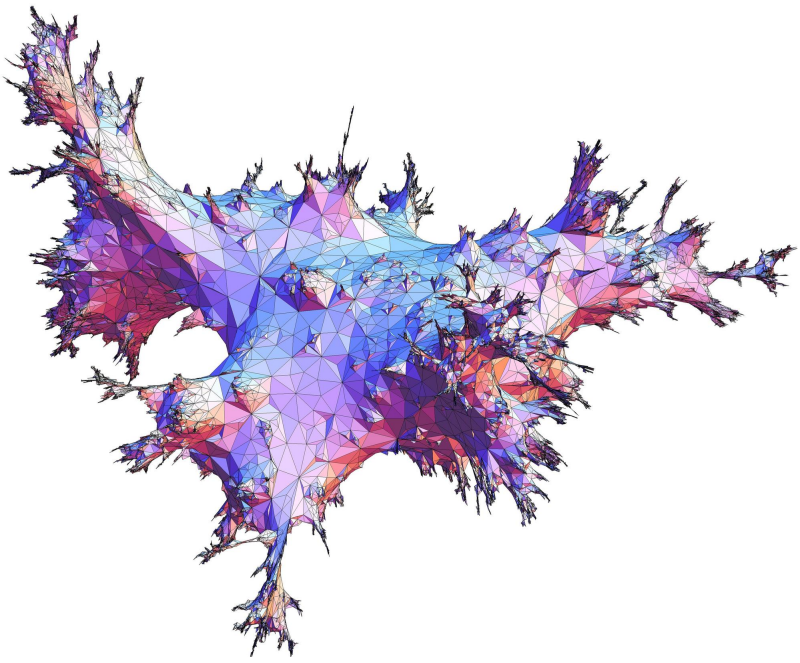

FIGURE 1.4 – Une triangulation uniforme de la sphère, avec environ 40000 sommets.

taille n et le genre g) sont connues [80], mais il ne semble pas possible d'en extraire des formules explicites. Les asymptotiques sont alors bien comprises quand $n \to +\infty$ à g fixé, mais pas quand n et g tendent vers l'infini simultanément.

1.1.2 Motivations

À quoi ressemble géométriquement une grande carte typique ? La principale question que se posent les probabilistes sur les cartes est la suivante : que peut-on dire des propriétés d'une grande carte typique ? Plus précisément, notons $\mathscr{T}_n$ l'ensemble des triangulations enracinées de la sphère à n sommets, et T_n une variable aléatoire uniforme sur $\mathscr{T}_n$. On peut faire de T_n un espace métrique en munissant l'ensemble de ses sommets de la distance de graphe : pour tous sommets u et v, on note $d_{T_n}(u, v)$ le nombre d'arêtes du plus court chemin dans T_n reliant u à v. Phillippe Chassaing et Gilles Schaeffer ont montré que les distances dans T_n sont typiquement de l'ordre de $n^{1/4}$ [54].

Surfaces aléatoires. Un des buts de la théorie des cartes aléatoires est de construire, par une approche discrète, des surfaces continues. En effet, les physiciens voient dans les cartes un modèle discret de gravité quantique bi-dimensionnelle, et un moyen de construire un analogue des intégrales de chemin de Feynman, mais qui porterait sur des surfaces (voir par exemple [9]).

Signalons également que d'autres approches mathématiques, purement continues, ont été développées pour construire ces objets, sous le nom de *gravité quantique de Liouville* [74]. Ces approches sont basées sur le *champ libre gaussien*, analogue naturel du mouvement brownien pour les fonctions de $\mathbb{R}^2$ dans $\mathbb{R}$. L'approche discrète et l'approche continue sont censées donner les mêmes objets. Les liens entre les deux ont fait l'objet de travaux récents par Jason Miller et Scott Sheffield [118, 119, 120], mais sont encore assez mal compris.

Relations KPZ. Une autre motivation, peut-être plus "terre-à-terre", pour s'intéresser aux cartes aléatoires, est l'étude de modèles de physique statistique, comme la percolation ou le modèle d'Ising. Comme on le verra un peu plus loin (Section 1.2.1), les cartes aléatoires possèdent une très agréable *propriété de Markov spatiale*, qui rend parfois l'étude de modèles de physique statistique plus facile sur des cartes aléatoires que sur des réseaux réguliers comme $\mathbb{Z}^2$. Certains exposants critiques peuvent ainsi être facilement obtenus sur des cartes aléatoires. De plus, d'après les physiciens, les exposants critiques sur des réseaux aléatoires et euclidiens sont liés par les *relations KPZ* [90], qui permettent donc en théorie de retrouver les exposants euclidiens. Pour cette raison, les cartes aléatoires couplées à des modèles de physique statistique ont elles aussi été beaucoup étudiées ces dernières années. Voir par exemple [39] pour le modèle $O(n)$ ou [128] pour la FK-percolation. Cependant, un grand nombre de questions dans cette direction restent encore ouvertes.

1.1.3 Limites locales

On rappelle que T_n est choisie uniformément parmi les triangulations de la sphère à n sommets. Une manière d'étudier les propriétés de T_n quand n est grand est de construire une "limite" de T_n quand n tend vers l'infini. Il existe deux manières naturelles de définir une telle limite : la *limite locale* (où T_n est "vue de près"), et la *limite d'échelle* (où T_n est "vue de loin"). On commence par décrire la limite locale. Intuitivement, si n est très grand et si un observateur se place à la racine de T_n, il ne voit plus que T_n vit sur la sphère, mais voit plutôt un plan autour de lui. La limite locale de T_n est donc une triangulation *infinie* du plan.

Distance locale. Pour parler de limite locale, on commence par définir la *topologie locale* sur l'ensemble des cartes. Cette topologie a été introduite par Itai Benjamini et Oded Schramm [31] (voir aussi l'article [8] de David Aldous et John Steele).

Soit m une carte planaire finie enracinée, et soit $r \geq 0$. La *boule de rayon* r de m, notée $B_r(m)$, est la carte formée par les sommets de m à distance de graphe au plus r de l'origine de l'arête racine, ainsi que les arêtes reliant ces sommets entre eux [1]. On définit alors la *distance locale* entre deux cartes m et m' par

$$d_{\mathrm{loc}}(m, m') = \left(1 + \max\{r \geq 0 | B_r(m) = B_r(m')\}\right)^{-1}.$$

Autrement dit, deux cartes sont proches si elles coïncident sur un grand voisinage de la racine. On peut alors vérifier que d_{loc} est une distance, et que le complété pour d_{loc} de l'ensemble des cartes finies est un espace polonais. Les éléments de ce complété qui ne sont pas des cartes finies sont appelés *cartes infinies*, et peuvent effectivement être vus comme des cartes possédant une infinité de sommets et d'arêtes (mais où tout sommet n'a qu'un nombre fini de voisins).

Enfin, on dit qu'une carte infinie m *n'a qu'un bout* si pour tout ensemble fini K de sommets de m, le complémentaire de K dans m n'a qu'une seule composante connexe infinie. C'est équivalent à ce qu'on puisse plonger m dans le plan de telle manière que chaque compact n'intersecte qu'un nombre fini de sommets, arêtes et faces.

L'UIPT. Le théorème suivant est le point de départ de la théorie des cartes aléatoires infinies. Il a été démontré par Omer Angel et Oded Schramm en 2003.

Théorème 1.1 (Angel–Schramm, [20])**.** Quand $n \to +\infty$, les cartes T_n convergent en loi pour la distance locale vers une triangulation infinie T_∞ à un seul bout.

1. Pour être tout à fait exact, la définition des boules qu'on utilise en pratique pour les triangulations est légèrement différente : c'est l'ensemble des faces adjacentes à au moins un sommet à distance au plus $r-1$ de la racine, ainsi que toutes leurs arêtes et tous leurs sommets.

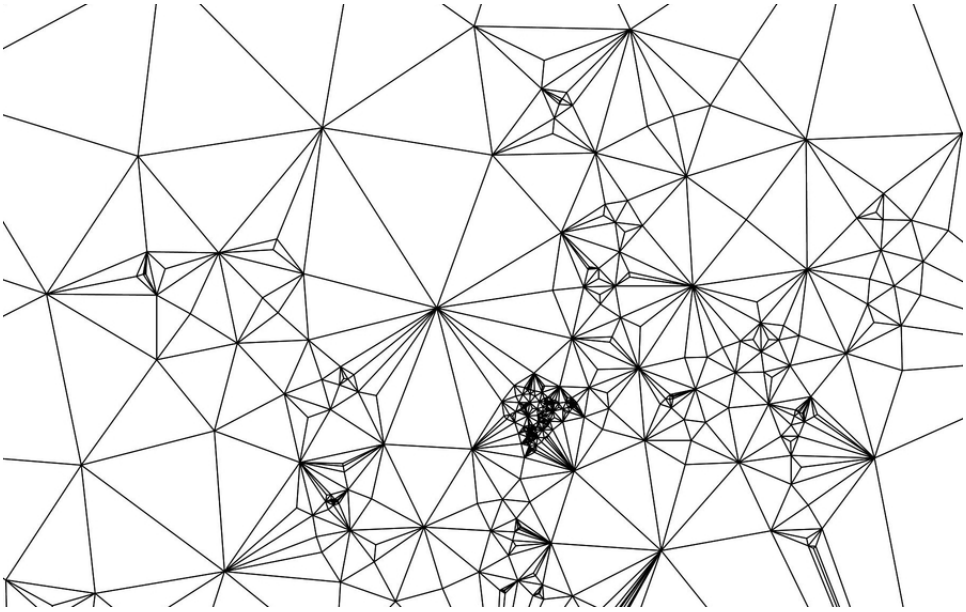

FIGURE 1.5 – Une simulation de l'UIPT. Le plongement a été calculé par le logiciel CirclePack de Ken Stephenson [3], puis on a effacé les cercles pour ne plus voir que les arêtes.

La triangulation T_∞ est appelée UIPT, pour *Uniform Infinite Planar Triangulation*. La preuve de ce résultat repose essentiellement sur les résultats d'énumération de Tutte cités plus haut. De plus, Angel et Schramm décrivent explicitement la loi de l'UIPT au sens où, pour toute triangulation t avec un trou de périmètre p, la probabilité d'observer t autour de la racine de T_∞ est calculée explicitement.

Propriétés de l'UIPT. De nombreuses propriétés de l'UIPT ont depuis été étudiées. Certaines de ces propriétés la rapprochent des réseaux euclidiens classiques, tandis que d'autres l'en éloignent fortement.

Commençons par un tour d'horizon des propriétés métriques connues de l'UIPT. On rappelle que la *boule de rayon* r d'une carte m, notée $B_r(m)$, est l'ensemble des sommets à distance de graphe au plus r de la racine. Notons que dans l'UIPT, les boules ne sont en général pas simplement connexes (voir Figure 1.6) : le complémentaire d'une boule a en fait de nombreuses composantes. On appelle *enveloppe* [4] *de rayon* r d'une carte infinie m l'union de $B_r(m)$ et de toutes les composantes connexes finies de son complémentaire (en violet sur la Figure 1.6). On la note $B_r^\bullet(m)$. Le fait que l'UIPT n'ait qu'un bout garantit que le complémentaire des enveloppes est bien connexe.

Omer Angel a montré [12] que les volumes de $B_r(m)$ et de $B_r^\bullet(m)$ pour r grand sont de l'ordre de r^4, à corrections logarithmiques près. De plus, le bord de l'enveloppe de rayon r (en bleu sur la Figure 1.6) est de longueur d'ordre r^2. En particulier, l'union des bords de ces enveloppes est d'ordre r^3 et non pas r^4, ce qui montre que la grande majorité des sommets se trouvent dans les "tentacules" de la Figure 1.6. Ces résultats ont depuis été précisés par Nicolas Curien et Jean-François Le Gall [64]. Par ailleurs, Maxim Krikun a montré que le plus petit cycle séparant la boule de rayon r de l'infini (en rouge sur la Figure 1.6) est de longueur d'ordre r [94], et Jean-François Le Gall et Thomas Lehéricy ont établi des inégalités isopérimétriques précises

3. Téléchargeable sur cette page : `http://www.circlepack.com/software.html`
4. *hull* en anglais

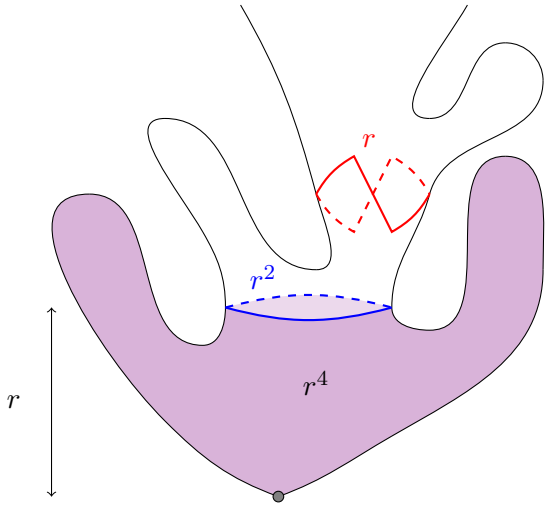

FIGURE 1.6 – L'UIPT représentée en "cactus" (la hauteur représente la distance à la racine).

[101]. Enfin, un phénomène de *coalescence des géodésiques* a été observé dans l'UIPQ (analogue de l'UIPT pour les quadrangulations) par Nicolas Curien, Laurent Ménard et Grégory Miermont [68], puis dans l'UIPT [67]. Cela signifie par exemple que deux géodésiques infinies ont toujours une infinité de points d'intersection.

Pour les raisons évoquées dans la Section 1.1.2, la percolation sur T_∞ a elle aussi été étudiée : Omer Angel a montré [12] que la probabilité critique pour la percolation par sites vaut $\frac{1}{2}$. Notons que c'est l'un des premiers résultats sur l'UIPT qui a été montré après son introduction, alors que la même question pour le réseau triangulaire est restée ouverte pendant 20 ans ! Plusieurs exposants critiques ont également été obtenus par Matthias Gorny, Édouard Maurel-Ségala et Arvind Singh [78]. La géométrie des clusters de percolation critique est également assez bien comprise, grâce notamment à des travaux de Nicolas Curien et Igor Kortchemski [61], et d'Olivier Bernardi, Nicolas Curien et Grégory Miermont [33]. À chaque fois, les méthodes utilisées sont complètement différentes de celles utilisées dans les réseaux euclidiens, et les preuves sont souvent plus faciles.

Enfin, une autre manière d'étudier la géométrie de T_∞ est de s'intéresser au comportement de la marche aléatoire simple (X_n) sur ce graphe : un marcheur démarre du sommet racine et, à chaque instant, saute sur un sommet choisi uniformément parmi les voisins de celui où il se trouve. Ori Gurel-Gurevich et Asaf Nachmias ont montré [82] que l'UIPT était récurrente (le marcheur revisite une infinité de fois la racine), tranchant une question qui était restée ouverte pendant une dizaine d'année. Leur résultat montre en fait que toute limite locale de graphe uniformément enracinés est récurrente, à condition que le degré de la racine ait une queue exponentielle. Itai Benjamini et Nicolas Curien ont également montré [28] que la marche aléatoire est *sous-diffusive* : la distance entre X_n et la racine est au plus d'ordre $n^{1/3}$. Cet exposant $\frac{1}{3}$ a été légèrement amélioré par Nicolas Curien et Cyril Marzouk [66]. L'exposant exact conjecturé est $\frac{1}{4}$, et fait l'objet d'un travail en cours d'Ewain Gwynne et Tom Hutchcroft.

Quelques variantes. À la suite de l'UIPT, plusieurs autres modèles de cartes infinies ont été introduits et étudiés, tels l'UIPQ, analogue pour les quadrangulations, introduite par Maxim Krikun [92]. Ses propriétés à grande échelle sont généralement les mêmes que celles de l'UIPT, mais les outils utilisés pour étudier les deux peuvent différer. Des modèles plus généraux de cartes infinies (cartes de Boltzmann infinies où, pour tout $p \geq 1$, les faces de degré p apparaissent avec

une fréquence fixée) ont été construits par Jakob Björnberg et Sigurdur Örn Stefànsson dans le cas biparti [38] (où toutes les faces sont de degré pair), puis par Robin Stephenson dans le cas général [129]. Tous ces modèles peuvent être vus comme les limites locales de modèles de cartes aléatoires finies. Enfin, tous ces modèles ont des analogues dans le demi-plan, qui ont également été étudiés (voir [11] pour les triangulations, [69] pour les quadrangulations, et [42, 57] pour les cartes de Boltzmann biparties générales). Ces modèles sont les limites locales de cartes aléatoires à bord, vues depuis un sommet du bord (par exemple de triangulations uniformes du p-gone avec n sommets internes, où on a fait tendre successivement n puis p vers l'infini).

1.1.4 Limites d'échelle

Une autre manière de définir une limite de cartes est de regarder la carte "de loin" : de même que la marche aléatoire sur $\mathbb{Z}$, une fois renormalisée, converge vers un mouvement brownien, on espère obtenir un objet compact, mais continu : c'est une *limite d'échelle*.

Distance de Gromov–Hausdorff. Les cartes seront vues comme des espaces métriques discrets. Il est donc d'abord nécessaire de préciser ce que signifie la convergence d'espaces métriques. On commence par définir la distance de Hausdorff sur l'ensemble des parties compactes d'un espace métrique (X, d). Soient $K_1, K_2 \subset X$ compacts. On pose

$$\mathbf{d}_H(K_1, K_2) = \max\left(\max_{x_1 \in K_1} d(x_1, K_2), \max_{x_2 \in K_2} d(x_2, K_1)\right).$$

Si (X_1, d_1) et (X_2, d_2) sont deux espaces métriques compacts, on définit leur *distance de Gromov–Hausdorff* par

$$\mathbf{d}_{GH}(X_1, X_2) = \inf_{\varphi_1, \varphi_2} \mathbf{d}_H\left(\varphi_1(X_1), \varphi_2(X_2)\right),$$

où l'infimum porte sur toutes les paires de plongements isométriques φ_1 et φ_2 de X_1 et X_2 dans un même espace métrique X. On peut vérifier que $\mathbf{d}_{GH}$ définit une distance sur l'ensemble des (classes d'isométrie d') espaces métriques compacts, et en fait un espace polonais (voir [48] pour plus de détails). Signalons également qu'il existe une variante permettant de prendre également en compte une mesure borélienne sur les espaces métriques. C'est la *distance de Gromov–Hausdorff–Prokhorov*, dont on ne détaillera pas la définition ici (voir par exemple [2]).

Signalons au passage que les chapitres 3 et 4 de cette thèse contiennent des résultats techniques sur ces topologies, qui pourraient être utiles dans d'autres contextes. Nous montrons dans l'appendice du chapitre 3 que, si $r > 0$ et si une suite d'espaces pointés (X_n, ρ_n) converge [5] vers (X, ρ), alors sous certaines hypothèses, l'enveloppe de rayon r et de centre ρ_n dans X_n converge vers l'enveloppe de rayon r et de centre ρ dans X. Dans l'appendice du chapitre 4, nous montrons la fermeture pour la distance de Gromov–Hausdorff de nombreuses propriétés liées aux géodésiques des espaces métriques.

La carte brownienne. On rappelle que T_n est une triangulation uniforme de la sphère à n sommets. La question de la convergence de T_n vers un objet continu a été posée pour la première fois par Oded Schramm en 2003. Les premiers résultats dans cette direction sont ceux de Philippe Chassaing et Gilles Schaeffer [54], qui montrent que les distances entre les sommets sont de l'ordre de $n^{1/4}$. Plus précisément, le profil des distances à la racines, renormalisées par $n^{1/4}$, converge vers une certaine mesure aléatoire. La tension pour la distance de Gromov–Hausdorff de T_n (en divisant les distances par $n^{1/4}$) a ensuite été montrée par Jean-François Le Gall [97], puis

5. Il faut pour cela étendre la distance de Gromov–Hausdorff aux espace pointés ou bipointés, voir Chapitre 3.

l'unicité de la limite indépendamment par Grégory Miermont (pour les quadrangulations, [116]) et Jean-François Le Gall [99].

Théorème 1.2 (Le Gall, [99])**.** On a la convergence en loi suivante pour la distance de Gromov–Hausdorff–Prokhorov :

$$\left(T_n, \frac{3^{1/4}}{n^{1/4}} d_{T_n}\right) \xrightarrow[n\to+\infty]{(loi)} \mathbf{m},$$

où $\mathbf{m}$ est un espace métrique mesuré compact aléatoire appelé *carte brownienne*, ou *sphère brownienne*.

Notons que [99] démontre que le résultat reste vrai en remplaçant T_n par une p-angulation uniforme pour p pair. L'objet limite reste le même, seule la constante $3^{1/4}$ diffère alors. Un certain nombre de résultats d'*universalité*, démontrant la convergence de nombreux modèles de cartes aléatoires vers la carte brownienne, ont depuis été établis (voir par exemple [112]). La Figure 1.4, bien que discrète, donne une idée de ce à quoi ressemble cet objet.

Propriétés de la carte brownienne. La géométrie de la carte brownienne a également été étudiée en détail : on sait qu'elle est presque sûrement homéomorphe à la sphère [103, 115], mais de dimension de Hausdorff 4 [97]. La structure des géodésiques dans la carte brownienne a également été étudiée [16, 98]. Ces géodésiques possèdent des propriétés de coalescence : presque sûrement, deux géodésiques issues de la racine vers deux points différents coïncident sur un segment initial.

Le plan brownien. Une variante non-compacte de la carte brownienne, appelée *plan brownien*, a été introduite par Nicolas Curien et Jean-François Le Gall [62]. Le plan brownien $\mathcal{P}$ possède les mêmes propriétés locales que la carte brownienne, mais est homéomorphe au plan. Il peut être vu à la fois comme l'espace obtenu en "zoomant à l'infini" vers la racine de la carte brownienne, ou comme la limite d'échelle de l'UIPT [6].

Contrairement à la carte brownienne, le plan brownien a l'avantage d'être invariant par changement d'échelle : pour tout $\alpha > 0$, la loi de $\mathcal{P}$ est la même que celle de $\alpha\mathcal{P}$, espace obtenu en multipliant par α les distances et par α^4 la mesure. Grâce à cette invariance, certains calculs sont plus faciles à effectuer sur $\mathcal{P}$ que sur la carte brownienne. Par exemple, Nicolas Curien et Jean-François Le Gall ont donné un sens au périmètre P_r et au volume V_r de l'enveloppe de rayon r dans le plan brownien [7]. Ils ont alors caractérisé complètement la loi du processus $(P_r, V_r)_{r\geq 0}$ [63]. En particulier, le processus (P_r) a la loi d'un processus de branchement (analogue continu d'un processus de Galton–Watson) avec un mécanisme de branchement stable d'indice 3/2, où le temps a été "retourné". Les processus (P_r) et (V_r) jouent un rôle important dans le chapitre 3.

Quelques autres variantes. Des surfaces browniennes avec d'autres topologies ont également été construites. Les disques browniens [37] et le demi-plan brownien [22] sont ainsi des limites d'échelle de cartes à bord. Citons également les surfaces browniennes de genre supérieur [35]. Enfin, signalons l'existence d'autres modèles continus appelés *cartes stables* [102], plus fondamentalement différents de la carte brownienne. Elles sont censées être les limites d'échelles de cartes où les degrés des faces ont une queue lourde, et où des faces macroscopiques apparaissent.

6. Pour être exact, Curien et Le Gall montrent que $\mathcal{P}$ est la limite d'échelle de l'UIPQ. L'adaptation de la preuve pour l'UIPT figure dans le chapitre 3.

7. C'est loin d'être évident pour le périmètre car le bord d'une enveloppe est une courbe de dimension fractale 2.

Leur géométrie reste cependant assez mal comprise (voir par exemple [125] pour une étude du bord des grandes faces)

1.2 Boîte à outils

Le but de cette section est de présenter la "boîte à outil" de la théorie des cartes aléatoires. La plupart des résultats cités dans la section précédente ont été montrés en utilisant ces outils [8].

1.2.1 Épluchage

L'épluchage, ou *peeling*, est une méthode d'exploration des cartes introduite par le physicien Yoshiyuki Watabiki [139], et développée rigoureusement pour l'UIPT par Omer Angel [12].

Propriété de Markov spatiale. Le point de départ de cette idée est la *propriété de Markov spatiale* de l'UIPT T_∞ : si t est une triangulation finie avec un trou de périmètre p, conditionnellement à $t \subset T_\infty$, la loi de $T_\infty \backslash t$ ne dépend que de p. Cette propriété est un analogue de la propriété de Markov "temporelle" usuelle au sens où le passé correspond à la triangulation t, le présent au périmètre p (c'est-à-dire au bord du passé), et le futur à la partie restante $T\backslash t$. Une propriété similaire apparaît également dans d'autres objets aléatoires planaires, comme les processus SLE ou le champ libre gaussien.

Épluchage. L'idée est alors d'explorer T_∞ par "petits morceaux", en utilisant de manière répétée la propriété de Markov spatiale. Plus précisément, soit t une triangulation finie avec un trou de périmètre p, et supposons qu'on sache que $t \subset T_\infty$. On choisit une arête e sur le bord de t et on note f la face adjacente à e à l'extérieur de t. On peut alors distinguer trois cas décrivant la forme de f, comme sur la Figure 1.7 :

(a) le troisième sommet de f n'est pas sur le bord de t,
(b) la face f sépare de l'infini un segment de ∂t de longueur i situé sur sa droite, avec $0 \leq i \leq p-1$,
(c) la face f sépare de l'infini un segment de ∂t de longueur i situé sur sa gauche, avec $0 \leq i \leq p-1$.

D'après la propriété de Markov spatiale, la probabilité de chacun de ces cas conditionnellement à $t \subset T_\infty$ ne dépend que de p.

Le processus d'épluchage consiste à découvrir progressivement T_∞ de la manière suivante : on démarre avec une simple face et, à chaque étape, on choisit une arête e sur le bord de la partie explorée. On découvre alors la face qui se trouve de l'autre côté de e. On dit qu'on a *épluché* l'arête e. Si cette face sépare une zone finie de l'infini (en vert sur la Figure 1.7), on découvre également cette zone. On note $\mathfrak{peel}_n(T_\infty)$ la triangulation finie découverte après n étapes. La propriété de Markov spatiale implique que le périmètre de $\mathfrak{peel}_n(T_\infty)$ est une chaîne de Markov à valeurs dans $\mathbb{N}$, dont on peut calculer explicitement les transitions. De plus, chaque fois qu'on est dans un des cas (b) et (c) de la Figure 1.7, la loi de la triangulation qui remplit la zone verte est connue explicitement.

Applications. La puissance de cette technique réside dans le fait qu'on peut choisir à chaque étape l'arête à éplucher [9]. Par conséquent, il est possible d'adapter l'algorithme d'épluchage (c'est-à-dire la manière de choisir e sur le bord à chaque étape) au type de propriétés de T_∞

8. Signalons dès maintenant une exception majeure : la preuve de la récurrence de l'UIPT [82] repose sur la très jolie théorie des empilements de cercles, qu'on n'abordera pas ici. Voir par exemple [126] pour un aperçu.

9. Du moment que ce choix ne dépend pas de la partie inexplorée de T_∞.

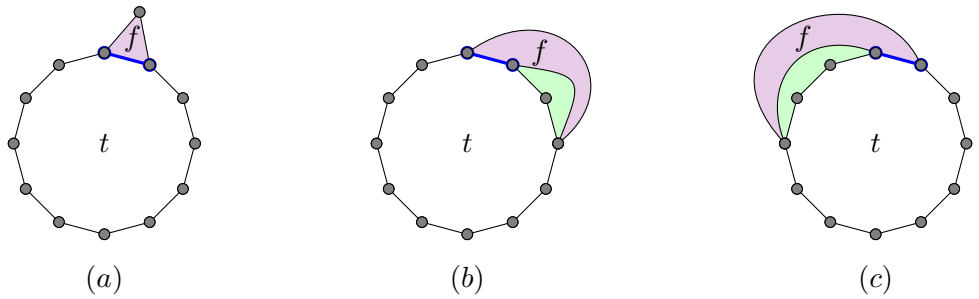

FIGURE 1.7 – Les trois formes possibles de la face f (avec $i = 2$ dans le cas (b), et $i = 3$ dans le cas (c)).

auxquelles on s'intéresse. L'épluchage a donc été utilisé pour étudier les propriétés de croissance [12, 64], la percolation [12, 61], la marche aléatoire simple [28] ou encore la structure conforme [56] de T_∞. Signalons également un joli résultat de Nicolas Curien et Jean-François Le Gall [64, Corollaire 6] : quel que soit l'algorithme d'épluchage choisi, presque sûrement, toute l'UIPT sera découverte (c'est-à-dire que $\bigcup_{n\geq 0} \mathfrak{peel}_n(T_\infty) = T_\infty$ p.s.).

Robustesse. La version du peeling "à la Angel" présentée ici peut également être utilisée pour explorer les quadrangulations, mais le nombre de cas à traiter explose rapidement avec les degrés des faces. Pour traiter des modèles de cartes beaucoup plus généraux, une variante plus robuste a été développée par Timothy Budd [42]. Cette variante permet d'étendre certains résultats connus pour l'UIPT à de nombreux autres modèles, voir [57] pour de nombreuses applications. Enfin, signalons que l'épluchage est particulièrement simple pour les cartes du demi-plan, puisque le bord de la partie restant à explorer est toujours le même (il est toujours infini), ce qui permet de se débarasser de la dépendance en p des probabilités de transitions.

1.2.2 Décomposition de Krikun

Un autre outil très utile pour étudier l'UIPT est la *décomposition de Krikun*, dont l'idée est de coder T_∞ par un arbre à l'envers, qu'on appelle son *squelette*, et où les hauteurs des sommets dans l'arbre correspondent à leur distance à la racine dans T_∞. Cette décomposition a été introduite par Maxim Krikun pour les triangulations de type II [94], et décrite par Nicolas Curien et Jean-François Le Gall pour les triangulations de type I [65].

Squelette d'une triangulation infinie. Plus précisément, considérons une triangulation infinie du plan T et, pour tout $r \geq 0$, traçons le bord $\partial B_r^\bullet$ de l'enveloppe de rayon r (en bleu sur la Figure 1.6). Les faces à l'intérieur de $B_r^\bullet$ adjacentes à des arêtes de $\partial B_r^\bullet$ sont appelées *triangles pointant vers le bas à hauteur* r.

La construction du squelette de Krikun est alors la suivante : on place un sommet au milieu de chaque arête de $\partial B_r^\bullet$, pour tout $r \geq 0$. Ces sommets seront ceux du squelette. De plus, soit e une arête d'un cycle $\partial B_r^\bullet$. On suit $\partial B_r^\bullet$ vers la droite en partant du milieu de e, jusqu'à rencontrer un triangle pointant vers le bas à hauteur $r + 1$. Alors le *parent* de e est l'arête de $\partial B_{r+1}^\bullet$ adjacente à ce triangle. Le squelette $\mathrm{Skel}(T)$ de T est alors la forêt qui décrit cette généalogie (voir Figure 1.8).

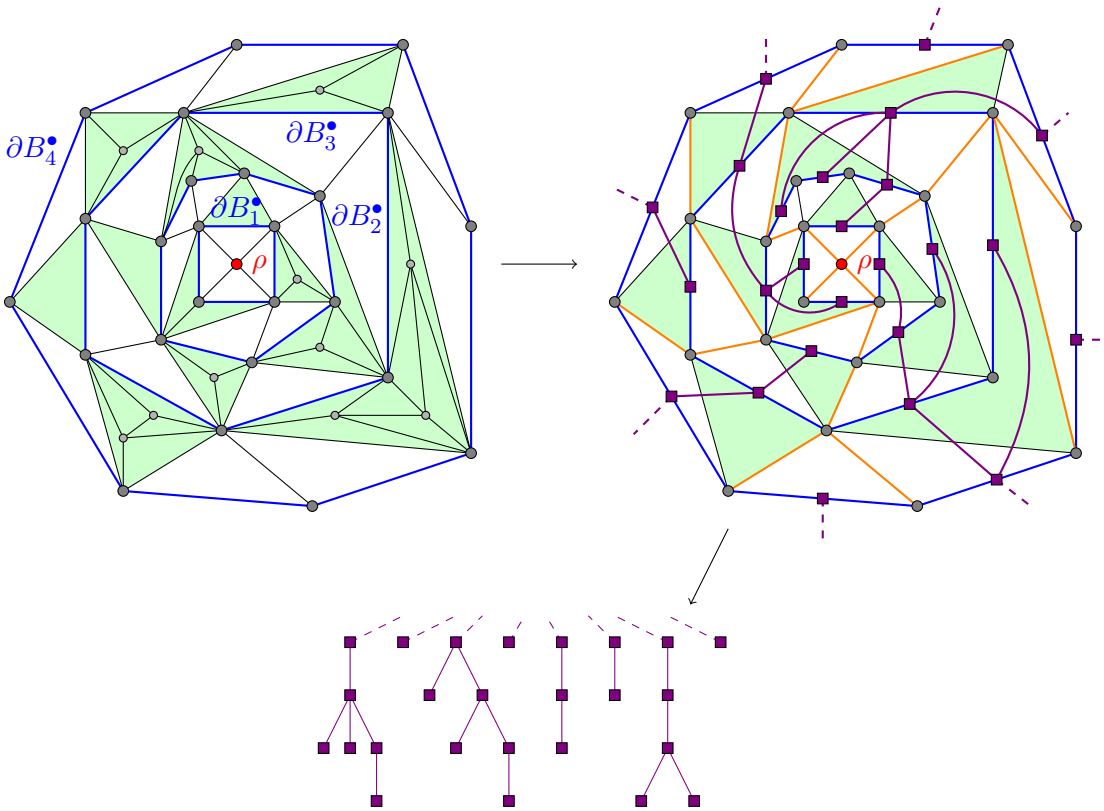

FIGURE 1.8 – La décomposition de Krikun de T_∞. En bleu, les cycles $\partial B_r^\bullet$ pour $r = 1, 2, 3, 4$. En blanc, les triangles pointant vers le bas. En violet, les branches du squelette qui descendent à hauteur 3 ou moins. En orange, les géodésiques les plus à gauche reliant les sommets de $\partial B_4^\bullet$ à la racine.

Décomposition de Krikun de T_∞. La loi du squelette de T_∞ peut être calculée explicitement. Cette forêt ne contient presque sûrement qu'un seul arbre. De plus, sa loi est essentiellement celle d'un arbre de Galton–Watson à l'envers, dont la racine serait à l'infini, et conditionné à s'éteindre exactement à la hauteur 0. En particulier, $(|\partial B_r^\bullet|)_{r\geq 0}$ peut être vu comme un processus de Galton–Watson à l'envers, ce qui permet d'obtenir des informations précises sur ce processus (voir [94, 122]). De plus, conditionnellement au squelette, les "trous" (en vert sur la Figure 1.8) sont remplis par des triangulations indépendantes, de lois connues explicitement.

Enfin, cette décomposition code des informations sur les géodésiques issues de la racine : les chemins "zigzaguant" entre les branches du squelette (en orange sur la Figure 1.8) sont les géodésiques les plus à gauche reliant les sommets de $\partial B_r^\bullet$ à la racine ρ. Le fait que $\mathrm{Skel}(T_\infty)$ soit connexe implique par exemple qu'il n'existe pas deux géodésiques infinies oranges disjointes. La décomposition de Krikun permet donc d'obtenir des propriétés de *confluence des géodésiques* (voir [67], ainsi que [93] sur les quadrangulations).

1.2.3 Bijection de Cori–Vauquelin–Schaeffer

Terminons l'inventaire de cette "boîte à outils" par une rapide description de la bijection de Cori–Vauquelin–Schaeffer. Même si cette bijection ne sera pas explicitement utilisés dans les travaux qui suivent, elle est le point de départ de tous les résultats connus actuellement sur les limites d'échelle de cartes aléatoires. Il s'agit d'une bijection entre d'une part les quadrangulations de la sphère à n faces, et d'autre part les arbres étiquetés à $n+1$ sommets.

Arbres bien étiquetés. Un *arbre plan* est une carte planaire (finie ou infinie) sans cycle ou, si l'on préfère, un arbre enraciné muni, pour chaque sommet v, d'un ordre cyclique sur l'ensemble des voisins de v. Si τ est un arbre plan à $n+1$ sommets, on peut faire le tour de τ dans le sens horaire, en longeant ses arêtes et en commençant par l'arête racine. On note alors $(v_0, v_1, \ldots, v_{2n})$ la liste des sommets rencontrés, de sorte que $v_0 = v_{2n}$ et v_1 sont respectivement les sommets de départ et d'arrivée de l'arête racine. Un même sommet peut apparaître plusieurs fois dans la liste (voir la partie gauche de la Figure 1.9). La suite $(v_0, v_1, \ldots, v_{2n})$ peut être étendue à $\mathbb{Z}$ par $2n$-périodicité, et est appelée *liste de contour* de τ.

Un *arbre étiqueté* est un couple (τ, ℓ) où τ est un arbre plan et ℓ une *fonction d'étiquetage* associant à chaque sommet de τ un entier relatif. Enfin, on dit que (τ, ℓ) est *bien étiqueté* si :

- l'étiquette $\ell(\rho)$ de la racine vaut 0,
- pour tous sommets voisins v et w de τ, on a $\ell(w) - \ell(v) \in \{-1, 0, 1\}$.

La bijection. Soit (τ, ℓ) un arbre bien étiqueté à $n+1$ sommets, et $(v_i)_{i\in\mathbb{Z}}$ sa liste de contour. On ajoute également un sommet v_∞ portant l'étiquette $\min(\ell) - 1$. Pour tout indice i, on définit le *successeur* de i, noté $s(i)$, comme le plus petit indice $j > i$ tel que $\ell(v_j) = \ell(v_i) - 1$. Notons que si $\ell(v_i) = \min(\ell)$, alors $s(i) = \infty$. Pour tout $0 \leq i \leq 2n-1$, on trace alors une arête reliant v_i à $v_{s(i)}$. Il est possible de tracer ces arêtes de manière à ce qu'elles ne se recoupent pas et n'intersectent pas les arêtes de τ. En effaçant ensuite les arêtes de τ, on obtient ainsi une carte planaire, qu'on enracine sur l'arête reliant v_0 à $v_{s(0)}$ et qu'on note $\Phi(\tau, \ell)$. On peut alors vérifier que $\Phi(\tau, \ell)$ est une quadrangulation à n faces, et que cette construction définit une bijection entre l'ensemble des arbres bien étiquetés à $n+1$ sommets et l'ensemble des quadrangulations de la sphère enracinées à n faces [10]. Notons également que la bijection réciproque Φ^{-1} peut être explicitée.

10. Pour être plus précis, la bijection prend en paramètre un arbre bien étiqueté et un bit $\varepsilon \in \{-1, +1\}$ (qui indique dans quel sens on oriente l'arête racine), et donne une quadrangulation enracinée et *pointée* (munie d'un sommet distingué, en plus de l'arête racine, qui correspond au sommet v_∞).

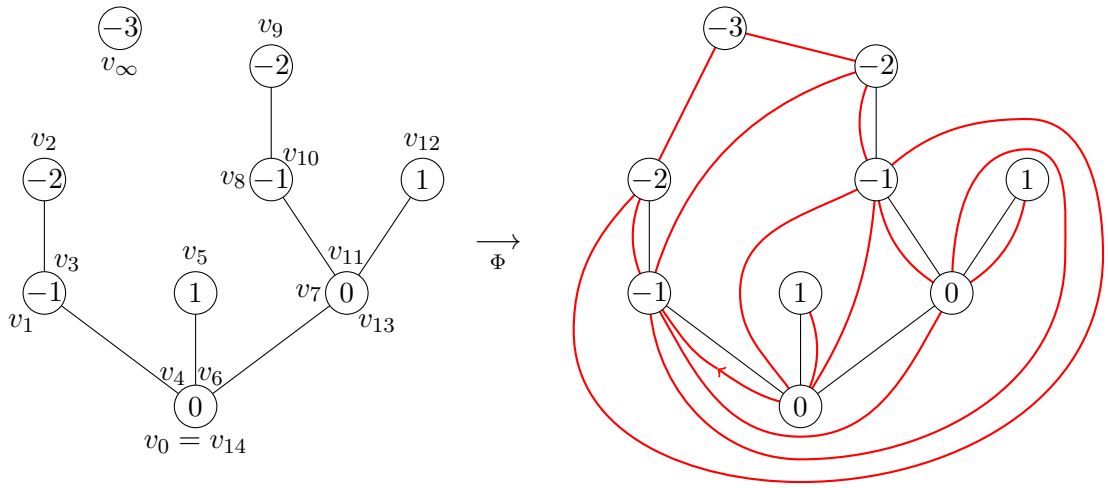

FIGURE 1.9 – La bijection de Cori–Vauquelin–Schaeffer sur un exemple avec $n = 7$. À gauche, un arbre bien étiqueté (τ, ℓ) et sa liste de contour. À droite, la quadrangulation $\Phi(\tau, \ell)$.

Bijection CVS et distances. La bijection a été introduite par Robert Cori et Bernard Vauquelin [55] dans le but d'obtenir une preuve bijective du dénombrement des quadrangulations de la sphère de taille fixée (1.1). Gilles Schaeffer a remarqué [127] qu'elle n'avait pas seulement un intérêt combinatoire, mais codait également des informations métriques. En effet, si (τ, ℓ) est un arbre bien étiqueté et v_i est un sommet de τ, alors la distance dans $\Phi(\tau, \ell)$ entre v_i et v_∞ est égale à

$$\ell(v_i) - \ell(v_\infty).$$

En particulier, les distances de graphe dans $\Phi(\tau, \ell)$ sont du même ordre de grandeur que les étiquettes de (τ, ℓ). Si (τ, ℓ) est choisi uniformément parmi les arbres bien étiquetés à $n + 1$ sommets, alors les branches de τ ont une longueur d'ordre $\sqrt{n}$, et les étiquettes le long d'une branche décrivent une marche aléatoire. L'ordre de grandeur des étiquettes est donc de $\sqrt{\sqrt{n}}$, ce qui donne l'origine de l'exposant 1/4 dans les résultats de limites d'échelle comme le Théorème 1.2.

Variantes. De nombreuses variantes de la bijection CVS ont depuis été introduites. En particulier, une version très robuste a été développée par Jérémie Bouttier, Philippe di Francesco et Emmanuel Guitter [40], et permet de traiter des contraintes arbitraires sur les degrés des faces [11]. D'autres versions ont été utilisées pour étudier des quadrangulations infinies du plan [53, 68] ou du demi-plan [22, 50], ou encore de genre supérieur [52]. Enfin, la plupart des constructions connues de la carte brownienne et d'autres limites d'échelles de cartes sont des variantes continues de la bijection CVS.

1.2.4 Flips sur des triangulations de la sphère (Chapitre 2, ou [44])

Comment réaliser des simulations ? Pour réaliser de jolies images comme celle de la Figure 1.4, il est nécessaire de savoir tirer uniformément au hasard une triangulation de la sphère avec n sommets. Or, d'après la formule (1.2), le nombre de ces triangulations croît exponentiellement en

11. Pour des triangulations, elle est cependant assez lourde à utiliser, car elle les code par des arbres de Galton–Watson à 4 types. C'est pour cette raison qu'on ne l'a décrite que pour des quadrangulations, bien que le reste de cette introduction se concentre principalement sur les triangulations.

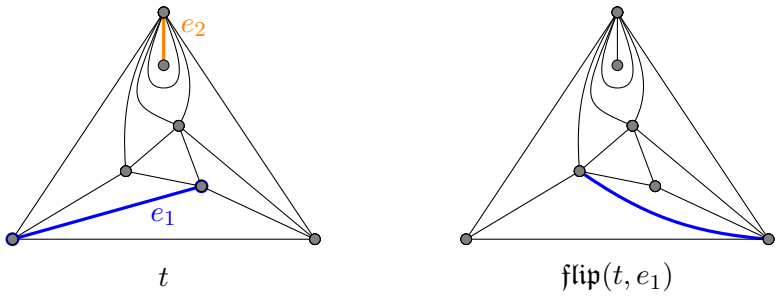

FIGURE 1.10 – Un exemple de flip d'une arête. Il est impossible de flipper l'arête orange e_2.

n. Il n'est donc pas envisageable de lister toutes les triangulations une par une pour en tirer une au hasard. Les outils décrits précédemment fournissent des méthodes permettant de construire des triangulations uniformes. Cependant, quand les physiciens ont commencé à s'intéresser aux cartes aléatoires dans les années 80, aucun de ces outils n'était disponible.

Une méthode de Monte-Carlo. Les physiciens ont donc utilisé des méthodes dites *de Monte-Carlo*. Commençons par rappeler le principe général de ces méthodes. On se donne une mesure μ sur un ensemble fini E, et on cherche à simuler une variable aléatoire sur E de loi μ. Pour cela, on considère une chaîne de Markov (X_n) irréductible sur E dont μ est la mesure stationnaire. Pour n suffisamment grand, la loi de X_n sera alors très proche de celle de μ. Dans notre cas, E est l'ensemble $\mathscr{T}_n$ des triangulations (de type I) de la sphère à n sommets, et μ est la mesure uniforme sur $\mathscr{T}_n$.

Pour qu'une telle méthode soit facile à implémenter, on souhaite que chaque étape de notre chaîne de Markov ne change qu'une petite partie de la triangulation. Une transformation locale semble alors très naturelle : les *flips*. Soit t une triangulation de la sphère, et soit e une arête de t. Si on efface l'arête e, les deux faces qu'elle séparait fusionnent, et forment une face de degré 4 dont on vient de supprimer une diagonale. On peut alors remettre l'autre diagonale de cette face. On dit qu'on a *flippé* l'arête e, et on note $\mathfrak{flip}(t, e)$ la nouvelle triangulation obtenue (voir Figure 1.10). Si jamais les deux côtés de e sont adjacents au même triangle (arête e_2 sur la Figure 1.10), on pose $\mathfrak{flip}(t, e) = t$.

La chaîne de Markov qu'on considère est alors la suivante : on se donne une triangulation initiale t_0 puis, à chaque étape, on choisit uniformément au hasard une des $3n - 6$ arêtes de la triangulation courante et on flippe cette arête. On note $T_n(k)$ la triangulation obtenue après k flips. Il est alors facile de vérifier que $(T_n(k))_{k\geq 0}$ est bien une chaîne de Markov sur $\mathscr{T}_n$, et qu'elle admet la mesure uniforme comme mesure stationnaire (et même réversible). Elle est de plus irréductible d'après un résultat de Klaus Wagner [138] [12] et apériodique, ce qui garantit qu'elle converge bien vers la mesure uniforme. La plupart des simulations présentées dans cette introduction ont été réalisées de cette manière.

Temps de mélange. En pratique, la convergence de $T_n(k)$ vers la mesure uniforme n'est cependant pas suffisante : il faut s'assurer que cette convergence est suffisamment rapide pour que les simulations puissent être effectuées en temps raisonnable. La quantité utilisée habituellement pour décrire cette vitesse de convergence est le *temps de mélange*. Notons $T_n(\infty)$ une variable

12. Pour l'anecdote, l'article de Wagner s'intitule *Bemerkungen zum Vierfarbenproblem*, soit *Remarques sur le théorème des quatre couleurs*. On peut donc supposer que son but était d'utiliser le fait que certaines triangulations sont 4-coloriables, puis de "propager" cette propriété vers toutes les triangulations en utilisant des flips.

aléatoire uniforme sur $\mathscr{T}_n$. Étant donnés n et $0 < \varepsilon < 1$, on note $t_{\text{mix}}(\varepsilon, n)$ le plus petit k tel que

$$\max_{t_0 \in \mathscr{T}_n} \max_{A \subset \mathscr{T}_n} |\mathbb{P}(T_n(k) \in A) - \mathbb{P}(T_n(\infty) \in A)| \leq \varepsilon, \tag{1.3}$$

et on nomme cette quantité *temps de mélange de T_n au niveau ε*. Autrement dit, $t_{\text{mix}}(\varepsilon, n)$ est le premier temps k pour lequel quel que soit la condition initiale t_0 choisie, la loi de $T_n(k)$ approche la mesure uniforme à ε près pour la distance en variation totale.

Théorème 1

Soit $\varepsilon > 0$. Il existe une constante $c > 0$ telle que, pour tout n, on ait $t_{\text{mix}}(\varepsilon, n) \geq cn^{5/4}$.

Nous n'avons réussi qu'à obtenir une borne inférieure sur ce temps de mélange, ce qui est en général plus facile que d'obtenir une borne supérieure : pour minorer un temps de mélange, il suffit de trouver une propriété qui est lente à mélanger, tandis que pour le majorer, il faut prouver que toutes les propriétés possibles se mélangent bien. Il ne semble même pas évident que $t_{\text{mix}}(\varepsilon, n)$ soit polynomial en n. Cependant, on conjecture que ce temps de mélange est bien d'ordre $n^{5/4}$.

Idée de la démonstration. La propriété à laquelle on s'intéresse est l'existence de petits cycles séparants, c'est-à-dire de cycles de longueur $o(n^{1/4})$ qui séparent une triangulation en deux parties macroscopiques. Jean-François Le Gall et Frédéric Paulin ont montré [103] qu'avec grande probabilité, de tels cycles n'existent pas dans une triangulation uniforme [13]. On montre alors que si la condition initiale possède un petit cycle séparant, elle ne pourra pas s'en débarrasser en $o(n^{5/4})$ flips, ce qui montre qu'on est alors loin de la mesure uniforme. La preuve de la "persistance" de ces petits cycles repose sur une méthode d'épluchage : on explore notre triangulation en même temps qu'elle est transformée par les flips. Notons que le résultat de Le Gall et Paulin repose sur des bijections de type CVS, et qu'on utilise à la fin un résultat de Krikun sur les cycles séparants, montré grâce à la décomposition de Krikun [94]. La démonstration combine donc en quelque sorte les trois outils vus précédemment !

1.3 Cartes aléatoires hyperboliques

1.3.1 Définitions

Triangulations markoviennes du demi-plan. La puissance et la souplesse de la propriété de Markov spatiale de l'UIPT (voir Section 1.2.1) amènent naturellement à se poser la question suivante : quelles sont les triangulations infinies aléatoires vérifiant une propriété similaire ? Cette question a dans un premier temps été étudiée par Omer Angel et Gourab Ray [18], pour les triangulations du demi-plan (de type II).

Théorème 1.3 ([18])**.** Les triangulations du demi-plan vérifiant une propriété de Markov spatiale forment une famille à un paramètre $(\mathbb{H}_\alpha)_{0 \leq \alpha < 1}$.

La construction des triangulations $\mathbb{H}_\alpha$ repose sur des algorithmes d'épluchage, comme ceux évoqués précédemment. De plus, le paramètre α a une interprétation probabiliste assez simple : fixons une arête du bord $\partial\mathbb{H}_\alpha$ de $\mathbb{H}_\alpha$, et soit f_0 la face triangulaire de $\mathbb{H}_\alpha$ adjacente à cette arête. Alors α est la probabilité que le troisième sommet de f_0 ne soit pas sur $\partial\mathbb{H}_\alpha$. Intuitivement, plus α est grand, plus $\mathbb{H}_\alpha$ contient de sommets en dehors de son bord. Il semble donc naturel que

13. Pour être exact, le résultat énoncé par Le Gall et Paulin porte sur les *quadrangulations*. Cependant, il repose sur la convergence vers la carte brownienne, qui est vraie également pour les triangulations.

plus α est grand, plus la croissance de $\mathbb{H}_\alpha$ est rapide. Angel et Ray mettent en effet en évidence une transition de phase en $\alpha = \frac{2}{3}$:

— pour $\alpha < \frac{2}{3}$, la triangulation $\mathbb{H}_\alpha$ ressemble de loin à un arbre, et peut aussi être vue comme la limite locale de triangulations à bord où la longueur du bord est proportionnelle au nombre de total de sommets,

— pour $\alpha = \frac{2}{3}$, la triangulation $\mathbb{H}_\alpha$ n'est autre que l'UIHPT, version semi-planaire de l'UIPT, et avec laquelle elle partage de nombreuses propriétés,

— pour $\alpha > \frac{2}{3}$, la triangulation $\mathbb{H}_\alpha$ est un graphe "hyperbolique".

Les propriétés de ces trois régimes seront précisées dans la Section 1.3.2.

Triangulations markoviennes du plan. Par la suite, Nicolas Curien a généralisé ces résultats aux triangulations du plan complet [58]. Pour être exact, l'article [58] se concentre sur les triangulations *de type II*, c'est-à-dire sans boucles. Cependant, toutes les preuves s'adaptent au type I qui, pour des raisons techniques, est celui dont on aura besoin dans la suite. On se contentera donc dans cette introduction du type I. Le lien entre les triangulations markoviennes de type I et de type II est clarifié dans l'Appendice A de cette thèse.

La définition choisie de la propriété de Markov spatiale est cette fois légèrement plus forte. Soit $\lambda \geq 0$. On dit qu'une triangulation aléatoire du plan $\mathbb{T}$ est λ-markovienne s'il existe des constantes $(C_p)_{p\geq 1}$ telles que, pour toute triangulation finie t avec un trou de périmètre p, on a

$$\mathbb{P}\left(t \subset T\right) = C_p \lambda^{|t|},$$

où $|t|$ est le nombre total de sommets de t. Commençons par expliquer le lien entre cette propriété de Markov spatiale et celle énoncée plus haut (Section 1.2.1). Soient t et t' deux triangulations avec des trous de périmètres respectivement p et p', et vérifiant $t \subset t'$. Soit aussi $\mathbb{T}$ une triangulation aléatoire λ-markovienne. Alors on a

$$\mathbb{P}\left(t' \subset \mathbb{T} | t \subset \mathbb{T}\right) = \frac{\mathbb{P}\left(t' \subset \mathbb{T}\right)}{\mathbb{P}\left(t \subset \mathbb{T}\right)} = \frac{C_{p'}\lambda^{|t'|}}{C_p \lambda^{|t|}} = \frac{C_{p'}}{C_p}\lambda^{|t'\setminus t|}.$$

Si on fixe t et fait varier t', cette formule caractérise entièrement la loi de $\mathbb{T}\setminus t$ conditionnellement à $t \subset \mathbb{T}$. Or, la seule dépendance en t est en fait en le périmètre de t, donc la loi de $\mathbb{T}\setminus t$ conditionnellement à $t \subset \mathbb{T}$ ne dépend que du périmètre de t.

Une famille à un paramètre. De manière analogue à ce qui se passe dans le demi-plan, Curien obtient le résultat suivant. Dans toute la suite, on notera $\lambda_c = \frac{1}{12\sqrt{3}}$.

Théorème 1.4. Si $0 < \lambda \leq \lambda_c$, alors il existe une unique (en loi) triangulation aléatoire λ-markovienne du plan, notée $\mathbb{T}_\lambda$. Si $\lambda > \lambda_c$, il n'en existe pas. De plus, la triangulation $\mathbb{T}_{\lambda_c}$ est en fait l'UIPT.

Comme pour les triangulations du demi-plan, la preuve repose sur des arguments d'épluchage. Notons que le régime sous-critique ($\alpha < \frac{2}{3}$) disparaît dans le plan complet. De plus, si $0 < \lambda < \lambda_c$, il existe $\frac{2}{3} < \alpha < 1$ pour lequel on peut coupler $\mathbb{T}_\lambda$ avec $\mathbb{H}_\alpha$ de telle manière que $\mathbb{H}_\alpha \subset \mathbb{T}_\lambda$ [14]. Les triangulations $\mathbb{T}_\lambda$ pour $0 < \lambda < \lambda_c$ ont donc elles aussi un comportement "hyperbolique". Elles sont appelées *PSHIT* (Planar Stochastic Hyperbolic Infinite Triangulations), et constituent l'objet d'étude central de cette thèse. La Figure 1.11 donne une idée de ce à quoi ressemblent ces cartes.

14. Ici aussi, on triche un peu : en toute rigueur, il faudrait remplacer les $\mathbb{H}_\alpha$ par un analogue de type I.

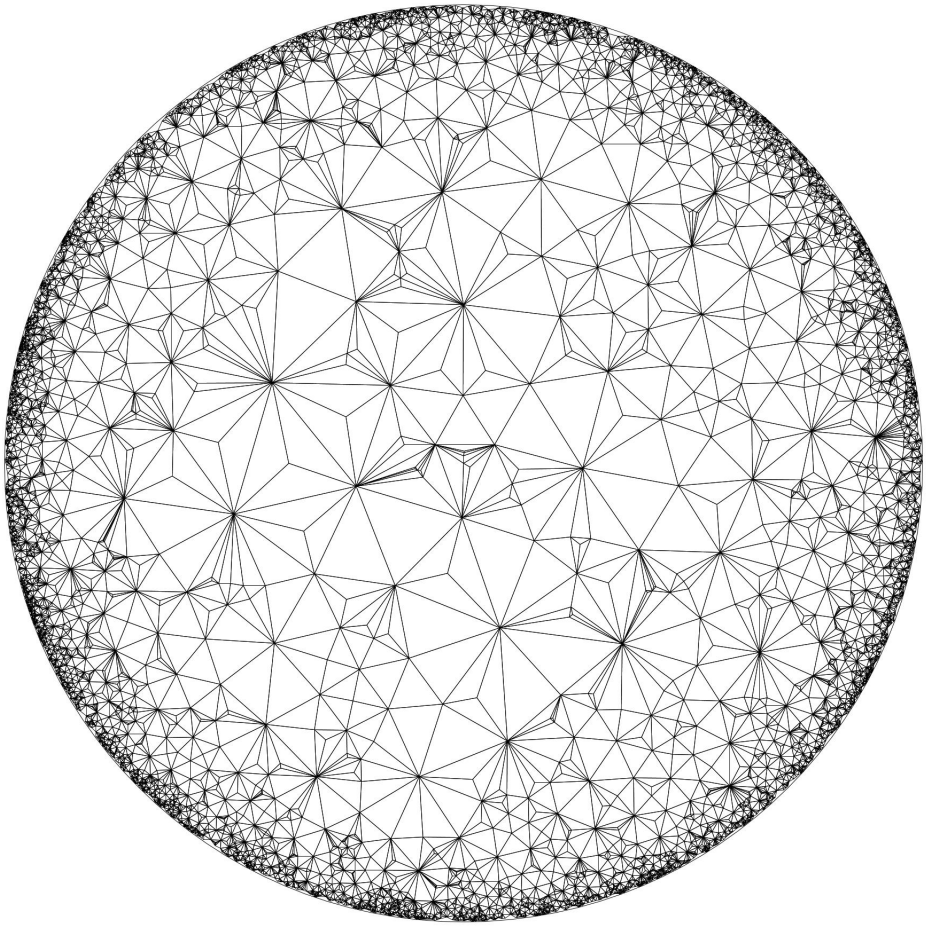

FIGURE 1.11 – Une simulation de PSHIT. Comme pour la Figure 1.5, le plongement a été calculé par le logiciel CirclePack, et on a effacé les cercles pour ne garder que les arêtes. Elle est ici représentée dans le disque unité, ce qui est plus naturel pour des objets de nature hyperbolique.

Une conjecture. De nombreux modèles naturels de cartes aléatoires infinies peuvent être vus comme des limites locales de cartes planaires aléatoires finies dont la taille tend vers l'infini. Pour les PSHIT, ce n'est pas possible. En effet, le degré moyen des sommets dans les PSHIT est strictement supérieur à six [15]. En revanche, dans une triangulation de la sphère à n sommets, la formule d'Euler montre que le nombre d'arêtes vaut $3n-6$, donc le degré moyen d'un sommet choisi uniformément vaut $6-\frac{12}{n}$. Par conséquent, dans une limite locale de triangulations de la sphère, le degré moyen vaudra au plus 6. Pour voir les PSHIT comme une limite locale de cartes finies, d'après la formule d'Euler, il est donc nécessaire de considérer des cartes de genre supérieur. Plus précisément, il faut que le genre soit linéaire en le nombre de faces de la triangulation. Pour tous n et g, soit $T_{n,g}$ une variable aléatoire uniforme sur l'ensemble des triangulations enracinées à n sommets du tore de genre g. La conjecture suivante est posée dans [58], et est une des motivations pour l'introduction des PSHIT.

Conjecture 1.4. Soit $0<\theta<+\infty$. Alors il existe $0<\lambda<\lambda_c$ tel que $T_{n,\lfloor\theta n\rfloor}$ converge en loi pour la topologie locale vers $\mathbb{T}_\lambda$ quand $n\to+\infty$.

Il peut paraître surprenant que des graphes hautement non-planaires convergent vers une triangulation du plan. Ce phénomène peut cependant être facilement observé, par exemple chez les graphes aléatoires : la limite locale de graphes 3-réguliers uniformes, objets très loin d'être planaires, est l'arbre binaire complet. Ici, cela signifie que le genre ne se voit pas au voisinage de la racine dans $T_{n,\lfloor\theta n\rfloor}$.

Établir la Conjecture 1.4 exigerait des résultats très précis sur les asymptotiques du nombre de triangulations à n sommets du tore de genre g quand n et g tendent simultanément vers l'infini. De tels résultats ne sont pour l'instant pas disponibles. Signalons cependant qu'un résultat similaire a été établi par Omer Angel, Guillaume Chapuy, Nicolas Curien et Gourab Ray pour des cartes unicellulaires (c'est-à-dire ne possédant qu'une seule face) de genre grand [13]. La limite locale est alors un arbre de Galton–Watson surcritique conditionné à survivre.

Enfin, par la formule d'Euler, on connaît le degré moyen des sommets dans $T_{n,\lfloor\theta n\rfloor}$. Un moyen d'établir un lien entre les paramètres θ et λ est donc de comprendre le "degré moyen" dans les triangulations infinies $\mathbb{T}_\lambda$. Dans l'Appendice B de cette thèse, nous effectuons des calculs sur le degré de la racine dans les PSHIT, et nous proposons une version plus précise (Conjecture B.1) de la Conjecture 1.4.

Au-delà des triangulations. Comme pour l'UIPT, il est naturel de se demander s'il existe des analogues des PSHIT chez les quadrangulations, ou bien chez des cartes plus générales. Nous traitons partiellement cette question dans l'Appendice C de cette thèse : on peut définir des analogues des PSHIT pour les $2p$-angulations pour tout $p\geq 2$, ainsi que des modèles de cartes plus généraux où les degrés des faces sont bornés. On peut également énoncer des analogues de la Conjecture 1.4 pour ces modèles. En revanche, l'existence ou non de cartes hyperboliques avec des faces de degrés à queue lourde reste ouverte.

1.3.2 Propriétés d'hyperbolicité

Le terme "hyperbolique" pour un graphe peut recouvrir des propriétés de natures très variées. Le but de cette section est de préciser ce qu'on entend par "hyperbolicité", et de résumer les propriétés connues des triangulations hyperboliques définies ci-dessus. On se concentre sur le cas surcritique $\alpha>\frac{2}{3}$ pour les $\mathbb{H}_\alpha$, et $\lambda<\lambda_c$ pour les $\mathbb{T}_\lambda$.

15. Ceci est à rapprocher du fait que le degré des sommets vaut 6 dans le réseau triangulaire euclidien, et au moins 7 dans les triangulations régulières hyperboliques.

Croissance. Une première manière de définir l'hyperbolicité pour un graphe est par sa croissance, c'est-à-dire la manière dont la taille de la boule de rayon r varie en fonction de r. Un graphe hyperbolique se doit d'avoir une croissance très rapide.

Gourab Ray a montré [124] que le volume et le périmètre des enveloppes de rayon r dans $\mathbb{H}_\alpha$ croît exponentiellement en r. Nicolas Curien a obtenu des résultats de croissance plus précis dans les PSHIT [58] : on a les convergences

$$\frac{|\partial B_r^\bullet(\mathbb{T}_\lambda)|}{m_\lambda^r} \xrightarrow[r\to+\infty]{p.s.} \Pi_\lambda \qquad \text{et} \qquad \frac{|B_r^\bullet(\mathbb{T}_\lambda)|}{|\partial B_r^\bullet(\mathbb{T}_\lambda)|} \xrightarrow[r\to+\infty]{p.s.} a_\lambda, \tag{1.4}$$

où Π_λ est une variable aléatoire strictement positive, et $a_\lambda, m_\lambda > 0$ sont des constantes calculables explicitement en fonction de λ.

Expansion enracinée. Les PSHIT vérifient également une propriété d'*expansion enracinée*[16] : il existe une constante $c > 0$ telle que presque sûrement, pour tout ensemble A de sommets fini, assez grand, connexe et contenant la racine, on a

$$|\partial A| \geq c|A|, \tag{1.5}$$

où $|\partial A|$ est le nombre d'arêtes de $\mathbb{T}_\lambda$ dont exactement une extrémité est dans A. Cette propriété a été montrée dans [124] pour les triangulations du demi-plan, puis dans [58] pour les $\mathbb{T}_\lambda$. La notion d'expansion enracinée est très proche de celle de *non-moyennabilité*, très utilisée pour étudier des graphes plus réguliers (graphes de Cayley de groupes infinis par exemple) : un graphe est dit non-moyennable s'il vérifie une inégalité isopérimétrique du même type que (1.5), uniformément pour tous les ensembles finis de sommets. Cependant, cette notion n'est pas adaptée à l'étude de graphes aléatoires : comme $\mathbb{T}_\lambda$ contient presque sûrement toutes les triangulations finies, il est en fait moyennable.

Marche aléatoire simple. Soit (X_n) la marche aléatoire simple sur $\mathbb{T}_\lambda$. Par un résultat de Balint Virág [137], l'expansion enracinée implique que (X_n) est transiente, c'est-à-dire qu'elle ne visite la racine qu'un nombre fini de fois. Une question naturelle est alors celle de la *vitesse* à laquelle la marche simple s'éloigne de la racine. Cette vitesse est strictement positive [58], c'est-à-dire que $\frac{d_{\mathbb{T}_\lambda}(X_0,X_n)}{n}$ converge vers une constante strictement positive (dépendant de λ) quand $n \to +\infty$. La preuve de ce résultat mélange des arguments d'épluchage et de théorie ergodique. Notons qu'Omer Angel, Asaf Nachmias et Gourab Ray ont également établi un résultat de vitesse positive pour les triangulations du demi-plan $\mathbb{H}_\alpha$ [17], et ce par des méthodes différentes (utilisation plus précise de l'expansion enracinée). Ils ont également montré que la probabilité de retour à l'origine au temps n était d'ordre $\exp(-n^{1/3})$.

Frontière de Poisson. Les PSHIT étant transientes, une autre question naturelle est la suivante : la marche aléatoire peut-elle tendre vers l'infini dans plusieurs directions différentes ? La notion naturelle pour formuler cette question plus précisément est celle de *frontière de Poisson*. Considérons une compactification $\widehat{G}$ d'un graphe transient G, c'est-à-dire un espace topologique compact dans lequel G est dense, et notons $\partial\widehat{G} = \widehat{G}\backslash G$ (par exemple, cela peut être le bord du disque unité sur la Figure 1.11). Supposons que presque sûrement, la marche aléatoire simple (X_n) sur G converge dans $\widehat{G}$ vers un point (aléatoire) X_∞ de $\partial\widehat{G}$. Alors pour toute fonction mesurable bornée f sur $\partial\widehat{G}$, on peut définir sur G la fonction suivante :

$$h(x) = \mathbb{E}_x\left[f(X_\infty)\right],$$

16. *anchored expansion* en anglais, la traduction n'est pas canonique.

où $\mathbb{E}_x$ est la loi de la marche aléatoire simple issue de x. Il est alors facile de vérifier que h est toujours harmonique et bornée. Si de plus toutes les fonctions harmoniques bornées sur G peuvent être construites de la sorte, on dit que $\partial\widehat{G}$ est une *réalisation de la frontière de Poisson* de G. On dit également que G vérifie la *propriété de Liouville* si sa frontière de Poisson est réduite à un point, c'est-à-dire que toutes les fonctions harmoniques bornées sur G sont constantes. Notons que tous les graphes récurrents vérifient la propriété de Liouville. De plus, $\mathbb{Z}^d$, bien que transient pour $d \geq 3$, vérifie aussi toujours cette propriété. La non-propriété de Liouville est donc plus forte que la transience, et peut-être vue comme une propriété d'hyperbolicité.

Nicolas Curien a prouvé que les PSHIT ne vérifient pas la propriété de Liouville [58], puis Omer Angel, Tom Hutchcroft, Asaf Nachmias et Gourab Ray ont construit une réalisation de leur frontière de Poisson en utilisant la théorie des empilements de cercles [14]. Cette réalisation est homéomorphe au cercle. Les résultats de [14] sont en fait valides pour n'importe quelle triangulation infinie *unimodulaire*, c'est-à-dire où "le sommet racine a été choisi uniformément" (voir par exemple [110, Chapitre 8] pour une définition précise).

1.3.3 Limite d'échelle quasi-critique (Chapitre 3, ou [45])

Limite quasi-critique. On a vu dans la Section 1.1.4 que l'UIPT admet une limite d'échelle continue, le *plan brownien* $\mathcal{P}$. Les PSHIT admettent-elles également une limite d'échelle ? Une première tentative naturelle serait de fixer $0 < \lambda < \lambda_c$, de multiplier les distances dans $\mathbb{T}_\lambda$ par $\alpha > 0$, puis de faire tendre α vers 0. Cependant, cette tentative est vouée à l'échec car les objets à croissance exponentielle se comportent très mal vis-à-vis des limites d'échelle. Par exemple, quand r est grand, le volume de la boule de rayon $2r$ croît beaucoup plus vite que celui de la boule de rayon r. Dans une éventuelle limite continue, la masse de la boule de rayon 2 devrait donc être infiniment plus grande que la masse de la boule de rayon 1.

Pour obtenir une limite d'échelle non-dégénérée, il faut donc ruser en "amortissant" cette croissance trop rapide. Nous proposons une approche *quasi-critique* : en même temps qu'on renormalise les distances, on fait tendre le paramètre λ vers sa valeur critique λ_c. Il est alors important trouver la bonne vitesse à laquelle faire tendre λ vers λ_c : si la convergence est trop lente, on retrouve le même problème que précédemment, et si elle est trop rapide, on ne voit plus la différence avec le cas critique et on obtient une convergence vers le plan brownien.

Le plan brownien hyperbolique. On peut voir $\mathbb{T}_\lambda$ comme un espace métrique mesuré, muni de la distance de graphe et de la mesure de comptage sur les sommets. Si X est un espace métrique mesuré et $\alpha > 0$, on note donc αX l'espace métrique mesuré obtenu en multipliant les distances par α et la mesure par α^4.

Théorème 2

Soit (λ_n) une suite de nombres dans $]0, \lambda_c]$ vérifiant

$$\lambda_n = \lambda_c \left(1 - \frac{2}{3n^4} + o\left(\frac{1}{n^4}\right)\right).$$

Alors on a la convergence suivante pour la distance de Gromov–Hausdorff–Prokhorov :

$$\frac{1}{n}\mathbb{T}_{\lambda_n} \xrightarrow[n\to+\infty]{(loi)} \mathcal{P}^h,$$

où $\mathcal{P}^h$ est un espace métrique mesuré aléatoire localement compact appelé *plan brownien hyperbolique*.

Absolue continuité entre $\mathcal{P}$ et $\mathcal{P}^h$. La preuve de ce résultat repose sur la convergence de l'UIPT vers le plan brownien et sur les relations d'absolue continuité entre les enveloppes dans l'UIPT et dans les PSHIT. Plus précisément, pour tout r, l'enveloppe $B_r^\bullet(\mathbb{T}_\lambda)$ est absolument continue par rapport à $B_r^\bullet(\mathbb{T}_{\lambda_c})$. De plus, la dérivée de Radon-Nikodym ne dépend que du périmètre et du volume de l'enveloppe, et peut être calculée explicitement. Il n'est alors pas surprenant que de telles relations d'absolue continuité existent également entre $\mathcal{P}^h$ et $\mathcal{P}$. Plus précisément, si X est un espace métrique mesuré, on note $\overline{B_r}(X)$ l'*enveloppe* de rayon r de X, c'est-à-dire l'union de la boule de rayon r et des composante connexes bornées de son complémentaire. Muni de la restriction de la distance et de la mesure sur X, c'est un espace métrique mesuré [17]. Notons P_r le périmètre de $\overline{B_r}(\mathcal{P})$ comme défini dans [63], et V_r son volume.

Théorème 3

Pour tout $r \geq 0$, l'espace métrique mesuré $\overline{B_r}(\mathcal{P}^h)$ est absolument continu par rapport à $\overline{B_r}(\mathcal{P})$, et la dérivée de Radon–Nikodym est donnée par

$$e^{-2V_{2r}} e^{P_{2r}} \int_0^1 e^{-3P_{2r}x^2} \mathrm{d}x. \tag{1.6}$$

Notons que la densité dépend du périmètre et du volume de l'enveloppe de rayon $2r$, et non pas r. C'est lié au fait que des "raccourcis" entre deux points de $\overline{B_r}(\mathcal{P})$ peuvent exister en-dehors de $\overline{B_r}(\mathcal{P})$, mais ces raccourcis ne peuvent pas sortir de $\overline{B_{2r}}(\mathcal{P})$. Signalons aussi que le facteur $\frac{2}{3}$ dans le Théorème 2 a été choisi pour que l'expression de la densité (1.6) soit la plus simple possible.

Propriétés de $\mathcal{P}^h$. Les propriétés de $\mathcal{P}^h$ peuvent grossièrement se résumer de la manière suivante :

- localement, $\mathcal{P}^h$ ressemble à $\mathcal{P}$,
- à une échelle r fixée, $\mathcal{P}^h$ est absolument continu par rapport à $\mathcal{P}$,
- à grande échelle, $\mathcal{P}^h$ a un comportement hyperbolique.

En particulier, les propriétés locales de confluence des géodésiques dans la carte brownienne et le plan brownien restent vraies dans $\mathcal{P}^h$. Nous montrons également que $\mathcal{P}^h$ est presque sûrement homéomorphe au plan, et vérifie une propriété d'"invariance par réenracinement" [18].

Périmètres et volumes dans $\mathcal{P}^h$. Par absolue continuité, on peut également définir le périmètre et le volume de l'enveloppe de rayon r dans $\mathcal{P}^h$, qu'on note respectivement P_r^h et V_r^h. Le Théorème 3 combiné aux résultats de Nicolas Curien et Jean-François le Gall [63] sur la loi de (P, V) sont en théorie suffisants pour caractériser entièrement la loi de (P^h, V^h). Nous donnons cependant une description plus directe et plus explicite de cette loi, similaire à celle donnée dans [63]. Pour tout $\delta > 0$, on définit la mesure de probabilité ν_δ sur $\mathbb{R}_+^*$ comme suit :

$$\nu_\delta(\mathrm{d}x) = \frac{\delta^3 e^{2\delta}}{1+2\delta} \frac{e^{-\frac{\delta^2}{2x}-2x}}{\sqrt{2\pi x^5}} \mathbb{1}_{x>0}\, \mathrm{d}x.$$

17. La notation $\overline{B_r}$ est différente de la notation $B_r^\bullet$ utilisée dans le cadre discret. Ce changement de notation souligne le fait que $\overline{B_r}$ n'est pas muni de sa distance *intrinsèque*, mais de la distance *induite*. En particulier, si X est un graphe, la distance entre deux sommets de $\overline{B_r}(X)$ ne dépend pas seulement de $B_r^\bullet(X)$.

18. Pour être plus précis, il s'agit d'un analogue continu de la propriété d'*unimodularité*, qui est elle-même la généralisation naturelle aux graphes infinis de l'invariance par réenracinement uniforme.

Théorème 4

- Le processus P^h est un processus de branchement sous-critique démarré de $+\infty$ au temps $-\infty$ et conditionné à atteindre 0 au temps 0, où le temps a été inversé.
- De plus, conditionnellement à P^h, la loi de V^h est celle du processus

$$\Big(\sum_{s_i \leq r} \xi_i^h\Big)_{r \geq 0},$$

où (s_i) est une énumération des sauts de P^h, les variables ξ_i^h sont indépendantes, et ξ_i^h a pour loi $\nu_{|\Delta P^h_{s_i}|}$ pour tout i.

En particulier, le processus P^h est donc un processus discontinu, dont tous les sauts sont négatifs. De plus, le processus des volumes V^h est croissant, et ne progresse que par sauts : chaque fois que P^h effectue un saut négatif, le volume V^h effectue un saut positif. Ces instants où le périmètre diminue et le volume augmente correspondent aux hauteurs où l'enveloppe avale une des "tentacules" de la Figure 1.6. Cette description est très proche de celle de P et V donnée dans [63].

Les deux seules différences sont la nature du processus de branchement (critique pour $\mathcal{P}$, sous-critique pour $\mathcal{P}^h$), et les lois ν_δ. Il est logique que le processus de branchement soit sous-critique : on s'attend à une croissance rapide de P_r^h, donc à une décroissance rapide une fois le temps inversé. Plus précisément, le Théorème 4 nous permet d'obtenir des asymptotiques très précises sur les processus P^h et V^h, analogues continus des résultats obtenus dans [58] sur les PSHIT (1.4). Dans le cas continu, on peut même calculer la loi de la variable limite qui apparaît.

Corollaire 5

On a les convergences

$$\frac{P_r^h}{e^{2\sqrt{2}\,r}} \xrightarrow[n\to+\infty]{p.s.} \mathcal{E} \qquad \text{et} \qquad \frac{V_r^h}{P_r^h} \xrightarrow[n\to+\infty]{p.s.} \frac{1}{4},$$

où $\mathcal{E}$ est une variable exponentielle de paramètre 12.

Signalons enfin que des résultats de limite d'échelle ont été montrés pour des analogues du cas sous-critique chez les *quadrangulations du demi-plan* par Erich Baur et Loïc Richier [23], et par Erich Baur, Grégory Miermont et Gourab Ray [22]. D'une part, la limite d'échelle de ces quadrangulations infinies sous-critiques (à paramètre fixé) est un arbre continu non-compact (version infinie de l'arbre brownien d'Aldous). D'autre part, il existe aussi une limite d'échelle quasi-critique, qui ressemble localement au demi-plan brownien mais asymptotiquement à l'arbre brownien infini. Les méthodes utilisées sont alors complètement différentes (utilisation de bijections du type CVS).

1.3.4 Géodésiques infinies (Chapitre 4, ou [46])

Géodésiques. Considérons une triangulation $\mathbb{T}_\lambda$ avec $0 < \lambda < \lambda_c$. Dans [46], on s'intéresse principalement aux géodésiques infinies de $\mathbb{T}_\lambda$ issues de la racine. L'étude des géodésiques (ou chemins les plus courts) d'un espace métrique ou d'un graphe est généralement un bon moyen d'acquérir des informations sur sa géométrie.

Définition 1.5. Soit I un intervalle de $\mathbb{N}$. Une géodésique dans un graphe G est une famille $(\gamma(i))_{i\in I}$ de sommets de G telle que pour tous $i, j \in I$, la distance de graphe entre $\gamma(i)$ et $\gamma(j)$

vaut $|i-j|$. Si $I = \mathbb{N}$ et $\gamma(0) = \rho$, on dit que γ est une *géodésique infinie issue de* ρ.

En particulier, un chemin de longueur minimale entre deux sommets est toujours une géodésique finie. Rappelons que dans l'UIPT, des propriétés de confluence des géodésiques assez fortes ont été établies par Nicolas Curien et Laurent Ménard [67] : presque sûrement, il existe une infinité de sommets distincts $(v_i)_{i\in\mathbb{N}}$ tels que toute géodésique infinie issue de la racine passe par tous les v_i. Le même résultat avait été montré auparavant par les mêmes auteurs et Grégory Miermont pour l'UIPQ [68]. À l'inverse, dans un graphe de nature hyperbolique, on s'attend à ce qu'il existe de nombreuses géodésiques infinies qui s'éloignent très vite les unes des autres.

Géodésiques les plus à gauche. On s'intéresse plus spécifiquement à une classe de géodésiques particulière : les *géodésiques les plus à gauche*.

Définition 1.6.
- Étant donnés une carte planaire M dessinée dans le plan et deux sommets x, y de M, on dit qu'une géodésique finie γ de x vers y est une *géodésique la plus à gauche* si pour toute géodésique γ' de x vers y, le chemin γ reste à gauche de γ'.
- On dit qu'une géodésique infinie γ est une *géodésique infinie la plus à gauche* si pour tous i et j, la portion de γ reliant $\gamma(i)$ à $\gamma(j)$ est une géodésique la plus à gauche de $\gamma(i)$ vers $\gamma(j)$.

Il est assez facile de voir qu'il existe toujours une unique géodésique la plus à gauche d'un sommet vers un autre, et un argument de compacité montre qu'il existe toujours au moins une géodésique infinie la plus à gauche issue d'un sommet. De plus, une fois que deux géodésiques infinies les plus à gauche sont séparées, elles ne peuvent plus se réintersecter (car alors une des deux serait forcément à droite de l'autre). Les géodésiques infinies les plus à gauche issues de la racine forment donc un arbre infini sans feuille, qu'on note $\boldsymbol{\tau}_\lambda^g$. Cet arbre sépare de plus $\mathbb{T}_\lambda$ en triangulations infinies qu'on appelle des *bandes*, et qui sont bordées par deux géodésiques infinies (voir Figure 1.12).

Arbres de Galton–Watson. Afin d'énoncer un résultat précis, nous avons besoin de définir les *arbres de Galton–Watson*, qui sont certainement le modèle le plus naturel et le plus étudié d'arbres aléatoires. Ces arbres, finis ou infinis, sont des *arbres plans* comme ceux définis dans la Section 1.2.3. Un tel arbre peut toujours être vu comme un arbre généalogique, où le début ρ de l'arête racine est l'ancêtre commun à tous les sommets, et le parent d'un sommet est son unique voisin qui le rapproche de ρ. Soit μ une mesure de probabilité sur $\mathbb{N}$. On dit qu'un arbre aléatoire (fini ou infini) $\boldsymbol{\tau}$ est un *arbre de Galton–Watson de loi de reproduction* μ si les nombres d'enfants de chacun de ses sommets sont i.i.d. de loi μ. On note alors sa loi GW_μ. Cela revient à dire que le nombre d'enfants de la racine a pour loi μ et que, conditionnellement à ce nombre d'enfants, les arbres des descendants de ces enfants sont i.i.d. de loi GW_μ. Soit $\boldsymbol{\tau}$ un arbre de Galton–Watson de loi de reproduction μ, et notons $m = \sum_{i\geq 0} i\mu(i)$ le nombre moyen d'enfants par sommet. On dit que $\boldsymbol{\tau}$ est *sous-critique* si $m < 1$, *critique* si $m = 1$ et *surcritique* si $m > 1$. Il est alors bien connu que $\boldsymbol{\tau}$ est presque sûrement fini si et seulement si il est critique ou sous-critique (sauf dans le cas trivial $\mu(1) = 1$, où l'arbre est infini). De plus, si $\boldsymbol{\tau}$ est surcritique et survit, le nombre de sommets à la génération r croît exponentiellement en r.

Décomposition en bandes de $\mathbb{T}_\lambda$. La description que nous donnons fait intervenir un paramètre m_λ, qui est une fonction explicite de λ vérifiant $m_\lambda \leq 1$, avec égalité si et seulement si $\lambda = \lambda_c$. Il coïncide avec le taux de croissance exponentielle des PSHIT (1.4).

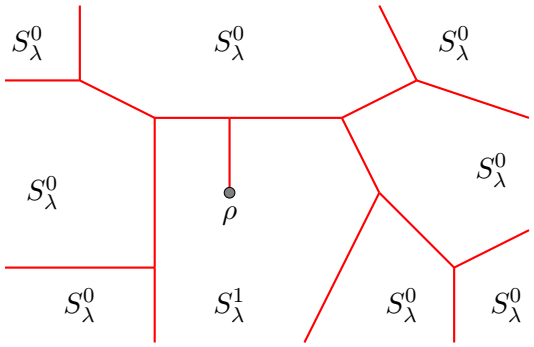

FIGURE 1.12 – La décomposition en bandes de $\mathbb{T}_\lambda$. L'arbre $\boldsymbol{\tau}^g_\lambda$ est en rouge, et sépare $\mathbb{T}_\lambda$ en une infinité de bandes.

Théorème 6

- L'arbre $\boldsymbol{\tau}^g_\lambda$ un arbre de Galton–Watson surcritique de loi de reproduction μ_λ, où $\mu_\lambda(0) = 0$ et
$$\mu_\lambda(k) = m_\lambda(1 - m_\lambda)^{k-1}$$
pour tout $k \geq 1$.
- Il existe deux bandes aléatoires S^0_λ et S^1_λ telles que, conditionnellement à $\boldsymbol{\tau}^g_\lambda$:
  — les bandes délimitées par $\boldsymbol{\tau}^g_\lambda$ sont indépendantes,
  — la bande adjacente à l'arête racine a la même loi que S^1_λ,
  — toutes les autres bandes ont la même loi que S^0_λ.

En particulier, pour $\lambda = \lambda_c$, l'arbre $\boldsymbol{\tau}^g_\lambda$ consiste en une seule géodésique infinie, et on retrouve une conséquence du résultat de Curien et Ménard sur l'UIPT [67]. En revanche, dès que $\lambda < \lambda_c$, l'arbre $\boldsymbol{\tau}^g_\lambda$ est surcritique : il existe une infinité de géodésiques infinies les plus à gauche, et leur nombre à hauteur r est exponentiel en r.

Décomposition de Krikun de $\mathbb{T}_\lambda$. L'outil principal pour prouver le Théorème 6 est la décomposition de Krikun (voir Section 1.2.2). En adaptant les calculs effectués dans [65] pour l'UIPT, on calcule la loi du squelette de $\mathbb{T}_\lambda$. Rappelons que le squelette de l'UIPT peut être vu comme un arbre de Galton–Watson critique à l'envers, conditionné à s'éteindre exactement à la hauteur 0. Le squelette de $\mathbb{T}_\lambda$ pour $\lambda < \lambda_c$ présente deux différences notables :

- la loi de reproduction est maintenant sous-critique,
- le squelette n'est plus un arbre, mais une forêt composée d'une infinité d'arbres infinis.

Les branches infinies de $\boldsymbol{\tau}^g_\lambda$ sont alors les chemins passant entre les arbres du squelette. De plus, chacun de ces arbres correspond à une des bandes de la Figure 1.12. L'indépendance des bandes dans le Théorème 6 vient alors de l'indépendance des arbres du squelette. Enfin, comme la loi de reproduction est sous-critique, les arbres ont tendance à rester très "fins" (largeur constante quand la hauteur tend vers l'infini). Les bandes S^0_λ et S^1_λ seront donc également "fines".

Hyperbolicité métrique. Signalons également dès maintenant que, grâce à des résultats que nous montrons dans le chapitre 5, l'existence d'un arbre de Galton–Watson surcritique à l'intérieur de $\mathbb{T}_\lambda$ permet d'établir des propriétés d'hyperbolicité métrique plus fortes que la croissance exponentielle (voir Théorème 9 plus loin).

Frontière de Poisson. On se restreint ici au cas $\lambda < \lambda_c$. La construction de $\mathbb{T}_\lambda$ à partir de l'arbre $\boldsymbol{\tau}_\lambda^g$ permet de munir $\mathbb{T}_\lambda$ d'une notion assez naturelle de *bord*. Plus précisément, étant donné un arbre infini T, on note ∂T l'ensemble des *rayons* de T, c'est-à-dire des chemins auto-évitants infinis issus de la racine. Dans $\mathbb{T}_\lambda$, si $\gamma_1, \gamma_2 \in \partial\boldsymbol{\tau}_\lambda^g$ sont respectivement le bord gauche et le bord droit d'une même bande, ils ne correspondent pas vraiment à deux directions différentes (par exemple, on s'attend à ce qu'ils tendent vers le même point du bord sur la Figure 1.11). On note donc $\widehat{\partial} T$ le bord de T, où on a identifié deux rayons à chaque fois qu'il n'y a aucun autre rayon entre les deux. On peut alors munir $\mathbb{T}_\lambda \cup \widehat{\partial}\boldsymbol{\tau}_\lambda^g$ d'une topologie qui en fait un espace métrique compact, et pour laquelle $\widehat{\partial}\boldsymbol{\tau}_\lambda^g$ est homéomorphe au cercle [19].

Par conséquent, on peut voir $\widehat{\partial}\boldsymbol{\tau}_\lambda^g$ comme un bord de $\mathbb{T}_\lambda$. Il est alors assez naturel de comparer cette notion de bord avec d'autres, comme la frontière de Poisson définie dans la Section 1.3.2.

Théorème 7

Le bord $\widehat{\partial}\boldsymbol{\tau}_\lambda^g$ est une réalisation de la frontière de Poisson de $\mathbb{T}_\lambda$.

D'après un résultat de Tom Hutchcroft et Yuval Peres [83], la démonstration de ce résultat se ramène à montrer que la marche aléatoire simple (X_n) sur $\mathbb{T}_\lambda$ converge presque sûrement vers un point X_∞ de $\widehat{\partial}\boldsymbol{\tau}_\lambda^g$, et que la loi de X_∞ conditionnellement à $\mathbb{T}_\lambda$ est p.s. non-atomique. La preuve utilise des arguments d'épluchage. Plus précisément, on adapte un argument de [58] pour montrer que deux marches aléatoires indépendantes sont toujours séparées par plusieurs rayons de $\boldsymbol{\tau}_\lambda^g$.

Géodésiques infinies dans le plan brownien hyperbolique. Enfin, nous nous intéressons aux géodésiques infinies dans le plan brownien hyperbolique $\mathcal{P}^h$. Les résultats montrés par Jean-François Le Gall sur la confluence des géodésiques dans la carte brownienne [98] permettent de montrer que les géodésiques infinies (pas forcément les plus à gauche !) dans $\mathcal{P}^h$ forment un arbre continu, qu'on note $\boldsymbol{\tau}^g(\mathcal{P}^h)$. Dans le plan brownien, cet arbre consiste en une unique géodésique infinie [62]. Dans le cas hyperbolique, la situation est une nouvelle fois bien différente. D'après le Théorème 2, on peut s'attendre à ce que $\boldsymbol{\tau}^g(\mathcal{P}^h)$ soit une limite d'échelle quasi-critique des arbres $\boldsymbol{\tau}_\lambda^g$. Nous montrons que c'est effectivement le cas. Plus précisément, soit $\mathbf{B}$ l'arbre infini où tous les sommets ont exactement deux enfants, sauf la racine qui n'en a qu'un.

Théorème 8

L'arbre $\boldsymbol{\tau}^g(\mathcal{P}^h)$ a la loi de l'arbre $\mathbf{B}$ où les longueurs des arêtes sont des variables exponentielles i.i.d. de paramètre $2\sqrt{2}$.

La démonstration de ce résultat repose sur les Théorèmes 2 et 6. Soit (λ_n) une suite vérifiant les hypothèses du Théorème 2, c'est-à-dire

$$\lambda_n = \lambda_c \left(1 - \frac{2}{3n^4} + o\left(\frac{1}{n^4}\right)\right).$$

Il est alors assez facile de vérifier que les arbres $\frac{1}{n}\boldsymbol{\tau}_{\lambda_n}^g$ convergent vers l'arbre continu voulu. Cependant, ce n'est pas suffisant pour établir le Théorème 8. En effet, il faut s'assurer que deux cas problématiques ne se produisent pas :

(i) deux géodésiques infinies distinctes pourraient être trop proches, et fusionner dans la limite d'échelle,

19. Il est assez facile de définir une telle topologie sur $\boldsymbol{\tau}_\lambda^g \cup \partial\boldsymbol{\tau}_\lambda^g$, qui donne à $\partial\boldsymbol{\tau}_\lambda^g$ la topologie d'un ensemble de Cantor. Le quotientage transforme alors l'ensemble de Cantor en cercle en "fermant les trous", ce qui est assez naturel.

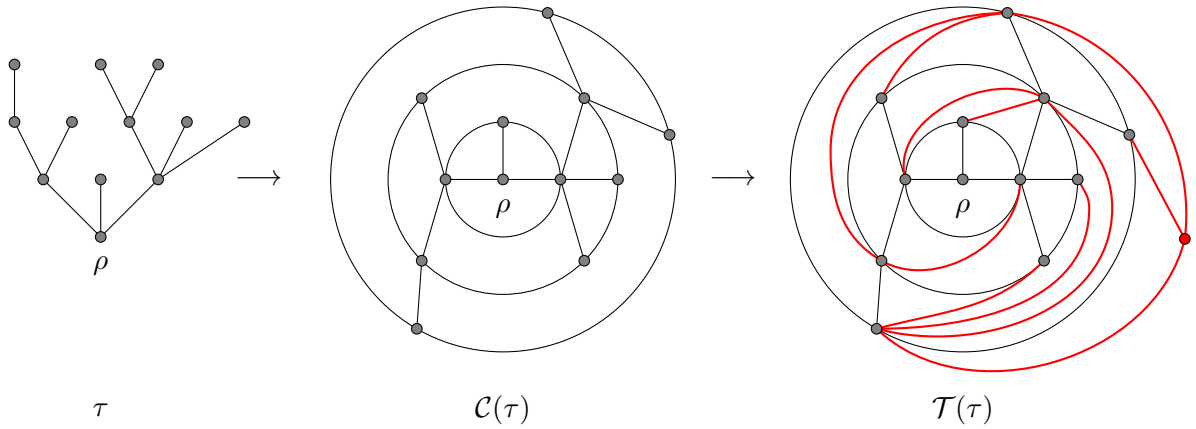

FIGURE 1.13 – De gauche à droite : un arbre plan τ, la carte causale associée à τ, et la triangulation causale associée à τ.

(ii) des chemins discrets proches d'être des géodésiques pourraient donner lieu, dans la limite d'échelle, à des géodésiques qui n'apparaissent pas au niveau discret.

1.4 Cartes causales

1.4.1 Définitions et motivations

Les triangulations causales ont été introduites par les physiciens théoriciens Jan Ambjørn et Renate Loll [10] pour décrire un espace-temps à deux dimensions où le temps et l'espace jouent des rôles asymétriques, contrairement aux triangulations usuelles de la sphère, où aucune des deux dimensions ne joue un rôle privilégié. Les triangulations causales sont des triangulations de la sphère particulières où les sommets sont répartis en couches successives, de telle manière que les sommets de chaque couche forment un cycle, et que tous les sommets de la i-ème couche se trouvent à distance exactement i de la racine.

Les triangulations causales à $n+1$ sommets sont alors en bijection avec les arbres plans à n sommets, et peuvent être construites de la manière suivante (voir aussi Figure 1.13). Étant donné un arbre plan τ, on ajoute à chaque hauteur un cycle reliant entre eux les sommets voisins. On obtient une carte planaire finie, qu'on appelle *carte causale associée à* τ, et qu'on note $\mathcal{C}(\tau)$. Puis on triangule les faces de $\mathcal{C}(\tau)$ en reliant, pour chaque face, le sommet en haut à droite de la face à tous les sommets de la face qui ne lui sont pas déjà adjacents. Enfin, on ajoute un sommet (en rouge sur la Figure 1.13) qu'on relie à tous les sommets de la dernière génération. La triangulation obtenue est appelée *triangulation causale associée à* τ. On s'attend à ce que la carte causale $\mathcal{C}(\tau)$ et la triangulation causale $\mathcal{T}(\tau)$ associées à un arbre τ aient des propriétés similaires. Dans la suite, on s'intéressera donc généralement à $\mathcal{C}(\tau)$, même si les preuves d'une bonne partie des résultats énoncés s'étendent aux triangulations causales.

Vers des dimensions supérieures ? Un autre intérêt des triangulations causales par rapport aux triangulations générales réside dans leur généralisation possible à des dimensions supérieures (le but final étant de modéliser un espace-temps de dimension autre que 2). L'analogue le plus naturel des triangulations aléatoires en dimension 3 serait de considérer des recollements de tétraèdres homéomorphes à la sphère de dimension 3. Un analogue des triangulations causales serait un modèle découpé en couches de dimension 2, où chaque couche est une triangulation

d'une surface et où des tétraèdres sont collés entre ces couches. Les simulations numériques réalisées par les physiciens semblent montrer que le premier modèle n'a pas de limite d'échelle intéressante : soit les distances entre les sommets sont d'ordre constant quand la taille tend vers l'infini ("*crumpled phase*"), soit l'objet obtenu ressemble à un arbre ("*branched polymer phase*"). Les limites d'échelle du second modèle semblent plus prometteuses. Voir par exemple la thèse de Timothy Budd [41, Introduction et Chapitre 4] pour plus de détails et de nombreuses références.

1.4.2 Propriétés des cartes causales critiques

Avant même de s'intéresser à des objets aléatoires, notons que la géométrie des cartes causales est très différente de celle des triangulations aléatoires uniformes. Par exemple, les boules de centre ρ et de rayon r dans $\mathcal{C}(\tau)$ et dans τ coïncident, donc le complémentaire des boules est connexe. Par conséquent, les "tentacules" (voir Figure 1.6), qui contenaient dans les triangulations uniformes l'essentiel de la masse, n'existent maintenant plus !

Cartes causales critiques infinies. On se fixe une loi de reproduction critique μ. Pour tout n, soit $\boldsymbol{\tau}_n$ un arbre de Galton–Watson critique de loi de reproduction μ, conditionné à avoir n sommets. À quoi ressemble la carte causale $\mathcal{C}(\boldsymbol{\tau}_n)$ quand n tend vers l'infini ? Comme pour les cartes aléatoires uniformes, une manière de répondre à cette question est d'étudier leur limite locale quand n tend vers l'infini. Cette limite locale est la carte $\mathcal{C}(\boldsymbol{\tau}_\infty)$, où $\boldsymbol{\tau}_\infty$ est un arbre aléatoire infini appelé *l'arbre de Galton–Watson de loi de reproduction μ conditionné à survivre*, introduit par Harry Kesten [88]. Cet arbre consiste essentiellement en une épine dorsale infinie, sur laquelle on a "branché" des arbres de Galton–Watson finis [20]. Deux exemples de telles cartes causales sont visibles sur la Figure 1.14.

Propriétés métriques. Bien que les cartes causales aient été étudiées de manière numérique par les physiciens, leur étude rigoureuse est très récente. On se concentrera dans la suite sur le cas où μ est à variance finie, soit $\sum_{i\geq 0} i^2\mu(i) < +\infty$. Certaines propriétés de $\mathcal{C}(\boldsymbol{\tau}_\infty)$ peuvent alors être déduites directement de celles de $\boldsymbol{\tau}_\infty$. Par exemple, le nombre de sommets à hauteur r dans $\boldsymbol{\tau}_\infty$ est de l'ordre de r, ce qui implique que la boule de rayon r dans $\mathcal{C}(\boldsymbol{\tau}_\infty)$ contient environ r^2 sommets. On peut ensuite se demander quel est l'ordre de grandeur des distances horizontales, c'est-à-dire entre sommets de la même hauteur r. Notons que ces distances valent au plus $2r$, car on peut redescendre à la racine et remonter. Nicolas Curien, Tom Hutchcroft et Asaf Nachmias ont montré [60] que ces distances horizontales sont en fait d'ordre $o(r)$, mais $r^{1-o(1)}$. Cela signifie qu'ajouter des cycles à chaque hauteur réduit beaucoup les distances horizontales, même si les distances verticales restent inchangés. La géométrie de $\mathcal{C}(\boldsymbol{\tau}_\infty)$ est donc très différente de celle de $\boldsymbol{\tau}_\infty$. En particulier, si on renormalise toutes les distances par r, tous les sommets à hauteur r sont à distance $o(1)$ les uns des autres, donc la limite d'échelle de $\mathcal{C}(\boldsymbol{\tau}_\infty)$ est une simple demi-droite.

Marche aléatoire simple. On peut également se demander quel est le comportement de la marche aléatoire simple sur $\mathcal{C}(\boldsymbol{\tau}_\infty)$. Pour tout $r \geq 1$, l'ensemble E_r des arêtes reliant un des sommets de hauteur r à son parent sépare la racine de l'infini. De plus, les ensembles E_r sont disjoints. Comme $|E_r|$ a pour ordre de grandeur r, on peut montrer que $\sum_{r\geq 1} \frac{1}{|E_r|} = +\infty$ p.s.. D'après le critère de Nash–Williams, ceci implique que $\mathcal{C}(\boldsymbol{\tau}_\infty)$ est p.s. récurrente. Curien, Hutchcroft et Nachmias ont également montré que la marche aléatoire simple est diffusive, c'est-à-dire qu'elle se trouve à hauteur d'ordre $\sqrt{n}$ après n pas. Du point de vue de la marche aléatoire, $\mathcal{C}(\boldsymbol{\tau}_\infty)$ est donc assez proche des réseaux réguliers comme le réseau carré. Une nouvelle fois, son

20. En particulier il n'a qu'une branche infinie, et a donc une géométrie très différente des arbres de Galton–Watson surcritiques.

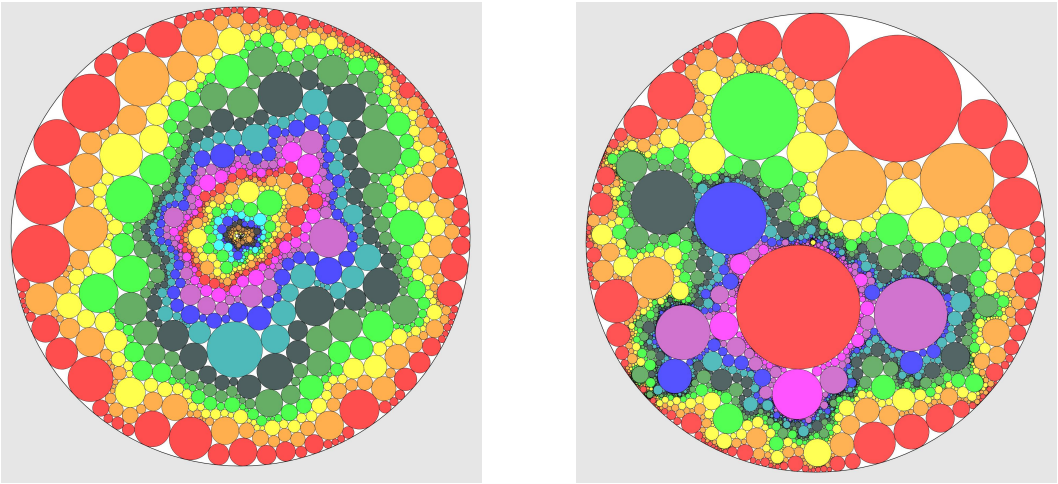

FIGURE 1.14 – Deux cartes causales obtenues à partir d'arbres de Galton–Watson critiques conditionnés à survivre. Dans celle de gauche, la loi de reproduction est à variance finie. Dans celle de droite, c'est une loi critique à queue lourde (dans le domaine d'attraction d'une loi stable d'indice 3/2). Les couleurs représentent les couches, c'est-à-dire les distances à la racine. Les deux cartes sont représentées par des empilements de cercles : chaque cercle représente un sommet, et deux cercles sont tangents si et seulement si les sommets correspondants sont reliés par une arête.

comportement est diffère de celui de $\boldsymbol{\tau}_\infty$ où la marche aléatoire, piégée dans de grands arbres finis, n'avance qu'à vitesse d'ordre $n^{1/3}$ [88].

Modèle à queue lourde. Enfin, certains résultats de [60] portent sur le cas critique à queue lourde, c'est-à-dire où μ est d'espérance 1 mais de variance infinie (cas de droite sur la Figure 1.14). Dans ce cas, la distance entre deux sommets typiques de hauteur r est d'ordre r : pour se déplacer horizontalement, la meilleure chose à faire est en quelque sorte de se déplacer verticalement. L'étude de ce modèle à queue lourde est motivée par la décomposition de Krikun de l'UIPT (voir Section 1.2.2) : le squelette de l'UIPT est essentiellement un arbre de Galton–Watson de loi de reproduction critique μ, avec $\mu(i) \sim ci^{-5/2}$. L'UIPT ressemble donc en quelque sorte à une carte causale, mais où on aurait inséré entre les branches de l'arbre plus que de simples arêtes (remplissages verts sur la Figure 1.8). Le fait que les branches de l'arbre soient les chemins les plus efficaces même pour se déplacer horizontalement évoque alors certaines propriétés des géodésiques dans la carte brownienne.

1.4.3 Cartes causales surcritiques (Chapitre 5, ou [47])

De même que pour les triangulations aléatoires uniformes, nous avons étudié un analogue hyperbolique de ces cartes causales. Le modèle est alors assez naturel : on considère un arbre de Galton–Watson surcritique conditionné à survivre[21], et on étudie les propriétés de la carte causale associée (voir Figure 1.15 pour un exemple). Pour toute cette section, on se fixe donc une loi de reproduction surcritique μ, et on note $\boldsymbol{\tau}$ un arbre de loi GW_μ, conditionné à survivre.

Motivations. Les arbres de Galton–Watson surcritiques sont aujourd'hui très bien compris, mais les preuves de nombreux résultats (en particulier sur la marche aléatoire, voir un peu plus

21. Cette fois, le conditionnement à survivre est non-dégénéré.

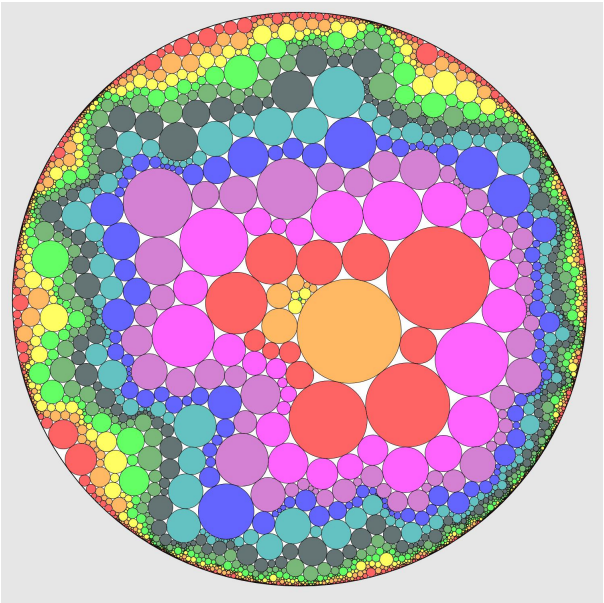

FIGURE 1.15 – La triangulation causale associée à un arbre de Galton–Watson surcritique, représentée par un empilement de cercles comme sur la Figure 1.14.

loin) reposent de manière très importante sur la structure d'arbre. Comme on va le voir, la carte $\mathcal{C}(\boldsymbol{\tau})$ partage de nombreuses propriétés avec l'arbre $\boldsymbol{\tau}$, ce qui montre que ces propriétés de $\boldsymbol{\tau}$ sont "robustes", et ne dépendent pas de la structure d'arbre.

Une autre motivation pour étudier des cartes causales est le fait qu'elles constituent un modèle-jouet de carte contenant un certain arbre. Ainsi, certains des résultats énoncés ci-dessous se généralisent à des modèles plus généraux, qui incluent les PSHIT étudiées dans la partie précédente. Plus précisément, rappelons qu'on appelle *bande* une carte planaire avec un bord infini (et une racine sur ce bord), et soit (s_i) une suite de telles bandes. Si on dessine un arbre infini τ dans le plan, il délimite des faces infinies. On note $\mathcal{M}(\tau, (s_i))$ la carte planaire obtenue en remplissant ces faces par les bandes s_i. Certains des résultats qui suivent porteront sur des cartes de la forme $\mathcal{M}(\boldsymbol{\tau}, (s_i))$, où on rappelle que $\boldsymbol{\tau}$ est un arbre de Galton–Watson surcritique. Le Théorème 8 montre alors que les PSHIT sont de cette forme, où les s_i sont indépendantes, et bordées par deux géodésiques infinies. La carte $\mathcal{C}(\boldsymbol{\tau})$ est également de cette forme, les bandes s_i étant déterministes et toutes égales.

Hyperbolicité métrique. Il est immédiat que toutes les cartes étudiées ont une croissance exponentielle. On va donc s'intéresser à des propriétés métriques plus fines. Une telle propriété est l'hyperbolicité de Gromov : on dit qu'un graphe G est Gromov-hyperbolique s'il existe une constante $k \geq 0$ telle que tous les triangles sont k-fins au sens suivant. Pour tous sommets x, y et z et toutes géodésiques γ_{xy}, γ_{yz} et γ_{zx} reliant ces trois sommets deux à deux, pour tout sommet v sur γ_{zx}, on a

$$d_G(v, \gamma_{xy} \cup \gamma_{yz}) \leq k.$$

En particulier, tout arbre est 0-hyperbolique. En revanche, $\mathbb{Z}^d$ pour $d \geq 2$ ne l'est pas. Comme c'est souvent le cas (voir par exemple la discussion sur l'expansion enracinée dans la Section 1.3.2), cette notion n'est pas adaptée à l'étude de graphes aléatoires [22]. Nous proposons donc

22. Par exemple, si $\mu(1) > 0$, alors $\mathcal{C}(\boldsymbol{\tau})$ contient des portions arbitrairement grandes de réseau carré, donc n'est pas Gromov-hyperbolique.

une notion plus faible.

Définition 1.7. On dit qu'une carte planaire M enracinée en ρ est *faiblement hyperbolique* s'il existe une constante $k \geq 0$ telle que, pour tous sommets x, y et z et toutes géodésiques γ_{xy}, γ_{yz} et γ_{zx} reliant ces trois sommets deux à deux, si le triangle formé par γ_{xy}, γ_{yz} et γ_{zx} entoure la racine, alors on a

$$d_M(\rho, \gamma_{xy} \cup \gamma_{yz} \cup \gamma_{zx}) \leq k.$$

Une autre propriété à laquelle on s'est intéressé est l'existence de géodésiques bi-infinies.

Définition 1.8. Une *géodésique bi-infinie* dans un graphe G est une famille $(\gamma(i))_{i\in\mathbb{Z}}$ telle que $d_G(\gamma(i), \gamma(j)) = |i - j|$ pour tous $i, j \in \mathbb{Z}$.

Contrairement aux géodésiques infinies, les géodésiques bi-infinies n'existent pas dans tous les graphes. Par exemple, les résultats de confluence des géodésiques de [67] impliquent qu'elles n'existent pas dans l'UIPT. De manière plus générale, même si de telles géodésiques peuvent exister dans des graphes non-hyperboliques comme $\mathbb{Z}^2$, on s'attend à ce qu'elles soient "instables", et disparaissent si la métrique est perturbée [23]. À l'inverse, dans les graphes Gromov-hyperboliques, un résultat d'Itai Benjamini et Romain Tessera montre que ces géodésiques existent et sont très stables [32].

Théorème 9

Toute carte de la forme $\mathcal{M}(\boldsymbol{\tau}, (s_i))$ est presque sûrement faiblement hyperbolique, et admet des géodésiques bi-infinies.

En particulier, ces deux propriétés sont vérifiées par les PSHIT. Signalons enfin qu'un résultat intermédiaire important dans notre démonstration est le suivant. Considérons deux rayons infinis (distincts et non-voisins) γ_1 et γ_2 de $\boldsymbol{\tau}$. Alors il existe k tel que pour tout $i \geq 0$, toute géodésique de $\gamma_1(i)$ vers $\gamma_2(i)$ passe à hauteur au plus k. Cela montre que la distance typique entre deux sommets de hauteur r est non seulement d'ordre r, mais même $2r - O(1)$, ce qui peut être vu comme une version forte des résultats de [60] pour le cas surcritique.

Frontière de Poisson. De même que pour les PSHIT, on s'est également intéressé à la marche aléatoire simple sur des cartes causales. Le fait que l'arbre $\boldsymbol{\tau}$ soit transient (c'est un résultat de Russell Lyons [106]) implique que toute carte planaire contenant $\boldsymbol{\tau}$ est transiente. On s'est donc intéressé à la frontière de Poisson des cartes causales (voir Section 1.3.2 pour la définition). Rappelons que $\widehat{\partial}\boldsymbol{\tau}$ est l'ensemble des rayons infinis de $\boldsymbol{\tau}$, quotienté par la relation d'équivalence "être le bord gauche et le bord droit d'une même bande". Ici aussi, on peut munir $\mathcal{C}(\boldsymbol{\tau}) \cup \widehat{\partial}\boldsymbol{\tau}$ d'une topologie qui en fait un espace compact et rend $\widehat{\partial}\boldsymbol{\tau}$ homéomorphe au cercle.

Théorème 10

- Presque sûrement, $\widehat{\partial}\boldsymbol{\tau}$ est une réalisation de la frontière de Poisson de $\mathcal{C}(\boldsymbol{\tau})$.
- Soient (S_i) des bandes i.i.d.. Alors presque sûrement, la carte $\mathcal{M}(\boldsymbol{\tau}, (S_i))$ ne vérifie pas la propriété de Liouville.

Par ailleurs, il est faux en général que si les (S_i) sont i.i.d., alors $\widehat{\partial}\boldsymbol{\tau}$ est une réalisation de la frontière de Poisson de $\mathcal{M}(\boldsymbol{\tau}, (S_i))$: si les S_i ont elles-même une frontière de Poisson non-triviale, alors celle de $\mathcal{M}(\boldsymbol{\tau}, (S_i))$ sera "plus grande" que $\widehat{\partial}\boldsymbol{\tau}$. Par conséquent, ce résultat ne donne pas la frontière de Poisson des PSHIT. C'est pourquoi nous avons utilisé des arguments différents dans le chapitre 4. Enfin, signalons que la démonstration du Théorème 10 repose sur des estimées de résistances électriques, qui peuvent également être intéressantes en tant que telles.

23. C'est un problème important et toujours ouvert en percolation de premier passage, voir par exemple [87].

Vitesse positive. En général, on s'attend dans un graphe hyperbolique à ce que la marche aléatoire simple s'éloigne très vite de l'origine. Par exemple, dans le cas $\mu(0) = 0$ (arbre de Galton–Watson sans feuille), Russell Lyons, Robin Pemantle et Yuval Peres ont montré [108] que la marche aléatoire (X_n) sur l'arbre $\boldsymbol{\tau}$ vérifie

$$\frac{d_{\boldsymbol{\tau}}(\rho, X_n)}{n} \xrightarrow[n\to+\infty]{p.s.} \sum_{i\geq 0} \frac{i-1}{i+1}\mu(i) > 0.$$

En particulier, la distance de X_n à la racine croît linéairement en n : on dit que la marche aléatoire simple sur $\boldsymbol{\tau}$ a une *vitesse positive*. Cependant, la preuve repose de manière très importante sur la structure d'arbre, et ne se généralise pas du tout aux cartes causales. On note maintenant X_n la marche aléatoire simple sur $\mathcal{C}(\boldsymbol{\tau})$.

Théorème 11

Supposons $\mu(0) = 0$. Alors il existe une constante $v_\mu > 0$ telle que

$$\frac{d_{\mathcal{C}(\boldsymbol{\tau})}(\rho, X_n)}{n} \xrightarrow[n\to+\infty]{p.s.} v_\mu.$$

Contrairement à certains des résultats précédents, ce résultat n'est pas du tout robuste : nous n'avons pu le montrer que pour la carte $\mathcal{C}(\boldsymbol{\tau})$, dans le cas où l'arbre sous-jacent n'a pas de feuille. Les preuves de vitesse positive pour les arbres de Galton–Watson [108] ou les PSHIT [58] utilisaient des outils de théorie ergodique reposant sur la stationnarité de la marche aléatoire (le graphe vu depuis X_n a la même loi que le graphe vu depuis X_0). Les cartes causales n'étant pas stationnaires, notre preuve utilise des méthodes assez différentes. Les deux principaux ingrédients sont les suivants.

- Une méthode d'exploration de $\mathcal{C}(\boldsymbol{\tau})$ simultanément avec la marche aléatoire permet de montrer que (X_n) ne visite pas de grande région où tous les sommets n'ont qu'un seul enfant chacun.
- On conclut en étudiant les *temps de renouvellement*, c'est-à-dire les instants où la marche aléatoire atteint une certaine hauteur pour la première fois, pour ne plus jamais redescendre en-dessous. Ces temps permettent de découper la marche aléatoire en blocs indépendants, et d'appliquer une loi des grands nombres. Ils ont été utilisés pour étudier plusieurs modèles comme la marche aléatoire en milieu aléatoire [130].

1.5 Perspectives

1.5.1 Sur les flips

La première question laissée ouverte sur les flips est bien sûr de trouver une borne supérieure sur le temps de mélange. Une borne supérieure polynomiale en n a très récemment été obtenue par Alessandra Caraceni et Alexandre Stauffer [51]. Plus précisément, la borne obtenue est d'ordre $n^{13/2}$. Trouver le bon exposant reste donc ouvert.

Par ailleurs, il est également possible de définir une dynamique de flips sur des triangulations infinies en munissant chaque arête d'un processus de Poisson, et en la flippant chaque fois que le processus saute. On peut montrer que l'UIPT ainsi que les PSHIT sont stationnaires pour cette dynamique. Il est alors naturel de se demander si les mélanges de PSHIT sont les seules triangulations stationnaires pour cette dynamique.

Cette dynamique sur des triangulations infinies soulève par ailleurs d'autres interrogations. Par exemple, peut-on trouver des propriétés presque sûres de l'UIPT pour lesquelles il existe des

temps exceptionnels auxquels la propriété n'est plus vérifiée ? La croissance en r^4 semble assez rigide, mais on peut penser par exemple à l'absence de composante infinie pour la percolation critique.

1.5.2 Sur les PSHIT

La principale question restant ouverte sur les PSHIT est la Conjecture 1.4 (dont la Conjecture B.1 est une version plus précise). Comme pour l'UIPT, une preuve de cette conjecture devrait reposer sur des résultats précis sur l'énumération des cartes de genre grand. Des relations de récurrence existent pour compter ces cartes [80], et il faudrait en extraire des asymptotiques quand la taille et le genre tendent simultanément vers l'infini.

Une autre question naturelle est la suivante : les PSHIT et le plan brownien hyperbolique admettent-ils, à l'instar de l'UIPT et du plan brownien, des constructions "à la Schaeffer" à partir d'arbres étiquetés ? Notons qu'avant l'introduction des PSHIT, une telle construction partant d'un arbre de Galton–Watson surcritique avait été proposée par Itai Benjamini [25], mais l'objet obtenu semble différent. Il y a des raisons de penser que de telles constructions n'existent pas forcément [24]. En revanche, nous pensons que les bandes délimitées par les géodésiques infinies que nous étudions dans le chapitre 4 peuvent être vues comme des limites de cartes browniennes, conditionnées à avoir un diamètre proportionnel à leur volume, puis découpées le long d'une géodésique.

Les modèles du demi-plan ne sont pas abordés dans le chapitre 3. On pourrait cependant renormaliser les triangulations du demi-plan $\mathbb{H}_\alpha$ de [18] avec $\alpha > \frac{2}{3}$ tout en faisant tendre α vers sa valeur critique $\frac{2}{3}$, et se demander si elles admettent une limite d'échelle quasi-critique qu'on appellerait le *demi-plan brownien hyperbolique*. Une approche naturelle serait alors d'adapter les arguments du chapitre 3 en utilisant le demi-plan brownien [22] au lieu du plan brownien. Cependant, cela nécessiterait une meilleure compréhension des périmètres et volumes des enveloppes dans le demi-plan brownien (analogue des résultats de [63]). Ces questions sont l'objet d'un travail en cours avec Armand Riera. L'outil naturel pour les étudier est la construction du demi-plan brownien proposée par Nicolas Curien et Alessandra Caraceni [50].

Enfin, une étude plus complète des cartes markoviennes biparties du plan serait très intéressante. En particulier, comme noté dans l'Appendice C, il est possible qu'il existe des analogues hyperboliques des cartes stables, c'est-à-dire des cartes hyperboliques possédant de grandes faces, et peut-être une géométrie différente des PSHIT.

1.5.3 Sur les cartes causales

Nous laissons dans le chapitre 5 un certain nombre de questions ouvertes sur les cartes causales surcritiques. Ces questions sont majoritairement liées au comportement de la marche aléatoire simple. La plus naturelle est de chercher à supprimer l'hypothèse $\mu(0) = 0$ dans le Théorème 1.4.3. De même que sur les arbres de Galton–Watson [109], on pourrait également étudier la marche biaisée sur ces cartes, où la probabilité d'aller vers le parent du sommet courant est différente de la probabilité d'aller vers un de ses enfants. En particulier, sur les arbres, si $\mu(0) > 0$, un drift trop fort vers le haut peut paradoxalement ralentir la marche en la "coinçant" dans de petits arbres finis. Nous pensons que ce phénomène disparaît dans les cartes causales.

Une étude plus précise de la mesure harmonique sur les cartes causales serait aussi intéressante. Par exemple, on peut la comparer à la mesure limite quand $n \to +\infty$ de la mesure

24. Dans les bijections entre arbres et cartes, l'arbre correspond usuellement au *cut-locus* de la carte, c'est-à-dire à l'ensemble des points admettant plusieurs géodésiques vers la racine (qui peut être à l'infini). Pour "remplir" la carte, ce cut-locus devrait croiser les géodésiques infinies que nous étudions dans le chapitre 4, ce qui est impossible d'après les résultats de [98].

uniforme sur l'ensemble des sommets de la génération n. Il est montré dans [108] que pour la marche sur l'arbre, ces deux mesures sont singulières, ce qui signifie que la marche aléatoire se concentre en fait sur une "petite" partie de l'arbre. Nous nous attendons à ce que ce phénomène se produise également dans les cartes causales.

Une autre quantité liée à la marche aléatoire qu'il peut être intéressant d'étudier est la vitesse de décroissance de la probabilité p_n de retour en 0 au temps n. Cette quantité décroît comme $\exp\left(-n^{1/3}\right)$ dans les arbres de Galton–Watson surcritiques et dans les versions semi-planaires des PSHIT [17]. On s'attend à un comportement similaire dans les cartes causales surcritiques (voir la dernière section du chapitre 5). Nos arguments dans la preuve du Théorème 1.4.3 permettent au mieux de montrer que p_n tend vers 0 plus vite que n'importe quel polynôme, ce qui reste un résultat assez faible. Les preuves pour les modèles précédents utilisent la propriété d'*expansion enracinée*, qui ne semble cependant pas évidente à montrer chez les cartes causales.

Chapter 2

On the mixing time of the flip walk on triangulations of the sphere

ou *Une suite de flips.*

This chapter is adapted from [44], published in Comptes-Rendus Mathématiques de l'Académie des Sciences.

A simple way to sample a uniform triangulation of the sphere with a fixed number n of vertices is a Monte-Carlo method: we start from an arbitrary triangulation and flip repeatedly a uniformly chosen edge. We give a lower bound of order $n^{5/4}$ on the mixing time of this Markov chain.

Contents

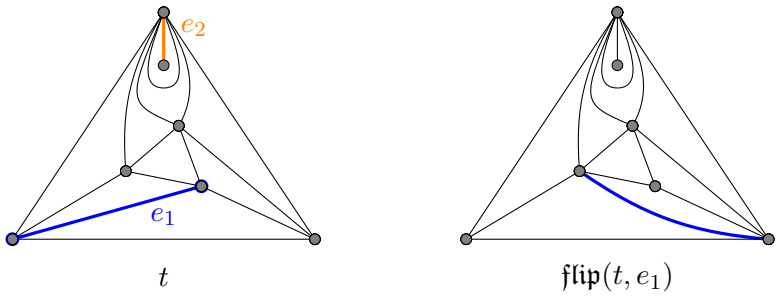

Figure 2.1 – An example of flip of an edge. The orange edge e_2 is not flippable.

2.1 Introduction

Much attention has been given recently to the study of large uniform triangulations of the sphere. Historically, these triangulations have been first considered by physicists as a discrete model for quantum gravity. Before the introduction of more direct tools (bijection with trees or peeling process), the first simulations [84, 86] were made using a Monte-Carlo method based on flips of triangulations.

More precisely, for all $n \geq 3$, let $\mathscr{T}_n$ be the set of rooted type-I triangulations of the sphere with n vertices (that is, triangulations that may contain loops and multiple edges, equipped with a distinguished oriented edge). If t is a triangulation we write $V(t)$ for the set of its vertices and $E(t)$ for the set of its edges. If $t \in \mathscr{T}_n$ and $e \in E(t)$, we write $\mathfrak{flip}(t, e)$ for the triangulation obtained by removing the edge e from t and drawing the other diagonal of the face of degree 4 that appears. We say that $\mathfrak{flip}(t, e)$ is obtained from t by *flipping* the edge e (cf. Figure 2.1). Note that it is possible to flip a loop and to flip the root edge. The only case in which an edge cannot be flipped is if both of its sides are adjacent to the same face like the edge e_2 on Figure 2.1. In this case $\mathfrak{flip}(t, e) = t$. Note that there is a natural bijection between $E(t)$ and $E(\mathfrak{flip}(t, e))$. When there is no ambiguity, we shall sometimes treat an element of one of these two sets as if it belonged to the other.

The graph of triangulations of the sphere in which two triangulations are related if one can pass from one to the other by flipping an edge has already been studied in the type-III setting (that is, triangulations with neither loops nor multiple edges): it is connected [138] and its diameter is linear in n [91]. We extend these results to our setup in Lemma 2.12.

We define a Markov chain $(T_n(k))_{k\geq 0}$ on $\mathscr{T}_n$ as follows: conditionally on $(T_n(0), \dots, T_n(k))$, let e_k be a uniformly chosen edge of $T_n(k)$. We take $T_n(k+1) = \mathfrak{flip}(T_n(k), e_k)$. It is easy to see that the uniform measure on $\mathscr{T}_n$ is reversible, thus stationary for $(T_n(k))_{k\geq 0}$, so this Markov chain will converge to the uniform distribution (the irreducibility is guaranteed by the connectedness results described above and the aperiodicity by the possible existence of non flippable edges). It is then natural to estimate the mixing time of $(T_n(k))_{k\geq 0}$ (see Chapter 4.5 of [104] for a proper definition of the mixing time). Our theorem provides a lower bound.

Theorem 2.1. There is a constant $c > 0$ such that for all $n \geq 3$ the mixing time of the Markov chain $(T_n(k))_{k\geq 0}$ is at least $cn^{5/4}$.

Mixing times for other types of flip chains have also been investigated. For triangulations of a convex n-gon without inner vertices it is known that the mixing time is polynomial and at least of order $n^{3/2}$ (see [113, 121]). In particular, our proof was partly inspired by the proof of the lower bound in [121]. Finally, see [49] for estimates on the mixing time of the flip walk on *lattice triangulations*, that is, triangulations whose vertices are points on a lattice and with Boltzmann weights depending on the total length of their edges.

The strategy of our proof is as follows: we start with two independent uniform triangulations with a boundary of length 1 and $\frac{n}{2}$ inner vertices and glue them together along their boundaries. We obtain a triangulation of the sphere with a cycle of length 1 such that half of the vertices lie on each side of this cycle. We then start our Markov chain from this triangulation and discover one of the two sides of the cycle gradually by a peeling procedure. By using the estimates of Curien and Le Gall [64] and a result of Krikun about separating cycles in the UIPT [94], we show that after $o(n^{5/4})$ flips, with high probability, the triangulation still has a cycle of length $o(n^{1/4})$, on each side of which lie a proportion at least $\frac{1}{4}$ of the vertices. But by a result of Le Gall and Paulin [103], this is not the case in a uniform triangulation (this is the discrete counterpart of the homeomorphicity of the Brownian map to the sphere), which shows that a time $o(n^{5/4})$ is not enough to approach the uniform distribution.

Acknowledgements: I thank Nicolas Curien for carefully reading earlier versions of this manuscript. I also thank the anonymous referee for his useful comments. I acknowledge the support of ANR Liouville (ANR-15-CE40-0013) and ANR GRAAL (ANR-14-CE25-0014).

2.2 Combinatorial preliminaries and couplings

For all $n \geq 3$, we recall that $\mathscr{T}_n$ is the set of rooted type-I triangulations of the sphere with n vertices. For $n \geq 0$ and $p \geq 1$ we also write $\mathscr{T}_{n,p}$ for the set of triangulations with a boundary of length p and n inner vertices, that is, planar maps with $n+p$ vertices in which all faces are triangles except one called the *outer face* whose boundary is a simple cycle of length p, equipped with a root edge such that the outer face touches the root edge on its right. We will sometimes refer to n and p as the *volume* and the *perimeter* of the triangulation.

The number of triangulations with fixed volume and perimeter can be computed by a result of Krikun. Here is a special case of the main theorem of [95] (the full theorem deals with triangulations with $r+1$ boundaries but we only use the case $r=0$):

$$\#\mathscr{T}_{n,p} = \frac{p(2p)!}{(p!)^2}\frac{4^{n-1}(2p+3n-5)!!}{n!(2p+n-1)!!} \underset{n\to+\infty}{\sim} C(p)\lambda_c^{-n}n^{-5/2}, \tag{2.1}$$

where $\lambda_c = \frac{1}{12\sqrt{3}}$ and $C(p) = \frac{3^{p-2}p(2p)!}{4\sqrt{2\pi}(p!)^2}$. In particular, a triangulation of the sphere with n vertices is equivalent after a root transformation to a triangulation with a boundary of length 1 and $n-1$ inner vertices (more precisely we need to duplicate the root edge, add a loop inbetween and root the map at this new loop, see for example Figure 2 in [65]), so

$$\#\mathscr{T}_n = \#\mathscr{T}_{n-1,1} = 2\frac{4^{n-2}\,(3n-6)!!}{(n-1)!\,n!!}. \tag{2.2}$$

For $n \geq 0$ and $p \geq 1$ we write $T_{n,p}$ for a uniform triangulation with a boundary of length p and n inner vertices, and T_n for a uniform triangulation of the sphere with n vertices. We also recall that the UIPT, that we write T_∞, is an infinite rooted planar triangulation whose distribution is characterized by the following equality. For any rooted triangulation t with a hole of perimeter p,

$$\mathbb{P}\left(t \subset T_\infty\right) = C(p)\lambda_c^{|t|}, \tag{2.3}$$

where λ_c and the $C(p)$ are as above, $|t|$ is the total number of vertices of t and by $t \subset T_\infty$ we mean that T_∞ can be obtained by filling the hole of t with an infinite triangulation with a boundary of length p.

In what follows we will use several times peeling explorations of random triangulations, see section 4.1 of [64] for a general definition. Let t be a triangulation and $\mathscr{A}$ be a peeling algorithm,

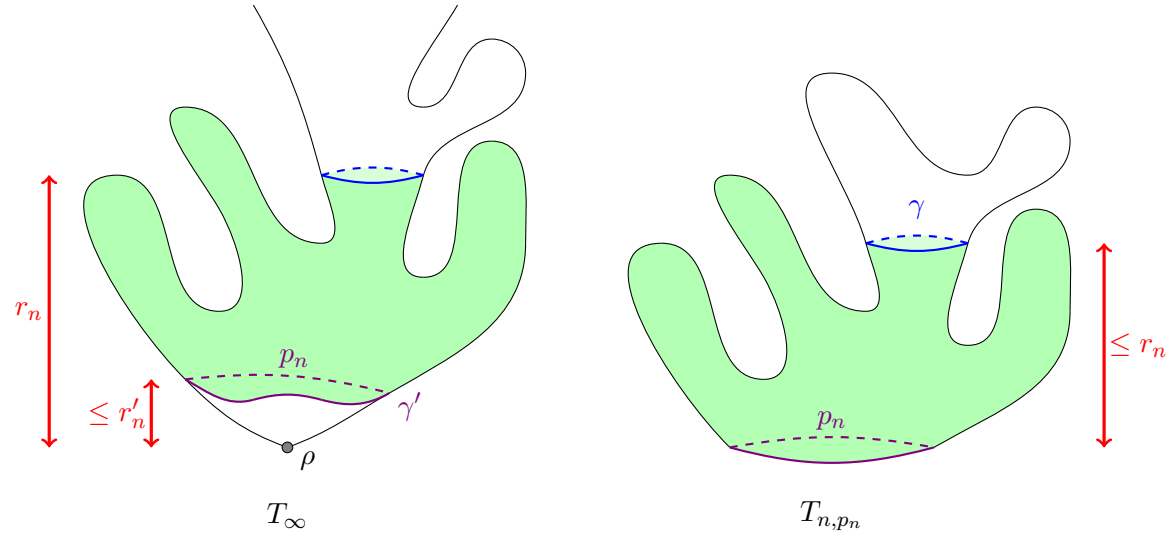

Figure 2.2 – Illustration of Lemma 2.1. With high probability, there are two cycles γ' and γ such that the two green parts coincide.

that is, a way to assign to every finite triangulation with one hole an edge on the boundary of the hole. We write $t_j^{\mathscr{A}}(t)$ for the part of t discovered after j steps of filled-in peeling following algorithm $\mathscr{A}$. By "filled-in" we mean that everytime the peeled face separates the unknown part of the map in two connected components we reveal the one with fewer vertices (if the two components have the same number of vertices we reveal one component picked deterministically). If the map is infinite and one-ended, we reveal the bounded component.

From the enumeration formulas it is possible to deduce precise coupling results between finite and infinite maps. The result we will need is similar to Proposition 12 of Chapter 3 but a bit more general since it deals with triangulations with a boundary. We recall that in a triangulation t of the sphere or the plane, the *ball* of radius r, that we write $B_r(t)$, is the triangulation with holes formed by those faces adjacent to at least one vertex lying at distance at most $r-1$ from the root, along with all their edges and vertices. If t is infinite, the *hull* of radius r, that we write $B_r^\bullet(t)$, is the union of $B_r(t)$ and all the bounded connected components of its complement. If t is finite, it is the union of $B_r(t)$ and all the connected components of its complement except the one that contains the most vertices (if there is a tie, we pick deterministically a component among those which contain the most vertices). If T is a triangulation with a boundary, we adopt the same definitions but we replace the distance to the root by the distance to the boundary.

Lemma 2.1. Let $p_n = o(\sqrt{n})$ and $r_n = o(n^{1/4})$ with $p_n = o(r_n^2)$. Then there are $r'_n = o(r_n)$ and couplings between T_{n,p_n} and T_∞ such that

$$\mathbb{P}\left(B^\bullet_{r_n}(T_\infty) \backslash B^\bullet_{r'_n}(T_\infty) \subset B^\bullet_{r_n}(T_{n,p_n})\right) \xrightarrow[n\to+\infty]{} 1.$$

The above lemma follows from the following. There is a cycle γ' of length p_n around the root of T_∞ that lies inside of its hull of radius r'_n and a cycle γ in T_{n,p_n} that stays at distance at most r_n from its boundary, such that the part of the hull of radius r_n of T_∞ that lies outside of γ' is isomorphic to the part of T_{n,p_n} that lies between its boundary and γ (see Figure 2.2).

Proof. We start by describing a coupling between the UIPT and the UIPT with a boundary of length p_n, that we write T_{∞,p_n}. We consider the peeling by layers $\mathscr{L}$ of the UIPT (see section 4.1 of [64]) and we write τ_{p_n} for the first time at which the perimeter of the discovered region is

equal to p_n (note that this time is always finite since the perimeter can increase by at most 1 at each peeling step). By the spatial Markov property of the UIPT, the part that is still unknown at time τ_{p_n} has the distribution of T_{∞,p_n}. Moreover, by the results of Curien and Le Gall (Theorem 1 of [64]), since $p_n = o(r_n^2)$, we have $\tau_{p_n} = o(r_n^3)$. By using Proposition 9 of [64] (more precisely the convergence of H), we obtain that the smallest hull of T_∞ containing $t^{\mathscr{L}}_{\tau_{p_n}}(T_\infty)$ has radius $o(r_n)$ in probability. Hence, our result holds if we replace T_{n,p_n} by T_{∞,p_n}.

Hence, it is enough to prove that there are couplings between T_{∞,p_n} and T_{n,p_n} such that

$$\mathbb{P}\left(B^\bullet_{r_n}(T_{n,p_n}) = B^\bullet_{r_n}(T_{\infty,p_n})\right) \xrightarrow[n\to+\infty]{} 1.$$

The proof relies on asymptotic enumeration results and is essentially the same as that of Proposition 12 of Chapter 3: by using the above coupling of T_{∞,p_n} and T_∞ we can show that

$$\left(\frac{1}{\sqrt{n}}|\partial B^\bullet_{r_n}(T_{\infty,p_n})|, \frac{1}{n}|B^\bullet_{r_n}(T_{\infty,p_n})|\right) \xrightarrow[n\to+\infty]{(P)} (0,0).$$

Moreover, if $q_n = o(\sqrt{n})$ and $v_n = o(n)$ and if t_n is a triangulation with two holes of perimeters p_n and q_n (rooted on the boundary of the p_n-gon) and v_n vertices that is a possible value of $B^\bullet_{r_n}(T_{\infty,p_n})$ for all $n \geq 0$, then

$$\frac{\mathbb{P}\left(B^\bullet_{r_n}(T_{n,p_n}) = t_n\right)}{\mathbb{P}\left(B^\bullet_{r_n}(T_{\infty,p_n}) = t_n\right)} \xrightarrow[n\to+\infty]{} 1$$

by the enumeration results, and we can conclude as in Proposition 12 of Chapter 3. □

We will also need another coupling lemma where we do not compare hulls of a fixed radius, but rather the parts of triangulations that have been discover after a fixed number of peeling steps.

Lemma 2.2. Let $j_n = o(n^{3/4})$, and let $\mathscr{A}$ be a peeling algorithm. Then there are couplings between T_n and T_∞ such that

$$\mathbb{P}\left(t^{\mathscr{A}}_{j_n}(T_n) = t^{\mathscr{A}}_{j_n}(T_\infty)\right) \xrightarrow[n\to+\infty]{} 1.$$

Proof. We write $P_\infty(j)$ and $V_\infty(j)$ for respectively the perimeter and volume of $t^{\mathscr{A}}_j(T_\infty)$. By the results of [64] we have the convergences

$$\frac{1}{\sqrt{n}} \sup_{0\leq j\leq j_n} P_\infty(j) \xrightarrow[n\to+\infty]{} 0 \quad \text{and} \quad \frac{1}{n} \sup_{0\leq j\leq j_n} V_\infty(j) \xrightarrow[n\to+\infty]{} 0 \tag{2.4}$$

in probability, so there are $p_n = o(\sqrt{n})$ and $v_n = o(n)$ such that

$$\mathbb{P}\left(P_\infty(j_n) \leq p_n \text{ and } V_\infty(j_n) \leq v_n\right) \to 1.$$

But by the enumeration results (2.1), (2.2) and by (2.3), if t_n is a rooted triangulation with perimeter at most p_n and volume at most v_n, we have

$$\frac{\mathbb{P}\left(t^{\mathscr{A}}_{j_n}(T_n) = t_n\right)}{\mathbb{P}\left(t^{\mathscr{A}}_{j_n}(T_\infty) = t_n\right)} = \frac{\mathbb{P}\big(t_n \subset T_n\big)}{\mathbb{P}\big(t_n \subset T_\infty\big)} \xrightarrow[n\to+\infty]{} 1.$$

As in Proposition 12 of Chapter 3, this proves that the total variation distance between the distributions of $t^{\mathscr{A}}_{j_n}(T_n)$ and $t^{\mathscr{A}}_{j_n}(T_\infty)$ goes to 0 as $n \to +\infty$, which proves our claim and the lemma. □

By combining this last lemma and the estimates (2.4), we immediately obtain estimates about the peeling process on finite uniform triangulations. We write $P_n(j)$ and $V_n(j)$ for the perimeter and volume of $t_j^{\mathscr{A}}(T_n)$.

Corollary 2.3. Let $j_n = o(n^{3/4})$. Then we have the following convergences in probability:

$$\frac{1}{\sqrt{n}} \sup_{0\leq j\leq j_n} P_n(j) \xrightarrow[n\to+\infty]{} 0 \quad \text{and} \quad \frac{1}{n} \sup_{0\leq j\leq j_n} V_n(j) \xrightarrow[n\to+\infty]{} 0.$$

Finally, we show a result about small cycles surrounding the boundary in uniform triangulations with a perimeter small enough compared to their volume.

Lemma 2.4. Let $p_n = o(\sqrt{n})$ and $r_n = o(n^{1/4})$ be such that $p_n = o(r_n^2)$. Then for all $\varepsilon > 0$, the probability of the event

"there is a cycle γ in T_{n,p_n} of length at most r_n such that the part of T_{n,p_n} lying between $\partial T_{n,p_n}$ and γ contains at most εn vertices"

goes to 1 as $n \to +\infty$.

This result is not surprising. In the context of quadrangulations with a non-simple boundary, it is a consequence of the convergence of quadrangulations with boundaries to Brownian disks, see [37]. However, no scaling limit result is known yet for triangulations with boundaries. Hence, we will rely on a result of Krikun about small cycles in the UIPT, that we will combine with Lemma 2.1. Here is a restatement of Theorem 6 of [94].

Theorem 2.2 (Krikun)**.** For all $\varepsilon > 0$, there is a constant C such that for all r, with probability at least $1 - \varepsilon$ there is a cycle of length at most Cr surrounding $B_r^\bullet(T_\infty)$ and lying in $B_{2r}^\bullet(T_\infty)$.

Note that Krikun deals with type-II triangulations, i.e. with multiple edges but no loops, but the decomposition used in [94] is still valid and even a bit simpler in the type-I setting, see [65]. The fact that the cycle stays in $B_{2r}^\bullet(T_\infty)$ is not in the statement of the theorem in [94] but it is immediate from its proof.

Proof of Lemma 2.4. By Lemma 2.1 it is possible to couple T_∞ and T_{n,p_n} in such a way that

$$\mathbb{P}\left(B_{r_n}^\bullet(T_\infty)\backslash B_{r_n'}^\bullet(T_\infty) \subset B_{r_n}^\bullet(T_{n,p_n})\right) \xrightarrow[n\to+\infty]{} 1, \tag{2.5}$$

where $r_n' = o(r_n)$. On the other hand, by Theorem 2.2, we have

$$\mathbb{P}\left(\text{there is a cycle } \gamma \text{ of length} \leq r_n \text{ in } B_{2r_n'}^\bullet(T_\infty) \text{ that surrounds } B_{r_n'}^\bullet(T_\infty)\right) \xrightarrow[n\to+\infty]{} 1.$$

For n large enough we have $r_n \geq 2r_n'$ so if such a γ exists in then it must stay in $B_{r_n}^\bullet(T_\infty)$. Since $r_n = o(n^{1/4})$, the probability that the number of vertices lying inside of γ is greater than εn goes to 0 by Theorem 2 of [64]. But if the event of (2.5) holds and if such a cycle exists in T_∞, then in T_{n,p_n} there is a cycle γ of length at most r_n such that the part of T_{n,p_n} lying between $\partial T_{n,p_n}$ and γ contains at most εn vertices. □

2.3 Proof of Theorem 2.1

Our main task will be to prove the following proposition.

Proposition 2.5. Let $k_n = o(n^{5/4})$. Then there are $t_n \in \mathscr{T}_n$ and $\ell_n = o(n^{1/4})$ such that conditionally on $T_n(0) = t_n$, the probability that there is a cycle of length at most ℓ_n that separates $T_n(k_n)$ in two parts of volume at least $\frac{n}{4}$ goes to 1 as $n \to +\infty$.

We first define the initial triangulation $T_n(0)$ we will be interested in: let $T_n^1(0)$ and $T_n^2(0)$ be two independent uniform triangulations with a boundary of length 1 and with respectively $\lfloor \frac{n-1}{2} \rfloor$ and $\lceil \frac{n-1}{2} \rceil$ inner vertices. We write $T_n(0)$ for the triangulation obtained by gluing together the boundaries of $T_n^1(0)$ and $T_n^2(0)$.

We will now perform an exploration of the triangulation while it gets flipped: the part T_n^1 will be considered as the "discovered" part and T_n^2 as the "unknown" part of the map. More precisely, we define by induction $T_n^1(k)$ and $T_n^2(k)$ such that $T_n(k)$ is obtained by gluing together the boundaries of $T_n^1(k)$ and $T_n^2(k)$. The two triangulations for $k = 0$ are defined above. Now assume we have constructed $T_n^1(k)$ and $T_n^2(k)$. Then:

— if e_k lies inside of $T_n^1(k)$ then $T_n^1(k+1) = \mathfrak{flip}(T_n^1(k), e_k)$ and $T_n^2(k+1) = T_n^2(k)$,
— if e_k lies inside of $T_n^2(k)$ then $T_n^1(k+1) = T_n^1(k)$ and $T_n^2(k+1) = \mathfrak{flip}(T_n^2(k), e_k)$,
— if $e_k \in \partial T_n^1(k)$, we write f_k for the face of $T_n^2(k)$ that is adjacent to e_k, and we let $T_n^2(k+1)$ be the connected component of $T_n^2(k) \backslash f_k$ with the largest volume and $T_n^1(k+1) = \mathfrak{flip}\big(T_n(k) \backslash T_n^2(k+1), e_k\big)$.

We now set $\widetilde{P}_n(k) = |\partial T_n^1(k)|$ and $\widetilde{V}_n(k) = \big|V\left(T_n^1(k)\right)\big| - \big|V\left(T_n^1(0)\right)\big| + 1$. Note that $\widetilde{V}_n(k)$ is nondecreasing in k.

For $k \geq 0$, we define a random variable $e_k^* \in E(T_n^1(k)) \cup \{\star\}$, where $\star$ is an additional state corresponding to all the edges not in $E(T_n^1(k))$, as follows: if e_k lies inside or on the boundary of $T_n^1(k)$ then $e_k^* = e_k$, and if not then $e_k^* = \star$. We also define $\mathcal{F}_k$ as the σ-algebra generated by the variables $\big(T_n^1(i)\big)_{0 \leq i \leq k}$ and $(e_i^*)_{0 \leq i \leq k-1}$.

Lemma 2.6. For all k, conditionally on $\mathcal{F}_k$, the triangulation $T_n^2(k)$ is a uniform triangulation with a boundary of length $\widetilde{P}_n(k)$ and $\lceil \frac{n+1}{2} \rceil - \widetilde{V}_n(k)$ inner vertices.

Proof. We prove the lemma by induction on k. For $k = 0$ it is obvious by the definition of $T_n^2(0)$. Let $k \geq 0$ be such that the lemma holds for k.

- If e_k^* lies inside $T_n^1(k)$, the result follows from the fact that $T_n^2(k) = T_n^2(k+1)$ and that conditionally on $\mathcal{F}_k$, the triangulation $T_n^2(k)$ is independent of e_k^*.
- If $e_k^* = \star$, it follows from the invariance of the uniform measure on $\mathscr{T}_{n,p}$ under flipping of a uniform edge among those which do not lie on the boundary.
- If $e_k^* \in \partial T_n^1(k)$, this is a standard peeling step: by invariance under rerooting of a uniform triangulation with fixed perimeter and volume, conditionally on $\mathcal{F}_k$ and e_k, the triangulation $T_n^2(k)$ rooted at e_k is uniform. Hence, if the third vertex of the face f_k of $T_n^2(k)$ adjacent to e_k lies inside of $T_n^2(k)$, the remaining part of $T_n^2(k)$ is a uniform triangulation with a boundary of length $\widetilde{P}_n(k) + 1$ and $\lceil \frac{n+1}{2} \rceil - \widetilde{V}_n(k) - 1$ inner vertices. If the third vertex of f_k lies on $\partial T_n^2(k)$, then the face f_k separates $T_n^2(k)$ in two independent uniform triangulations with fixed perimeters and volumes, and the lemma follows.

□

We now define the stopping times τ_j as the times at which the flipped edge lies on the boundary of the unknown part of the map, that is, the times k at which we discover new parts of $T_n^2(k)$: we set $\tau_0 = 0$ and $\tau_{j+1} = \inf\{k > \tau_j | e_k \in \partial T_n^1(k)\}$ for $j \geq 0$. We also write $P_n(j) = \widetilde{P}_n(\tau_j + 1)$ and $V_n(j) = \widetilde{V}_n(\tau_j + 1)$.

Then Lemma 2.6 shows that (P_n, V_n) is a Markov chain with the same transitions as the perimeter and volume processes associated to the peeling process of a uniform triangulation with a boundary of length 1 and $\lceil \frac{n-1}{2} \rceil$ inner vertices. Hence, Corollary 2.3 provides estimates for this process. Our next lemma will allow us to estimate the times τ_j.

Lemma 2.7. Let $k_n = o(n^{5/4})$. Then for all $\varepsilon > 0$ we have

$$\mathbb{P}\left(\tau_{\varepsilon n^{3/4}} > k_n\right) \xrightarrow[n\to+\infty]{} 1.$$

Proof. Conditionally on P_n, the variables $\tau_{j+1} - \tau_j$ are independent geometric variables with respective parameters $\frac{P_n(j)}{n}$. Hence, $\tau_{\varepsilon n^{3/4}}$ dominates the sum S_n of $\varepsilon n^{3/4}$ i.i.d. geometric variables with parameter $Q_n = \frac{1}{n}\max_{0\le j\le \varepsilon n^{3/4}} P_n(j)$. We have

$$\mathbb{E}[S_n|P_n] = \varepsilon n^{3/4} Q_n = \varepsilon n^{5/4} \times \frac{1}{\sqrt{n}} \max_{0\le j\le \varepsilon n^{3/4}} P_n(j).$$

By the results of [64], the factor $\frac{1}{\sqrt{n}}\max_{0\le j\le \varepsilon n^{3/4}} P_n(j)$ converges in distribution, so $\frac{\mathbb{E}[S_n|P_n]}{\varepsilon n^{5/4}}$ converges in distribution so $\frac{\mathbb{E}[S_n|P_n]}{k_n} \to +\infty$ in probability. By the weak law of large numbers we get $\frac{S_n}{k_n} \to +\infty$ in probability so $\frac{\tau_{\varepsilon n^{3/4}}}{k_n} \to +\infty$ in probability. □

By combining Corollary 2.3 and Lemma 2.7 we get the following result.

Lemma 2.8. Let $k_n = o(n^{5/4})$. Then we have the convergences

$$\frac{1}{\sqrt{n}}\widetilde{P}_n(k_n) \xrightarrow[n\to+\infty]{} 0 \quad \text{and} \quad \frac{1}{n}\widetilde{V}_n(k_n) \xrightarrow[n\to+\infty]{} 0$$

in probability.

Proof. By Lemma 2.7 there is a deterministic sequence $j_n = o(n^{3/4})$ such that $\mathbb{P}\left(\tau_{j_n} > k_n\right) \to 1$. This means that with probability going to 1 as $n \to +\infty$ there is $J \le j_n$ such that $\tau_J < k_n \le \tau_{J+1}$ so

$$\widetilde{P}_n(k_n) = P_n(J) \le \sup_{0\le j\le j_n} P_n(j) \quad \text{and} \quad \widetilde{V}_n(k_n) = V_n(J) \le \sup_{0\le j\le j_n} V_n(j).$$

But we know from Corollary 2.3 that

$$\left(\frac{1}{\sqrt{n}}\sup_{0\le j\le j_n} P_n(j), \frac{1}{n}\sup_{0\le j\le j_n} V_n(j)\right) \xrightarrow[n\to+\infty]{(P)} 0,$$

which proves Lemma 2.8. □

So $T_n^2(k_n)$ has the distribution of $T_{n/2-\widetilde{V}_n(k_n),\widetilde{P}_n(k_n)}$ and there is $p_n = o(\sqrt{n})$ such that

$$\mathbb{P}\left(\widetilde{P}_n(k_n) < p_n \text{ and } n/2 - \widetilde{V}_n(k_n) > \frac{n}{3}\right) \xrightarrow[n\to+\infty]{} 1.$$

Let r_n be such that $r_n = o(n^{1/4})$ and $p_n = o(r_n^2)$ (take for example $r_n = n^{1/8}p_n^{1/4}$). By Lemma 2.4, with probability going to 1 as $n \to +\infty$, there is a cycle γ in $T_n^2(k_n)$ of length at most r_n such that the part of $T_n^2(k_n)$ lying between $\partial T_n^2(k_n)$ and γ has volume at most $\frac{n}{6}$. Moreover we have $\widetilde{V}_n(k_n) = o(n)$ in probability by Lemma 2.8, so the two parts of $T_n(k_n)$ separated by γ both have volume at least $\frac{n}{4}$, which proves Proposition 2.5.

The proof of our main theorem is now easy: let $\mathscr{T}_n^{\bowtie}$ be the set of the triangulations t of the sphere with n vertices in which there is a cycle of length at most ℓ_n that separates t in two parts of volume at least $\frac{n}{4}$. Let also $k_n = o(n^{5/4})$. By Proposition 2.5 we have

$$\mathbb{P}\left(T_n(k_n) \in \mathscr{T}_n^{\bowtie}\right) \xrightarrow[n\to+\infty]{} 1,$$

whereas by Corollary 1.2 of [103], if $T_n(\infty)$ denotes a uniform variable on $\mathscr{T}_n$ we have

$$\mathbb{P}\left(T_n(\infty) \in \mathscr{T}_n^{\bowtie}\right) \xrightarrow[n\to+\infty]{} 0.$$

Hence, the total variation distance between the distributions of $T_n(k_n)$ and $T_n(\infty)$ goes to 1 as $n \to +\infty$ so the mixing time is greater than k_n for n large enough. Since this is true for any $k_n = o(n^{5/4})$, the mixing time must be at least $cn^{5/4}$ with $c > 0$.

We end this paper by a few remarks about our lower bound and an open question.

Remark 2.9. We proved a lower bound on the mixing time in the worst case, but our proof still holds for the mixing time from a typical starting point. We just need to fix $\varepsilon > 0$ small, take as initial condition a uniform triangulation $T_n(0)$ conditioned on $\left|\partial B^{\bullet}_{n^{1/4}}(T_n(0))\right| \leq \varepsilon\sqrt{n}$ and $\frac{n}{3} \leq \left|B^{\bullet}_{n^{1/4}}(T_n(0))\right| \leq \frac{2n}{3}$ and let $T_n^1(0) = B^{\bullet}_{n^{1/4}}(T_n(0))$. The event on which we condition has probability bounded away from 0 (by the results of [64] and coupling arguments) and after time $o(n^{5/4})$ there is still a seperating cycle of length $O(\varepsilon^{1/2}n^{1/4})$.

Remark 2.10. Here is a back-of-the-enveloppe computation that leads us to believe the lower bound we give is sharp if we start from a typical triangulation. The lengths of the geodesics in a uniform triangulation of volume n are of order $n^{1/4}$, so if we fix two vertices x and y the probability that a flip hits the geodesic from x to y is roughly $n^{-3/4}$. Hence, if we do $n^{5/4}$ flips, about $n^{1/2}$ of them will affect the distance between x and y. If we believe that this distance evolves roughly like a random walk, it will vary of about $\sqrt{n^{1/2}} = n^{1/4}$, which shows we are at the right scale. Of course, there are many reasons why this computation seems hard to be made rigourous, but it does not seem to be contradicted by numerical simulations.

Finally, note that even in the simpler case of triangulations of a polygon, the lower bound $n^{3/2}$ is believed to be sharp but the best known upper bound [113] is only $n^{5+o(1)}$. In our case we were not even able to prove the following.

Conjecture 2.11. The mixing time of $(T_n(k))_{k\geq 0}$ is polynomial in n.

2.A Appendix: Connectedness of the flip graph for type-I triangulations

In this appendix, we show that the Markov chain we study is indeed irreducible.

Lemma 2.12. Let $\mathscr{G}_n$ be the graph whose vertex set is $\mathscr{T}_n$ and where two triangulations are related if one can pass from one to the other by a flip. Then $\mathscr{G}_n$ is connected and its diameter is linear in n.

Proof. It is proved in [138] that the flip graph for type-III triangulations is connected, and in [91] that its diameter is linear in n. Hence, it is enough to show that any triangulation is connected to a type-III triangulation in $\mathscr{G}_n$ by a linear number of edges. If t is a finite triangulation with loops, it contains a *minimal* loop, that is, a loop dividing the sphere in two parts, one of which contains no loop. By flipping a minimal loop we delete a loop without to create any new one, so

we make the number of loops decrease and we can delete all loops in a linear number of flips. Moreover, if t contains no loop and there are two edges e_1, e_2 between the same pair of vertices, then flipping e_1 does not create any loop or additional multiple edges, so we can also delete all multiple edges in a linear number of flips. □

Chapter 3

The hyperbolic Brownian plane

ou *Un plan bien engraissé.*

This chapter is adapted from [45], published in Probability Theory and Related Fields.

We introduce and study a new random surface, which we call the *hyperbolic Brownian plane* and which is the near-critical scaling limit of the hyperbolic triangulations constructed in [58]. The law of the hyperbolic Brownian plane is obtained after biasing the law of the Brownian plane [62] by an explicit martingale depending on its perimeter and volume processes studied in [63]. Although the hyperbolic Brownian plane has the same local properties as those of the Brownian plane, its large scale structure is much different since we prove e.g. that is has exponential volume growth.

Contents

3.1 Introduction

The construction and the study of random surfaces as scaling limits of random planar maps has been a very active field of research in the last years, see [96, 114] for survey. The first such random surface that was built is the Brownian map [99, 116], which is now known to be the scaling limit of a wide class of finite planar maps conditioned to be large [1, 3, 24, 36, 65]. Curien & Le Gall introduced the Brownian plane in [62], which can be seen as a non-compact version of the Brownian map. They showed that it is the scaling limit of the Uniform Infinite Planar Quadrangulation (UIPQ). They also conjectured it to be the scaling limit of several other random infinite lattices such as the Uniform Infinite Planar Triangulation (UIPT) of Angel & Schramm [20] (we verify this fact below for type-I triangulations). The goal of this paper is to introduce and to study a new random surface which we call *the hyperbolic Brownian plane*. This surface is obtained as a near-critical scaling limit of the hyperbolic triangulations of [58].

The Brownian plane as the near-critical limit of the PSHIT. The spatial Markov property of random maps is a key feature of random lattices like the UIPT and UIPQ and has been used a lot in recent years to study their geometric structure, see e.g. [12, 43, 64]. Recently, Angel & Ray characterized all the triangulations of the half-plane enjoying a spatial Markov property and discovered a new family of triangulations of the half-plane having hyperbolic flavor [18]. This has been extended to cover the case of the full-plane in [58]. More precisely, [58] constructs a one-parameter family $(\mathbf{T}_\kappa)_{0<\kappa\leq\kappa_c}$ of Markovian random triangulations of the plane, where the value κ_c is equal to $\frac{2}{27}$. The triangulations $\mathbf{T}_\kappa$ for $\kappa < \kappa_c$ are called type-II Planar Stochastic Hyperbolic Triangulations (PSHIT). At the critical value $\kappa = \kappa_c$, the random triangulation $\mathbf{T}_{\kappa_c}$ is the UIPT of Angel & Schramm, whereas $\mathbf{T}_\kappa$ has hyperbolic features when $\kappa < \kappa_c$. Note that if $\kappa < \kappa_c$ is fixed, then it is impossible to rescale $\mathbf{T}_\kappa$ to get a scaling limit in the Gromov–Hausdorff sense [1]. Hence, in order to get a proper scaling limit, it is necessary to let the parameter $\kappa \to \kappa_c$ at the right speed as we renormalize the distances. If we let $\kappa \to \kappa_c$ too slow, then there is no scaling limit as above and if $\kappa \to \kappa_c$ too fast, then the scaling limit is just the Brownian plane: our approach is *near-critical*.

Our main tool for proving such a convergence will be the absolute continuity relations between the hyperbolic triangulations and the UIPT. These relations allow us to deduce convergence results for hyperbolic maps from the analogous results for the UIPT. The above works [20, 58] deal with type-II triangulations where loops are forbidden. Unfortunately, as of today, no scaling limit result is available in the literature for type-II triangulations. This forces us to work with type-I triangulations (i.e. where loops are allowed), for which the convergence to the Brownian map has been established [99]. Our first (easy) task is then to generalize the results of [58] and to introduce the type-I PSHIT, which we denote by

$$\big(\mathbb{T}_\lambda\big)_{0<\lambda\leq\lambda_c}, \quad \text{where} \quad \lambda_c = \frac{1}{12\sqrt{3}}.$$

As above, in the critical case $\lambda = \lambda_c$, the random lattice $\mathbb{T}_{\lambda_c}$ is just the type-I UIPT [65, 129].

We denote by $\mathcal{P}$ the Brownian plane of [62]. We recall that it is equipped with a volume measure $\mu_\mathcal{P}$. For $r \geq 0$, we write $B_r(\mathcal{P})$ for the closed ball of radius r centered at the origin point of $\mathcal{P}$. We also write $\overline{B_r}(\mathcal{P})$ for the hull of radius r, that is, the union of $B_r(\mathcal{P})$ together with all the bounded connected components of its complementary. We equip $\overline{B_r}(\mathcal{P})$ with the induced metric (i.e. the restriction of the metric on $\mathcal{P}$), and the restriction of $\mu_\mathcal{P}$. We will therefore consider $\overline{B_r}(\mathcal{P})$ as a measured metric space. The last ingredients we need before stating our

1. we can find in the ball of radius r of $\mathbf{T}_\kappa$ a number of points at distance at least $\frac{r}{10}$ from each other that goes to $+\infty$ as $r \to +\infty$, so the sequence $\big(\frac{1}{r}B_r(\mathbf{T}_\kappa)\big)_{r\geq 1}$ is not tight for the Gromov–Hausdorff topology

main theorem are the perimeter and volume processes of $\mathcal{P}$ introduced by Curien & Le Gall [63]. If $A \subset \mathcal{P}$, we will write $|A| = \mu_{\mathcal{P}}(A)$ for the measure of A. For $r > 0$, the *volume* of the hull of radius r of $\mathcal{P}$ is $|\overline{B_r}(\mathcal{P})|$. Moreover, Proposition 1.1 of [63] states that the limit

$$\lim_{\varepsilon \to 0} \frac{1}{\varepsilon^2} \left| B_{r+\varepsilon}(\mathcal{P}) \backslash \overline{B_r}(\mathcal{P}) \right| \tag{3.1}$$

is a.s well-defined and positive. We call it the *perimeter* of $\overline{B_r}(\mathcal{P})$ and denote it by $|\partial \overline{B_r}(\mathcal{P})|$. Our main theorem is the following.

Theorem 3.1 ($\mathcal{P}^h$ as a near-critical scaling limit of the PSHIT)**.** For $n \geq 0$, consider $\mathbb{T}_{\lambda_n}$ the type-I planar stochastic hyperbolic triangulation of parameter $\lambda_n \to \lambda_c$ in such a way that

$$\lambda_n = \lambda_c \left(1 - \frac{2}{3n^4}\right) + o\left(\frac{1}{n^4}\right). \tag{3.2}$$

Then we have the following convergence for the local Gromov–Hausdorff–Prokhorov distance:

$$\frac{1}{n} \mathbb{T}_{\lambda_n} \xrightarrow[n \to \infty]{(d)} \mathcal{P}^h,$$

where $\mathcal{P}^h$ is a random locally compact metric space that we call the *hyperbolic Brownian plane*. Its distribution is characterized by the fact that for every $r \geq 0$, the random measured metric space $\overline{B_r}(\mathcal{P}^h)$ has density

$$e^{-2|\overline{B_{2r}}(\mathcal{P})|} e^{|\partial \overline{B_{2r}}(\mathcal{P})|} \int_0^1 e^{-3|\partial \overline{B_{2r}}(\mathcal{P})| x^2} \mathrm{d}x \tag{3.3}$$

with respect to $\overline{B_r}(\mathcal{P})$.

The choice of the constant $\frac{2}{3}$ in (3.2) was made so that the expression (3.3) looks as simple as possible. Of course another choice woud have resulted in a scaling limit just obtained by dilating $\mathcal{P}^h$. The fact that we need to bias $\overline{B_r}(\mathcal{P})$ by a function of the perimeter and volume of the hull of radius $2r$ instead of r may seem surprising. It is due to the fact that we equip $\overline{B_r}(\mathcal{P})$ with the *induced* distance. Hence, the distance between two points in the hull of radius r of a map may depend on the part of the map that lies outside of this hull (but not outside of the hull of radius $2r$). In order to obtain a similar result with $|\partial \overline{B_r}(\mathcal{P})|$ and $|\overline{B_r}(\mathcal{P})|$, we would need to equip $\overline{B_r}(\mathcal{P})$ with its *intrinsic* distance instead of the *induced* one (see Section 2.1 for more details about this distinction). We would also need to prove an analog of Proposition 3.10 for the intrinsic distance, which we have not been able to do.

In the study of a one-parameter family of models exhibiting a critical behavior, it has become quite usual to study *near-critical* (scaling) limits. By near-critical, we mean that the parameter converges to its critical value at the right speed as the distances in the graph or the mesh of the lattice are going to zero. Understanding the near-critical limit usually sheds some light on the critical model because of the existence of scaling relations between near-critical and critical exponents. See for example the works on near-critical percolation [77, 89, 123], on the Ising model [73] or on the Erdös-Rényi random graph [4, 5].

Techniques. As said above, the idea of the proof of Theorem 3.1 is to use the absolute continuity relations between the hulls of the type-I UIPT and of the hyperbolic triangulations $\mathbb{T}_\lambda$. Our main technical tool is a reinforcement of the convergence of the UIPT towards the Brownian plane. In the result below, $|\overline{B}_r(\mathbb{T}_{\lambda_c})|$ and $|\partial \overline{B}_r(\mathbb{T}_{\lambda_c})|$ respectively stand for the volume (number of vertices) and perimeter of the hull of radius r in the UIPT.

Theorem 3.2 (Extended convergence towards the Brownian plane)**.** We have the joint convergences

$$\left(\frac{1}{n}\mathbb{T}_{\lambda_c}, \left(\frac{1}{n^4}|\overline{B_{rn}}(\mathbb{T}_{\lambda_c})|\right)_{r\geq 0}, \left(\frac{1}{n^2}|\partial\overline{B_{rn}}(\mathbb{T}_{\lambda_c})|\right)_{r\geq 0}\right) \xrightarrow[n\to\infty]{(d)} \left(\mathcal{P}, \left(3|\overline{B_r}(\mathcal{P})|\right)_{r\geq 0}, \left(\frac{3}{2}|\partial\overline{B_r}(\mathcal{P})|\right)_{r\geq 0}\right)$$

in the local Gromov–Hausdorff–Prokhorov sense for the first marginal and in the Skorokhod sense for the last two.

The convergence of the first marginal has been established in the case of quadrangulations in [62] and we extend the proof to cover our case. On the other hand, the joint convergence of the last two marginals follows from the work [64] (both in the case of quadrangulations and type-I triangulations). But it is important in our Theorem 3.2 that those convergences hold jointly, which requires some additional work. The convergence of near-critical PSHIT towards the hyperbolic Brownian plane then follows from Theorem 3.2 and a couple of asymptotic enumeration results gathered in Section 1.

Properties of $\mathcal{P}^h$. We also establish some properties of the hyperbolic Brownian plane. Since the density (3.3) goes to 1 as r goes to 0, the hyperbolic Brownian plane is "locally isometric" to the Brownian plane (and hence also to the Brownian map). More precisely, for all $\varepsilon > 0$, there is a $\delta > 0$ and a coupling between $\mathcal{P}$ and $\mathcal{P}^h$ such that, with probability at least $1-\varepsilon$, they have the same ball of radius δ around the origin. We also prove that $\mathcal{P}^h$ almost surely has Hausdorff dimension 4 and is homeomorphic to the plane.

The Brownian map is known to be invariant under uniform re-rooting. This means that, if we resample its root uniformly according to its volume measure, the rooted metric space we obtain has the same distribution as the Brownian map. This property has played an important role in the proof of universality results in [1, 3, 24, 36, 99], and in the axiomatic characterization of the Brownian map given by [117]. Unfortunately, it makes no sense anymore when the volume measure is infinite. However, we prove the following property of $\mathcal{P}^h$: for every measurable, nonnegative function f we have

$$\mathbb{E}\Big[\int_{\mathcal{P}^h} f(\mathcal{P}^h, \rho, y)\mu_{\mathcal{P}^h}(\mathrm{d}y)\Big] = \mathbb{E}\Big[\int_{\mathcal{P}^h} f(\mathcal{P}^h, y, \rho)\mu_{\mathcal{P}^h}(\mathrm{d}y)\Big],$$

where ρ is the origin of $\mathcal{P}^h$ and $\mu_{\mathcal{P}^h}$ its volume measure. This property is a continuous analog of the discrete property of unimodularity, which is a natural substitute for infinite random graphs to invariance under uniform rerooting (see for example [7] for the discrete case). More precisely, the two properties are equivalent for finite random graphs. Our result shows that the hyperbolic Brownian plane is a natural surface to look at even from the purely continuum point of view.

Also, since the hyperbolic Brownian plane is a biased version of the Brownian plane, it is possible to define the perimeter P_r^h and the volume V_r^h of its hull of radius r as Curien and Le Gall did for the Brownian plane in [63]. We identify the joint distribution of these two processes in a similar way as in [63]. Let Z^h be the subcritical continuous-state branching process with branching mechanism

$$\psi(\lambda) = \sqrt{\frac{8}{3}}\lambda\sqrt{\lambda+3}, \qquad \lambda > 0.$$

Moreover, for all $\delta > 0$, let ν_δ be the following measure on $\mathbb{R}^+$:

$$\nu_\delta(\mathrm{d}x) = \frac{\delta^3 e^{2\delta}}{1+2\delta}\frac{e^{-\frac{\delta^2}{2x}-2x}}{\sqrt{2\pi x^5}}\mathbb{1}_{x>0}\,\mathrm{d}x.$$

Theorem 3.3 (Perimeter and volume processes of $\mathcal{P}^h$)**.**

1) The perimeter process P^h of the hyperbolic Brownian plane has the same distribution as the time-reversal of Z^h, started from $+\infty$ at time $-\infty$, and conditioned to die at time 0.

2) Conditionally on P^h, the process V^h has the same distribution as the process

$$\Big(\sum_{s_i \le r} \xi_i^h\Big)_{r\ge 0},$$

where (s_i) is a measurable enumeration of the jumps of P^h, the random variables ξ_i^h are independent and ξ_i^h has distribution $\nu_{|\Delta P^h_{s_i}|}$ for all i.

This allows us to compute the asymptotics of these processes as in the discrete case in [58].

Corollary 3.1 (Exponential growth)**.** We have the convergences

$$\frac{P_r^h}{e^{2\sqrt{2}r}} \xrightarrow[r\to+\infty]{a.s.} \mathcal{E} \quad \text{and} \quad \frac{V_r^h}{P_r^h} \xrightarrow[r\to+\infty]{a.s.} \frac{1}{4},$$

where $\mathcal{E}$ is an exponential variable of parameter 12.

The structure of the paper is as follows. In Section 1 we introduce the type-I analog of the PSHIT, and show they are the only type-I triangulations enjoying a similar domain Markov property as that defined by Curien in [58]. We also gather a few enumeration results. Section 2 is devoted to the proof of Theorem 3.1 and 3.2 and Section 3 to the study of the perimeter and volume processes. Appendix A contains a technical result about the Gromov–Hausdorff–Prokhorov convergence. It shows that under some technical assumptions, if a sequence (X_n) of metric spaces converges to X, then the hulls $\overline{B_r}(X_n)$ converge to $\overline{B_r}(X)$.

Acknowledgments: I thank Nicolas Curien for suggesting me to study this object, and for carefully reading many earlier versions of this manuscript. I also thank the anonymous referee for his useful comments. I acknowledge the support of ANR Liouville (ANR-15-CE40-0013) and ANR GRAAL (ANR-14-CE25-0014).

3.2 Prerequisites: enumeration and type-I PSHIT

3.2.1 Combinatorial preliminaries

A *type-I triangulation of a p-gon* is a planar map equipped with a distinguished oriented edge called the *root*, in which the face to the right of the root has a simple boundary of length p and every other face has degree 3. It may contain multiple edges and self-loops. In this whole work we will make repeated use of the results of Krikun [95] about the enumeration of type-I triangulations. For $p \ge 1$ and $n \ge 0$, we write $\mathscr{T}_{n,p}$ for the set of type-I triangulations of a p-gon with n inner vertices, and $\#\mathscr{T}_{n,p}$ for its cardinal. By Euler's formula, a triangulation of a p-gon with n inner vertices has $3n+2p-3$ edges. Hence, the main theorem of [95] in the case $r=0$ (that is, triangulations with only one hole) can be rewritten

$$\#\mathscr{T}_{n,p} = \frac{p(2p)!}{(p!)^2}\frac{4^{n-1}(2p+3n-5)!!}{n!(2p+n-1)!!} \underset{n\to+\infty}{\sim} C(p)\lambda_c^{-n}n^{-5/2}, \tag{3.4}$$

where we recall that $\lambda_c = \frac{1}{12\sqrt{3}}$, and where

$$C(p) = \frac{3^{p-2}p(2p)!}{4\sqrt{2\pi}(p!)^2} \underset{p\to+\infty}{\sim} \frac{1}{36\pi\sqrt{2}}12^p\sqrt{p}. \tag{3.5}$$

A triangulation of the sphere with n vertices can be seen after a root transformation as a triangulation of a 1-gon with $n-1$ inner vertices (see Figure 2 in [65]). Hence, the number of triangulations of the sphere with n vertices is

$$\#\mathscr{T}_{n-1,1} \underset{n\to+\infty}{\sim} \frac{1}{72\sqrt{6\pi}}\lambda_c^{-n}n^{-5/2}. \tag{3.6}$$

We also write $Z_p(\lambda)=\sum_{n\geq 0}\#\mathscr{T}_{n,p}\lambda^n$. Note that, by the asymptotics (3.4), we have $Z_p(\lambda)<+\infty$ iff $\lambda\leq\lambda_c$. We finally write $G_\lambda(x)=\sum_{p\geq 1}Z_p(\lambda)x^p$. Formula (4) of [95] computes G_λ after a simple change of variables:

$$G_\lambda(x)=\frac{\lambda}{2}\Big(\Big(1-\frac{1+8h}{h}x\Big)\sqrt{1-4(1+8h)x}-1+\frac{x}{\lambda}\Big), \tag{3.7}$$

where $h\in\big(0,\frac{1}{4}\big]$ is such that

$$\lambda=\frac{h}{(1+8h)^{3/2}}. \tag{3.8}$$

Note that our h corresponds to the h^3 of [95]. From (3.7) we easily get

$$Z_1(\lambda)=\frac{1}{2}-\frac{1+2h}{2\sqrt{1+8h}} \tag{3.9}$$

and, for $p\geq 2$,

$$Z_p(\lambda)=(2+16h)^p\frac{(2p-5)!!}{p!}\frac{(1-4h)p+6h}{4(1+8h)^{3/2}}. \tag{3.10}$$

We now prove a combinatorial estimate that we will use later in the proof of the convergence of the type-I UIPT to the Brownian plane.

Lemma 3.2. When $n,p\to+\infty$ with $p=O\big(\sqrt{n}\big)$, we have

$$\#\mathscr{T}_{n,p}\sim\frac{1}{36\pi\sqrt{2}}\lambda_c^{-n}n^{-5/2}12^p\sqrt{p}\,\exp\Big(-\frac{2p^2}{3n}\Big).$$

Proof. This follows from developping (3.4) asymptotically using the Stirling formula. The same estimate for type-II triangulations can be found in the proof of Proposition 8 of [64]. Only some constants differ, and these constants for type-I triangulations are given in Section 6.1 of [64]. We omit the details here. □

3.2.2 Definition of the type-I PSHIT

The goal of this section is to construct the analog of the hyperbolic triangulations of [58] in the case of type-I triangulations. Since the construction is roughly the same, we only stress the differences. If t is a finite, rooted triangulation with a simple hole of perimeter p, we write $|t|$ for its number of vertices. By $t\subset T$, we mean that T may be obtained by filling the hole of t with an infinite triangulation of perimeter p.

Definition 3.3. Let $\lambda>0$. A random (rooted) infinite type-I triangulation of the plane T is λ-*Markovian* if there are constants $\big(C_p(\lambda)\big)_{p\geq 1}$ such that, for all finite rooted triangulations t with a hole of perimeter p, we have

$$\mathbb{P}\big(t\subset T\big)=C_p(\lambda)\lambda^{|t|}.$$

Remark 3.4. Like Curien [58], we choose a stronger definition of the Markov property than that of Angel & Ray [18]. Although these two definitions should coincide for type-II triangulations, this is important in the context of type-I triangulations. Indeed, a weaker definition would allow a much larger class of Markovian triangulations (see [18], Section 3.4 for a precise discussion in the half-planar case).

Proposition 3.5. If $\lambda > \lambda_c$, there is no λ-Markovian type-I triangulation. If $\lambda \leq \lambda_c$, there is a unique one (in distribution). Besides we have

$$C_p(\lambda) = \frac{1}{\lambda}\Big(8 + \frac{1}{h}\Big)^{p-1} \sum_{q=0}^{p-1} \binom{2q}{q} h^q, \tag{3.11}$$

where h is like in (3.8). We will write $\mathbb{T}_\lambda$ for this triangulation and $\mathbb{T} = \mathbb{T}_{\lambda_c}$, which coincides with the type-I UIPT [65, 129].

Proof. The uniqueness can be proved along the same lines as in Section 1 of [58]. The analog of relation (5) in [58] is, for all $p \geq 1$,

$$C_p(\lambda) = \lambda C_{p+1}(\lambda) + 2\sum_{i=0}^{p-1} C_{p-i}(\lambda) Z_{i+1}(\lambda). \tag{3.12}$$

Note that in our case, the sum starts at 0 and ends at $p-1$ (instead of 1 and $p-2$ in [58]) because of the possible presence of loops. Hence, the $\lambda \leq \lambda_c$ condition comes from the fact that the radius of convergence of Z_p is λ_c by (3.4). If we write $F_\lambda(x) = \sum_{p\geq 0} C_p(\lambda) x^p$, then (3.12) becomes

$$F_\lambda(x) = \frac{\lambda}{x}\Big(F_\lambda(x) - C_1(\lambda)x\Big) + \frac{2}{x} G_\lambda(x) F_\lambda(x),$$

so

$$F_\lambda(x) = \frac{\lambda C_1(\lambda) x}{\lambda - x + 2G_\lambda(x)}. \tag{3.13}$$

By combining (3.13) and (3.7) we get

$$F_\lambda(x) = \frac{C_1(\lambda) x}{\Big(1 - \frac{1+8h}{h} x\Big)\sqrt{1 - 4(1+8h)x}}. \tag{3.14}$$

Finally, for all $p \geq 1$,

$$\begin{aligned} C_p(\lambda) &= C_1(\lambda) \sum_{q+r+1=p} \Big(\frac{1+8h}{h}\Big)^r \binom{-1/2}{q} (-4)^q (1+8h)^q \\ &= C_1(\lambda) \sum_{q=0}^{p-1} \Big(\frac{1+8h}{h}\Big)^{p-1-q} \frac{(-1)^q}{4^q} \binom{2q}{q} (-4)^q (1+8h)^q \\ &= C_1(\lambda)\Big(8 + \frac{1}{h}\Big)^{p-1} \sum_{q=0}^{p-1} \binom{2q}{q} h^q. \end{aligned}$$

To prove the uniqueness and obtain the desired formula, it only remains to prove that we must have $C_1(\lambda) = \frac{1}{\lambda}$. Let t_0 be the map consisting of a single loop. Since any triangulation of the sphere can be seen as a triangulation of a 1-gon (see Figure 2 in [65]), we must have $t_0 \subset \mathbb{T}_\lambda$ with probability 1. Hence, $C_1(\lambda) = \frac{1}{\lambda}$.

The proof of the existence is essentially the same as in [58]: consider the sequence $\big(C_p(\lambda)\big)_{p\geq 1}$ given by (3.11) with $C_p(\lambda) > 0$ for all p. It verifies (3.12), so for all $p \geq 1$, we have

$$\lambda \frac{C_{p+1}(\lambda)}{C_p(\lambda)} + 2\sum_{i=0}^{p-1} Z_{i+1}(\lambda)\frac{C_{p-i}(\lambda)}{C_p(\lambda)} = 1.$$

The last display can be interpreted as transition probabilities for the peeling process of $\mathbb{T}_\lambda$. This allows us to construct a random triangulation by peeling like in [58]. The same arguments as in [58] prove that we get a triangulation $\mathbb{T}_\lambda$ of the plane, that the distribution of $\mathbb{T}_\lambda$ is independent of the peeling algorithm used for the construction, and that $\mathbb{T}_\lambda$ is λ-Markovian. □

We note that in the critical case, we have a more explicit expression of $C_p(\lambda)$: we have $h = \frac{1}{4}$, so

$$C_p\big(\lambda_c\big) = \lambda_c^{-1} \times 12^{p-1} \sum_{q=0}^{p-1} \frac{1}{4^q}\binom{2q}{q} = 2\sqrt{3} \times 3^p \frac{p(2p)!}{p!^2}, \tag{3.15}$$

as easily proved by induction.

We will later need precise asymptotics of the numbers $C_p(\lambda)$. For that purpose, we already state the following estimate.

Lemma 3.6. Let $(p_n)_{n\geq 0}$ be a sequence of positive integers such that $\frac{p_n}{n^2} \to \frac{3}{2}p$ where $p \geq 0$. Let also $(h_n)_{n\geq 0}$ be a sequence of numbers in $\big(0, \frac{1}{4}\big]$ such that

$$h_n = \frac{1}{4} - \frac{1}{2n^2} + o\Big(\frac{1}{n^2}\Big).$$

Then we have

$$\frac{1}{\sqrt{p_n}} \sum_{q=0}^{p_n-1} \binom{2q}{q} h_n^q \xrightarrow[n\to+\infty]{} \frac{2}{\sqrt{\pi}} \int_0^1 e^{-3px^2}\mathrm{d}x.$$

Proof. Note that if $q = \lfloor xn^2 \rfloor$ with $x \in \big(0, \frac{3}{2}p\big)$, then

$$\frac{1}{\sqrt{p_n}}\binom{2q}{q} h_n^q \underset{n\to+\infty}{\sim} \frac{1}{n^2}\frac{1}{\sqrt{\pi x}}\sqrt{\frac{3}{2p}}e^{-2x}.$$

The Riemann sums are easily seen to converge to

$$\sqrt{\frac{2}{3p}} \int_0^{3p/2} \frac{1}{\sqrt{\pi y}} e^{-2y}\mathrm{d}y,$$

which is equal to the desired integral after the change of variables $y = \frac{3p}{2}x^2$. The details are left to the reader. □

3.3 Convergence to the hyperbolic Brownian plane

3.3.1 About the Gromov–Hausdorff–Prokhorov convergence

We first recall from [2] the definition of the (bipointed) Gromov–Hausdorff–Prokhorov distance.

Definition 3.7. Let $\big((X_1,d_1),x_1,y_1,\mu_1\big)$ and $\big((X_2,d_2),x_2,y_2,\mu_2\big)$ be two compact bipointed measured metric spaces. We assume the measures μ_1 and μ_2 are finite (but they do not have to be probability measures). The *Gromov–Hausdorff–Prokhorov distance* (we will sometimes write GHP distance and denote it by d_{GHP}) between X_1 and X_2 is the infimum of all $\varepsilon>0$ for which there are isometrical embeddings Ψ_1 and Ψ_2 of X_1 and X_2 in the same metric space (Z,d) such that:

a) $\Psi_1(x_1)=\Psi_2(x_2)$,
b) $d\big(\Psi_1(y_1),\Psi_2(y_2)\big)\leq\varepsilon$,
c) the Hausdorff distance between $\Psi_1(X_1)$ and $\Psi_2(X_2)$ is not greater than ε,
d) the Lévy–Prokhorov distance between $\mu_1\circ\Psi_1^{-1}$ and $\mu_2\circ\Psi_2^{-1}$ is not greater than ε.

The same definition holds for pointed measured compact metric spaces. We just need to withdraw condition b).

If $\big((X,d),x,\mu\big)$ is a pointed measured metric space and $r\geq 0$, we write $B_r(X)$ for the closed ball of radius r centered at x in X, equipped with the restrictions of the distance d and of the measure μ.

Definition 3.8. Let $\big((X_1,d_1),x_1,\mu_1\big)$ and $\big((X_2,d_2),x_2,\mu_2\big)$ be two locally compact pointed measured metric spaces. The *local Gromov–Hausdorff–Prokhorov distance* between X_1 and X_2, which we will denote by $d_{LGHP}(X_1,X_2)$, is the sum

$$\sum_{r\geq 1}\frac{1}{2^r}\max\Big(1,d_{GHP}\big(B_r(X_1),B_r(X_2)\big)\Big).$$

Definition 3.9. Let $\big((X,d),x,y,\mu\big)$ be a locally compact bi-pointed measured metric space. The *hull* of center x and radius r with respect to y is the union of the closed ball of radius r centered at x and all the connected components of its complementary that do not contain y. It is denoted by $\overline{B_r}(X,x,y)$. We will write $\overline{B_r}(X)$ when there is no ambiguity. Equipped with the restrictions of d and μ, it is a compact measured metric space.

If $\big((X,d),x\big)$ is unbounded and one-ended we will omit the second distinguished point: the hull will be the union of the ball of radius r centered at x and all the bounded connected components of its complementary. It means that y is at infinity.

Recall that there are two natural ways to equip a part A of (X,d) with a metric: the *induced* metric, i.e. the restriction of d to A, and the *intrisic* one, which makes A a geodesic space when it is well-defined (see [48], Chapter 2.3). In order to avoid further confusions, we insist that $\overline{B_r}(X)$ is equipped with the induced distance. If m is a map, we will also write $B_r(m)$ for the map consisting of all the faces of m having at least one vertex at distance at most $r-1$ from the root vertex, along with all their vertices and edges. We write $B_r^\bullet(m)$ for the map that is the union of $B_r(m)$ and all the bounded connected components of its complement. When m is seen as a metric space, the hull $\overline{B_r}(m)$ has the same set of vertices as $B_r^\bullet(m)$. However, the distances inherited from m are not the same as those in $B_r^\bullet(m)$. We will always see $\overline{B_r}(m)$ as a metric space and $B_r^\bullet(m)$ as a map.

We will need several times to deduce properties of one of these distances from properties of the other. To this end, we point out that if m is a map, then $\overline{B_r}(m)$ is a measurable function of $B_{2r}^\bullet(m)$ for all $r\geq 0$. Indeed, any geodesic in m between two vertices x and y of $\overline{B_r}(m)$ must stay in $B_{2r}^\bullet(m)$, so the distances between x and y in m and in $B_{2r}^\bullet(m)$ coincide.

We will also need the following technical result that is proved in Appendix A.

Proposition 3.10. Let $\big((X_n,d_n),x_n,\mu_n\big)$ be a sequence of unbounded, locally compact, pointed measured metric spaces. Assume that $\big((X_n,d_n),x_n,\mu_n\big)$ converges for the local GHP distance to a measured metric space $\big((X,d),x,\mu\big)$. Let $r\geq 0$. We assume that:

(i) X and the X_n are one-ended length spaces,
(ii) every non-empty open subset of X has positive measure,
(iii) the function $V : s \to \mu\big(\overline{B_s}(X)\big)$ is continuous at r.

Then:

1) $\overline{B_r}(X_n)$ converge for the GHP distance to $\overline{B_r}(X)$,
2) in particular, we have the convergence

$$\mu_n\big(\overline{B_r}(X_n)\big) \longrightarrow \mu\big(\overline{B_r}(X)\big).$$

Moreover, the proposition also holds for bipointed, compact spaces $\big((X_n, d_n), x_n, y_n, \mu_n\big)$ and $\big((X, d), x, y, \mu\big)$.

3.3.2 Convergence of the type-I UIPT to the Brownian plane

If (X, d) is a metric space and $\alpha > 0$, we will write αX for the metric space $(X, \alpha d)$. We recall that $\mathbb{T} = \mathbb{T}_{\lambda_c}$ is the type-I UIPT. If t is a (possibly infinite) triangulation, recall that $B_r(t)$ denotes its ball of radius r around the origin of its root edge and $\overline{B_r}(t)$ its hull, endowed with the induced metric. We denote by $\mathcal{P}$ the Brownian plane defined in [62]. Our goal in this section is to prove Theorem 3.2. We start with the first two marginals, whereas the convergence of the third one will be the content of Section 2.3.

Proposition 3.11. Let $\mu_{\mathbb{T}}$ be the measure on $\mathbb{T}$ giving mass 1 to each vertex, and let $\mu_{\mathcal{P}}$ be the volume measure on $\mathcal{P}$ [62, 63]. We have the convergence

$$\Big(\frac{1}{n^{1/4}}\mathbb{T}, \frac{1}{3n}\mu_{\mathbb{T}}\Big) \xrightarrow[n\to+\infty]{(d)} \big(\mathcal{P}, \mu_{\mathcal{P}}\big)$$

for the local GHP distance.

We note that this result has been proved for quadrangulations in [62] for the Gromov–Hausdorff distance and in [140] for the (stronger) GHP convergence. Our main tool will be the following theorem by Curien and Le Gall. It is a refinement of the convergence of uniform type-I triangulations proved by Le Gall in [99].

Theorem 3.4 ([65], Appendix A1, Theorem 6)**.** Let T_n be a uniform type-I triangulation of the sphere with n vertices and μ_{T_n} the counting measure on the set of its vertices. Let also $\mathbf{m}_\infty$ be the Brownian map and $\mu_{\mathbf{m}_\infty}$ its volume measure ([99]). The following convergence holds for the GHP distance:

$$\Big(\frac{1}{n^{1/4}}T_n, \frac{1}{n}\mu_{T_n}\Big) \xrightarrow[n\to+\infty]{(d)} \Big(\frac{1}{3^{1/4}}\mathbf{m}_\infty, \mu_{\mathbf{m}_\infty}\Big).$$

To prove Proposition 3.11 we need to invert the local and the scaling limit. Hence, we need the local convergence $T_n \to \mathbb{T}$ to be "uniform in the scale", which is the point of the next lemma. It parallels Proposition 1 of [62] in the case of type-I triangulations.

Proposition 3.12. Let $n \geq 1$. Let also $T_n^* = (T_n, y)$ be a uniform type-I triangulation of the sphere with n vertices, equipped with a uniform distinguished vertex y. We write $B_r^\bullet(T_n^*)$ for the hull of radius r of T_n, centered at the root, with respect to y. Then, for all $\varepsilon > 0$, there is a constant $A > 0$ such that if $n > Ar^4$, there is a coupling between T_n^* and $\mathbb{T}$ in which

$$\mathbb{P}\big(B_r^\bullet(T_n^*) = B_r^\bullet(\mathbb{T})\big) \geq 1 - \varepsilon.$$

Proof. The proof of Proposition 1 in [62] uses the Schaeffer bijection between maps and trees. Although a similar bijection exists for triangulations, it is more complicated. Hence, we do the computations directly on maps instead of trees as in Section 6 of [65].

Let $\delta > 0$. We know from Section 6.1 of [64] that $\frac{1}{r^4}|B_r^\bullet(\mathbb{T})|$ and $\frac{1}{r^2}|\partial B_r^\bullet(\mathbb{T})|$ converge in distribution to a.s. positive random variables. Hence, there are positive constants c_δ and C_δ such that, for r large enough, we have

$$\mathbb{P}\left(c_\delta\, r^2 \leq |\partial B_r^\bullet(\mathbb{T})| \leq C_\delta\, r^2 \text{ and } c_\delta\, r^4 \leq |B_r^\bullet(\mathbb{T})| \leq C_\delta\, r^4\right) \geq 1-\delta.$$

Now take m and p such that $c_\delta\, r^2 \leq p \leq C_\delta r^2$ and $c_\delta\, r^4 \leq m \leq C_\delta r^4$. Let t be a triangulation of a p-gon with m vertices (including the boundary) such that t is a possible value of $B_r^\bullet(\mathbb{T})$. On the one hand, we have

$$\mathbb{P}\big(B_r^\bullet(\mathbb{T}) = t\big) \underset{\text{Def 3.3}}{=} C_p(\lambda_c)\lambda_c^m \underset{(3.15)}{=} 2\sqrt{3}\times 3^p \frac{p(2p)!}{p!^2}\cdot \lambda_c^m \underset{p\to+\infty}{\sim} \frac{2\sqrt{3}}{\sqrt{\pi}}\lambda_c^m \times 12^p\sqrt{p}.$$

On the other hand, we fix $A > C_\delta$ and we take $n \geq Ar^4$. There are $n\#\mathscr{T}_{n-1,1}$ pointed triangulations of the sphere with n vertices. Moreover, if $B_r^\bullet(T_n^*) = t$, there are $\#\mathscr{T}_{n-m,p}$ ways to fill the p-gon to complete T_n^* and $n-m$ ways to choose the distinguished vertex in it (the distinguished vertex cannot lie in $B_r(T_n^*)$, since then $B_r^\bullet(T_n^*)$ would be the full T_n^*). Hence, we have $\mathbb{P}\big(B_r^\bullet(T_n^*) = t\big) = \frac{(n-m)\#\mathscr{T}_{n-m,p}}{n\#\mathscr{T}_{n-1,1}}$. When we let $r \to +\infty$, we have $n-m, p \to +\infty$ with $p = O\big(\sqrt{n-m}\big)$. By Lemma 3.2, when r goes to $+\infty$, the probability $\mathbb{P}\big(B_r^\bullet(T_n^*) = t\big)$ is equivalent to

$$\frac{2\sqrt{3}}{\sqrt{\pi}}\lambda_c^m \times 12^p\sqrt{p}\exp\Big(-\frac{2p^2}{3(n-m)}\Big) \underset{r\to+\infty}{\sim} \exp\Big(-\frac{2p^2}{3(n-m)}\Big)\mathbb{P}\big(B_r^\bullet(\mathbb{T}) = t\big).$$

Hence, if we have chosen A large enough, the following holds:

$$(1-\delta)\mathbb{P}\big(B_r^\bullet(\mathbb{T}) = t\big) \leq \mathbb{P}\big(B_r^\bullet(T_n^*) = t\big) \leq (1+\delta)\mathbb{P}\big(B_r^\bullet(T) = t\big),$$

as soon as $c_\delta r^2 \leq |\partial t| \leq C_\delta r^2$ and $c_\delta r^4 \leq |t| \leq C_\delta r^4$. But we know that $B_r^\bullet(\mathbb{T})$ satisfies these assumptions with probability at least $1-\delta$. Hence, we can easily prove that, for $n \geq Ar^4$ and any set B of finite maps, we have

$$\big|\mathbb{P}\big(B_r^\bullet(\mathbb{T}) \in B\big) - \mathbb{P}\big(B_r^\bullet(T_n^*) \in B\big)\big| \leq 4\delta.$$

This shows that, for r large enough and $n \geq Ar^4$, the total variation distance between the distributions of $B_r^\bullet(\mathbb{T})$ and $B_r^\bullet(T_n^*)$ is less than 4δ, which proves the proposition. □

Proof of Proposition 3.11. We use Proposition 3.12 with $2r$ instead of r. The metric spaces $B_r(T_n^*)$ and $B_r(\mathbb{T})$ (equipped with the induced distance) are measurable functions of respectively $B_{2r}^\bullet(T_n^*)$ and $B_{2r}^\bullet(\mathbb{T})$. Hence, Proposition 3.12 still holds if we replace the maps $B_r^\bullet$ by the metric spaces B_r. The proof is now the same as the proof of Theorem 2 in [62], with two small modifications:

— we deal with Gromov–Hausdorff–Prokhorov convergence and not only Gromov–Hausdorff, but this does not change anything in the details of the proof, see [140] for details,

— the constant factors are not the same: because of the factor $\frac{1}{3^{1/4}}$ in Theorem 3.4, the measured spaces $\Big(\frac{1}{n^{1/4}}\mathbb{T}, \frac{1}{n}\mu_{\mathbb{T}}\Big)$ converge to $\big(\frac{1}{3^{1/4}}\mathcal{P}, \mu_{\mathcal{P}}\big)$. This has the same distribution as $\big(\mathcal{P}, 3\mu_{\mathcal{P}}\big)$ by the scaling property of the Brownian plane.

□

We can now prove the joint convergence of the first two marginals in Theorem 3.2.

Proposition 3.13. We have the joint convergence

$$\left(\frac{1}{n}\mathbb{T}, \left(\frac{1}{n^4}|\overline{B_{rn}}(\mathbb{T})|\right)_{r\geq 0}\right) \xrightarrow[n\to+\infty]{(d)} \left(\mathcal{P}, \left(3|\overline{B_r}(\mathcal{P})|\right)_{r\geq 0}\right),$$

where the convergence of the first marginal is for the local GHP distance, and the second one for the Skorokhod topology.

We will deduce Proposition 3.13 from Proposition 3.11 thanks to the second point of Proposition 3.10. Let us check this carefully.

Proof of Proposition 3.13. By the Skorokhod representation theorem, we may assume the convergence in Proposition 3.11 is almost sure. Theorem 1.4 of [63] computes $\mathbb{E}\left[e^{-|\overline{B_r}(\mathcal{P})|}\right]$. In particular, it is a continuous function of r. Since the process $\left(|\overline{B_r}(\mathcal{P})|\right)_{r\geq 0}$ has only positive jumps, it means that for all $r \geq 0$, it is almost surely continuous at r. Finally, the Brownian plane is defined in [62] as a quotient of $\mathbb{R}$. This means that there is a continuous surjection from $\mathbb{R}$ to $\mathcal{P}$, and the volume measure on $\mathcal{P}$ is the push-forward of the Lebesgue measure under this surjection. The inverse-image of a non-empty open subset of $\mathcal{P}$ is a non-empty open subset of $\mathbb{R}$. Hence, it has positive measure, which means that any non-empty open subset of $\mathcal{P}$ has positive measure.

Instead of $\mathbb{T}$, we can consider the metric space $\mathbb{T}^e$, which is the union of all the vertices and edges of $\mathbb{T}$. It is equipped with the metric that makes it a geodesic space in which all edges have length 1. We also equip this space with the counting measure $\mu_{\mathbb{T}}$ on the set of vertices. We have $d_{GHP}(B_r(\mathbb{T}), B_r(\mathbb{T}^e)) \leq 1$ for all r, so $d_{LGHP}\left(\left(\frac{1}{n^{1/4}}\mathbb{T}, \frac{1}{3n}\mu_{\mathbb{T}}\right), \left(\frac{1}{n^{1/4}}\mathbb{T}^e, \frac{1}{3n}\mu_{\mathbb{T}}\right)\right) \leq \frac{2}{n^{1/4}}$. Hence, Proposition 3.11 still holds if we replace $\mathbb{T}$ by $\mathbb{T}^e$. The sequence $\left(\frac{1}{n}\mathbb{T}^e, \frac{1}{n^4}\mu_{\mathbb{T}}\right)_{n\geq 0}$ satisfies the assumptions of Proposition 3.10, so for all $(r_1, \dots, r_k) \in (\mathbb{R}^+)^k$ we have

$$\forall i \in [\![1, k]\!],\ \frac{1}{3n^4}\mu_{\mathbb{T}}\left(\overline{B_{r_i n}}(\mathbb{T})\right) \xrightarrow[n\to+\infty]{a.s.} \mu_{\mathcal{P}}\left(\overline{B_{r_i}}(\mathcal{P})\right).$$

This gives the joint convergence of $\frac{1}{n}\mathbb{T}$ and of the finite-dimensional marginals of the process of volumes.

Hence, to complete the proof of Proposition 3.13, we only need tightness. The tightness of the first marginal is given by Proposition 3.11. On the other hand, Theorem 2 of [64] shows that the volume process converges. In particular, it is tight. This concludes the proof. □

3.3.3 Joint convergence of the perimeter process

The goal of this subsection is to prove the joint convergence of the last marginal in Theorem 3.2.

Proof of Theorem 3.2. By Proposition 3.13, the first two marginal converge in distribution to $\left(\mathcal{P}, \left(3V_r\right)_{r\geq 0}\right)$, where $V_r = |\overline{B_r}(\mathcal{P})|$. We also know by Theorem 2 of [64] that the processes $\left(\left(\frac{1}{n^2}|\partial\overline{B_{rn}}(\mathbb{T})|\right)_{r\geq 0}\right)_{n\geq 0}$ converge, so they are tight, so the triplet in Theorem 3.2 is tight. Hence, it is enough to prove uniqueness of the limit. Let (n_k) be a subsequence along which it converges in distribution to a triplet

$$\left(\mathcal{P}, \left(3V_r\right)_{r\geq 0}, \left(\frac{3}{2}\widetilde{P}_r\right)_{r\geq 0}\right),$$

where $\widetilde{P} = \big(\widetilde{P}_r\big)_{r\geq 0}$ is a càdlàg process. We also write $P_r = |\partial\overline{B_r}(\mathcal{P})|$, and we want to show $\widetilde{P} = P$.

On the one hand, Theorem 1.3 of [63] describes the joint distribution of P and V. There is a measurable enumeration (s_i) of the jumps of P and an i.i.d. sequence (ξ_i) of variables with distribution $\nu(\mathrm{d}x) = \frac{e^{-1/2x}}{\sqrt{2\pi x^5}}\mathbb{1}_{x>0}\,\mathrm{d}x$ such that, for all $r \geq 0$, we have

$$V_r = \sum_{s_i \leq r} (\Delta P_{s_i})^2 \xi_i.$$

On the other hand, Theorem 2 of [64] shows that the second and third marginals of Theorem 3.2 converge, and identifies the limit. Hence, it gives the distribution of the couple $(\widetilde{P}_r, V_r)_{r\geq 0}$ (see Section 6.1 of [64] for the computation of the constants for type-I triangulations). We get

$$(\widetilde{P}_r, V_r)_{r\geq 0} \overset{(d)}{=} (P_r, V_r)_{r\geq 0}.$$

To prove that $\widetilde{P}_r = P_r$, we show that it is possible to "track back" P_r from the process $(V_r)_{r\geq 0}$, which is done in the following lemma. We will need the following notation: for a nondecreasing function f and $a, b, h \in \mathbb{R}^+$ with $a \leq b$, we write $N_a^b(f, h)$ for the number of jumps of f in $[a, b]$ of height at least h.

Lemma 3.14. For all $r \geq 0$, we have

$$\lim_{\delta\to 0} \delta^{-1} \lim_{\varepsilon\to 0} \varepsilon^{3/4} N_r^{r+\delta}(V, \varepsilon) = cP_r$$

almost surely, where $c = \frac{2^{1/4}}{\pi\sqrt{3}}\Gamma\big(\frac{3}{4}\big)$.

Once this lemma is proved, Theorem 3.2 follows easily. Indeed, since $(\widetilde{P}, V) \overset{(d)}{=} (P, V)$, for any $r \geq 0$ the variables $\widetilde{P}_r$ and P_r are both the a.s. limits of the same quantity. Hence, almost surely, $\widetilde{P}_r = P_r$ for every $r \in \mathbb{Q}^+$. Since $\widetilde{P}$ and P are càdlàg we have $\widetilde{P} = P$ a.s. Hence, the sequence of triplets has only one subsequential limit, which proves the theorem. □

Remark 3.15. Note that Proposition 1.1 of [63] provides another way to "read" P on the measured metric space $\mathcal{P}$. However, it involves the volumes of the balls $B_{r+\varepsilon}$, which are not given by the process V. Our lemma is also quite similar to the main theorem of [100], although much easier to prove.

Proof of Lemma 3.14. Let S^+ be the stable spectrally negative Lévy process of index $\frac{3}{2}$ conditioned to stay positive. We normalize it in such a way that its characteristic exponent is $\psi(\lambda) = \sqrt{8/3}\,\lambda^{3/2}$. Let also (t_i) be a measurable enumeration of the jumps of S^+, and let (ξ_i') be a sequence of i.i.d. random variables with distribution ν. We write

$$Q_r^+ = \sum_{t_i \leq r} \xi_i' (\Delta S_{t_i}^+)^2.$$

We first show that almost surely, for any $0 < a < b$, we have

$$\lim_{\varepsilon\to 0} \varepsilon^{3/4} N_a^b(Q^+, \varepsilon) = c(b - a). \tag{3.16}$$

By monotonicity, it is enough to prove it for $a, b \in \mathbb{Q}_+^*$. Hence, it is enough to prove it for fixed a and b. Let Q be the same process as Q^+ but constructed from a non-conditioned stable Lévy process S instead of S^+. Then Q is a subordinator whose Lévy measure σ is the image of

$\mu \otimes \nu$ under $(x, y) \to x^2 y$, where μ is the Lévy measure of S. An easy computation shows that $\sigma([\varepsilon, +\infty[) = c\varepsilon^{-3/4}$ for all ε. Hence, equation (3.16) for Q instead of Q^+ follows from a law of large numbers. But since S^+ is absolutely continuous with respect to S on $[a, b]$, equation (3.16) also holds for Q^+.

We now recall the law of (P, V) as described in Section 4.4 of [64]. It has the same distribution as $(S^+_{\tau_r}, Q^+_{\tau_r})_{r\geq 0}$, where $\tau_r = \inf\{s \geq 0 | \int_0^s \frac{\mathrm{d}u}{S^+_u} \geq r\}$ for every $r \geq 0$. Hence, we have $N_r^{r+\delta}(V, \varepsilon) = N_{\tau_r}^{\tau_{r+\delta}}(Q^+, \varepsilon)$. By (3.16), for any r and δ we have (since a.s. (3.16) holds for any $0 < a < b$, we can take a and b random)

$$\lim_{\varepsilon \to 0} \varepsilon^{3/4} N_{\tau_r}^{\tau_{r+\delta}}(Q^+, \varepsilon) = c(\tau_{r+\delta} - \tau_r).$$

Now, it is easy to see, by the right-continuity of S^+ at τ_r, that $\delta^{-1}(\tau_{r+\delta} - \tau_r) \xrightarrow[\delta\to 0]{a.s.} S^+_{\tau_r} = P_r$. This finishes the proof. □

3.3.4 The hyperbolic Brownian plane

The goal of this section is to prove Theorem 3.1. We will write

$$\varphi(p, v) = e^{-2v} e^p \int_0^1 e^{-3px^2} \mathrm{d}x$$

in the whole proof of the theorem. Moreover, if t is a triangulation with a simple hole of perimeter p' and with v' vertices in total and $\lambda \in (0, \lambda_c]$, we write

$$\varphi_\lambda(p', v') = \frac{\mathbb{P}\left(t \subset \mathbb{T}_{\lambda_n}\right)}{\mathbb{P}\left(t \subset \mathbb{T}\right)}. \tag{3.17}$$

Since $\mathbb{T}_{\lambda_n}$ and $\mathbb{T}$ are Markovian, this only depends on p' and v' and not on t.

Proposition 3.16. Let (λ_n) be a sequence of numbers in $\left(0, \lambda_c\right]$ that satisfies (3.2). Let $r > 0$ and let (p'_n) and (v'_n) be two sequences of positive integers such that $\frac{p'_n}{n^2} \to \frac{3}{2}p$ and $\frac{v'_n}{n^4} \to 3v$. Then

$$\varphi_{\lambda_n}(p'_n, v'_n) \xrightarrow[n\to+\infty]{} \varphi(p, v).$$

Proof. This is just a matter of gathering our estimates in Section 1.1 and 1.2 together. Let us proceed: for every n, let $h_n \in \left(0, \frac{1}{4}\right]$ be such that $\lambda_n = \frac{h_n}{(1+8h_n)^{3/2}}$. It is easy to check that

$$h_n = \frac{1}{4} - \frac{1}{2n^2} + o\Big(\frac{1}{n^2}\Big). \tag{3.18}$$

We know from (3.11) that

$$\begin{aligned}
\varphi_{\lambda_n}(p'_n, v'_n) &\underset{\text{Def. 3.3}}{=} \frac{\lambda_n^{v_n} C_{p_n}(\lambda_n)}{\lambda_c^{v_n} C_{p_n}(\lambda_c)} \\
&\underset{(3.15)}{=} \Big(\frac{\lambda_n}{\lambda_c}\Big)^{v_n} \frac{{p_n!}^2}{2\sqrt{3} \times 3^{p_n} p_n (2p_n)!} \frac{1}{\lambda_n} \Big(8 + \frac{1}{h_n}\Big)^{p_n - 1} \sum_{q=0}^{p_n-1} \binom{2q}{q} h_n^q \\
&= \Big(\frac{\lambda_n}{\lambda_c}\Big)^{3vn^4 + o(n^4)} \Big(\frac{2}{3} + \frac{1}{12h_n}\Big)^{p_n - 1} \frac{\lambda_c}{\lambda_n} \frac{4^{p_n} p_n!^2}{2p_n(2p_n)!} \sum_{q=0}^{p_n-1} \binom{2q}{q} h_n^q.
\end{aligned}$$

The first factor converges to e^{-2v} because $\frac{\lambda_n}{\lambda_c} = 1 - \frac{2}{3n^4} + o\Big(\frac{1}{n^4}\Big)$. The second factor is also easy to estimate:

$$\Big(\frac{2}{3} + \frac{1}{12h_n}\Big)^{p_n - 1} = \Big(1 + \frac{2}{3n^2} + o\Big(\frac{1}{n^2}\Big)\Big)^{\frac{3p}{2}n^2 + o(n^2)} \xrightarrow[n\to+\infty]{} e^p.$$

The third one converges to one. By the Stirling formula, we have $\frac{4^{p_n} p_n!^2}{2p_n(2p_n)!} \sim \frac{\sqrt{\pi}}{2\sqrt{p_n}}$. Finally, Lemma 3.6 gives an asymptotic equivalent of the last factor and we are done. □

The last proposition is more or less equivalent to vague convergence. We now need to show that no mass "escapes" at infinity, i.e. that the total mass of the limit measure is 1.

Lemma 3.17. Recall that $P_r = |\partial\overline{B_r}(\mathcal{P})|$ and $V_r = |\overline{B_r}(\mathcal{P})|$. For every $r \geq 0$, we have

$$\mathbb{E}\big[\varphi(V_r, P_r)\big] = 1.$$

Proof. We use the expression of the Laplace transform of (P_r, V_r) that is computed in [63]. First, Theorem 1.4 of [63] computes $\mathbb{E}\left[e^{-\mu V_r} | P_r = \ell\right]$ for $\mu, \ell \geq 0$. We apply it with $\mu = 2$:

$$\mathbb{E}\Big[\exp(-2V_r)\Big|P_r = \ell\Big] = 2\sqrt{2}r^3 \frac{\cosh(\sqrt{2}r)}{\sinh^3(\sqrt{2}r)} \exp\Big(-\ell\Big(3\coth^2(\sqrt{2}r) - 2 - \frac{3}{2r^2}\Big)\Big).$$

We now apply Fubini's theorem, and use the Laplace transform of P_r given by Theorem 1.2 of [63]:

$$\begin{aligned}
\mathbb{E}\big[\varphi(V_r, P_r)\big] &= \mathbb{E}\Big[\mathbb{E}\big[\exp(-2V_r)\big|P_r\big]\exp\big(P_r\big)\int_0^1 \exp\big(-3P_r x^2\big)\mathrm{d}x\Big] \\
&= 2\sqrt{2}r^3\frac{\cosh(\sqrt{2}r)}{\sinh^3(\sqrt{2}r)}\int_0^1 \mathbb{E}\left[\exp\Big(-P_r\Big(3\coth^2(\sqrt{2}r) - 2 - \frac{3}{2r^2}\Big)\Big)\exp\Big((1-3x^2)P_r\Big)\right]\mathrm{d}x \\
&= 2\sqrt{2}r^3\frac{\cosh(\sqrt{2}r)}{\sinh^3(\sqrt{2}r)}\int_0^1 \left(1 + \frac{2r^2}{3}\Big(3\coth^2(\sqrt{2}r) - 2 - \frac{3}{2r^2} - 1 + 3x^2\Big)\right)^{-3/2}\mathrm{d}x \\
&= 2\sqrt{2}r^3\frac{\cosh(\sqrt{2}r)}{\sinh^3(\sqrt{2}r)}\int_0^1 \left(2r^2\coth^2(\sqrt{2}r) - 2r^2 + 2r^2x^2\right)^{-3/2}\mathrm{d}x \\
&= \frac{\cosh(\sqrt{2}r)}{\sinh^3(\sqrt{2}r)}\int_0^1 \Big(\frac{1}{\sinh^2(\sqrt{2}r)} + x^2\Big)^{-3/2}\mathrm{d}x \\
&= \cosh(\sqrt{2}r)\int_0^1 \Big(1 + \sinh^2(\sqrt{2}r)x^2\Big)^{-3/2}\mathrm{d}x \\
&= \cosh(\sqrt{2}r)\frac{1}{\sqrt{1+\sinh^2(\sqrt{2}r)}} \\
&= 1,
\end{aligned}$$

where at the end we used the fact that $(1+ax^2)^{-3/2}$ is the derivative of $\frac{x}{\sqrt{1+ax^2}}$. □

Our main theorem is now easy to prove.

Proof of Theorem 3.1. In this proof, if X is a measured metric space, we will write $\frac{1}{n}X$ for the metric space obtained from X by multiplying all distances by $\frac{1}{n}$ and the measure by $\frac{1}{3n^4}$.

Let $r > 0$ and let $\mathbb{T}$ be a type-I UIPT. Almost surely, the processes P and V have no jump at r and at $2r$, so we have convergence of the one-dimensional marginal at $2r$ in Theorem 3.2. For

all $n \geq 1$, let $P'_n = |\partial B^\bullet_{2rn}(\mathbb{T})|$ and $V'_n = |B^\bullet_{2rn}(\mathbb{T})|$. By the Skorokhod representation theorem, we may assume, as $n \to +\infty$, the convergences

$$\begin{cases} \frac{1}{n}\mathbb{T} & \xrightarrow[LGHP]{a.s.} & \mathcal{P} \\ \frac{1}{n^2}P'_n & \xrightarrow{a.s.} & \frac{3}{2}P_{2r} \\ \frac{1}{n^4}V'_n & \xrightarrow{a.s.} & 3V_{2r}. \end{cases} \tag{3.19}$$

We have already verified in the proof of Theorem 3.13 that the assumptions of Proposition 3.10 are satisfied. By the first point of Proposition 3.10, the convergence of $\frac{1}{n}\mathbb{T}$ to $\mathcal{P}$ implies

$$\frac{1}{n}\overline{B_{rn}}(\mathbb{T}) \xrightarrow[GHP]{a.s.} \overline{B_r}(\mathcal{P}). \tag{3.20}$$

Let $\mathscr{G}$ denote the GHP space and let f be a bounded, continuous function from $\mathscr{G}$ to $\mathbb{R}^+$. We recall that the space $\overline{B_{rn}}(\mathbb{T})$ is a measurable function of the map $B^\bullet_{2rn}(\mathbb{T})$. Hence, there is a function $\widetilde{f}$ (that depends on n) from the set of finite triangulations to $\mathbb{R}^+$ such that $f\Big(\frac{1}{n}\overline{B_{rn}}(t)\Big) = \widetilde{f}\big(B^\bullet_{2rn}(t)\big)$ for any one-ended, infinite triangulation t. Hence, by the definition (3.17) of φ_λ, we can write

$$\begin{aligned} \mathbb{E}\Big[f\Big(\frac{1}{n}\overline{B_{rn}}(\mathbb{T}_{\lambda_n})\Big)\Big] &= \mathbb{E}\Big[\widetilde{f}\big(B^\bullet_{2rn}(\mathbb{T}_{\lambda_n})\big)\Big] \\ &= \mathbb{E}\Big[\varphi_{\lambda_n}(P'_n, V'_n)\,\widetilde{f}\big(B^\bullet_{2rn}(\mathbb{T})\big)\Big] \\ &= \mathbb{E}\Big[\varphi_{\lambda_n}(P'_n, V'_n)\,f\Big(\frac{1}{n}\overline{B_{rn}}(\mathbb{T})\Big)\Big]. \end{aligned}$$

By the convergences (3.19) and (3.20), Proposition 3.16 and the continuity of f, the expression inside the expectation converges a.s. to $\varphi\big(P_{2r}, V_{2r}\big)f\big(\overline{B_r}(\mathcal{P})\big)$. By Fatou, we have

$$\liminf_n \mathbb{E}\Big[f\Big(\frac{1}{n}\overline{B_{rn}}(\mathbb{T}_{\lambda_n})\Big)\Big] \geq \mathbb{E}\Big[\varphi\big(P_{2r}, V_{2r}\big)f\big(\overline{B_r}(\mathcal{P})\big)\Big]. \tag{3.21}$$

Let $M \in \mathbb{R}^+$ be such that $f \leq M$. By Lemma 3.17 we have

$$\mathbb{E}\Big[M\varphi\big(P_{2r}, V_{2r}\big)\Big] = M,$$

so by applying (3.21) to $M - f$ we get the reverse inequality. This shows that $\frac{1}{n}\overline{B_{rn}}(\mathbb{T}_{\lambda_n})$ converges in distribution to the random metric space having density $\varphi\big(P_{2r}, V_{2r}\big)$ with respect to $\overline{B_r}(\mathcal{P})$. We denote this space by $\overline{B_r}^h$.

Moreover, let $0 < s \leq r$. We can take points y_n on the boundary of $\frac{1}{n}\overline{B_{rn}}(\mathbb{T}_{\lambda_n})$ and $y \in \overline{B_r}^h$ such that $\Big(\frac{1}{n}\overline{B_{rn}}(\mathbb{T}_{\lambda_n}), y_n\Big)$ converges in distribution to $\big(\overline{B_r}^h, y\big)$. By Proposition 3.10, we have

$$\overline{B_s}\Big(\frac{1}{n}\overline{B_{rn}}(\mathbb{T}_{\lambda_n})\Big) \xrightarrow[n\to+\infty]{(d)} \overline{B_s}\big(\overline{B_r}^h\big),$$

where the hulls are taken with respect to y_n and y respectively. More precisely, we use here Proposition 3.10 in the compact, bipointed case. But $\overline{B_s}\Big(\frac{1}{n}\overline{B_{rn}}(\mathbb{T}_{\lambda_n})\Big) = \frac{1}{n}\overline{B_{sn}}(\mathbb{T}_{\lambda_n})$, which converges in distribution to $\overline{B_s}^h$. Hence, we have

$$\overline{B_s}\big(\overline{B_r}^h\big) \overset{(d)}{=} \overline{B_s}^h,$$

i.e. the $\overline{B_r}^h$ are "consistent". By the Kolmogorov extension theorem, there is a random metric space, that we write $\mathcal{P}^h$, such that $\overline{B_r}(\mathcal{P}^h) \overset{(d)}{=} \overline{B_r}^h$ for all $r \geq 0$, and we have

$$\frac{1}{n}\mathbb{T}_{\lambda_n} \xrightarrow[n\to+\infty]{(d)} \mathcal{P}^h$$

for the local GHP distance. □

3.4 Properties of the hyperbolic Brownian plane

3.4.1 Local properties of the hyperbolic Brownian plane

The absolute continuity relation between the Brownian plane $\mathcal{P}$ and the hyperbolic Brownian plane $\mathcal{P}^h$ implies that they have the same almost sure "local" properties. This gives us the two following properties of $\mathcal{P}^h$.

Proposition 3.18. Almost surely, $\mathcal{P}^h$ has Hausdorff dimension 4 and is homeomorphic to $\mathbb{R}^2$.

Proof. First, by the absolute continuity relation with the Brownian plane, for all $r \geq 0$, the space $\overline{B_r}(\mathcal{P}^h)$ has a.s. Hausdorff dimension 4 and so does $\mathcal{P}^h$.

If (X, d, ρ) is a pointed metric space and $r > 0$, we write $U_r(X) = \bigcup_{\varepsilon>0} \overline{B_{r-\varepsilon}}(X)$. Let $r > 0$. The set $U_r(\mathcal{P})$ is a connected, open subset of $\mathcal{P}$, and it is quite easy to prove that $\mathcal{P}\backslash U_r(\mathcal{P})$ is connected. Indeed, $\overline{B_{r+1}}(\mathcal{P})\backslash U_r(\mathcal{P})$ is connected because it is the decreasing intersection of the $\overline{B_{r+1}}(\mathcal{P})\backslash\overline{B_{r-\varepsilon}}(\mathcal{P})$, which are compact, connected subsets.

The set $U_r(\mathcal{P})$ is a connected, open subset of the plane whose complement is connected, so it is homeomorphic to the open unit disk (this is a consequence of the Riemann mapping theorem). In particular, almost surely, the two following hold:

1. for all $x \in U_r(\mathcal{P})$, there is a neighbourhood of x that is homeomorphic to the unit disk,
2. any loop in $U_r(\mathcal{P})$ is contractible in $U_r(\mathcal{P})$.

By absolute continuity and since U_r is always a subset of $\overline{B_r}$, the two above items a.s. hold for $U_r(\mathcal{P}^h)$. Almost surely, they hold for any $r \in \mathbb{N}^*$. Hence, $\mathcal{P}^h$ is a noncompact, simply connected topological surface. Therefore, it is homeomorphic to the plane (for example, it is a consequence of the Riemann uniformization theorem and the fact that any topological surface may be equipped with a Riemann surface structure). □

Note that in this proof, it was important to consider $U_r(\mathcal{P})$ and not $\overline{B_r}(\mathcal{P})$. Indeed, a metric space may not be homeomorphic to the plane, even if all its hulls are homeomorphic to the closed unit disk (for exemple a closed half-plane, rooted at an interior point).

For the same reason as above, the results about the local confluence of geodesics in the Brownian map and the Brownian plane also hold for the hyperbolic Brownian plane.

Proposition 3.19. A.s., for any $\varepsilon > 0$, there is a $\delta > 0$ such that the following holds. All the geodesics in $\mathcal{P}^h$ from the root to a point at distance at least ε from the root share a common initial segment of length at least δ.

Finally, the Brownian map, the Brownian plane and the hyperbolic Brownian plane are locally isometric in the following sense.

Proposition 3.20. For any $\varepsilon > 0$, there is a $\delta > 0$ such that the following holds. It is possible to couple the Brownian map, the Brownian plane and the hyperbolic Brownian plane in such a way that their hulls of radius δ coincide with probability at least $1 - \varepsilon$.

Proof. The result for the Brownian map and the Brownian plane is Theorem 1 of [62], so we only need to prove it for $\mathcal{P}$ and $\mathcal{P}^h$. Let $\delta > 0$, and let A be a mesurable subset of the Gromov–Hausdorff–Prokhorov space. Then

$$\begin{aligned}\left|\mathbb{P}\big(\overline{B_\delta}(\mathcal{P}^h) \in A\big) - \mathbb{P}\big(\overline{B_\delta}(\mathcal{P}) \in A\big)\right| &= \left|\mathbb{E}\big[\varphi(P_\delta, V_\delta)\mathbb{1}_{\overline{B_\delta}(\mathcal{P})\in A}\big] - \mathbb{P}\big(\overline{B_\delta}(\mathcal{P}) \in A\big)\right| \\ &\leq \mathbb{E}\big[|\varphi(P_\delta, V_\delta) - 1|\big],\end{aligned}$$

which goes to 0 as $\delta \to 0$. Hence, the total variation distance between the distributions of $\overline{B_\delta}(\mathcal{P}^h)$ and $\overline{B_\delta}(\mathcal{P})$ goes to 0 as $\delta \to 0$. This proves the result by the maximal coupling theorem (see e.g. Section 2 of [70]). □

3.4.2 Unimodularity

The goal of this section is to prove that $\mathcal{P}^h$ satisfies a property that is the continuum analog of unimodularity for random graphs.

Proposition 3.21. Let ρ be the root of $\mathcal{P}^h$ and $\mu_{\mathcal{P}^h}$ its volume measure. Let f be a measurable function from the space of locally compact, bipointed measured metric spaces to $\mathbb{R}^+$. We have

$$\mathbb{E}\Big[\int_{\mathcal{P}^h} f(\mathcal{P}^h, \rho, y)\, \mu_{\mathcal{P}^h}(\mathrm{d}y)\Big] = \mathbb{E}\Big[\int_{\mathcal{P}^h} f(\mathcal{P}^h, y, \rho)\, \mu_{\mathcal{P}^h}(\mathrm{d}y)\Big].$$

To prove this result, we will need the following lemma. It roughly means that the degree of the root is independent of the large-scale geometry of the map.

Lemma 3.22. As earlier, we write $\frac{1}{n}\mathbb{T}_{\lambda_n}$ for the set of the vertices of $\mathbb{T}_{\lambda_n}$, equipped with $\frac{1}{n}$ times its graph distance and the measure giving mass $\frac{1}{3n^4}$ to each vertex. Let ρ_n be the origin of the root edge of $\mathbb{T}_{\lambda_n}$. We have the convergence

$$\Big(\deg(\rho_n), \frac{1}{n}\mathbb{T}_{\lambda_n}\Big) \xrightarrow[n\to+\infty]{(d)} \big(D, \mathcal{P}^h\big),$$

where D has the same distribution as the degree of the root in the UIPT and is independent of $\mathcal{P}^h$.

Proof. The convergence of the first marginal is obvious. Hence, it is enough to prove that for all $d \geq 1$, the spaces $\frac{1}{n}\mathbb{T}_{\lambda_n}$ conditioned on $\deg(\rho_n) = d$ converge in distribution to $\mathcal{P}^h$. More generally, we prove that for each finite triangulation t that is a possible value of $B_1^\bullet(\mathbb{T})$, the spaces $\frac{1}{n}\mathbb{T}_{\lambda_n}$ conditioned on $B_1^\bullet(\mathbb{T}_{\lambda_n}) = t$ converge to $\mathcal{P}^h$. For all $s \geq 0$, the GHP distance between $B_s(\mathbb{T}_{\lambda_n})$ and $B_s(\mathbb{T}_{\lambda_n})\backslash B_1^\bullet(\mathbb{T}_{\lambda_n})$ is tight (it is bounded by $|B_1^\bullet(\mathbb{T}_{\lambda_n})|$). Hence, for all $r \geq 0$, the GHP distance between $B_r\Big(\frac{1}{n}\mathbb{T}_{\lambda_n}\Big)$ and $B_r\Big(\frac{1}{n}\big(\mathbb{T}_{\lambda_n}\backslash B_1^\bullet(\mathbb{T}_{\lambda_n})\big)\Big)$ goes to 0 in probability as $n \to +\infty$. Hence, it is enough to prove the lemma for $\frac{1}{n}\big(\mathbb{T}_{\lambda_n}\backslash B_1^\bullet(\mathbb{T}_{\lambda_n})\big)$ instead of $\frac{1}{n}\mathbb{T}_{\lambda_n}$. But by the spatial Markov property of $\mathbb{T}_{\lambda_n}$, the distribution of $\mathbb{T}_{\lambda_n}\backslash B_1^\bullet(\mathbb{T}_{\lambda_n})$ conditionally on $B_1^\bullet(\mathbb{T}_{\lambda_n})$ only depends on $|\partial B_1^\bullet(\mathbb{T}_{\lambda_n})|$.

We recall that a *peeling algorithm* $\mathscr{A}$ is a way to assign, to every finite triangulation t with a hole, an edge $\mathscr{A}(t)$ on the boundary of the hole. Informally, $\mathscr{A}(t)$ is the next edge to be explored once the explored part of the map is equal to t. See [12] or [64] for more details. We now fix a deterministic peeling algorithm, and we explore $\mathbb{T}_{\lambda_n}$ using this algorithm. For every $p \geq 1$, we write τ_p for the first time at which the perimeter of the discovered map is equal to p. Note that, since $\mathbb{T}_{\lambda_n}$ may be identified with an infinite triangulation of a 1-gon, the times τ_p are a.s. finite even for $p = 1$ or $p = 2$. We also write $D_p(\mathbb{T}_{\lambda_n})$ for the triangulation discovered at time τ_p. By the spatial Markov property, the map $\mathbb{T}_{\lambda_n}\backslash B_1^\bullet(\mathbb{T}_{\lambda_n})$ conditionally on $|\partial B_1^\bullet(\mathbb{T}_{\lambda_n})| = p$ has the same distribution as $\mathbb{T}_{\lambda_n}\backslash D_p(\mathbb{T}_{\lambda_n})$. Hence, it is enough to prove that for any $p \geq 0$, we have

$$\frac{1}{n}\big(\mathbb{T}_{\lambda_n}\backslash D_p(\mathbb{T}_{\lambda_n})\big) \xrightarrow[n\to+\infty]{(d)} \mathcal{P}^h.$$

This easily follows from Theorem 3.1 and the fact that, for all $r \geq 0$, the GHP distance between $B_r(\mathbb{T}_{\lambda_n})$ and $B_r(\mathbb{T}_{\lambda_n})\backslash D_p(\mathbb{T}_{\lambda_n})$ is at most $\big|D_p(\mathbb{T}_{\lambda_n})\big|$, which is tight. □

We now move on to the proof of Proposition 3.21.

Proof. We claim that it is enough to prove the result for functions f such that:

(i) there is an $r \geq 0$ such that, if the distance between x and y is greater than r, then $f(X, x, y) = 0$,

(ii) there is a $v > 0$ such that, if one of the balls of radius r centered at x or y has a volume greater than v, then $f(X, x, y) = 0$,

(iii) there is an $s > r$ such that $f(X, x, y)$ only depends on the intersection of the balls of radius s around x and y in X, bipointed at x and y,

(iv) f is bounded and uniformly continuous for the (bipointed) GHP distance.

Note that assumption (iv) makes sense even if we consider the GHP (and not local GHP) distance. Indeed, by assumption (iii), we may see $f(X, x, y)$ as a function of $B_s(X, x)$, so we do not need the local GHP distance. To show our claim, assume the theorem is true for any f satisfying (i), (ii), (iii) and (iv). By the monotone convergence theorem, it is true for all indicator functions of open events satisfying (i), (ii) and (iii). By the monotone class theorem, it is true for the indicator functions of any event satisfying (i), (ii) and (iii). Now let A be an event whose indicator function satisfies (i) and (ii). By the monotone class theorem again, the event A can be approximated by events whose indicator function satisfies (i), (ii) and (iii). More precisely, for any $\varepsilon > 0$, there is an event B whose indicator function satisfies (i), (ii) and (iii), and such that

$$\mathbb{E}\Big[\int_{\mathcal{P}^h} \mathbb{1}_{A\Delta B}(\mathcal{P}^h, \rho, y)\, \mu_{\mathcal{P}^h}(\mathrm{d}y)\Big] \leq \varepsilon.$$

Hence, we can get rid of assumption (iii). We can then get rid of assumptions (i) and (ii) by monotone convergence. Therefore, the result is true for any indicator function, and finally for any nonnegative measurable function.

We now prove the theorem for a function f satisfying (i), (ii), (iii) and (iv). The idea is to use the unimodularity of (a variant of) $\mathbb{T}_{\lambda_n}$, and take the scaling limit. Let (λ_n) be a sequence satisfying the assumptions of Theorem 3.1. We first remark that the triangulations $\mathbb{T}_{\lambda_n}$ are invariant by rerooting along the simple random walk. This is the type-I analog of (a part of) Proposition 9 in [58], and the proof is exactly the same. Moreover, invariance along the simple random walk and unimodularity are closely related by Proposition 2.5 of [27]. More precisely, we write $\widehat{\mathbb{T}}_{\lambda_n}$ for the map $\mathbb{T}_{\lambda_n}$ biased by the inverse of the degree of its root ρ_n. Then $\widehat{\mathbb{T}}_{\lambda_n}$ is unimodular. Hence, we have

$$\mathbb{E}\left[\frac{1}{3n^4}\sum_{y\in\mathbb{T}_{\lambda_n}} f\Big(\frac{1}{n}\widehat{\mathbb{T}}_{\lambda_n}, \rho_n, y\Big)\right] = \mathbb{E}\left[\frac{1}{3n^4}\sum_{y\in\mathbb{T}_{\lambda_n}} f\Big(\frac{1}{n}\widehat{\mathbb{T}}_{\lambda_n}, y, \rho_n\Big)\right],$$

where ρ_n is the root vertex of $\widehat{\mathbb{T}}_{\lambda_n}$. We may restrict ourselves to $y \in B_{rn}(\widehat{\mathbb{T}}_{\lambda_n})$ thanks to assumption (i). By the definition of $\widehat{\mathbb{T}}_{\lambda_n}$, the last equation can be rewritten

$$\mathbb{E}\left[\frac{1}{\deg(\rho_n)}\frac{1}{3n^4}\sum_{y\in B_{rn}(\mathbb{T}_{\lambda_n})} f\Big(\frac{1}{n}\mathbb{T}_{\lambda_n}, \rho_n, y\Big)\right] = \mathbb{E}\left[\frac{1}{\deg(\rho_n)}\frac{1}{3n^4}\sum_{y\in B_{rn}(\mathbb{T}_{\lambda_n})} f\Big(\frac{1}{n}\mathbb{T}_{\lambda_n}, y, \rho_n\Big)\right].$$

Hence, in order to prove the proposition, it is enough to prove that the left-hand side converges as $n \to +\infty$ to the left-hand side in the statement of the proposition, multiplied by $\mathbb{E}\big[\frac{1}{D}\big]$. The proof of the same fact for the right-hand side will be similar. By the Skorokhod representation theorem, we may assume the convergence in Lemma 3.22 is almost sure. By the dominated convergence theorem (the domination follows from assumption (ii) and the fact that f is bounded), it is enough to prove

$$\frac{1}{3n^4}\sum_{y\in B_{rn}(\mathbb{T}_{\lambda_n})} f\Big(\frac{1}{n}\mathbb{T}_{\lambda_n}, \rho_n, y\Big) \xrightarrow[n\to+\infty]{a.s.} \int_{B_r(\mathcal{P}^h)} f(\mathcal{P}^h, \rho, y)\mu_{\mathcal{P}^h}(dy). \tag{3.22}$$

This follows from assumptions (iii) and (iv) and the GHP convergence of $\frac{1}{n}B_{(r+s)n}(\mathbb{T}_{\lambda_n})$ to $B_{r+s}(\mathcal{P}^h)$. □

Remark 3.23. The same result is true for the Brownian plane. The proof is essentially the same, the only change is that we need to use Proposition 3.11 instead of Theorem 3.1.

3.4.3 Identification of the perimeter and volume processes

We first explain what we mean by biasing a process by a martingale, which will be used a lot in what follows. Let X be an adapted process, and let M be a martingale for the underlying filtration with $\mathbb{E}\big[M_0\big] = 1$. We say that a process Y is the process X biased by M if for all $r_0 > 0$, the process $(Y_r)_{0\le r\le r_0}$ has the same distribution as $(X_r)_{0\le r\le r_0}$ biased by M_{r_0}. The martingale property of M shows the consistency for different values of r_0.

The hyperbolic Brownian plane is a biased version of the Brownian plane. Hence, it is naturally equipped with a perimeter and a volume process inherited from those of the Brownian plane. We denote by $\big(P^h_r\big)_{r\ge0}$ and $\big(V^h_r\big)_{r\ge0}$ the perimeter and volume processes of $\mathcal{P}^h$. More precisely, the proof of Theorem 3.1 gives the following joint convergences in distribution:

$$\begin{cases} \frac{1}{n}\mathbb{T}_{\lambda_n} & \xrightarrow[n\to+\infty]{} \quad \mathcal{P}^h, \\ \Big(\frac{1}{n^2}\big|\partial B^\bullet_{rn}(\mathbb{T}_{\lambda_n})\big|\Big)_{r\ge0} & \xrightarrow[n\to+\infty]{} \quad \big(\frac{3}{2}P^h_r\big)_{r\ge0}, \\ \Big(\frac{1}{n^4}\big|B^\bullet_{rn}(\mathbb{T}_{\lambda_n})\big|\Big)_{r\ge0} & \xrightarrow[n\to+\infty]{} \quad \big(3V^h_r\big)_{r\ge0}. \end{cases}$$

As previously, the first convergence is for the local GHP distance and the other two for the Skorokhod topology. Moreover, for all $r_0 > 0$, the triplet $\Big(\overline{B_{r_0}}(\mathcal{P}^h), \big(P^h_r\big)_{0\le r\le r_0}, \big(V^h_r\big)_{0\le r\le r_0}\Big)$ has the same distribution as $\Big(\overline{B_{r_0}}(\mathcal{P}), \big(P_r\big)_{0\le r\le r_0}, \big(V_r\big)_{0\le r\le r_0}\Big)$ biased by $\varphi(P_{2r_0}, V_{2r_0})$.

This description implies that these two triplets have the same a.s. properties as for the Brownian plane. In particular, P^h and V^h are both càdlàg processes and can be expressed as measurable functions of $\mathcal{P}^h$. Indeed, we have $V^h_r = |\overline{B_r}(\mathcal{P}^h)|$ for all $r \ge 0$, and Proposition 1.1 of [63] (already recalled in (3.1)) gives the convergence

$$\frac{1}{\varepsilon^2}|B_{r+\varepsilon}(\mathcal{P}^h)\backslash\overline{B_r}(\mathcal{P}^h)| \xrightarrow[\varepsilon\to0]{} P^h_r$$

in probability.

On the other hand, the perimeter and volume of the map $B^\bullet_{rn}(\mathbb{T}_{\lambda_n})$ only depend on $B^\bullet_{rn}(\mathbb{T}_{\lambda_n})$. Hence, if we apply the proof of Theorem 3.1 to the pair $\Big(\frac{1}{n^2}\big|\partial B^\bullet_{rn}(\mathbb{T}_{\lambda_n})\big|, \frac{1}{n^4}\big|B^\bullet_{rn}(\mathbb{T}_{\lambda_n})\big|\Big)_{r\ge0}$ instead of the triplet, we only need to bias by $\varphi(P_{r_0}, V_{r_0})$ instead of $\varphi(P_{2r_0}, V_{2r_0})$. We obtain the following result.

Lemma 3.24. The pair of processes (P^h, V^h) has the same distribution as (P, V) biased by $\varphi(P, V)$.

The goal of this section is to identify the two processes P^h and V^h in a more convenient way as expressed in Theorem 3.3. Before moving on to the proof of Theorem 3.3, we recall the description of (P, V) given by [63]. Let Z be the critical continuous-state branching process with branching mechanism $\sqrt{\frac{8}{3}}\lambda^{3/2}$ for $\lambda > 0$. We also recall that the measure ν is defined by $\nu(\mathrm{d}x) = \frac{e^{-1/2x}}{\sqrt{2\pi x^5}}\mathbb{1}_{x>0}\,\mathrm{d}x$. Then we can restate some results of [63] as follows.

Theorem 3.5 ([63], Proposition 1.2 (ii) and Theorem 1.3)**.**

1) The perimeter process P of the Brownian plane has the same distribution as the time-reversal of Z, started from $+\infty$ at time $-\infty$, and conditioned to die at time 0.

2) Conditionally on P, the process V has the same distribution as the process

$$\Big(\sum_{t_i\le r}(\Delta P_{t_i})^2\,\xi_i\Big)_{r\ge 0},$$

where (t_i) is a measurable enumeration of the jumps of P, and the variables ξ_i are i.i.d. with distribution ν.

We will first prove the second part of Theorem 3.3, that is, we determine the distribution of V^h conditionally on P^h. We recall that for all $\delta>0$, the measure ν_δ on $\mathbb{R}$ is defined by

$$\nu_\delta(\mathrm{d}x)=\frac{\delta^3 e^{2\delta}}{1+2\delta}\frac{e^{-\frac{\delta^2}{2x}-2x}}{\sqrt{2\pi x^5}}\mathbb{1}_{x>0}\,\mathrm{d}x.$$

If $\xi^h(\delta)$ is a random variable with distribution ν_δ, we have, for all $\beta\ge 0$,

$$\mathbb{E}\big[e^{-\beta\xi^h(\delta)}\big]=\frac{(1+\delta\sqrt{4+2\beta})e^{-\delta\sqrt{4+2\beta}}}{(1+2\delta)e^{-2\delta}}. \tag{3.23}$$

Notice for further use that $\xi^h(\delta)$ has the same distribution as $\delta^2\xi$ biased by $e^{-2\delta^2\xi}$, where ξ is a random variable with density ν.

Proof of Theorem 3.3. We fix $r_0>0$. We write $D([0,r_0])$ for the Skorokhod space on $[0,r_0]$. We write $g(p)=e^p\int_0^1 e^{-3x^2p}\mathrm{d}x$. Let $t_1,t_2,\dots$ (resp. $t_1^h,t_2^h,\dots$) be the (random) jump times of the process P (resp. P^h) up to time r_0, ordered by decreasing size of jumps $|\Delta P_{t_i}|$. Let f be any measurable function $f:D([0,r_0])\longrightarrow\mathbb{R}^+$ and let $u_1,u_2,\dots\ge 0$. Since V is a pure jump process and only jumps at jump times of P, by Lemma 3.24 we have

$$\mathbb{E}\Big[f\big((P^h_r)_{0\le r\le r_0}\big)\exp\Big(-\sum_{i\ge 0}u_i|\Delta V^h_{t_i^h}|\Big)\Big]=\mathbb{E}\Big[f\big((P_r)_{0\le r\le r_0}\big)g(P_{r_0})\exp\Big(-\sum_{i\ge 0}(u_i+2)|\Delta V_{t_i}|\Big)\Big]. \tag{3.24}$$

By Theorem 3.5, conditionally on P, the jumps ΔV_{t_i} are independent and are distributed as $\big(\Delta P_{t_i}\big)^2\xi$, where ξ has law ν. Hence, the last display can be rewritten as

$$\mathbb{E}\Bigg[f\big((P_r)_{0\le r\le r_0}\big)g(P_{r_0})\prod_{i\ge 0}\mathbb{E}\big[e^{-2(\Delta P_{t_i})^2\xi}|\Delta P_{t_i}\big]\frac{\mathbb{E}\big[e^{-(u_i+2)(\Delta P_{t_i})^2\xi}|\Delta P_{t_i}\big]}{\mathbb{E}\big[e^{-2(\Delta P_{t_i})^2\xi}|\Delta P_{t_i}\big]}\Bigg]. \tag{3.25}$$

From the distribution of ξ we compute easily, for $\alpha\ge 0$,

$$\mathbb{E}\big[e^{-2\alpha^2\xi}\big]=(1+2\alpha)e^{-2\alpha}.$$

By combining this with Equation (3.23), (3.25) becomes

$$\mathbb{E}\Big[f\big((P_r)_{0\le r\le r_0}\big)g(P_{r_0})\prod_{i\ge 0}(1+2|\Delta P_{t_i}|)e^{-2|\Delta P_{t_i}|}\prod_{i\ge 0}\mathbb{E}\big[e^{-u_i\xi^h(|\Delta P_{t_i}|)}|\Delta P_{t_i}\big]\Big],$$

where $\xi^h(\delta)$ has law ν_δ.

This expression shows two things. First, it proves point 2) of Theorem 3.3, that is, conditionally on $(P^h_r)_{0\le r\le r_0}$, the jumps $\Delta V^h_{t_i^h}$ are independent and of law $\nu_{|\Delta P^h_{t_i^h}|}$. Second, by setting $u_i=0$ for every i, it proves that the density of the process $(P^h_r)_{0\le r\le r_0}$ with respect to $(P_r)_{0\le r\le r_0}$ is given by

$$g(P_{r_0})\prod_{i\ge 0}(1+2|\Delta P_{t_i}|)e^{-2|\Delta P_{t_i}|}. \tag{3.26}$$

We now move on to characterizing the law of P^h in a nicer way. Since we know by Theorem 3.5 that P is a reversed branching process with mechanism $\sqrt{\frac{8}{3}}u^{3/2}$, it seems natural to first study the effect of the density (3.26) on its associated Lévy process. So let S be the spectrally *positive* $\frac{3}{2}$-stable Lévy process. We normalize it in such a way that, for all $t, u \geq 0$, we have

$$\mathbb{E}\big[e^{-uS_t}\big] = \exp\Big(\sqrt{\frac{8}{3}}tu^{3/2}\Big) = \exp\Big(t\int_0^{+\infty}(e^{-ux} - 1 + ux)\mu(\mathrm{d}x)\Big),$$

where $\mu(\mathrm{d}x) = \sqrt{\frac{3}{2\pi}}x^{-5/2}\mathbb{1}_{x>0}\,\mathrm{d}x$.

For every $t \geq 0$, let $M_t = e^{-S_t}\prod_{t_i\leq t}(1+2|\Delta S_{t_i}|)e^{-2|\Delta S_{t_i}|}$, where (t_i) is a measurable enumeration of the jumps of S.

Lemma 3.25. The process M is a martingale with respect to the natural filtration associated to S.

Proof. In the whole proof we will write $f(x) = (1+2x)e^{-2x}$. To prove the lemma, we first note that $f(x) \leq 1$ for all $x \geq 0$, so the product defining M is always well-defined, and $\mathbb{E}\big[M_t\big] \leq \mathbb{E}\big[e^{-S_t}\big] < +\infty$ for all t. Let $(\mathscr{F}_t)_{t\geq 0}$ be the natural filtration associated to S. Since S is a Lévy process, it is easy to see that, if $s \leq t$, then $\mathbb{E}\big[M_t|\mathscr{F}_s\big] = M_s\mathbb{E}\big[M_{t-s}\big]$. Hence, it is enough to prove that $\mathbb{E}\big[M_t\big] = 1$ for all $t \geq 0$.

We claim that, for all $u \geq 0$ and $t \geq 0$, we have

$$\mathbb{E}\Big[e^{-uS_t}\prod_{t_i\leq t} f(\Delta S_{t_i})\Big] = e^{t\psi(u)}, \tag{3.27}$$

where $\psi(u) = \int_0^{+\infty}(e^{-ux}f(x) - 1 + ux)\mu(\mathrm{d}x) = \sqrt{\frac{8}{3}}\frac{u^2+u-2}{\sqrt{u+2}}$. Once this is proved, we have in particular $\psi(1) = 0$ and the lemma follows. To prove (3.27), for $\varepsilon > 0$, we write

$$S_t^\varepsilon = \sum_{\substack{t_i\leq t\\ \Delta S_{t_i}\geq\varepsilon}} \Delta S_{t_i} - t\int_\varepsilon^{+\infty} x\,\mu(\mathrm{d}x).$$

We know that S_t is the a.s. limit of S_t^ε as $\varepsilon \to 0$. Moreover, for every ε, we have

$$\mathbb{E}\big[e^{-2uS_t^\varepsilon}\big] = \exp\Big(t\int_\varepsilon^{+\infty}(e^{-2ux} - 1 + 2ux)\mu(\mathrm{d}x)\Big) \leq \mathbb{E}\big[e^{-2uS_t}\big].$$

Hence, since $f \leq 1$, the variables $e^{-uS_t^\varepsilon}\prod_{\substack{t_i\leq t\\ \Delta S_{t_i}\geq\varepsilon}} f(\Delta S_{t_i})$ are bounded in L^2 as $\varepsilon \to 0$. Therefore, they are uniformly integrable, so the left-hand side of (3.27) is equal to

$$\lim_{\varepsilon\to 0}\mathbb{E}\Big[e^{-uS_t^\varepsilon}\prod_{\substack{t_i\leq t\\ \Delta S_{t_i}\geq\varepsilon}} f(\Delta S_{t_i})\Big] = \lim_{\varepsilon\to 0} e^{tu\int_\varepsilon^{+\infty} x\mu(\mathrm{d}x)}\mathbb{E}\Big[\prod_{\substack{t_i\leq t\\ \Delta S_{t_i}\geq\varepsilon}} e^{-u\Delta S_{t_i}}f(\Delta S_{t_i})\Big].$$

By the exponential formula for Poisson point processes, the expectation in the right-hand side is equal to

$$\exp\Big(t\int_\varepsilon^{+\infty}\big(e^{-ux}f(x) - 1\big)\mu(\mathrm{d}x)\Big),$$

which proves (3.27) by letting $\varepsilon \to 0$. □

Since M is a martingale with $M_0 = e^{-S_0}$, we may consider the process S biased by $e^{S_0}M$. We denote it by S^h. From the form of M it is easy to prove that S^h is a Lévy process. Moreover, by (3.27), it holds for all $u \geq 0$ that

$$\mathbb{E}\big[e^{-uS_t^h}\big] = \mathbb{E}\big[M_t e^{-uS_t}\big] = e^{t\psi(u+1)} = \exp\big(\sqrt{\frac{8}{3}}u\sqrt{u+3}\big). \tag{3.28}$$

We can also compute the jump measure of S^h, which is given by

$$\mu^h(\mathrm{d}x) = (1+2x)e^{-2x}\mu(dx) = \sqrt{\frac{3}{2\pi}}\frac{1+2x}{x^{5/2}}e^{-2x}\mathbb{1}_{x>0}\,\mathrm{d}x. \tag{3.29}$$

In order to study the continuous-state branching process (CSBP) associated to S^h via the Lamperti transform, we consider the process S^h started from $x > 0$ and write $\tau = \inf\{t | S_t^h = 0\}$. Note that τ is a.s. finite since S^h drifts to $-\infty$ and has only positive jumps. We claim that $(S^h_{t\wedge\tau})_{t\geq 0}$ under $\mathbb{P}_x$ has density $e^x M_\tau$ with respect to $(S_{t\wedge\tau})_{t\geq 0}$ under $\mathbb{P}_x$. To prove it, we write $\tau_n = \frac{\lceil 2^n\tau\rceil}{2^n}$. For all $t_1,\dots,t_k \geq 0$ and $f:\mathbb{R}^k \to \mathbb{R}$ bounded, we have

$$\begin{aligned}
\mathbb{E}\big[f(S^h_{t_1\wedge\tau},\dots,S^h_{t_k\wedge\tau})\big] &= \lim_{n\to+\infty}\mathbb{E}\big[f(S^h_{t_1\wedge\tau_n},\dots,S^h_{t_k\wedge\tau_n})\big] \\
&= \lim_{n\to+\infty}\sum_{i=1}^{\infty}\mathbb{E}\big[f(S^h_{t_1\wedge\tau_n},\dots,S^h_{t_k\wedge\tau_n})\mathbb{1}_{\tau_n=i/2^n}\big] \\
&= \lim_{n\to+\infty}\sum_{i=1}^{\infty}\mathbb{E}\big[e^x M_{i/2^n} f(S_{t_1\wedge\tau_n},\dots,S_{t_k\wedge\tau_n})\mathbb{1}_{\tau_n=i/2^n}\big] \\
&= \lim_{n\to+\infty}\mathbb{E}\big[e^x M_{\tau_n} f(S_{t_1\wedge\tau_n},\dots,S_{t_k\wedge\tau_n})\big] \\
&= \mathbb{E}\big[e^x M_\tau f(S_{t_1\wedge\tau},\dots,S_{t_k\wedge\tau})\big],
\end{aligned}$$

which proves the claim.

We now introduce Z^h, the CSBP with branching mechanism $\psi^h(u) = \sqrt{\frac{8}{3}}u\sqrt{u+3}$ that is associated with S^h via the Lamperti transform. We also recall that Z is the CSBP with branching mechanism $\sqrt{\frac{8}{3}}u^{3/2}$. Since the Lamperti transform is a measurable function of the Lévy process, the process Z^h started from x has density

$$e^x\prod_{t_i}(1+2|\Delta Z_{t_i}|)e^{-2|\Delta Z_{t_i}|} \tag{3.30}$$

with respect to Z started from x.

We will now do the same construction as Curien and Le Gall in Section 2.1 of [63] with this new branching mechanism. The semigroup of Z^h is characterized as follows: for all $\lambda > 0$ and $x, t \geq 0$ we have

$$\mathbb{E}_x\big[e^{-\lambda Z_t^h}\big] = e^{-xu_t^h(\lambda)},$$

where

$$\frac{\mathrm{d}u_t^h(\lambda)}{\mathrm{d}t} = -\sqrt{\frac{8}{3}}u_t^h(\lambda)\sqrt{u_t^h(\lambda)+3} \quad \text{and} \quad u_0(\lambda) = \lambda. \tag{3.31}$$

The solution of this equation is

$$u_t^h(\lambda) = \frac{3}{\sinh^2\Big(\operatorname{argsh}\sqrt{\frac{3}{\lambda}} + \sqrt{2}t\Big)}.$$

This gives

$$\mathbb{E}_x\big[e^{-\lambda Z_t^h}\big] = \exp\Big(-\frac{3x}{\sinh^2\Big(\operatorname{argsh}\sqrt{\frac{3}{\lambda}}+\sqrt{2}t\Big)}\Big)$$

and, by differentiating with respect to λ,

$$\mathbb{E}_x\big[Z_t^h e^{-\lambda Z_t^h}\big] = \frac{3\sqrt{3}x}{\lambda\sqrt{\lambda+3}}\frac{\cosh\Big(\operatorname{argsh}\sqrt{\frac{3}{\lambda}}+\sqrt{2}t\Big)}{\sinh^3\Big(\operatorname{argsh}\sqrt{\frac{3}{\lambda}}+\sqrt{2}t\Big)}\exp\Big(-\frac{3x}{\sinh^2\Big(\operatorname{argsh}\sqrt{\frac{3}{\lambda}}+\sqrt{2}t\Big)}\Big). \quad (3.32)$$

Let $\tau' = \inf\{t \geq 0 | Z_t^h = 0\}$ be the extinction time of Z^h. By the above calculation, we have

$$\mathbb{P}_x\big(\tau' \leq t\big) = \mathbb{P}\big(Z_t^h = 0\big) = \lim_{\lambda\to+\infty}\mathbb{E}\big[e^{-\lambda Z_t^h}\big] = \exp\Big(-\frac{3x}{\sinh^2(\sqrt{2}t)}\Big).$$

Hence, τ' has density

$$\Phi_t^h(x) = 6\sqrt{2}x\frac{\cosh(\sqrt{2}t)}{\sinh^3(\sqrt{2}t)}\exp\Big(-\frac{3x}{\sinh^2(\sqrt{2}t)}\Big) \quad (3.33)$$

with respect to the Lebesgue measure.

We now introduce the process Z^h conditioned on extinction at a fixed time ρ. We denote by $q_t(x, \mathrm{d}y)$ the transition kernels of Z^h. The process Z^h conditioned on extinction at time ρ is the time-inhomogeneous Markov process indexed by $[0, \rho]$ whose transition kernel between times s and t is

$$\pi_{s,t}(x, \mathrm{d}y) = \frac{\Phi_{\rho-t}^h(y)}{\Phi_{\rho-s}^h(x)}q_{t-s}(x, \mathrm{d}y)$$

for $x > 0$ and $0 \leq s < t < \rho$, and

$$\pi_{s,\rho}(x, \mathrm{d}y) = \delta_0(\mathrm{d}y).$$

As in [63], this interpretation can be justified by the fact that, for all $0 < s_1 < \cdots < s_k$, the conditional distribution of $(Z_{s_1}^h, \ldots, Z_{s_k}^h)$ on $\mathbb{P}_x\big(\,\cdot\,|\rho \leq \tau' \leq \rho+\varepsilon\big)$ converges to

$$\pi_{0,s_1}(x, \mathrm{d}y_1)\pi_{s_1,s_2}(y_1, \mathrm{d}y_2)\cdots\pi_{s_{k-1},s_k}(y_{k-1}, \mathrm{d}y_k)$$

as $\varepsilon \to 0^+$.

We also recall that the extinction time of Z has density $\Phi_t(x) = \frac{3x}{t^3}\exp\Big(-\frac{3x}{2t^2}\Big)$ (see Section 2.1 of [63]). By combining this with the density (3.30) for the nonconditioned processes, we get an absolute continuity relation for conditioned versions of Z and Z^h. The process Z^h started from x and conditioned to die at time ρ has density

$$e^x\frac{\Phi_\rho(x)}{\Phi_\rho^h(x)}\prod_{t_i\geq 0}(1+2|\Delta Z_{t_i}|)e^{-2|\Delta Z_{t_i}|} \quad (3.34)$$

with respect to Z started from x and conditioned to die at time ρ. To prove this properly, we just need to condition on $\rho \leq \tau' \leq \rho+\varepsilon$ and let ε go to 0.

We finally define X^h, which is a version of Z^h, starting from $+\infty$ at time $-\infty$, and conditioned to die exactly at time 0. We can construct a process $(X_t^h)_{t\leq 0}$ with càdlàg paths and no negative jumps such that:

(i) $X_t^h > 0$ for all $t < 0$ and $X_0^h = 0$ a.s.,

(ii) $X_t^h \longrightarrow +\infty$ almost surely as $t \to -\infty$,

(iii) for all x, if $T_x = \inf\{t \leq 0 | X_t \leq x\}$, the process $(X^h_{(T_x+t)\wedge 0})_{t\geq 0}$ has the same distribution as the process Z^h started from x.

Note that [63] describes a process X that is obtained from Z in the same way X^h is obtained from Z^h. Our Z corresponds to the X of [63], whereas our X corresponds to $\widetilde{X}$. As in [63], we can get an explicit construction of X^h by concatenating independent copies of Z^h started from $n+1$ and killed when hitting n for every $n \in \mathbb{N}$.

Proposition 4.4 of [63] states that for any $\rho, x > 0$, the process $(X_{t-\rho})_{0\leq t\leq \rho}$ conditioned on $X_{-\rho} = x$ has the same distribution as Z started from x and conditioned to die exactly at time ρ. By the same proof, this also holds for X^h and Z^h. Hence, $(X^h_{-t})_{0\leq t\leq \rho}$ conditioned on $X^h_{-\rho} = x$ has the density (3.34) with respect to $(X_{-t})_{0\leq t\leq \rho}$ conditioned on $X_{-\rho} = x$. Therefore, by the first point of Theorem 3.5, the process $(X^h_{-t})_{0\leq t\leq \rho}$ conditioned on $X^h_{-\rho} = x$ has the density (3.34) with respect to $(P_r)_{0\leq r\leq \rho}$ conditioned on $P_\rho = x$. But by (3.26), the process $(P^h_r)_{0\leq r\leq \rho}$ conditioned on $P^h_\rho = x$ has a density of the form

$$f(x,\rho)\prod_{t_i}(1+2|\Delta Z_{t_i}|)e^{-2|\Delta Z_{t_i}|}$$

with respect to $(P_r)_{0\leq r\leq \rho}$ conditioned on $P_\rho = x$. Since the density must have expectation one, we must have $f(x,\rho) = e^x \frac{\Phi_\rho(x)}{\Phi^h_\rho(x)}$.

Hence, in order to prove that P^h has the same distribution as $(X^h_{-t})_{t\geq 0}$, we only need to prove that these two processes have the same one-dimensional marginals. To this end, we will now compute the Laplace transform of the one-dimensional marginals of X^h.

Lemma 3.26. For all $t \geq 0$ and $\lambda \geq 0$, we have

$$\mathbb{E}\big[e^{-\lambda X^h_{-t}}\big] = \Big(1+\frac{\lambda}{3}\sinh^2(\sqrt{2}t)\Big)^{-1}\Big(1+\frac{\lambda}{3}\tanh^2(\sqrt{2}t)\Big)^{-1/2}.$$

Proof. Once again, the same computation for X is performed in [63]. By the exact same proof as in the beginning of the proof of Proposition 1.2 (ii) (section 4.1) in [63], we have

$$\mathbb{E}\big[e^{-\lambda X^h_{-t}}\big] = \lim_{x\to+\infty}\mathbb{E}_x\Big[\int_0^{+\infty} e^{-\lambda Z^h_s}\Phi^h_t(Z^h_s)\mathrm{d}s\Big] = \lim_{x\to+\infty}\int_0^{+\infty}\mathbb{E}_x\big[e^{-\lambda Z^h_s}\Phi^h_t(Z^h_s)\big]\mathrm{d}s.$$

By (3.33) and (3.32), one can compute $\mathbb{E}_x\big[e^{-\lambda Z^h_s}\Phi^h_t(Z^h_s)\big]$ exactly:

$$\begin{aligned}
\mathbb{E}_x\big[e^{-\lambda Z^h_s}\Phi^h_t(Z^h_s)\big] &= 6\sqrt{2}\frac{\cosh(\sqrt{2}t)}{\sinh^3(\sqrt{2}t)}\mathbb{E}_x\Big[Z^h_s\exp\Big(-\Big(\lambda+\frac{3}{\sinh^2(\sqrt{2}t)}\Big)Z^h_s\Big)\Big]\\
&= \frac{18\sqrt{6}x}{\Big(\lambda+\frac{3}{\sinh^2(\sqrt{2}t)}\Big)\sqrt{\lambda+3\coth^2(\sqrt{2}t)}}\frac{\cosh(\sqrt{2}t)}{\sinh^3(\sqrt{2}t)}\\
&\times \frac{\cosh\Big(\mathrm{argsh}\sqrt{\frac{3\sinh^2(\sqrt{2}t)}{\lambda\sinh^2(\sqrt{2}t)+3}}+\sqrt{2}s\Big)}{\sinh^3\Big(\mathrm{argsh}\sqrt{\frac{3\sinh^2(\sqrt{2}t)}{\lambda\sinh^2(\sqrt{2}t)+3}}+\sqrt{2}s\Big)}\\
&\times \exp\left(-\frac{3x}{\sinh^2\Big(\mathrm{argsh}\sqrt{\frac{3\sinh^2(\sqrt{2}t)}{\lambda\sinh^2(\sqrt{2}t)+3}}+\sqrt{2}s\Big)}\right).
\end{aligned}$$

We can now integrate over $s \geq 0$ to get

$$\begin{aligned}\int_0^{+\infty} \mathbb{E}_x\big[e^{-\lambda Z_s^h}\Phi_t^h(Z_s^h)\big]\mathrm{d}s &= \frac{3\sqrt{3}}{\Big(\lambda+\frac{3}{\sinh^2(\sqrt{2}t)}\Big)\sqrt{\lambda+3\coth^2(\sqrt{2}t)}}\frac{\cosh(\sqrt{2}t)}{\sinh^3(\sqrt{2}t)}\\ &\times \left(1-\exp\Big(-x\Big(\lambda+\frac{3}{\sinh^2(\sqrt{2}t)}\Big)\Big)\right).\end{aligned}$$

As $x \to +\infty$, the last factor goes to 1 and we get the claimed result. □

In remains to check that P_t^h has indeed the same Laplace transform. This is obtained by combining Lemma 3.24 and the Laplace transforms of the variables P_r and V_r that are computed in Proposition 1.2 and 1.4 of [63]. The computation is essentially the same as the proof of Lemma 3.17 and we omit it here. This completes the proof of Theorem 3.3. □

3.4.4 Asymptotics for the perimeter and volume processes

Proof of Corollary 3.1. Section 4.4 of [64] gives another construction of the process (P, V) via the Lamperti transform. We mimic this construction for (P^h, V^h). We consider the Lévy process S^h described by (3.28), started from $x > 0$. We write γ for the first time at which S^h hits 0. For every $t \geq 0$, we also set

$$\tau_t = \inf\left\{s \geq 0 \,\middle|\, \int_0^s \frac{\mathrm{d}u}{S_u^h} \geq t\right\}.$$

Finally, let $T = \int_0^\gamma \frac{\mathrm{d}u}{S_u^h}$. The process $\big(S_{\tau_t}^h\big)_{0\leq t\leq T}$ is the *Lamperti transform* of S^h, and has the same distribution as the branching process Z^h started from x. We now introduce the time-reversal $\mathcal{Q}$ of S^h (the notation $\mathcal{Q}$ is the mirror of S). It is the Lévy process with no positive jumps whose distribution is characterized by

$$\mathbb{E}\big[e^{u\mathcal{Q}_t}\big] = \exp\Big(t\sqrt{\frac{8}{3}}u\sqrt{u+3}\Big) \text{ for } u \geq 0.$$

Let also $\mathcal{Q}^+$ be the process $\mathcal{Q}$ conditionned to stay positive. We recall a classical time-reversal theorem (e.g. Theorem VII.18 of [34]). The process S^h started from $x > 0$ and killed when hitting 0 is the time-reversal of the process $\mathcal{Q}^+$ stopped when hitting x for the last time. By applying the Lamperti tranform, the time reversal of X^h, between its first passage at x and 0, is the Lamperti transform of $\mathcal{Q}^+$, stopped when hitting x for the last time. But by item 1) of Theorem 3.3, the time-reversal of X^h is P^h. Hence, the process P^h is the Lamperti transform of $\mathcal{Q}^+$. More precisely, for every $t \geq 0$, let

$$\eta_t = \inf\left\{s \geq 0 \,\middle|\, \int_0^s \frac{\mathrm{d}u}{\mathcal{Q}_u^+} \geq t\right\}.$$

Then P^h has the same distribution as $\big(\mathcal{Q}_{\eta_t}^+\big)_{t\geq 0}$. Moreover, let (s_i) be a measurable enumeration of the jumps of $\mathcal{Q}$. Conditionally on $\mathcal{Q}$, let $\big(\xi_i^h\big)$ be independent variables such that, for all i, the variable ξ_i^h has distribution $\nu_{|\Delta\mathcal{Q}_{s_i}|}$. We write

$$L_t = \sum_{s_i\leq t} \xi_i^h.$$

Let also L^+ be the process obtained by performing the exact same operation on $\mathcal{Q}^+$ instead of $\mathcal{Q}$. Since the construction of L^+ from $\mathcal{Q}^+$ is the same as the construction of V^h from P^h, we get

$$\big(P_r^h, V_r^h\big)_{r\geq 0} \overset{(d)}{=} \big(\mathcal{Q}_{\eta_t}^+, L_{\eta_t}^+\big)_{t\geq 0}.$$

It is now easy to obtain an asymptotic estimation of (P_r^h) through the study of $\mathcal{Q}$ and its conditioned version $\mathcal{Q}^+$. For all $t \geq 0$, we have

$$\begin{aligned}\mathbb{E}\big[\mathcal{Q}_t\big] &= \frac{d}{du}\Big|_{u=0} \mathbb{E}\big[e^{u\mathcal{Q}_t^+}\big] \\ &= \frac{d}{du}\Big|_{u=0} \exp\Big(t\sqrt{\frac{8}{3}}u\sqrt{u+3}\Big) \\ &= 2\sqrt{2}t.\end{aligned}$$

Moreover, it holds that $\mathbb{E}\big[e^{u\mathcal{Q}_t}\big] < +\infty$ for all $t \geq 0$ and $u \geq -3$ (this follows easily from the definition of S^h as S biased by an explicit martingale and the fact that $\mathcal{Q}$ is its time-reversal). Hence, by a classic moderate deviation argument, we have almost surely $\frac{\mathcal{Q}_n}{n} = 2\sqrt{2} + O(n^{-1/4})$ as $n \to +\infty$ for $n \in \mathbb{N}$. Moreover, let $c > 0$ be such that $\mathbb{P}\big(|\mathcal{Q}_t| \leq c\big) \geq \frac{1}{2}$ for all $0 \leq t \leq 1$. By the strong Markov property, we have $\mathbb{P}\big(|\mathcal{Q}_1| \geq x - c\big) \geq \frac{1}{2}\mathbb{P}\big(\sup_{t\in[0,1]} |\mathcal{Q}_t| \geq x\big)$ for all $x > 0$. Hence, the random variable $\sup_{t\in[0,1]} |\mathcal{Q}_t|$ has exponential tails on both sides, so for $n \in \mathbb{N}$ large enough and $0 \leq t \leq 1$, we have $|\mathcal{Q}_{n+t} - \mathcal{Q}_n| \leq n^{1/4}$. Hence, we have $\frac{\mathcal{Q}_t}{t} \underset{t\to+\infty}{\overset{a.s.}{=}} 2\sqrt{2} + O(t^{-1/4})$.

But the distribution of $(\mathcal{Q}_{t+1}^+ - \mathcal{Q}_1^+)_{t\geq 0}$ is just that of $\mathcal{Q}$, conditioned on an event of positive probability ($\mathcal{Q}$ drifts to $+\infty$ so it has a positive probability never to hit 0 after time 1). Hence, we deduce from above that

$$\frac{\mathcal{Q}_t^+}{t} \underset{t\to+\infty}{\overset{a.s.}{=}} 2\sqrt{2} + O(t^{-1/4}) \text{ , and so } \frac{1}{\mathcal{Q}_t^+} \underset{t\to+\infty}{\overset{a.s.}{=}} \frac{1}{2\sqrt{2}\,t} + O(t^{-5/4}).$$

The integral of the error term converges, so there is a random variable X_∞ such that $\int_0^s \frac{\mathrm{d}u}{\mathcal{Q}_u^+} = \frac{\ln s}{2\sqrt{2}} + X_\infty + o(1)$ a.s. It follows that $\frac{\ln \eta_t}{2\sqrt{2}} = t - X_\infty + o(1)$, so $\eta_t \sim e^{-2\sqrt{2}X_\infty} e^{2\sqrt{2}t}$ a.s. We finally get

$$\mathcal{Q}_{\eta_t}^+ \sim 2\sqrt{2}\eta_t \sim 2\sqrt{2}e^{-2\sqrt{2}X_\infty} e^{2\sqrt{2}t}$$

a.s. when $t \to +\infty$, so there is a random variable $0 < \mathcal{E} < +\infty$ such that $e^{-2\sqrt{2}r} P_r^h \xrightarrow[r\to+\infty]{a.s.} \mathcal{E}$.

To prove that $\frac{V_r^h}{P_r^h}$ converges a.s. to a deterministic constant, we first notice that L is a nondecreasing Lévy process. By construction, we have

$$\mathbb{E}\big[L_1\big] = \int_0^{+\infty} \mathbb{E}\big[\xi^h(x)\big]\mu^h(\mathrm{d}x),$$

where μ^h is the Lévy measure of S^h that is computed in (3.29), and $\xi^h(x)$ has distribution ν_x. By derivating (3.23) at $\beta = 0$, one can compute $\mathbb{E}\big[\xi^h(x)\big] = \frac{x^2}{2x+1}$. Since μ^h integrates x^2 near 0 and has exponential tail, the last display is finite. By the law of large numbers, we get

$$\frac{L_t}{t} \xrightarrow[t\to+\infty]{a.s.} c,$$

where $c = \mathbb{E}\big[L_1\big]$. By absolute continuity on $[1, +\infty[$, the same holds for L^+ instead of L.

In order to complete the proof of Corollary 3.1, we only need to find the distribution of $\mathcal{E}$ and the constant c. We will do this with Laplace transforms. We know that, for all $\lambda, \mu \geq 0$, we have

$$\mathbb{E}\big[e^{-\lambda\mathcal{E} - c\mu\mathcal{E}}\big] = \lim_{r\to+\infty} \mathbb{E}\Big[\exp\big(-\lambda e^{-2\sqrt{2}r} P_r^h - e^{-2\sqrt{2}r} V_r^h\big)\Big].$$

But we can compute these Laplace transforms thanks to Lemma 3.24 and Proposition 1.2 and 1.4 of [63]:

$$\begin{aligned}\mathbb{E}\big[e^{-\lambda P_r^h-\mu V_r^h}\big] &= \mathbb{E}\big[e^{-\lambda P_r}e^{-\mu V_r}e^{P_r}e^{-2V_r}\int_0^1 e^{-3x^2P_r}\mathrm{d}x\big]\\ &= \int_0^1 \mathbb{E}\Big[e^{(1-\lambda-3x^2)P_r}\mathbb{E}\big[e^{-(\mu+2)V_r}|P_r\big]\Big]\mathrm{d}x\\ &= \Big(1+\frac{2\lambda-2+\sqrt{2\mu+4}}{3\sqrt{2\mu+4}}\sinh^2\left((2\mu+4)^{1/4}r\right)\Big)^{-1}\\ &\quad\times\Big(1+\frac{2\lambda+4-2\sqrt{2\mu+4}}{3\sqrt{2\mu+4}}\tanh^2\left((2\mu+4)^{1/4}r\right)\Big)^{-1/2}.\end{aligned}$$

This gives the value of $\mathbb{E}\Big[\exp\Big(-e^{-2\sqrt{2}r}\lambda P_r^h - e^{-2\sqrt{2}r}\mu V_r^h\Big)\Big]$. When we let r go to $+\infty$, the second factor goes to 1. We also have

$$\frac{2\lambda e^{-2\sqrt{2}r}-2+\sqrt{2\mu e^{-2\sqrt{2}r}+4}}{3\sqrt{2\mu e^{-2\sqrt{2}r}+4}} \sim \frac{(2\lambda+\mu/2)e^{-2\sqrt{2}r}}{6}$$

and $\sinh^2\left((2\mu e^{-2\sqrt{2}r}+4)^{1/4}r\right) \sim \left(\frac{e^{\sqrt{2}r}}{2}\right)^2$, so the first factor goes to $\left(1+\frac{\lambda}{12}+\frac{\mu}{48}\right)^{-1}$. This is the Laplace transform of the couple $\left(X, \frac{1}{4}X\right)$, where X is an exponential variable of parameter 12. This ends the proof. □

3.A Appendix: Proof of Proposition 3.10

The goal of this appendix is to prove Proposition 3.10. We note that similar ideas to those below appear in Section 2 of [117], and more precisely in the proofs of Proposition 2.17 and 2.18. In particular, Lemma 3.30 is essentially proved in the proof of Proposition 2.18 there. If A is a subset of a metric space X and $\varepsilon > 0$, we will write A^ε for the union of all the open balls of radius ε centered at an element of A. Note that this is an open subset of X. We also recall that $B_r(X)$ and $\overline{B_r}(X)$ are respectively the closed ball and the hull centered at the root of X. In particular, they are both closed subsets of X. To prove Proposition 3.10, we will also need several times the following definition.

Definition 3.27. Let (X, d) be a metric space, $x, y \in X$ and $\varepsilon > 0$. An *ε-chain from x to y* is a finite sequence of points $(z_i)_{0\le i\le k}$ of X such that $z_0 = x$, $z_k = y$ and $d(z_i, z_{i+1}) \le \varepsilon$ for all $0 \le i \le k-1$.

Lemma 3.28. Under the assumptions of Proposition 3.10, the application $s \to \overline{B_s}(X)$ is continuous at r for the Hausdorff distance on the set of the compact subsets of $\overline{B_{r+1}}(X)$.

Proof. Let $\delta > 0$. It is enough to show that, for some $\varepsilon > 0$, we have:
a) $\overline{B_{r+\varepsilon}}(X) \subset \overline{B_r}(X)^\delta$,
b) $\overline{B_r}(X) \subset \overline{B_{r-\varepsilon}}(X)^\delta$.

We start with the first point. Let $A = \Big(\bigcap_{\varepsilon>0}\overline{B_{r+\varepsilon}}(X)\Big)\backslash\overline{B_r}(X)$. If $x \in A$, then there is a geodesic from x to a point $y \in X\backslash\overline{B_{r+1}}(X)$ that stays outside $B_r(X)$. However, γ has to intersect $B_{r+\varepsilon}(X)$ for all $\varepsilon > 0$. This is clearly impossible, so $A = \emptyset$. In particular, the decreasing intersection of the compact sets $\overline{B_{r+1/n}}(X)\backslash\overline{B_r}(X)^\delta$ is empty, so one of these compact sets is empty, which proves item a).

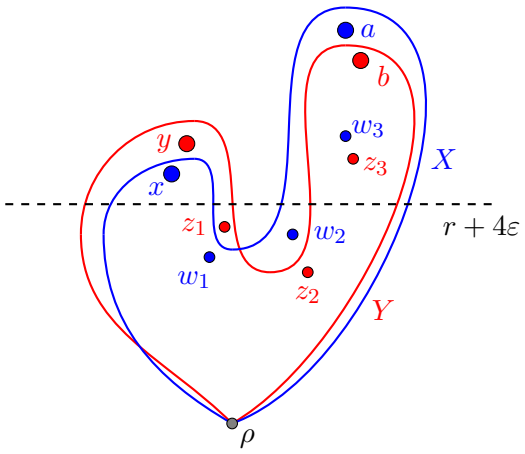

Figure 3.1 – Illustration of the proof of Lemma 3.30

For item b), let $A' = \left(\overline{B_r}(X)\backslash B_r(X)\right)\backslash\left(\bigcup_{\varepsilon>0}\overline{B_{r-\varepsilon}}(X)\right)$. By assumption (iii) we have $\mu(A') = \emptyset$ and A' is open, so by assumption (ii) we get $A' = \emptyset$. This implies $\overline{B_r}(X)\backslash B_r(X) \subset \bigcup_{n\geq 0}\overline{B_{r-1/n}}(X)^\delta$. Moreover, we have $B_r(X) \subset \overline{B_{r-1/n}}(X)^\delta$ for $\frac{1}{n} < \delta$. Hence, the increasing family of open sets $\overline{B_{r-1/n}}(X)^\delta$ covers the compact space $\overline{B_r}(X)$, so there is an $\varepsilon > 0$ such that $\overline{B_r}(X) \subset \overline{B_{r-\varepsilon}}(X)^\delta$. □

Remark 3.29. Note that the first point is true even without assumptions (ii) and (iii).

Lemma 3.30. Let $\varepsilon > 0$. Let (X, ρ, a) and (Y, ρ, b) be two bipointed connected, compact subsets of a locally compact metric space (Z, d). Assume that the Hausdorff distance between X and Y is less than ε and that $d(a, b) \leq \varepsilon$. Then, for all r such that $r + 4\varepsilon < d(\rho, b)$, we have

1) $\overline{B_r}(X)^\varepsilon \cap Y \subset \overline{B_{r+4\varepsilon}}(Y)$,
2) $\overline{B_r}(X) \subset \overline{B_{r+4\varepsilon}}(Y)^\varepsilon$.

Proof. We first notice that the connectedness of X implies that, for any two points x and x' in X, there is an ε-chain in X from x to x'. The same is true for Y.

Let y be a point in $\overline{B_r}(X)^\varepsilon \cap Y$. We want to show that $y \in \overline{B_{r+4\varepsilon}}(Y)$. We can assume $d(\rho, y) > r + 4\varepsilon$ (if it is not the case, then $y \in \overline{B_{r+4\varepsilon}}(Y)$ is obvious). Let $(z_i)_{0\leq i\leq k}$ be an ε-chain from y to b in Y. We know that $d(a, b) \leq \varepsilon$ and that there is a point $x \in X$ such that $d(x, y) \leq \varepsilon$. We write $w_0 = x$, $w_k = a$ and, for all $1 \leq i \leq k-1$, we take $w_i \in X$ such that $d(w_i, z_i) \leq \varepsilon$ (see Figure 3.1 for an illustration). For all i, we have $d(w_i, w_{i+1}) \leq d(w_i, z_i) + d(z_i, z_{i+1}) + d(z_{i+1}, w_{i+1}) \leq 3\varepsilon$, so (w_i) is a 3ε-chain from x to a in X. Since $x \in \overline{B_r}(X)$, there must be at least one i such that $d(\rho, w_i) \leq r + 3\varepsilon$, which implies $d(\rho, z_i) \leq r + 4\varepsilon$. Hence, every ε-chain from y to b has at least one point at distance from the root less than $r + 4\varepsilon$, so there is no ε-chain from y to b in $Y\backslash B_{r+4\varepsilon}(Y)$. This implies that y and b do not lie in the same connected component of $Y\backslash B_{r+4\varepsilon}(Y)$, so $y \in \overline{B_{r+4\varepsilon}}(Y)$, which proves the first point of the lemma.

The second point is easily obtained from the first one: if $x \in \overline{B_r}(X)$, then there is a $y \in Y$ such that $d(x, y) \leq \varepsilon$. By the first point, we have $y \in \overline{B_{r+4\varepsilon}}(Y)$, so $x \in \overline{B_{r+4\varepsilon}}(Y)^\varepsilon$. □

Proof of Proposition 3.10. First, we need to prove that the radii of the $\overline{B_r}(X_n)$ are bounded. Let $R = \max\{d(\rho, x) | x \in \overline{B_{r+4}}(X)\}$ be the radius of $\overline{B_{r+4}}(X)$. For n large enough, we have

$$d_{GH}(B_{R+3}(X_n), B_{R+3}(X)) < 1.$$

We write $X' = B_{R+3}(X)$ and $X'_n = B_{R+3}(X_n)$. The above inequality means that we can embed X' and X'_n isometrically in the same space (Z, d) in such a way that $X'_n \subset (X')^1$ and $\rho_n = \rho$.

Let $b \in X'_n \backslash B_{R+2}(X_n)$ and $a \in X'$ such that $d(a, b) \leq 1$. We have $d(a, \rho) \geq R + 1$, so $a \notin \overline{B_r}(X)$ and $\overline{B_r}(X) = \overline{B_r}(X')$ is the hull of radius r with center ρ with respect to a in X'. By item 2) of Lemma 3.30 for $\varepsilon = 1$, we have $\overline{B_r}(X'_n) \subset \overline{B_{r+4}}(X')^1$. In particular, the radius of $\overline{B_r}(X'_n)$ is less than $R + 1$. This means that only one connected component of $X' \backslash B_r(X'_n)$ contains points at distance greater than $R + 1$ from the root. In other words, the radius of $\overline{B_r}(X_n)$ is at most $R + 1$.

We now move on to the proof of our proposition. Let $\delta > 0$. By Lemma 3.28, there is $\varepsilon > 0$ such that $\overline{B_{r+4\varepsilon}}(X) \subset \overline{B_r}(X)^\delta$ and $\overline{B_r}(X) \subset \overline{B_{r-4\varepsilon}}(X)^\delta$. For n large enough, we have $d_{GHP}(X', X'_n) \leq \varepsilon$. This means that we can embed X' and X'_n in the same space Z in such a way that

a) $\rho = \rho_n$,
b) $X'_n \subset (X')^\varepsilon$,
c) $X' \subset (X'_n)^\varepsilon$,
d) for every $A \subset X'_n$ that is measurable we have $\mu_n(A) \leq \mu(A^\varepsilon) + \varepsilon$,
e) for every $B \subset X'$ that is measurable we have $\mu(B) \leq \mu_n(B^\varepsilon) + \varepsilon$.

This embedding provides a natural way to embed the measured metric spaces $\overline{B_r}(X_n)$ and $\overline{B_r}(X)$ in Z. We will deduce an upper bound for the GHP distance between these two hulls.

For all $y \in \overline{B_r}(X_n)$ there is an $x \in X'$ such that $d(x, y) \leq \varepsilon$. By Lemma 3.30 we have $x \in \overline{B_{r+4\varepsilon}}(X)$, so by Lemma 3.28 and our choice of ε, we have $x \in \overline{B_r}(X)^\delta$. This proves $\overline{B_r}(X_n) \subset \overline{B_r}(X)^{\delta+\varepsilon}$.

Similarly, let $x \in \overline{B_r}(X)$. We have $x \in \overline{B_{r-4\varepsilon}}(X)^\delta$ by Lemma 3.28 and our choice of ε. Let $z \in \overline{B_{r-4\varepsilon}}(X)$ be such that $d(x, z) \leq \delta$. There is a $y \in X'_n$ such that $d(y, z) \leq \varepsilon$. By Lemma 3.30 we have $y \in \overline{B_{r-4\varepsilon+4\varepsilon}}(X_n)$ and $d(x, y) \leq \delta + \varepsilon$, which proves $\overline{B_r}(X) \subset \overline{B_r}(X_n)^{\delta+\varepsilon}$. Hence, in our embedding, the Hausdorff distance between $\overline{B_r}(X)$ and $\overline{B_r}(X_n)$ is less than $\varepsilon + \delta$.

For all $A \subset \overline{B_r}(X_n)$ measurable, we have $\mu_n(A) \leq \mu(A^\varepsilon) + \varepsilon = \mu\big(A^\varepsilon \cap X'\big) + \varepsilon$. By Lemma 3.30 we have the inclusion $A^\varepsilon \cap X' \subset A^\varepsilon \cap \overline{B_{r+4\varepsilon}}(X)$, so we get

$$\begin{aligned} \mu_n(A) &\leq \varepsilon + \mu\big(A^\varepsilon \cap \overline{B_{r+4\varepsilon}}(X)\big) \\ &\leq \varepsilon + \mu\big(A^\varepsilon \cap \overline{B_r}(X)\big) + V(r + 4\varepsilon) - V(r), \end{aligned}$$

where we recall that $V(s) = \mu(\overline{B_s}(X))$ for all s.

Similarly, for all $B \subset \overline{B_r}(X)$ measurable, we have

$$\begin{aligned} \mu(B) &\leq \mu\big(B \cap \overline{B_{r-4\varepsilon}}(X)\big) + V(r) - V(r - 4\varepsilon) \\ &\leq \varepsilon + \mu_n\big((B \cap \overline{B_{r-4\varepsilon}}(X))^\varepsilon \cap X'_n\big) + V(r) - V(r - 4\varepsilon) \\ &\leq \varepsilon + \mu_n\big(B^\varepsilon \cap \overline{B_{r-4\varepsilon}}(X)^\varepsilon \cap X'_n\big) + V(r) - V(r - 4\varepsilon) \\ &\leq \varepsilon + \mu_n\big(B^\varepsilon \cap \overline{B_r}(X_n)\big) + V(r) - V(r - 4\varepsilon), \end{aligned}$$

where the last inequality uses Lemma 3.30.

Hence, our embedding of $\overline{B_r}(X)$ and $\overline{B_r}(X_n)$ gives the following bound for n large enough:

$$d_{GHP}\big(\overline{B_r}(X), \overline{B_r}(X_n)\big) \leq \max\Big(\varepsilon + \delta, \varepsilon + V(r + 4\varepsilon) - V(r), \varepsilon + V(r) - V(r - 4\varepsilon)\Big).$$

By assumption (iii) in the statement of the proposition, the right-hand side can be made arbitrarily small, which proves the first point of Proposition 3.10. The second point is an obvious consequence of the first one.

Finally, we note that in the case of compact, bipointed metric spaces $\big((X_n, d_n), x_n, y_n, \mu_n\big)$ and $\big((X, d), x, y, \mu\big)$ with $r < d(x, y)$, the above proof still works. The first part of the proof

(i.e. proving that the radii of the $\left(\overline{B_r}(X_n)\right)_{n\geq 0}$ are bounded) is not necessary anymore. We may apply the second part of the proof directly to X and X_n instead of the compact subsets X' and X'_n. Note that if $r > d(x, y)$, then $\overline{B_r}(X) = X$ and $\overline{B_r}(X_n) = X_n$ for n large enough, so the conclusion is immediate. □

Chapter 4

Infinite geodesics in random hyperbolic triangulations

ou *Où l'on se branche à la sortie d'une boule.*

This chapter is adapted from the preprint [46].

We study the structure of infinite geodesics in the Planar Stochastic Hyperbolic Triangulations $\mathbb{T}_\lambda$ introduced in [58]. We prove that these geodesics form a supercritical Galton–Watson tree with geometric offspring distribution. The tree of infinite geodesics in $\mathbb{T}_\lambda$ provides a new notion of boundary, which is a realization of the Poisson boundary. By scaling limits arguments, we also obtain a description of the tree of infinite geodesics in the hyperbolic Brownian plane. Finally, by combining our main theorem with results from Chapter 5, we obtain new hyperbolicity properties of $\mathbb{T}_\lambda$: they satisfy a weaker form of Gromov-hyperbolicity and admit bi-infinite geodesics.

Contents

4.1 Introduction

The construction and study of random infinite triangulations has been a very active field of reasearch in the last years. The first such triangulation that was built is the UIPT [20, 12]. A key feature in the study of this object is its spatial Markov property, which motivated the introduction of a one-parameter family $(\mathbb{T}_\lambda)_{0<\lambda\leq\lambda_c}$ of triangulations with $\lambda_c = \frac{1}{12\sqrt{3}}$, satisfying a similar property [58, 45] (see also [18] for similar constructions in the halfplanar case). The case $\lambda = \lambda_c$ corresponds to the UIPT, whereas for $\lambda < \lambda_c$ the triangulation $\mathbb{T}_\lambda$ has a hyperbolic behaviour. For example, it was proved that $\mathbb{T}_\lambda$ has a.s. exponential volume growth and that the simple random walk on it has positive speed [58]. The goal of this work is to establish hyperbolicity properties of these maps related to their geodesics.

Leftmost geodesic rays. Our first goal in this work is to describe precisely the structure of infinite geodesics in the triangulations $\mathbb{T}_\lambda$. More precisely, all the triangulations considered here are *rooted*, that is, equipped with a distinguished oriented edge called the *root edge*. The *root vertex*, that we write ρ, is the starting point of the root edge. For any vertices x and y in $\mathbb{T}_\lambda$, we call a geodesic γ from x to y *leftmost* if for any geodesic γ' from x to y, the union of γ and γ' cuts $\mathbb{T}_\lambda$ in two parts, and the part on the left of γ is infinite. A *leftmost geodesic ray* is a sequence of vertices $(\gamma(n))_{n\geq 0}$ such that $\gamma(0) = \rho$ and for any $n \geq 0$, the path $(\gamma(i))_{0\leq i\leq n}$ is a leftmost geodesic from ρ to $\gamma(n)$. We denote by $\mathbf{T}^g_\lambda$ the union of all the leftmost geodesic rays of $\mathbb{T}_\lambda$. We can see this set of vertices as a graph by relating two vertices if they are consecutive on a same geodesic ray. By uniqueness of the leftmost geodesic between two points, the graph $\mathbf{T}^g_\lambda$ is an infinite tree with no leaf. Moreover, the tree $\mathbf{T}^g_\lambda$ divides $\mathbb{T}_\lambda$ into infinite maps with geodesic boundaries, that we call *strips* (see the left part of Figure 4.1). Our first main result describes the distribution of $\mathbf{T}^g_\lambda$ and of these strips. Let $0 < \lambda \leq \lambda_c$. Let $h \in \left(0, \frac{1}{4}\right]$ be such that

$$\lambda = \frac{h}{(1+8h)^{3/2}}, \tag{4.1}$$

and let

$$m_\lambda = \frac{1-2h-\sqrt{1-4h}}{2h} \leq 1. \tag{4.2}$$

Theorem 4.1.
— The tree $\mathbf{T}^g_\lambda$ is a Galton–Watson tree with offspring distribution μ_λ, where $\mu_\lambda(0) = 0$ and $\mu_\lambda(k) = m_\lambda(1-m_\lambda)^{k-1}$ for $k \geq 1$.
— There are two random infinite strips S^0_λ and S^1_λ such that the following holds. Conditionally on $\mathbf{T}^g_\lambda$:
1. the strips delimited by $\mathbf{T}^g_\lambda$ are independent,
2. the strip containing the root edge has the same distribution as S^1_λ,
3. all the other strips have the same distribution as S^0_λ.

Note that for $\lambda = \lambda_c$, we have $h = \frac{1}{4}$ so $m_{\lambda_c} = 1$ and $\mu_{\lambda_c}(1) = 1$, so the tree $\mathbf{T}^g_{\lambda_c}$ consists of a single ray. This is reminiscent from geodesics confluence properties already observed in [68] for the UIPQ (the natural analog of $\mathbb{T}_{\lambda_c}$ for quadrangulations): there are infinitely many points that lie on every geodesic ray. See also [67] for similar results in the UIPT. However, our result for $\lambda = \lambda_c$ does not obviously imply those of [68] since we only deal with leftmost geodesic rays and not general geodesic rays. On the other hand, for $\lambda < \lambda_c$, we have $m_\lambda < 1$, so the offspring distribution μ_λ is supercritical and there are infinitely many leftmost geodesic rays. Moreover, the rate of exponential growth m_λ^{-1} of $\mathbf{T}^g_\lambda$ is the same as the rate of exponential volume growth of $\mathbb{T}_\lambda$.

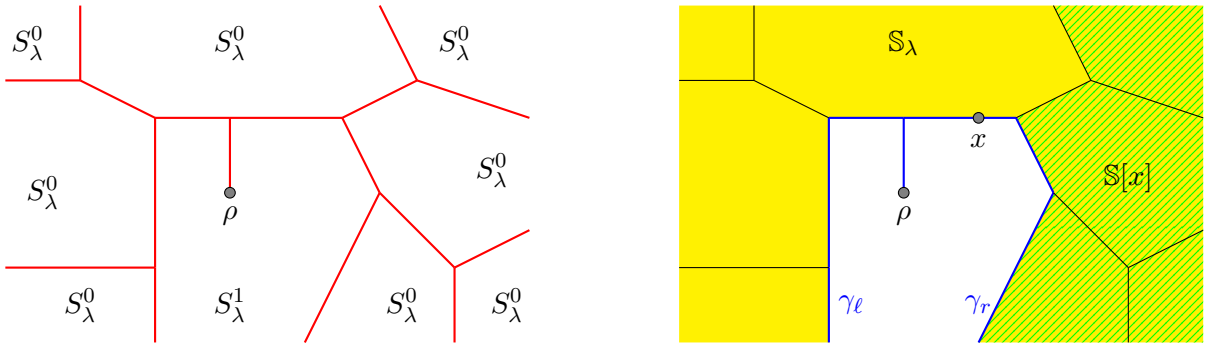

Figure 4.1 – The strip decomposition of $\mathbb{T}_\lambda$. On the left, the tree $\mathbf{T}^g_\lambda$ is in red. On the right $\mathbb{S}_\lambda = \mathbb{S}[\rho]$ is colored in yellow. The part $\mathbb{S}[x]$ of $\mathbb{S}_\lambda$ that is shaded in green has the same distribution as $\mathbb{S}_\lambda$.

We will also describe the distributions of S^0_λ and S^1_λ explicitely in terms of reverse Galton–Watson trees. For $\lambda < \lambda_c$, these strips should be thought of as "thin", in the sense that their width is of constant order as the distance from the root goes to $+\infty$.

We also state right now a consequence of Theorem 4.1 that will be useful later. Let γ_ℓ (resp. γ_r) be the leftmost (resp. rightmost) path in $\mathbf{T}^g_\lambda$. We write $\mathbb{S}_\lambda$ for the part of $\mathbb{T}_\lambda$ lying between γ_ℓ and γ_r, including the initial segment that γ_ℓ and γ_r have in common (cf. Figure 4.1). Then $\mathbb{S}_\lambda$ can be seen as a gluing of infinitely many independent copies of S^0_λ along $\mathbf{T}^g_\lambda$. This implies that $\mathbb{S}_\lambda$ has an interesting self-similarity property. Indeed, let $r > 0$. We condition on $B_r\left(\mathbf{T}^g_\lambda\right)$, the finite subtree of $\mathbf{T}^g_\lambda$ formed by those vertices lying at distance at most r from ρ. Let x be a vertex of $\mathbf{T}^g_\lambda$ such that $d(\rho, x) = r$. Let γ^x_ℓ (resp. γ^x_r) be the leftmost (resp. rightmost) infinite path in $\mathbf{T}^g_\lambda$ started from ρ and passing through x. Then the part of $\mathbb{T}_\lambda$ lying between γ^x_ℓ and γ^x_r above x has the same distribution as $\mathbb{S}_\lambda$ (see the right part of Figure 4.1). Indeed, it is also a gluing of independent copies of S^0_λ and the tree of descendants of x in $\mathbf{T}^g_\lambda$ has the same distribution as $\mathbf{T}^g_\lambda$. We will denote this part of $\mathbb{T}_\lambda$ by $\mathbb{S}[x]$. In all the rest of this work, "thin" maps such as S^0_λ and S^1_λ will be refered to as *strips*, and "thick" maps like $\mathbb{S}_\lambda$ as *slices*.

Hyperbolicity properties related to geodesics. By Theorem 4.1, for $\lambda < \lambda_c$, the map $\mathbb{T}_\lambda$ contains a supercritical Galton–Watson tree. By combining this with the results of Chapter 5, we obtain that $\mathbb{T}_\lambda$ satisfies two metric hyperbolicity properties: a weak form of Gromov-hyperbolicity, and the existence of bi-infinite geodesics.

More precisely, we recall that a graph G is hyperbolic in the sense of Gromov if there is a constant $k \geq 0$ such that all the triangles are k-thin in the following sense. Let x, y and z be vertices of G and γ_{xy}, γ_{yz}, γ_{zx} be geodesics from x to y, from y to z and from z to x. Then for any vertex v on γ_{xy}, the graph distance between v and $\gamma_{yz} \cup \gamma_{zx}$ is at most k. As usual, such a strong, uniform statement cannot hold for $\mathbb{T}_\lambda$ since any finite triangulation appears somewhere in $\mathbb{T}_\lambda$. Therefore, we need to study an "anchored" version.

Definition 4.1. Let M be a planar map. We say that M is *weakly anchored hyperbolic* if there is $k \geq 0$ such that the following holds. Let x, y and z be three vertices of M, and let γ_{xy} (resp. γ_{yz}, γ_{zx}) be a geodesic from x to y (resp. y to z, z to x). Assume the triangle formed by γ_{xy}, γ_{yz} and γ_{zx} surrounds the root vertex ρ. Then

$$d\left(\rho, \gamma_{xy} \cup \gamma_{yz} \cup \gamma_{zx}\right) \leq k.$$

Another property studied in Chapter 5 is the existence of bi-infinite geodesics, i.e. paths

$\big(\gamma(i)\big)_{i\in\mathbb{Z}}$ such that for any i and j, the graph distance between $\gamma(i)$ and $\gamma(j)$ is exactly $|i-j|$. This is not strictly speaking a hyperbolicity property, since such geodesics exist e.g. in $\mathbb{Z}^2$. However, they are expected to disappear after perturbations of the metric like first-passage percolation (see for example [87]). On the other hand, bi-infinite geodesics are much more stable in hyperbolic graphs [32]. In Chapter5, we prove that any random planar map containing a supercritical Galton–Watson tree with no leaf is a.s. weakly anchored hyperbolic, and contains bi-infinite geodesics. In particular, the following result follows from Theorem 4.1.

Corollary 4.2. Let $0 < \lambda < \lambda_c$. Almost surely, the map $\mathbb{T}_\lambda$ is weakly anchored hyperbolic and admits bi-infinite geodesics.

The existence of bi-infinite geodesics answers a question of Benjamini and Tessera [32]. Once again, there is a sharp contrast between the hyperbolic setting and "usual" random planar maps. For example, the results of [68] imply that such bi-infinite geodesics do not exist in the UIPQ.

Poisson boundary. Another goal of this work is to give a new description of the Poisson boundary of $\mathbb{T}_\lambda$ for $0 < \lambda < \lambda_c$ in terms of the tree $\mathbf{T}_\lambda^g$. Let G be an infinite, locally finite graph, and let $G \cup \partial G$ be a compactification of G, i.e. a compact metric space in which G is dense. Let also (X_n) be the simple random walk on G. We say that ∂G is a *realization of the Poisson boundary* of G if the following two properties hold:

— (X_n) converges a.s. to a point $X_\infty \in \partial G$,

— every bounded harmonic function h on G can be written in the form

$$h(x) = \mathbb{E}_x\left[g(X_\infty)\right],$$

where g is a bounded measurable function from ∂G to $\mathbb{R}$.

A first realization of the Poisson boundary of $\mathbb{T}_\lambda$ is given by a work of Angel, Hutchcroft, Nachmias and Ray [14]: let $\partial_{CP}\mathbb{T}_\lambda$ be the boundary of the circle packing of $\mathbb{T}_\lambda$ in the unit disk. We may equip $\mathbb{T}_\lambda \cup \partial_{CP}\mathbb{T}_\lambda$ with the topology induced by the usual topology on the unit disk. Then almost surely, $\partial_{CP}\mathbb{T}_\lambda$ is a realization of the Poisson boundary of $\mathbb{T}_\lambda$. Moreover, almost surely, the distribution of the limit point X_∞ has full support and no atoms in $\partial_{CP}\mathbb{T}_\lambda$.

We write $\partial\mathbf{T}_\lambda^g$ for the space of infinite rays of $\mathbf{T}_\lambda^g$. If $\gamma, \gamma' \in \partial\mathbf{T}_\lambda^g$, we write $\gamma \sim \gamma'$ if $\gamma = \gamma'$ or if γ and γ' are the left and right boundaries of the same strip. Then $\sim$ is a.s. an equivalence relation in which countably many equivalence classes have cardinal 2, and all the others have cardinal 1. We write $\widehat{\partial}\mathbf{T}_\lambda^g = \partial\mathbf{T}_\lambda^g/\sim$. There is a natural way to equip $\mathbb{T}_\lambda \cup \widehat{\partial}\mathbf{T}_\lambda^g$ with a topology that makes it a compact space, see Section 3.1 for more details. Hence, $\mathbb{T}_\lambda \cup \widehat{\partial}\mathbf{T}_\lambda^g$ can be seen as a compactification of the infinite graph $\mathbb{T}_\lambda$. We show that $\widehat{\partial}\mathbf{T}_\lambda^g$ is also a realization of the Poisson boundary.

Theorem 4.2. Let $0 < \lambda < \lambda_c$. Then almost surely:

1. the limit $\lim X_n = X_\infty$ exists, and its distribution has full support and no atoms in $\widehat{\partial}\mathbf{T}_\lambda^g$,
2. $\widehat{\partial}\mathbf{T}_\lambda^g$ is a realization of the Poisson boundary of $\mathbb{T}_\lambda$.

Note that, by a result of Hutchcroft and Peres [83], the second point will follow from the first one. Finally, we show in Chapter 5 that the existence of the limit X_∞ and the fact that it has full support are true in the more general setting of planar maps obtained by "filling the faces" of a supercritical Galton–Watson tree with i.i.d. strips. However, we did not manage to prove non-atomicity and to obtain a precise description of the Poisson boundary in this general setting. Our proof of non-atomicity here uses an argument specific to $\mathbb{T}_\lambda$, based on its peeling process.

Geodesic rays in the hyperbolic Brownian plane. Finally, the last goal of this work is to take the scaling limit of Theorem 4.1 to obtain results about geodesics in continuum objects. Indeed, another purpose of the theory of random planar maps is to build continuum random surfaces. The first such surface that was introduced is the Brownian map [99, 116], which is now known to be the scaling limit of a wide class of finite planar maps conditioned to be large [1, 3, 24, 36, 65, 112]. A noncompact version $\mathcal{P}$ called the Brownian plane was introduced in [62] and is the scaling limit of the UIPQ and also of the UIPT (see Chapter 3). Finally, it was shown in Chapter 3 that the hyperbolic random triangulations have a near-critical scaling limit called the hyperbolic Brownian plane and denoted $\mathcal{P}^h$. More precisely, let (λ_n) be a sequence of numbers in $(0, \lambda_c]$ satisfying

$$\lambda_n = \lambda_c \left(1 - \frac{2}{3n^4} + o\left(\frac{1}{n^4}\right)\right).$$

Then $\frac{1}{n}\mathbb{T}_{\lambda_n}$ converges for the local Gromov–Hausdorff distance to $\mathcal{P}^h$. By taking the scaling limit of Theorem 4.1 and checking that the tree of infinite leftmost geodesics behaves well in the scaling limit, we obtain a precise description of the geodesic rays in $\mathcal{P}^h$. Let $\mathbf{B}$ be the infinite tree in which every vertex has exactly two children, except the root which has only one.

Theorem 4.3. The infinite geodesic rays of $\mathcal{P}^h$ form a tree $\mathbf{T}^g(\mathcal{P}^h)$ that is distributed as a Yule tree with parameter $2\sqrt{2}$, i.e. the tree $\mathbf{B}$ in which the lengths of the edges are i.i.d. exponential variables with parameter $2\sqrt{2}$.

Once again, this behaviour is very different from the non-hyperbolic setting: in the Brownian plane, there is only one geodesic ray (this is Proposition 15 of [62], and an easy consequence of the local confluence of geodesics proved in [98] for the Brownian map). We also note that the rate $2\sqrt{2}$ of exponential growth of $\mathbf{T}^g(\mathcal{P}^h)$ is the same as the rate of exponential growth of the perimeters and volumes of the hulls of $\mathcal{P}^h$ [45, Corollary 1].

The skeleton decomposition. Our main tool for proving these results will be the skeleton decomposition of planar triangulations introduced by Krikun [94] for the type-II UIPT. See also [65] for the adaptation to the (slightly easier) type-I case. This decomposition encodes a triangulation by a reverse forest, where leftmost geodesics from the root pass between the trees and between their branches. The infinite forest describing the UIPT consists of a single tree, which can be seen as a reverse Galton–Watson tree with critical offspring distribution, started at time $-\infty$, and conditioned to have exactly 1 vertex at time 0 and to die at time 1. We obtain a similar description for the infinite forest encoding $\mathbb{T}_\lambda$ for $0 < \lambda < \lambda_c$, but here the offspring distribution is subcritical. A key feature is that the forest now contains infinitely many infinite trees. The parts of $\mathbb{T}_\lambda$ described by each of these trees are the strips delimited by the leftmost geodesic rays.

Structure of the paper. The structure of the paper is as follows. In Section 1, we recall some combinatorial results about planar triangulations, and we recall the definition and some basic properties of the maps $\mathbb{T}_\lambda$ and their halfplanar analogs. In Section 2, we prove Theorem 4.1 by computing the skeleton decomposition of $\mathbb{T}_\lambda$. Section 3 is devoted to the proof of Theorem 4.2, and Section 4 to the proof of Theorem 4.3. In Appendix A, we prove a technical result needed in Section 4, which shows that a wide class of events related to geodesics are closed for the Gromov–Hausdorff distance.

Acknowledgments: I thank Nicolas Curien for carefully reading several earlier versions of this manuscript, and Itai Benjamini for his question about bi-infinite geodesics. I acknowledge the support of ANR Liouville (ANR-15-CE40-0013) and ANR GRAAL (ANR-14-CE25-0014).

4.2 Combinatorics and preliminaries

4.2.1 Combinatorics

For $n \geq 0$ and $p \geq 1$, a *triangulation of the p-gon with n inner vertices* is a planar map with $n+p$ vertices in which all faces are triangles except one called the *outer face*, such that the boundary of the outer face is a simple cycle of length p. It is equipped with a root edge such that the outer face touches the root edge on its right. We consider type-I triangulations, which means we allow triangulations containing loops and multiple edges. We denote by $\mathscr{T}_{n,p}$ the set of triangulations of the p-gon with n inner vertices.

The number of triangulations with fixed volume and perimeter can be computed by a result of Krikun, as a special case of the main theorem of [95]:

$$\#\mathscr{T}_{n,p} = \frac{p(2p)!}{(p!)^2}\frac{4^{n-1}(2p+3n-5)!!}{n!(2p+n-1)!!} \underset{n\to+\infty}{\sim} c(p)\lambda_c^{-n}n^{-5/2}, \tag{4.3}$$

where $\lambda_c = \frac{1}{12\sqrt{3}}$ and

$$c(p) = \frac{3^{p-2}p(2p)!}{4\sqrt{2\pi}(p!)^2} \underset{p\to+\infty}{\sim} \frac{1}{36\pi\sqrt{2}}12^p\sqrt{p}. \tag{4.4}$$

For $p \geq 1$ and $\lambda \geq 0$, we write $w_\lambda(p) = \sum_{n\geq 0}\#\mathscr{T}_{n,p}\lambda^n$. Note that by the asymptotics (4.3), we have $w_\lambda(p) < +\infty$ if and only if $\lambda \leq \lambda_c$. We finally write $W_\lambda(x) = \sum_{p\geq 1} w_\lambda(p)x^p$. Formula (4) of [95] computes W_λ after a simple change of variables:

$$W_\lambda(x) = \frac{\lambda}{2}\left(\left(1-\frac{1+8h}{h}x\right)\sqrt{1-4(1+8h)x} - 1 + \frac{x}{\lambda}\right), \tag{4.5}$$

where $h \in \left(0, \frac{1}{4}\right]$ is given [1] by (4.1). From this formula, we easily get

$$w_\lambda(1) = \frac{1}{2} - \frac{1+2h}{2\sqrt{1+8h}} \tag{4.6}$$

and, for $p \geq 2$,

$$w_\lambda(p) = (2+16h)^p\frac{(2p-5)!!}{p!}\frac{(1-4h)p+6h}{4(1+8h)^{3/2}}. \tag{4.7}$$

We also define a *Boltzmann triangulation of the p-gon with parameter λ* as a random triangulation T such that $\mathbb{P}(T=t) = \frac{\lambda^n}{w_\lambda(p)}$ for every $n \geq 0$ and $t \in \mathscr{T}_{n,p}$.

4.2.2 Planar and halfplanar hyperbolic type-I triangulations

We recall from Chapter 3 the definition of the random triangulations $\mathbb{T}_\lambda$ for $0 < \lambda \leq \lambda_c$. A *finite triangulation with a hole of perimeter p* is a rooted map in which all the faces are triangles, except one whose boundary is a simple cycle of length p (the difference with a triangulation of the p-gon is that we do not require the root to lie on the boundary). Let t be a finite triangulation with a hole of perimeter p and let T be an infinite, one-ended triangulation of the plane. We write $t \subset T$ if T can be obtained by filling the hole of t with an infinite triangulation of the p-gon. For $0 < \lambda \leq \lambda_c$, the distribution of $\mathbb{T}_\lambda$ can be characterized as follows. For any finite triangulation t with a hole of perimeter p, we have

$$\mathbb{P}(t \subset \mathbb{T}_\lambda) = c_\lambda(p)\lambda^{|t|},$$

1. Note that our h corresponds to the h^3 of Krikun.

where $|t|$ is the total number of vertices of t and

$$c_\lambda(p) = \frac{1}{\lambda}\Big(8+\frac{1}{h}\Big)^{p-1}\sum_{q=0}^{p-1}\binom{2q}{q}h^q, \tag{4.8}$$

where h is as in (4.1). Equivalently, we can compute the generating function

$$C_\lambda(x) = \sum_{p\geq 1} c_\lambda(p)x^p = \frac{x}{\lambda\Big(1-\frac{1+8h}{h}x\Big)\sqrt{1-4(1+8h)x}}. \tag{4.9}$$

Note that the numbers $c_{\lambda_c}(p)$ are equal to the $c(p)$ defined by (4.4) and $\mathbb{T}_{\lambda_c}$ corresponds to the type-I UIPT [65, 129]. As in the type-II case [58], the triangulations $\mathbb{T}_\lambda$ exhibit a spatial Markov property similar to that of the UIPT: if t is a finite triangulation with a hole of perimeter p, conditionally on $t \subset \mathbb{T}_\lambda$, the distribution of the infinite triangulation that fills the hole of t only depends on p. We denote by $\mathbb{T}_\lambda^p$ a triangulation with this distribution. By a simple root transformation (more precisely duplicating the root edge, adding a loop inbetween and rooting the map at this new loop, see Figure 2 of [65]), triangulations of the plane are equivalent to infinite triangulations of the 1-gon. In particular, the image of $\mathbb{T}_\lambda$ under this root transformation is $\mathbb{T}_\lambda^1$, so studying one or the other are equivalent. In particular, the root transformation does not affect leftmost geodesics from the root, so all the results we will first obtain about geodesic rays in $\mathbb{T}_\lambda^1$ are immediate to transfer to $\mathbb{T}_\lambda$.

We also define the halfplanar analog of $\mathbb{T}_\lambda$, that we will denote by $\mathbb{H}_\lambda$. This will only be used in Section 3 to study the random walk on $\mathbb{T}_\lambda$, so all the rest of Section 1 can be skipped in first reading. A *triangulation of the halfplane* is an infinite planar map in which all the faces are triangles except one called the *outer face*, whose boundary is simple and infinite. Triangulations of the halfplane are rooted in such a way that the root edge touches the outer face on its right. We note that Angel and Ray build in [18] a family $(\mathbb{H}_\alpha^{\mathrm{II}})_{2/3<\alpha<1}$ of hyperbolic triangulations of the halfplane in the type-II setting and explain in Section 3.4 how to "add loops" to obtain type-I triangulations. The triangulations $\mathbb{H}_\lambda$ we define are a particular case of their construction. However, as in Chapter 3 and in order to limit the computations, we prefer to construct the maps $\mathbb{H}_\lambda$ directly instead of relying on the type-II case.

The law of $\mathbb{H}_\lambda$ is characterized by the following. Let t be a triangulation of the p-gon with a marked segment of edges $\partial_{\mathrm{out}}t$ on its boundary, such that $\partial_{\mathrm{out}}t$ contains the root edge. Such triangulations will be called *marked triangulations*. Let $\partial_{\mathrm{in}}t = \partial t \backslash \partial_{\mathrm{out}}t$. We write $|t|_{\mathrm{in}}$ for the number of vertices of t that do not lie on $\partial_{\mathrm{out}}t$. We also write $|\partial_{\mathrm{in}}t|$ (resp. $|\partial_{\mathrm{out}}t|$) for the number of edges on $\partial_{\mathrm{in}}t$ (resp. $\partial_{\mathrm{out}}t$). We write $t \subset \mathbb{H}_\lambda$ if $\mathbb{H}_\lambda$ can be obtained by gluing a triangulation of the halfplane H to t in such a way that $\partial t \cap \partial H = \partial_{\mathrm{in}}t$.

Proposition 4.3. For any $0 < \lambda \leq \lambda_c$, there is a random triangulation $\mathbb{H}_\lambda$ of the halfplane such that for every marked triangulation t, we have

$$\mathbb{P}\left(t \subset \mathbb{H}_\lambda\right) = \left(8+\frac{1}{h}\right)^{|\partial_{\mathrm{in}}t|-|\partial_{\mathrm{out}}t|}\lambda^{|t|_{\mathrm{in}}}.$$

Proof. We construct this triangulation by peeling along the same lines as in [18]. If $\mathbb{H}_\lambda$ exists, let f_0 be its triangular face that is adjacent to the root. Then f_0 has one of the three forms described by Figure 4.2. Moreover, we have $\mathbb{P}\left(\text{Case } I \text{ occurs}\right) = \lambda\left(8+\frac{1}{h}\right) = \frac{1}{\sqrt{1+8h}}$. By summing over all possible ways to fill the green zone, we also have

$$\mathbb{P}\left(\text{Case } II_i \text{ occurs}\right) = \mathbb{P}\left(\text{Case } III_i \text{ occurs}\right) = \left(8+\frac{1}{h}\right)^{-i} w_\lambda(i+1)$$

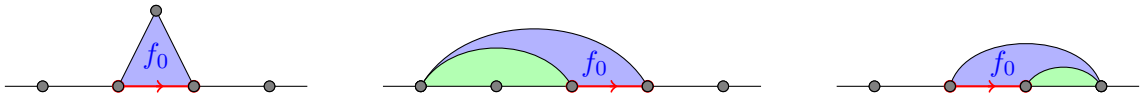

Case I Case II_i with $i = 2$ Case III_i with $i = 1$

Figure 4.2 – The three cases of peeling. In the two last cases, the parameter $i \geq 0$ corresponds to the number of edges of $\partial\mathbb{H}_\lambda$ that f_0 separates from infinity.

for all $i \geq 0$. If we sum up these probabilities, we obtain

$$\frac{1}{\sqrt{1+8h}} + 2\sum_{i\geq 0}\left(8+\frac{1}{h}\right)^{-i} w_\lambda(i+1) = \frac{1}{\sqrt{1+8h}} + 2\left(8+\frac{1}{h}\right) W_\lambda\left(\frac{h}{1+8h}\right) = 1 \qquad (4.10)$$

by (4.5). Since these probabilities sum up to 1, we can construct $\mathbb{H}_\lambda$ by peeling with the transitions described above. Everytime case II_i or III_i occurs, we fill the green bounded region with a Boltzmann triangulation of the $(i+1)$-gon with parameter λ. As in [18] (see also [58] in the planar case), we can check that we indeed obtain a triangulation of the halfplane, that its distribution does not depend on the choice of the peeling algorithm, and that the random triangulation we obtain has the right distribution. □

We now state a coupling result between $\mathbb{T}_\lambda$ and $\mathbb{H}_\lambda$ similar to the one stated in [17] and (implicitly) [58] in the type-II case. We recall that a *peeling algorithm* is a way to assign to every triangulation with a hole an edge on its boundary (see e.g. [58, Section 1.3]). To any peeling algorithm is naturally associated a filled-in exploration of $\mathbb{T}_\lambda$. By *filled-in*, we mean that every time the face just explored separates a finite region from infinity, the interior of the finite region is entirely discovered.

Lemma 4.4. (i) For $0 < \lambda \leq \lambda_c$, the triangulation $\mathbb{H}_\lambda$ is the local limit as $p \to +\infty$ of $\mathbb{T}_\lambda^p$.

(ii) For $0 < \lambda < \lambda_c$, consider a filled-in peeling algorithm $\mathscr{A}$ with infinitely many peeling steps, and let T be the part of $\mathbb{T}_\lambda$ that is discovered by $\mathscr{A}$. Then conditionally on T, the infinite connected components of $\mathbb{T}_\lambda \backslash T$ (rooted on their boundary according to a rooting convention depending only on T) are independent copies of $\mathbb{H}_\lambda$.

Proof. (i) If $p \geq 1$ and t is a marked triangulation with $|\partial_{\text{out}} t| \leq p$, let t_0 be a triangulation with a hole of perimeter p and let $t_0 + t$ be a triangulation obtained by gluing $\partial_{\text{out}} t$ to a segment of ∂t_0. Then by the definition of $\mathbb{T}_\lambda^p$, we have

$$\begin{aligned}
\mathbb{P}\left(t \subset \mathbb{T}_\lambda^p\right) = \frac{\mathbb{P}\left(t_0 + t \subset \mathbb{T}_\lambda\right)}{\mathbb{P}\left(t_0 \subset \mathbb{T}_\lambda\right)} &= \frac{c_\lambda(p + |\partial_{\text{in}} t| - |\partial_{\text{out}} t|)}{c_\lambda(p)} \lambda^{|t|_{\text{in}}} \\
&\xrightarrow[p\to+\infty]{} \left(8+\frac{1}{h}\right)^{|\partial_{\text{in}} t| - |\partial_{\text{out}} t|} \lambda^{|t|_{\text{in}}} \\
&= \mathbb{P}\left(t \subset \mathbb{H}_\lambda\right),
\end{aligned}$$

which is enough to conclude.

(ii) We first note that there are infinitely many peeling steps and all the finite holes are filled-in, so every connected component of $\mathbb{T}_\lambda \backslash T$ is halfplanar. The second point then

follows from the first one since the perimeter of the region discovered after n peeling steps a.s. goes to $+\infty$ as $n \to +\infty$. See the proof of Lemma 2.16 in [17] for the same result in the type-II case. Note that for $\lambda = \lambda_c$, we have $T = \mathbb{T}_\lambda$ a.s. by Corollary 7 of [64], so the statement (ii) is irrelevant.

□

4.3 The skeleton decomposition of hyperbolic triangulations

The aim of this section is to prove Theorem 4.1. It is organized as follows. In Section 2.1, we describe the finite skeleton decomposition, which associates to every finite triangulation a finite forest, and its infinite counterpart. In Section 2.2, we compute the distribution of the skeletons of the hulls of $\mathbb{T}_\lambda$. This characterizes the skeleton of $\mathbb{T}_\lambda$ entirely, but in a form that is not convenient for the proof of Theorem 4.1. In Section 2.3, we explain why the infinite skeleton of $\mathbb{T}_\lambda$ is related to infinite leftmost geodesics in the triangulation. Section 2.4 contains a description of the strips S^0_λ and S^1_λ by the distribution of their hulls, without a proof of their existence. In Section 2.5, we use all that precedes to prove Theorem 4.1. Finally, Section 2.6 is devoted to the construction of S^0_λ and S^1_λ.

4.3.1 The skeleton decomposition of finite and infinite triangulations

The finite setting: skeleton decomposition of triangulations of the cylinder. We first recall the skeleton decomposition of triangulations introduced by Krikun [94, 92] for type-II triangulations and quadrangulations, and described in [65] for type-I triangulations (see also [19, 122]). This decomposition applies to so-called triangulations of the cylinder. Most of the presentation here is adapted from [65].

Definition 4.5. Let $r \geq 1$. A *triangulation of the cylinder of height* r is a rooted planar map in which all faces are triangles except two distinguished faces called the top and the bottom faces, such that the following properties hold. The boundaries of the top and bottom faces are simple cycles. The bottom face lies on the right of the root edge. Finally, every vertex incident to the top face is at distance r from the boundary of the bottom face, and every edge adjacent to the top face is also adjacent to a face whose third vertex is at distance $r-1$ from the boundary of the bottom face.

If Δ is a triangulation of the cylinder of height r, we write $\partial\Delta$ and $\partial_*\Delta$ for the boundaries of the bottom and top faces. Let $p = |\partial\Delta|$ and $q = |\partial_*\Delta|$. The skeleton decomposition encodes Δ by a forest of q plane trees and a family of triangulations of polygons indexed by the vertices of this forest. For $1 \leq j \leq r-1$, we define the ball $B_j(\Delta)$ as the map formed by all the faces of Δ having at least one vertex at distance at most $j-1$ from $\partial\Delta$, along with their vertices and edges. We also define the hull $B^\bullet_j(\Delta)$ as the union of $B_j(\Delta)$ and all the connected components of its complementary, except the one that contains $\partial_*\Delta$. It is easy to see that $B^\bullet_j(\Delta)$ is a triangulation of the cylinder of height j. We denote by $\partial_j\Delta$ the top boundary of $B^\bullet_j(\Delta)$, with the conventions $\partial_0\Delta = \partial\Delta$ and $\partial_r\Delta = \partial_*\Delta$.

If $1 \leq j \leq r$, every edge of $\partial_j\Delta$ is incident to exactly one triangle whose third vertex belongs to $\partial_{j-1}\Delta$. Such triangles are called *downward triangles at height* j. We can define a genealogy on $\bigcup_{j=0}^r \partial_j\Delta$ by saying that $e \in \partial_j\Delta$ for $j \geq 1$ is the parent of $e' \in \partial_{j-1}\Delta$ if the downward triangle adjacent to e is the first one that one encounters when moving along $\partial_{j-1}\Delta$ in clockwise order starting from the middle of the edge e' (see Figure 4.3). We obtain q trees rooted on $\partial_*\Delta$. Let $(t_1, \ldots, t_q)$ be the forest obtained by listing these trees in clockwise order in such a way that the root edge of Δ lies in t_1. This forest is called the *skeleton of* Δ and we denote it by $\mathrm{Skel}(\Delta)$.

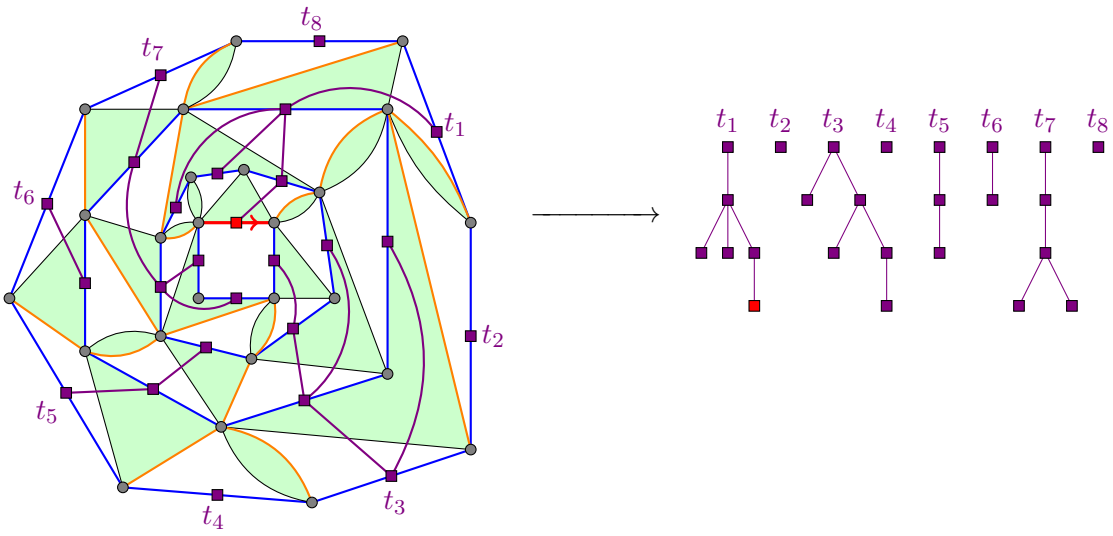

Figure 4.3 – A triangulation of the cylinder Δ and its skeleton $\mathrm{Skel}(\Delta) \in \mathscr{F}'_{4,8,3}$ in purple. The cycles $\partial_j\Delta$ are in blue, and the root edge is in red. The leftmost geodesics from the vertices of $\partial_*\Delta$ to $\partial\Delta$ are in orange. The green holes must be filled by triangulations of polygons.

Note that t_1 has a distinguished vertex at height r. The set of possible values of $\mathrm{Skel}(\Delta)$ is called the set of (p,q,r)-*admissible forests* and is described by the next definition.

Definition 4.6. Let $p,q,r \geq 1$: a (p,q,r)-pre-admissible plane forest is a sequence of plane trees $f = (t_1,\dots,t_q)$, equipped with a distinguished vertex ρ, such that:
- the maximal height of the trees t_i is r,
- the total number of vertices at height r in the trees t_i is p,
- ρ lies at height r.

We write $\mathscr{F}_{p,q,r}$ for the set of (p,q,r)-pre-admissible forests. If furthermore $\rho \in t_1$, we say that f is (p,q,r)-admissible, and we write $\mathscr{F}'_{p,q,r}$ for the set of (p,q,r)-admissible forests.

Let $f = (t_1,\dots,t_p) \in \mathscr{F}_{p,q,r}$. Most of the time, we will represent f with the roots of the trees t_i on the top. Hence, if $x \in t_i$ is a vertex of f, we define the *reverse height of x in f* as r minus the height of x in t_i, and we write it $h_f^{\mathrm{rev}}(v)$. In particular, the roots of the trees t_i have reverse height r and ρ has reverse height 0. Although quite unusual, this convention is natural because the reverse heights in $\mathrm{Skel}(\Delta)$ match the distances to the root in the triangulation Δ. It will also be more convenient when we will deal with reverse forests with infinite height.

The forest $\mathrm{Skel}(\Delta)$ is not enough to completely describe Δ: if we consider all the downward triangles of Δ, there is a family of holes, each of which is naturally associated to an edge of $\bigcup_{j=0}^{r} \partial_j\Delta$. If $e \in \partial_j\Delta$ with $1 \leq j \leq r$, the associated hole is bounded by the edges of $\partial_{j-1}\Delta$ that are children of e and by two vertical edges connecting the initial vertex of e to two vertices of $\partial_{j-1}\Delta$. This hole has perimeter $c_e + 2$, where c_e is the number of children of e, so it must be filled by a triangulation of a (c_e+2)-gon. If $c_e = 0$, it is possible that the hole of perimeter 2 is filled by the triangulation of the 2-gon consisting of a single edge, which means that the two vertical edges are simply glued together.

If f is a (p,q,r)-admissible forest, let f^* be the set of those vertices v of f such that $h_f^{\mathrm{rev}}(v) > 0$. The decomposition we just described is a bijection between triangulations of the cylinder Δ with height r such that $\partial\Delta = p$ and $\partial_*\Delta = q$, and pairs consisting of a (p,q,r)-admissible forest f and a family $(M_v)_{v\in f^*}$ of maps such that M_v is a finite triangulation of a (c_v+2)-gon for every v.

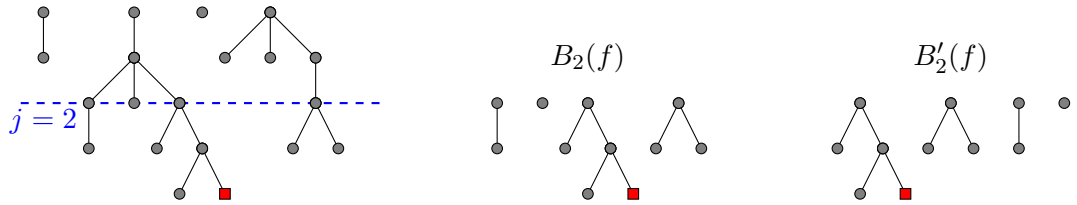

Figure 4.4 – A forest f of height 4 and the forests $B_2(f)$ and $B'_2(f)$. The ancestors lie on the top and the distinguished vertex is in red.

Moreover, this decomposition encodes informations about leftmost geodesics from the vertices of $\partial_*\Delta$ to $\partial\Delta$: these geodesics (in orange on Figure 4.3) are the paths going between the trees of Skel(Δ). These geodesics cut Δ into q slices, each of which contains an edge of $\partial_*\Delta$. Moreover, the slice containing the i-th edge of $\partial_*\Delta$ can be completely described by the i-th tree t_i of Skel(Δ) and the maps M_v for $v \in t_i$.

The infinite setting: infinite reverse forests. If $f \in \mathscr{F}_{p,q,r}$ and $0 \leq j \leq r$, let $x_1^j, \dots, x_s^j$ be the vertices of f lying at reverse height j in f, from left to right. For every $1 \leq i \leq s$, let t_i^j be the tree of descendants of x_i^j. Let also i_0 be the index such that the distinguished vertex belongs to $t_{i_0}^j$. We define $B_j(f)$ as the forest $(t_1^j, \dots, t_s^j)$, with the same distinguished vertex as f (see Figure 4.4), and $B'_j(f)$ as the forest $(t_{i_0}^j, t_{i_0+1}^j, \dots, t_s^j, t_1^j, \dots, t_{i_0-1}^j)$. Note that $B_j(f) \in \mathscr{F}_{p,s,j}$ and $B'_j(f) \in \mathscr{F}'_{p,s,j}$, but it is not always the case that $B_j(f) \in \mathscr{F}'_{p,s,j}$. Note also that if v is a vertex of $B_j(f)$, then $h^{\mathrm{rev}}_{B_j(f)}(x) = h^{\mathrm{rev}}_{B'_j(f)}(x) = h^{\mathrm{rev}}_f(x)$.

Definition 4.7. Let $p \geq 1$. A *p-pre-admissible infinite forest* is a sequence $\mathbf{f} = (f_r)_{r\geq 1}$ of plane forests such that

(i) for every $r \geq 1$, there is $q \geq 1$ such that the forest f_r is (p, q, r)-pre-admissible,

(ii) for every $s \geq r \geq 1$, we have $B_r(f_s) = f_r$.

We write $\mathscr{F}_{p,\infty,\infty}$ for the set of p-pre-admissible infinite forests. We will also write $B_r(\mathbf{f}) = f_r$ for $r \geq 1$.

Note that $\mathbf{f} \in \mathscr{F}_{p,\infty,\infty}$ can also be seen as an increasing sequence of finite graphs, and therefore as an infinite graph. We call *infinite reverse trees* the connected components of $\mathbf{f}$. If v is a vertex of $\mathbf{f}$, we have $v \in f_r$ for r large enough. Note that $h^{\mathrm{rev}}_{f_r}(v)$ does not depend on r as long as $v \in f_r$. We call it the *reverse height of v in* $\mathbf{f}$ and denote it by $h^{\mathrm{rev}}_{\mathbf{f}}(v)$. We also write $\mathbf{f}^*$ for the set of vertices of $\mathbf{f}$ that do not lie at reverse height 0.

Definition 4.8. A p-pre-admissible forest $\mathbf{f}$ is called *p-admissible* if the distinguished vertex lies in the leftmost infinite tree of $\mathbf{f}$, i.e. if for every $r \geq 0$, the leftmost vertex of $\mathbf{f}$ at reverse height r is in the same infinite tree as the distinguished vertex. We write $\mathscr{F}'_{p,\infty,\infty}$ for the set of p-admissible infinite forests.

Skeleton decomposition of infinite triangulations of the p-gon. We now introduce the skeleton decomposition of infinite triangulations of the p-gon. Let T be an infinite, one-ended triangulation of the p-gon. For every $r \geq 1$, the hull $B_r^\bullet(T)$ is a triangulation of the cylinder of height r, so we can define its skeleton $f'_r = \mathrm{Skel}\,(B_r^\bullet(T)) \in \mathscr{F}'_{p,q,r}$ for some q. It is also easy to see that the forests f'_r are consistent in the sense that $B'_r(f'_s) = f'_r$ for every $s \geq r \geq 1$. We claim that such a family (f'_r) always defines an infinite p-admissible forest. More precisely, if $\mathbf{f} \in \mathscr{F}'_{p,\infty,\infty}$ and $r \geq 1$, let $B'_r(\mathbf{f})$ be the reordered ball of radius r in $\mathbf{f}$ (that is, the ball $B_r(\mathbf{f})$ in which the

trees have been cyclically permutated so that the distinguished vertex lies in the first tree). Then there is a unique $\mathbf{f} \in \mathscr{F}'_{p,\infty,\infty}$ such that $B'_r(\mathbf{f}) = f'_r$ for every $r \geq 1$. We do not prove this formally, but explain how to build $\mathbf{f}$ from (f'_r): for any $s \geq r$, the forest $B_r(f'_s)$ is a cyclic permutation of f'_r, and this cyclic permutation does not depend on s for s large enough. Hence, we can set $f_r = B_r(f'_s)$ for s large enough and $\mathbf{f} = (f_r)_{r\geq 1}$. Therefore, there is a unique p-admissible forest, that we denote by $\mathrm{Skel}(T)$ and call the *skeleton* of T, such that $B'_r(\mathrm{Skel}(T)) = \mathrm{Skel}(B^\bullet_r(T))$ for every $r \geq 1$. As in the finite case, the skeleton decomposition establishes a bijection between one-ended infinite triangulations of the p-gon and pairs consisting of a p-admissible infinite forest $\mathbf{f}$ and a family $(M_v)_{v\in\mathbf{f}^*}$ of maps such that M_v is a finite triangulation of a (c_v+2)-gon for every v.

Remark 4.9. In order to define the skeleton decomposition, it would have been more convenient to define an infinite reverse forest $\mathbf{f}$ as the sequence $(B'_r(\mathbf{f}))$ instead of $(B_r(\mathbf{f}))$. The reason why we chose this definition is that it will later make the decomposition of an infinite forest in infinite reverse trees much more convenient to define.

4.3.2 Computation of the skeleton decomposition of the hulls of $\mathbb{T}_\lambda$

We now compute the law of the skeletons of the hulls of $\mathbb{T}^p_\lambda$. The map $\mathbb{T}^p_\lambda$ can be seen as a triangulation of the cylinder with infinite height. For every $r \geq 1$, the map $B^\bullet_r(\mathbb{T}^p_\lambda)$ is a triangulation of the cylinder of height r with bottom boundary length equal to p.

Lemma 4.10. Let $0 < \lambda \leq \lambda_c$. Let Δ be a triangulation of the cylinder of height r. We write p (resp. q) for the length of its bottom (resp. top) boundary. Let $f = \mathrm{Skel}(\Delta) \in \mathscr{F}'_{p,q,r}$. Then

$$\mathbb{P}\left(B^\bullet_r(\mathbb{T}^p_\lambda) = \Delta\right) = \frac{q\,h_\lambda(q)}{p\,h_\lambda(p)} \prod_{v\in f^*} \theta_\lambda(c_v) \prod_{v\in f^*} \frac{\lambda^{|M_v|}}{w_\lambda(c_v+2)},$$

where:
- c_v is the number of children of v in f,
- $h \in \left(0, \frac{1}{4}\right]$ is given by (4.1),
- $h_\lambda(p) = \frac{1}{p}\left(8+\frac{1}{h}\right)^{-p} c_\lambda(p)$,
- θ_λ is the offspring distribution whose generating function g_λ is given by

$$g_\lambda(x) = \sum_{i\geq 0} \theta_\lambda(i) x^i = \frac{1}{x} - \frac{(1-x)(1-\sqrt{1-4hx})}{2hx^2}. \tag{4.11}$$

Proof. The proof is exactly the same as in the case $\lambda = \lambda_c$ (Lemma 2 in [65]), up to changes of notation (the ρ of [65] corresponds to our λ and the α corresponds to $8+\frac{1}{h}$). The same computations yield

$$\theta_\lambda(i) = \frac{1}{\sqrt{1+8h}} \left(\frac{h}{1+8h}\right)^i w_\lambda(i+2), \tag{4.12}$$

and the computation of g_λ follows from (4.5). □

Let $p, q, r \geq 1$ and $f \in \mathscr{F}'_{p,q,r}$. We sum the formula of Lemma 4.10 over all families $(M_v)_{v\in f^*}$ such that M_v is a triangulation of a (c_v+2)-gon for every v. By the definition of $w_\lambda(i+2)$, we have $\sum_{n\geq 0} \#\mathscr{T}_{i+2,n} \frac{\lambda^n}{w_\lambda(i+2)} = 1$ for every $i \geq 0$, so we get

$$\mathbb{P}\left(\mathrm{Skel}\left(B^\bullet_r(\mathbb{T}^p_\lambda)\right) = f\right) = \frac{q h_\lambda(q)}{p h_\lambda(p)} \prod_{v\in f^*} \theta_\lambda(c_v). \tag{4.13}$$

Note that (4.13) describes explicitly the distribution of $B'_r\left(\mathrm{Skel}(\mathbb{T}_\lambda^p)\right)$, so we completely know the law of $\mathrm{Skel}(\mathbb{T}_\lambda^p)$. As we will see in Section 2.3, this is in theory enough to prove Theorem 4.1. However, infinite leftmost geodesics are not very tractable in this characterization. Hence, we will need to find another construction of the $\mathrm{Skel}\left(\mathbb{T}_\lambda^p\right)$ and prove it is equivalent to (4.13). This will be the main goal of the rest of Section 2. Before moving on to the proof of Theorem 4.1, we end this subsection with a few remarks about the perimeter process of $\mathbb{T}_\lambda$.

We notice that (4.13) can be used to study the perimeter process of $\mathbb{T}_\lambda$ in the same way as the perimeter process of $\mathbb{T}_{\lambda_c}$ is studied in [94] and [122]. More precisely, by the same computation as in the proof of Lemma 3 in [65], by summing (4.13) over all (p,q,r)-admissible forests, we obtain

$$\mathbb{P}\left(|\partial B_r^\bullet\left(\mathbb{T}_\lambda^p\right)| = q\right) = \frac{h_\lambda(q)}{h_\lambda(p)}\mathbb{P}_q\left(X_\lambda(r) = p\right),$$

where X_λ is the Galton–Watson process with offspring distribution θ_λ. Since we know that $(|\partial B_r^\bullet(\mathbb{T}_\lambda)|)_{r\geq 0}$ is a Markov chain (by the spatial Markov property) and has the same distribution as the perimeter process of $\mathbb{T}_\lambda^1$, we even get, for every $p,q\geq 1$ and $s\geq r\geq 0$,

$$\mathbb{P}\left(|\partial B_s^\bullet\left(\mathbb{T}_\lambda\right)| = q \middle| |\partial B_r^\bullet\left(\mathbb{T}_\lambda\right)| = p\right) = \frac{h_\lambda(q)}{h_\lambda(p)}\mathbb{P}_q\left(X_\lambda(s-r) = p\right). \tag{4.14}$$

Let $m_\lambda = \sum_{i\geq 0} i\theta_\lambda(i)$ be the mean number of children. By (4.11), we can compute $m_\lambda = g'_\lambda(1)$ and obtain (4.2). In particular, we have $m_\lambda \leq 1$, with equality if and only if $\lambda = \lambda_c$. Hence, the Galton–Watson process X_λ is subcritical for $\lambda < \lambda_c$. We can therefore see the perimeter process of $\mathbb{T}_\lambda$ for $\lambda < \lambda_c$ as a time-reversed subcritical branching process. Note that the perimeters and volumes of the hulls in $\mathbb{T}_\lambda$ are already quite well-known. Sharp exponential growth for a fixed $\lambda < \lambda_c$ is proved in Section 2 of [58], whereas the near-critical scaling limit as $\lambda \to \lambda_c$ is studied in Section 3 of Chapter 3. Equation (4.14) together with Lemma 4.12 can give explicit formulas for the generating function of the perimeters of the hull, so it should be possible to recover these results by using the same techniques as in [122], but we do not do this in this work.

4.3.3 Slicing the skeleton

The goal of this subsection is twofold. First, we explain the link between the skeleton decomposition and the infinite leftmost geodesics of an infinite triangulation. Second, we introduce some formalism, that will later allow us to obtain a construction of $\mathrm{Skel}(\mathbb{T}_\lambda)$ that is more suitable for our purpose than (4.13).

Decomposition of a 1-admissible forest in reverse trees. An infinite 1-admissible forest $\mathbf{f}$ may contain several infinite reversed trees, and we will need to study the way these trees are placed with respect to each other. This can be encoded by a genealogical structure (see Figure 4.5). If $\mathbf{t}$ is one of the infinite trees of $\mathbf{f}$, let $h_{\min}^{\mathrm{rev}}(\mathbf{t})$ be the reverse height of the lowest vertex of $\mathbf{t}$. We consider the set of pairs $(\mathbf{t},i)$ where $\mathbf{t}$ is an infinite tree of $\mathbf{f}$ and $i \geq h_{\min}^{\mathrm{rev}}(\mathbf{t})$. If $i > h_{\min}^{\mathrm{rev}}(\mathbf{t})$, the parent of $(\mathbf{t},i)$ is $(\mathbf{t},i-1)$. If $i = h_{\min}^{\mathrm{rev}}(\mathbf{t}) > 0$, let $\mathbf{t}'$ be the first infinite tree on the left of $\mathbf{t}$ such that $h_{\min}^{\mathrm{rev}}(\mathbf{t}') \leq i-1$ (note that $\mathbf{t}'$ always exists because $\mathbf{f}$ is admissible). Then the parent of $(\mathbf{t},i)$ is $(\mathbf{t}',i-1)$. Finally, if $i=0$, then $(\mathbf{t},i)$ has no parent. This genealogy is encoded in an infinite plane tree with no leaf that we denote by $\mathbf{U}(\mathbf{f})$ (see Figure 4.5). Note that the genealogy in $\mathbf{U}(\mathbf{f})$ is "reversed" compared to the genealogy in $\mathbf{f}$: the parent of a vertex x of $\mathbf{U}(\mathbf{f})$ lies below x, whereas in the forest $\mathbf{f}$, the parent of a vertex lies above it. Therefore, the heights in $\mathbf{U}(\mathbf{f})$ match the reverse heights in $\mathbf{f}$. We also write $B_r(\mathbf{U}(\mathbf{f}))$ for the subtree of $\mathbf{U}(\mathbf{f})$ whose vertices are the $(\mathbf{t},i)$ with $i\leq r$. Note that the tree $B_r(\mathbf{U}(\mathbf{f}))$ is not a measurable function of $B_r(\mathbf{f})$

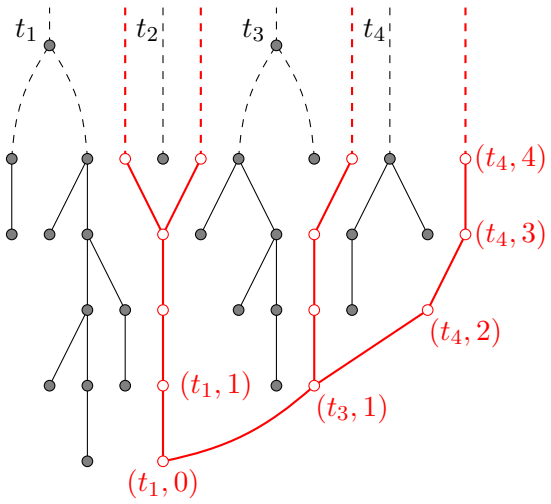

Figure 4.5 – An infinite 1-admissible forest $\mathbf{f}$ in which 4 trees t_1, t_2, t_3 and t_4 reach reverse height 4. In red, the tree $\mathbf{U}(\mathbf{f})$ and the names of some of its vertices. In the proof of Proposition 4.17 for $r = 4$, the p_j^r are equal to 2, 1, 2 and 1 and the $h(f_j)$ are equal to 4, 0, 3 and 2.

(it is impossible by looking at $B_r(\mathbf{f})$ to know if two vertices belong to the same infinite tree). Finally, it is easy to see that a 1-admissible forest is completely described by the tree $\mathbf{U}(\mathbf{f})$ and the infinite trees it contains.

Leftmost infinite geodesics and decomposition of the skeleton in reverse trees. Let T be an infinite triangulation of a 1-gon, let ρ be its root vertex (i.e. the unique point on its boundary), and let $\mathbf{f} = \mathrm{Skel}(T)$. We have seen that for every $r \geq 0$, the paths going between the trees of $B_r(\mathbf{f})$ correspond to the leftmost geodesics from ρ to the vertices of $\partial B_r^\bullet(T)$ (cf. Figure 4.3). Therefore, infinite paths started from ρ in $\mathbf{U}(\mathbf{f})$ correspond to leftmost geodesic rays in T, so the tree of leftmost infinite geodesics in T is isomorphic to $\mathbf{U}(\mathbf{f})$.

The skeleton decomposition of infinite strips. We will also need to describe the skeleton decomposition of *strips*, which are infinite triangulations with two infinite geodesic boundaries. They correspond to the S^0 and S^1 appearing in Theorem 4.1.

Definition 4.11. An *infinite strip* is a one-ended planar triangulation bounded by two infinite geodesics γ_ℓ (on its left) and γ_r (on its right), and equipped with a root vertex ρ, such that:

(i) ρ is the only common point of γ_ℓ and γ_r,
(ii) for every $i, j \geq 0$, the path γ_r is the only geodesic from $\gamma_r(i)$ to $\gamma_r(j)$,
(iii) γ_r and γ_ℓ are the only leftmost geodesic rays in S.

Exactly as for infinite triangulations of the p-gon, if S is an infinite strip, we define its ball $B_r(S)$ of radius r as the union of all its faces containing a vertex at distance at most $r-1$ from ρ. We also define its hull $B_r^\bullet(S)$ of radius r as the union of $B_r(S)$ and all the finite connected components of its complement. To define the skeleton of an infinite strip S, we note that there is a simple transformation that associates to S an infinite triangulation of the p-gon, where p is the number of edges on $\partial B_1^\bullet(S)$. More precisely, we write $\widetilde{S}$ for the map obtained by rooting $S \backslash B_1^\bullet(S)$ at the leftmost edge of $\partial B_1^\bullet(S)$ and gluing γ_ℓ and γ_r together. We define the skeleton of S as the skeleton of $\widetilde{S}$ (see Figure 4.6). Here again, the skeleton decomposition is a bijection

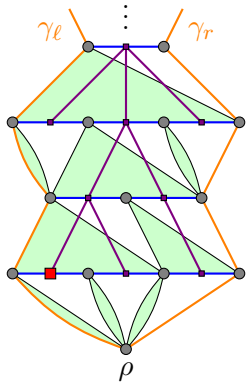

Figure 4.6 – An infinite strip S and its skeleton.

between infinite strips and pairs consisting on the one hand of an infinite reverse tree $\mathbf{t}$ that is rooted at its leftmost vertex of reverse height 0, and on the other hand of a family of maps $(M_v)_{v\in\mathbf{t}}$ such that M_v is a triangulation of a (c_v+2)-gon for every v. The fact that $\mathbf{t}$ must be connected follows from the uniqueness of the leftmost geodesic rays γ_ℓ and γ_r in a strip. Note that this time, the triangulations filling the holes are indexed by $\mathbf{t}$ and not $\mathbf{t}^*$ since the M_v for v at reverse height 0 are used to encode $B_1^\bullet(S)$ (see Figure 4.6).

Finally, let T be an infinite triangulation of a 1-gon, and let $\mathbf{f}$ be its skeleton. As we have seen above, the tree $\mathbf{U}(\mathbf{f})$ can be seen as the tree of leftmost infinite geodesics of T. Moreover, $\mathbf{U}(\mathbf{f})$ cuts T into strips, whose skeletons are the infinite reverse trees of $\mathbf{f}$. This reduces the proof of Theorem 4.1 to the study of $\mathrm{Skel}(\mathbb{T}_\lambda^1)$.

4.3.4 The distribution of S_λ^0 and S_λ^1

The goal of this subsection is to describe (without building them explicitely) the strips S_λ^0 and S_λ^1 in a way that will allow us to prove Theorem 4.1. This description will involve the quasi-stationary distribution of the branching process X_λ. Hence, we start with two explicit computations about X_λ. First, we give an explicit formula for the iterates of the generating function g_λ. We define $g_\lambda^{\circ r}$ by $g_\lambda^{\circ 0}=\mathrm{Id}$ and $g_\lambda^{\circ(r+1)}=g_\lambda\circ g_\lambda^{\circ r}$ for every $r\geq 0$. Note that $g_{\lambda_c}^{\circ r}$ is explicitely computed in [65]: we have

$$g_{\lambda_c}^{\circ r}(x)=1-\left(r+\frac{1}{\sqrt{1-x}}\right)^{-2} \tag{4.15}$$

for every $x\in[0,1]$. The iterates of g_λ in the subcritical case are also computed in [122], with different notations. Note that when $\lambda\to\lambda_c$ in the formula below, we recover (4.15).

Lemma 4.12. Let $0<\lambda<\lambda_c$. For every $r\geq 0$, we have

$$g_\lambda^{\circ r}(x)=1-\frac{1-4h}{4h\sinh^2\left(\mathrm{argsh}\sqrt{\frac{1-4h}{4h(1-x)}}+rb_\lambda\right)},$$

where $b_\lambda=\mathrm{argch}\,\frac{1}{\sqrt{4h}}=-\frac{1}{2}\ln m_\lambda$, and h is as in (4.1).

Proof. This computation is already done in Section 3.1 of [122], with different notation. In [122], the function $\varphi_t(u)$ for $0\leq t\leq 1$ is defined by

$$\varphi_t(u)=\frac{\rho s}{\alpha t}\sum_{i\geq 0}(\alpha t)^i w_{\rho s}(i+2)u^i,$$

where $\alpha = \frac{1}{12}$, $\rho = \frac{1}{12\sqrt{3}} = \lambda_c$ and $s = t\sqrt{3-2t}$. For $t = \frac{12h}{1+8h}$, we have $s = \frac{\lambda}{\lambda_c}$, so by (4.12), we have $\frac{\rho s}{\alpha t}(\alpha t)^i w_{\rho s}(i+2) = \theta_\lambda(i)$ for every $i \geq 0$, and our function g_λ corresponds to the φ_t of [122] for $t = \frac{12h}{1+8h}$. Our formula follows then immediately from Lemma 3 of [122]. Finally, using (4.2), it is easy to check that $e^{-2b_\lambda} = m_\lambda$. □

Remark 4.13. It is also possible to prove the last lemma directly by using (4.14). Indeed, the probability must add up to 1 when we sum them over q, which shows that $(h_\lambda(p))$ is an invariant measure for X_λ. If we write $H_\lambda(x) = \sum_{p\geq 1} h_\lambda(p)x^p$, this easily gives

$$H_\lambda(g_\lambda(x)) = H_\lambda(x) + H_\lambda(g_\lambda(0)),$$

so $g_\lambda^{\circ r}(x) = H_\lambda^{-1}\left(H_\lambda(x) + rH_\lambda(g_\lambda(0))\right)$. Since H_λ can be explicitly computed, this also gives the result.

Our second computation deals with the quasi-stationary measure of X_λ. We first recall a few facts about quasi-stationary measures of Galton–Watson processes (see for example Chapter 7 of [21]). For every $p \geq 1$, the ratio $\frac{\mathbb{P}_1(X_\lambda(n)=p)}{\mathbb{P}_1(X_\lambda(n)=1)}$ is nondecreasing in n and converges to $\pi_\lambda(p) < +\infty$. Moreover, let Π_λ be the generating function of $(\pi_\lambda(p))_{p\geq 1}$. In our case, it is explicit:

Lemma 4.14. For $0 < \lambda < \lambda_c$, we have

$$\Pi_\lambda(x) = \frac{1}{\sqrt{1-4h}}\left(1 - (1-x)\left(\frac{\sqrt{1-4h}+1}{\sqrt{1-4h}+\sqrt{1-4hx}}\right)^2\right).$$

We also have

$$\Pi_{\lambda_c}(x) = 2\left(\frac{1}{\sqrt{1-x}} - 1\right).$$

In particular, we have $\Pi_\lambda(\theta_\lambda(0)) = \Pi_\lambda(1-h) = \frac{1-\sqrt{1-4h}}{2h}$.

Proof. By the definition of $\pi_\lambda(p)$, we have

$$\Pi_\lambda(x) = \lim_{n\to+\infty} \frac{g_\lambda^{\circ n}(x) - g_\lambda^{\circ n}(0)}{(g_\lambda^{\circ n})'(0)}.$$

Hence, by using Lemma 4.12, the computation of $\Pi_\lambda(x)$ is straightforward. Note that the case $\lambda = \lambda_c$ already appears in [65] (in the proof of Lemma 3). □

We can now describe the distribution of S_λ^0 and S_λ^1. We will first admit the existence of reverse trees with a certain distribution (described by the next lemma), and later (Section 2.6) build these trees by a spine decomposition approach.

Lemma 4.15. There are two infinite reverse trees $\boldsymbol{\tau}_\lambda^0$ and $\boldsymbol{\tau}_\lambda^1$ whose distributions are characterized as follows. For every $r \geq 0$ and every forest $(t_1, \ldots, t_p)$ of height exactly r, we have

$$\mathbb{P}\left(B_r(\boldsymbol{\tau}_\lambda^0) = (t_1, \ldots, t_p)\right) = \frac{\pi_\lambda(p) m_\lambda^{-r}}{\Pi_\lambda(\theta_\lambda(0))} \prod_{i=1}^{p} \prod_{v\in t_i} \theta_\lambda(c_v) \tag{4.16}$$

and, if f has only one vertex at height r,

$$\mathbb{P}\left(B_r(\boldsymbol{\tau}_\lambda^1) = (t_1, \ldots, t_p)\right) = \frac{\pi_\lambda(p) m_\lambda^{-r}}{\theta_\lambda(0)} \prod_{i=1}^{p} \prod_{v\in t_i} \theta_\lambda(c_v), \tag{4.17}$$

where c_v is the number of children of v for every vertex v.

Moreover, let $\boldsymbol{\tau}_\lambda^{1,*}$ be the tree $\boldsymbol{\tau}_\lambda^1$ in which we have cut the only vertex at reverse height 0. The reverse heights in $\boldsymbol{\tau}^{1,*}$ are shifted by 1, so that the minimal reverse height in $\boldsymbol{\tau}^{1,*}$ is 0. This allows us to define the two strips that appear in Theorem 4.1.

Definition 4.16. We denote by S_λ^0 (resp. S_λ^1) the random infinite strip whose skeleton is $\boldsymbol{\tau}_\lambda^0$ (resp. $\boldsymbol{\tau}_\lambda^{1,*}$) and where conditionally on the skeleton, all the holes are filled with independent Boltzmann triangulations with parameter λ.

The reason why we need to replace $\boldsymbol{\tau}_\lambda^1$ by $\boldsymbol{\tau}_\lambda^{1,*}$ in this last definition is linked to the root transformation between $\mathbb{T}_\lambda$ and $\mathbb{T}_\lambda^1$, and will be explained in details in the end of Subsection 2.5 (see Figure 4.7).

4.3.5 Proof of Theorem 4.1

The goal of this subsection is to prove Theorem 4.1. For this, we build a random infinite 1-admissible forest $\mathbf{F}_\lambda$ directly in terms of its decomposition in reverse infinite trees, and we show that it has the same distribution as $\mathrm{Skel}(\mathbb{T}_\lambda^1)$.

We recall that $\mu_\lambda(0) = 0$ and $\mu_\lambda(k) = m_\lambda(1-m_\lambda)^{k-1}$ for $k \geq 1$. We have seen that an infinite 1-admissible forest $\mathbf{f}$ is completely described by the tree $\mathbf{U}(\mathbf{f})$ and the infinite trees that $\mathbf{f}$ contains. Therefore, there is a unique (in distribution) random 1-admissible forest $\mathbf{F}_\lambda$ such that $\mathbf{U}(\mathbf{F}_\lambda)$ is a Galton–Watson tree with offspring distribution μ_λ and, conditionally on $\mathbf{U}(\mathbf{F}_\lambda)$:

(i) the trees of $\mathbf{F}_\lambda$ are independent,

(ii) the unique tree that reaches reverse height 0 has the same distribution as $\boldsymbol{\tau}_\lambda^1$,

(iii) all the other trees have the same distribution as $\boldsymbol{\tau}_\lambda^0$ described above.

A more rigorous (but heavier) way to define $\mathbf{F}$ would be to build explicitly $B_r(\mathbf{F})$ by concatenating independent forests of the form $B_j(\boldsymbol{\tau}^0)$ and $B_j(\boldsymbol{\tau}^1)$.

In order to prove that $\mathbf{F}_\lambda$ has indeed the same distribution as $\mathrm{Skel}(\mathbb{T}_\lambda^1)$, we need to introduce one last notation. Let $\mathbf{f}$ be a 1-admissible forest, and let $r \geq 1$. Let ℓ be the number of infinite reverse trees of $\mathbf{f}$ that intersect $B_r(\mathbf{f})$, and let $\mathbf{t}_1^r, \dots, \mathbf{t}_\ell^r$ be these trees, from left to right. For every $1 \leq j \leq \ell$, we denote by $p_j^r(\mathbf{f})$ the number of vertices of $\mathbf{t}_j^r$ whose reverse height in $\mathbf{f}$ is exactly r (see Figure 4.5 for an example). If we already know $B_r(\mathbf{f})$, knowing the values $p_j^r(\mathbf{f})$ is equivalent to knowing which of the trees of $B_r(\mathbf{f})$ lie in the same infinite reverse tree of $\mathbf{f}$. We write $\widetilde{B}_r(\mathbf{f})$ for the pair $(B_r(\mathbf{f}), (p_1^r(\mathbf{f}), \dots, p_\ell^r(\mathbf{f})))$. The reason why we introduce this object is that its distribution for $\mathbf{f} = \mathbf{F}_\lambda$ will be easier to compute.

The key result is the next proposition. As explained in Section 2.1, a 1-admissible forest $\mathbf{f}$ is characterized by its reordered balls $B_r'(\mathbf{f})$, so the distribution of $\mathrm{Skel}(\mathbb{T}_\lambda^1)$ is characterized by the distribution of its reordered balls. Therefore, the next proposition will imply that $\mathbf{F}_\lambda$ and $\mathrm{Skel}(\mathbb{T}_\lambda^1)$ have the same distribution, which implies Theorem 4.1.

Proposition 4.17. For every $r \geq 0$, the forests $B_r'(\mathbf{F}_\lambda)$ and $\mathrm{Skel}\left(B_r^\bullet(\mathbb{T}_\lambda^1)\right)$ have the same distribution.

Proof. In all this proof, we will fix $0 < \lambda \leq \lambda_c$ and omit the parameter λ in the notation. The idea of the proof is the following: we will first compute the law of $\widetilde{B}_r(\mathbf{F})$, then that of $B_r(\mathbf{F})$ and finally that of $B_r'(\mathbf{F})$.

First, we know that $B_r(\mathbf{U}(\mathbf{F}))$ is a tree with height r in which all the leaves lie at height r. Moreover, if t is such a tree, we have

$$\mathbb{P}\left(B_r(\mathbf{U}(\mathbf{F})) = t\right) = \prod_{v \in t_{<r}} m(1-m)^{c_v - 1} = m^{|t_{<r}|}(1-m)^{|t|-1-|t_{<r}|}, \tag{4.18}$$

where $t_{<r}$ is the set of vertices of t at height strictly less than r, and c_v is the number of children of a vertex v.

Now let $(f,(p_1,\dots,p_\ell))$ be a possible value of $\widetilde{B}_r(\mathbf{F})$, i.e. $f=(t_1,\dots,t_p)$ is a forest of height r, the positive integers $p_1,\dots,p_\ell$ satisfy $p_1+\cdots+p_\ell=p$, and f has a unique vertex at reverse height 0 which lies in one of the trees $t_1,\dots,t_{p_1}$. For every $1\le j\le \ell$, we write $f_j=(t_{p_1+\cdots+p_{j-1}+1},\dots,t_{p_1+\cdots+p_j})$ and $h(f_j)$ for the height of the forest f_j. Each of these forests corresponds to one of the infinite trees that reach reverse height r.

We now check that the tree $B_r(\mathbf{U}(\mathbf{F}))$ is a measurable function of $\widetilde{B}_r(\mathbf{F})$ (although not of $B_r(\mathbf{F})$). The reader may find helpful to look at Figure 4.5 while reading what follows. We consider the tree u whose vertices are the pairs (j,i) with $1\le j\le\ell$ and $r-h(f_j)\le i\le r$, and in which:

(i) the pair $(1,0)$ is the root vertex,

(ii) if $i>r-h(f_j)$, then the parent of (j,i) is $(j,i-1)$,

(iii) if $i=r-h(f_j)$, then the parent of (j,i) is $(k,i-1)$, where k is the greatest integer such that $k<j$ and $h(f_k)\ge r-i+1$.

This is the natural analog of $\mathbf{U}(\mathbf{f})$ for the finite forest f (cf. Figure 4.5). Note that for every $1\le j\le\ell$, there are exactly $h(f_j)+1$ vertices of u of the form (j,i), exactly one of which lies at height r. Hence, we have

$$\sum_{j=1}^{\ell} h(f_j)=|u_{<r}|. \tag{4.19}$$

It is easy to see that if $\widetilde{B}_r(\mathbf{F})=(f,(p_1,\dots,p_\ell))$, then $B_r(\mathbf{U}(\mathbf{F}))=u$. It is also possible to read the heights $h(f_j)$ on the tree u: first, we have $h(f_1)=r$. Moreover, if $2\le j\le\ell$, let x_j be the j-th leaf in u starting from the left. Then $h(f_j)$ is the smallest h such that the ancestor of x_j at height $r-h$ is not the leftmost child of its parent (cf. Figure 4.5). Hence, conditionally on $B_r(\mathbf{U}(\mathbf{F}))=u$, the forest $\widetilde{B}_r(\mathbf{F})$ is a concatenation of ℓ independent forests of heights $h(f_1),\dots,h(f_\ell)$. The first forest has the same distribution as $B_r(\boldsymbol{\tau}^1)$ and, for $j\ge 2$, the j-th forest has the same distribution as $B_{h(f_j)}(\boldsymbol{\tau}^0)$. By combining this observation with (4.18) and Lemma 4.15, we obtain

$$\begin{aligned}
\mathbb{P}\left(\widetilde{B}_r(\mathbf{F})=(f,(p_i)_{1\le i\le\ell})\right) &= m^{|u_{<r}|}(1-m)^{|u|-|u_{<r}|-1}\times\frac{\pi(p_1)m^{-r}}{\theta(0)}\prod_{v\in f_1}\theta(c_v)\\
&\qquad\times\prod_{j=2}^{\ell}\frac{\pi(p_j)m^{-h(f_j)}}{\Pi(\theta(0))}\prod_{v\in f_j}\theta(c_v)\\
&= \frac{m^{|u_{<r}|-\sum_{j=1}^{\ell}h(f_j)}(1-m)^{|u|-|u_{<r}|-1}}{\theta(0)\,\Pi(\theta(0))^{\ell-1}}\prod_{j=1}^{\ell}\pi(p_j)\times\prod_{v\in f}\theta(c_v)\\
&= \frac{1}{\theta(0)}\left(\frac{1-m}{\Pi(\theta(0))}\right)^{\ell-1}\prod_{j=1}^{\ell}\pi(p_j)\times\prod_{v\in f}\theta(c_v),
\end{aligned}\tag{4.20}$$

where in the end, we used (4.19). Moreover, by Lemma 4.14, we can compute $\frac{1-m}{\Pi(g(0))}=\sqrt{1-4h}$.

We now compute the distribution of $B_r(\mathbf{F})$: let $f=(t_1,\dots,t_p)$ be a possible value of $B_r(\mathbf{F})$, and let i_0 be the index such that the only vertex of reverse height 0 belongs to t_{i_0}. We need to sum (4.20) over all the possible values of $\ell\ge 1$ and $p_1,\dots,p_\ell\ge 1$ such that $\sum_{j=1}^{\ell}p_j=p$ and $p_1\ge i_0$ (by construction of $\mathbf{F}$, the lowest vertex always belongs to the leftmost infinite tree). We obtain

$$\mathbb{P}\left(B_r(\mathbf{F})=f\right)=\frac{1}{\theta(0)}\left(\sum_{\ell\ge1}(1-4h)^{\frac{\ell-1}{2}}\sum_{\substack{p_1+\cdots+p_\ell=p\\ p_1\ge i_0}}\prod_{j=1}^{\ell}\pi(p_j)\right)\prod_{v\in f}\theta(c_v). \tag{4.21}$$

Now let $f' = (t_1, \dots, t_p)$ be a possible value of $B'_r(\mathbf{F})$, i.e. a forest of height r in which the only vertex of reverse height 0 lies in t_1. To obtain $\mathbb{P}\left(B'_r(\mathbf{F}) = f'\right)$, we need to sum Equation (4.21) over all the forests one may get by applying a cyclic permutation to the trees of f'. The values of p and $\prod_{v\in f}\theta(c_v)$ are the same for all these forests, but the value of i_0 ranges from 1 to p, so we have

$$\begin{aligned}\mathbb{P}\left(B'_r(\mathbf{F}) = f'\right) &= \frac{1}{\theta(0)}\Bigg(\sum_{i_0=1}^{p}\sum_{\ell\geq 1}(1-4h)^{\frac{\ell-1}{2}}\sum_{\substack{p_1+\cdots+p_\ell=p\\ p_1\geq i_0}}\prod_{j=1}^{\ell}\pi(p_j)\Bigg)\prod_{v\in f}\theta(c_v)\\ &= \frac{1}{\theta(0)}\Bigg(\sum_{\ell\geq 1}(1-4h)^{\frac{\ell-1}{2}}\sum_{p_1+\cdots+p_\ell=p} p_1\prod_{j=1}^{\ell}\pi(p_j)\Bigg)\prod_{v\in f}\theta(c_v).\end{aligned}$$

By comparing this to (4.13), we only need to prove that for every $p \geq 1$,

$$\sum_{\ell\geq 1}(1-4h)^{\frac{\ell-1}{2}}\sum_{p_1+\cdots+p_\ell=p} p_1\prod_{j=1}^{\ell}\pi(p_j) = \frac{p\,h_\lambda(p)}{h_\lambda(1)}. \tag{4.22}$$

It is enough to show that the generating functions of both sides coincide. But the generating function of the left-hand side is

$$\begin{aligned}\sum_{p\geq 1}\Bigg(\sum_{\ell\geq 1}(1-4h)^{\frac{\ell-1}{2}}\sum_{p_1+\cdots+p_\ell=p} p_1\prod_{j=1}^{\ell}\pi(p_j)\Bigg)y^p &= \sum_{\ell\geq 1}(1-4h)^{\frac{\ell-1}{2}}\, y\,\Pi'(y)\Pi(y)^{\ell-1}\\ &= \frac{y\,\Pi'(y)}{1-\sqrt{1-4h}\,\Pi(y)},\end{aligned}$$

which is explicitly known by Lemma 4.14. On the other hand, we recall from Lemma 4.10 that $h_\lambda(p) = \frac{1}{p}\left(8+\frac{1}{h}\right)^{-p} c_\lambda(p)$. Hence, the generating function of the right-hand side of (4.22) is

$$\frac{1}{c_\lambda(1)}\left(8+\frac{1}{h}\right)\sum_{p\geq 1}\left(8+\frac{1}{h}\right)^{-p} c_\lambda(p)\,y^p = \frac{1}{y}\frac{1}{c_\lambda(1)}\left(8+\frac{1}{h}\right)C_\lambda\left(\frac{h}{1+8h}y\right),$$

where C_λ is given by (4.9). Hence, it is easy to check that the generating functions of both sides of (4.22) coincide, which concludes the proof. □

End of the proof of Theorem 4.1. By Proposition 4.17, the skeleton of $\mathbb{T}^1_\lambda$ has the same distribution as $\mathbf{F}_\lambda$. In particular, the infinite leftmost geodesics are the paths separating infinite trees in $\mathbf{F}_\lambda$, so their union is isomorphic to $\mathbf{U}(\mathbf{F}_\lambda)$. Moreover, consider the strips delimited by $\mathbf{U}(\mathbf{F}_\lambda)$ in $\mathbb{T}^1_\lambda$. Let e_1 be the root edge of $\mathbb{T}^1_\lambda$ (i.e. the loop on its boundary). The skeletons of the strips that are not adjacent to e_1 are independent copies of $\boldsymbol{\tau}^0$ and all the holes are filled with independent Boltzmann triangulations with parameter λ, so these strips are independent copies of S^0.

The case of the strip containing e_1 needs to be handled more carefully because of the presence of the bottom boundary (actually, this is not exactly a strip in the sense of Definition 4.11). More precisely, let $S^{1,+}$ be the random strip obtained from $\boldsymbol{\tau}^1$ as on the left part of Figure 4.7, where all the green holes are filled with independent Boltzmann triangulations with parameter λ. Then the strip of $\mathbb{T}^1_\lambda$ adjacent to e_1 has skeleton $\boldsymbol{\tau}^1$ and its holes are filled with Boltzmann triangulations, but it has a bottom boundary of length one, so it has the same distribution as $S^{1,+}$.

Finally, we recall that $\mathbb{T}_\lambda$ is the image of $\mathbb{T}^1_\lambda$ by the following root transformation. Let f_1 be the face of $\mathbb{T}^1_\lambda$ adjacent to its boundary loop e_1. We obtain $\mathbb{T}_\lambda$ by contracting e_1 and gluing

Figure 4.7 – On the left, the strip $S^{1,+}$ and its skeleton $\boldsymbol{\tau}^1$. On the right, the strip S^1 and its skeleton $\boldsymbol{\tau}^{1,*}$. The root transformation contracting the face f_1 (whose sides are in red on the left) sends the first strip to the second. The skeletons are in violet.

together the two other sides of f_1. This transformation does not affect the tree of infinite leftmost geodesics and the strips that are not adjacent to e_1. Its effect on the strip $S^{1,+}$ adjacent to e_1 is described on Figure 4.7, and the strip we obtain is then S^1. This proves Theorem 4.1.

□

4.3.6 Construction of reverse Galton–Watson trees and infinite strips

For all this subsection, we fix $0 < \lambda \leq \lambda_c$. Most of the notation we use and introduce here will depend on λ, but the index λ will be implicit when there is no ambiguity. Our goal here is to construct the reverse subcritical trees $\boldsymbol{\tau}^0$ and $\boldsymbol{\tau}^1$ and to prove Lemma 4.15 and a few useful estimates. We will also show that $\boldsymbol{\tau}^0$ is the local limit as $n \to +\infty$ of a subcritical Galton–Watson tree conditioned to die at generation n, seen from its last generation. This is different from the more usual Galton–Watson tree conditioned to survive [88], where the tree is seen from its root. Moreover, $\boldsymbol{\tau}^1$ is just $\boldsymbol{\tau}^0$, conditioned to have only one vertex at reverse height 0. Our trees will be built by a spine decomposition approach, which will be useful to obtain geometric estimates on these trees in Section 4.

We start with a vertical half-line which is infinite on the top side, that we call the *spine*. The root of the tree will be the lowest point on the spine. For every $r \geq 1$, we sample a random pair (L_r, R_r) of integers such that all these pairs are independent and, for every $i, j \geq 0$, we have

$$\mathbb{P}\left(L_r = i, R_r = j\right) = \frac{g^{\circ r}(0) - g^{\circ(r-1)}(0)}{g^{\circ(r+1)}(0) - g^{\circ r}(0)} \theta(i+j+1)\, g^{\circ(r-1)}(0)^i\, g^{\circ r}(0)^j. \tag{4.23}$$

Lemma 4.18. Equation (4.23) defines a probability measure on $\mathbb{N}^2$, and $\mathbb{E}\left[L_r + R_r\right]$ is bounded if $\lambda < \lambda_c$, and $O(r)$ if $\lambda = \lambda_c$.

Proof. To show that (4.23) defines a probability measure on $\mathbb{N}^2$, we just need to sum the right-hand side over pairs (i, j) with a fixed value of $k = i + j$ and then sum over k. Similarly, by a straightforward computation, we obtain

$$\mathbb{E}\left[L_r + R_r\right] \;=\; \frac{g^{\circ r}(0) g'(g^{\circ r}(0)) - g^{\circ(r-1)}(0) g'(g^{\circ(r-1)}(0))}{g^{\circ(r+1)}(0) - g^{\circ r}(0)} - 1.$$

This can be rewritten $\frac{\widetilde{g}(x_r) - \widetilde{g}(y_r)}{g(x_r) - g(y_r)}$, where $x_r = g^{\circ r}(0)$, $y_r = g^{\circ(r-1)}(0)$ and $\widetilde{g}(x) = x g'(x)$ for $x \in [0, 1]$. By convexity of g and $\widetilde{g}$, we obtain

$$\mathbb{E}\left[L_r + R_r\right] \leq \frac{(x_r - y_r)\, \widetilde{g}'(x_r)}{(x_r - y_r)\, g'(y_r)} - 1 = \frac{g''(x_r)}{g'(y_r)}. \tag{4.24}$$

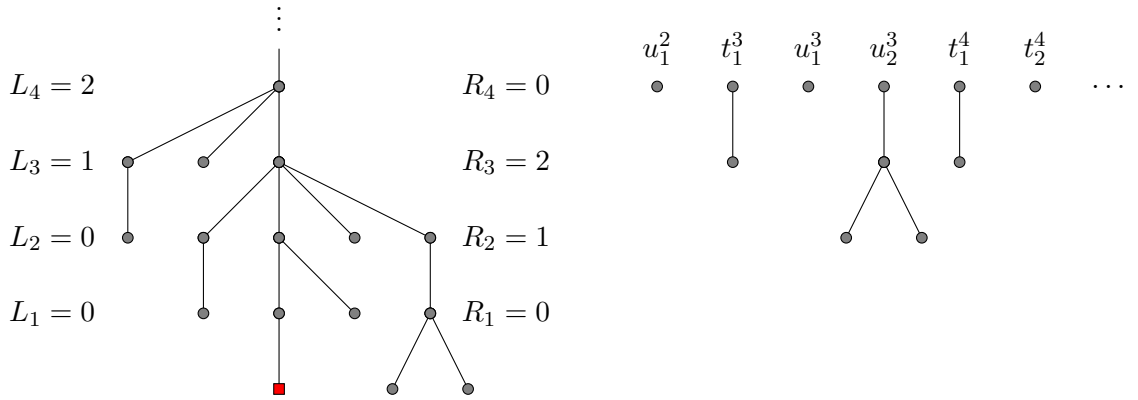

Figure 4.8 – Construction of the tree $\boldsymbol{\tau}^0$. Every vertex lies below its parent.

If $\lambda < \lambda_c$, this converges to $\frac{g''(1)}{g'(1)}$, so $\mathbb{E}[L_r + R_r]$ is bounded. For $\lambda = \lambda_c$, this is not true since $g''_{\lambda_c}(1) = +\infty$. However, by (4.15) we have $x_r = 1 - \frac{1}{(r+1)^2}$ and $y_r = 1 - \frac{1}{r^2}$. Moreover, we have $g''_{\lambda_c}(x) \sim \frac{3}{2\sqrt{1-x}}$ when $x \to 1$, so the right-hand side of (4.24) is equivalent to $\frac{3}{2}r$, and $\mathbb{E}[L_r + R_r] = O(r)$. □

We call s_r the vertex at height r in the spine. For all $r \geq 2$, let $t_1^r, \dots, t_{L_r}^r$ be L_r independent Galton–Watson trees with offspring distribution θ conditioned on having height at most $r-2$. We graft L_r edges to the left of the vertex s_r, and the trees $t_1^r, \dots, t_{L_r}^r$ to the other ends of these edges. Similarly, for all $r \geq 1$, let $u_1^r, \dots, u_{R_r}^r$ be R_r independent Galton–Watson trees with offspring distribution θ conditioned on having height at most $r-1$. We graft R_r edges to the right of s_r and the trees $u_1^r, \dots, u_{R_r}^r$ to the other ends of these edges (see Figure 4.8).

We denote by $\boldsymbol{\tau}^0$ the infinite tree we obtain, and we define a genealogy on it: for every $r \geq 1$, the children of s_r are s_{r-1} and the roots of the trees t_i^r and u_j^r. Inside the trees t_i^r and u_j^r, the genealogy is the usual one in a Galton–Watson tree. We fix the reverse height of the root at 0 and declare that the parent of a vertex of reverse height r has reverse height $r+1$ (it corresponds to the height of the vertices on Figure 4.8). By the conditioning we have chosen for the trees t_i^r and u_j^r, every vertex of the tree has a nonnegative reverse height, and the root is the leftmost vertex with reverse height 0.

Lemma 4.19. For every $r \geq 0$, the tree $\boldsymbol{\tau}^0$ has a finite number of vertices at height r.

Proof. We fix $r \geq 0$ and take $r' > r$, and we define $A_r^{r'}$ as the event

$$\{\text{one of the trees grafted at } s_{r'} \text{ reaches reverse height } r \text{ in } \boldsymbol{\tau}^0\}.$$

We want to show that $\mathbb{P}\left(A_r^{r'}\right)$ decreases fast enough in r'. Consider one of the trees $t_i^{r'}$ grafted on the left of the spine at $s_{r'}$. This is a Galton–Watson tree conditioned to have height at most $r'-2$, so the probability that $t_i^{r'}$ has height at least $r'-r-1$ is $\frac{g^{\circ(r'-1)}(0) - g^{\circ(r'-r-1)}(0)}{g^{\circ(r'-1)}(0)}$. The denominator is larger than $\frac{1}{2}$ for r' large enough. Hence, the probability that one of the $L_{r'}$ trees grafted on the left of $s_{r'}$ reaches reverse height r is at most $2\left(g^{\circ(r'-1)}(0) - g^{\circ(r'-r-1)}(0)\right)\mathbb{E}[L_{r'}]$. Similarly, the probability of the analog event on the right is at most $2\left(g^{\circ r'}(0) - g^{\circ(r'-r-1)}(0)\right)\mathbb{E}[R_{r'}]$. Therefore, for r' large enough, we have

$$\mathbb{P}\left(A_r^{r'}\right) \leq 2\left(g_\lambda^{\circ r'}(0) - g^{\circ(r'-r-1)}(0)\right)\mathbb{E}[L_{r'} + R_{r'}].$$

If $\lambda < \lambda_c$, the first factor decreases exponentially fast in r' (because $g^{\circ r'}(0) - g^{\circ(r'-r-1)} \leq 1 - g^{\circ(r'-r-1)}$, which decreases exponentially fast). If $\lambda = \lambda_c$, we have $g^{\circ r'}(0) = 1 - \frac{1}{(1+r')^2}$, so the first factor is $O\left(\frac{1}{r'^3}\right)$. By combining this with Lemma 4.18, we obtain $\mathbb{P}\left(A_r^{r'}\right) = O\left(\frac{1}{r'^2}\right)$. In all cases, we have $\sum_{r'>r} \mathbb{P}\left(A_r^{r'}\right) < +\infty$, so a.s. $A_r^{r'}$ does not occur for r' large enough, which proves the lemma. □

For $r \geq 0$, let $d_r(\boldsymbol{\tau}^0)$ be the tree of descendants of the r-th vertex s_r of the spine.

Lemma 4.20. The tree $d_r(\boldsymbol{\tau}^0)$ has the same distribution as a Galton–Watson tree with offspring distribution θ, conditioned to have height exactly r.

Proof. Let t be a finite plane tree of height r. Let $s_0(t)$ be its leftmost vertex of reverse height 0 and let $s_r(t)$ be its root. Let $(s_0(t), s_1(t), \ldots, s_r(t))$ be the unique geodesic path from $s_0(t)$ to $s_r(t)$. For every $1 \leq i \leq r$, let $\ell_i(t)$ (resp. $r_i(t)$) be the number of children of $s_i(t)$ on the left (resp. on the right) of $s_{i-1}(t)$, and let $\left(v_j^i(t)\right)_{1\leq j\leq \ell_i(t)}$ (resp. $\left(w_j^i(t)\right)_{1\leq j\leq r_i(t)}$) be those children. Finally, for every i and j, let $t_j^i(t)$ (resp. $u_j^i(t)$) be the tree of descendants of $v_j^i(t)$ (resp. $w_j^i(t)$). Then we have

$$\mathbb{P}\left(d_r(\boldsymbol{\tau}^0) = t\right) = \prod_{i=1}^{r} \mathbb{P}\left(L_i = \ell_i(t), R_i = r_i(t)\right) \times \prod_{i=1}^{r} \left(\prod_{j=1}^{\ell_i(t)} \mathbb{P}\left(t_j^i = t_j^i(t)\right) \times \prod_{j=1}^{r_i(t)} \mathbb{P}\left(u_j^i = u_j^i(t)\right) \right). \tag{4.25}$$

Moreover, by definition of t_j^i and u_j^i, we have

$$\mathbb{P}\left(t_j^i = t_j^i(t)\right) = \frac{1}{g^{\circ(r-1)}(0)} \prod_{v\in t_j^i(t)} \theta(c_v) \quad \text{and} \quad \mathbb{P}\left(u_j^i = u_j^i(t)\right) = \frac{1}{g^{\circ r}(0)} \prod_{v\in u_j^i(t)} \theta(c_v).$$

By combining this with (4.23) and (4.25), we obtain

$$\mathbb{P}\left(d_r(\boldsymbol{\tau}^0) = t\right) = \frac{1}{g^{\circ(r+1)}(0) - g^{\circ r}(0)} \prod_{v\in t} \theta(c_v),$$

which concludes. □

Now let $r \geq 0$. By Lemma 4.19, there is $r' \geq r$ such that all the vertices of reverse height at most r lie in $d_{r'}(\boldsymbol{\tau}^0)$. Hence, $\boldsymbol{\tau}^0$ is the a.s. local limit as $r' \to +\infty$ of $d_{r'}(\boldsymbol{\tau}^0)$ rooted at its leftmost vertex of height r'. By Lemma 4.20, this proves that $\boldsymbol{\tau}^0$ is the local limit (in distribution) as $r' \to +\infty$ of Galton–Watson trees conditioned on extinction at time r', seen from the last generation.

In particular, for every $r \geq 0$, let $Y(r)$ be the number of vertices of $\boldsymbol{\tau}^0$ at reverse height r. Then $(Y(-r))_{r\geq 0}$ is the limit in distribution as $n \to +\infty$ of a Galton–Watson process with offspring distribution θ, started from 1 at time $-n$ and conditioned on extinction at time exactly 1. Hence, Y has the same distribution as the reverse Galton–Watson process described by Esty in [76], started from 0 at time -1. In particular, by Equation 3 of [76], we obtain explicitly the distribution of $Y(r)$.

Lemma 4.21. For all $r \geq 0$ and $p \geq 1$, we have

$$\mathbb{P}\left(Y(r) = p\right) = \frac{\pi(p) m^{-r}}{\Pi(\theta(0))} \left(g^{\circ(r+1)}(0)^p - g^{\circ r}(0)^p \right).$$

We also define $\boldsymbol{\tau}^1$ as the tree $\boldsymbol{\tau}^0$, conditioned on $Y(0) = 1$. We can now prove Lemma 4.15.

Proof of Lemma 4.15. We start with the first part of the lemma (i.e. the part related to $\boldsymbol{\tau}^0$).

Let R be the smallest r' such that $s_{r'}$ is an ancestor of all the vertices at reverse height r in $\boldsymbol{\tau}^0$. For every $r' > r$, the event $\{R \leq r'\}$ depends only on the trees t_j^i and u_j^i for $i > r'$, so it is independent of the tree $d_{r'}(\boldsymbol{\tau}^0)$. We write $Y(r, r')$ for the number of descendants of $s_{r'}$ at reverse height r, and $B_r^{r'}(\boldsymbol{\tau}^0)$ for the forest consisting of the trees of descendants of these vertices. We have

$$\begin{aligned}
\mathbb{P}\left(B_r(\boldsymbol{\tau}^0) = (t_1, \dots, t_p)\right) &= \lim_{r' \to +\infty} \mathbb{P}\left(R \leq r', B_r(\boldsymbol{\tau}^0) = (t_1, \dots, t_p)\right) \\
&= \lim_{r' \to +\infty} \mathbb{P}\left(R \leq r', B_r^{r'}(\boldsymbol{\tau}^0) = (t_1, \dots, t_p)\right) \\
&= \lim_{r' \to +\infty} \mathbb{P}\left(R \leq r'\right) \mathbb{P}\left(Y(r, r') = p\right) \\
&\qquad \times \mathbb{P}\left(B_r^{r'}(\boldsymbol{\tau}^0) = (t_1, \dots, t_p) \middle| Y(r, r') = p\right).
\end{aligned}$$

But by Lemma 4.20, the tree $d_{r'}(\boldsymbol{\tau}^0)$ has the same distribution as a Galton–Watson tree conditioned to have height r'. From here, it is easy to show that the distribution of $B_r^{r'}(\boldsymbol{\tau}^0)$ conditioned on $Y(r, r') = p$ is that of a Galton–Watson forest of p trees conditioned to have height r, i.e.

$$\mathbb{P}\left(B_r^{r'}(\boldsymbol{\tau}^0) = (t_1, \dots, t_p) \middle| Y(r, r') = p\right) = \frac{1}{g^{\circ(r+1)}(0)^p - g^{\circ r}(0)^p} \prod_{i=1}^{p} \prod_{v \in t_i} \theta(c_v).$$

On the other hand, by using again the independence of $\{R \leq r'\}$ and $d_{r'}(\boldsymbol{\tau}^0)$, we have

$$\begin{aligned}
\lim_{r' \to +\infty} \mathbb{P}\left(R \leq r'\right) \mathbb{P}\left(Y(r, r') = p\right) &= \lim_{r' \to +\infty} \mathbb{P}\left(R \leq r', Y(r, r') = p\right) \\
&= \lim_{r' \to +\infty} \mathbb{P}\left(R \leq r', Y(r) = p\right) \\
&= \mathbb{P}\left(Y(r) = p\right) \\
&= \frac{\pi(p) m^{-r}}{\Pi(\theta(0))} \left(g^{\circ(r+1)}(0)^p - g^{\circ r}(0)^p\right)
\end{aligned}$$

by Lemma 4.21, which proves the first part of the lemma.

Given the definition of $\boldsymbol{\tau}^1$, the proof of the second part is now easy. For every forest $(t_1, \dots, t_p)$ of height r with exactly one vertex at reverse height 0, we have

$$\begin{aligned}
\mathbb{P}\left(B_r(\boldsymbol{\tau}^1) = (t_1, \dots, t_p)\right) &= \mathbb{P}\left(B_r(\boldsymbol{\tau}^0) = (t_1, \dots, t_p) \middle| Y(0) = 1\right) \\
&= \frac{\mathbb{P}\left(B_r(\boldsymbol{\tau}^0) = (t_1, \dots, t_p)\right)}{\mathbb{P}\left(Y(0) = 1\right)}.
\end{aligned}$$

By Lemma 4.21 we have $\mathbb{P}\left(Y(0) = 1\right) = \frac{\theta(0)}{\Pi(\theta(0))}$, so the second part of the lemma follows from the first one. □

We end up this subsection by showing that the strips S^0 and S^1 constructed from $\boldsymbol{\tau}^0$ and $\boldsymbol{\tau}^1$ are in some sense very close to each other. This will allow us in Section 4 to conclude a strip verifies a property if the other does, and therefore to avoid annoying case distinctions.

Lemma 4.22. Let $\lambda_n \to \lambda_c$ and let (A_n) be measurable events. Then the following two assertions are equivalent:

(i) $\mathbb{P}\left(S^0_{\lambda_n} \in A_n\right) \xrightarrow[n \to +\infty]{} 0$,

(ii) $\mathbb{P}\left(S^1_{\lambda_n} \in A_n\right) \xrightarrow[n\to+\infty]{} 0$.

Proof. We recall that $\boldsymbol{\tau}^{1,*}$ is the tree $\boldsymbol{\tau}^1$ in which we have cut the only vertex at reverse height 0, and we have shifted the reverse heights by 1. The main part of the proof consists in showing that the trees $\boldsymbol{\tau}^0_{\lambda_n}$ and $\boldsymbol{\tau}^{1,*}_{\lambda_n}$ are absolutely continuous with respect to each other, uniformly in n. We claim that $\boldsymbol{\tau}^{1,*}$ has the same distribution as $\boldsymbol{\tau}^0$ biased by $Y(0)$.

Indeed, for any forest $f \in \mathscr{F}_{p,q,r}$, we have

$$\mathbb{P}\left(B_r(\boldsymbol{\tau}^{1,*}) = f\right) = \sum_{v \in f\setminus f^*} \mathbb{P}\left(B_{r+1}(\boldsymbol{\tau}^1) = f_v^+\right),$$

where $f\setminus f^*$ is the set of vertices of f at reverse height 0, and f_v^+ is the forest of $\mathscr{F}_{1,q,r+1}$ obtained by adding a unique child to v in f. By the second part of Lemma 4.15, we obtain

$$\mathbb{P}\left(B_r(\boldsymbol{\tau}^{1,*}) = f\right) = \frac{\theta(1)}{\theta(0)} p\pi(q) m^{-r-1} \prod_{v\in f} \theta(c_v).$$

Combined with the first part of Lemma 4.15, this yields

$$\frac{\mathbb{P}\left(B_r(\boldsymbol{\tau}^{1,*}) = f\right)}{\mathbb{P}\left(B_r(\boldsymbol{\tau}^0) = f\right)} = \frac{\theta(1)}{\theta(0)} \frac{\Pi\left(\theta(0)\right)}{m} p. \tag{4.26}$$

In other words, the forest $B_r(\boldsymbol{\tau}^{1,*})$ has the distribution of $B_r(\boldsymbol{\tau}^0)$, biased by $Y(0)$. Since it is true for all r, we can conclude that $\boldsymbol{\tau}^{1,*}$ has the distribution of $\boldsymbol{\tau}^0$ biased by $Y(0)$.

From here, we can deduce a version of the lemma for the trees $\boldsymbol{\tau}^0$ and $\boldsymbol{\tau}^{1,*}$. More precisely, we claim that if (A'_n) are measurable events, the following two assertions are equivalent:

(i') $\mathbb{P}\left(\boldsymbol{\tau}^0_{\lambda_n} \in A'_n\right) \xrightarrow[n\to+\infty]{} 0$,

(ii') $\mathbb{P}\left(\boldsymbol{\tau}^{1,*}_{\lambda_n} \in A'_n\right) \xrightarrow[n\to+\infty]{} 0$.

Indeed, for any $n \geq 0$, we have

$$\mathbb{P}\left(\boldsymbol{\tau}^{1,*}_{\lambda_n} \in A'_n\right) = \frac{\mathbb{E}\left[Y_{\lambda_n}(0)\mathbb{1}_{\boldsymbol{\tau}^0_{\lambda_n}\in A'_n}\right]}{\mathbb{E}\left[Y_{\lambda_n}(0)\right]}. \tag{4.27}$$

Moreover, (4.26) shows that

$$\mathbb{E}\left[Y_\lambda(0)\right] = \frac{\theta_\lambda(0)}{\theta_\lambda(1)} \frac{m_\lambda}{\Pi_\lambda\left(\theta_\lambda(0)\right)} = \frac{1-h}{2h(1-h)} \frac{1-2h-\sqrt{1-4h}}{1-\sqrt{1-4h}},$$

which is continuous at $\lambda = \lambda_c$. Therefore, the denominator in (4.27) converges to a positive limit, so it is enough to prove that $\mathbb{P}\left(\boldsymbol{\tau}^0_{\lambda_n} \in A'_n\right) \to 0$ if and only if $\mathbb{E}\left[Y_{\lambda_n}(0)\mathbb{1}_{\boldsymbol{\tau}^0_{\lambda_n}\in A'_n}\right] \to 0$. The indirect implication is immediate since $Y_{\lambda_n}(0) \geq 1$. To prove the direct one, it is enough to check the variables $Y_{\lambda_n}(0)$ are uniformly integrable. We know that they converge to $Y_{\lambda_c}(0)$ in distribution. By the Skorokhod representation theorem, we may assume the convergence is almost sure. Since we also have convergence of the expectations by (4.27), by Scheffé's Lemma we have $Y_{\lambda_n}(0) \to Y_{\lambda_c}(0)$ in L^1. In particular, the family $(Y_{\lambda_n}(0))$ is uniformly integrable, which proves the direct implication, and the equivalence of (i') and (ii').

The lemma now just follows from the fact that $S^1_{\lambda_n}$ is constructed from $\boldsymbol{\tau}^{1,*}_{\lambda_n}$ in the exact same way as $S^0_{\lambda_n}$ is constructed from $\boldsymbol{\tau}^0_{\lambda_n}$. Writing it properly would require to condition the strip on its skeleton. □

4.4 The Poisson boundary of $\mathbb{T}_\lambda$

The goal of this section is to prove Theorem 4.2. We fix $0 < \lambda < \lambda_c$ until the end of the section and omit the parameter λ in most of the notation.

4.4.1 Construction of the geodesic boundary

We start by defining precisely the compactification of $\mathbb{T}$ that we will afterwards prove to be a realization of its Poisson boundary. We recall that $\partial\mathbf{T}^g$ is the set of ends of the tree $\mathbf{T}^g$. If $\gamma, \gamma' \in \partial\mathbf{T}^g$, we write $\gamma \sim \gamma'$ if $\gamma = \gamma'$ or if γ and γ' are the left and right boundaries of one of the copies of S^0 or S^1 (cf. left part of Figure 4.1). Note that a.s., every ray of $\mathbf{T}^g$ branches infinitely many times, so no ray is equivalent to two distinct other rays. It follows that $\sim$ is a.s. an equivalence relation, for which countably many equivalence classes have cardinal 2, and all the others have cardinal 1. We write $\widehat{\partial}\mathbf{T}^g = \partial\mathbf{T}^g/\sim$, i.e. we identify two geodesic rays if they are the left and right boundaries of the same strip. We also write $\gamma \to \widehat{\gamma}$ for the canonical projection from $\partial\mathbf{T}^g$ to $\widehat{\partial}\mathbf{T}^g$. If S is one the copies of S^0 or S^1 appearing in the strip decomposition, then the two geodesics bounding S correspond to the same point of $\widehat{\partial}\mathbf{T}^g$, that we denote by $\widehat{\gamma}_S$.

Our goal is now to define a topology on $\mathbb{T} \cup \widehat{\partial}\mathbf{T}^g$. It should be possible to define it by an explicit distance, but such a distance would be tedious to write down, so we prefer to give an "abstract" construction. Let S and S' be two distinct strips appearing in the strip decomposition of $\mathbb{T}$ (cf. left part of Figure 4.1). Consider the smallest r such that S and S' both intersect $B_r^\bullet(\mathbb{T})$. Then $\mathbb{T}\backslash\left(B_r^\bullet(\mathbb{T}) \cup S \cup S'\right)$ has two connected components, that we denote by (S, S') and (S', S) (the vertices on the geodesics bounding S and S' do not belong to (S, S') and (S', S)). This notation means that we see (S, S') and (S', S) as the two "open intervals" between S and S'. We also write

$$\partial_g\left(S, S'\right) = \{\widehat{\gamma} |\ \gamma \text{ is a ray of } \mathbf{T}^g \text{ such that } \gamma(i) \in (S, S') \text{ for } i \text{ large enough}\}$$

We define $\partial_g(S', S)$ similarly. Note that $\partial_g(S, S')$ and $\partial_g(S', S)$ are disjoint, and their union is $\widehat{\partial}\mathbf{T}^g\backslash\{\widehat{\gamma}_S, \widehat{\gamma}_{S'}\}$.

Definition 4.23. The *geodesic compactification* of $\mathbb{T}$ is the set $\mathbb{T} \cup \widehat{\partial}\mathbf{T}^g$, equipped with the topology generated by

- the singletons $\{x\}$, where x is a vertex of $\mathbb{T}$,
- the sets $(S, S') \cup \partial_g(S, S')$, where S and S' are two distinct strips appearing in the strip decomposition of $\mathbb{T}$.

This topology is separated (if $\widehat{\gamma}_1 \neq \widehat{\gamma}_2$, then there are two strips separating γ_1 and γ_2) and has a countable basis, so it is induced by a distance. Moreover, any open set of our basis intersects $\mathbb{T}$, so $\mathbb{T}$ is dense in $\mathbb{T} \cup \widehat{\partial}\mathbf{T}^g$. The end of this subsection is devoted to the proof of two very intuitive topological properties of the geodesic compactification.

Lemma 4.24. The space $\mathbb{T} \cup \widehat{\partial}\mathbf{T}^g$ is compact.

Proof. Let (x_n) be a sequence with values in $\mathbb{T} \cup \widehat{\partial}\mathbf{T}^g$. We first assume $x_n \in \mathbb{T}$ for every n. We may assume $d(\rho, x_n) \to +\infty$. If there is a strip S that contains infinitely many x_n, then $\widehat{\gamma}_S$ is a subsequential limit of (x_n). We now assume it is not the case. We recall that for every vertex v of $\mathbf{T}^g$, the slice $\mathbb{S}[v]$ is the part of $\mathbb{T}$ lying between the leftmost and the rightmost rays passing through v, above v (see Figure 4.1). We will construct a ray γ of $\mathbf{T}^g$ step by step, in such a way that for every $k \geq 0$, there are infinitely many points x_n in $\mathbb{S}[\gamma(k)]$.

Assume we have already built $\gamma(0), \ldots, \gamma(k)$. If $\gamma(k)$ has only one child in $\mathbf{T}^g$, then $\gamma(k+1)$ must be this child. If $\gamma(k)$ has $d \geq 2$ children, we call them $y_1, \ldots, y_d$. Then $\mathbb{S}[\gamma(k)]$ is the union

of the slices $\mathbb{S}[y_i]$ for $1 \leq i \leq d$ and of the $d-1$ strips whose lowest vertex is $\gamma(k)$. We know that $\mathbb{S}[\gamma(k)]$ contains infinitely many of the vertices x_n, but the $d-1$ strips contain finitely many of them. Therefore, there is an index $1 \leq i_0 \leq d$ such that $\mathbb{S}[y_{i_0}]$ contains infinitely many of them. We choose $\gamma(k+1) = y_{i_0}$. We can check that the class $\widehat{\gamma}$ of the geodesic we built is a subsequential limit of (x_n), which concludes the case where $x_n \in \mathbb{T}$ for every n.

Finally, let δ be a distance on $\mathbb{T} \cup \widehat{\partial}\mathbf{T}^g$ that generates its topology, and let (x_n) be any sequence in $\mathbb{T} \cup \widehat{\partial}\mathbf{T}^g$. If $x_n \in \widehat{\partial}\mathbf{T}^g$, let $y_n \in \mathbb{T}$ be such that $\delta(x_n, y_n) \leq \frac{1}{n}$ (it exists by density). If $x_n \in \mathbb{T}$, we take $y_n = x_n$. By the first case (y_n) has a subsequential limit, so (x_n) also has one. □

Lemma 4.25. The boundary $\widehat{\partial}\mathbf{T}^g$ is homeomorphic to the circle.

Proof. We build an explicit homeomorphism. Consider a ray γ of $\mathbf{T}^g$. For every $i \geq 0$, let $c_\gamma(i)$ be the number of children of $\gamma(i)$ in $\mathbf{T}^g$. We denote these children by $x_0, \dots, x_{c_\gamma(i)-1}$ and we denote by $j_\gamma(i)$ the index $j \in \{0, 1, \dots, c_\gamma(i)-1\}$ such that $\gamma(i+1) = x_j$. We also define

$$\begin{array}{rcl} \Psi : \partial\mathbf{T}^g & \longrightarrow & \mathbb{R}/\mathbb{Z} \\ \gamma & \longrightarrow & \sum_{k\geq 0} \frac{j_\gamma(i)}{\prod_{i=0}^{k} c_\gamma(i)} \pmod 1. \end{array}$$

If γ_ℓ and γ_r are the left and right boundaries of a copy of S^0, then there is i_0 such that:
— for $i < i_0$, we have $c_{\gamma_r}(i) = c_{\gamma_\ell}(i)$ and $j_{\gamma_r}(i) = j_{\gamma_\ell}(i)$,
— we have $c_{\gamma_r}(i_0) = c_{\gamma_\ell}(i_0)$ and $j_{\gamma_r}(i_0) = j_{\gamma_\ell}(i_0) + 1$,
— for $i > i_0$, we have $j_{\gamma_\ell}(i) = c_{\gamma_\ell}(i) - 1$ and $j_{\gamma_r}(i) = 0$,
which implies $\Psi(\gamma_\ell) = \Psi(\gamma_r)$. Moreover, if γ_ℓ and γ_r are respectively the leftmost and rightmost rays of $\mathbf{T}^g$, then $\Psi(\gamma_\ell) = \Psi(\gamma_r)$ (the sum is equal to 0 for γ_ℓ and to 1 for γ_r). Hence, we have $\Psi(\gamma) = \Psi(\gamma')$ as soon as $\gamma \sim \gamma'$, so Ψ defines an application from $\widehat{\partial}\mathbf{T}^g$ to $\mathbb{R}/\mathbb{Z}$. The verification that this application is a homeomorphism is easy and left to the reader. □

4.4.2 Proof of Theorem 4.2

We first argue why the second point of Theorem 4.2 is an easy consequence of the first one. Assume the first point is proved. Since $\widehat{\partial}\mathbf{T}^g$ is homeomorphic to the circle, we can embed $\mathbb{T} \cup \widehat{\partial}\mathbf{T}^g$ in the unit disk $\mathbb{D}$ in such a way that $\widehat{\partial}\mathbf{T}^g$ is sent to $\partial\mathbb{D}$ (we do not describe the embedding explicitly). In this embedding, the simple random walk converges to a point of $\partial\mathbb{D}$ and the law of the limit point is a.s. non-atomic. Therefore, by Theorem 1.3 of [83], $\widehat{\partial}\mathbf{T}^g$ is a realization of the Poisson boundary of $\mathbb{T}$.

The idea of the proof of the first point is to first show that two independent random walks are quite well separated in terms of $\widehat{\partial}\mathbf{T}^g$. Let X^1 and X^2 be two independent simple random walks started from ρ. By Proposition 11 of [58] we know that $\mathbb{T}$ is transient and does not have the intersection property (although [58] deals with type-II triangulations, all the arguments of the proof still hold in our case). Hence, the complementary of $\{X^1_n | n \in \mathbb{N}\} \cup \{X^2_n | n \in \mathbb{N}\}$ has two infinite connected components with infinite boundaries. We denote them by $\left[X^1, X^2\right]$ and $\left[X^2, X^1\right]$. By point (ii) of Lemma 4.4 applied to the peeling along X_1 and X_2 (see also Section 3.2 of [58] for a similar argument), the halfplanar triangulations $\left[X^1, X^2\right]$ and $\left[X^2, X^1\right]$ are independent copies of $\mathbb{H} = \mathbb{H}_\lambda$.

Therefore, we will need to study geodesics in halfplanar triangulations.

Definition 4.26. Let H be a halfplanar triangulation and ∂H its boundary. An *infinite geodesic away from the boundary* is a sequence $(\gamma(n))_{n\geq 0}$ of vertices of H such that for any $n \geq 0$:
(i) the vertices $\gamma(n)$ and $\gamma(n+1)$ are neighbours,
(ii) we have $d(\gamma(n), \partial H) = n$.

Remark 4.27. Contrary to geodesics away from a point, the existence of such geodesics is not obvious. For example, they do not exist in the UIHPT $\mathbb{H}_{\lambda_c}$.

Lemma 4.28. Almost surely, there is a geodesic away from the boundary in $\mathbb{H}_\lambda$.

Proof. We write $P_r = |\partial B_r^\bullet(\mathbb{T})|$ and $L_r = |\partial B_r^\bullet(\mathbb{T}) \cap \mathbf{T}^g|$. We recall that $\mathbf{T}^g$ is a Galton–Watson tree with offspring distribution μ given by Theorem 4.1. In particular, we have $\sum_{i\geq 0} i\mu(i) = m_\lambda^{-1}$ and $\sum_{i\geq 0}(i\log i)\mu(i) < +\infty$, so by the Kesten–Stigum Theorem $m_\lambda^r L_r$ converges a.s. to a positive random variable. On the other hand, as in Section 2 of [58], it holds that $m_\lambda^r P_r$ converges a.s. to a positive random variable ([58] only deals with type-II triangulations but we can either apply the same proof to the type-I case, or use the core decomposition of the Appendix A of this thesis to deduce the type-I exponential growth rate from the type-II one). Hence, there is a constant $c > 0$ such that, for r large enough,

$$\mathbb{P}\left(L_r \geq cP_r\right) \geq \frac{1}{2}.$$

Therefore, if z_r is a random vertex chosen uniformly on $\partial B_r^\bullet(\mathbb{T})$, for r large enough we have $\mathbb{P}\left(z_r \in \mathbf{T}^g\right) \geq \frac{c}{2}$. Hence, for any $s > 0$, with probability at least $\frac{c}{2}$, there is a point $x \in \partial B_{r+s}^\bullet(\mathbb{T})$ at distance exactly s from z_r.

But $\mathbb{H}$ is the local limit as $r \to +\infty$ of $\mathbb{T}\backslash B_r^\bullet(\mathbb{T})$ (rooted at a uniform edge on its boundary), so if ρ denotes the root vertex of $\mathbb{H}$, for any $s \geq 0$, we have

$$\mathbb{P}\left(\text{there is } x \in \mathbb{H} \text{ such that } d(x,\rho) = d(x,\partial\mathbb{H}) = s\right) \geq \frac{c}{2}.$$

This event is nonincreasing in s, so with probability at least $\frac{c}{2}$ it occurs for every s. By a compactness argument, with probability at least $\frac{c}{2}$, there is an infinite geodesic away from the boundary γ with $\gamma(0) = \rho$. Finally, we claim that $\mathbb{H}$ is invariant under root translation, and that the root translation is ergodic, which is enough to conclude. Indeed, by local limit, the invariance is a consequence of the invariance of $\mathbb{T}^p$ for every p under re-rooting along the boundary. The ergodicity is proved in the type-II case in [18] (this is Proposition 1.3), and the proof adapts well here. □

We can now show that X^1 and X^2 are a.s. separated by an infinite leftmost geodesic.

Lemma 4.29. Almost surely, there is a ray $\widehat{\gamma}$ of $\mathbf{T}^g$ such that for n large enough, we have $\widehat{\gamma}(n) \in \left[X^1, X^2\right]$, and the same is true for $\left[X^2, X^1\right]$.

Proof. In this proof, we will write $H = \left[X^1, X^2\right]$ to avoid too heavy notations. Let γ be an infinite geodesic away from the boundary in H, which exists by Lemma 4.28. For $n \geq 0$, let $(\gamma_n(i))_{0\leq i\leq d(\rho,\gamma(n))}$ be the leftmost geodesic from ρ to $\gamma(n)$ (cf. Figure 4.9). For $n \geq i \geq d(\rho,\gamma(0))$, we have $i \leq n = d(\gamma(n),\partial H) \leq d(\gamma(n),\rho)$ so $\gamma_n(i)$ is well defined, and

$$\begin{aligned} d(\gamma(n),\gamma_n(i)) &= d(\rho,\gamma(n)) - i && \text{(since } \gamma \text{ is a geodesic)} \\ &\leq n + d(\rho,\gamma(0)) - i && \text{(by the triangular inequality)} \\ &\leq n \\ &= d(\gamma(n),\partial H), \end{aligned}$$

so $\gamma_n(i) \in H$. But by a compactness argument, there is an infinite path $\widetilde{\gamma}$ in $\mathbb{T}$ such that for any i, there are infinitely many n such that $\widetilde{\gamma}(i) = \gamma_n(i)$. It easy to check that $\widetilde{\gamma}$ is an infinite leftmost geodesic in $\mathbb{T}$. Moreover, since $\gamma_n(i) \in H$ for $n \geq i \geq d(\rho,\gamma(0))$, we have $\widetilde{\gamma}(i) \in H$ for i large enough, so there is an infinite leftmost geodesic of $\mathbb{T}$ that lies in H eventually. □

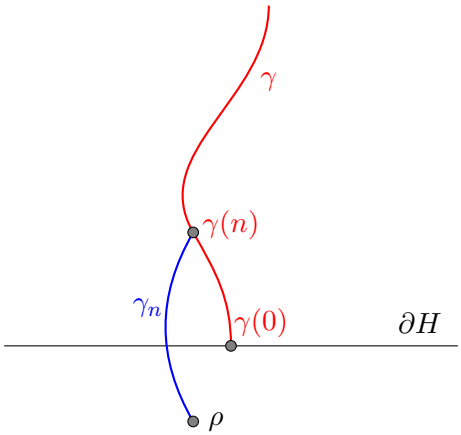

Figure 4.9 – From an infinite geodesic away from the boundary in H, we can build an infinite leftmost geodesic in $\mathbb{T}$, that lies in H eventually.

Proof of Theorem 4.2. We can now prove the convergence of the simple random walk to a point of $\widehat{\partial}\mathbf{T}^g$. If X is a simple random walk, we write $\mathrm{Acc}(X)$ for the set of accumulation points of X on $\widehat{\partial}\mathbf{T}^g$. By Lemma 4.24, it is enough to prove that $\mathrm{Acc}(X)$ is reduced to a point. We first claim that $\mathrm{Acc}(X)$ is a circle arc of $\widehat{\partial}\mathbf{T}^g$. Indeed, assume $\widehat{\gamma}_1 \neq \widehat{\gamma}_2$ are two points of $\mathrm{Acc}(X)$. Then $\widehat{\partial}\mathbf{T}^g \backslash \{\widehat{\gamma}_1, \widehat{\gamma}_2\}$ has two connected components, that we denote by $(\widehat{\gamma}_1, \widehat{\gamma}_2)$ and $(\widehat{\gamma}_2, \widehat{\gamma}_1)$. To oscillate between $\widehat{\gamma}_1$ and $\widehat{\gamma}_2$, the walk X must intersect infinitely many times either all the γ such that $\widehat{\gamma} \in (\widehat{\gamma}_1, \widehat{\gamma}_2)$ or all the γ such that $\widehat{\gamma} \in (\widehat{\gamma}_2, \widehat{\gamma}_1)$ (see Figure 4.10). In both cases, $\mathrm{Acc}(X)$ contains one of the two arcs from $\widehat{\gamma}_1$ to $\widehat{\gamma}_2$. Hence, $\mathrm{Acc}(X)$ is closed and connected, so it is a circle arc.

Now let ν be a probability measure with no atom and full support on $\widehat{\partial}\mathbf{T}^g$ (for example, one can consider the exit measure of the nonbacktracking random walk on $\mathbf{T}^g$). Either $\mathrm{Acc}(X)$ is a singleton, or it has positive measure, so it is enough to show that $\mathbb{P}\left(\nu\left(\mathrm{Acc}(X)\right) > 0\right) = 0$. By Lemma 4.29, we know that if X^1 and X^2 are two independent simple random walks started from ρ, then there are two rays of $\mathbf{T}^g$ lying respectively in $[X_1, X_2]$ and $[X_2, X_1]$ eventually, so they separate X^1 and X^2. Therefore, $\mathrm{Acc}(X^1) \cap \mathrm{Acc}(X^2)$ contains at most two points, so $\nu\left(\mathrm{Acc}(X^1) \cap \mathrm{Acc}(X^2)\right) = 0$ a.s.. Let X^i for $i \in \mathbb{N}$ be infinitely many independent simple random walks started from ρ. We have

$$\sum_{i \geq 0} \nu\left(\mathrm{Acc}(X^i)\right) = \nu\left(\bigcup_{i \geq 0} \mathrm{Acc}(X^i)\right) \leq 1.$$

But the $\nu\left(\mathrm{Acc}(X^i)\right)$ are i.i.d. (conditionally on $\mathbb{T}$), so they must be 0 a.s.. Therefore, $\mathrm{Acc}(X)$ cannot have positive measure so it is a.s. a point.

Hence, the simple random walk a.s. converges to a point of $\widehat{\partial}\mathbf{T}^g$, so it defines an exit measure on $\widehat{\partial}\mathbf{T}^g$, that we denote by ν_∂. We now prove that ν_∂ is nonatomic. Once again, our main tool is Lemma 4.29. Assume ν_∂ has an atom with positive probability, and let $(X^i)_{1 \leq i \leq 4}$ be four independent SRW started from ρ. For every $1 \leq i \leq 4$, let X^i_∞ be the limit of X^i in $\widehat{\partial}\mathbf{T}^g$. Then $\mathbb{P}\left(X^1_\infty = X^2_\infty = X^3_\infty = X^4_\infty\right) > 0$. If this happens, we can assume (up to a factor $\frac{1}{24}$) that they lie in clockwise order, and that $\widehat{\partial}\mathbf{T}^g \backslash \{X^1_\infty\}$ lies in the part between X^4 and X^1. By Lemma 4.29, there are at least one ray of $\mathbf{T}^g$ between X^1 and X^2, one between X^2 and X^3 and one between X^3 and X^4 (cf. Figure 4.11). Hence, there are at least three rays of $\mathbf{T}^g$ in the part between X^1 and X^4 that does not contain $\widehat{\partial}\mathbf{T}^g \backslash \{X^1_\infty\}$. In particular, two of them are not equivalent for $\sim$

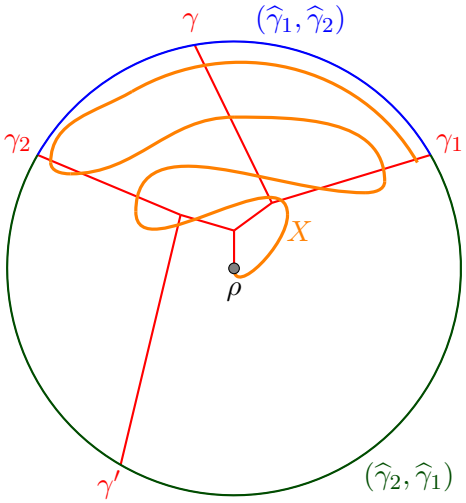

Figure 4.10 – Proof that $\mathrm{Acc}(X)$ is an arc circle: assume X oscillates between γ_1 and γ_2 and there is $\gamma' \in (\widehat{\gamma}_2, \widehat{\gamma}_1)$ that X intersects only finitely many times. Then X intersects infinitely many times every γ with $\widehat{\gamma} \in (\widehat{\gamma}_1, \widehat{\gamma}_2)$.

(we recall that the equivalence classes have cardinal at most 2). Hence, the part lying between X^1 and X^4 containing X^2 and X^3 (the left part on Figure 4.11) also contains a slice of the form $\mathbb{S}[y]$ (see Figure 4.1), so $X^1_\infty \neq X^4_\infty$. We get a contradiction.

We finally show that ν_∂ has full support. Since ν_∂ is nonatomic, we have $\nu_\partial(\{\widehat{\gamma}_{S^1}\}) = 0$ a.s.. Hence, a.s., we have $X_n \in \mathbb{S}[\rho] = \mathbb{S}$ for n large enough. Therefore, we also have

$$\mathbb{P}(\forall k \geq 0, X_k \in \mathbb{S}[\rho] | X_0 = \rho) > 0 \qquad \text{a.s.}$$

Now fix $r > 0$, condition on $B^\bullet_r(\mathbb{T})$ and take $x \in \mathbf{T}^g$ such that $d(\rho, x) = r$. Since $\mathbb{S}[x]$ has the same distribution as $\mathbb{S}[\rho]$ (see Figure 4.1), we have

$$\mathbb{P}(\forall k \geq 0, X_k \in \mathbb{S}[x] | X_0 = x) > 0 \qquad \text{a.s..}$$

Hence, we have $\mathbb{P}(X_k \in \mathbb{S}[x] \text{ for } k \text{ large enough} | X_0 = \rho) > 0$ a.s.. If this occurs, then X_∞ is of the form $\widehat{\gamma}$, where $\gamma(k) \in \mathbb{S}[x]$ for k large enough. Therefore, we have

$$\nu_\partial(\{\widehat{\gamma} | \gamma \text{ lie in } \mathbb{S}[x] \text{ eventually}\}) > 0 \qquad \text{a.s..}$$

Almost surely, this holds for every $x \in \mathbf{T}^g$, which is enough to ensure that ν_∂ has full support. This ends the proof of Theorem 4.2. □

Remark 4.30. We end this section with a remark about the Gromov boundary [81], which is another natural notion of boundary for an infinite graph $G = (V, E)$. Let $\mathcal{C}(G)$ be the space of functions $f : V \to \mathbb{R}$ equipped with the product topology. We say that two functions of $\mathcal{C}(G)$ are equivalent if they are equal up to an additive constant. We quotient $\mathcal{C}(G)$ by this equivalence relation to obtain the quotient space $\mathcal{C}(G)/\mathbb{R}$. If $x \in V$, we define $f_x \in \mathcal{C}(G)/\mathbb{R}$ by $f_x(y) = d_G(x, y)$ for any $y \in V$. The *Gromov compactification* $\widehat{G}$ of G is the closure of $\{f_x | x \in V\}$ in $\mathcal{C}(G)/\mathbb{R}$ and the *Gromov boundary* $\partial_{Gr}G$ of G is the set $\widehat{G} \backslash \{f_x | x \in V\}$.

It is easy to show that for any geodesic ray γ, the sequence $f_{\gamma(n)}$ converges in $\mathcal{C}(G)/\mathbb{R}$, so it defines a point $f_\gamma \in \partial_{Gr}\mathbb{T}$. A natural question is to ask whether there is a natural correspondence between $\widehat{\partial}\mathbf{T}^g$ and $\partial_{Gr}\mathbb{T}$. The answer is no.

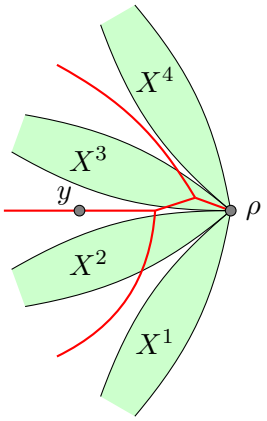

Figure 4.11 – Sketch of the proof that ν_∂ is nonatomic: if $X^1, \dots, X^4$ are four independent simple random walks, then there are three infinite leftmost geodesics (in red) between X^1 and X^4, so there is a slice $\mathbb{S}[y]$ there. The hulls of the four random walks are in green.

To prove it, we show that if γ_1 and γ_2 are respectively the left and right boundary of the same strip, then $f_{\gamma_1} \neq f_{\gamma_2}$. Indeed, let n be such that $\gamma_1(n) \notin \gamma_2$. We have $f_{\gamma_1}(\gamma_1(n)) - f_{\gamma_1}(\rho) = -n$ by definition of f_{γ_1}. But if we had $f_{\gamma_2}(\gamma_1(n)) - f_{\gamma_2}(\rho) = -n$, this would mean that there is $m > n$ such that $d(\gamma_2(m), \gamma_1(n)) - d(\gamma_2(m), \rho) = -n$, i.e. $d(\gamma_2(m), \gamma_1(n)) = m - n$. Take such an m minimal. Then by concatenating γ_1 from ρ to $\gamma_1(n)$ and a geodesic from $\gamma_1(n)$ to $\gamma_2(m)$, we obtain a geodesic from ρ to $\gamma_2(m)$ that lies strictly to the left of γ_2. This contradicts the fact that γ_2 is a leftmost geodesic in $\mathbb{T}$. This suggests that $\partial_{Gr}\mathbb{T}$ should not be homeomorphic to the circle, but rather to a Cantor set.

4.5 The tree of infinite geodesics in the hyperbolic Brownian plane

4.5.1 The tree $\mathbf{T}^g(\mathcal{P}^h)$

The goal of this section is to take the scaling limit of Theorem 4.1 and to prove Theorem 4.3. For all this section, we fix a sequence (λ_n) of numbers in $(0, \lambda_c]$ such that $\lambda_n = \lambda_c\left(1 - \frac{2}{3n^4}\right) + o\left(\frac{1}{n^4}\right)$. We know, by the main result of Chapter 3, that $\frac{1}{n}\mathbb{T}_{\lambda_n}$ converges for the local Gromov–Hausdorff distance to $\mathcal{P}^h$. Therefore, it seems reasonable that $\mathbf{T}^g(\mathcal{P}^h)$ should be the scaling limit of the trees $\mathbf{T}^g_{\lambda_n}$. This scaling limit is easy to describe. We recall that $\mathbf{B}$ is the infinite tree in which every vertex has exactly two children, except the root which has one. For $\alpha > 0$, we denote by $\mathbf{Y}_\alpha$ the Yule tree of parameter α, i.e. the tree $\mathbf{B}$ in which the lengths of the edges are i.i.d. exponential variables with parameter α.

Lemma 4.31. The trees $\frac{1}{n}\mathbf{T}^g_{\lambda_n}$ converge for the local Gromov–Hausdorff distance to $\mathbf{Y}_{2\sqrt{2}}$.

Sketch of proof. We recall that $\mathbf{T}^g_{\lambda_n}$ is a Galton–Watson tree with offspring distribution μ_{λ_n}, where $\mu_{\lambda_n}(k) = m_{\lambda_n}(1 - m_{\lambda_n})^{k-1}$ for every $k \geq 1$, and m_λ is explicitly given by (4.2). We can compute $m_{\lambda_n} = 1 - \frac{2\sqrt{2}}{n} + O\left(\frac{1}{n^2}\right)$. Let $\widetilde{\mu}_\lambda$ be the distribution defined by $\widetilde{\mu}_\lambda(1) = m_\lambda$ and $\widetilde{\mu}_\lambda(2) = 1 - m_\lambda$, and let $\widetilde{\mathbf{T}}^g_\lambda$ be a Galton–Watson tree with offspring distribution $\widetilde{\mu}_\lambda$. Then $\widetilde{\mathbf{T}}^g_\lambda$ is a copy of $\mathbf{B}$ where the length of each edge is a geometric variable of parameter $1 - m_\lambda$, so $\frac{1}{n}\widetilde{\mathbf{T}}^g_{\lambda_n}$ converges to $\mathbf{Y}_{2\sqrt{2}}$. Moreover, $\mathbf{T}^g_\lambda$ can be obtained by adding children to some of the vertices of $\widetilde{\mathbf{T}}^g_\lambda$ with two children. Since $\mu_{\lambda_n}([3, +\infty[) = O\left(\frac{1}{n^2}\right)$, the probability to affect a vertex is of order

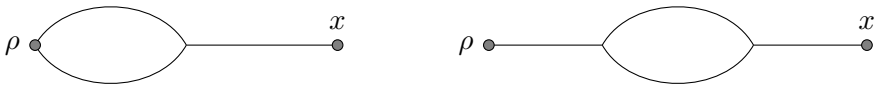

Figure 4.12 – Proposition 4.32 : almost surely, for any x, none of these two cases occurs.

$\frac{1}{n^2}$. The number of vertices at height of order n in $\mathbf{T}^g_{\lambda_n}$ is of order n, so the difference bewteen $\mathbf{T}^g_{\lambda_n}$ and $\widetilde{\mathbf{T}}^g_{\lambda_n}$ does not affect the scaling limit. □

However, taking the scaling limit of $\mathbf{T}^g_{\lambda_n}$ is not enough to obtain a description of infinite geodesics in $\mathcal{P}^h$. Three different kinds of problems could prevent this:

(i) it is not completely clear that the infinite geodesics in $\mathcal{P}^h$ form a tree,

(ii) two different discrete leftmost geodesics might be too close and collide in the scaling limit,

(iii) discrete paths that are not infinite leftmost geodesics might become infinite geodesics in the scaling limit.

We take care of item (i) right now, while the goal of Lemmas 4.33 and 4.35 will be to rule out items (ii) and (iii).

The fact that the infinite geodesics of $\mathcal{P}^h$ indeed form a tree is a quite strong result, that follows from the confluence of geodesics properties in the Brownian map. More precisely, it will be a consequence of [16, Proposition 28], which we recall here. We write m_∞ for the Brownian map, and denote its root by ρ.

Proposition 4.32. [16] Almost surely, for any $x \in m_\infty$, if γ and γ' are two geodesics from ρ to x that coincide on a neighbourhood of x, then $\gamma = \gamma'$.

We write $\mathbf{T}^g(\mathcal{P}^h)$ for the set of those points of $\mathcal{P}^h$ that lie on an infinite geodesic of $\mathcal{P}^h$ started from ρ. By local isometry of the Brownian plane and the Brownian map and scale invariance of the Brownian plane [62], Proposition 4.32 also holds for the Brownian plane. By the absolute continuity relations of Chapter 3, it also holds for $\mathcal{P}^h$. We claim that it implies that for every $x \in \mathbf{T}^g(\mathcal{P}^h)$, there is a unique geodesic from ρ to x in $\mathcal{P}^h$. Indeed, let $x \in \mathbf{T}^g(\mathcal{P}^h)$ and let γ be an infinite geodesic of $\mathcal{P}^h$ passing through x. Let γ' be a geodesic from ρ to x. Finally, let y be a point on the ray γ such that $d(\rho, y) > d(\rho, x)$. The concatenation γ'' of γ' and the part of γ between x and y is a geodesic from ρ to y that coincides with γ in a neighbourhood of y. By Proposition 4.32, the path γ'' must coincide with γ, so γ' must coincide with γ, which proves our claim.

We equip $\mathbf{T}^g(\mathcal{P}^h)$ with its natural tree metric: if $x, y \in \mathbf{T}^g(\mathcal{P}^h)$, the intersection of the geodesics from ρ to x and y is compact so there is a unique point z on it that maximizes $d(\rho, z)$. We write $d_{\mathbf{T}^g(\mathcal{P}^h)}(x, y) = d_{\mathcal{P}^h}(x, z) + d_{\mathcal{P}^h}(z, y)$. Equipped with this distance $\mathbf{T}^g(\mathcal{P}^h)$ is a real tree. However, it is not obvious that it is locally compact (for example it is not if $\mathcal{P}^h$ is replaced by $\mathbb{R}^2$ equipped with the Euclidean norm).

4.5.2 Two lemmas about near-critical strips

We first show that two disjoint geodesics in $\mathbf{T}^g_\lambda$ are quite well-separated, which rules out problem (ii). We denote by γ_ℓ and γ_r the left and right boundaries of the infinite strip $S^1_{\lambda_n}$.

Lemma 4.33. Let $b > a > 0$ and $\varepsilon > 0$. Then there is $\delta > 0$ such that, for n large enough, the following holds:

$$\mathbb{P}\left(\forall i, j \in [an, bn], d_{S^1_{\lambda_n}}(\gamma_\ell(i), \gamma_r(j)) \geq \delta n\right) \geq 1 - \varepsilon.$$

Note that this lemma is very similar to Lemma 3.4 of [99]. However, the slices considered here are not exactly the same, and the lemma of [99] only gives the result with positive probability.

Proof. The idea of the proof is as follows: assume that two points on γ_r and γ_ℓ at height in $[a, b]$ are too close from each other. Then with positive probability S^1_λ is the only strip of $\mathbb{T}_{\lambda_n}$ that reaches height b. In this case, we have a small separating cycle in $\mathbb{T}_{\lambda_n}$, which becomes a pinch point in the scaling limit. This contradicts the homeomorphicity of $\mathcal{P}^h$ to the plane (this is Proposition 18 of Chapter 3, and a consequence of the homeomorphicity of the Brownian map to the sphere [103]).

More precisely, assume the lemma is not true. Then up to extraction, for all $\delta > 0$ and for n large enough, we have

$$\mathbb{P}\left(\exists i, j \in [an, bn], d_{S^1_{\lambda_n}}(\gamma_\ell(i), \gamma_r(j)) < \delta n\right) \geq \varepsilon.$$

Note that if this happens, then by the triangle inequality we must have $|i - j| < \delta n$.

On the other hand, by Theorem 4.1, the probability that the tree $\mathbf{T}^g_{\lambda_n}$ has only one vertex at height $(b+1)n$ is $\mu_{\lambda_n}(1)^{(b+1)n} = m_{\lambda_n}^{(b+1)n} \longrightarrow e^{-2\sqrt{2}(b+1)}$. Since the tree $\mathbf{T}^g_{\lambda_n}$ is independent of the strip $S^1_{\lambda_n}$, for n large enough, the following event occurs with probability at least $\frac{1}{2}e^{-2\sqrt{2}(b+1)}\varepsilon$:

$$\{\text{there are } i, j \in [an, bn] \text{ with } d_{S^1_{\lambda_n}}(\gamma_\ell(i), \gamma_r(j)) < \delta n\}$$
$$\text{AND}$$
$$\{S^1_{\lambda_n} \text{ is the only strip of } \mathbb{T}_{\lambda_n} \text{ at height } (b+1)n\}.$$

If this event occurs, the hull of radius $(b + 1)n$ in $\mathbb{T}_{\lambda_n}$ is the hull of radius $(b + 1)n$ in $S^1_{\lambda_n}$, where the boundary geodesics γ_ℓ and γ_r have been glued together. Hence, by concatenating the geodesic in $S^1_{\lambda_n}$ between $\gamma_\ell(i)$ and $\gamma_r(j)$ and a portion of $\gamma_\ell = \gamma_r$ of length $|i - j|$, we get a cycle of length not greater than $2\delta n$, at height at least $(a - \delta)n$, and that separates ρ from infinity. In other words, for all $\delta > 0$, with probability at least $\frac{1}{2}e^{-2\sqrt{2}(b+1)}\varepsilon$, the following event occurs:

$$\{\text{there is a point } z \in \mathbb{T}_{\lambda_n} \text{ with } a - \delta \leq \tfrac{1}{n}d(\rho, z) \leq b + \delta \text{ such that for any continuous path } p$$
$$\text{from } \rho \text{ to infinity, there is a point } y \text{ of } p \text{ such that } \tfrac{1}{n}d(y, z) \leq 2\delta\}.$$

This last property is closed for the Gromov–Hausdorff topology. To prove it properly, we would need to replace the path to infinity by a path to a point x at distance from ρ large enough so that x cannot lie in $B^\bullet_{(b+1)n}(\mathbb{T}_{\lambda_n})$, and then to replace the continuous path by a δ-chain. We omit the details here, see Appendix of Chapter 3 for something very similar.

Hence, by the main Theorem of Chapter 3, for any $\delta > 0$, with probability at least $\frac{1}{2}e^{-2\sqrt{2}(b+1)}\varepsilon$, there is a point $z \in \mathcal{P}^h$ with $a - \delta \leq d(\rho, z) \leq b + \delta$ such that any continuous path from ρ to infinity contains a point at distance at most δ from z. Since this event is nondecreasing in δ, with probability at least $\frac{1}{2}e^{-2\sqrt{2}(b+1)}\varepsilon$ it occurs for every $\delta > 0$. If it does, let z_n be such a point for $\delta = \frac{1}{n}$, and let z be a subsequential limit of (z_n). Then we have $a \leq d(\rho, z) \leq b$ and every infinite path from ρ to infinity must contain z. Hence, there is a single point that separates the origin from infinity in $\mathcal{P}^h$, which is impossible by homeomorphicity of $\mathcal{P}^h$ to the plane. □

The next lemma shows that for any $x \in \mathbb{T}_{\lambda_n}$, there is a geodesic from x to ρ that coincides with a leftmost infinite geodesic on a quite long distance. Combined with the uniqueness of geodesics between ρ and points of $\mathbf{T}^g(\mathcal{P}^h)$, this will rule out problem (iii). We consider a strip $S^0_{\lambda_n}$ and we denote by γ_ℓ and γ_r its left and right geodesic boundaries.

Definition 4.34. Let $x \in S^0_{\lambda_n}$ and $a > 0$. We say that x is *a-close to the boundary* if there is a geodesic γ from ρ to x that contains either $\gamma_\ell(i)$ for every $0 \leq i \leq a$, or $\gamma_r(i)$ for every $0 \leq i \leq a$.

Lemma 4.35. Let $\varepsilon, r > 0$.

(i) There is $C > 0$ such that for n large enough

$$\mathbb{P}\left(\text{every } x \in \partial B^{\bullet}_{Cn}\left(S^0_{\lambda_n}\right) \text{ is } (rn)\text{-close to the boundary}\right) \geq 1 - \varepsilon.$$

(ii) There is $C' > 0$ such that for n large enough

$$\mathbb{P}\left(\text{every } x \in S^0_{\lambda_n} \setminus B_{C'n}\left(S^0_{\lambda_n}\right) \text{ is } (rn)\text{-close to the boundary}\right) \geq 1 - 2\varepsilon.$$

Note that we gave two quite similar versions of the lemma. The version (i) is the most natural to prove, whereas (ii) passes more easily to the Gromov–Hausdorff limit and is the one we will use later.

Proof. (i) We recall that $\boldsymbol{\tau}^0_{\lambda_n}$ is the skeleton of $S^0_{\lambda_n}$. We first argue that showing point (i) of the lemma is equivalent to bounding the heights of the trees grafted on the spine of $\boldsymbol{\tau}^0$. Let $k > 0$ (we will precise its value later). We denote by $x_1, \dots, x_{p+1}$ the vertices of $\partial B^{\bullet}_k\left(S^0_{\lambda_n}\right)$ that lie on the right of the spine of $\boldsymbol{\tau}^0$, from left to right. For $1 \leq i \leq p$, let also t_i be the subtree of descendants of the edge $\{x_i, x_{i+1}\}$ in $\boldsymbol{\tau}^0$. For $1 \leq i \leq p$, let also γ_i be the leftmost geodesic from x_i to ρ. It is clear (see Figure 4.6) that the distance between x_i and the point at which γ_i and γ_r merge is equal to the maximum of the heights of the trees starting between x_i and γ_r. Hence, we have $\gamma_i(j) \in \gamma_r$ as soon as j is greater than the heights of all the trees $t_i, t_{i+1}, \dots, t_p$. Therefore, if we denote by $H^r_{\lambda_n}(k)$ the height of the forest $(t_1, \dots, t_p)$, then for any i, there is a geodesic from x_i to ρ that contains $\gamma_r(j)$ for $0 \leq j \leq k - H^r_{\lambda_n}(k)$. We can do the same reasoning for vertices on the left of the spine. We denote by $H^{\ell}_{\lambda_n}(k)$ the height of the forest that is defined similarly on the left of the spine, and write $H_{\lambda_n}(k) = \max\left(H^{\ell}_{\lambda_n}(k), H^r_{\lambda_n}(k)\right)$. It is then enough to find C such that for n large enough:

$$\mathbb{P}\left(H_{\lambda_n}(Cn) \leq (C - r)n\right) \geq 1 - \varepsilon. \tag{4.28}$$

We write $P_{\lambda_n}(Cn) = |\partial B^{\bullet}_{Cn}\left(S^0_{\lambda_n}\right)| - 2$ (this is the number of vertices of $\boldsymbol{\tau}^0$ at height Cn that are not on the spine). Conditionally on $P_{\lambda_n}(Cn)$, all the trees of descendants of the vertices x_i are independent Galton–Watson trees conditioned on extinction before a finite time $\lfloor Cn \rfloor$, so we have

$$\mathbb{P}\left(H_{\lambda_n}(Cn) \leq (C - r)n | P_{\lambda_n}(Cn) = p\right) = \left(\frac{g^{\circ \lfloor (C-r)n \rfloor}_{\lambda_n}(0)}{g^{\circ \lfloor Cn \rfloor}_{\lambda_n}(0)}\right)^p \geq g^{\circ \lfloor (C-r)n \rfloor}_{\lambda_n}(0)^p.$$

By Lemma 4.12, we get

$$g^{\circ \lfloor (C-r)n \rfloor}_{\lambda_n}(0) = 1 - \frac{2}{\sinh^2\left(\sqrt{2}(C - r)\right)} \frac{1}{n^2} + o\left(\frac{1}{n^2}\right),$$

so $g^{\circ \lfloor (C-r)n \rfloor}_{\lambda_n}(0) \geq 1 - \frac{z}{n^2}$ for n large enough, where $z = \frac{3}{\sinh^2\left(\sqrt{2}(C-r)\right)}$. Hence, we get

$$\mathbb{P}\left(H_{\lambda_n}(Cn) \leq (C - r)n\right) \geq \mathbb{E}\left[\left(1 - \frac{z}{n^2}\right)^{P_{\lambda_n}(Cn)}\right].$$

By using the distribution of P_{Cn} given by Lemma 4.21, we obtain

$$\begin{aligned}\mathbb{E}\left[\left(1-\frac{z}{n^2}\right)^{P_{\lambda_n}(Cn)}\right] &= \frac{m_{\lambda_n}^{-\lfloor Cn\rfloor}}{\Pi_{\lambda_n}(\theta_{\lambda_n}(0))} \\ &\qquad \times\left(\Pi_{\lambda_n}\left(\left(1-\frac{z}{n^2}\right)g_{\lambda_n}^{\circ(\lfloor Cn\rfloor+1)}(0)\right)-\Pi_{\lambda_n}\left(\left(1-\frac{z}{n^2}\right)g_{\lambda_n}^{\circ\lfloor Cn\rfloor}(0)\right)\right) \\ &\geq \frac{1}{\Pi_{\lambda_n}(\theta_{\lambda_n}(0))}\times m_{\lambda_n}^{-\lfloor Cn\rfloor}\times\left(1-\frac{z}{n^2}\right)\times\left(g_{\lambda_n}^{\circ(\lfloor Cn\rfloor+1)}(0)-g_{\lambda_n}^{\circ\lfloor Cn\rfloor}(0)\right) \\ &\qquad \times\Pi'_{\lambda_n}\left(\left(1-\frac{z}{n^2}\right)g_{\lambda_n}^{\circ\lfloor Cn\rfloor}(0)\right).\end{aligned}$$

by convexity of Π_{λ_n}. We can now compute everything using Lemmas 4.12 and 4.14. As $n\to+\infty$, the first factor goes to $\frac{1}{2}$, the second one goes to $e^{2\sqrt{2}C}$, the third one goes to 1. Moreover, by Lemma 4.12 we have

$$g_{\lambda_n}^{\circ\lfloor Cn\rfloor}(0)=1-\frac{2}{\sinh^2(\sqrt{2}C)}\frac{1}{n^2}+O\left(\frac{1}{n^3}\right)$$

and, if $x_n=1-\frac{c}{n^2}+o\left(\frac{1}{n^2}\right)$, by (4.11) we have

$$g_{\lambda_n}(x_n)-x_n\underset{n\to+\infty}{\sim}\frac{2c\sqrt{c+2}}{n^3},$$

so the fourth factor is equivalent to $4\sqrt{2}\frac{\cosh\left(\sqrt{2}C\right)}{\sinh^3\left(\sqrt{2}C\right)}\frac{1}{n^3}$. Finally, by taking the derivative of Lemma 4.14, we get

$$\Pi'_{\lambda_n}\left(1-\frac{c}{n^2}+o\left(\frac{1}{n^2}\right)\right)=\frac{1}{\sqrt{2}\left(\sqrt{2}+\sqrt{c+2}\right)^2}n^3+o\left(n^3\right).$$

By putting all these estimates together we obtain

$$\mathbb{E}\left[\left(1-\frac{z}{n^2}\right)^{P_{\lambda_n}(Cn)}\right]\xrightarrow[n\to+\infty]{}e^{2\sqrt{2}C}\frac{\cosh\left(\sqrt{2}C\right)}{\sinh^3\left(\sqrt{2}C\right)}\left(1+\sqrt{z+\coth^2\left(\sqrt{2}C\right)}\right)^{-2}.$$

This goes to 1 as $C\to+\infty$ (remember that $z=\frac{3}{\sinh^2\left(\sqrt{2}(C-r)\right)}$ with r fixed), so if C is chosen large enough, then $\mathbb{E}\left[\left(1-\frac{z}{n^2}\right)^{P_{\lambda_n}(Cn)}\right]\geq 1-\varepsilon$ for n large enough. This proves (4.28) and the version (i) of the lemma.

(ii) This is quite easy using version (i). Let C be given by point (i). Note that if any $x\in\partial B^\bullet_{Cn}\left(S^0_{\lambda_n}\right)$ is (rn)-close to the boundary, then so is any $x'\in S^0_{\lambda_n}\backslash B^\bullet_{Cn}\left(S^0_{\lambda_n}\right)$. Indeed, any geodesic γ from x' to ρ must contain a point $x\in\partial B^\bullet_{Cn}\left(S^0_{\lambda_n}\right)$, and we can replace the portion of γ between x and ρ by a geodesic that coincides with γ_ℓ or γ_r between height 0 and rn.

Hence, it is enough to find C' such that with probability $1-\varepsilon$, any point $x'\in S^0_{\lambda_n}$ such that $d(x',\rho)\geq C'n$ is not in $B^\bullet_{Cn}\left(S^0_{\lambda_n}\right)$. In other words, we want to prove that the radius of $\frac{1}{n}B^\bullet_{Cn}\left(S^0_{\lambda_n}\right)$ from ρ is tight as $n\to+\infty$. Since $S^0_{\lambda_n}$ can be embedded in $\mathbb{T}_{\lambda_n}$in a way that preserves the distances from ρ, this is a consequence of the local Gromov–Hausdorff tightness of $\frac{1}{n}\mathbb{T}_{\lambda_n}$.

□

Note that by Lemma 4.22, Lemma 4.33 holds if we replace $S^1_{\lambda_n}$ by $S^0_{\lambda_n}$ and Lemma 4.35 also holds if we replace $S^0_{\lambda_n}$ by $S^1_{\lambda_n}$. We will use these results for both $S^0_{\lambda_n}$ and $S^1_{\lambda_n}$.

4.5.3 Identification of the geodesic tree via Gromov–Hausdorff-closed events

The last two lemmas together with Lemma 4.31 and the fact that $\mathbf{T}^g(\mathcal{P}^h)$ is a tree are basically enough to prove Theorem 4.3. However, to prove it properly, we need to express the distribution of $\mathbf{T}^g(\mathcal{P}^h)$ in terms of closed events for the Gromov–Hausdorff topology, which turns out to be a bit technical.

Let $\mathbf{t}$ be a (finite or infinite) plane tree with a root vertex ρ. If $v \in V(\mathbf{t})\backslash\{\rho\}$, we write p_v for its parent. Let $(h_v)_{v\in V(\mathbf{t})}$ be a family of nonnegative numbers satisfying $h_\rho = 0$ and $h_v > h_{p_v}$ for every $v \in V(\mathbf{t})\backslash\{\rho\}$. We write $\mathbf{t}[h]$ for the metric space obtained from $\mathbf{t}$ by giving, for every $v \in V(\mathbf{t})\backslash\{\rho\}$, a length $h_v - h_{p_v}$ to the edge between p_v and v. We also recall that $\mathbf{B}$ is the infinite tree in which every vertex has two children, except the root which has only one.

We will now define a large family of events, whose probability will characterize the distribution of a random tree of the form $\mathbf{B}[H]$. Let t be a finite binary tree (that is, a tree in which every vertex has 0 or 2 children, except the root which has exactly one). We write $V^*(t)$ for the set of vertices of t that are not leaves and not ρ. Let $r > 0$, and let $(a_v)_{v\in V^*(t)}$, $(b_v)_{v\in V^*(t)}$ be such that $0 < a_v < b_v < r$ for every $v \in V^*(t)$. We write $\mathscr{A}_r^t(a, b)$ for the set of unbounded trees $\mathbf{T}$ (considered as metric spaces) such that $B_r(\mathbf{T})$ is of the form $t[h]$, where $h_\rho = 0$, $h_v = r$ if v is a leaf of t and $a_v \le h_v \le b_v$ for every $v \in V^*(t)$.

In order to prove Theorem 4.3, we will estimate the probability that $\mathbf{T}^g(\mathcal{P}^h)$ belongs to $\mathscr{A}_r^t(a, b)$. Unfortunately, for the reasons listed in of Section 4.1, the events $\{\mathbf{T}^g(X) \in \mathscr{A}_r^t(a, b)\}$ are not closed for the Gromov–Hausdorff distance. To compute $\mathbb{P}\left(\mathbf{T}^g(\mathcal{P}^h) \in \mathscr{A}_r^t(a, b)\right)$ from our discrete estimates, we need to approximate the event $\{\mathbf{T}^g(X) \in \mathscr{A}_r^t(a, b)\}$ by closed events. Since such approximations are tedious to write down explicitly in the general case, we will focus on the case where $t = t_0$ is the binary tree with two leaves and one vertex of degree 3. Note that $|V^*(t_0)| = 1$, so a and b are just two real numbers.

Let $C > r$ and let $R \ge C + 1$. If $\delta, \varepsilon > 0$, we write $\mathscr{A}_{r,C,R}^{\delta,\varepsilon}(a, b)$ for the set of compact metric spaces (X, d) satisfying the following property.

"There are points x_0, x_1 and x_2 in X and geodesics γ_1 (resp. γ_2) from x_0 to x_1 (resp. to x_2) such that:

(i) $d(\rho, x_0) = a$,
(i) $d(\rho, x_1) = d(\rho, x_0) + d(x_0, x_1) = r$ and $d(\rho, x_2) = d(\rho, x_0) + d(x_0, x_2) = r$,
(iii) for every $x \in X$ with $d(\rho, x) > C$, the distance $d(\rho, x)$ is equal to $d(\rho, x_1) + d(x_1, x)$ or to $d(\rho, x_2) + d(x_2, x)$ (it may be equal to both),
(iv) there are two points $y, z \in X$ with $d(\rho, y) \ge R$ and $d(\rho, z) \ge R$ such that $d(\rho, y) = d(\rho, x_1) + d(x_1, y)$ and $d(\rho, z) = d(\rho, x_2) + d(x_2, z)$,
(v) if $u \in \gamma_1$ and $v \in \gamma_2$ with $d(\rho, u) > b + 2\varepsilon$ and $d(\rho, v) > b + 2\varepsilon$, then $d(u, v) \ge \delta$."

We refer to the Appendix for the proof that this event is closed for the pointed Gromov–Hausdorff topology. More precisely, it is easy to check that $\mathscr{A}_{r,C,R}^{\delta,\varepsilon}(a, b)$ is simply generated by geodesics as in Definition 4.39, so by Proposition 4.40 it is closed. By the convergence of $\frac{1}{n}\mathbb{T}_{\lambda_n}$ to $\mathcal{P}^h$, we have

$$\mathbb{P}\left(B_R(\mathcal{P}^h) \in \mathscr{A}_{r,C,R}^{\delta,\varepsilon}(a, b)\right) \ge \limsup_{n\to+\infty} \mathbb{P}\left(\frac{1}{n} B_{Rn}\left(\mathbb{T}_{\lambda_n}\right) \in \mathscr{A}_{r,C,R}^{\delta,\varepsilon}(a, b)\right). \tag{4.29}$$

We now try to estimate the right-hand side. By Lemma 4.31, we have

$$\lim_{n\to+\infty} \mathbb{P}\left(\frac{1}{n}\mathbf{T}_{\lambda_n}^g \in \mathscr{A}_r^{t_0}(a, b)\right) = \mathbb{P}\left(\mathbf{Y}_{2\sqrt{2}} \in \mathscr{A}_r^{t_0}(a, b)\right). \tag{4.30}$$

Note that to deduce (4.30) from Lemma 4.31, we need to show $\mathbb{P}\left(\mathbf{Y}_{2\sqrt{2}} \in \partial\mathscr{A}_r^{t_0}(a, b)\right) = 0$, where $\partial\mathscr{A}_r^{t_0}(a, b)$ is the boundary of $\mathscr{A}_r^{t_0}(a, b)$ in the space of rooted metric trees, equipped with

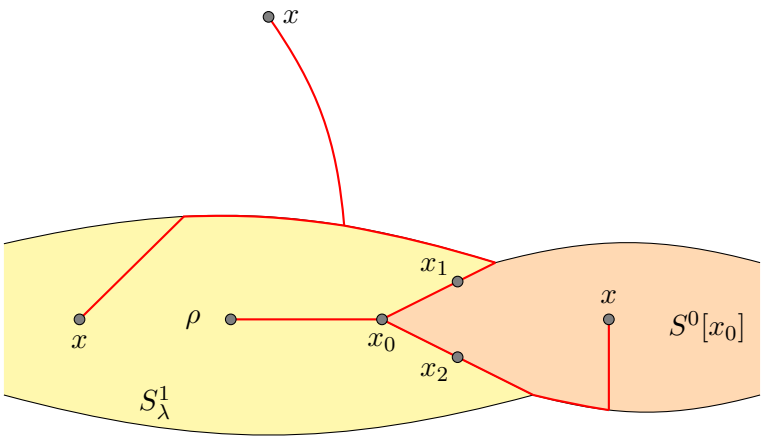

Figure 4.13 – If $x \notin S^1_\lambda \cup S^0[x_0]$, then any geodesic from x to ρ must cross the boundary of one of the two strips S^1_λ and $S^0[x_0]$ above x_1 or x_2. Therefore, for every x with $d(\rho, x) \geq Cn$, we have a geodesic (in red) from x to ρ that passes through x_1 or x_2.

the local Gromov–Hausdorff distance. This is true because if $T \in \partial\mathscr{A}_r^{t_0}(a,b)$, then T must have a branching point at height exactly a, b or r, which a.s. does not happen.

If the event in the left-hand side of (4.30) occurs, let x_0 be the unique point of $\mathbf{T}^g_{\lambda_n}$ at height an, and let x_1 (resp. x_2) be the vertex on the left (resp. right) branch of $\mathbf{T}^g_{\lambda_n}$ at height rn. Then the vertices x_0, x_1 and x_2 and the geodesics γ_1 and γ_2 joining x_1 and x_2 to x_0 in $\mathbf{T}^g_{\lambda_n}$ satisfy assumptions (i) and (ii) in the definition of $\mathscr{A}^{\delta,\varepsilon}_{r,C,R}(a,b)$. Since the tree $\mathbf{T}^g_{\lambda_n}$ is infinite, they also satisfy assumption (iv) for any $R > 0$.

We now fix $\varepsilon > 0$ and apply Lemma 4.35 (version (ii)) to the two strips $S^1_{\lambda_n}$ and $S^0[x_0]$ (the strip whose lowest point is x_0). Lemma 4.35 shows that there is $C > 0$ such that, with probability at least $1 - 4\varepsilon$, for any x in one of the two strips $S^1_{\lambda_n}$ and $S^0(x_0)$ such that $d(x,\rho) > Cn$, there is a geodesic from ρ to x that coincides with γ_1 between ρ and x_1 or with γ_2 between ρ and x_2. We claim that this is also the case if x does not belong to one of these two strips. Indeed, a geodesic from x to ρ must hit the boundary of one of the two strips above x_1 or x_2 (cf. Figure 4.13). Therefore, the probability that $\frac{1}{n}\mathbf{T}^g_{\lambda_n} \in \mathscr{A}_r^{t_0}(a,b)$ but assumption (iii) is not satisfied for C is at most 4ε.

Finally, if $\frac{1}{n}\mathbf{T}^g_{\lambda_n} \in \mathscr{A}_r^{t_0}(a,b)$ but assumption (v) is not satisfied for some $0 < \delta < \varepsilon$, there are two vertices v_1 on γ_1 and v_2 on γ_2, at distance from ρ between $(b+2\varepsilon)n$ and rn, such that $d_{\mathbb{T}_{\lambda_n}}(v_1, v_2) < \delta n$. A geodesic from v_1 to v_2 must cross either $S^1_{\lambda_n}$ or $S^0[x_0]$ and, since $\delta < \varepsilon$, it must stay at distance at least $(b+\varepsilon)n$ from ρ. Therefore, there are two vertices v_1' on γ_1 and v_2' on γ_2, at distance from ρ between $(b+\varepsilon)n$ and $(r+\varepsilon)n$, such that $d_{S^1_{\lambda_n}}(v_1', v_2') < \delta n$ or $d_{S^0[x_0]}(v_1', v_2') < \delta n$. By applying Lemma 4.33 to $S^1_{\lambda_n}$ between heights $b+\varepsilon$ and $r+\varepsilon$, and to $S^0[x_0]$ between heights ε and $r+\varepsilon$, we can find $\delta > 0$ such that this occurs with probability at most 2ε. Hence, for every $\varepsilon > 0$, there is $\delta > 0$ such that the probability that $\frac{1}{n}\mathbf{T}^g_{\lambda_n} \in \mathscr{A}_r^{t_0}(a,b)$ but assumption (v) is not satisfied is at most 2ε.

Therefore, for all $\varepsilon > 0$, there are $C > r$ and $\delta > 0$ such that, for any $R \geq C + 1$,

$$\liminf_{n\to+\infty} \mathbb{P}\left(\frac{1}{n} B_{Rn}\left(\mathbb{T}_{\lambda_n}\right) \in \mathscr{A}^{\delta,\varepsilon}_{r,C,R}(a,b)\right) \geq \mathbb{P}\left(\mathbf{Y}_{2\sqrt{2}} \in \mathscr{A}_r^{t_0}(a,b)\right) - 6\varepsilon,$$

so by (4.29),

$$\mathbb{P}\left(B_R(\mathcal{P}^h) \in \mathscr{A}^{\delta,\varepsilon}_{r,C,R}(a,b)\right) \geq \mathbb{P}\left(\mathbf{Y}_{2\sqrt{2}} \in \mathscr{A}_r^{t_0}(a,b)\right) - 6\varepsilon.$$

Since the event $\left\{B_R(X) \in \mathscr{A}^{\delta,\varepsilon}_{r,C,R}(a,b)\right\}$ is nonincreasing in R, we obtain that for any $\varepsilon > 0$, we have

$$\mathbb{P}\left(\exists C > r, \exists \delta > 0, \forall R \geq C+1, B_R(\mathcal{P}^h) \in \mathscr{A}^{\delta,\varepsilon}_{r,C,R}(a,b)\right) \geq \mathbb{P}\left(\mathbf{Y}_{2\sqrt{2}} \in \mathscr{A}^{t_0}_r(a,b)\right) - 6\varepsilon.$$

Finally, the event above is increasing in ε so

$$\mathbb{P}\left(\forall \varepsilon > 0, \exists \delta > 0, \exists C > 0, \forall R \geq C+1, B_R(\mathcal{P}^h) \in \mathscr{A}^{\delta,\varepsilon}_{r,C,R}(a,b)\right) \geq \mathbb{P}\left(\mathbf{Y}_{2\sqrt{2}} \in \mathscr{A}^{t_0}_r(a,b)\right). \tag{4.31}$$

Lemma 4.36. Almost surely, if

$$\forall \varepsilon > 0, \exists \delta > 0, \exists C > 0, \forall R \geq C+1, B_R(\mathcal{P}^h) \in \mathscr{A}^{\delta,\varepsilon}_{r,C,R}(a,b),$$

then $\mathbf{T}^g(\mathcal{P}^h) \in \mathscr{A}^{t_0}_r(a,b)$.

Proof. Fix C, δ and ε and assume that $B_R(\mathcal{P}^h) \in \mathscr{A}^{\delta,\varepsilon}_{r,C,R}(a,b)$ for any $R \geq C+1$. Let x_0, x_1 and x_2 be given by the definition of $\mathscr{A}^{\delta,\varepsilon}_{r,C,R}(a,b)$. We first check that these points do not depend on the parameters δ, ε, C and R. By assumption (iv), the points x_1 and x_2 lie on geodesics of length $C+1$ started from ρ. We claim they are the only two points at distance r from ρ satisfying this property. Indeed, if y is a point with $d(\rho,y) = C+1$ and γ a geodesic from ρ to y, let z be the point of γ such that $d(\rho,z) = C$. By assumption (iii) and the fact that $\mathcal{P}^h$ is a length space, there is a geodesic γ' from z to ρ passing through x_1 or x_2. By concatenating γ' from ρ to z and γ from z to y, we obtain a geodesic from ρ to y that coincides with γ between z and y. By Proposition 4.32, this geodesic must be equal to γ, so γ must pass through x_1 or x_2.

Hence, if $B_R(\mathcal{P}^h) \in \mathscr{A}^{\delta,\varepsilon}_{r,C,R}(a,b)$, then x_1 and x_2 are the only two points at distance r from ρ that lie on a geodesic of length $C+1$ started from ρ. In particular, they do not depend on δ, ε and R. Moreover, let $C' \geq C$, and assume that $B_R(\mathcal{P}^h) \in \mathscr{A}^{\delta,\varepsilon}_{r,C,R}(a,b)$ and $B_R(\mathcal{P}^h) \in \mathscr{A}^{\delta,\varepsilon}_{r,C',R}(a,b)$ for all $R \geq C'$. Then the two points at distance r from ρ that lie on a geodesic of length $C'+1$ from ρ are the same as the two points that lie on a geodesic of length $C+1$ started from ρ, so the points x_1 and x_2 do not depend on C. Similarly, the point x_0 is the only point at distance a from the root that lies on a geodesic of length $C+1$, so it does not depend on $\delta, \varepsilon, C, R$. Hence, we can find x_0, x_1 and x_2 in $\mathcal{P}^h$ and γ_1, γ_2 such that:

— assumptions (i) and (ii) in the definition of $\mathscr{A}^{\delta,\varepsilon}_{r,C,R}(a,b)$ are satisfied,
— there is $C > 0$ such that assumption (iii) is satisfied,
— for every $R \geq C+1$, assumption (iv) is satisfied,
— for every $\varepsilon > 0$, there is $\delta > 0$ such that assumption (v) is satisfied, which means that the geodesics from ρ to x_1 and x_2 are disjoint between heights b (excluded) and r (included).

Since assumption (iv) is satisfied for any R large enough, there are arbitrarily large geodesics started from ρ passing through x_1. By a compactness argument, there are infinite geodesics started from ρ and passing through x_1, and the same is true for x_2. By assumption (iii), the points x_1 and x_2 are the only ones with this property. By assumptions (i) and (v), the branching point between the geodesics from ρ to x_1 and x_2 lies between heights a and b, so $\mathbf{T}^g(\mathcal{P}^h) \in \mathscr{A}^{t_0}_r(a,b)$. □

The end of the proof of Theorem 4.3 is now easy. By Lemma 4.36 and (4.31), we get

$$\mathbb{P}\left(\mathbf{T}^g(\mathcal{P}^h) \in \mathscr{A}^{t_0}_r(a,b)\right) \geq \mathbb{P}\left(\mathbf{Y}_{2\sqrt{2}} \in \mathscr{A}^{t_0}_r(a,b)\right).$$

The general case for t can be treated along the same lines. This shows that the distribution of $\mathbf{T}^g(\mathcal{P}^h)$ dominates that of $\mathbf{Y}_{2\sqrt{2}}$. Since they are both probability measures, they are the same.

4.A Appendix: A Gromov–Hausdorff closedness result

The goal of this appendix is to prove Proposition 4.40. It shows that a wide class of events related to geodesics are closed for the Gromov–Hausdorff distance. We believe it might be of interest in other settings. We write $\mathscr{G}$ for the space of pointed compact metric spaces, equipped with the Gromov–Hausdorff distance. We will be interested in some events depending on a metric space $(X, d, \rho) \in \mathscr{G}$.

Definition 4.37. We say that a subset $\mathscr{A}$ of $\mathscr{G}$ is *simply generated by points* if it has the following form. Let $k \geq 1$, and let $F \subset \mathbb{R}^{(k+2)^2}$ be closed. Then $\mathscr{A}$ is the set of those $(X, d, \rho) \in \mathscr{G}$ for which there are $(x_i)_{0 \leq i \leq k}$ in X with $x_0 = \rho$ such that, for any $x_{k+1} \in X$, the matrix

$$(d(x_i, x_j))_{0 \leq i,j \leq k+1}$$

lies in F.

Lemma 4.38. Any subset of $\mathscr{G}$ that is simply generated by points is closed.

Proof. Assume that $\mathscr{A}$ is simply generated by points and let k and F be as above. Let (X_n, d_n, ρ_n) converge to a space (X, d, ρ) with $X_n \in \mathscr{A}$ for every n. By Gromov–Hausdorff convergence, we can embed X and all the X_n isometrically in a space (Z, d_Z) such that the Hausdorff distance D_n between X_n and X goes to 0.

For every n, let $x_0^n, \dots, x_k^n \in X_n$ satisfy the condition given by Definition 4.37. We take $y_0^n, \dots, y_k^n \in X$ such that $d_Z(x_i^n, y_i^n) \leq 2D_n$. For all $0 \leq i \leq k$, let y_i be a subsequential limit of $(y_i^n)_{n \geq 0}$ in X (which exists by compactness). To complete the proof that $X \in \mathscr{A}$, all we need to show is that $y_0 = \rho$ and that for any $y_{k+1} \in X$, we have

$$(d(y_i, y_j))_{0 \leq i,j \leq k+1} \in F.$$

The first point is easy because the distances $d_Z(\rho, \rho_n)$, $d_Z(\rho_n, y_0^n)$ and $d(y_0^n, y_0)$ all go to 0 along some subsequence. Moreover, let $y_{k+1} \in X$. There is $x_{k+1}^n \in X_n$ such that $d_Z(x_{k+1}^n, y_{k+1}) \leq 2D_n$. For every $0 \leq i, j \leq k+1$, we then have

$$d(y_i, y_j) = \lim_{n \to +\infty} d(x_i^n, x_j^n)$$

along some subsequence. But we know that for all n we have $\left(d(x_i^n, x_j^n)\right)_{0 \leq i,j \leq k+1} \in F$, so we can conclude. □

Definition 4.39. We say that a subset $\mathscr{A}$ of $\mathscr{G}$ is *simply generated by geodesics* if it has the following form. Let $k \geq 1$, and let $F \subset \mathbb{R}^{(2k+2)^2}$ be closed. Then $\mathscr{A}$ is the set of those $(X, d, \rho) \in \mathscr{G}$ for which there are $(x_i)_{0 \leq i \leq k}$ in X with $x_0 = \rho$ and geodesics $(\gamma_i)_{1 \leq i \leq k}$ from ρ to x_i, satisfying the following property. For any $(x_{k+i})_{1 \leq i \leq k}$ such that $x_{k+i} \in \gamma_i$ for every i, and for every $x_{2k+1} \in X$, the matrix

$$(d(x_i, x_j))_{0 \leq i,j \leq 2k+1}$$

lies in F.

Proposition 4.40. Any subset of $\mathscr{G}$ that is simply generated by geodesics is closed.

To go from Lemma 4.38 to Proposition 4.40, we will need the following definition.

Definition 4.41. Let $\ell \geq 0$, and let x, y be two points of a metric space (X, d). An *ℓ-geodesic chain from x to y* is a finite sequence $(x(i))_{0 \leq i \leq 2^\ell}$ of points of X such that

(i) $x(0) = x$ and $x(2^\ell) = y$,

(ii) $d(x(i), x(i+1)) = \frac{1}{2^\ell} d(x, y)$ for any $0 \leq i \leq 2^\ell - 1$.

Proof of Proposition 4.40. Let $\mathscr{A}$ be a subset of $\mathscr{G}$ that is simply generated by geodesics. For $\ell > 0$, we write $\mathscr{A}^\ell$ for the subset of $\mathscr{G}$ we obtain if we replace continuous geodesics by ℓ-geodesic chains in Definition 4.39. Then $\mathscr{A}^\ell$ is simply generated by points (because the conditions in the definition of an ℓ-geodesic chain are closed), so $\mathscr{A}^\ell$ is closed by Lemma 4.38. Hence, to conclude, it is enough to show

$$\mathscr{A} = \bigcap_{\ell \geq 0} \mathscr{A}^\ell.$$

The inclusion from left to right is immediate since any continuous geodesic contains an ℓ-geodesic chain. Now let $(X, d, \rho) \in \bigcap_{\ell \geq 0} \mathscr{A}^\ell$. For every $\ell \geq 0$ and $1 \leq i \leq k$, let (x_i^ℓ) in X and let $\left(\gamma_i^\ell(j)\right)_{0 \leq j \leq 2^\ell}$ be ℓ-geodesic chains from ρ to x_i^ℓ satisfying the assumptions of definition 4.39. Up to extraction, we may assume that for every $1 \leq i \leq k$, the points x_i^ℓ converge to a point $x_i \in X$. Up to further extraction, by a diagonal argument, for every t of the form $\frac{j}{2^m}$ with $0 \leq j \leq 2^m$, the sequence $\left(\gamma_i^{m+\ell}(2^\ell j)\right)_{\ell \geq 0}$ converges to a point $\gamma_i(t)$. Moreover, for all such t, t', we have $d(\gamma_i(t), \gamma_i(t')) = |t - t'| d(\rho, x_i)$, so we can extend $(\gamma_i(t))_{0 \leq t \leq 1, t = j/2^m}$ to a continuous geodesic from ρ to x_i. It is then easy to check that the geodesics γ_i satisfy the required hypothesis. $\square$

Chapter 5

Supercritical causal maps

ou *Même sans tentacules !*

This chapter is adapted from the preprint [47].

We study the random planar maps obtained from supercritical Galton–Watson trees by adding the horizontal connections between successive vertices at each level. These are the hyperbolic analog of the maps studied by Curien, Hutchcroft and Nachmias in [60], and a natural model of random hyperbolic geometry. We first establish metric hyperbolicity properties of these maps: we show that they admit bi-infinite geodesics and satisfy a weak version of Gromov-hyperbolicity. We also study the simple random walk on these maps: we identify their Poisson boundary and, in the case where the underlying tree has no leaf, we prove that the random walk has positive speed. Some of the methods used here are robust, and allow us to obtain more general results about planar maps containing a supercritical Galton–Watson tree.

Contents

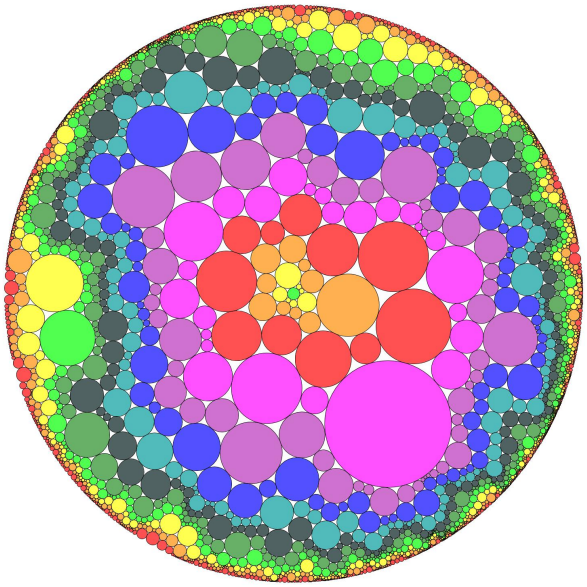

Figure 5.1 – The circle packing of a causal triangulation constructed from a Galton–Watson tree with geometric offspring distribution of mean 3/2. This was made with the help of the software CirclePack by Ken Stephenson.

Introduction

Causal maps and random hyperbolic geometry. *Causal triangulations* were introduced by theoretical physicists Ambjørn and Loll [10], and have been the object of a lot of numerical investigations. However, their rigorous study is quite recent [75, 60]. They are a discrete model of Lorentzian quantum gravity with one time and one space dimension where, in contrast with uniform random planar maps, time and space play asymetric roles.

Here is the definition of the model. For any (finite or infinite) plane tree t, we denote by $\mathcal{C}(t)$ the planar map obtained from t by adding at each level the horizontal connections between consecutive vertices, as on Figure 5.2 (this includes an horizontal connection between the leftmost and rigtmost vertices at each level). Our goal here is to study the graph $\mathcal{C}(T)$, where T is a supercritical Galton–Watson tree conditioned to survive.

This defines a new model of random "hyperbolic" graph. Several other such models have been investigated so far, such as supercritical Galton–Watson trees [108], Poisson–Voronoi tesselations of the hyperbolic plane [29], or the Planar Stochastic Hyperbolic Infinite Triangulations (PSHIT) of [58]. Many notions appearing in the study of these models are adapted from the study of Cayley graphs of nonamenable groups, and an important idea is to find more general versions of the useful properties of these Cayley graphs. Let us mention two such tools.

- For example, the three aforementioned models are all *stationary*, which means their distribution is invariant under rerooting along the simple random walk [1]. This property generalizes the transitivity of Cayley graphs, and is a key tool to prove positive speed for the simple random walk on supercritical Galton–Watson trees [108, 6] or on the PSHIT [58]. More generally, in the context of stationary random graphs, general relations are known between the exponential growth rate, the speed of the random walk, and its asymp-

1. This is not exactly true for Galton–Watson trees, but it is true for the closely related *augmented Galton–Watson trees*.

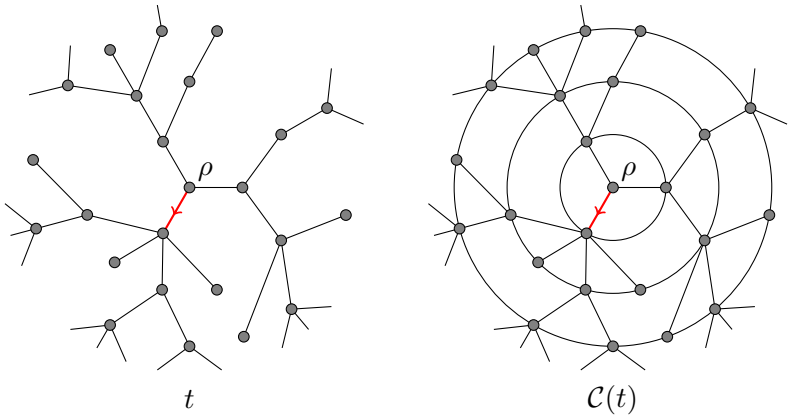

Figure 5.2 – An infinite plane tree t and the associated causal map $\mathcal{C}(t)$. The edge in red joins the root vertex to its leftmost child.

totic entropy, which is itself related to the Poisson boundary and the Liouville property. See [27, Proposition 3.6], which adapts classical results about Cayley graphs [85]. On the other hand, supercritical causal maps are not stationary, and it seems hard to find a stationary environment for the simple random walk. In absence of stationarity, we will be forced to use other properties of our graphs such as the independence properties given by the structure of Galton–Watson trees.

- Another important property in the study of random hyperbolic graphs is the anchored expansion, which is a weaker version of nonamenability, and its natural generalization to random graphs. It is known to imply positive speed and heat kernel decay bounds of the form $\exp(-n^{1/3})$ for bounded-degree graphs [137]. This property also played an important role in the study of non-bounded-degree graphs such as Poisson-Voronoi tesselations of the hyperbolic plane [29], and the half-planar versions of the PSHIT [17]. However, we have not been able to establish this property for causal maps, and need once again to use other methods.

Supercritical causal maps. In all that follows, we fix an offspring distribution μ with $\sum_{i=0}^{\infty} i\mu(i) > 1$. We denote by T a Galton–Watson tree with offspring distribution μ conditioned to survive. The goal of this work is to study the maps $\mathcal{C}(T)$. We will study both large-scale metric properties of $\mathcal{C}(T)$, and the simple random walk on this map. All the results that we will prove show that $\mathcal{C}(T)$ has a hyperbolic flavour, which is also true for the tree T.

Metric hyperbolicity properties. The first goal of this work is to establish two metric hyperbolicity properties of $\mathcal{C}(T)$. We recall that a graph G is hyperbolic in the sense of Gromov if there is a constant $k \geq 0$ such that all the triangles are k-thin in the following sense. Let x, y and z be three vertices of G and γ_{xy}, γ_{yz}, γ_{zx} be geodesics from x to y, from y to z and from z to x. Then for any vertex v on γ_{xy}, the graph distance between v and $\gamma_{yz} \cup \gamma_{zx}$ is at most k. However, such a strong, uniform statement usually cannot hold for random graphs. For example, if $\mu(1) > 0$, then $\mathcal{C}(T)$ contains arbitrarily large portions of the square lattice, which is not hyperbolic. Therefore, we suggest a weaker, "anchored" definition [2].

2. The most natural definition would be to require that any geodesic triangle surrounding the root is k-thin, but this is still too strong (consider the triangle formed by root vertex and two vertices x, y in a large portion of square lattice).

Definition 5.1. Let M be a rooted planar map. We say that M is *weakly anchored hyperbolic* if there is a constant $k \geq 0$ such that the following holds. Let x, y and z be three vertices of M and γ_{xy} (resp. γ_{yz}, γ_{zx}) be a geodesic from x to y (resp. y to z, z to x). Assume the triangle formed by γ_{xy}, γ_{yz} and γ_{zx} surrounds the root vertex ρ. Then

$$d_M\left(\rho, \gamma_{xy} \cup \gamma_{yz} \cup \gamma_{zx}\right) \leq k.$$

Theorem 5.1 (Metric hyperbolicity of $\mathcal{C}(T)$)**.** Let T be a supercritical Galton–Watson tree conditioned to survive, and let $\mathcal{C}(T)$ be the associated causal map.

1. The map $\mathcal{C}(T)$ is a.s. weakly anchored hyperbolic.
2. The map $\mathcal{C}(T)$ a.s. admits bi-infinite geodesics, i.e. paths $(\gamma(i))_{i\in\mathbb{Z}}$ such that for any i and j, the graph distance between $\gamma(i)$ and $\gamma(j)$ is exactly $|i-j|$.

These two results are very robust and hold in a much more general setting that includes the PSHIT. In particular, the second point for the PSHIT answers a question of Benjamini and Tessera [32]. More general results are discussed in the end of this introduction.

Poisson boundary. The second goal of this work is to study the simple random walk on $\mathcal{C}(T)$ and to identify its Poisson boundary. First note that $\mathcal{C}(T)$ contains as a subgraph the supercritical Galton–Watson tree T, which is transient, so $\mathcal{C}(T)$ is transient as well. We recall the general definition of the Poisson boundary. Let G be an infinite, locally finite graph, and let $G \cup \partial G$ be a compactification of G, i.e. a compact metric space in which G is dense. Let also (X_n) be the simple random walk on G started from ρ. We say that ∂G is a *realization of the Poisson boundary* of G if the following two properties hold:

— (X_n) converges a.s. to a point $X_\infty \in \partial G$,

— every bounded harmonic function h on G can be written in the form

$$h(x) = \mathbb{E}_x\left[g\left(X_\infty\right)\right],$$

where g is a bounded measurable function from ∂G to $\mathbb{R}$.

We denote by ∂T the space of infinite rays of T. If $\gamma, \gamma' \in \partial T$, we write $\gamma \sim \gamma'$ if $\gamma = \gamma'$ or if γ and γ' are two "consecutive" rays in the sense that there is no ray between them. Then $\sim$ is a.s. an equivalence relation for which countably many equivalence classes have cardinal 2 and all the others have cardinal 1. We write $\widehat{\partial} T = \partial T / \sim$. There is a natural way to equip $\mathcal{C}(T) \cup \widehat{\partial} T$ with a topology that makes it a compact space. We refer to Section 5.3.1 for the construction of this topology, but we mention right now that $\widehat{\partial} T$ is homeomorphic to the circle, whereas ∂T is homeomorphic to a Cantor set. The space $\mathcal{C}(T) \cup \widehat{\partial} T$ can be seen as a compactification of the infinite graph $\mathcal{C}(T)$. We show that this is a realization of its Poisson boundary.

Theorem 5.2 (Poisson boundary of $\mathcal{C}(T)$)**.** Almost surely:

1. the limit $\lim(X_n) = X_\infty$ exists and its distribution has full support and no atoms in $\widehat{\partial} T$,
2. $\widehat{\partial} T$ is a realization of the Poisson boundary of $\mathcal{C}(T)$.

Note that, by a result of Hutchcroft and Peres [83], the second point will follow from the first one.

Positive speed. A natural and strong property shared by many models of hyperbolic graphs is the positive speed of the simple random walk. See for example [108] for supercritical Galton–Watson trees, and [58, 17] for the PSHIT or their half-planar analogs. The third goal of this work is to prove that the simple random walk on $\mathcal{C}(T)$ has a.s. positive speed. Unfortunately,

we have only been able to prove it in the case where $\mu(0) = 0$, i.e. when the tree T has no leaf. We recall that (X_n) is the simple random walk on $\mathcal{C}(T)$, and denote by $d_{\mathcal{C}(T)}$ the graph distance on $\mathcal{C}(T)$.

Theorem 5.3 (Positive speed on $\mathcal{C}(T)$)**.** If $\mu(0) = 0$ and $\mu(1) < 1$, then there is $v_\mu > 0$ such that

$$\frac{d_{\mathcal{C}(T)}(\rho, X_n)}{n} \xrightarrow[n\to+\infty]{\text{a.s.}} v_\mu.$$

However, we expect to still have positive speed if $\mu(0) = 0$. As mentioned above, this result is not obvious because of the lack of stationarity (for stationary graphs, the results of [27] show that positive speed is equivalent to being non-Liouville under some mild assumptions).

The critical case. We note that similar properties have been studied in the critical case in [60]. The results of [60] show that the geometric properties of causal maps are closer to those of uniform random maps than to those of the trees from which they were built. This contrasts sharply with the supercritical case, where the properties of the causal map are very close to those of the associated tree. More precisely, in the finite variance case, the distance between vertices at some fixed height r is $o(r)$, but $r^{1-o(1)}$. Moreover, the exponents describing the behaviour of the simple random walk are the same as for the square lattice, and different from the exponents we would obtain in a tree.

Robustness of the results and applications to other models. Another motivation to study causal maps is that many other models of random planar maps can be obtained by adding connections (and, in some cases, vertices) to a random tree. For example, the UIPT [20] or its hyperbolic variants the PSHIT [58] can be constructed from a reverse Galton–Watson tree or forest via the Krikun decomposition [94, 65, 46]. Among all the maps that can be obtained from a tree t without to add vertical "shortcuts", the causal map is the one with the "closest" connections, which makes it a useful toy model. The causal map may even provide general bounds for any map obtained from a fixed tree (we will see such applications in this paper, see also [59] for applications to uniform planar maps via the Krikun decomposition).

Here, the causal maps $\mathcal{C}(T)$ fit in a more general framework. We define a *strip* as an infinite, one-ended planar map s with exactly one infinite face, such that the infinite face has a simple boundary ∂s, and equipped with a root vertex on the boundary on the infinite face. If t is an infinite tree with no leaf and $(s_i)_{i\in\mathbb{N}}$ is a sequence of strips, let $\mathcal{M}(t, (s_i))$ be the map obtained by filling the (infinite) faces of t with the strips s_i (see Section 5.1 for a more careful construction). Some of our results can be generalized to random maps of the form $\mathcal{M}(\mathbf{T}, (s_i))$, where $\mathbf{T}$ is a supercritical Galton–Watson tree with no leaf, and the s_i are (random or not) strips.

By the backbone decomposition for supercritical Galton–Watson trees (that we recall in Section 5.1), the maps of the form $\mathcal{C}(T)$, where T is a supercritical Galton–Watson tree with leaves, are a particular case of this construction. The results of Chapter 4 prove that the PSHIT $\mathbb{T}_\lambda$ can also be obtained by this construction: the tree $\mathbf{T}$ is then the tree of infinite leftmost geodesics of $\mathbb{T}_\lambda$ and has geometric offspring distribution.

We will show that Theorem 5.1 is very robust and applies to this general context, see Theorem 5.1 bis. A particular case of interest are the PSHIT. In particular, point 2 of Theorem 5.1 for the PSHIT answers a question of Benjamini and Tessera [32].

As for causal maps, any map of the form $\mathcal{M}(\mathbf{T}, (s_i))$ contains the transient graph $\mathbf{T}$, so it is transient itself. Most of our proof of Theorem 5.2 can also be adapted to the general setting where the strips s_i are i.i.d. and independent of $\mathbf{T}$. However, Theorem 5.2 cannot be true if the strips s_i are too large (for example if themselves have a non-trivial Poisson boundary). On the

	non-Liouville	$\widehat{\partial}\mathbf{T}$ is the Poisson boundary	positive speed
T	✓	✓	✓
$\mathcal{C}(T)$ (if $\mu(0) = 0$)	✓	✓	✓
$\mathcal{C}(T)$ (if $\mu(0) > 0$)	✓	✓	?
PSHIT	✓	✓	✓
$\mathcal{M}$ with (S_i) i.i.d., recurrent	✓	?	✗
$\mathcal{M}$ with (S_i) i.i.d.	✓	✗	✗
general $\mathcal{M}$	✗	✗	✗

Figure 5.3 – The symbol ✓ means that the property is proved in an earlier work or in this one. The symbol ? indicates properties that we believe to be true but did not prove in this paper, and the symbol ✗ means the property is false in general. See Section 5.5 for a quick description of some counterexamples.

other hand, we can still show that the Poisson boundary is non-trivial. See Theorem 5.2 bis for a precise statement, and Figure 5.3 for a summary of the results proved in this paper and the results left to prove.

As we will see later, Theorem 5.2 bis is not strictly speaking more general than Theorem 5.2, since the strips used to construct $\mathcal{C}(T)$ from the backbone of T are not completely independent. On the other hand, once again, the PSHIT satisfy these assumptions (up to a root transformation, since the strip containing the root has a slightly different distribution). However, it was already known that the PSHIT are non-Liouville (see [58], or [14] for another identification of the Poisson boundary via circle packings). We also prove in Chapter 4, by a specific argument based on the peeling process, that $\widehat{\partial}\mathbf{T}$ is indeed a realization of the Poisson boundary in the case of the PSHIT.

Structure of the paper. The paper is structured as follows. In Section 5.1, we fix some definitions and notations that will be used in all the rest of this work, and recall the backbone decomposition of supercritical Galton–Watson trees. In Section 5.2, we investigate metric properties and establish Theorem 5.1 bis, of which Theorem 5.1 is a particular case. Section 5.3 is devoted to the study of the Poisson boundary and to the proof of Theorems 5.2 and 5.2 bis. In Section 5.4, we prove Theorem 5.3 about positive speed. Finally, in Section 5.5, we discuss some counterexamples related to Figure 5.3, and state a few conjectures.

Acknowledgments: I thank Nicolas Curien for his comments on earlier versions of this work, Arvind Singh for his explanations about renewal theory, and Itai Benjamini for providing the reference [26]. I acknowledge the support of ANR Liouville (ANR-15-CE40-0013), ANR GRAAL (ANR-14-CE25-0014) and ERC GeoBrown (740943).

5.1 General framework and the backbone decomposition

The goal of this first section is to give definitions and notations, and to make a few useful remarks that will be needed in all the paper. All our constructions will be based on infinite, locally finite plane trees. We insist that the plane tree structure is important to define the associated causal map. We will use normal letters to denote general infinite trees, and bold letters like $\mathbf{t}$ for trees with no leaf. All the trees will be rooted at a vertex ρ. If v is a vertex of a tree t, we denote by $h(v)$ its distance to the root, which we will sometimes call its *height*. A *ray*

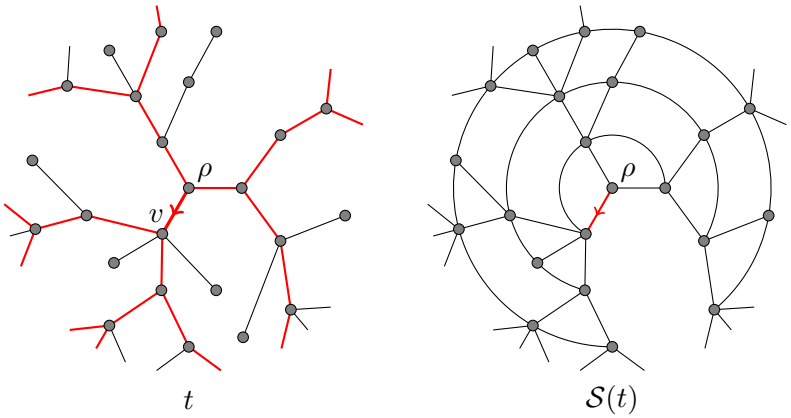

Figure 5.4 – The same infinite tree t as on Figure 5.2 and the associated causal slice $\mathcal{S}(t)$. Note that two vertices have been deleted. On the left part, the backbone of t is in red. We have $c(v) = 4$, $c_{\mathbf{B}}(v) = 2$ and $A_v = \{2, 4\}$.

in a tree t is an infinite sequence $(\gamma(i))_{i\geq 0}$ of vertices such that $\gamma(0) = \rho$, and $\gamma(i+1)$ is a child of $\gamma(i)$ for every $i \geq 0$. If t is an infinite tree, the *backbone* of t is the union of its rays, i.e. the set of the vertices of t that have infinitely many descendants. We will denote it by $\mathbf{B}(t)$, and we note that $\mathbf{B}(t)$ is always an infinite tree with no leaf.

We recall that if t is an infinite plane tree, then $\mathcal{C}(t)$ is the map obtained from t by adding horizontal edges at every height between consecutive vertices. We also define the *causal slice* $\mathcal{S}(t)$ associated to t, which will be used a lot in all that follows. Let γ_ℓ (resp. γ_r) be the leftmost (resp. rightmost) infinite ray of $\mathbf{B}(t)$. Then $\mathcal{S}(t)$ is the map obtained from t by deleting all the vertices on the left of γ_ℓ and on the right of γ_r, and by adding the same horizontal edges as for $\mathcal{C}(t)$ between the remaining vertices, except the edge between γ_ℓ and γ_r at each level (cf. Figure 5.4). The union of γ_ℓ and γ_r is the *boundary* of $\mathcal{S}$, and is written $\partial\mathcal{S}$.

In all this work, μ will denote a supercritical offspring distribution, i.e. satisfying $\sum_{i\geq 0} i\mu(i)$, and T will be a Galton–Watson tree with offspring distribution μ conditioned to survive. For every $n \geq 0$, we will denote by Z_n the number of vertices of T at height n. We will write $\mathcal{C}$ for $\mathcal{C}(T)$ and $\mathcal{S}$ for $\mathcal{S}(T)$, unless stated otherwise.

If v is a vertex of $\mathbf{B}(T)$, we will denote by $T[v]$ the tree of descendants of v in T, and by $\mathcal{S}[v]$ the causal slice associated to $T[v]$. An important consequence of the backbone decomposition stated below is that for each $v \in T$, conditionally on $v \in \mathbf{B}(T)$, the slice $\mathcal{S}[v]$ has the same distribution as $\mathcal{S}$. Moreover, these slices are independent for base points that are not ancestors of each other.

If G is a graph rooted at a vertex ρ, we will denote by d_G its graph distance and by $B_r(G)$ (resp. $\partial B_r(G)$) the set of vertices of G at distance at most r (resp. exactly r) from ρ. Note that the vertices of $B_r(T)$ and of $B_r(\mathcal{C})$ are the same. For any vertex v of T, we also denote by $c_T(v)$ (or by $c(v)$ when there is no ambiguity) the number of children of v in T. Note that the degree of v in $\mathcal{C}$ is equal to $c_T(v)+3$ if $v \neq \rho$, and to $c_T(v)$ if $v = \rho$. For every graph G and every vertex v of G, we will denote by $P_{G,v}$ the distribution of the simple random walk on G started from v.

We now recall the backbone decomposition for supercritical Galton–Watson trees conditioned to survive, as it appears e.g. in [107]. Let f be the generating function of μ, i.e. $f(x) = \sum_{i\geq 0} \mu(i)x^i$. Let also q be the extinction probability of a Galton–Watson tree with offspring

distribution μ, i.e. the smallest fixed point of f in $[0,1]$. We define $\boldsymbol{f}$ and $\widetilde{f}$ by

$$\boldsymbol{f}(s) = \frac{f(q+(1-q)s)-q}{1-q} \quad \text{and} \quad \widetilde{f}(s) = \frac{f(qs)}{q} \tag{5.1}$$

for $s \in [0,1]$. Then $\boldsymbol{f}$ is the generating function of a supercritical offspring distribution $\boldsymbol{\mu}$ with $\boldsymbol{\mu}(0) = 0$, and $\widetilde{f}$ is the generating function of a subcritical offspring distribution $\widetilde{\mu}$. For every vertex x of $\mathbf{B}(T)$, we denote by A_x the set of indices $1 \leq i \leq c(x)$ such that the i-th child of x is in the backbone, and we write $c_{\mathbf{B}}(x) = |A_x|$. Finally, let $\mathbf{B}'(T)$ be the set of vertices y of T such that the parent of y is in $\mathbf{B}(T)$, but y is not. The following result characterizes entirely the distribution of T.

Theorem 5.4. 1. The tree $\mathbf{B}(T)$ is a Galton–Watson tree with offspring distribution $\boldsymbol{\mu}$.

2. Conditionally on $\mathbf{B}(T)$, the variables $c(x) - c_{\mathbf{B}}(x)$ are independent, with distribution characterized by
$$\mathbb{E}\left[s^{c(x)-c_{\mathbf{B}}(x)} \middle| \mathbf{B}(T)\right] = \frac{f^{(c_{\mathbf{B}}(x))}(qs)}{f^{(c_{\mathbf{B}}(x))}(q)}$$
for every $s \in [0,1]$, where $f^{(k)}$ stands for the k-th derivative of f.
3. Conditionally on $\mathbf{B}(T)$ and the variables $c(x)$ for $x \in \mathbf{B}(T)$, the sets A_x for $x \in \mathbf{B}(T)$ are independent and, for every x, the set A_x is uniformly distributed among all the subsets of $\{1, 2, \ldots, c(x)\}$ with $c_{\mathbf{B}}(x)$ elements.
4. Conditionally on everything above, the trees $T[y]$ for $y \in \mathbf{B}'(T)$ are independent Galton–Watson trees with offspring distribution $\widetilde{\mu}$.

In particular, this decomposition implies that, for every $h \geq 1$, conditionally on $B_h(T)$ and on the set $\partial B_h(T) \cap \mathbf{B}(T)$, the trees $T[x]$ for $x \in \partial B_h(T) \cap \mathbf{B}(T)$ are i.i.d. copies of T. Therefore, the slices $\mathcal{S}[x]$ for $x \in \partial B_h(T) \cup \mathbf{B}(T)$ are i.i.d. copies of $\mathcal{S}$. This "self-similarity" property of $\mathcal{S}$ will be used a lot later on.

We end this section by adapting these notions to the more general setting of strips glued in the faces of a tree with no leaf. We recall that a *strip* is an infinite, one-ended planar map with an infinite, simple boundary, such that all the faces except the outer face have finite degree. A strip is also rooted at a root vertex on its boundary. Let $\mathbf{t}$ be an infinite plane tree with no leaf. We draw $\mathbf{t}$ in the plane in such a way that its edges do not intersect (except at a common endpoint), and every compact subset of the plane intersects finitely many vertices and edges. Then $\mathbf{t}$ separates the plane into a countable family $(f_i)_{i\geq 0}$ of faces, where f_0 is the face delimited by the leftmost and the rightmost rays of $\mathbf{t}$, and the other faces are enumerated in a deterministic fashion. For every index $i \geq 0$, we denote by ρ_i the lowest vertex of $\mathbf{t}$ adjacent to f_i, and by h_i its height. Note that this vertex is always unique. On the other hand, for every vertex v of $\mathbf{t}$, there are exactly $c_{\mathbf{t}}(v) - 1$ faces f_i such that $\rho_i = v$.

Let $(s_i)_{i\geq 0}$ be a family of random strips. We denote by $\mathcal{M}(\mathbf{t}, (s_i)_{i\geq 0})$ the infinite planar map obtained by gluing s_i in the face f_i for every $i \geq 0$, in such a way that the root vertex of s_i coincides with ρ_i for every i. We also denote by $\mathcal{S}(\mathbf{t}, (s_i)_{i\geq 0})$ the map obtained by gluing s_i in the face f_i for every $i > 0$ (this is a map with an infinite boundary analog to the slice $\mathcal{S}$). If v is a vertex of $\mathbf{t}$, we also define the "slice of descendants" of v as the map enclosed between the leftmost and the rightmost rays of $\mathbf{t}$ started from v. We denote it by $\mathcal{S}(\mathbf{t}, (s_i))[v]$.

We note that causal maps are a particular case of this construction. This is trivial for supercritical Galton–Watson trees with no leaf. Thanks to Theorem 5.4, this can be extended to the case $\mu(0) > 0$, with $\mathbf{B}(T)$ playing the role of $\mathbf{t}$. This time, however, the strips are random, but they are not independent. Indeed, if v is a vertex of $\mathbf{B}(T)$ and w one of its children in $\mathbf{B}(T)$, the children of v on the left of w and the children on the right of w belong to different strips.

However, by points 2 and 3 of Theorem 5.4, the numbers of children on the left and on the right are not independent, except in some very particular cases (for example if μ is geometric).

In what follows, we will study maps of the form $\mathcal{M}(\mathbf{T}, (s_i)_{i\geq 0})$, where $\mathbf{T}$ is a supercritical Galton–Watson tree with no leaf. We notice right now that if the strips s_i are random and i.i.d., then the slice $\mathcal{S}(\mathbf{T}, (s_i))$ has the same self-similarity property as the causal slices of the form $\mathcal{S}(T)$. Let $h > 0$. We condition on $B_h(\mathbf{T})$ and on all the strips s_i such that $h_i \leq h-1$. Then the trees $\mathbf{T}[v]$ for $v \in \partial B_h(\mathbf{T})$ are independent copies of $\mathbf{T}$, so the slices $\mathcal{S}(\mathbf{T}, (s_i))[v]$ for $v \in \partial B_h(\mathbf{T})$ are i.i.d. copies of $\mathcal{S}(\mathbf{T}, (s_i))$. This will be useful in Section 5.3.

5.2 Metric hyperbolicity properties

The goal of this section is to prove the following result, of which Theorem 5.1 is a particular case. Note that we make no assumption about the strips s_i below.

Theorem 5.1 bis. Let $\mathbf{T}$ be a supercritical Galton–Watson tree with no leaf, and let (s_i) be a sequence of strips. Then:

1. the map $\mathcal{M}(\mathbf{T}, (s_i))$ is a.s. weakly anchored hyperbolic,
2. the map $\mathcal{M}(\mathbf{T}, (s_i))$ a.s. admits bi-infinite geodesics.

In all this section, we will only deal with the general case of $\mathcal{M}(\mathbf{T}, (s_i))$ where $\mathbf{T}$ is a supercritical Galton–Watson tree with no leaf and the s_i are strips. We will write $\mathcal{M}$ for $\mathcal{M}(\mathbf{T}, (s_i))$ and $\mathcal{S}$ for $\mathcal{S}(\mathbf{T}, (s_i))$. Our main tool will be the forthcoming Proposition 5.2, which roughly shows that $\mathcal{S}$ is hard to cross horizontally at large heights.

5.2.1 A hyperbolicity result about slices

We call γ_ℓ and γ_r the left and right boundaries of $\mathcal{S}$, and ρ its root (note that γ_ℓ and γ_r may have an initial segment in common near ρ). Both points of Theorem 5.1 bis will be consequences of the following hyperbolicity result about $\mathcal{S}$.

Proposition 5.2. There is a (random) $K \geq 0$ such that any geodesic in $\mathcal{S}$ from a point on γ_ℓ to a point on γ_r contains a point at distance at most K from ρ.

We first give a very short proof of this proposition in the particular case of a causal slice of the form $\mathcal{S}(T)$. Let $i, j > 0$ and let γ be a geodesic in $\mathcal{S}(T)$ from $\gamma_\ell(i)$ to $\gamma_r(j)$. Let v_0 be the lowest point of γ, and let h_0 be the height of v_0. By the structure of $\mathcal{S}(T)$, each step of γ is either horizontal or vertical. Since the height varies by at most 1 at each vertical step, we need at least $(i-h_0)+(j-h_0)$ vertical steps. Moreover, let $Z_h^{\mathbf{B}}$ be the number of vertices of $\mathbf{B}(T)$ at height h. Then γ needs to cross all the trees $T[x]$ for $x \in \mathbf{B}(T)$ at height h_0, so γ contains at least $Z_{h_0}^{\mathbf{B}}$ horizontal steps. On the other hand, γ is a geodesic so it is shorter than the "obvious" path following γ_ℓ from $\gamma_\ell(i)$ to ρ and then γ_r until $\gamma_r(j)$. Therefore, we have

$$i + j - 2h_0 + Z_{h_0}^{\mathbf{B}} \geq |\gamma| \geq i + j,$$

so $Z_{h_0}^{\mathbf{B}} \leq 2h_0$. However, $Z^{\mathbf{B}}$ has a.s. exponential growth, so this inequality only holds for finitely many values of h_0, so h_0 is bounded independently of i and j.

To generalize this proof, there are two obstacles: first, γ_ℓ and γ_r are no longer geodesics in $\mathcal{S}$ in the general case. Second, an edge of $\mathcal{S}$ can play the role both of a vertical and a horizontal step if it crosses a strip and joins two vertices of $\mathbf{T}$ at different heights. However, we can still cross at most one strip per step in this way.

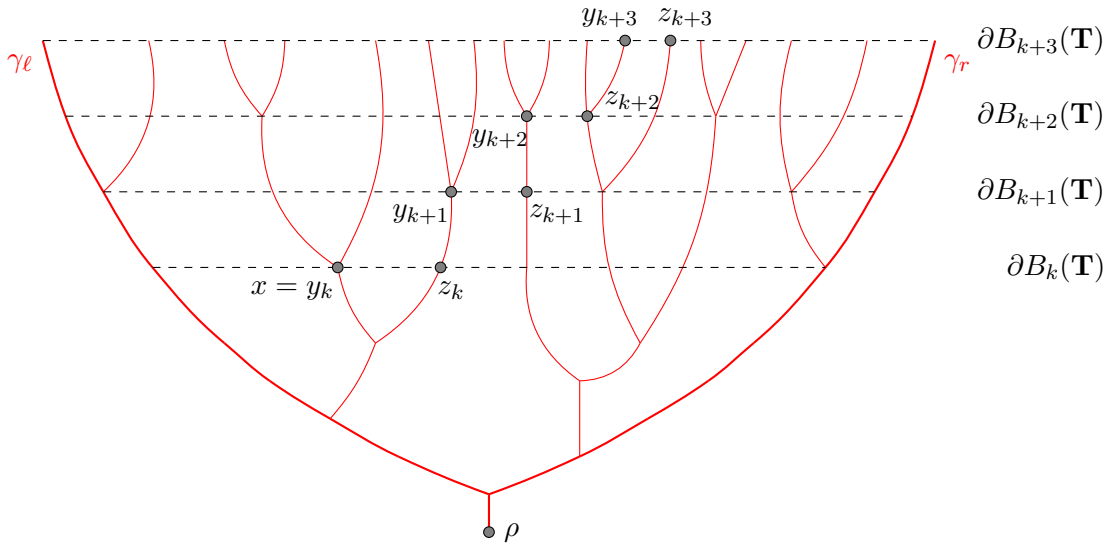

Figure 5.5 – The sequences $(y_i)_{i\geq k}$ and $(z_i)_{i\geq k}$. Here we have taken $u_i = 1$ for every i. The tree $\mathbf{T}$ is in red.

In order to prove the general Proposition 5.2, we first state a lemma showing roughly that if a path in $\mathbf{T}$ with nondecreasing height stays at height h during a time subexponential in h, it cannot cross $\mathcal{S}$.

More precisely, we fix a sequence of positive integers $(u_i)_{i\geq 0}$ and a height $k \geq 0$. Let $x \in \partial B_k(\mathbf{T})$. We define by induction two sequences $(y_i)_{i\geq k}$ and $(z_i)_{i\geq k}$ of vertices of $\mathbf{T}$ with $y_i, z_i \in \partial B_i(\mathbf{T})$ as follows (see Figure 5.5 for an example):

(i) $y_k = x$,

(ii) for every $i \geq k$, if there are at least u_i vertices on the right of y_i on $\partial B_i(\mathbf{T})$ (y_i excluded), then z_i is the u_i-th such vertex,

(iii) if there are less than u_i vertices of $\partial B_i(\mathbf{T})$ on the right of y_i, the sequences (y_i) and (z_i) are killed at time i,

(iv) for every $i \geq k$, if $z_i \notin \gamma_r$, the vertex y_{i+1} is the rightmost child of z_i in $\mathbf{T}$. If $z_i \in \gamma_r$, both sequences are killed.

We call $(y_i)_{i\geq k}$ and $(z_i)_{i\geq k}$ the *sequences escaping from x on the right*. We say that x is *u-far from γ_r* if the sequences (y_i) and (z_i) survive, that is, y_i and z_i are well-defined and do not hit γ_r for all $i \geq k$.

Note that being u-far from γ_r is a monotonic property: if we shift the point x to the left, then the points y_i and z_i are also shifted to the left. Hence, if a point $x \in \partial B_k(\mathbf{T})$ is u-far from γ_r and $x' \in \partial B_k(\mathbf{T})$ lies on the left of x, then x' is also u-far from γ_r. We can similarly define the sequences escaping on the left, and a vertex u-far from γ_ℓ.

Lemma 5.3. Assume u is subexponential, i.e. $u_i = o(c^i)$ for every $c > 1$. Then there is a (random) K such that for any $k \geq K$, the vertex $\gamma_\ell(k)$ is u-far from γ_r and the vertex $\gamma_r(k)$ is u-far from γ_ℓ.

Proof. The idea is to reduce the proof to the study of a supercritical Galton–Watson process where u_i individuals are killed at generation i. It is enough to show that $\gamma_\ell(k)$ is u-far from γ_r for k large enough. Note that if $\gamma_\ell(k)$ is u-far from γ_r, then some point at height $k+1$ is also u-far from γ_r and, by monotonicity, so is $\gamma_\ell(k+1)$. Therefore, it is enough to show that there is $k \geq 0$ such that $\gamma_\ell(k)$ is u-far from γ_r.

Let $k \geq 0$ and $x = \gamma_\ell(k)$. Let $(y_i)_{i\geq k}$ and $(z_i)_{i\geq k}$ be the sequences escaping from x on the right. We also denote by Z_i^k the number of vertices of $\partial B_i(\mathbf{T})$ lying (strictly) on the right of z_i. We first remark that the evolution of the process Z^k can be described explicitly. We have $Z_k^k = Z_k$. Moreover, we recall that μ is the offspring distribution of $\mathbf{T}$. Conditionally on $(Z_k^k, Z_{k+1}^k, \ldots, Z_i^k)$, the variable Z_{i+1}^k has the same distribution as

$$\left(\sum_{j=1}^{Z_i^k} X_{i,j}\right) - u_{i+1},$$

where the $X_{i,j}$ are i.i.d. with distribution μ. To prove our lemma, it is enough to show that $\mathbb{P}\left(\forall i \geq k, Z_i^k > 0\right)$ goes to 1 as k goes to $+\infty$. Since the process Z describing the number of individuals at each generation is a supercritical Galton–Watson process, there is a constant $c > 1$ such that

$$\mathbb{P}\left(\partial B_k(\mathbf{T}) > c^k\right) \xrightarrow[k\to+\infty]{} 1.$$

Therefore, by a monotonicity argument, it is enough to prove that

$$\mathbb{P}\left(\forall i \geq k, Z_i^k > 0 \middle| Z_k^k = \lceil c^k \rceil\right) \xrightarrow[k\to+\infty]{} 1.$$

To prove this, we show that Z^k dominates a supercritical Galton–Watson process. Let $\delta > 0$ be such that $(1-\delta)\sum_{i\geq 0} i\mu(i) > 1$. Let also Z^* be the Markov chain defined by

$$Z_k^* = \lceil c^k \rceil \quad \text{and} \quad Z_{i+1}^* = \left(\sum_{j=1}^{Z_i^*} X_{i,j}\right) - N_{i+1},$$

where the $X_{i,j}$ are i.i.d. with distribution μ and, conditionally on $(X_{i,j})_{i,j\geq 0}$ and $(N_j)_{1\leq j\leq i}$, the variable N_{i+1} has binomial distribution with parameters δ and $\sum_{j=1}^{Z_i^*} X_{i,j}$. In other words, Z^* is a Galton–Watson process in which at every generation, right after reproduction, every individual is killed with probability δ. By our choice of δ, the process Z^* is a supercritical Galton–Watson process and, on survival, grows exponentially. By an easy large deviation argument, we have

$$\mathbb{P}\left(\forall i \geq k, N_i \geq u_i | Z_k^* = \lceil c^k \rceil\right) \xrightarrow[k\to+\infty]{} 1.$$

If this occurs, then $Z_i^* \leq Z_i^k$ for every $i \geq k$ (by an easy induction on i), so

$$\mathbb{P}\left(\forall i \geq k, Z_i^k \geq Z_i^* > 0 | Z_k^k = k\right) \xrightarrow[k\to+\infty]{} 1$$

and the lemma follows. □

The proof of Proposition 5.2 given Lemma 5.3 only relies on deterministic considerations.

Proof of Proposition 5.2. We note that for any $a \in \mathbf{T}$, the slice $\mathcal{S}[a]$ is of the form $\mathcal{S}\left(\mathbf{T}[a], (s'_i)\right)$ where $\mathbf{T}[a]$ is a supercritical Galton–Watson tree with no leaf, so we can apply Lemma 5.3 to $\mathcal{S}[a]$.

Now let a, a' be two vertices of $\mathbf{T}$, neither of which is an ancestor of the other. Then $\mathcal{S}[a]$ and $\mathcal{S}[a']$ are disjoint. Without loss of generality, we may assume that $\mathcal{S}[a]$ lies on the right of $\mathcal{S}[a']$. Let K (resp. K') be given by the conclusion of Lemma 5.3 for $\mathcal{S}[a]$ (resp. $\mathcal{S}[a']$) and $u_i = 2\left(i + \max(d(\rho,a), d(\rho,a'))\right) + 1$. We take $K'' = \max\left(d(\rho,a) + K + 1, d(\rho,a') + K' + 1\right)$.

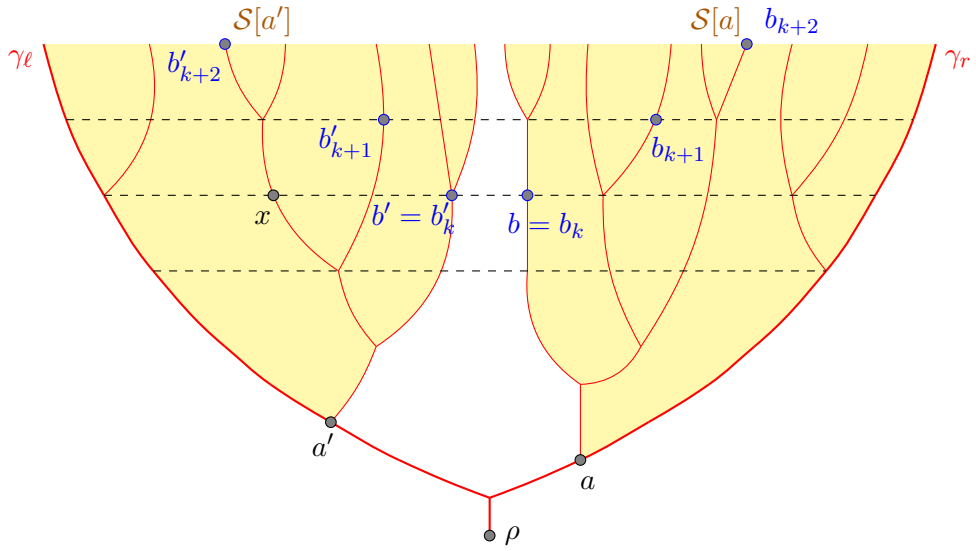

Figure 5.6 – Proof of Proposition 5.2: the point b is u-far from γ_r and b' is u-far from γ_ℓ, so any point on $\partial B_k(\mathbf{T})$ is u-far from either γ_r or γ_ℓ. Here we have taken $u_i = 1$.

We consider a geodesic γ from a vertex $\gamma_\ell(m)$ to a vertex $\gamma_r(n)$ in $\mathcal{S}$. Let k be the minimal height of $\gamma \cap \mathbf{T}$. We assume $k > K''$, and we will get a contradiction.

Let b be the leftmost vertex of $\mathcal{S}[a] \cap \partial B_k(\mathbf{T})$, and let b' be the rightmost vertex of $\mathcal{S}[a'] \cap \partial B_k(\mathbf{T})$ (cf. Figure 5.6). By Lemma 5.3 and our choice of K'', the point b is $\widetilde{u}$-far from γ_r in $\mathcal{S}$ and b' is $\widetilde{u}$-far from γ_ℓ in $\mathcal{S}$, where $\widetilde{u}_i = 2i + 1$ (the change of sequence u is due to the fact that the distances to the root are not the same in $\mathcal{S}[a]$ and $\mathcal{S}$, so the sequence needs to be shifted). But any vertex of $\partial B_k(\mathbf{T})$ lies either on the left of b or on the right of b', so it is either $\widetilde{u}$-far from γ_ℓ or from γ_r (see Figure 5.6). In particular, let x be the first point of γ lying on $\mathbf{T}$ at height k. We may assume that x is u-far from γ_r, the other case can be treated in the same way.

Recall that n is the height of the endpoint of γ. For every $k \leq i \leq n$, let

$$j_i = \max\left\{j \in [\![0, |\gamma|]\!] \middle| \gamma(j) \in \mathbf{T} \text{ and } h\left(\gamma(j)\right) \leq i\right\},$$

and let $x_i = \gamma(j_i)$. Note that we have $h(x_i) \leq i$, but since the height can increase by more than 1 in one step, the inequality may be strict.

Let also $(y_i)_{i \geq k}$ and $(z_i)_{i \geq k}$ be the sequences escaping from x on the right in $\mathcal{S}$ (for $\widetilde{u}_i = 2i+1$). By our assumption that x is $\widetilde{u}$-far from γ_r, these sequences are well-defined and do not hit γ_r. Moreover, for every $k \leq i \leq n$, the vertices x_i and $z_{h(x_i)}$ are both in $\mathbf{T}$ and at the same height $h(x_i)$. We claim that for every $i \geq k$, the vertex x_i lies strictly on the left of the vertex $z_{h(x_i)}$. This is enough to prove the proposition, since then x_n cannot lie on γ_r.

We show this claim by induction on i, and we start with the case $i = k$. The vertices $x = \gamma(j_0)$ and $x_k = \gamma(j_k)$ both lie on $\partial B_k(\mathbf{T})$, so the distance in $\mathcal{S}$ between them is at most $2k$. Hence, since γ is a geodesic, we have $|j_k - j_0| \leq 2k$. Now, we consider the slices $\mathcal{S}[v]$ for $v \in \partial B_k(\mathbf{T})$. These slices are disjoint and, by definition of k, the path $(\gamma_j)_{j_0 \leq j \leq j_k}$ does not cross $\mathbf{T}$ below height k, so it cannot intersect more than $2k < \widetilde{u}_k$ of these slices, which implies that x_k lies on the left on z_k.

We now move on to the induction step. We assume x_i lies strictly on the left of $z_{h(x_i)}$, and split the proof in two cases.

— If $h(x_{i+1}) < i + 1$, then x_{i+1} is the last point of γ at height at most $i + 1$, so it is also the last point of γ at height at most i, so $x_{i+1} = x_i$. In particular, it is on the left of

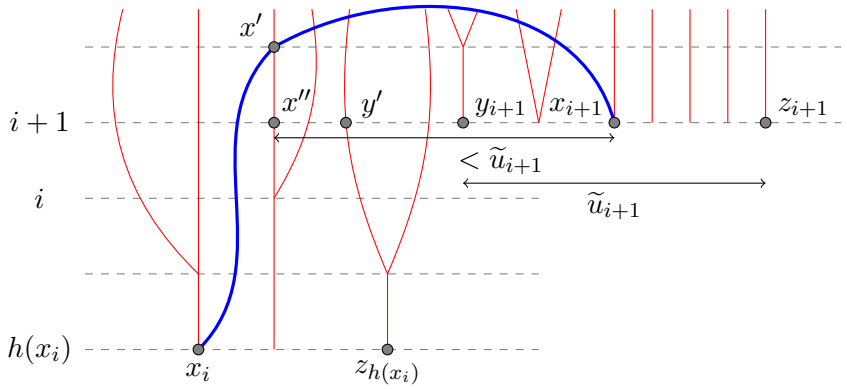

Figure 5.7 – The induction step in the proof of Proposition 5.2. The path γ between x_i and x_{i+1} is in blue. Branches of $\mathbf{T}$ are in red. The dashed lines are not edges of $\mathcal{S}$, but indicate the height. We see that x'' is on the left of y' which is on the left of y_{i+1}, so x_{i+1} is on the left of z_{i+1}.

$z_{h(x_{i+1})} = z_{h(x_i)}$.

— If $h(x_{i+1}) = i+1$, we need to introduce some more notation that is summed up in Figure 5.7. We denote by x' the first point of γ after x_i that belongs to $\mathbf{T}$ (note that $h(x') \geq i+1$ by definition of x_i), and by x'' the ancestor of x' at height $i+1$. Let also y' be the leftmost descendant of $z_{h(x_i)}$ at height $i+1$. Note that by construction of the sequences (y_i) and (z_i), the vertex y' is on the left of y_{i+1}. We know that between x_i and x', the path γ does not cross $\mathbf{T}$, so x_i and x' must be adjacent to the same strip. Since x_i is strictly on the left of $z_{h(x_i)}$, it implies that x' lies on the left of any descendant of z', so x'' is on the left of y', and therefore on the left of y_{i+1}. Moreover, by the definition of x_i, all the vertices of γ between x_i and x_{i+1} that belong to $\mathbf{T}$ have height at least $i+1$. By the same argument as before, the length of the part of γ between x_i and x_{i+1} is at most $h(x_i) + h(x_{i+1}) \leq 2(i+1) < \widetilde{u}_{i+1}$. Hence, this part cannot cross $\widetilde{u}_{i+1}$ of the slices $\mathcal{S}[v]$ with $v \in \partial B_{i+1}(\mathbf{T})$, so the distance between x'' and x_{i+1} along $\partial B_{i+1}(\mathbf{T})$ is less than $\widetilde{u}_{i+1}$. Since x'' is on the left of y_{i+1} and z_{i+1} is at distance $\widetilde{u}_{i+1}$ on the right of y_{i+1}, it follows that x_{i+1} is strictly on the left of z_{i+1}, which concludes the induction and the proof of the proposition.

□

5.2.2 Weak anchored hyperbolicity and bi-infinite geodesics

We can now deduce Theorem 5.1 bis from Proposition 5.2.

Proof of point 1 of Theorem 5.1 bis. Let $(a_i)_{1\leq i\leq 4}$ be four points of $\mathbf{T}$, neither of which is an ancestor of another. The slices $\mathcal{S}[a_i]$ are disjoint and satisfy the assumptions of Proposition 5.2. For every $1 \leq i \leq 4$, let K_i be given by Proposition 5.2 for $\mathcal{S}[a_i]$. Now consider three vertices x, y, z of $\mathcal{M}$ and three geodesics γ_{xy} (resp. γ_{yz}, γ_{zx}) from x to y (resp. y to z, z to x) that surround ρ. There is an index $1 \leq i \leq 4$ such that $\mathcal{S}[a_i]$ contains none of the points x, y and z. Assume it is $\mathcal{S}[a_1]$. Since the triangle formed by γ_{xy}, γ_{yz} and γ_{zx} surrounds ρ, one of these three geodesics must either intersect the path in $\mathbf{T}$ from ρ to a_1, or cross the slice $\mathcal{S}[a_1]$, as on Figure 5.8. We assume this geodesic is γ_{xy}. In the first case, γ_{xy} contains a point at distance at most $h(a_1)$ from ρ. In the second case, assume γ_{xy} crosses $\mathcal{S}[a_1]$ from left to right. Let γ_ℓ and γ_r be respectively the left and right boundaries of $\mathcal{S}[a_1]$. Let v be the last point of γ_{xy} that lies on

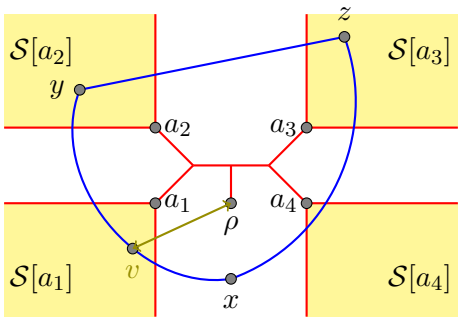

Figure 5.8 – Illustration of the proof of point 1 of Theorem 5.1 bis. Here $\mathcal{S}[a_1]$ contains none of the vertices x, y and z, so it is crossed by the geodesic from x to y, and contains a point v at bounded distance from ρ.

γ_ℓ and let w be the first point of γ_{xy} after v that lies on γ_r. Then the portion of γ_{xy} between v and w is a geodesic in $\mathcal{M}$ so it is also a geodesic in $\mathcal{S}[a_1]$ that crosses $\mathcal{S}[a_1]$. Hence, it contains a point z such that $d(z, a_1) \leq K_1$. This concludes the proof by taking

$$K = \max_{1 \leq i \leq 4} (K_i + h(a_i)) .$$

□

Proof of point 2 of Theorem 5.1 bis. Let $a_1, a_2 \in \mathbf{T}$, neither of which is an ancestor of the other, so that $\mathcal{S}[a_1]$ and $\mathcal{S}[a_2]$ are disjoint. Let γ_ℓ and γ_r be the left and right boundaries of $\mathcal{S}[a_1]$. The idea of our construction is the following: we first "approximate" the paths γ_ℓ and γ_r by two infinite geodesics $\widetilde{\gamma}_\ell$ and $\widetilde{\gamma}_r$, and we then try to connect $\widetilde{\gamma}_\ell$ to $\widetilde{\gamma}_r$ in the shortest possible way. Before expliciting this construction, we need to reinforce slightly Proposition 5.2.

By Proposition 5.2, we know that any geodesic γ in $\mathcal{S}[a_1]$ between a point of γ_ℓ and a point of γ_r contains a vertex at bounded distance from ρ. We claim that this is also the case if we consider geodesics in $\mathcal{M}$ instead of $\mathcal{S}[a_1]$. Indeed, let K_1 (resp. K_2) be given by Proposition 5.2 for $\mathcal{S}[a_1]$ (resp. $\mathcal{S}[a_2]$). Let $i, j \geq 0$ and let γ be a geodesic from $\gamma_\ell(i)$ to $\gamma_r(j)$ in $\mathcal{M}$. We are in one of the three following cases (cf. Figure 5.9) :

(i) γ intersects the path in $\mathbf{T}$ from ρ to a_1 or from ρ to a_2,

(ii) γ crosses $\mathcal{S}[a_1]$,

(iii) γ crosses $\mathcal{S}[a_2]$.

In all three cases, γ contains a point at distance at most K from ρ, where

$$K = \max\left(h(a_1) + K_1, h(a_2) + K_2\right). \tag{5.2}$$

We can now build our infinite geodesics $\widetilde{\gamma}_\ell$ and $\widetilde{\gamma}_r$. For every n, let γ_ℓ^n be a geodesic from ρ to $\gamma_\ell(n)$. By an easy compactness argument, there is an infinite geodesic $\widetilde{\gamma}_\ell$ such that, for every $i \geq 0$, there are infinitely many n such that $\widetilde{\gamma}_\ell(i) = \gamma_\ell^n(i)$. We build an infinite geodesic $\widetilde{\gamma}_r$ from γ_r in a similar way.

For every $i, j \geq 0$, we define

$$a_{i,j} = i + j - d_{\mathcal{M}}\left(\widetilde{\gamma}_\ell(i), \widetilde{\gamma}_r(j)\right).$$

This quantity measures "by how much" the concatenation of $\widetilde{\gamma}_\ell$ and $\widetilde{\gamma}_r$ between $\widetilde{\gamma}_\ell(i)$ and $\widetilde{\gamma}_r(j)$ is not a geodesic. We note that for any $i, j \geq 0$, we have

$$d_{\mathcal{M}}\left(\widetilde{\gamma}_\ell(i+1), \widetilde{\gamma}_r(j)\right) \leq d_{\mathcal{M}}\left(\widetilde{\gamma}_\ell(i+1), \widetilde{\gamma}_\ell(i)\right) + d_{\mathcal{M}}\left(\widetilde{\gamma}_\ell(i), \widetilde{\gamma}_r(j)\right) = 1 + d_{\mathcal{M}}\left(\widetilde{\gamma}_\ell(i), \widetilde{\gamma}_r(j)\right),$$

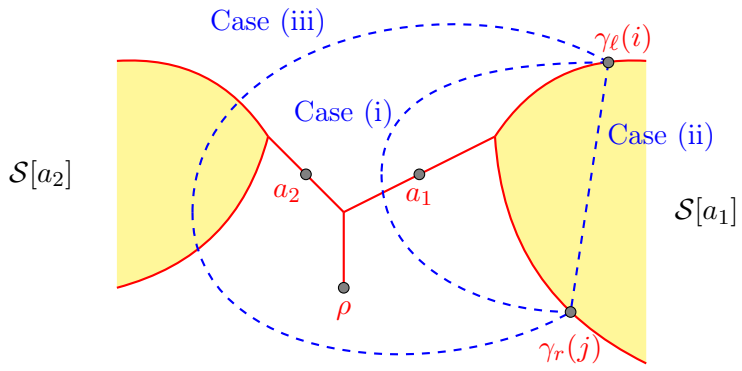

Figure 5.9 – Reinforcement of Proposition 5.2. In blue, the geodesic γ. It must intersect a geodesic from ρ to a_1 or a_2, or cross $\mathcal{S}[a_1]$ or $\mathcal{S}[a_2]$.

so $a_{i+1,j} \geq a_{i,j}$, so $a_{i,j}$ is nondecreasing in i. Similarly, it is nondecreasing in j. We claim the following.

Lemma 5.4. Almost surely, $(a_{i,j})_{i,j\geq 0}$ is bounded.

Proof of Lemma 5.4. Let $i, j \geq 0$. By the definition of $\widetilde{\gamma}_\ell$ and $\widetilde{\gamma}_r$, there are two indices m and n such that $\widetilde{\gamma}_\ell(i)$ lies on a geodesic from ρ to $\gamma_\ell(m)$ and $\widetilde{\gamma}_r(j)$ lies on a geodesic from ρ to $\gamma_r(n)$. Therefore, we have

$$\begin{aligned} d_{\mathcal{M}}(\gamma_\ell(m), \gamma_r(n)) &\leq d_{\mathcal{M}}(\gamma_\ell(m), \widetilde{\gamma}_\ell(i)) + d_{\mathcal{M}}(\widetilde{\gamma}_\ell(i), \widetilde{\gamma}_r(j)) + d_{\mathcal{M}}(\widetilde{\gamma}_r(j), \gamma_r(n)) \\ &= d_{\mathcal{M}}(\rho, \gamma_\ell(m)) - i + d_{\mathcal{M}}(\widetilde{\gamma}_\ell(i), \widetilde{\gamma}_r(j)) + d_{\mathcal{M}}(\rho, \gamma_r(n)) - j \\ &= d_{\mathcal{M}}(\rho, \gamma_\ell(m)) + d_{\mathcal{M}}(\rho, \gamma_r(n)) - a_{i,j}. \end{aligned}$$

On the other hand, we know that for any $m, n \geq 0$, any geodesic from $\gamma_\ell(m)$ to $\gamma_r(n)$ contains a vertex v_0 at distance at most K from ρ, where K is given by (5.2). Therefore, we have

$$\begin{aligned} d_{\mathcal{M}}(\gamma_\ell(m), \gamma_r(n)) &= d_{\mathcal{M}}(\gamma_\ell(m), v_0) + d_{\mathcal{M}}(v_0, \gamma_r(n)) \\ &\geq d_{\mathcal{M}}(\gamma_\ell(m), \rho) + d_{\mathcal{M}}(\gamma_r(n), \rho) - 2d_{\mathcal{M}}(\rho, v_0) \\ &\geq d_{\mathcal{M}}(\gamma_\ell(m), \rho) + d_{\mathcal{M}}(\gamma_r(n), \rho) - 2K. \end{aligned}$$

By combining the last two equations, we obtain $a_{i,j} \leq 2K$ for every $i, j \geq 0$. □

The construction of a bi-infinite geodesic is now easy. Let i_0, j_0 be two indices such that $a_{i_0,j_0} = \sup\{a_{i,j} | i, j \geq 0\}$, and let $d = d_{\mathcal{M}}(\widetilde{\gamma}_\ell(i_0), \widetilde{\gamma}_r(j_0))$. Let also $\widehat{\gamma}$ be a geodesic from $\widetilde{\gamma}_\ell(i_0)$ to $\widetilde{\gamma}_r(j_0)$ in $\mathcal{M}$. We define a bi-infinite path γ as follows :

$$\gamma(i) = \begin{cases} \widetilde{\gamma}_\ell(i_0 - i) & \text{if } i \leq 0, \\ \widehat{\gamma}(i) & \text{if } 0 \leq i \leq d, \\ \widetilde{\gamma}_r(i - d + j_0) & \text{if } i \geq d. \end{cases}$$

We finally check that this is indeed a bi-infinite geodesic. Let $i \geq i_0$ and $j \geq j_0$. We have $a_{i,j} = a_{i_0,j_0}$, so

$$d_{\mathcal{M}}(\widetilde{\gamma}_\ell(i), \widetilde{\gamma}_r(j)) = d + (i - i_0) + (j - j_0),$$

so $d(\gamma(i_0 - i), \gamma(d + j - j_0)) = (d + j - j_0) - (i_0 - i)$. Therefore, we have $d(\gamma(i'), \gamma(j')) = j' - i'$ for $i' \leq 0$ small enough and $j' \geq 0$ large enough, so γ is a bi-infinite geodesic. □

5.3 Poisson boundary

5.3.1 General setting

The goal of this subsection is to build a compactification of maps that, as we will later prove, is under some assumptions a realization of their Poisson boundary. We will perform this construction directly in the general framework $\mathcal{M} = \mathcal{M}(\mathbf{T}, (s_i))$. This construction is exactly the same as the construction performed for the PSHIT in Section 3.1 of Chapter 4.

We recall that $\partial\mathbf{T}$ is the set of infinite rays from ρ in $\mathbf{T}$. If $\gamma, \gamma' \in \partial\mathbf{T}$, we write $\gamma \sim \gamma'$ if $\gamma = \gamma'$ or if γ and γ' are "consecutive" in the sense that there is no ray between them (in particular, if γ_ℓ and γ_r are the leftmost and rightmost rays of $\mathbf{T}$, then $\gamma_\ell \sim \gamma_r$). It is equivalent to saying that γ and γ' are the left and right boundaries of some strip s_i in the map $\mathcal{M}$. Note that a.s., every ray of $\mathbf{T}$ contains infinitely many branching points, so no ray is equivalent to two distinct other rays. It follows that $\sim$ is a.s. an equivalence relation for which countably many equivalence classes have cardinal 2, and all the others have cardinal 1. We write $\widehat{\partial}\mathbf{T} = \partial\mathbf{T}/\sim$ and we denote by $\gamma \to \widehat{\gamma}$ the canonical projection from $\partial\mathbf{T}$ to $\widehat{\partial}\mathbf{T}$. Finally, for every strip s_i, the left and right boundaries of s_i correspond to the same point of $\widehat{\partial}\mathbf{T}$, that we denote by $\widehat{\gamma}_i$.

Our goal is now to define a topology on $\mathcal{M} \cup \widehat{\partial}\mathbf{T}$. It should be possible to define it by an explicit distance, but such a distance would be tedious to write down, so we prefer to give an "abstract" construction. Let s_i and s_j be two distinct strips of $\mathcal{M}$, and fix $h > 0$ such that both s_i and s_j both intersect $B_h(\mathbf{T})$. Then $\mathcal{M} \setminus (B_h(\mathcal{M}) \cup s_i \cup s_j)$ has two infinite connected components, that we denote by (s_i, s_j) and (s_j, s_i) (the vertices on the boundaries of s_i and s_j do not belong to (s_i, s_j) and (s_j, s_i)). We also write

$$\widehat{\partial}(s_i, s_j) = \{\widehat{\gamma} |\ \gamma \text{ is a ray of } \mathbf{T} \text{ such that } \gamma(k) \in (s_i, s_j) \text{ for } k \text{ large enough}\}.$$

We define $\widehat{\partial}(s_j, s_i)$ similarly. Note that $\widehat{\partial}(s_i, s_j)$ and $\widehat{\partial}(s_j, s_i)$ are disjoint subsets of $\widehat{\partial}\mathbf{T}$, and their union is $\widehat{\partial}\mathbf{T} \setminus \{\widehat{\gamma}_i, \widehat{\gamma}_j\}$.

We can now equip the set $\mathcal{M} \cup \widehat{\partial}\mathbf{T}$ with the topology generated by the following open sets:
— the singletons $\{v\}$, where v is a vertex of $\mathcal{M}$,
— the sets $(s_i, s_j) \cup \widehat{\partial}(s_i, s_j)$, where s_i and s_j are two distinct strips of $\mathcal{M}$.

This topology is separated (if $\widehat{\gamma}_1 \neq \widehat{\gamma}_2$, then there are two strips separating γ_1 and γ_2) and has a countable basis, so it is induced by a distance. Moreover, any open set of our basis intersects $\mathcal{M}$, so $\mathcal{M}$ is dense in $\mathcal{M} \cup \widehat{\partial}\mathbf{T}$. Finally, we state an intuitive result about the topology of $\mathcal{M} \cup \widehat{\partial}\mathbf{T}$. Its proof in the particular case of the PSHIT can be found in Chapter 4, and adapts without any change to the general case.

Lemma 5.5. The space $\mathcal{M} \cup \widehat{\partial}\mathbf{T}$ is compact, and $\widehat{\partial}\mathbf{T}$ is homeomorphic to the unit circle.

5.3.2 Transience away from the boundary in causal slices

The goal of this section is to prove Proposition 5.6, which is the main tool in the proof of Theorem 5.2. We recall that $\mathcal{S}$ is the causal slice associated to a supercritical Galton–Watson tree T, and $\partial\mathcal{S}$ is the *boundary* of $\mathcal{S}$, i.e. the set of vertices of $\mathcal{S}$ that are either the leftmost or the rightmost vertex of their generation. We also write $\tau_{\partial\mathcal{S}} = \min\{n \geq 0 | X_n \in \partial\mathcal{S}\}$, where (X_n) is the simple random walk on $\mathcal{S}$.

Proposition 5.6. Almost surely, there is a vertex $x \in \mathcal{S}$ such that

$$P_{\mathcal{S},x}(\tau_{\partial\mathcal{S}} = +\infty) > 0.$$

Note that if such a vertex x exists, then we have $P_{\mathcal{S},v}(X_n \notin \partial\mathcal{S} \text{ for } n \text{ large enough}) > 0$ for every vertex $v \in \mathcal{S}$. The proof of Proposition 5.6 is based on estimates of effective resistances. We will use the following inequality, that holds for every graph and every vertex x:

$$P_{\mathcal{S},x}\left(\tau_{\partial\mathcal{S}} < +\infty\right) \leq \frac{R^{\mathcal{S}}_{\text{eff}}(x \leftrightarrow \{\partial\mathcal{S}, \infty\})}{R^{\mathcal{S}}_{\text{eff}}(x \leftrightarrow \partial\mathcal{S})} \leq \frac{R^{\mathcal{S}}_{\text{eff}}(x \leftrightarrow \infty)}{R^{\mathcal{S}}_{\text{eff}}(x \leftrightarrow \partial\mathcal{S})}. \tag{5.3}$$

For example, this is a particular case of Exercise 2.36 of [110]. We will find a sequence (x_n) of vertices satisfying the following two properties:

1. we have $R^{\mathcal{S}}_{\text{eff}}(x_n \leftrightarrow \partial\mathcal{S}) \to +\infty$ a.s. when $n \to +\infty$,
2. for every $n \geq 0$, the resistance $R^{\mathcal{S}}_{\text{eff}}(x_n \leftrightarrow \infty)$ is stochastically dominated by $R^{\mathcal{S}}_{\text{eff}}(\rho \leftrightarrow \infty)$. In particular, a.s., $\left(R^{\mathcal{S}}_{\text{eff}}(x_n \leftrightarrow \infty)\right)$ has a bounded subsequence.

By (5.3), this will guarantee that

$$P_{\mathcal{S},x_n}\left(\tau_{\partial\mathcal{S}} < +\infty\right) \xrightarrow[n\to+\infty]{} 0$$

along some subsequence, which is enough to prove Proposition 5.6.

We choose for the sequence (x_n) the nonbacktracking random walk on the backbone of T. More precisely, we take $x_0 = \rho$ and, for every $n \geq 0$, conditionally on $\mathcal{S}$ and $x_0, \ldots, x_n$, the vertex x_{n+1} is chosen uniformly among the children of x_n in $\mathbf{B}(T)$. We can give a "spinal decomposition" of $\mathbf{B}(T)$ along (x_n). We recall that $\boldsymbol{\mu}$ is the offspring distribution of $\mathbf{B}(T)$, cf. (5.1). For every $n \geq 0$, let L_n (resp. R_n) be the number of children of x_n in $\mathbf{B}(T)$ on the left (resp. on the right) of x_{n+1}. A vertex v of $\mathbf{B}(T)$ will be called a *spine brother* if the parent of v is equal to x_n for some n but $v \neq x_{n+1}$. Then the pairs (L_n, R_n) are i.i.d. with distribution ν given by

$$\mathbb{P}\left(L_n = \ell, R_n = r\right) = \nu\left(\{(\ell, r)\}\right) = \frac{1}{r+\ell+1}\boldsymbol{\mu}(r+\ell+1). \tag{5.4}$$

Moreover, conditionally on (L_n) and (R_n), the backbones of the trees of descendants of the spine brothers are i.i.d. Galton–Watson trees with offspring distribution $\boldsymbol{\mu}$. The distribution of T conditionally on this backbone is then given by Theorem 5.4. In particular, for every $n \geq 0$, the tree of descendants of x_n has the same distribution as T, so $\mathcal{S}[x_n]$ has the same distribution as $\mathcal{S}$.

Therefore, for every n, we have

$$R^{\mathcal{S}}_{\text{eff}}(x_n \leftrightarrow \infty) \leq R^{\mathcal{S}[x_n]}_{\text{eff}}(x_n \leftrightarrow \infty),$$

where $R^{\mathcal{S}[x_n]}_{\text{eff}}(x_n \leftrightarrow \infty)$ has the same distribution as $R^{\mathcal{S}}_{\text{eff}}(\rho \leftrightarrow \infty)$. This proves the second property that we wanted (x_n) to satisfy. Hence, it only remains to prove that $R^{\mathcal{S}}_{\text{eff}}(x_n \leftrightarrow \partial\mathcal{S})$ almost surely goes to $+\infty$, which is the goal of the next lemma.

Lemma 5.7. There are disjoint vertex sets $(A_k)_{k\geq 0}$, satisfying the following properties:

(i) for any $k \geq 1$, the set A_k separates $\partial\mathcal{S}$ from all the sets A_i with $i > k$, and from x_n for n large enough,

(ii) the parts of $\mathcal{S}$ lying between A_{2k} and A_{2k+1} for $k \geq 0$ are i.i.d.,

(iii) we have $R^{\mathcal{S}}_{\text{eff}}(A_{2k} \leftrightarrow A_{2k+1}) > 0$ a.s. for every $k \geq 0$.

Proof of Proposition 5.6 given Lemma 5.7. Fix k and choose n large enough, so that x_n is separated from $\partial\mathcal{S}$ by $A_0, A_1, \ldots, A_{2k-1}$. Since $A_0, \ldots, A_{2k-1}$ are disjoint cutsets separating x_n from

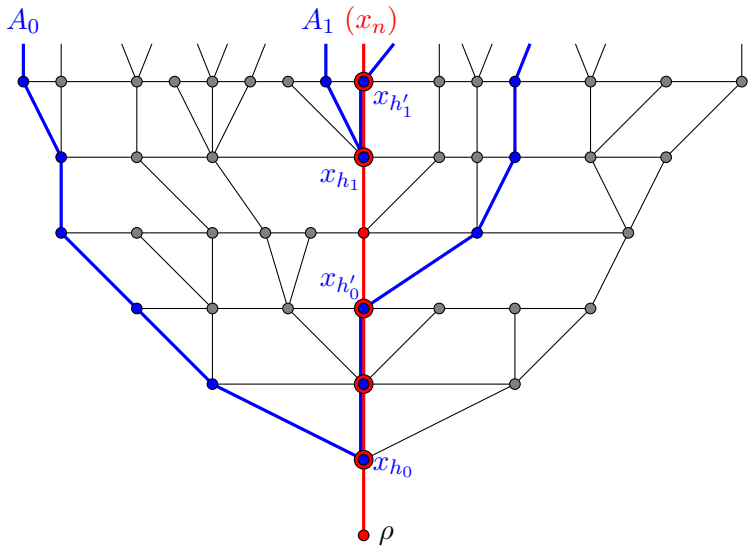

Figure 5.10 – The slice $\mathcal{S}$ with the sequences of vertices (x_n) (in red), and the separating sets A_k (in blue). Here we have $h_0 = 1$, $h'_0 = 3$, $h_1 = 5$ and $h'_1 = 6$.

$\partial\mathcal{S}$, we have

$$\begin{aligned} R^{\mathcal{S}}_{\text{eff}}(x_n \leftrightarrow \partial\mathcal{S}) &\geq R^{\mathcal{S}}_{\text{eff}}(\partial\mathcal{S} \leftrightarrow A_0) + R^{\mathcal{S}}_{\text{eff}}(x_n \leftrightarrow A_{2k-1}) + \sum_{i=0}^{2k-2} R^{\mathcal{S}}_{\text{eff}}(A_i \leftrightarrow A_{i+1}) \\ &\geq \sum_{i=0}^{k-1} R^{\mathcal{S}}_{\text{eff}}(A_{2i} \leftrightarrow A_{2i+1}). \end{aligned}$$

Since the variables $R^{\mathcal{S}}_{\text{eff}}(A_{2i} \leftrightarrow A_{2i+1})$ for $0 \leq i \leq k-1$ are i.i.d. and a.s. positive, this goes to $+\infty$ as $k \to +\infty$, so we have $R^{\mathcal{S}}_{\text{eff}}(x_n \leftrightarrow A_0) \to +\infty$ a.s. when $n \to +\infty$, which ends the proof. □

Remark 5.8. The law of large numbers even shows that there is a constant $c > 0$ such that, for n large enough, we have $R^{\mathcal{S}}_{\text{eff}}(x_n \leftrightarrow \partial\mathcal{S}) \geq cn$. This is quite similar to the resistance estimates proved in [26] in the case where T is the complete binary tree, and it might be interesting to apply our results to the study of the Gaussian free field on causal maps. Unfortunately, our estimates only hold in "typical" directions, and not uniformly for all the vertices. We also note that the idea to "cut" $\mathcal{S}$ along the sets A_k was inspired by the proof of Lemma 1 of [26].

We now build the subsets A_k. We define by induction the heights h_k and h'_k for $k \geq 0$ by

$$\begin{aligned} h_0 &= \min\{n \geq 0 | L_n > 0\}, \\ h'_k &= \min\{n > h_k | R_n > 0\}, \\ h_{k+1} &= \min\{n > h'_k | L_n > 0\}. \end{aligned}$$

Note that the pairs (L_n, R_n) are i.i.d. with $\mathbb{P}(L_n > 0) > 0$ and $\mathbb{P}(R_n > 0) > 0$, so h_k and h'_k are a.s. well-defined for every k. We define A_k as the union of the leftmost ray of $\mathbf{B}(T)$ from x_{h_k}, the rightmost ray of $\mathbf{B}(T)$ from $x_{h'_k}$ and the vertices x_i with $h_k \leq i \leq h'_k$ (see Figure 5.10).

It is easy to see that the sets A_k are disjoint and that A_k separates $\partial\mathcal{S}$ from A_i for every $i > k$, and from x_n for n large enough, so they satisfy property (i) of Lemma 5.7.

For every $k \geq 0$, let $\mathcal{U}_k$ be the sub-map of $\mathcal{S}$ whose vertices are the vertices between A_k and A_{k+1} (the vertices of A_k and A_{k+1} are included), rooted at x_{h_k}. By using (5.4) and the backbone decomposition, it is quite straightforward to prove that the maps $\mathcal{U}_{2k}$ for $k \geq 0$ are i.i.d.. We do not give a precise description of the distribution of $\mathcal{U}_k$, the only property of $\mathcal{U}_k$ that we will need later is that it contains a copy of T on each side of the spine.

Remark 5.9. It is still true that the $\mathcal{U}_k$ for $k \geq 0$ are identically distributed. However, $\mathcal{U}_k$ and $\mathcal{U}_{k+1}$ are not independent. Indeed, for $h_{k+1} \leq n < h'_{k+1}$, the variables L_n and R_n are not independent, and L_n affects $\mathcal{U}_{k+1}$, whereas R_n affects $\mathcal{U}_k$. This is why we restrict ourselves to even values of k.

It only remains to prove property (iii) in Lemma 5.7. Let $\mathcal{U}$ be distributed as the $\mathcal{U}_k$, and let A_b (resp. A_t) be its bottom (resp. top) boundary, playing the same role as A_k (resp. A_{k+1}). The proof that $R^{\mathcal{U}}_{\text{eff}}(A_b \leftrightarrow A_t) > 0$ relies on a duality argument. We first recall a classical result about duality of resistances in planar maps. Let M be a finite planar map drawn in the plane, and let a and z be two vertices adjacent to its outer face. We draw two infinite half-lines from a and z that split the outer face in two faces a^* and z^*. Now let M^* be the dual planar map whose vertices are a^*, z^* and the internal faces of M. Then we have

$$R^{M}_{\text{eff}}(a \leftrightarrow z) = \left(R^{M^*}_{\text{eff}}(a^* \leftrightarrow z^*)\right)^{-1}. \tag{5.5}$$

In our case, the infinite graph $\mathcal{U}$ has two ends ∞_ℓ (on the left of the spine) and ∞_r (on the right). Informally, we would like to write

$$R^{\mathcal{U}}_{\text{eff}}(A_b \leftrightarrow A_t) = \left(R^{\mathcal{U}^*}_{\text{eff}}(\infty^*_\ell \leftrightarrow \infty^*_r)\right)^{-1}, \tag{5.6}$$

which would reduce the problem to the proof of $R^{\mathcal{U}^*}_{\text{eff}}(\infty^*_\ell \leftrightarrow \infty^*_r) < +\infty$. Our first job will be to state and prove (a proper version of) (5.6).

More precisely, we denote by $\mathcal{U}_\ell$ (resp. $\mathcal{U}_r$) the part of $\mathcal{U}$ lying on the left (resp. on the right) of the spine. We also define $\mathcal{U}^*$ as the dual map of $\mathcal{U}$ in the following sense: the vertices of $\mathcal{U}^*$ are the finite faces of $\mathcal{U}$, and for every edge of $\mathcal{U}$ that does not link two vertices of A_b or two vertices of A_t, we draw an edge e^* between the two faces adjacent to e. Let θ^* be a flow on $\mathcal{U}^*$ with no source. We assume that θ^* is unitary in the sense that the mass of θ^* crossing the spine from left to right is equal to 1. We recall that the *energy* $\mathcal{E}(\theta^*)$ of θ^* is the sum over all edges e^* of $\mathcal{U}^*$ of $\theta^*(e^*)^2$. For every $n \geq 0$, let $A_b(n)$ (resp. $A_t(n)$) be the set of vertices of A_b (resp. A_t) at height at most n. We consider the map $\mathcal{U}(n)$ obtained by cutting $\mathcal{U}$ above height n. The restriction of θ to the dual of this map becomes a unitary flow crossing $\mathcal{U}(n)^*$ from left to right. Therefore, the dual resistance from left to right in $\mathcal{U}(n)^*$ is at most $\mathcal{E}(\theta^*)$ so, by (5.5), we obtain

$$R^{\mathcal{U}}_{\text{eff}}(A_b(n) \leftrightarrow A_t(n)) \geq \mathcal{E}(\theta^*)^{-1}$$

and, by letting $n \to +\infty$, we get $R^{\mathcal{U}}_{\text{eff}}(A_b \leftrightarrow A_t) \geq \mathcal{E}(\theta^*)^{-1}$. In particular, if there is such a flow θ^* with finite energy, then $R^{\mathcal{U}}_{\text{eff}}(A_b \leftrightarrow A_t) > 0$ and Lemma 5.7 is proved.

We now define $\mathcal{U}_\ell$ (resp. $\mathcal{U}_r$) as the part of $\mathcal{U}$ lying on the left (resp. on the right) of the spine. Let f_ℓ and f_r be two faces of $\mathcal{U}$ lying respectively on the left and on the right of the same edge of the spine. A simple way to construct a unitary flow θ^* with no sources is to concatenate a flow θ^*_ℓ from infinity to f_ℓ in $\mathcal{U}^*_\ell$, a flow of mass 1 in the dual edge from f_ℓ to f_r and a flow θ^*_r from f_r to infinity in $\mathcal{U}^*_r$. For this flow to have finite energy, we need θ^*_ℓ and θ^*_r to have finite energy, so we need both $\mathcal{U}^*_\ell$ and $\mathcal{U}^*_r$ to be transient.

We now define $\mathcal{S}^*$ as the dual of the slice $\mathcal{S}$ (as above, the vertices of $\mathcal{S}^*$ are the inner faces of $\mathcal{S}$). We note that, for every vertex $v_0 \in \mathcal{U}_\ell \cap \mathbf{B}(T)$ that does not belong to the spine, the tree

of descendants of v_0 has the same distribution as T and is entirely contained in $\mathcal{U}_\ell$. Therefore, $\mathcal{U}_\ell$ contains a copy of $\mathcal{S}$, so $\mathcal{U}_\ell^*$ contains a copy of $\mathcal{S}^*$, and the same is true for $\mathcal{U}_r^*$. Hence, we have reduced the proof of Proposition 5.6 to the next result.

Lemma 5.10. The dual slice $\mathcal{S}^*$ is a.s. transient.

Proof. We will show that we can embed a transient tree in $\mathcal{S}^*$. The idea will be to follow the branches of the tree T in the dual, to obtain a tree T^* that is similar to T. However, vertices of high degree become obstacles: if a vertex v of T has degree d in T, we need d dual edges to "move around" v in $\mathcal{S}^*$. Therefore, it becomes difficult to control the ratio between resistances in T^* and in T. To circumvent this problem, we will use the fact that T contains a supercritical Galton–Watson tree with bounded degrees.

More precisely, we fix a constant $c_{\max}$ large enough to have

$$\sum_{i=0}^{c_{\max}} i\mu(i) > 1.$$

If t is a (finite or infinite) tree, for every vertex v with more than $c_{\max}$ children, we remove all the edges between v and its children, and we call t_{bdd} the connected component of the root. If T' is a Galton–Watson tree with distribution μ, then T'_{bdd} is a Galton–Watson tree with offspring distribution μ_{bdd} given by

$$\mu_{\mathrm{bdd}}(i) = \begin{cases} 0 & \text{if } i > c_{\max}, \\ \mu(i) & \text{if } 0 < i \leq c_{\max}, \\ \mu(0) + \sum_{j > c_{\max}} \mu(j) & \text{if } i = 0. \end{cases}$$

In particular, we have $\sum_i i\mu_{\mathrm{bdd}}(i) = \sum_{i=0}^{c_{\max}} i\mu(i) > 1$, so T'_{bdd} is supercritical, and it survives with positive probability. But T is a Galton–Watson tree conditioned to survive, so it contains infinitely many i.i.d. copies of T (take for examples the trees of descendants of the children of the right boundary). Therefore, there is a.s. a vertex $v_0 \in T$ that is not on the left boundary of $\mathcal{S}$, such that $T[v_0]_{\mathrm{bdd}}$ is a Galton—Watson tree and survives. In particular, it is transient and, for every $v \in T[v_0]_{\mathrm{bdd}}$, the number of children of v in T is bounded by $c_{\max}$.

From here, can can build a tree T^* in $\mathcal{S}^*$ whose branches follow the branches of $T[v_0]_{\mathrm{bdd}}$ on their left, and which circumvents branching points of $T[v_0]_{\mathrm{bdd}}$ by the top. See Figure 5.11 for the construction of this tree. The tree T^* is then a subgraph of $\mathcal{S}^*$. Therefore, it is enough to prove that T^* is transient. But since the vertex degrees in $T[v_0]_{\mathrm{bdd}}$ are bounded, it is easy to see that $T[v_0]_{\mathrm{bdd}}$ and T^* are quasi-isometric, so T^* is also transient (by e.g. Section 2.4.4 of [72]). □

5.3.3 Consequences on the Poisson boundary

We recall that (X_n) is the simple random walk on $\mathcal{C}$ started from ρ. By a result of Hutchcroft and Peres (Theorem 1.3 of [83]), the first point of Theorem 5.2 implies the second.

Proof of Theorem 5.2. We first show the almost sure convergence of (X_n). By compactness (Lemma 5.5), it is enough to prove that (X_n) a.s. has a unique subsequential limit in $\widehat{\partial}T$. Note that if $\widehat{\gamma}_1 \neq \widehat{\gamma}_2$ are two distinct points of $\widehat{\partial}T$, then there are two vertices $v, v' \in \mathbf{B}(T)$ such that the slices $\mathcal{S}[v]$ and $\mathcal{S}[v']$ separate γ_1 from γ_2. Therefore, if γ_1 and γ_2 are subsequential limits of (X_n), by transience of $\mathcal{C}(T)$, the walk (X_n) crosses infinitely many times either $\mathcal{S}[v]$ or $\mathcal{S}[v']$ horizontally. Therefore, it is enough to prove that for every $v_0 \in \mathbf{B}(T)$, the walk (X_n) cannot cross infinitely many times $\mathcal{S}[v_0]$ horizontally.

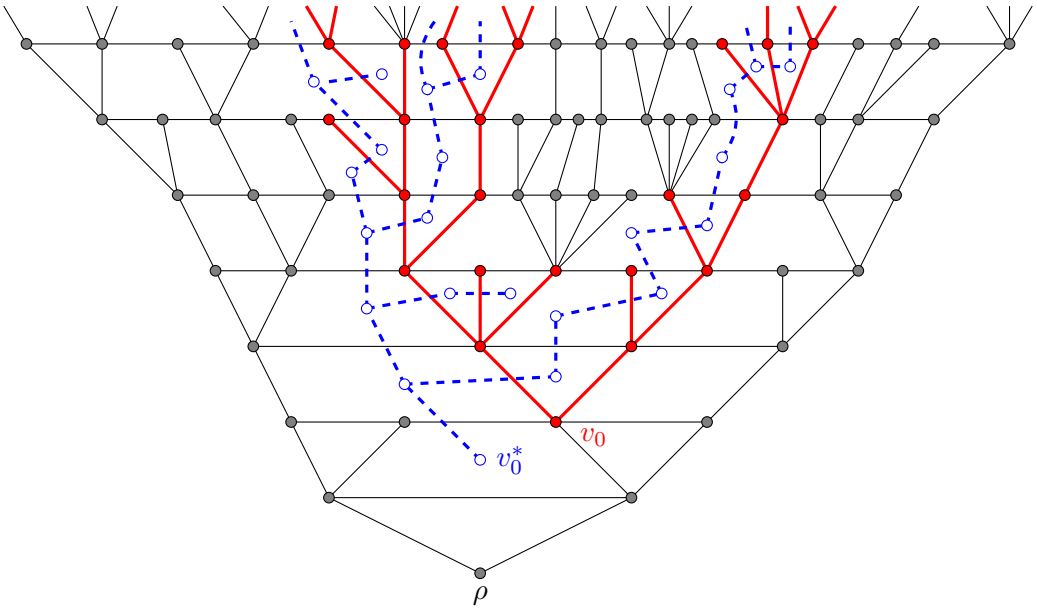

Figure 5.11 – The construction of the dual tree T^* (in blue) from the tree $T[v_0]_{\text{bdd}}$ (in red). Here, we have taken $c_{\max} = 3$. The vertices of T^* are the faces adjacent to the vertices of $T[v_0]_{\text{bdd}}$ at their bottom-left corners.

For every $v \in \mathcal{S}[v_0]$, let $f(v) = P_{\mathcal{S}[v_0],v}\left(\tau_{\partial\mathcal{S}[v_0]} < +\infty\right)$. The function f is harmonic on $\mathcal{S}[v_0]\backslash\partial\mathcal{S}[v_0]$. Moreover, by Proposition 5.6, there is a vertex $v_1 \in \mathcal{S}[v_0]$ such that $f(v_1) < 1$. Let (Y_n) be a simple random walk started from v_1 and killed when it hits $\partial\mathcal{S}[v_0]$. Then $f(Y_n)$ is a martingale and, by the martingale convergence theorem, it converges a.s. to $\mathbb{1}_{\tau_{\partial\mathcal{S}[v_0]}<+\infty}$. In particular, it has limit zero with positive probability, so there is an infinite path (w_k) going to infinity in $\mathcal{S}[v_0]$, such that $f(w_k) \to 0$.

We fix $k_0 > 0$. Everytime the walk (X_n) crosses $\mathcal{S}[v_0]$ horizontally at a large enough height, it must cross the path $(w_k)_{k\geq 0}$. Since $\mathcal{C}$ is transient, if X crosses $\mathcal{S}[v_0]$ infinitely many times, it must cross $(w_k)_{k\geq k_0}$, and then hit $\partial\mathcal{S}[v_0]$. If this happens, let K be such that w_K is the first of the points $(w_k)_{k\geq k_0}$ to be hit by X (if none of these points is hit, we take $K = +\infty$). We have

$$\mathbb{P}\left(X \text{ hits } (w_k)_{k\geq k_0} \text{ and then } \partial\mathcal{S}[v_0]\right) = \mathbb{E}\left[\mathbb{1}_{K<+\infty} f(w_K)\right] \leq \mathbb{E}\left[\sup_{k\geq k_0} f(w_k)\right].$$

Since $f(w_k) \to 0$, by dominated convergence, this goes to 0, which proves that X cannot cross $\mathcal{S}[v_0]$ infinitely many times. This implies the almost sure convergence of X to a point X_∞ of $\widehat{\partial}T$.

The proof that X_∞ has full support is quite easy. Let $v_0 \in \mathbf{B}(T)$. Then (X_n) has a positive probability to visit the slice $\mathcal{S}[v_0]$ and, by Proposition 5.6, it a.s. has a positive probability to stay there ever after. But if $X_n \in \mathcal{S}[v_0]$ for n large enough, then X_∞ must correspond to a ray of descendants of v_0, so the distribution of X_∞ gives a positive mass to rays that are descendants of v_0. This is almost surely true for any $v_0 \in \mathbf{B}(T)$, so the distribution of X_∞ has a.s. full support.

Finally, to prove the almost sure nonatomicity, it is enough to prove that if X and Y are two independent simple random walks on $\mathcal{C}$, then $X_\infty \neq Y_\infty$ almost surely. The idea of the proof is that everytime X and Y reach a new height for the first time, by Proposition 5.6, they have a positive probability to get "swallowed" in two different slices of the form $\mathcal{S}[x]$ and $\mathcal{S}[y]$, so this will almost surely happen at some height.

More precisely, **in this proof and in this proof only**, until the end of Section 5.3.3, we

assume that T is a **non-conditioned** Galton-Watson tree with offspring distribution μ. We recall that Z_h is the number of vertices of T at height h. For every $h \geq 0$, let

$$\tau_h^X = \min\{n \geq 0 | h(X_n) = h\} \qquad \text{and} \qquad \tau_h^Y = \min\{n \geq 0 | h(Y_n) = h\}.$$

Note that if T survives, then $\tau_h^X, \tau_h^Y < +\infty$ for every h. Let also $\mathcal{F}_h$ be the σ-algebra generated by $B_h(\mathcal{C}(T))$, $(X_n)_{0 \leq n \leq \tau_h^X}$ and $(Y_n)_{0 \leq n \leq \tau_h^Y}$, and let $\mathcal{F}_\infty$ be the σ-algebra generated by $\bigcup_{h \geq 0} \mathcal{F}_h$. We note right now that $(\mathcal{F}_h)_{h \geq 0}$ is nondecreasing and that $\mathcal{F}_\infty$ is the σ-algebra generated by $(\mathcal{C}, X, Y)$. Finally, for every $h > 0$, let A_h be the event

$\Big\{$There are four distinct vertices $(x_i)_{1 \leq i \leq 4}$ of T at height h such that:
— the vertices x_1, x_2, x_3 and x_4 lie in this cyclic order,
— the trees $T[x_i]$ all survive,
— for every $n \geq \tau_h^X$, we have $X_n \in T[x_1]$
— for every $n \geq \tau_h^Y + 2$, we have $Y_n \in T[x_3]$.$\Big\}$.

Lemma 5.11. There is a constant $\delta > 0$ such that for every h, if $Z_h \geq 4$, then

$$\mathbb{P}(A_h | \mathcal{F}_h) \geq \delta.$$

Once this lemma is known, the end of the proof is quite easy: let $A = \bigcup_{h \geq 0} A_h$. If T survives, then $Z_h \geq 4$ for h large enough, so

$$\delta \leq \mathbb{P}(A_h | \mathcal{F}_h) \leq \mathbb{P}(A | \mathcal{F}_h) \xrightarrow[h \to +\infty]{a.s.} \mathbb{P}(A | \mathcal{F}_\infty) = \mathbb{1}_A,$$

by the martingale convergence theorem and the fact that $(\mathcal{C}, X, Y)$ is $\mathcal{F}_\infty$-measurable. Therefore, almost surely, if T survives, there is an h such that A_h occurs. But if it does, the slices $\mathcal{S}[x_2]$ and $\mathcal{S}[x_4]$ separate X and Y eventually, so they separate X_∞ from Y_∞, so $X_\infty \neq Y_\infty$, which ends the proof of Theorem 5.2. □

Remark 5.12. It is easy to show by using Proposition 5.6 that $\mathbb{P}(A_h | \mathcal{F}_h) > 0$ a.s.. However, this is not sufficient to prove Lemma 5.11, and this is actually in the proof of Lemma 5.11 that our argument fails in a more general setting. More precisely, in the setting of a tree $\mathbf{T}$ filled with i.i.d. strips, the lower degree of $X_{\tau_h^X}$ (i.e. the number of edges joining this vertex to a lower vertex) is not constant but depends on $\mathcal{F}_h$, and we might imagine that it goes to $+\infty$ as $h \to +\infty$. In this case, we might have $\mathbb{P}(A_h | \mathcal{F}_h) \to 0$ (with high probability, X goes back down right after τ_h^X). This problem does not occur in $\mathcal{C}(T)$, where the lower degree is always equal to 1, but it explains why the proof of Lemma 5.11 needs to be treated with some care.

Proof of Lemma 5.11. The proof will be split into three cases: the case where $X_{\tau_h^X} = Y_{\tau_h^Y}$, the case where $X_{\tau_h^X}$ and $Y_{\tau_h^Y}$ are distinct but neighbours, and the case where they are not neighbours. We treat carefully the first one, which is slightly more complicated than the others.

In the first case, we write $x_1 = X_{\tau_h^X} = Y_{\tau_h^Y}$. We also denote by x_2 and x_3 the two vertices at height h on the right of x_1, and by x_4 the left neighbour of x_1. Let also A'_h be the following event:

$\Big\{$The trees $T[x_i]$ for $1 \leq i \leq 4$ survive. Moreover, we have $X_n \in T[x_1]$ for every $n \geq \tau_h^X$ and
$$Y_{\tau_h^Y + 1} = x_2,\ Y_{\tau_h^Y + 2} = x_3 \text{ and } Y_n \in T[x_3] \text{ for every } n \geq \tau_h^Y + 2.\Big\}.$$

If A'_h occurs, then so does A_h. Moreover, we claim that the probability for A'_h to occur is independent of h and $\mathcal{F}_h$. The reason why this is true is that for every vertex v in one of these trees (say $T[x_1]$), the number of neighbours of v in $\mathcal{C}$ that are not in $\mathcal{S}[x_1]$ is fixed: there are 3

such neighbours if $v = x_1$, there is 1 such neighbour if $v \neq x_1$ is on the boundary of $\mathcal{S}[x_1]$ and 0 if it is not. Therefore, the probability given $\mathcal{C}$ that X stays in $\mathcal{S}[x_1]$ after time τ_h^X only depends on $T[x_1]$. Similarly, the probability for Y to perform the right first two steps after time τ_h^Y only depends on the numbers of children of x_1 and x_2, and the probability to stay in $T[x_3]$ ever after only depends on $T[x_3]$. All this is independent of h and $\mathcal{F}_h$, which proves our claim. If we write $\delta_1 = \mathbb{P}(A'_h|\mathcal{F}_h)$, we have $\delta_1 > 0$ by Proposition 5.6 and $\mathbb{P}(A_h|\mathcal{F}_h) \geq \delta_1$ for every h in this first case.

The other two cases can be treated similarly with minor adaptations in the choices of the vertices x_i, and the first two steps of Y after τ_h^Y. While we needed to control exactly the first two steps of Y in the first case, we only need one step for the second case and zero step for the third one. The other two cases yield two constants δ_2 and δ_3, which proves the lemma by taking $\delta = \min(\delta_1, \delta_2, \delta_3)$.

□

5.3.4 Robustness of Proposition 5.6

The goal of this subsection is to explain why Proposition 5.6 still holds in a quite general setting and to deduce the following result. We recall that a graph G is *Liouville* if its Poisson boundary is trivial, i.e. if every bounded harmonic function on G is constant.

Theorem 5.2 bis. Let $\mathbf{T}$ be a supercritical Galton–Watson tree with offspring distribution μ such that $\mu(0) = 0$, and let $(S_i)_{i\geq 0}$ be an i.i.d. sequence of random strips. We assume that (S_i) is also independent from $\mathbf{T}$.

1. Proposition 5.6 holds if we replace $\mathcal{S}$ by $\mathcal{S}(\mathbf{T}, (S_i)_{i\geq 0})$.
2. The map $\mathcal{M}(\mathbf{T}, (S_i))$ is a.s. non-Liouville.

Note that we cannot expect that $\widehat{\partial}\mathbf{T}$ is always the Poisson boundary of $\mathcal{M}(\mathbf{T}, (S_i))$. For example, if some strips S_i have a non-trivial Poisson boundary, then the Poisson boundary of $\mathcal{M}(\mathbf{T}, (S_i))$ is larger than $\widehat{\partial}\mathbf{T}$. See Section 5.5 for a more developed discussion.

Proof of the first point. Most of the proof works exactly along the same lines as the proof of Proposition 5.6, with $\mathbf{T}$ playing the same role as the backbone tree. In particular, we choose for (x_n) a nonbacktracking random walk on $\mathbf{T}$, and the sets A_k are built in the same way as in the original proof, but from $\mathbf{T}$ instead of $\mathbf{B}(T)$. The proof of Proposition 5.6 from Lemma 5.7 is very similar, as well as points (i) and (ii) of Lemma 5.7. The only difference is that the proof of point (ii) of Lemma 5.7 is easier in our new framework, because of the independence of the strips, and we do not need anymore to restrict ourselves to even values of k.

Exactly as in the first proof, by using the "self-similarity" property of $\mathcal{S}(\mathbf{T}, (S_i)_{i\geq 0})$, the proof of point (iii) of Lemma 5.7 can be reduced to the proof that the dual map of $\mathcal{S}(\mathbf{T}, (S_i)_{i\geq 0})$ is transient (it is also important that the sets A_k do not touch each other, which is why we have required that the boundaries of the strips are simple). The adaptation of the proof of Lemma 5.10 (transience of the dual slice), however, is not obvious.

More precisely, let $\mathcal{S}^*$ be the graph whose vertices are the finite faces of $\mathcal{S}(\mathbf{T}, (S_i)_{i\geq 0})$ and where, for every edge e of $\mathcal{S}(\mathbf{T}, (S_i)_{i\geq 0})$ that is adjacent to two finite faces, we draw an edge e^* between these two faces. We note right now that, since all the strips have only finite faces, the graph $\mathcal{S}^*$ is a connected graph, and we need to prove that it is transient. The idea of the proof is the following: we will build a genealogy on the set of strips, which contains a complete binary tree. As in the proof of Lemma 5.10, we will then kill the strips whose root is "too far" from its children, in order to preserve some quasi-isometry.

We first build a genealogy on the set of strips. We recall that the *root* of a strip S_i is the lowest vertex of its boundary, and is denoted by ρ_i. The *height* of S_i is the height of ρ_i. We

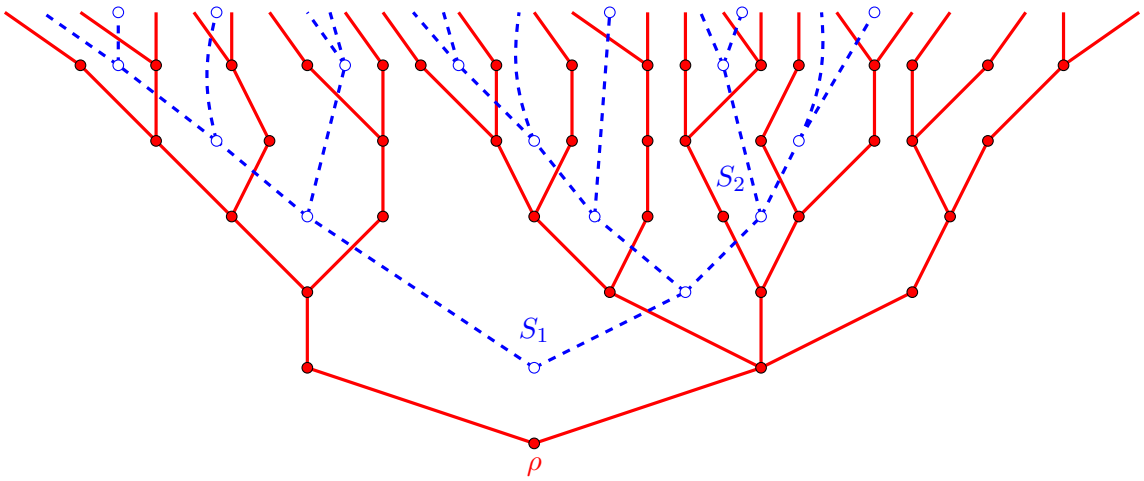

Figure 5.12 – The tree $\mathbf{T}$ (in red), and the tree T_{str} of descendants of the strip S_1 (in blue). Note that the strip S_2 has one parent on each side, but the right parent is not a descendant of S_1.

call two strips *adjacent* if their respective boundaries share at least one edge. If S_i is a strip, we consider the first vertex on the left boundary of S_i (apart from ρ_i) that is also a branching point of $\mathbf{T}$. This vertex is also the root of some strips, exactly one of which is adjacent to S_i. We call this strip the *left child* of S_i (cf. Figure 5.12). We can similarly define its *right child*. Note that almost surely, every branch of $\mathbf{T}$ branches eventually, so these childs always exist. We now fix a strip S_1. We claim that the restriction of this genealogy to the set of descendants of S_1 is encoded by a complete binary tree, which we denote by T_{str}. Indeed, all the descendants of the left child of a strip S lie on the left of S, whereas all the descendants of its right child lie on its right. Therefore, it is not possible to obtain the same strip by two different genealogical lines from S_1 (see Figure 5.12).

We now kill some of the strips. We fix a constant $\ell_{\max} > 0$. For every strip S_i, let e_i^ℓ (resp. e_i^r) be the first edge on the left (resp. right) boundary of S_i that is also adjacent to its left (resp. right) child. We call a strip S_i *good* if, for every face f of S_i that is adjacent to ρ_i, the dual of S_i contains a path of length at most $\ell_{\max}$ from f to e_i^ℓ, and similarly for e_i^r.

Note that the fact that S_i is good or not only depends on the internal geometry of S_i, and on the numbers of children of the vertices on the part of ∂S_i lying between e_i^ℓ and e_i^r. These parts for different values of $i \in T_{\text{str}}$ are disjoint. Hence, since the strips are i.i.d. and $\mathbf{T}$ is a Galton–Watson tree, the events

$$\{S_i \text{ is good}\}$$

for $i \in T_{\text{str}}$ are independent, and have the same probability. Therefore, removing from T_{str} all the strips that are not good is equivalent to performing a Bernoulli site percolation on the complete binary tree T_{str}. Moreover, the probability for a strip to be good goes to 1 as $\ell_{\max}$ go to $+\infty$, so we can find $\ell_{\max}$ such that this percolation is supercritical. We fix such an $\ell_{\max}$ until the end of the proof.

Let T'_{str} be an infinite connected component of T_{str} containing only good strips. Then T'_{str} is a supercritical Galton–Watson tree and survives, so it is transient. We can now define a submap $\mathcal{S}^*_{\text{bdd}}$ of $\mathcal{S}^*$. For every strip $S_i \in T'_{\text{str}}$, let p_i^ℓ (resp. p_i^r) be a dual path of length at most $\ell_{\max}$ joining the face of S_i that is adjacent to its parent to the face adjacent to e_i^ℓ (resp. e_i^r). Then the edges of $\mathcal{S}^*_{\text{bdd}}$ are the edges of these paths for all i such that $S_i \in T'_{\text{str}}$, as well as the dual edges of the edges e_i^ℓ and e_i^r (cf. Figure 5.13). The vertices of $\mathcal{S}^*_{\text{bdd}}$ are simply the vertices adjacent to these edges. Since the lengths of the paths p_i^ℓ and p_i^r are bounded by $\ell_{\max}$, it is easy to see

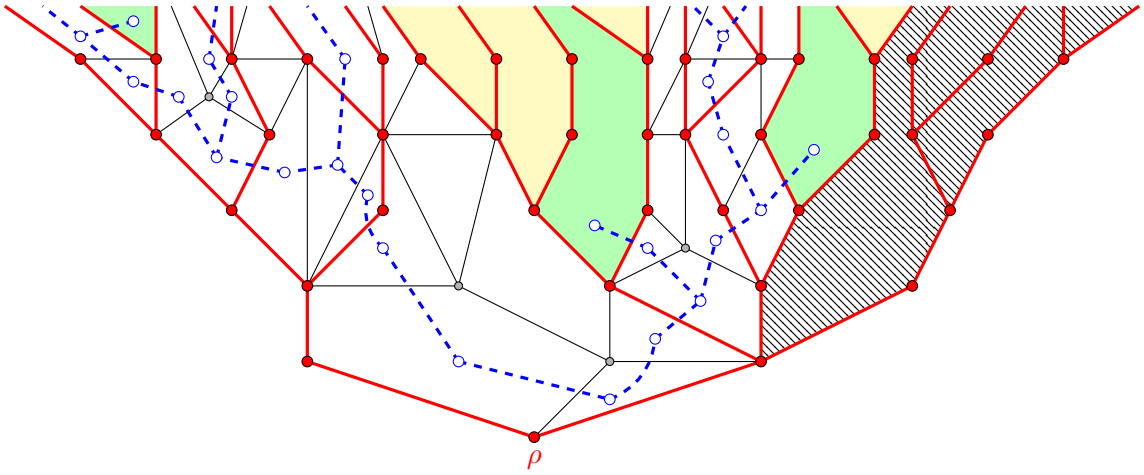

Figure 5.13 – Construction of $\mathcal{S}^*_{\text{bdd}}$ (in blue). The tree $\mathbf{T}$ is in red. The hatched strips do not belong to T_{str}. The green strips are the strips that are not good, and the yellow ones are the descendants of the green ones, so they are not in T'_{str}. For the sake of clarity, we have not drawn the interiors of the strips that do not belong to T'_{str}.

that $\mathcal{S}^*_{\text{bdd}}$ is quasi-isometric to the tree T'_{str}, so it is transient. This implies that $\mathcal{S}^*$ is transient as well, which concludes the proof of the first point.

□

Proof of point 2 of Theorem 5.2 bis. We mimic the beginning of the proof of Theorem 5.2, but we use the first point of Theorem 5.2 bis instead of Proposition 5.6. The proof of the first two points is robust, and shows that almost surely, the simple random walk X on $\mathcal{M}(\mathbf{T}, (S_i)_{i\geq 0})$ converges a.s. surely to a point X_∞ of $\partial\mathbf{T}$, and that the distribution of X_∞ has a.s. full support. In particular, for every $x \in \mathbf{T}$, let $\widehat{\partial}\mathbf{T}[y]$ be the set of the classes $\widehat{\gamma}$, where γ is a ray passing through y. Let y_1, y_2 be two vertices such that $\widehat{\partial}\mathbf{T}[y_1] \cap \widehat{\partial}\mathbf{T}[y_2] = \emptyset$. We define the function h on $\mathcal{M}$ by

$$h(x) = P_{\mathcal{M},x}\left(X_\infty \in \partial\mathbf{T}[y_1]\right).$$

Then h is harmonic and bounded on $\mathcal{M}$. Moreover, by the same argument as in the beginning of the proof of Theorem 5.2, there is a sequence (x_n) of vertices in $\mathcal{S}[y_1]$ such that $h(x_n) \to 1$. On the other hand, by the first point of Theorem 5.2 bis, there is a positive probability that X stays in $\mathcal{S}[y_2]$ eventually, so $h(x) < 1$ for every x, so h is non-constant. It follows that $\mathcal{M}$ is non-Liouville. □

5.4 Positive speed of the simple random walk

5.4.1 Sketch of the proof and definition of the half-plane model $\mathcal{H}$

The goal of this section is to prove Theorem 5.3. In all this section, we assume $\mu(0) = 0$. We first give a quick sketch of the proof. Our main task will be to prove positive speed in a half-planar model $\mathcal{H}$, constructed from an infinite forest of supercritical Galton–Watson trees, and where we have more vertical stationarity than in $\mathcal{C}$. The results of Section 5.3 will allow us to pass from $\mathcal{H}$ to $\mathcal{C}$. Note that Theorem 5.3 is very easy in the case where $\mu(0) = \mu(1) = 0$, since then the height of the simple random walk dominates a random walk on $\mathbb{Z}$ with positive drift. Therefore, we need to make sure that the vertices with only one child do not slow down the walk too much. Our proof in $\mathcal{H}$ relies on two ingredients:

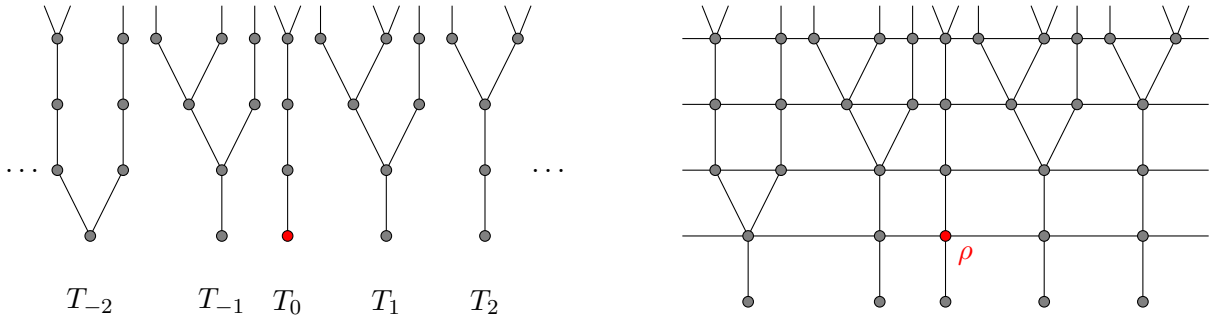

Figure 5.14 – On the left, an infinite forest $(T_i)_{i\in\mathbb{Z}}$. On the right, the half-planar map $\mathcal{H}$ obtained from this forest. The root vertex is in red.

— An exploration method of $\mathcal{H}$ will allow us to prove that the walk cannot be too far away from a point with at least two children. This will guarantee that the walk spends some time at vertices with at least two children. At these vertices, the height of the walk accumulates a positive drift. This will give a "quasi-positive speed" result: the height at time n is $n^{1-o(1)}$.

— Thanks to the stationarity properties of $\mathcal{H}$, we can study regeneration times, i.e. times at which the walk reaches some height for the first time, and stays above this height ever after. The estimates obtained in the first point are sufficient to prove that the first regeneration time has finite expectation. As in many other models (like random walks in random environments, see [130]), this is enough to ensure positive speed.

We now define our half-plane model. Let $(T_i)_{i\in\mathbb{Z}}$ be a family of i.i.d. Galton–Watson trees with offspring distribution μ. We draw the trees T_i with their roots on a horizontal line, and for every $h \geq 0$, we add horizontal connections between successive vertices of height h. Finally, for every vertex v of height 0, we add a parent of v at height -1, which is linked only to v. We root the obtained map at the root of T_0 (which has height 0), and denote it by $\mathcal{H}$ (see Figure 5.14). As for $\mathcal{C}$, if v is a vertex of $\mathcal{H}$, the height $h(v)$ of v in $\mathcal{H}$ is defined as its height in the forest (T_i). We denote by $\partial\mathcal{H}$ the set of vertices of $\mathcal{H}$ at height -1.

We first explain why it is enough to prove the result in $\mathcal{H}$ instead of $\mathcal{C}$. We denote by X^G the simple random walk on a graph G.

Lemma 5.13. If in $\mathcal{H}$ we have $\frac{1}{n}h(X_n^{\mathcal{H}}) \to v_\mu > 0$ a.s. as $n \to +\infty$, then Theorem 5.3 is true.

Proof. For every $i \in \mathbb{Z}$, let $\mathcal{S}(T_i)$ be the slice associated to the tree T_i. Then almost surely, if the walk (X_n) stays in $\mathcal{S}(T_i)$ eventually, the distance between X_n and the root of T_i is equivalent to $v_\mu n$. Since $\mathcal{S}(T_i)$ has the same distribution as $\mathcal{S}$, the simple random walk on $\mathcal{S}$ has a.s. speed v_μ on the event that it does not hit $\partial\mathcal{S}$.

Back in $\mathcal{C}$, by Theorem 5.2, we know that $X^{\mathcal{C}}$ converges to a point $X_\infty \in \widehat{\partial}T$ and that X_∞ is a.s. not the point of $\widehat{\partial}T$ corresponding to the leftmost and rightmost rays of T. Therefore, almost surely, the walk $X^{\mathcal{C}}$ only hits the boundary of $\mathcal{S}(T)$ finitely many times, so it has a.s. speed $v_\mu > 0$. □

5.4.2 An exploration method of $\mathcal{H}$

Let $k > 0$ and let x be a vertex of $\mathcal{H}$ at height $h \geq 0$. We write $x_0 = x$. Let also $x_{-1}, \dots, x_{-k}$ be the k first neighbours at height h on the left of x, and let $x_1, \dots, x_k$ be the k first neighbours of x on its right.

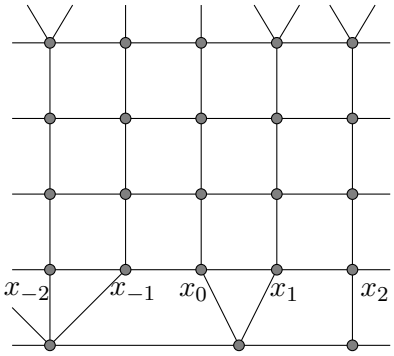

Figure 5.15 – An example for Definition 5.14: the point $x = x_0$ is 2-bad, but not 3-bad.

Definition 5.14. Let x be a vertex of $\mathcal{H}$ at height $h \geq 0$. We say that x is *k-bad* if all the descendants of $x_{-k}, x_{-(k-1)}, \ldots, x_k$ at heights $h, h+1, \ldots, h+k$ have only one child (cf. Figure 5.15).

The goal of this section is to show that the probability for the walk to visit a "very bad" point in its first n steps is very small. More precisely, we will prove the following result.

Lemma 5.15. There is a constant $c > 0$ such that for every $k, n > 0$, we have

$$\mathbb{P}\left(\text{one of the points } X_0, X_1, \ldots, X_n \text{ is } k\text{-bad}\right) \leq ck(n+1)^2 \mu(1)^{k^2}.$$

The presence of the factor $\mu(1)^{k^2}$ is not surprising. For example, the probability for the root vertex X_0 to be k-bad is exactly $\mu(1)^{(k+1)(2k+1)}$. The idea of the proof is to explore $\mathcal{H}$ at the same time as the random walk moves, in such a way that every time we discover a new vertex, either its k neighbours on the right or its k neighbours on the left have all their descendants undiscovered. If this is the case, the probability for the discovered vertex to be k-bad is at most $\mu(1)^{k^2}$. The factor $ck(n+1)^2$ means that our exploration method needs to explore at most $ck(n+1)^2$ vertices to discover $\{X_0, X_1, \ldots, X_n\}$.

The rest of Section 5.4.2 is devoted to the proof of Lemma 5.15. We first define our exploration method. We then state precisely the properties that we need this exploration to satisfy (Lemmas 5.16, 5.17 and 5.18), and explain how to conclude the proof from here. Finally, we prove these properties.

Exploration methods. By an *exploration method*, we mean a nondecreasing sequence $(E_i)_{i \geq 0}$ of finite sets of vertices of $\mathcal{H}$. For every $i \geq 0$, the part of $\mathcal{H}$ discovered at time i is the finite map formed by:

- the vertices of E_i,
- the edges of $\mathcal{H}$ whose two endpoints belong to E_i,
- for every edge of $\mathcal{H}$ with exactly one endpoint in E_i, the half-edge adjacent to E_i.

We denote this map by $\mathcal{E}_i$ (see Figure 5.16). In particular, when we explore a vertex, we know how many children it has, even if these children are yet undiscovered. Moreover, for every i, the map $\mathcal{E}_i$ will be equipped with an additional marked oriented edge or half-edge e_i describing the current position of the simple random walk. In all the explorations we will consider, we can pass from $\mathcal{E}_{i-1}$ to $\mathcal{E}_i$ by either adding one vertex to the explored set E_{i-1}, or moving the marked edge to a neighbour edge or half-edge. In the first case, we call i an *exploration step* and, in the second, we call i a *walk step*. In particular, the time of the exploration is not the same as the time of the random walk.

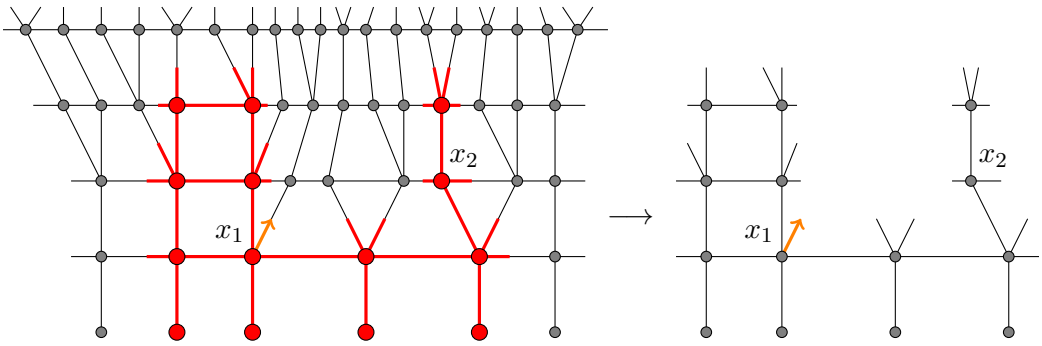

Figure 5.16 – On the left, the map $\mathcal{H}$ and the set of vertices E_i (in red). On the right, the explored map $\mathcal{E}_i$. The edge in orange means that at the current time, the random walk is leaving the vertex x_1 towards its right child. The vertex x_1 is 2-free on the right, but not 3-free. The vertex x_2 is k-free on the right for every $k \geq 0$. It is also 5-free on the left, but not 6-free. By only looking at $\mathcal{E}_i$, we can be sure that x_1 is 2-free on the right and x_2 is 4-free on the left, but not that x_2 is 5-free on the left.

We say that a vertex $v \in E_i$ is *k-free on the right* in E_i if none of the k first neighbours on the right of the rightmost child of E_i belongs to E_i. We define similarly a *k-free on the left* vertex. We would like to build an exploration of $\mathcal{H}$ such that at every exploration step i, the unique vertex of $E_{i+1} \backslash E_i$ is either k-free on the left or on the right in E_{i+1}. Note that it is not always possible by looking at $\mathcal{E}_i$ to decide whether a vertex $v \in E_i$ is k-free or not (cf. vertex v_2 on Figure 5.16). However, as we will see later (proof of Lemma 5.18), it is sometimes possible to be sure that a vertex is k-free.

Choice of the exploration method. The simplest exploration method coming into mind is to explore a vertex when it is hit for the first time by the walk (X_n). However, this is not suitable for our purpose. If for example we discover some descendants of a vertex v and explore v afterwards, then we have some partial information about the descendance of v, so we cannot control the probability for v to be k-bad. For this reason, we never want to discover a vertex before discovering all its ancestors. We will therefore require that the sets E_i are *stable*, which means that for every vertex $v \in E_i$, all the ancestors of v lie in E_i as well.

A second natural exploration method is now the following: everytime the walk (X_n) hits a vertex v for the first time, we discover all the ancestors of v that are yet undiscovered (including v), from the lowest to the highest. Although more convenient than the first one, this method is not sufficient either. Indeed, assume that at some point we explore a vertex v such that the k first neighbours on the left and on the right of v have already been discoverd, as well as many of their descendants. Then we have accumulated some partial information about the descendances of all the neighbours of v, so we cannot control the probability for v to be k-bad.

More generally, this problem occurs if we allow the formation of "narrow pits" in the explored part of $\mathcal{H}$. Therefore, the idea of our exploration is the following: everytime the walk (X_n) hits a vertex v for the first time, we explore all the ancestors of v that are yet undiscovered (including v), from the lowest to the highest. If by doing so we create a pit of width at most $2k$, then we explore completely the bottom of this pit, until it has width greater than $2k$.

To define this exploration more precisely, we need to define precisely a pit. We assume that the set E_i is stable, which implies that $\mathcal{E}_i$ is simply connected. Then there is a unique way to move around the map $\mathcal{E}_i$ from left to right by crossing all the half-edges exactly once, as on Figure 5.17. A *pit* is a sequence of consecutive half-edges $p_0, p_1, \ldots, p_{j+1}$ such that:

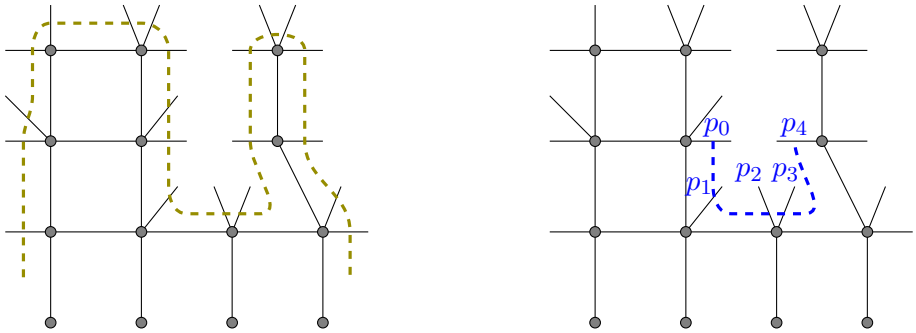

Figure 5.17 – On the left, we move around $\mathcal{E}_i$ by crossing all the half-edges. On the right, the unique pit of $\mathcal{E}_i$. It has width 3 and height 0. In particular, the map $\mathcal{E}_i$ is 1-flat, but not 2-flat.

— the half-edges $p_1, \dots, p_j$ point upwards, and start from the same height h,
— the half-edge p_0 points to the right and lies at height $h+1$,
— the half-edge p_{j+1} points to the left and lies at height $h+1$.

Note that all the pits are half-edge-disjoint. We call j and h the *width* and the *height* of the pit. Finally, we say that $\mathcal{E}_i$ is k-*flat* if it has no pit of width $j \leq 2k$ (see Figure 5.17).

If e is an oriented edge or half-edge, we will denote by e^- its starting point and e^+ its endpoint. We can now describe our exploration algorithm precisely. We take for E_0 the set formed by the root vertex ρ of $\mathcal{H}$ and its parent at height -1, and pick e_0 uniformly among all the edges and half-edges started from ρ. For every $i \geq 0$, we recall that e_i is the oriented edge or half-edge of $\mathcal{E}_i$ marking the position of the simple random walk. For every $i \geq 1$, given $(\mathcal{E}_{i-1}, e_{i-1})$, we construct $(\mathcal{E}_i, e_i)$ as follows.

(i) If the marked edge e_{i-1} is a full edge and $\mathcal{E}_{i-1}$ is k-flat, we perform a walk step: we set $E_i = E_{i-1}$ and pick e_i uniformly among all the edges and half-edges whose starting point is e_{i-1}^+.

(ii) If e_{i-1} is a half-edge, we perform an exploration step: we denote by v_i the lowest ancestor of e_{i-1}^+ that does not belong to E_{i-1}. If v_i lies at height -1, then E_i is the union of E_{i-1}, the vertex v_i and its child at height 0 (this is the only case where we explore two vertices at once, to make sure that $\mathcal{E}_i$ remains connected). If not, then $E_i = E_{i-1} \cup \{v_i\}$. Note that if $v_i = e_{i-1}^+$, then the marked half-edge becomes a full edge.

(iii) If e_{i-1} is a full edge but $\mathcal{E}_{i-1}$ is not k-flat, let $(p_0, p_1, \dots, p_{j+1})$ be the leftmost pit of width at most $2k$. We then take $E_i = E_{i-1} \cup \{p_1^+\}$. This means that we explore the endpoint of the leftmost vertical half-edge of the pit.

It is easy to check that the sets E_i we just defined are all stable.

The exploration is Markovian. An important feature of our exploration that we need to check is that it is Markovian. More precisely, for every $i \geq 0$, let ∂E_i be the set of vertices consisting of:

— the endpoints of the half-edges of $\mathcal{E}_i$ pointing upwards,
— the vertices of $\mathcal{H}$ of height 0 that do not belong to E_i.

Lemma 5.16. Conditionally on $(\mathcal{E}_j, e_j)_{0 \leq j \leq i}$, the trees of descendants of the vertices of ∂E_i are independent Galton–Watson trees with offspring distribution μ.

Proof. Given the Markovian structure of supercritical Galton–Watson tree, it is enough to check that at every step, our exploration is independent of the part of $\mathcal{H}$ that has not yet been

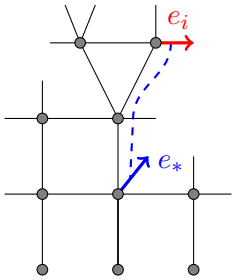

Figure 5.18 – The second case of our exploration algorithm. We move around the boundary of $\mathcal{E}_i$ towards the right: we encounter one horizontal half-edge, and then a vertical half-edge e_*. This e_* has descendents at the same height as e_i, and the leftmost such descendant is e_i^+.

discovered. This is easy in the case (i): conditionally on $\mathcal{E}_i$, the choice of e_i is independent of the rest of $\mathcal{H}$. We claim that in the other two cases, the discovered vertex is either the endpoint of a vertical half-edge which is a deterministic function of $(\mathcal{E}_i, e_i)$, or the root of the first infinite tree on the left or on the right of $\mathcal{E}_i$. This claim is obvious in the case (iii).

In the case (ii), the half-edge e_i points either to the top, the left or the right. If it points to the top, then all the ancestors of e_i^+ have been discovered, so the explored vertex is e_i^+. If e_i does not point to the top, we assume without loss of generality that it points to the right. We then start from the half-edge e_i and move around $\mathcal{E}_i$ towards the right. We first cross horizontal half-edges pointing to the right at decreasing heights, until either we reach height 0, or we cross a vertical half-edge pointing to the top. If we cross a first vertical half-edge e, then e must be an ancestor of e_i^+ (indeed, e has descendants at the same height as e_i^+, and there is no other half-edge pointing to the top between e_i^- and e_*^-, see Figure 5.18). Therefore, the explored vertex must be e_*^+. Finally, if we reach the bottom boundary, then the explored vertex is the root of the first tree on the right of $\mathcal{E}_i$ (and its parent at height -1).

□

We denote by $\varphi(n)$ the n-th walk step of our exploration, with $\varphi(0) = 0$. Note that the marked edge or half-edge from the time $\varphi(n)$ to the time $\varphi(n+1) - 1$ corresponds to the edge in $\mathcal{H}$ from X_n to X_{n+1}. Therefore, at time $\varphi(n)$, the explored part covers $\{X_0, X_1, \dots, X_n\}$. As explained in the beginning of this subsection, it is important to control φ. We can now state the two important properties that our exploration satisfies.

Lemma 5.17. There is a deterministic constant c such that for every $n \geq 0$, we have

$$\varphi(n+1) - \varphi(n) \leq ck(n+1).$$

Lemma 5.18. For every exploration step i, the unique vertex v_i of $E_i \backslash E_{i-1}$ that does not lie at height -1 is either k-free on the left or k-free on the right.

Note that our exploration method was precisely designed to satisfy Lemma 5.18. Both these lemmas are completely deterministic: they hold if we replace $\mathcal{H}$ by any infinite causal map with no leaf, and (X_n) by any infinite path. Before proving them, we explain how to conclude the proof of Lemma 5.15 given these two results.

Proof of Lemma 5.15 given Lemmas 5.17 and 5.18. For every $i \geq 0$, let $\mathcal{F}_i$ be the σ-algebra generated by $(\mathcal{E}_j, e_j)_{0 \leq j \leq i}$. We first show that, for every exploration step $i \geq 0$, we have

$$\mathbb{P}\left(v_i \text{ is } k\text{-bad} | \mathcal{F}_i\right) \leq \mu(1)^{k^2}. \tag{5.7}$$

Let $i \geq 0$ be an exploration step. By Lemma 5.18, without loss of generality, we may assume that v_i is k-free on the right. Let $v_i^1, \dots, v_i^k$ be the k first neighbours on the right of the rightmost child of v_i. Since v_i is k-free on the right, these vertices do not belong to E_i. By Lemma 5.16, conditionally on $\mathcal{F}_i$, their trees of descendants are i.i.d. Galton–Watson trees with offspring distribution μ. If the vertex v_i is k-bad, each of the vertices $v_i^1, \dots, v_i^k$ has only one descendant at height $h + k + 1$, so we have

$$\begin{aligned}
\mathbb{P}\left(v_i \text{ is } k\text{-bad} | \mathcal{F}_i\right) &\leq \mathbb{P}\left(v_i^1, \dots, v_i^k \text{ have one descendant each at height } h+k+1 | \mathcal{F}_i\right) \\
&= \prod_{j=1}^{k} \mathbb{P}\left(v_i^j \text{ has exactly one descendant at height } h+k+1 | \mathcal{F}_i\right) \\
&= \left(\mu(1)^k\right)^k,
\end{aligned}$$

and we obtain (5.7), which implies

$$\mathbb{P}\left(i \text{ is an exploration step and } v_i \text{ is } k\text{-bad}\right) \leq \mu(1)^{k^2}$$

for every $i \geq 0$. By summing Lemma 5.17, we obtain $\varphi(n+1) \leq ck(n+1)^2$. Therefore, the vertices $X_0, X_1, \dots, X_n$ all lie in $\mathcal{E}_{ck(n+1)^2}$. If one of these points is k-bad, it cannot lie at height -1 by definition of a bad point, so it is equal to v_i for some $0 \leq i \leq ck(n+1)^2$. Therefore, we have

$$\begin{aligned}
\mathbb{P}\left(\text{one of the points } X_0, X_1, \dots, X_n \text{ is } k\text{-bad}\right) &\leq \sum_{i=0}^{ck(n+1)^2} \mathbb{P}\left(v_i \text{ is } k\text{-bad}\right) \\
&\leq ck(n+1)^2 \mu(1)^{k^2},
\end{aligned}$$

which ends the proof. □

Proof of Lemma 5.17. To bound $\varphi(n+1) - \varphi(n)$, we describe precisely what happens between the times $\varphi(n)$ and $\varphi(n+1)$. Note that $\mathcal{E}_{\varphi(n)} = \mathcal{E}_{\varphi(n)-1}$ is k-flat (if it was not, $\varphi(n)$ would have to be an exploration step). Hence, if $e_{\varphi(n)}$ is a full edge, then $\varphi(n)+1$ is a walk step and we have $\varphi(n+1) = \varphi(n)+1$, so it is only necessary to treat the case where $e_{\varphi(n)}$ is a half-edge.

In this case, the first thing our algorithm does is to explore the vertex $e^+_{\varphi(n)}$ and all its undiscovered ancestors. We know that $e^+_{\varphi(n)} = X_{n+1}$, so in particular its height is at most $n+1$. Therefore, exploring all its ancestors takes at most $n+2$ steps.

We can now perform the $(n+1)$-th walk step, except if exploring $e^+_{\varphi(n)}$ and its ancestors has created a new pit of width at most $2k$. This can happen in two different ways, as on Figure 5.19:

— if $e_{\varphi(n)}$ is vertical, exploring $e^+_{\varphi(n)}$ may split an existing pit in two,

— if $e_{\varphi(n)}$ is horizontal (say it points to the right), exploring its ancestors may decrease the width of an existing pit on the right of $e_{\varphi(n)}$.

Note that in the first case, we can create at most two new pits, whereas in the second, we can shrink only one (cf. Figure 5.19). Hence, we will create at most 2 narrow pits, at the same height h.

In the first case (top of Figure 5.19), our algorithm will then fill the pit on the left if it has width at most $2k$, and then the pit on the right. The number of steps this takes is at most $2 \times 2k = 4k$. By doing so, we may create a new pit at height $h+1$. However, since the trees we work with have no leaf, this pit is at least as wide as the pit of $\mathcal{E}_{\varphi(n)}$ in which $e_{\varphi(n)}$ lies. Hence, the new pit at height $h+1$ has width greater than $2k$, and does not need to be filled. Therefore, in the first case, the number of exploration steps needed to fill all the narrow pits is at most $4k$.

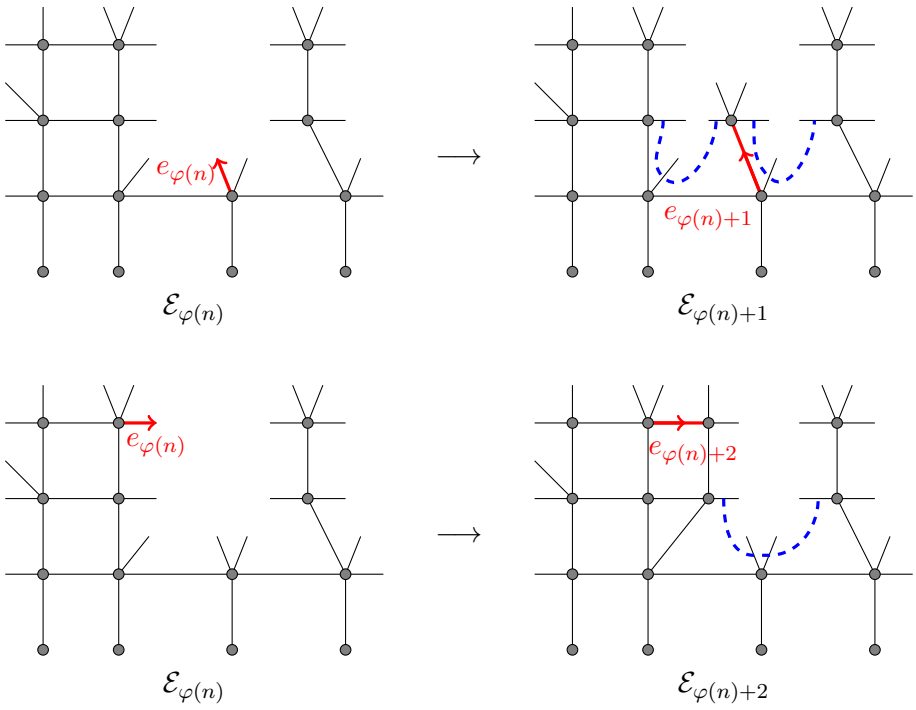

Figure 5.19 – The two ways our exploration can create new pits between two walk steps. On the top, a pit of width 3 is split into two pits of width 1. On the bottom, the width of a pit decreases from 3 to 2. The new pits of width at most 2 are indicated in blue.

In the second case (bottom of Figure 5.19), if the pit has width at most $2k$, our algorithm will explore all the vertical half-edges of this pit, from left to right. This takes at most $2k$ steps. Once again, this creates a new pit at height $h+1$, but this time this pit may have width $2k$ or less. If this is the case, our algorithm will explore all its half-edges and perhaps create a pit at height $h+2$, and so on. Note that the maximal height of $\mathcal{E}_{\varphi(n)}$ is at most n. Indeed, the only times at which this maximal height increases is when the random walk reaches some height for the first time, so the maximal height of $\mathcal{E}_{\varphi(n)}$ cannot be larger than the maximal height of the random walk during its first n steps. Therefore, we will need to fill at most n pits (at heights between 0 and $n-1$), each of which taking at most $2k$ steps. Hence, filling the narrow pits takes at most $2kn$ steps.

Therefore, in both cases, the number of steps needed to obtain a k-flat map and perform a new walk step is bounded by $\max(4k, 2kn)$. If we add the number of steps needed to explore the ancestors of $e^+_{\varphi(n)}$ and the walk step $\varphi(n+1)$, we obtain

$$\varphi(n+1) - \varphi(n) \leq \max(4k, 2kn) + (n+2) + 1 \leq ck(n+1),$$

with e.g. $c = 7$. □

We finally prove Lemma 5.18. The proof will make use of the "step by step" description of our exploration that we also used in the last proof. We recall that for every exploration step i, we call v_i the unique vertex of nonnegative height in $E_i \backslash E_{i-1}$, and e_i the oriented edge or half-edge marking the current position of the random walk.

Proof of Lemma 5.18. We fix an exploration step $i \geq 0$. Note that the vertex v_i is always the endpoint of some half-edge of $\mathcal{E}_{i-1}$, that we denote by e_*.

Before moving on to the details of the proof, we explain how it is possible, by only looking at the map $\mathcal{E}_{i-1}$, to be sure that the vertex v_i is k-free on the right in E_i. We move along the boundary of $\mathcal{E}_{i-1}$ from e_* towards the right, and stop when we encounter a vertex of height $h(v_i)+1$. If this never occurs, it means that in $\mathcal{E}_{i-1}$ and $\mathcal{E}_i$, there is no vertex at height $h(v_i)+1$ on the right of v_i, so v_i is k-free on the right. If this occurs, assume that by moving so, we cross at least k vertical half-edges. Since the trees we consider have no leaf, all these vertical half-edges have descendants at height $h(v_i)+1$, which lie on the right of all the children of v_i. Moreover, none of these descendants belongs to $\mathcal{E}_i$. Therefore, v_i must be k-free on the right in $\mathcal{E}_i$. Of course, this is also true for k-free on the left vertices (see the end of the caption of Figure 5.16 for an example). This remark will be implicitly used in all the cases below.

Let n be the integer such that $\varphi(n) < i < \varphi(n+1)$. We distinguish two cases, corresponding to the two "phases" of exploration between $\varphi(n)$ and $\varphi(n+1)$ that we described in the proof of Lemma 5.17. Both of these cases will be separated in a few subcases.

- We first treat the case where e_{i-1} is a half-edge, so the explored vertex v_i is an ancestor of e_{i-1}^+.
  - We start with the subcase where e_{i-1} is vertical, and lies in a pit p. Since e_{i-1} is vertical, exploring the ancestors of e_{i-1}^+ takes only one step, so $i = \varphi(n)+1$ and $v_i = e_{i-1}^+$. Since $i-1$ is a walk step, the pit p has width at least $2k+1$. Without loss of generality, we may assume that at least k of the vertical half-edges of p are on the right of e_{i-1}, so v_i is k-free on the right.
  - If e_{i-1} is vertical but is not in a pit, as in the previous case, we have $i = \varphi(n)+1$ and v_i is the endpoint of e_{i-1}. Moreover, there is a direction (left or right) such that when we start from e_{i-1} and move along the boundary of $\mathcal{E}_{i-1}$ in this direction, the height decreases before increasing for the first time (if not, e_{i-1} would be in a pit). Without loss of generality, this direction is the right. Let p be the first pit that we encounter on the right of e_{i-1}. If p does not exist, it means that the height never increases again, so there is no vertex on the right of e_{i-1} in $\mathcal{E}_{i-1}$ that is higher than e_{i-1}^-. Therefore, the vertex $v_i = e_{i-1}^+$ is k-free on the right in $\mathcal{E}_i$. If p exists, it has width at least $2k+1$, so v_i is $(2k+1)$-free on the right, and in particular k-free.
  - If e_{i-1} is horizontal, without loss of generality it points to the right. As explained earlier (see Figure 5.18), the edge e_* is in this case the first vertical half-edge we meet when we move around $\mathcal{E}_{i-1}$ from e_{i-1} towards the right (except if no such vertical half-edge exists, in which case v_i is the root of a new tree on the right of $\mathcal{E}_{i-1}$, and v_i is obviously k-free on the right). By the same argument as in the previous case, if there is no pit on the right of e_*, then v_i is k-free on the right (and even ∞-free). If there is one and p is the first such pit, note that between the times $\varphi(n)$ and i, the pit p has either been untouched, or has been shrunk by 1. Therefore, at time i it has width at least $2k$, so v_i is k-free in E_i.
- We now consider the second "phase", i.e. the case where $\mathcal{E}_{i-1}$ has a pit of width at most $2k$, and the goal of the exploration step i is to fill it.
  - If e_{i-1} points to the top, then the pit has been created at time $\varphi(n)+1$ as in the top part of Figure 5.19. Hence, the half-edge e_* belonged at time $\varphi(n)$ to a pit $(p_0, p_1, \dots, p_{j+1})$ of height h and width $j > 2k$. Therefore, either k of the vertical half-edges $p_1, \dots, p_j$ lie on the left of e_*, or k of them lie on its right (the two cases are not symmetric since the pit is filled from left to right). If k of these half-edges lie on the left of e_*, then their k endpoints (of height $h+1$) have been explored before v_i, but none of the descendants of these endpoints has been discovered. Therefore, the map $\mathcal{E}_i$ contains at least k vertical half-edges at height $h+1$ on the left of v_i, so v_i is k-free on the left. This case is the reason why, in the definition of a k-free vertex,

we asked the neighbours of the *children* of v to be undiscovered, and not simply the neighbours of v. If k of the half-edges of p lie on the right of e_*, the argument is similar (it is actually simpler since the half-edges on the right of e_* have not yet been explored).

— If e_{i-1} points to the right, then we are in the bottom case of Figure 5.19: a pit p of width $2k+1$ has been shrunk to width $2k$ during the first phase, resulting in a pit $(p_0, p_1, \dots, p_{2k+1})$ of width $2k$ at some height h. Let also $h' \geq h$ be the height of e_*. Since the pit is filled layer by layer from the bottom, the half-edge e_* must be a descendant of a half-edge p_{ℓ_0} with $1 \leq \ell_0 \leq 2k$. Moreover, our algorithm fills the layers from left to right. Therefore, at time i, for every $1 \leq \ell \leq 2k$, we have already explored the descendants of e_ℓ up to height $h'+1$ if $\ell \leq \ell_0$ and up to height h' if $\ell > \ell_0$. But the vertex v_i lies at height $h'+1$. Therefore, if $\ell_0 \geq k+1$, then v_i has k vertical half-edges on its left so it is k-free on the left. On the other hand, if $\ell_0 \leq k$, then v_i has k vertical half-edges on its right at height h, so it is k-free on the right. This concludes the proof.

□

5.4.3 Quasi-positive speed in $\mathcal{H}$

The goal of this subsection is to use Lemma 5.15 to prove that the walk has a speed $n^{1-o(1)}$, which is slightly weaker than positive speed. We will need to "bootstrap" this result in Section 5.4.4 to obtain positive speed. We denote by H_n the height of X_n. We also write

$$D_n = \max\{H_k - H_\ell | 0 \leq k < \ell \leq n\}$$

for the greatest "descent" of X before time n.

Proposition 5.19. Let $0 < \delta < 1$ and $\beta > 0$. Then we have

$$\mathbb{P}\left(H_n \leq n^{1-\delta}\right) = o\left(n^{-\beta}\right) \quad \text{and} \quad \mathbb{P}\left(D_n \geq n^{\delta}\right) = o\left(n^{-\beta}\right)$$

as $n \to +\infty$.

Proof. We start with the proof of the first estimate, the proof of the second will follow the same lines. We call a vertex of $\mathcal{H}$ *good* if it has height -1 or if it has at least two children. The idea of the proof is the following: by Lemma 5.15, with high probability, the walk does not visit any k-bad point before time n for some k. Hence, it is never too far from a good point. Therefore, the walk always has a reasonable probability to reach a good point in a near future. It follows that X will visit many good points, so (H_n) will accumulate a large positive drift.

More precisely, let $a > 0$ (we will take a large later). We define two events formalizing the ideas we just explained:

$$\begin{aligned} A_1 &= \{\text{none of the vertices } X_0, X_1, \dots, X_n \text{ is } a\sqrt{\log n}\text{-bad}\}, \\ A_2 &= \{\text{for every } 0 \leq m \leq n - n^{\delta/2}, \text{ one of the points } X_m, X_{m+1}, \dots, X_{m+n^{\delta/2}} \text{ is good}\}. \end{aligned}$$

Then we have

$$\mathbb{P}\left(H_n \leq n^{1-\delta}\right) \leq \mathbb{P}(A_1^c) + \mathbb{P}(A_1 \backslash A_2) + \mathbb{P}\left(A_2 \cap \{H_n \leq n^{1-\delta}\}\right). \tag{5.8}$$

We start with the first term. By Lemma 5.15, we have

$$\mathbb{P}(A_1^c) \leq ca\sqrt{\log n}\,(n+1)^2\,\mu(1)^{a^2 \log n}.$$

Hence, if we choose a large enough (i.e. $a^2 > \frac{\beta+2}{-\log \mu(1)}$), we have $\mathbb{P}(A_1^c) = o(n^{-\beta})$.

We now bound the second term of (5.8). For every $0 \leq m \leq n$, let $\mathcal{F}_m$ be the σ-algebra generated by $\mathcal{H}$ and $(X_0, X_1, \ldots, X_m)$. If X_m is not $a\sqrt{\log n}$-bad (which is an $\mathcal{F}_m$-measurable event), let Y be the closest good vertex from X_m (we may have $Y = X_m$). We have $d_{\mathcal{H}}(X_m, Y) \leq 2a\sqrt{\log n}$, so there is a path from X_m to Y of length at most $2a\sqrt{\log n}$, and visiting only vertices of degree 4 (except of course Y). Therefore, we have

$$\mathbb{P}\left(X \text{ visits the vertex } Y \text{ between time } m \text{ and time } m + 2a\sqrt{\log n} | \mathcal{F}_m\right) \geq \left(\frac{1}{4}\right)^{2a\sqrt{\log n}}$$

if X_m is not $a\sqrt{\log n}$-bad. By induction on i, we easily obtain, for every $i \geq 0$,

$$\begin{aligned}\mathbb{P}\left(X_m, X_{m+1}, \ldots, X_{m+2ia\sqrt{\log n}} \text{ are neither good nor } a\sqrt{\log n}\text{-bad}\right) &\leq \left(1 - \frac{1}{4^{2a\sqrt{\log n}}}\right)^i \\ &\leq \exp\left(-\frac{i}{4^{2a\sqrt{\log n}}}\right).\end{aligned}$$

In particular, by taking $i = \frac{n^{\delta/2}}{2a\sqrt{\log n}}$, we obtain, for every m:

$$\mathbb{P}\left(X_m, X_{m+1}, \ldots, X_{m+n^{\delta/2}} \text{ are neither good nor } a\sqrt{\log n}\text{-bad}\right) \leq \exp\left(-\frac{n^{\delta/2}}{2a\sqrt{\log n}\, 4^{2a\sqrt{\log n}}}\right).$$

If the event $A_1 \backslash A_2$ occurs, then there is an m with $0 \leq m \leq n - n^{\delta/2}$ such that the above event occurs. Therefore, by summing the last equation over $0 \leq m \leq n - n^{\delta/2}$, we obtain

$$\mathbb{P}(A_1 \backslash A_2) \leq n \exp\left(-\frac{n^{\delta/2}}{2a\sqrt{\log n}\, 4^{2a\sqrt{\log n}}}\right) = o(n^{-\beta}).$$

Finally, we bound the third term of (5.8) by the Azuma inequality. For every $n \geq 0$, let

$$M_n = H_n - \sum_{i=0}^{n-1} E_{\mathcal{H}}\left[H_{i+1} - H_i | X_0, X_1, \ldots, X_i\right].$$

It is clear that M is a martingale with $|M_{n+1} - M_n| \leq 2$ for every n, and $M_0 = 0$. Moreover, we have

$$E_{\mathcal{H}}\left[H_{i+1} - H_i | X_0, X_1, \ldots, X_i\right] = \frac{c(X_i) - 1}{c(X_i) + 3} \mathbb{1}_{X_i \notin \partial\mathcal{H}} + \mathbb{1}_{X_i \in \partial\mathcal{H}},$$

where we recall that $c(v)$ is the number of children of a vertex v. In particular, we have $E_{\mathcal{H}}\left[H_{i+1} - H_i | X_0, X_1, \ldots, X_i\right] \geq 0$, and $E_{\mathcal{H}}\left[H_{i+1} - H_i | X_0, X_1, \ldots, X_i\right] \geq \frac{1}{5}$ if X_i is a good vertex. If A_2 occurs, the walk X must visit at least $n^{1-\delta/2}$ good vertices before time n, so we have

$$\sum_{i=0}^{n-1} E_{\mathcal{H}}\left[H_{i+1} - H_i | X_0, X_1, \ldots, X_i\right] \geq \frac{1}{5} n^{1-\delta/2}.$$

Therefore, if the event in the third term of (5.8) occurs, we have

$$M_n \leq n^{1-\delta} - \frac{1}{5} n^{1-\delta/2} < 0.$$

On the other hand, the Azuma inequality applied to M gives

$$P_{\mathcal{H},\rho}\left(M_n \leq n^{1-\delta} - \frac{1}{5} n^{1-\delta/2}\right) \leq \exp\left(-\frac{1}{8n}\left(\frac{1}{5} n^{1-\delta/2} - n^{1-\delta}\right)^2\right),$$

so

$$\mathbb{P}\left(M_n \leq n^{1-\delta} - \frac{1}{5}n^{1-\delta/2}\right) = o(n^{-\beta}),$$

which bounds the third term of (5.8), and proves the first part of Proposition 5.19.

To prove the second part, we decompose the event $\{D_n \geq n^\delta\}$ in the same way as in (5.8). By the definition of D_n, it is enough to show

$$\max_{0\leq k\leq \ell\leq n} \mathbb{P}\left(A_2 \cap \{H_\ell - H_k \leq -n^\delta\}\right) = o(n^{-(\beta+2)}), \tag{5.9}$$

and then to sum over k and ℓ. To prove (5.9), note that if $\ell < k + n^\delta$, then $H_\ell - H_k > -n^\delta$ deterministically. If $\ell \geq k + n^\delta$, we use the same argument based on the Azuma inequality as for the first part. Let $0 \leq k < k + n^\delta \leq \ell \leq n$. If A_2 occurs, then X visits at least $\frac{\ell-k}{n^{\delta/2}}$ good vertices between times k and ℓ, so

$$\sum_{i=k}^{\ell-1} E_{\mathcal{H}}[H_{i+1} - H_i | X_0, \ldots, X_i] \geq \frac{1}{5}\frac{\ell-k}{n^{\delta/2}}.$$

Hence, if the event of (5.9) occurs for k and ℓ, we have

$$M_\ell - M_k \leq H_\ell - H_k - \frac{1}{5}\frac{\ell-k}{n^{\delta/2}} \leq -n^\delta - \frac{1}{5}\frac{\ell-k}{n^{\delta/2}} \leq -\frac{2}{\sqrt{5}}(\ell-k)^{1/2}n^{\delta/4}.$$

But the Azuma inequality gives

$$\begin{aligned}
\mathbb{P}\left(M_\ell - M_k \leq -\frac{2}{\sqrt{5}}(\ell-k)^{1/2}n^{\delta/4}\right) &\leq \exp\left(-\frac{(2(\ell-k)^{1/2}n^{\delta/4})^2}{8\times 5(\ell-k)}\right) \\
&= \exp\left(-\frac{n^{\delta/2}}{10}\right) \\
&= o(n^{-(\beta+2)}),
\end{aligned}$$

which proves (5.9) and the second point of Proposition 5.19. □

5.4.4 Positive speed in $\mathcal{H}$ via regeneration times

For every $0 \leq h < h' \leq +\infty$, we denote by $\mathcal{B}_{h,h'}$ the map formed by the vertices of $\mathcal{H}$ with height in $\{h, h+1, \ldots, h'\}$, in which for every vertex v at height h, we have added a vertex below v that is linked only to v. We root $\mathcal{B}_{h,h'}$ at the vertex ρ_h corresponding to the leftmost descendant of ρ at generation h. The *height* of a vertex in $\mathcal{B}_{h,h'}$ is its height in $\mathcal{H}$, minus h, and the height of the additional vertices is -1. We denote by $\partial\mathcal{B}_{h,h'}$ the set of these additional vertices. Note that for any $h \geq 0$, the rooted map $(\mathcal{B}_{h,\infty}, \rho_h)$ is independent of $\mathcal{B}_{0,h}$ and has the same distribution as $(\mathcal{H}, \rho)$. Since this distribution is invariant by horizontal root translation, this is still true for any choice of the root vertex of $\mathcal{B}_{h,h'}$ at height 0, as long as the choice of the root is independent of $\mathcal{B}_{h,\infty}$.

Definition 5.20. We say that $n > 0$ is a *regeneration time* if $H_i < H_n$ for every $i < n$, and $H_i \geq H_n$ for every $i \geq n$. We denote by $\tau^1 < \tau^2 < \ldots$ the list of regeneration times in increasing order.

We also denote by T_∂ the first time at which the simple random walk X on $\mathcal{H}$ hits $\partial\mathcal{H}$. The key of the proof of Theorem 5.3 will be to combine the two following results.

Proposition 5.21. We have $\mathbb{E}\left[\tau^1\right] < +\infty$. In particular, $\tau^1 < +\infty$ a.s..

Lemma 5.22. 1. Almost surely, $\tau^j < +\infty$ for every $j \geq 1$.

2. The path-decorated maps $\left(\mathcal{B}_{H_{\tau^j}, H_{\tau^{j+1}}}, (X_{\tau^j+i})_{0\leq i\leq \tau^{j+1}-\tau^j}\right)$ for $j \geq 1$ are i.i.d. and have the same distribution as $\left(\mathcal{B}_{0,H_{\tau^1}}, (X_i)_{0\leq i\leq \tau^1}\right)$ conditioned on $\{T_\partial = +\infty\}$.
3. In particular, the pairs $\left(\tau^{j+1} - \tau^j, H_{\tau^{j+1}} - H_{\tau^j}\right)$ for $j \geq 1$ are i.i.d. and have the same distribution as (τ^1, H_{τ^1}) conditioned on $\{T_\partial = +\infty\}$.

Proposition 5.21 will be deduced from the results of Sections 5.4.2 and 5.4.3. On the other hand, Lemma 5.22 is the reason why regeneration times have been used to prove positive speed for many other models. The same property has already been observed and used in various contexts such as random walks in random environments [130], or biased random walks on Galton–Watson trees [109]. Although the proof is basically the same for our model, we write it formally in Appendix 5.A.

Finally, we note that the finiteness of the times τ^i could be deduced directly from the results of Section 5.3, even in the case $\mu(0) > 0$. However, this is not sufficient to ensure positive speed.

We now explain how to conclude the proof of Theorem 5.3 from the last two results.

Proof of Theorem 5.3 given Proposition 5.21 and Lemma 5.22. By Lemma 5.13, it is enough to prove the result on $\mathcal{H}$. By item 3 of Lemma 5.22 and Proposition 5.21, we have

$$\mathbb{E}\left[\tau^2 - \tau^1\right] = \frac{\mathbb{E}[\tau^1 \mathbb{1}_{\forall n\geq 0,\, H_n\geq 0}]}{\mathbb{P}(\forall n \geq 0,\, H_n \geq 0)} < +\infty.$$

Moreover, $H_{\tau^2} - H_{\tau^1} \leq \tau^2 - \tau^1$, so $\mathbb{E}\left[H_{\tau^2} - H_{\tau^1}\right] < +\infty$ as well. By Lemma 5.22 and the law of large numbers, we have

$$\frac{\tau^j}{j} \xrightarrow[j\to+\infty]{a.s.} \mathbb{E}\left[\tau^2 - \tau^1\right] \qquad \text{and} \qquad \frac{H_{\tau^j}}{j} \xrightarrow[j\to+\infty]{a.s.} \mathbb{E}\left[H_{\tau^2} - H_{\tau^1}\right].$$

For every $n > \tau^1$, let $j(n)$ be the index such that $\tau^{j(n)} \leq n < \tau^{j(n)+1}$. Then we have $\frac{j(n)}{n} \to \mathbb{E}[\tau^2 - \tau^1]^{-1}$ a.s.. Moreover, we have $H_{\tau^{j(n)}} \leq H_n \leq H_{\tau^{j(n)+1}}$ by the definition of regeneration times, so $\frac{H_n}{j(n)} \to \mathbb{E}[H_{\tau^2} - H_{\tau^1}]$ a.s.. The result follows, with

$$v_\mu = \frac{\mathbb{E}\left[H_{\tau^2} - H_{\tau^1}\right]}{\mathbb{E}\left[\tau^2 - \tau^1\right]} > 0.$$

□

Proof of Proposition 5.21. We will actually show that τ^1 has a subpolynomial tail, i.e. for every $\beta > 0$, we have

$$\mathbb{P}\left(\tau^1 > n\right) = o(n^{-\beta}).$$

We first need to introduce a few notation. We define by induction stopping times τ_j and τ'_j for every $j \geq 1$:

— $\tau_1 = \inf\{n | H_n > 0\}$,
— $\tau'_j = \inf\{n \geq \tau_j | H_n < H_{\tau_j}\}$ for every $j \geq 1$,
— $\tau_{j+1} = \inf\left\{n > \tau'_j | H_n > \max\left(H_0, H_1, \ldots, H_{\tau'_j}\right)\right\}$.

Let also J be the largest index such that $\tau_J < +\infty$. We claim that J is a geometric variable. Indeed, on the one hand, we know that $H_n \to +\infty$ when $n \to +\infty$, so almost surely, if $\tau'_j < +\infty$, then $\tau_{j+1} < +\infty$. On the other hand, if $\tau_j < +\infty$, let $\mathcal{F}_{\tau_j}$ be the σ-algebra generated by $\mathcal{B}_{0,H_{\tau_j}}$ and $(X_0, X_1, \dots, X_{\tau_j})$. Then the variable

$$\left(\mathcal{B}_{H_{\tau_j},\infty}, (X_{\tau_j+i})_{0\le i\le \tau'_j-\tau_j}\right)$$

is independent of $\mathcal{F}_{\tau_j}$ and has the same distribution as $\left(\mathcal{B}_{1,\infty}, (X_{\tau_1+i})_{0\le i\le \tau'_1-\tau_1}\right)$. In particular, if $\tau_j < +\infty$, we have

$$\mathbb{P}\left(\tau'_j < +\infty | \mathcal{F}_{\tau_j}\right) = \mathbb{P}\left(\tau'_1 < +\infty\right) = \mathbb{P}\left(T_\partial = +\infty\right),$$

so $\mathbb{P}\left(\tau_{j+1} < +\infty | \tau_j < +\infty\right)$ does not depend on j. This shows that J is a.s. finite and geometric. Note that $\tau^1 = \tau_J$. For any $n > 0$, we also denote by J_n the largest index j such that $\tau_j \le n$. Finally, we recall that D_n is the greatest "descent" of X before time n.

In order to estimate the tail of τ^1, we partition the event $\{\tau^1 > n\}$ into several "bad" events. Let $\delta > 0$ be small (we will actually only need $\delta < 1/3$). We have

$$\begin{aligned}\mathbb{P}\left(\tau^1 > n\right) &= \mathbb{P}\left(\tau^1 \neq \tau_{J_n}\right)\\ &\le \mathbb{P}\left(J_n \ge n^\delta\right) + \mathbb{P}\left(D_n \ge n^\delta\right) + \mathbb{P}\left(J_n < n^\delta, D_n < n^\delta, \tau^1 \neq \tau_{J_n}\right). \end{aligned} \tag{5.10}$$

We now bound these terms one by one. First, we know that $J_n \le J$, which is a geometric variable. Hence, the first term is at most $\exp(-cn^\delta)$ for some constant c, so it is $o(n^{-\beta})$ for any $\beta > 0$. Moreover, the second part of Proposition 5.19 shows that the second term is $o(n^{-\beta})$ as well.

Finally, we study the third term of (5.10). We first show that if $D_n < n^\delta$ and $J_n < n^\delta$, then $H_{\tau_{J_n}} < n^{2\delta}$ (this is a deterministic statement). If $D_n < n^\delta$, let $1 \le j < J_n$. We have $\tau_{j+1} \le n$, so $\tau'_j \le n$ and $H_{\tau'_j} = H_{\tau_j} - 1$ by the definition of τ'_j. By the definitions of τ_{j+1} and of D_n, we have

$$H_{\tau_{j+1}} - H_{\tau_j} = 1 + \max_{[0,\tau'_j]} H - (H_{\tau'_j} + 1) \le D_n < n^\delta.$$

By summing over j (and remembering $H_{\tau_1} = 1$), we obtain

$$H_{\tau_{J_n}} \le 1 + n^\delta(J_n - 1) < 1 + n^\delta(n^\delta - 1) < n^{2\delta}.$$

Therefore, if the event in the third term of (5.10) occurs, we have $H_{\tau_{J_n}} < n^{2\delta}$ but $\tau'_{J_n} < +\infty$, so there is $k > n$ such that $H_k < n^{2\delta}$. On the other hand, if $\delta < 1/3$, we have

$$\mathbb{P}\left(\exists k > n, H_k \le n^{2\delta}\right) \le \mathbb{P}\left(\exists k > n, H_k \le k^{1-\delta}\right) \le \sum_{k>n} \mathbb{P}\left(H_k \le k^{1-\delta}\right) = \sum_{k>n} o\left(k^{-(\beta+1)}\right)$$

by the first point of Proposition 5.19. This proves that the third term of (5.10) decays super-polynomially, which concludes the proof. □

5.5 Counterexamples and open questions

We finally discuss the necessity of the various assumptions made in the results of this paper, and we state a few conjectures. See Figure 5.3 for a quick summary.

Liouville property. We first note that if we do not require the strips (S_i) to be i.i.d., then Theorem 5.2 bis fails. Indeed, we start from $\mathcal{C}(\mathbf{T})$ and choose a ray γ_0 of $\mathbf{T}$. We then duplicate many times the horizontal edges to add a very strong lateral drift towards γ_0. If we also duplicate the edges of γ_0 enough times, we can make sure that the simple random walk eventually stays on the path γ_0. This yields a map of the form $\mathcal{M}(\mathbf{T}, (S_i))$ which has the intersection property, so it is Liouville.

Poisson boundary. The description of the Poisson boundary given by Theorem 5.2 cannot be true for any map of the form $\mathcal{M}(\mathbf{T}, (S_i))$, even if the strips (S_i) are i.i.d.. Indeed, it is possible to choose S_i such that the walk (X_n) has a positive probability to stay in S_i forever, and such that S_i itself has a non-trivial Poisson boundary. In this case, the Poisson boundary of $\mathcal{M}(\mathbf{T}, (S_i))$ is larger than $\widehat{\partial}\mathbf{T}$, and nonatomicity in Theorem 5.2 is false. On the other hand, we conjecture that if we furthermore assume that all the slices S_i are recurrent graphs (and i.i.d.), then $\widehat{\partial}\mathbf{T}$ is a realization of the Poisson boundary. As explained in Remark 5.12, our arguments cannot handle this general setting.

Positive speed. The positive speed is also false in general maps of the form $\mathcal{M}(\mathbf{T}, (S_i))$ if the strips S_i are too large and do not add vertical drift. For example, if they are equal to the half-planar regular triangular lattice, then the random walk will spend long periods in the same strip, where it has speed zero. On the other hand, we conjecture that the assumption $\mu(0) = 0$ is not necessary in Theorem 5.3.

As for Galton–Watson trees, another process of interest on the maps $\mathcal{C}(T)$ is the λ-biased random walk X^λ. If a vertex x has $c(x)$ children and $X_n^\lambda = x$, then X_{n+1}^λ is equal to y with probability $\frac{1}{c(x)+3\lambda}$ for every child y of x, and to z with probability $\frac{\lambda}{c(x)+3\lambda}$ if z is the parent or one of the two neighbours of x.

If $\lambda > 1$, we expect that, whether $\mu(0) = 0$ or not, the process behaves in the same way as on trees [109]: the walk is recurrent for $\lambda > \lambda_c$ (as easily shown by the Nash–Williams criterion) and should have positive speed for $\lambda < \lambda_c$, where $\lambda_c = \sum i\mu(i)$. If $\lambda < 1$ and $\mu(0) = 0$, it is easy to see that the speed is positive on $\mathcal{C}(T)$ since the drift at every vertex is positive. For $\lambda < 1$ and $\mu(0) > 0$, the λ-biased walk on T has speed zero for λ small enough ($\lambda \leq f'(q)$, where q is the extinction probability of T and f the generating function of μ). We believe that this regime disappears on causal maps, and that the λ-biased walk on $\mathcal{C}(T)$ has positive speed for every $\lambda < 1$.

Other properties of the simple random walk (for $\mu(0) = 0$). As shown by Theorem 5.2, the harmonic measure of $\mathcal{C}(T)$ on $\widehat{\partial}T$ is a.s. nonatomic and has full support. It would be interesting to investigate finer properties of this measure, as it has been done for Galton–Watson trees [108, 105]. We believe that as for Galton–Watson trees, the harmonic measure is not absolutely continuous with respect to the mass measure, and should satisfy a dimension drop.

Another quantity of interest related to the simple random walk is the heat kernel decay, i.e. the probability of returning to the root at time n. Perhaps surprisingly, the annealed and quenched heat kernels might have different behaviours: if $\mu(1) > 0$, the possibility that T does not branch during the first $n^{1/3}$ steps gives an annealed lower bound of order $e^{-n^{1/3}}$. On the other hand, the worst possible traps after the first branching points seem to be large portions of square lattice, which yield a quenched lower bound of order $e^{-n^{1/2}}$. Our argument for quasi-positive speed could be adapted to prove that the heat kernel decays quicker than any polynomial, which seems far from optimal. On the other hand, a natural first step to show that the lower bounds

are tight would be to prove anchored expansion for $\mathcal{C}(T)$. However, this property does not seem well suited to the study of causal maps since connected subsets of $\mathcal{C}(T)$ can be quite nasty.

Other random processes. Finally, other random processes such as percolation on $\mathcal{C}(T)$ might be investigated. We expect that we should have $p_c < p_u$, i.e. there is a regime where infinitely many infinite components coexist, as it is generally conjectured for hyperbolic graphs [30]. We note that oriented percolation is studied in a work in progress of David Marchand.

More generally, for unimodular, planar graphs, other notions of hyperbolicity (including $p_c < p_u$) have been studied in [15] and proved to be equivalent to each other. It might be interesting to study the relation with our setting: if it is true that any hyperbolic (in the sense of [15]) unimodular map contains a supercritical Galton–Watson tree, then our results of Section 5.2 apply. On the other hand, it is clear that every unimodular planar map containing a supercritical Galton–Watson tree is hyperbolic in the sense of [15].

5.A Appendix: The regeneration structure

The goal of this appendix is to prove Lemma 5.22. We recall that $\partial\mathcal{H}$ is the set of vertices at height -1, and that T_∂ is the first time at which X hits $\partial\mathcal{H}$. We will first prove the following intermediate result.

Lemma 5.23. 1. We have $\tau^1 < +\infty$ a.s..

2. The path-decorated map $\left(\mathcal{B}_{H_{\tau^1},\infty}, (X_{\tau^1+i})_{i\geq 0}\right)$ is independent of $\left(\mathcal{B}_{0,H_{\tau^1}}, (X_i)_{0\leq i\leq \tau^1}\right)$ and has the same distribution as $(\mathcal{H}, (X_i)_{i\geq 0})$ conditioned on the event $\{T_\partial = +\infty\}$.

Note that the first point follows from Proposition 5.21, so we only need to focus on the second point.

Proof of Lemma 5.23. We first note that, by Proposition 5.19, we have $H_n \to +\infty$ a.s., so the conditioning on $\{T_\partial = +\infty\}$ is non-degenerate.

For every $h \geq 0$, let $T_h = \min\{n \geq 0 | H_n = h\}$, and let $T'_h = \min\{n \geq T_h | H_n < h\}$. By Proposition 5.19, we have $T_h < +\infty$ a.s.. We also know that the rooted map $(\mathcal{B}_{h,\infty}, X_{T_h})$ is independent of $(\mathcal{B}_{0,h}, (X_i)_{0\leq i\leq T_h})$ and has the same distribution as $\mathcal{H}$. Therefore, the path-decorated map $\left(\mathcal{B}_{h,\infty}, (X_{T_h+i})_{0\leq i\leq T'_h - T_h}\right)$ is independent of $(\mathcal{B}_{0,h}, (X_i)_{0\leq i\leq T_h})$ and has the same distribution as $(\mathcal{H}, (X_i)_{0\leq i\leq T_\partial})$.

It follows that, for any two measurable sets A and B of path-decorated maps, we have

$$\begin{aligned}
&\mathbb{P}\left(\left(\mathcal{B}_{0,H_{\tau^1}}, (X_i)_{0\leq i\leq \tau^1}\right) \in A \text{ and } \left(\mathcal{B}_{H_{\tau^1},\infty}, (X_{\tau^1+i})_{i\geq 0}\right) \in B\right)\\
&= \sum_{h\geq 0} \mathbb{P}\left((\mathcal{B}_{0,h}, (X_i)_{0\leq i\leq T_h}) \in A \text{ and } (\mathcal{B}_{h,\infty}, (X_{T_h+i})_{i\geq 0}) \in B \text{ and } \tau^1 = h\right)\\
&= \sum_{h\geq 0} \mathbb{P}\left((\mathcal{B}_{0,h}, (X_i)_{0\leq i\leq T_h}) \in A \text{ and } \forall i < h, T'_i \leq T_h \right.\\
&\qquad\qquad \left. \text{and } (\mathcal{B}_{h,\infty}, (X_{T_h+i})_{i\geq 0}) \in B \text{ and } T'_h = +\infty\right),
\end{aligned}$$

by noting that H_{τ^1} is the smallest height i such that $T'_i = +\infty$. Note that the event $\{\forall i < h, T'_i \leq T_h\}$ is a measurable function of $(\mathcal{B}_{0,h}, (X_i)_{0\leq i\leq T_h})$, and the event $\{T'_h = +\infty\}$ is a measurable function of $(\mathcal{B}_{h,\infty}, (X_{T_h+i})_{i\geq 0})$. Hence, by the independence and the distribution of

$(\mathcal{B}_{h,\infty}, (X_{T_h+i})_{i\geq 0})$ found above, we have

$$\begin{aligned}
&\mathbb{P}\left(\left(\mathcal{B}_{0,H_{\tau^1}}, (X_i)_{0\leq i\leq \tau^1}\right) \in A \text{ and } \left(\mathcal{B}_{H_{\tau^1},\infty}, (X_{\tau^1+i})_{i\geq 0}\right) \in B\right)\\
&= \sum_{h\geq 0} \mathbb{P}\left((\mathcal{B}_{0,h}, (X_i)_{0\leq i\leq T_h}) \in A \text{ and } \forall i < h, T_i' \leq T_h\right) \mathbb{P}\left((\mathcal{B}_{0,\infty}, (X_i)_{i\geq 0}) \in B \text{ and } T_\partial = +\infty\right)\\
&= \mathbb{P}\left((\mathcal{B}_{0,\infty}, (X_i)_{i\geq 0}) \in B \middle| T_\partial = +\infty\right) f(A),
\end{aligned}$$

where $f(A)$ is a function of A. Therefore, the path-decorated maps $\left(\mathcal{B}_{0,H_{\tau^1}}, (X_i)_{0\leq i\leq \tau^1}\right)$ and $\left(\mathcal{B}_{H_{\tau^1},\infty}, (X_{\tau^1+i})_{i\geq 0}\right)$ are independent and, by taking $A = \Omega$, we obtain that the distribution of the second is a multiple of the distribution of $(\mathcal{H}, (X_i)_{i\geq 0})$ conditioned on the event $\{T_\partial = +\infty\}$. Since both are probability measures, they coincide. □

Proof of Lemma 5.22. We define the shift operator θ as follows:

$$(\mathcal{H}, (X_i)_{i\geq 0}) \circ \theta = \left(B_{H_{\tau^1},\infty}(\mathcal{H}), (X_{\tau^1+i})_{i\geq 0}\right).$$

We first notice that Lemma 5.23 remains true if we consider $(\mathcal{H}, X)$ under the measure $\mathbb{P}(\cdot|T_\partial = +\infty)$ instead of $\mathbb{P}$. Indeed, conditioning on an event of positive probability does not change the fact that $\tau^1 < +\infty$ a.s.. Moreover, the event $\{T_\partial = +\infty\}$ only depends on the path-decorated map $\left(\mathcal{B}_{0,H_{\tau^1}}, (X_i)_{0\leq i\leq \tau^1}\right)$ and not on $\left(\mathcal{B}_{H_{\tau^1},\infty}, (X_{\tau^1+i})_{i\geq 0}\right)$, so conditioning on this event affects neither the independence of these two path-decorated maps, nor the distribution of the second.

But by Lemma 5.23, the map $(\mathcal{H}, (X_i)_{i\geq 0})\circ\theta$ has the same distribution as $(\mathcal{H}, (X_i)_{i\geq 0})$ under $\mathbb{P}(\cdot|T_\partial = +\infty)$, so Lemma 5.23 applies after composition by θ. In particular, we have $\tau^1\circ\theta < +\infty$ a.s., i.e. $\tau^2 < +\infty$ a.s.. Moreover, the two following path-decorated maps are independent :

— $\left(\mathcal{B}_{0,H_{\tau^1}}, (X_i)_{0\leq i\leq \tau^1}\right) \circ \theta = \left(\mathcal{B}_{H_{\tau^1},H_{\tau^2}}, (X_{\tau^1+i})_{0\leq i\leq \tau^2-\tau^1}\right)$,

— $\left(\mathcal{B}_{H_{\tau^1},\infty}, (X_{\tau^1+i})_{i\geq 0}\right) \circ \theta = \left(\mathcal{B}_{H_{\tau^2},\infty}, (X_{\tau^2+i})_{i\geq 0}\right)$,

and the second one has the same distribution as $(\mathcal{H}, (X_i)_{i\geq 0})$ under $\mathbb{P}(\cdot|T_\partial = +\infty)$. From here, an easy induction on j shows that for any $j \geq 1$, we have $\tau^j < +\infty$ and the path-decorated map $\left(\mathcal{B}_{H_{\tau^j},H_{\tau^{j+1}}}, (X_{\tau^j+i})_{0\leq i\leq \tau^{j+1}-\tau^j}\right)$ has indeed the right distribution and is independent of $\left(\mathcal{B}_{H_{\tau^{j+1}},\infty}, (X_{\tau^{j+1}+i})_{i\geq 0}\right)$. This proves Lemma 5.22 (the third item is a direct consequence of the first two). □

Appendix A

Correspondence between type-I and type-II PSHIT

We recall that a *type-I triangulation* is a triangulation that may contain loops and multiple edges, while a *type-II triangulation* is a triangulation that may contain multiple edges, but no loop. The goal of this appendix is to describe precisely the relation between the type-II PSHIT defined in [58] and the type-I PSHIT defined in Chapter 3. As an application, we also compute some constants related to the growth of type-I PSHIT. The type-II PSHIT will be denoted by $(\mathbb{T}_\kappa^{II})_{0<\kappa\leq\kappa_c}$ and the type-I PSHIT by $(\mathbb{T}_\lambda^{I})_{0<\lambda\leq\lambda_c}$, where $\kappa_c = \frac{2}{27}$ and $\lambda_c = \frac{1}{12\sqrt{3}}$.

Two-connected core. Consider an infinite, one-ended type-I triangulation T. Note that if ℓ is a loop in T and f is the face adjacent to ℓ on its unbounded side, then either ℓ is contained in a larger loop, or the two other sides of f are linking the same pair of distinct vertices. Therefore, if we remove all the loops of T and all the vertices and edges lying inside a loop, we obtain a map where all the faces have degree 2 or 3. By gluing multiple edges to close the faces of degree 2, we obtain a type-II triangulation that we call the *two-connected core* of T, or simply its *core*, and denote by $\mathrm{Core}(T)$. If the root edge of T lies inside a loop attached to a vertex v, we root $\mathrm{Core}(T)$ at the edge that has been contracted around this loop, oriented such that v is its starting point.

On the other hand, we cannot recover T from $\mathrm{Core}(T)$. However, it is easy to see that every triangulation T with core C can be obtained by the following procedure.

- We duplicate every edge e of C into $b_e \geq 1$ multiple edges. These multiple edges form $b_e - 1$ faces of degree 2 that we call *facets*.
- For every edge e, we insert a loop in each of the $b_e - 1$ facets lying between the copies of e (there are always two vertices adjacent to a facet, so two ways to insert the loop).
- We fill all these loops with type-I triangulations of the 1-gon (for every edge e of C and every $1 \leq i \leq b_e - 1$, we denote by t_e^i the triangulation filling the loop in the i-th facet formed by copies of e).

The core decomposition of $\mathbb{T}_\lambda^I$. We will denote by e_0^I the root edge of $\mathbb{T}_\lambda^I$ and by e_0^{II} the root edge of its core. Unfortunately, the complete description of the core decomposition of $\mathbb{T}_\lambda^I$ is quite long because the neighbourhood of the root edge is a bit tricky to handle.

We fix $\lambda \in (0, \lambda_c]$ and, as in Chapter 3, we take $h \in \left(0, \frac{1}{4}\right]$ such that $\lambda = \frac{h}{(1+8h)^{3/2}}$. We recall that the partition function of Boltzmann type-I triangulations of a 1-gon with parameter λ is

$$Z_1(\lambda) = \frac{1}{2} - \frac{1+2h}{2\sqrt{1+8h}}.$$

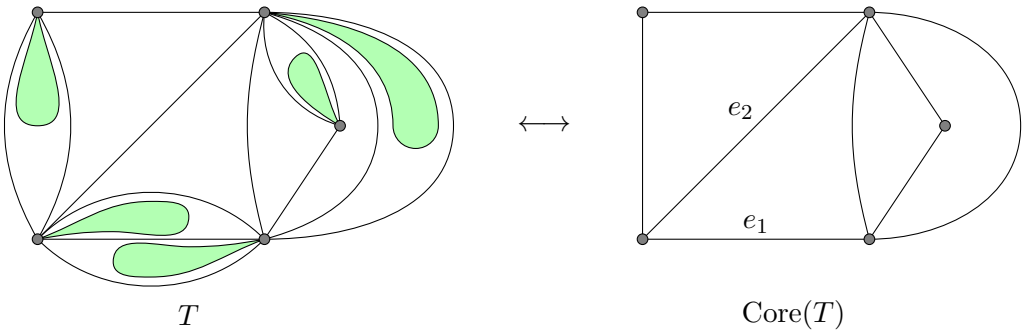

Figure A.1 – A triangulation T and its core. On this example, we have $b_{e_1} = 3$ and $b_{e_2} = 1$. The green zones are filled with triangulations of a 1-gon (which have not been drawn for the sake of clarity).

Let $\beta = 2Z_1(\lambda)$ and $\kappa = \frac{\lambda}{(1-2Z_1(\lambda))^3} = \frac{h}{(1+2h)^3}$. Finally, unless otherwise specified, by a *geometric variable B with parameter β*, we will mean a geometric variable starting at 1, i.e.

$$\mathbb{P}(B = i) = (1-\beta)\beta^{i-1}$$

for every $i \geq 1$.

Proposition A.1. The core of $\mathbb{T}^I_\lambda$ has the same distribution as $\mathbb{T}^{II}_\kappa$. Moreover, conditionally on Core($\mathbb{T}^I_\lambda$), the following holds.

- The variables b_e for the edges e of Core($\mathbb{T}^I_\lambda$) are independent.
- For every edge $e \neq e_0^{II}$ of Core($\mathbb{T}^I_\lambda$), the variable b_e is geometric with parameter β.
- The vertices at which the loops are inserted in the facets are picked according to independent fair coin flips.
- The triangulations t^i_e for $e \neq e_0^{II}$ and $1 \leq i \leq b_e - 1$ are independent type-I Boltzmann triangulations of the 1-gon with parameter λ.
- The probability that e_0^I is not a loop and is not inside a loop is equal to $\frac{1}{\sqrt{1+8h}}$.
- Conditionally on e_0^I not being in a loop:
  — the variable $b_{e_0^{II}}$ is a geometric variable of parameter β biased by its size,
  — the root e_0^I is chosen uniformly among the $b_{e_0^{II}}$ copies of e_0^{II} and oriented like e_0^{II},
  — the insertions of the loops "inbetween" e_0^{II} are again decided by independent fair coins,
  — the triangulations $t^i_{e_0^{II}}$ for $1 \leq i \leq b_{e_0^{II}} - 1$ are independent type-I Boltzmann triangulations of the 1-gon with parameter λ.
- conditionally on e_0^I being a loop or inside a loop:
  — the variable $b_{e_0^{II}} - 1$ is a geometric variable of parameter β biased by its size,
  — the index i_0 such that $e_0^I \in t_e^{i_0}$ is picked uniformly in $\{1, 2, \ldots, b_{e_0^{II}} - 1\}$,
  — the insertions of the loops containing $t^i_{e_0^{II}}$ for $i \neq i_0$ are decided by independent fair coins (the insertion for $t^{i_0}_{e_0^{II}}$ is fixed by the orientation of e_0^{II}),
  — the triangulations $t^i_{e_0^{II}}$ for $i \neq i_0$ are independent Boltzmann type-I triangulations of the 1-gon with parameter λ,
  — the triangulation $t^{i_0}_{e_0^{II}}$ is a Boltzmann type-I triangulation of the 1-gon with parameter λ, biased by its total number of edges, and e_0^I is picked uniformly among all the edges

of $t^{i_0}_{e_0^{II}}$, and oriented by a fair coin flip.

We note that a similar decomposition is described by Angel and Ray [18] to build Markovian type-I triangulations of the half-plane from Markovian type-II triangulations. However, in [18], the Markov property considered is weaker than the one from [58], and we obtain Markovian triangulations if we replace (b_e) by i.i.d. geometric variables of any parameter and fill the loops with i.i.d. triangulations of any law. This is not true anymore with the stronger Markov property of [58].

We recall that a *triangulation with a hole of perimeter p* is a triangulation with a simple boundary of length p, rooted at any oriented edge, while a *triangulation of a p-gon* is rooted in such a way that the external face lies on the right of the root edge. In all the proof below, if t is a triangulation with a hole, we will denote by $|t|$ its total number of vertices. If t is a triangulation of a p-gon, we will denote by $|t|$ the number of its vertices that *do not* lie on its boundary.

Proof. We will prove this decomposition by rebuilding $\mathbb{T}^I_\lambda$ from $\mathbb{T}^{II}_\kappa$. Clearly, there is a unique (in distribution) random triangulation $\widehat{T}$ that satisfies all the conclusions of Proposition A.1, and we only need to prove that $\widehat{T}$ has the same distribution as $\mathbb{T}^I_\lambda$. For this, we will prove that $\widehat{T}$ is λ-Markovian and conclude by Proposition 5 of Chapter 3. Let t be a finite triangulation with a hole of perimeter p. We split the event $\{t \subset \widehat{T}\}$ in two, and write

$$A_1 = \{t \subset \widehat{T} \text{ and } t \text{ is surrounded in } \widehat{T} \text{ by a loop}\} \quad \text{and} \quad A_2 = \{t \subset \widehat{T}\} \backslash A_1.$$

In the case where $p = 1$, we consider that the unique edge of ∂t is a loop surrounding t. We will denote by e_0^I the root edge of $\widehat{T}$, and by e_0^{II} the root edge of its core.

If t is surrounded by a loop, then in particular e_0^I is surrounded by a loop, which has probability $1 - \frac{1}{\sqrt{1+8h}}$. If this occurs, let ℓ be the maximal loop of $\widehat{T}$ surrounding t. By the construction of $\widehat{T}$, the triangulation T_1 filling ℓ is a Boltzmann type-I triangulation of a 1-gon of parameter λ, biased by its number of edges and rooted at a uniform oriented edge. Moreover, we have

$$\mathbb{P}(A_1) = \left(1 - \frac{1}{\sqrt{1+8h}}\right) \mathbb{P}(t \subset T_1). \tag{A.1}$$

If $t \subset T_1$, the part of T_1 lying between its boundary and t is then a triangulation t' of a p-gon with an additional hole of perimeter 1. If $p = 1$, then t' may be the trivial triangulation consisting of a single loop, which means that t and ℓ are simply glued together. For every possible triangulation t', let $t + t'$ be the triangulation obtained by gluing the boundary of t' to the hole of t, and rooted on the root of t. We have

$$\mathbb{P}(T_1 = t + t') = \frac{\lambda^{|t+t'|}}{2\sum_{t''} |E(t'')| \lambda^{|t''|}}, \tag{A.2}$$

where the sum in the denominator is over all triangulations $|t''|$ of a 1-gon, and $|E(t'')|$ is the number of edges of t''. Moreover, we have $|t+t'| = |t| + |t'|$ (the vertices of the cycle along which t and t' are glued are counted in $|t|$, but not in $|t'|$). By summing (A.2) over all possible values of t', we obtain that $\mathbb{P}(t \subset T_1)$ is of the form $c_p \lambda^{|t|}$. By (A.1), so is $\mathbb{P}(A_1)$. The values of c_p could be computed, but this will not be needed here.

We now move on to the event A_2, where $t \subset \widehat{T}$ is not surrounded by a loop. If this is the case, let $\mathrm{Core}(t)$ be the type-II triangulation obtained by deleting the loops and the parts of t that are separated from its hole by loops. Then we must have $\mathrm{Core}(t) \subset \mathrm{Core}(\widehat{T})$, which has probability $C_p^{II} \kappa^{|\mathrm{Core}(t)|}$. Conditionally on this, let b_e be the number of copies of e in t, and B_e its number

of copies in $\widehat{T}$. Finally, for every $e \in \mathrm{Core}(t)$ and $1 \leq i \leq b_e$, let t_e^i be the triangulation filling the i-th loop inserted "inbetween" e.

We first treat the case where the root edge of t is not separated from the hole by a loop. Then the root is not surrounded by a loop in $\widehat{T}$, so we are in the case of probability $\frac{1}{\sqrt{1+8h}}$ in our proposition. If furthermore the root of $\mathrm{Core}(t)$ does not lie on its boundary, we can write

$$\mathbb{P}(A_2) = \frac{C_p^{II}\kappa^{|\mathrm{Core}(t)|}}{\sqrt{1+8h}} \times \mathbb{P}\left(B_{e_0^{II}} = b_{e_0^{II}}\right) \times \prod_{e\in\partial E} \mathbb{P}(B_e \geq b_e) \times \prod_{e\in E\setminus(\partial E\cup\{e_0^{II}\})} \mathbb{P}(B_e = b_e)$$
$$\times \prod_{e\in E}\prod_{i=1}^{b_e-1}\left(\frac{1}{2}\frac{\lambda^{|t_e^i|}}{Z_1(\lambda)}\right) \times \frac{1}{b_{e_0^{II}}},$$

where E is the set of all edges of $\mathrm{Core}(t)$ and ∂E the set of its boundary edges. The factors $\frac{1}{2}$ come from the choice of the vertex at which we insert each loop in its facet, and the factor $\frac{1}{b_{e_0^{II}}}$ from the choice of the root of $\widehat{T}$ among the $b_{e_0^{II}}$ copies of e_0^{II}. By using the explicit distribution of the variables b_e, we obtain

$$\mathbb{P}(A_2) = \frac{C_p^{II}\kappa^{|\mathrm{Core}(t)|}}{\sqrt{1+8h}} \times \prod_{e\in E}\prod_{i=1}^{b_e-1} \lambda^{|t_e^i|} \times (1-2Z_1(\lambda))^{|E\setminus\partial E|+1}.$$

By the Euler formula, we have $|E\setminus\partial E|+1 = 3|t|-2p+2$. By replacing κ by its value and using the fact that $|t| = \mathrm{Core}(t) + \sum_{e,i}|t_e^i|$, we finally get

$$\mathbb{P}(A_2) = \frac{C_p^{II}}{\sqrt{1+8h}}\lambda^{|t|}(1-2Z_1(\lambda))^{-2p+2}. \tag{A.3}$$

If e_0^{II} lies on the boundary of $\mathrm{Core}(t)$, the computation is similar, but it is more convenient to use the following description of $b_{e_0^{II}}$: the number of copies of $b_{e_0^{II}}$ on the left and on the right of $b_{e_0^I}$ are i.i.d. geometric variables **started at 0** with parameter $2Z_1(\lambda)$. We obtain the same formula as (A.3) in the end.

We finally treat the case where e_0^I is separated from the hole of t by a loop. Note that in this case, the vertex attached to the loop around e_0^I is fixed by the orientation of e_0^{II}. By similar arguments as above, after using the construction of $\widehat{T}$ and the values of β and κ, we obtain

$$\begin{aligned}\mathbb{P}(A_2) &= C_p^{II}\lambda^{|t|}(1-2Z_1(\lambda))^{-2p+2} \times \left(1-\frac{1}{\sqrt{1+8h}}\right) \times \frac{1}{2(3\lambda Z_1'(\lambda)-Z_1(\lambda))} \\ &= \frac{C_p^{II}}{\sqrt{1+8h}}\lambda^{|t|}(1-2Z_1(\lambda))^{-2p+2}. \end{aligned}\tag{A.4}$$

The factor $2(3\lambda Z_1'(\lambda)-Z_1(\lambda))$ in the denominator is the partition function of Boltzmann triangulations of a 1-gon with a marked oriented edge. The last step uses the values of $Z_1(\lambda)$ and $Z_1'(\lambda)$ given by [95]. Note that (A.3) and (A.4) yield the same result, which is of the form $c_p'\lambda^{|t|}$. Therefore, the probability

$$\mathbb{P}\left(t \subset \widehat{T}\right) = \mathbb{P}(A_1) + \mathbb{P}(A_2)$$

is also of the form $c_p''\lambda^{|t|}$, so $\widehat{T}$ is λ-Markovian. By Proposition 5 of Chapter 3, it must have the same distribution as $\mathbb{T}_\lambda^I$. □

Applications. This decomposition allows to transfer to $\mathbb{T}_\lambda^I$ many properties of $\mathbb{T}_\kappa^{II}$ that are shown in [58]. For example, the simple random walk on $\mathbb{T}_\lambda^I$ has the same distribution as the simple random walk on its core, which has been delayed at some points. Therefore, they must have the same Poisson boundary.

On the other hand, the core of $\mathbb{T}_\lambda^I$ has the same infinite geodesics as $\mathbb{T}_\lambda^I$, so our results from Chapters 4 and 5 extend to the type-II PSHIT. In particular, they are also weakly hyperbolic and admit bi-infinite geodesics. The near-critical scaling limit result of Chapter 3 for the Gromov–Hausdorff distance can also be extended to the type-II PSHIT. However, for the Gromov–Hausdorff distance, this would require some additional work to prove that the mass inside the loops is well distributed.

Finally, we discuss the precise growth of $\mathbb{T}_\lambda^I$. If $v \in \mathrm{Core}\left(\mathbb{T}_\lambda^I\right)$, the distances of v to the root in $\mathbb{T}_\lambda^I$ and in its core are the same up to a random additive constant given by the distance from e_0^I to the largest loop surrounding it. Therefore, by using the sharp exponential growth results for $\mathbb{T}_\kappa^{II}$ (Theorem 2 of [58]), we obtain similar results for $\mathbb{T}_\lambda^I$.

Corollary A.2. Assume $\lambda < \lambda_c$. We have the almost sure convergences

$$\left(\frac{1-2h-\sqrt{1-4h}}{2h}\right)^r \left|\partial B_r^\bullet\left(\mathbb{T}_\lambda^I\right)\right| \xrightarrow[r\to+\infty]{a.s.} \Pi_\lambda^I \quad \text{and} \quad \frac{\left|B_r^\bullet\left(\mathbb{T}_\lambda^I\right)\right|}{\left|\partial B_r^\bullet\left(\mathbb{T}_\lambda^I\right)\right|} \xrightarrow[r\to+\infty]{a.s.} \frac{1}{1-4h},$$

where Π_λ^I is a positive random variable.

In particular, this result shows that the exponential growth rate of $\mathbb{T}_\lambda^I$ is the same as the exponential growth rate of its infinite tree of geodesics defined in Chapter 4.

Sketch of proof of Corollary A.2. We consider the coupling of $\mathbb{T}_\lambda^I$ with $\mathbb{T}_\kappa^{II}$ given by Proposition A.1. Note that there is a random $\theta \geq 0$ (corresponding to the distance from e_0^I to the largest loop around it) such that, for every $r \geq 0$, we have

$$\partial B_{r+\theta}^\bullet\left(\mathbb{T}_\lambda^I\right) = \partial B_r^\bullet\left(\mathbb{T}_\kappa^{II}\right).$$

Therefore, to prove the exponential growth of perimeters, we just need to translate Theorem 2 of [58] in terms of λ instead of κ: the exponential growth rate in [58] is

$$\frac{\alpha+\sqrt{\alpha(3\alpha-2)}}{\alpha-\sqrt{\alpha(3\alpha-2)}},$$

where $\alpha \in \left(\frac{2}{3}, 1\right)$ is such that $\kappa = \frac{\alpha^2(1-\alpha)}{2}$. We easily obtain $\alpha = \frac{1}{1+2h}$, so the exponential growth rate is

$$\frac{\alpha+\sqrt{\alpha(3\alpha-2)}}{\alpha-\sqrt{\alpha(3\alpha-2)}} = \frac{1+\sqrt{1-4h}}{1-\sqrt{1-4h}} = \left(\frac{2h}{1-2h-\sqrt{1-4h}}\right)^{-1}.$$

To estimate the volumes, we notice that the vertices of $B_{r+\theta}^\bullet\left(\mathbb{T}_\lambda^I\right)$ are exactly:

- the vertices of $B_r^\bullet\left(\mathbb{T}_\kappa^{II}\right)$,
- the vertices inside the loops inserted "inbetween" the edges of $B_r^\bullet\left(\mathbb{T}_\kappa^{II}\right)$ that *do not* belong to the boundary of the hull.

By Theorem 2 of [58], we have

$$\frac{\left|B_r^\bullet\left(\mathbb{T}_\kappa^{II}\right)\right|}{\left|\partial B_r^\bullet\left(\mathbb{T}_\kappa^{II}\right)\right|} \xrightarrow[r\to+\infty]{a.s.} v_{II},$$

where $v_{II} = \frac{2\alpha-1}{3\alpha-2} = \frac{1-2h}{1-4h}$. Therefore, by the Euler formula, the number of edges of $B_r^\bullet\left(\mathbb{T}_\kappa^{II}\right)$ that do not belong to its boundary is asymptotically $(3v_{II}-2)\left|\partial B_r^\bullet\left(\mathbb{T}_\kappa^{II}\right)\right|$. Moreover, the average

number of loops inserted in each of these edges is $\frac{\beta}{1-\beta} = \frac{2Z_1(\lambda)}{1-2Z_1(\lambda)}$, and the average number of additional vertices per loop is $\frac{\lambda Z_1'(\lambda)}{Z_1(\lambda)}$. By the law of large numbers, we obtain

$$\begin{aligned}\frac{\left|B_r^\bullet\left(\mathbb{T}_\lambda^I\right)\right|}{\left|\partial B_r^\bullet\left(\mathbb{T}_\lambda^I\right)\right|} &\xrightarrow[r\to+\infty]{a.s.} v_{II} + (3v_{II}-2) \times \frac{2Z_1(\lambda)}{1-2Z_1(\lambda)} \times \frac{\lambda Z_1'(\lambda)}{Z_1(\lambda)} \\ &= \frac{1}{1-4h}.\end{aligned}$$

□

Appendix B

The average degree in the PSHIT

The goal of this second appendix is to study the degree of the root in the type-I and type-II PSHIT $\mathbb{T}_\lambda^I$ and $\mathbb{T}_\kappa^{II}$. We recall that the *root vertex* ρ of a triangulation T is the starting point of the root edge, and the *root degree* $\deg_T(\rho)$ is the number of half-edges incident to ρ (in particular, loops incident to ρ are counted twice). Our motivation for this is to give a more precise form of Conjecture 1.4. More precisely, for $n \geq 1$ and $g \geq 0$, let $T_{n,g}^I$ (resp. $T_{n,g}^{II}$) be a uniform type-I (resp. type-II) triangulation of a torus of genus g with n vertices. Let also $\theta \geq 0$. As noted in Section 4 of [58], if $T_{n,\theta_{II}n}^{II}$ admits a local limit $\mathbb{T}$ as $n \to +\infty$, then we have

$$\mathbb{E}\left[\frac{1}{\deg_{\mathbb{T}}(\rho)}\right] = \lim_{n\to+\infty} \mathbb{E}\left[\frac{1}{\deg_{T_{n,\theta_{II}n}^{II}}(\rho)}\right] = \frac{1}{6(1+2\theta_{II})}. \tag{B.1}$$

The reason why the quantity of interest is the mean *inverse* root degree and not the mean root degree is that ρ is not picked uniformly in $T_{n,\theta_{II}n}^{II}$: each vertex is chosen with probability proportional to its degree, so we need to bias back by the inverse degree to obtain a "typical" vertex. In particular, if $\mathbb{T}$ is of the form $\mathbb{T}_\kappa^{II}$ as it is believed, then by computing the expected inverse of the root degree in $\mathbb{T}_\kappa^{II}$, we can identify the parameter κ in terms of θ_{II}. This last argument also holds for type-I triangulations. Our computations about the root degree in the PSHIT allow us to state a more precise version of Conjecture 1 of [58].

Conjecture B.1.

- Let $0 < \lambda \leq \lambda_c$, and let $h \in \left(0, \frac{1}{4}\right]$ be such that $\lambda = \frac{h}{(1+8h)^{3/2}}$. Then $\mathbb{T}_\lambda^I$ is the local limit as $n \to +\infty$ of $T_{n,\theta_I n}^I$, where

$$\theta_I = \frac{1}{2}\left(\frac{(1+8h)\sqrt{1-4h}}{6h\log\frac{1+\sqrt{1-4h}}{1-\sqrt{1-4h}}} - 1\right). \tag{B.2}$$

- Let $0 < \kappa \leq \kappa_c$, and let $\alpha \in \left[\frac{2}{3}, 1\right)$ be such that $\kappa = \frac{\alpha^2(1-\alpha)}{2}$. Then $\mathbb{T}_\kappa^{II}$ is the local limit as $n \to +\infty$ of $T_{n,\theta_{II}n}^{II}$, where

$$\theta_{II} = \frac{1}{2}\left(\frac{1}{3(1-\alpha)\left(-1+\frac{\sqrt{\alpha}}{\sqrt{3\alpha-2}}\log\frac{\sqrt{\alpha}+\sqrt{3\alpha-2}}{\sqrt{\alpha}-\sqrt{3\alpha-2}}\right)} - 1\right). \tag{B.3}$$

Note that in particular, we have $\theta_I = 0$ if and only if $\lambda = \lambda_c$, and $\theta_{II} = 0$ if and only if $\kappa = \kappa_c$. On the other hand, we have $\theta_I, \theta_{II} \to +\infty$ when $\kappa, \lambda \to 0$. We will first focus on the type-II case, where the computations are more tractable (although, surprisingly, the formula is slightly simpler for θ_I than for θ_{II}).

B.1 The root degree in $\mathbb{T}_\kappa^{II}$

As explained above, our goal here is to compute the mean inverse degree of the type-II PSHIT.

Proposition B.2. Let $0 < \kappa \leq \kappa_c$, and let $\alpha \in \left[\frac{2}{3}, 1\right)$ be such that $\kappa = \frac{\alpha^2(1-\alpha)}{2}$. Then we have

$$\mathbb{E}\left[\frac{1}{\deg_{\mathbb{T}_\kappa^{II}}(\rho)}\right] = -\frac{1-\alpha}{2} + \frac{(1-\alpha)\sqrt{\alpha}}{2\sqrt{3\alpha-2}} \log \frac{\sqrt{\alpha}+\sqrt{3\alpha-2}}{\sqrt{\alpha}-\sqrt{3\alpha-2}}.$$

This, together with (B.1), justifies the second part of Conjecture B.1. Note that in particular, when $\kappa \to \kappa_c$, we have $\alpha \to \frac{2}{3}$ and $\mathbb{E}\left[\frac{1}{\deg_{\mathbb{T}_\kappa}(\rho)}\right] \to \frac{1}{6}$, so we recover that the average degree in the type-II UIPT is 6. This can also be deduced directly from the local convergence of triangulations of the sphere to the UIPT.

Proof. We will compute the generating function of the root degree in the type-II PSHIT by a peeling approach. Note that it is possible that at some point of the peeling exploration, the root vertex is "trapped" in a finite Boltzmann triangulation. Therefore, we also need to know the distribution of the root degree in Boltzmann triangulations.

We fix $0 < \kappa \leq \kappa_c$. From now on, the type-II PSHIT will be denoted by $\mathbb{T}_\kappa$ (we omit the index II). We will denote by T_κ^p a Boltzmann type-II triangulation of a p-gon with parameter κ. By the spatial Markov property of $\mathbb{T}_\kappa$, for any finite triangulation t with a hole of length p, conditionally on $t \subset \mathbb{T}_\kappa$, the triangulation $\mathbb{T}_\kappa \backslash t$ is a infinite triangulation of the p-gon whose distribution only depends on p. We denote by $\mathbb{T}_\kappa^p$ a triangulation with this distribution. We will first study the root degree in T_κ^p, then in $\mathbb{T}_\kappa^p$. The root degree in $\mathbb{T}_\kappa$ can be easily deduced since $\mathbb{T}_\kappa$ is just $\mathbb{T}_\kappa^2$, where the two boundary edges have been glued together.

We denote by Z_p the partition function of type-II Boltzmann triangulations of a p-gon with parameter κ, and by C_p the constants describing the peeling process of $\mathbb{T}_\kappa$ [58]. We also introduce the following notation:

$$D_p(x) = \mathbb{E}\left[x^{\deg_{T_\kappa^p}(\rho)}\right], \quad \widetilde{D}_p(x) = Z_p D_p(x) \quad \text{and} \quad \mathcal{D}(x,y) = \sum_{p\geq 2} \widetilde{D}_p(x) y^{p-2}$$

and, similarly,

$$E_p(x) = \mathbb{E}\left[x^{\deg_{\mathbb{T}_\kappa^p}(\rho)}\right], \quad \widetilde{E}_p(x) = C_p E_p(x) \quad \text{and} \quad \mathcal{E}(x,y) = \sum_{p\geq 2} \widetilde{E}_p(x) y^{p-2}.$$

The reason why we prefer to work with $\widetilde{D}$ and $\widetilde{E}$ instead of D and E will be clear in a few lines. We also define

$$\mathcal{Z}(y) = \sum_{p\geq 2} Z_p y^{p-2} = \mathcal{D}(1,y) \quad \text{and} \quad \mathcal{C}(y) = \sum_{p\geq 2} C_p y^{p-2} = \mathcal{E}(1,y). \tag{B.4}$$

These last two functions can be computed. The first one is an intermediate step in the enumeration of triangulations of polygons (see e.g. [79]):

$$\mathcal{Z}(y) = \frac{1}{2y^2}\left((\kappa-\alpha y)\sqrt{1-\frac{4y}{\alpha^2}} + y - \kappa\right).$$

Moreover, by the same idea as in the type-I case (see the proof of Proposition 5 of Chapter 3), we can show

$$\mathcal{C}(y) = \frac{1}{\kappa^2 - \kappa y + 2\kappa y^2 \mathcal{Z}(y)}. \tag{B.5}$$

We can now start by computing $\mathcal{D}$. By peeling the edge on the left of the root vertex in T_κ^p, we obtain the following equation:

$$D_p(x) = \frac{\kappa Z_{p+1}}{Z_p} x D_{p+1}(x) + \frac{\mathbb{1}_{p=2}}{Z_p} x + \sum_{i=1}^{p-2} \frac{Z_{p-i} Z_{i+1}}{Z_p} x D_{i+1}(x). \tag{B.6}$$

The term $\frac{\mathbb{1}_{p=2}}{Z_p} x$ comes from the fact that T_κ^2 may just be the trivial triangulation where the two edges of the boundary are simply glued together. By multiplying both sides by Z_p, the last equation becomes a (nicer) equation about the numbers $\widetilde{D}_p(x)$. By multiplying by y^{p-2} and summing over p, we obtain

$$\mathcal{D}(x,y) = \frac{\kappa x}{y} \left(\mathcal{D}(x,y) - \mathcal{D}(x,0)\right) + x + xy\mathcal{Z}(y)\mathcal{D}(x,y). \tag{B.7}$$

We solve this equation by a kernel method: it can be rewritten

$$\left(1 - \frac{\kappa x}{y} + xy\mathcal{Z}(y)\right) \mathcal{D}(x,y) = x - \frac{\kappa x}{y} \mathcal{D}(x,0). \tag{B.8}$$

For every x, let $Y(x)$ be such that the first factor in the left-hand side vanishes, i.e.

$$1 - \frac{\kappa x}{Y(x)} + xY(x)\mathcal{Z}(Y(x)) = 0. \tag{B.9}$$

Note that Y can be explicitely computed by a computer algbra system, but the formula is a bit complicated and it will be more convenient not to do so. The right-hand side of (B.8) must vanish as well for $y = Y(x)$, so $\mathcal{D}(x,0) = \frac{Y(x)}{\kappa}$ and, by (B.7), we obtain

$$\mathcal{D}(x,y) = \frac{Y(x)}{\kappa - yY(x)\frac{y\mathcal{Z}(y) - Y(x)\mathcal{Z}(Y(x))}{y - Y(x)}}.$$

By letting $y \to Y(x)$, we also obtain

$$\mathcal{D}(x, Y(x)) = \frac{Y}{\kappa - Y^2\mathcal{Z}(Y) - Y^3\mathcal{Z}'(Y)}. \tag{B.10}$$

Similarly, by performing one peeling step on $\mathbb{T}_\kappa^p$, we obtain

$$E_p(x) = \frac{\kappa C_{p+1}}{C_p} x E_{p+1}(x) + \sum_{i=1}^{p-2} \frac{C_{p-i} Z_{i+1}}{C_p} x D_{i+1}(x) + \sum_{i=1}^{p-2} \frac{C_{i+1} Z_{p-i}}{C_p} x E_{i+1}(x),$$

where the first sum corresponds to the case where the discovered face separates ρ from infinity, and the second to the case where ρ still lies in the infinite part. Exactly like (B.6), this equation is nicer when rewritten in terms of $\widetilde{E}_p(x)$. We then multiply it by y^{p-2} and sum over p to get

$$\mathcal{E}(x,y) = \frac{\kappa x}{y} \left(\mathcal{E}(x,y) - \mathcal{E}(x,0)\right) + xy\mathcal{C}(y)\mathcal{D}(x,y) + xy\mathcal{Z}(y)\mathcal{E}(x,y).$$

We now apply the same kernel trick as for $\mathcal{D}$. Note that we are only interested in the root degree in $\mathbb{T}_\kappa^2$, so it is enough to compute $\mathcal{E}(x,0)$. The last equation can be rewritten

$$\left(1 - \frac{\kappa x}{y} + xy\mathcal{Z}(y)\right) \mathcal{E}(x,y) = xy\mathcal{C}(y)\mathcal{D}(x,y) - \frac{\kappa x}{y} \mathcal{E}(x,0).$$

Note that the factor in the left-hand side is the same as in (B.8). Therefore, if we replace y by $Y(x)$ as given by (B.9), we obtain

$$\mathcal{E}(x,0) = \frac{Y^2}{\kappa}\mathcal{C}(Y)\mathcal{D}(x,Y) = \frac{Y^3}{\kappa^2\left(\kappa - Y + 2Y^2\mathcal{Z}(Y)\right)\left(\kappa - Y^2\mathcal{Z}(Y) - Y^3\mathcal{Z}'(Y)\right)}, \tag{B.11}$$

by using (B.5) and (B.10).

From here, we can compute the expected inverse of the root degree. We recall that $\mathbb{T}_\kappa$ is equal to $\mathbb{T}_\kappa^2$ where we have glued together the two edges of the boundary, so $\deg_{\mathbb{T}_\kappa}(\rho) = \deg_{\mathbb{T}_\kappa^2}(\rho) - 1$, so the generating function of $\deg_{\mathbb{T}_\kappa}(\rho)$ is $\frac{\kappa^2}{x}\mathcal{E}(x,0)$ (the factor $\kappa^2 = 1/C_2$ comes from the distinction between $E_2(x)$ and $\widetilde{E}_2(x)$). Hence, we have

$$\mathbb{E}\left[\frac{1}{\deg_{\mathbb{T}_\kappa}(\rho)}\right] = \int_0^1 \mathbb{E}\left[x^{\deg_{\mathbb{T}_\kappa}(\rho)-1}\right]\mathrm{d}x = \int_0^1 \frac{\kappa^2}{x^2}\mathcal{E}(x,0)\,\mathrm{d}x.$$

To compute this integral, we use the change of variables $Y(x) = y$. By (B.9), we have

$$x = \frac{y}{\kappa - y^2\mathcal{Z}(y)} \quad \text{and} \quad \mathrm{d}x = \frac{\kappa + y^2\mathcal{Z}(y) + y^3\mathcal{Z}'(y)}{\left(\kappa - y^2\mathcal{Z}(y)\right)^2}\,\mathrm{d}y.$$

By using the expression (B.11) of $\mathcal{E}$, we obtain

$$\begin{aligned}
\mathbb{E}\left[\frac{1}{\deg_{\mathbb{T}_\kappa}(\rho)}\right] &= \int_{Y(0)}^{Y(1)} \frac{y}{\kappa - y + 2y^2\mathcal{Z}(y)}\mathrm{d}y \\
&= \int_0^{(1-\alpha)(3\alpha-1)/4} \frac{y}{(\kappa - \alpha y)\sqrt{1 - \frac{4y}{\alpha^2}}}\,\mathrm{d}y \\
&= -\frac{1-\alpha}{2} + \frac{(1-\alpha)\sqrt{\alpha}}{2\sqrt{3\alpha-2}}\log\frac{\sqrt{\alpha}+\sqrt{3\alpha-2}}{\sqrt{\alpha}-\sqrt{3\alpha-2}},
\end{aligned}$$

where in the end we use the formula (B.4) for $\mathcal{Z}$. □

B.2 From type-II to type-I

Before explaining where our formula for θ_I comes from, we first give two reasons why the computation performed in the type-II case is more difficult for the type-I PSHIT.

- As in the type-II case, we can pass from $\mathbb{T}_\lambda^I$ to a triangulation of a 1-gon by a simple root transformation: we duplicate the root edge, add a loop between these two copies, and root the obtained map at this new loop. However, if the root edge of $\mathbb{T}_\lambda^I$ is a loop, duplicating it adds 2 to the root degree, whereas if the root edge is not a loop, duplicating it adds only 1 to the degree. Therefore, we must distinguish two cases in the end of the computation.
- A more serious problem is the peeling case described on Figure B.1: in the type-II case, when the discovered face separates the triangulation in 2, all the edges adjacent to the root lie in the same part, whereas it is not the case here. In the case of Figure B.1, the generating functions of the root degrees in t_1 and t_2 must be multiplied, so the type-I analog of Equation (B.6) contains a factor $D_p(x)D_1(x)$ and becomes non-linear in D. Although the equation remains explicitely solvable, it seems to make all the computations much harder, and we did not manage to end the computation.

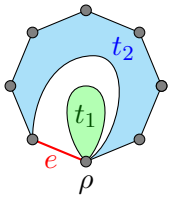

Figure B.1 – A "bad" peeling step: the degree of ρ depends on both parts t_1 and t_2. The peeled edge is e.

To obtain our formula for θ_I, instead of a rigorous computation of the mean inverse root degree, we use a heuristic argument based on the formula for θ_{II} and the core decomposition of the previous appendix. Let $0 < \lambda \leq \lambda_c$, and let h be as in Conjecture B.1. Let also $\kappa = \frac{h}{(1+2h)^3}$, as in Appendix A. Assume that, for some $\theta_I, \theta_{II} \geq 0$, the triangulations $T^I_{n,\theta_I n}$ converge locally to $\mathbb{T}^I_\lambda$, and $T^{II}_{n,\theta_{II} n}$ converge to $\mathbb{T}^{II}_\kappa$.

The idea is the following: we will express in two different ways the proportion of the vertices of $T^I_{n,\theta_I n}$ that belong to its core, and deduce a relation between θ_I and θ_{II}. By using our formula for θ_{II}, this will give a formula for θ_I. Note that the "core" of $T^I_{n,\theta_I n}$ is not precisely well-defined: for maps with a high genus, loops do not always separate the map in two disjoint parts. However, if we believe that $T^I_{n,\theta_I n}$ converges to an infinite planar triangulation, these "non-separating" loops are very rare (the probability to find such a loop in the neighbourhood of a typical vertex goes to 0 as $n \to +\infty$). For example, we could define the "core" of $T^I_{n,\theta_I n}$ as the set of vertices v that can be linked to a vertex at distance 1000 from v, without to cross any loop.

If the triangulations $T^I_{n,\theta_I n}$ converge locally to $\mathbb{T}^I_\lambda$, then the core of $T^I_{n,\theta_I n}$, seen from a typical vertex, looks like the core of $\mathbb{T}^I_\lambda$, which is $\mathbb{T}^{II}_\kappa$. It seems reasonable to expect that the core of $T^I_{n,\theta_I n}$ is roughly of the form $T^{II}_{n',g'}$ for some (random, but concentrated) n', g'. Since this core looks locally like $\mathbb{T}^{II}_\kappa$, if the second part of Conjecture B.1 is true, we should have $\frac{g'}{n'} \sim \theta_{II}$, where θ_{II} is given by (B.3). On the other hand, taking the core should not affect the genus much, so we expect $g' = \theta_I n$, so $n' = \frac{\theta_I}{\theta_{II}}$. Therefore, the core of $T^I_{n,\theta_I n}$ should look locally like $T^{II}_{(\theta_I/\theta_{II})n,\theta_I n}$. In particular, the proportion of the vertices of $T^I_{n,\theta_I n}$ lying in its core is $\frac{\theta_I}{\theta_{II}}$.

On the other hand, the core decomposition of triangulations of the form $T^I_{n,\theta_I n}$ should be similar to the decomposition of $\mathbb{T}^I_\lambda$ given by Proposition A.1. Therefore, such triangulations should be obtained by starting from $\mathbb{T}^{II}_{m,\theta_{II} m}$ for some m and inserting loops in the same way as in Proposition A.1. By the Euler formula, the number of edges of $\mathbb{T}^{II}_{m,\theta_{II} m}$ is asymptotically $3(1 + 2\theta_{II})m$. Moreover, as noted in the proof of Corollary A.2, the mean number of vertices inserted per edge of the core is

$$\frac{2\lambda Z_1'(\lambda)}{1 - 2Z_1(\lambda)} = \frac{2h}{1+2h}.$$

Therefore, the number of new vertices is equivalent to $\frac{6h}{1+2h}(1+2\theta_{II})m$, so the proportion of the final vertices that belong to the core is

$$\frac{1}{1 + \frac{6h}{1+2h}(1 + 2\theta_{II})}.$$

Therefore, we should have

$$\frac{\theta_I}{\theta_{II}} = \frac{1}{1 + \frac{6h}{1+2h}(1 + 2\theta_{II})}.$$

We can now replace θ_{II} by the formula of the second part of Conjecture B.1, and replace α by $\frac{1}{1+2h}$ in this formula. After a few simplifications, we obtain (B.2).

Appendix C

Hyperbolic Boltzmann bipartite maps

The goal of this appendix is to try to generalize the construction of the PSHIT to the general setting of bipartite planar maps with Boltzmann weights on the face degrees. We prove the existence of hyperbolic Boltzmann maps in a case including natural analogs of the PSHIT for $2p$-angulations for $p \geq 2$ (see Proposition C.2 below). However, we do not obtain a full characterization of the Boltzmann weights corresponding to an infinite Markovian map. In particular, we do not know whether the maps with large faces studied in [43] have a hyperbolic analog.

Maps with a general boundary. As for triangulations, to define an infinite Markovian bipartite map M, we first need to define precisely the events of the form $m \subset M$. We will use the definitions from [42] and [57], which are different from those we used with triangulations. We first recall these definitions. Intuitively, they mean that holes are filled by maps *with a general boundary*, whereas for triangulations the holes were filled by maps with a *simple* boundary.

Let M be an infinite, one-ended planar map. Let also E be a connected set of dual edges of M, including the dual edge of the root. Let $F(E)$ be the set of faces of M that are adjacent to at least one edge of E. We denote by $m(E)$ the finite map obtained by gluing the faces of $F(E)$ along the dual edges of the edges of E (see Figure C.1). This finite map is rooted at the edge corresponding to the root edge of M. It has additional faces that do not belong to $F(E)$. We call these faces the *holes* of $m(E)$, while its other faces are called *internal faces*. We write $m(E) \subset M$. We will only be interested in the case where $m(E)$ has only one hole.

The important difference between this definition and the one we used for triangulations is that if we know that $m \subset M$, it is still possible that two different edges of ∂m correspond to the same edge of M. This is equivalent to saying that the boundary of m does not have to be a simple path in M. With this definition, it is still possible to define a Markov property and a peeling process for bipartite maps. This new peeling process, introduced by Budd [42], is better suited to the study of general maps for two reasons:

- The holes that might be created by a peeling step must now be filled by maps with a general boundary, whose combinatorics are better understood than those of maps with a simple boundary.
- When a peeling step discovers a large new face, we do not need to know right now which of its vertices belong to the pre-existing boundary. This avoids an explosion of the number of peeling cases.

We refer to [42] or [57] for a precise definition of this peeling process.

Definition C.1. Let $\mathbf{q} = (q_i)_{i \geq 0}$ be a sequence of nonnegative numbers. We say that a random infinite, one-ended bipartite planar map M is $\mathbf{q}$-*Markovian* if there are constants $(C_p(\mathbf{q}))_{p \geq 1}$

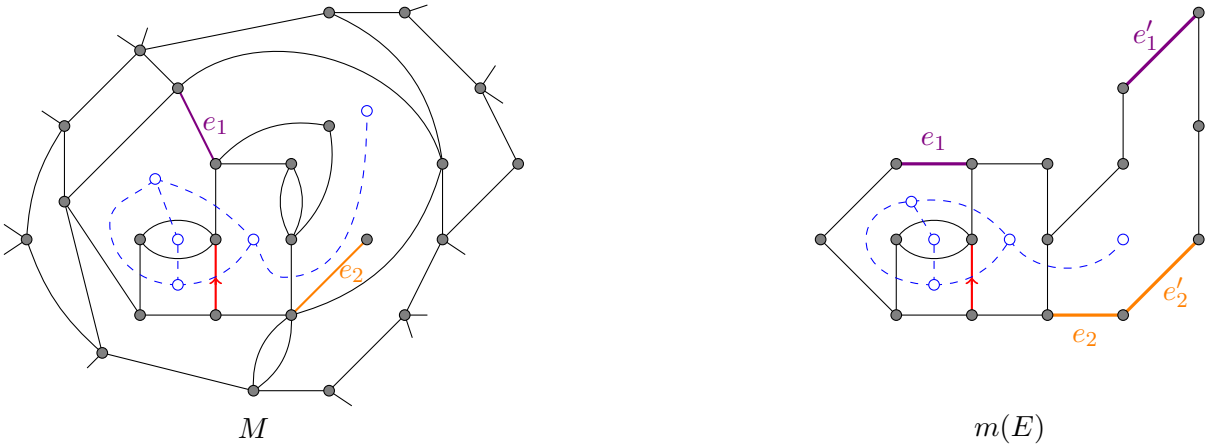

Figure C.1 – On the right, a infinite bipartite M and a set of dual edges E (in blue). On the right, the finite map $m(E)$. Note that the edges e_1, e'_1 of $m(E)$ are glued together in M, as well as the edges e_2, e'_2.

such that, for every map m with one hole of perimeter $2p$, we have

$$\mathbb{P}\left(m \subset M\right) = C_p(\mathbf{q}) \prod_{f \in m} q_{\deg(f)/2},$$

where the product is over all internal faces f of m.

The reason why we restrict ourselves to bipartite maps (i.e. maps where all the face degrees are even) is that the combinatorial results about these maps are simpler to express (see [42]).

The combinatorics of Boltzmann bipartite maps. Before stating our results, we recall a few facts about the combinatorics of Boltzmann bipartite maps. For $p \geq 1$, let $W_p(\mathbf{q})$ be the partition function of Boltzmann map of a p-gon with weight sequence $\mathbf{q}$, i.e.

$$W_p(\mathbf{q}) = \sum_m \prod_{f \in m} q_{\deg(f)/2},$$

where the sum is over all maps m of a p-gon with a general boundary. We say that the weight sequence $\mathbf{q}$ is *admissible* if all these partition functions are finite. Let $f_{\mathbf{q}}(x) = \sum_{k \geq 1} q_k \binom{2k-1}{k-1} x^{k-1}$. The criterion proved in [111] states that $\mathbf{q}$ is admissible if and only if the equation

$$f_{\mathbf{q}}(x) = 1 - \frac{1}{x} \tag{C.1}$$

has a positive solution. If this is the case, let $Z_{\mathbf{q}}$ be the smallest such solution. Then $f_{\mathbf{q}}(x) > 1 - \frac{1}{x}$ for every $0 \leq x < Z_{\mathbf{q}}$, so $f'_{\mathbf{q}}(Z_{\mathbf{q}}) \leq \frac{1}{Z_{\mathbf{q}}^2}$. We call $\mathbf{q}$ *critical* if $f'_{\mathbf{q}}(Z_{\mathbf{q}}) = \frac{1}{Z_{\mathbf{q}}^2}$ (this means that $\mathbf{q}$ is only "barely" admissible) and *sub-critical* if $f'_{\mathbf{q}}(Z_{\mathbf{q}}) < \frac{1}{Z_{\mathbf{q}}^2}$. Moreover, the partition functions $W_p(\mathbf{q})$ can be computed in an implicit way [57, Section 5.1].

Before stating our result, we need to define a measure on $\mathbb{Z}$ that appears naturally when studying the peeling process of finite Boltzmann maps. Let $\mathbf{q}$ be an admissible weight sequence, and let $g_{\mathbf{q}} = 4Z_{\mathbf{q}}$. For $i \in \mathbb{Z}$, we set

$$\nu(i) = \begin{cases} q_{i+1}\, g_{\mathbf{q}}^i & \text{if } i \geq 0, \\ 2W_{-1-i}(\mathbf{q})\, g_{\mathbf{q}}^i & \text{if } i \leq -1. \end{cases} \tag{C.2}$$

Then ν is a probability measure (see Section 3.3.1 of [57]). More precisely, it is the step distribution of the random walk describing the evolution of the "perimeter" of (critical or sub-critical) Boltzmann *half-planar* maps under their peeling process. Moreover, by [57, Proposition 14] (or [42, Proposition 5]), if $\mathbf{q}$ is sub-critical, then the walk with step distribution ν drifts to $-\infty$.

If $\mathbf{q}$ is admissible and critical, then $\mathbf{q}$-Markovian maps are constructed in [42, 57]. They are also proved to be the local limits of $\mathbf{q}$-Boltzmann maps of the sphere conditionned to be large. Therefore, we will mostly be interested in the sub-critical case.

Proposition C.2. (i) If a $\mathbf{q}$-Markovian infinite map exists, it is unique (in distribution).

(ii) If $\mathbf{q}$ is not admissible, then there is no $\mathbf{q}$-Markovian infinite map.

(iii) If $\mathbf{q}$ is admissible and sub-critical, then there is a $\mathbf{q}$-Markovian map if and only if the equation

$$\sum_{i\in\mathbb{Z}} \nu(i)\omega^i = 1 \tag{C.3}$$

has a solution $\omega > 1$.

(iv) In this case, for every $p \geq 1$, we have

$$C_p(\mathbf{q}) = \omega^{p-1} \sum_{i=0}^{p-1} (4\omega)^{-i} \binom{2i}{i}. \tag{C.4}$$

(v) In particular, if $\mathbf{q}$ is admissible and has an infinite convergence radius, then there is a $\mathbf{q}$-Markovian infinite map.

In particular, point (v) includes the case where $\mathbf{q}$ has a finite support. In the very particular case $q_i = q\mathbb{1}_{i=p}$, this implies that we can define non-critical Markovian $2p$-angulations for any $p \geq 2$. Note also that the constants $C_p(\mathbf{q})$ have a universal form, which is the same as for triangulations.

Proof. As for triangulations (Theorem 1 of [58] or Proposition 5 of Chapter 3), the idea is to perform a peeling step and deduce a recurrence relation for the numbers $C_p(\mathbf{q})$. If at some point the discovered part of a Markovian map is m and we peel an edge e of ∂m, the two peeling cases are as follows: either we discover a new face on the other side of e, or we find out that e was actually glued to another edge of ∂m. In the second case, the boundary ∂m is split in two, and one part must be filled with a finite $\mathbf{q}$-Boltzmann map with a general boundary. The equation we obtain is

$$C_p(\mathbf{q}) = \sum_{i\geq 1} q_i C_{p+i-1}(\mathbf{q}) + 2\sum_{i=1}^{p-1} W_{i-1}(\mathbf{q}) C_{p-i}(\mathbf{q}) \tag{C.5}$$

for every $p \geq 1$. In particular, if $W_i(\mathbf{q}) = +\infty$ for some $i \geq 0$, then the constants C_p must be infinite, so there is no $\mathbf{q}$-Markovian map, which proves point (ii).

If there are constants $(C_p(\mathbf{q}))_{p\geq 1}$ satisfying these equations, it is possible to build infinite $\mathbf{q}$-Markovian maps by peeling, along the same lines as in [58]. Moreover, the distribution of a $\mathbf{q}$-Markovian map is entirely characterized by the constants $C_p(\mathbf{q})$. Therefore, the problem is reduced to the existence of a solution to (C.5).

As in [58], we will interpret (C.5) as the search of a harmonic function for a certain random walk. We extend $(C_p(\mathbf{q}))$ to $\mathbb{Z}$ by setting $C_p(\mathbf{q}) = 0$ for $p \leq 0$, and set $\widetilde{C}_p(\mathbf{q}) = g_{\mathbf{q}}^{-p} C_p(\mathbf{q})$. Then (C.5) can be rewritten as

$$\widetilde{C}_p(\mathbf{q}) = \sum_{i\in\mathbb{Z}} \nu(i)\widetilde{C}_{p+i},$$

where ν is given by (C.2). In other words, there is a $\mathbf{q}$-Markovian infinite map if and only if the random walk with step distribution ν admits a non-trivial harmonic (on positive integers) function that vanishes on nonpositive integers.

This problem is precisely adressed by Doney in [71]: for a walk drifting to $-\infty$, such a harmonic function exists if and only if the equation

$$F_{\mathbf{q}}(\omega) := \sum_{i\in\mathbb{Z}} \nu(i)\omega^i = 1 \tag{C.6}$$

admits a solution $\omega > 1$. This proves point (iii). Moreover, if a harmonic function exists, it is unique up to a multiplicative constant, which proves point (i). The results of [71] give a formula for the harmonic functions (if it exists) in terms of the *ladder process* of a random walk S with step distribution ν. More precisely, we define the *decreasing ladder process* of S as the process H^- such that $H_0^- = 0$ and, for every $n \geq 0$, the first value strictly smaller than $-H_n^-$ taken by S is $-H_{n+1}^-$. Let also

$$u(i) = \sum_{n\geq 0} \mathbb{P}\left(H_n^- = i\right)$$

for every $i \geq 0$. Then by Theorem 1 of [71], we have

$$C_p(\mathbf{q}) = \sum_{i=0}^{p-1} \omega^{p-1-i} u(i).$$

The reason why the sum ranges up to $p-1$ instead of p is that we want a harmonic function vanishing on the set of nonpositive integers, whereas [71] deals with harmonic functions vanishing on negative integers, so all the indices must be shifted by 1. On the other hand, the function u is precisely the function $h^{\downarrow}$ in the proof of Proposition 14 in [57], so $u(i) = \frac{1}{4^i}\binom{2i}{i}$, which proves point (iv).

We finally prove point (v). The generating function $F_{\mathbf{q}}$ can be rewritten

$$F_{\mathbf{q}}(\omega) = \sum_{i\geq 1} q_i (g_{\mathbf{q}}\omega)^{i-1} + 2\sum_{i\geq 1} W_{i-1}(\mathbf{q})(g_{\mathbf{q}}\omega)^{-i} = 1. \tag{C.7}$$

Since ν is a probability measure, we have $F_{\mathbf{q}}(1) = 1$. In particular, the second sum in the right-hand side of (C.7) is finite for $\omega = 1$, so it is finite for any $\omega \geq 1$. Moreover, since $\mathbf{q}$ has an infinite convergence radius, the first sum is finite for any ω, so $F_{\mathbf{q}}(\omega)$ is finite for every $\omega \geq 1$. Moreover, the first sum in (C.7) has a finite derivative at $\omega = 1$, so $F'_{\mathbf{q}}(1)$ is either finite or $-\infty$. In the first case, $F'_{\mathbf{q}}(1)$ is the mean of ν, which must be negative since the walk S drifts to $-\infty$. Therefore, in both cases, we have $F_{\mathbf{q}}(x) < 1$ for $x > 1$ small enough. On the other hand, since $\mathbf{q}$ has an infinite convergence radius and $q_i > 0$ for some $i \geq 2$, the first sum in (C.7) goes to $+\infty$ as $\omega \to +\infty$, so there must be $\omega > 1$ such that $F_{\mathbf{q}}(\omega) = 1$. □

Note that the proof of point (v) still works if the convergence radius $R_{\mathbf{q}}$ of $\mathbf{q}$ is finite, as long as $\sum_{i\geq 0} q_i x^{i-1}$ takes the value 1 before $R_{\mathbf{q}}$ (and in particular if it diverges near $R_{\mathbf{q}}$).

The peeling process of non-critical Markovian infinite maps. Let $\mathbf{q}$ be a sub-critical, admissible sequence for which (C.3) has a solution $\omega > 1$. Then we can study $\mathbb{M}_{\mathbf{q}}$, the unique $\mathbf{q}$-Markovian map, by its peeling process. By the formula (C.4), we have

$$\frac{C_{p+i}(\mathbf{q})}{C_p(\mathbf{q})} \xrightarrow[p\to+\infty]{} \omega^i$$

for every $i \in \mathbb{Z}$. Therefore, a quick computation shows that, when p is large, the perimeter process of $\mathbb{M}_{\mathbf{q}}$ behaves like a random walk with step distribution $\widetilde{\nu}$, where

$$\widetilde{\nu}(i) = \omega^i \nu(i)$$

for every $i \in \mathbb{Z}$. In particular, the drift of $\widetilde{\nu}$ is equal to $\omega F'_{\mathbf{q}}(\omega)$, which is positive by convexity of $F_{\mathbf{q}}$ (it may be equal to $+\infty$). Therefore, the perimeter grows at least linearly in the number of peeling steps. By applying the arguments of [43], the growth of the dual of $\mathbb{M}_{\mathbf{q}}$ should be at least exponential (and quicker if $F'_{\mathbf{q}}(\omega) = +\infty$). Therefore, we can expect the non-critical Markovian infinite maps to have a hyperbolic behaviour. If $F'_{\mathbf{q}}(\omega) = +\infty$, we might obtain planar maps exhbiting both a hyperboic behaviour like the PSHIT and the presence of very large faces like the maps studied in [43]. However, we did not find any explicit example where $F'_{\mathbf{q}}(\omega) = +\infty$ (see also next paragraph). This might be a subject of future work.

A counter-example. Given Proposition C.2, it is natural to ask whether for every admissible, sub-critical weight sequence $\mathbf{q}$, there is a $\mathbf{q}$-Markovian map. The answer seems to be no.

Indeed, a natural family of critical, heavy-tailed weight sequences is the one described in Section 6 of [43]: we fix $a \in \left(\frac{3}{2}, \frac{5}{2}\right)$ and, for every $i \geq 1$, we set

$$q_i = c\gamma^{i-1} \frac{\Gamma(1/2 - a + i)}{\Gamma(1/2 + i)} \mathbb{1}_{i \geq 2},$$

where $c = \frac{\sqrt{\pi}}{-2\Gamma(3/2-a)}$ and $\gamma = \frac{1}{4a-2}$. The advantage of these sequences is that the step distribution ν can be expressed in a quite simple way. Moreover, they correspond to infinite maps with a very different geometry from the UIPT (the UIPQ corresponds to the case $a \to \frac{5}{2}$). To turn $\mathbf{q}$ into a sub-critical sequence, we fix $y \in (0, 1)$ and "twist" $\mathbf{q}$ by setting

$$\widetilde{q}_i = y^{i-1} q_i.$$

Then $f_{\widetilde{\mathbf{q}}}(x) = f_{\mathbf{q}}(yx)$ and it is easy to check that $\widetilde{\mathbf{q}}$ is an admissible, sub-critical weight sequence.

Moreover, by Proposition C.2 and the convexity of $F_{\widetilde{\mathbf{q}}}$, there is a $\widetilde{\mathbf{q}}$-Markovian map if and only if $\lim_{x \to R_{\widetilde{\mathbf{q}}}} F_{\widetilde{\mathbf{q}}}(x) \geq 1$, where $R_{\widetilde{\mathbf{q}}}$ is the convergence radius of $\widetilde{\mathbf{q}}$. We can then compute numerically the quantity $F_{\widetilde{\mathbf{q}}}(R_{\widetilde{\mathbf{q}}})$. It seems that for any choice of $a \in \left(\frac{3}{2}, \frac{5}{2}\right)$ and $y \in (0, 1)$, we have $F_{\widetilde{\mathbf{q}}}(R_{\widetilde{\mathbf{q}}}) < 1$, so there should be no $\widetilde{\mathbf{q}}$-Markovian map.

Bibliographie

[1] C. Abraham. Rescaled bipartite planar maps converge to the Brownian map. *Ann. Inst. H. Poincaré Probab. Statist.*, 52(2):575–595, 05 2016.

[2] R. Abraham, J.-F. Delmas, and P. Hoscheit. A note on the Gromov–Hausdorff–Prokhorov distance between (locally) compact metric measure spaces. *Electron. J. Probab.*, 18:21 pp., 2013.

[3] L. Addario-Berry and M. Albenque. The scaling limit of random simple triangulations and random simple quadrangulations. *Ann. Probab.*, 45(5):2767–2825, 09 2017.

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

[5] L. Addario-Berry, N. Broutin, C. Goldschmidt, and G. Miermont. The scaling limit of the minimum spanning tree of the complete graph. *Ann. Probab.*, 45(5):3075–3144, 09 2017.

[6] E. Aïdékon. Speed of the biased random walk on a Galton–Watson tree. *Probability Theory and Related Fields*, 159(3):597–617, Aug 2014.

[7] D. Aldous and R. Lyons. Processes on unimodular random networks. *Electron. J. Probab.*, 12:no. 54, 1454–1508 (electronic), 2007.

[8] D. Aldous and J. M. Steele. The objective method: probabilistic combinatorial optimization and local weak convergence. In *Probability on discrete structures*, volume 110 of *Encyclopaedia Math. Sci.*, pages 1–72. Springer, Berlin, 2004.

[9] J. Ambjørn, B. Durhuus, and T. Jonsson. *Quantum geometry: A statistical field approach.* Cambridge Monographs on Mathematical Physics. Cambridge University Press, Cambridge, 1997.

[10] J. Ambjørn and R. Loll. Non-perturbative lorentzian quantum gravity, causality and topology change. *Nuclear Physics B*, 536(1):407 – 434, 1998.

[11] O. Angel. Scaling of percolation on infinite planar maps, I. *arXiv:0501006.*

[12] O. Angel. Growth and percolation on the uniform infinite planar triangulation. *Geom. Funct. Anal.*, 13(5):935–974, 2003.

[13] O. Angel, G. Chapuy, N. Curien, and G. Ray. The local limit of unicellular maps in high genus. *Electron. Commun. Probab.*, 18(86):1–8, 2013.

[14] O. Angel, T. Hutchcroft, A. Nachmias, and G. Ray. Unimodular hyperbolic triangulations: circle packing and random walk. *Inventiones mathematicae*, 206(1):229–268, Oct 2016.

[15] O. Angel, T. Hutchcroft, A. Nachmias, and G. Ray. Hyperbolic and parabolic unimodular random maps. *Geometric and Functional Analysis*, Jun 2018.

[16] O. Angel, B. Kolesnik, and G. Miermont. Stability of geodesics in the Brownian map. *Ann. Probab.*, 45(5):3451–3479, 09 2017.

[17] O. Angel, A. Nachmias, and G. Ray. Random walks on stochastic hyperbolic half planar triangulations. *Random Structures and Algorithms*, 49(2):213–234, 2016.

[18] O. Angel and G. Ray. Classification of half planar maps. *Ann. Probab.*, 43(3):1315–1349, 2015.

[19] O. Angel and G. Ray. The half plane UIPT is recurrent. *Probability Theory and Related Fields*, 170(3):657–683, Apr 2018.

[20] O. Angel and O. Schramm. Uniform infinite planar triangulations. *Comm. Math. Phys.*, 241(2-3):191–213, 2003.

[21] K. B. Athreya and P. E. Ney. *Branching processes*, volume 196 of *Die Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, 1972.

[22] E. Baur, G. Miermont, and G. Ray. Classification of scaling limits of uniform quadrangulations with a boundary. *Ann. of Probab.*, to appear.

[23] E. Baur and L. Richier. Uniform infinite half-planar quadrangulations with skewness. *Electron. J. Probab.*, 23:43 pp., 2018.

[24] J. Beltran and J.-F. Le Gall. Quadrangulations with no pending vertices. *Bernoulli*, 19:1150–1175, 2013.

[25] I. Benjamini. *Coarse geometry and randomness*, volume 2100 of *Lecture Notes in Mathematics*. Springer, Cham, 2013. Lecture notes from the 41st Probability Summer School held in Saint-Flour, 2011.

[26] I. Benjamini. Gaussian free field on hyperbolic lattices. *Lecture Notes in Mathematics*, 2116:39–45, 08 2014.

[27] I. Benjamini and N. Curien. Ergodic theory on stationary random graphs. *Electron. J. Probab.*, 17(93):1–20, 2012.

[28] I. Benjamini and N. Curien. Simple random walk on the uniform infinite planar quadrangulation: Subdiffusivity *via* pioneer points. *Geom. Funct. Anal.*, 23(2):501–531, 2013.

[29] I. Benjamini, E. Paquette, and J. Pfeffer. Anchored expansion, speed and the Poisson–Voronoi tessellation in symmetric spaces. *Ann. Probab.*, 46(4):1917–1956, 07 2018.

[30] I. Benjamini and O. Schramm. Percolation beyond Z^d, many questions and a few answers. *Electron. Commun. Probab.*, 1:71–82, 1996.

[31] I. Benjamini and O. Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6:no. 23, 13 pp. (electronic), 2001.

[32] I. Benjamini and R. Tessera. First passage percolation on a hyperbolic graph admits bi-infinite geodesics. *Electron. Commun. Probab.*, 22:8 pp., 2017.

[33] O. Bernardi, N. Curien, and G. Miermont. A Boltzmann approach to percolation on random triangulations. *arXiv:1705.04064*, May 2017.

[34] J. Bertoin. *Lévy processes*, volume 121 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1996.

[35] J. Bettinelli. Scaling limits for random quadrangulations of positive genus. *Electron. J. Probab.*, 15:1594–1644, 2010.

[36] J. Bettinelli, E. Jacob, and G. Miermont. The scaling limit of uniform random plane maps, via the Ambjørn–Budd bijection. *Electronic J. Probab.*, 19, 2014.

[37] J. Bettinelli and G. Miermont. Compact Brownian surfaces I: Brownian disks. *Probability Theory and Related Fields*, 167(3):555–614, Apr 2017.

[38] J. E. Björnberg and S. O. Stefansson. Recurrence of bipartite planar maps. *Electron. J. Probab.*, 19(31):1–40, 2014.

[39] G. Borot, J. Bouttier, and E. Guitter. A recursive approach to the O(N) model on random maps via nested loops. *J. Phys. A: Math. Theor.*, 45, 2012.

[40] J. Bouttier, P. Di Francesco, and E. Guitter. Planar maps as labeled mobiles. *Electron. J. Combin.*, 11(1):Research Paper 69, 27 pp. (electronic), 2004.

[41] T. Budd. *Non-perturbative quantum gravity : a conformal perspective. PhD thesis.* PhD thesis, 2012.

[42] T. Budd. The peeling process of infinite Boltzmann planar maps. *Electronic Journal of Combinatorics*, 23, 06 2015.

[43] T. Budd and N. Curien. Geometry of infinite planar maps with high degrees. *Electron. J. Probab.*, 22:37 pp., 2017.

[44] T. Budzinski. On the mixing time of the flip walk on triangulations of the sphere. *Comptes Rendus Mathematique*, 355(4):464 – 471, 2017.

[45] T. Budzinski. The hyperbolic Brownian plane. *Probability Theory and Related Fields*, 171(1):503–541, Jun 2018.

[46] T. Budzinski. Infinite geodesics in hyperbolic random triangulations. *arXiv:1804.07711*, 2018.

[47] T. Budzinski. Supercritical causal maps: geodesics and simple random walk. *arXiv:1806.10588*, 2018.

[48] D. Burago, Y. Burago, and S. Ivanov. *A course in metric geometry*, volume 33 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2001.

[49] P. Caputo, F. Martinelli, A. Sinclair, and A. Stauffer. Random lattice triangulations: Structure and algorithms. *Ann. Appl. Probab.*, 25(3):1650–1685, 06 2015.

[50] A. Caraceni and N. Curien. Geometry of the uniform infinite half-planar quadrangulation. *Random Structures & Algorithms*, 52(3):454–494.

[51] A. Caraceni and A. Stauffer. Polynomial mixing time of edge flips on quadrangulations. *ArXiv e-prints*, Sept. 2018.

[52] G. Chapuy, M. Marcus, and G. Schaeffer. A bijection for rooted maps on orientable surfaces. *SIAM J. Discrete Math.*, 23(3):1587–1611, 2009.

[53] P. Chassaing and B. Durhuus. Local limit of labeled trees and expected volume growth in a random quadrangulation. *Ann. Probab.*, 34(3):879–917, 2006.

[54] P. Chassaing and G. Schaeffer. Random planar lattices and integrated superBrownian excursion. *Probab. Theory Related Fields*, 128(2):161–212, 2004.

[55] R. Cori and B. Vauquelin. Planar maps are well labeled trees. *Canad. J. Math.*, 33(5):1023–1042, 1981.

[56] N. Curien. A glimpse of the conformal structure of random planar maps. *Commun. Math. Phys.*, 333:1417–1463, 2015.

[57] N. Curien. Peeling random planar maps. *Cours Peccot, Collège de France*, 2016.

[58] N. Curien. Planar stochastic hyperbolic triangulations. *Probability Theory and Related Fields*, 165(3):509–540, 2016.

[59] N. Curien. Scaling limits of stable causal maps. *Work in preparation.*, 2018+.

[60] N. Curien, T. Hutchcroft, and A. Nachmias. Geometric and spectral properties of causal maps. *arXiv:1710.03137*, 2017.

[61] N. Curien and I. Kortchemski. Percolation on random triangulations and stable looptrees. *Probability Theory and Related Fields*, 163(1):303–337, Oct 2015.

[62] N. Curien and J.-F. Le Gall. The Brownian plane. *J. Theoret. Probab.*, 27(4):1249–1291, 2014.

[63] N. Curien and J.-F. Le Gall. The hull process of the brownian plane. *Probability Theory and Related Fields*, 166(1):187–231, Oct 2016.

[64] N. Curien and J.-F. Le Gall. Scaling limits for the peeling process on random maps. *Ann. Inst. H. Poincaré Probab. Statist.*, 53(1):322–357, 02 2017.

[65] N. Curien and J.-F. Le Gall. First-passage percolation and local modifications of distances in random triangulations. *Annales Scientifiques de l'ENS*, to appear.

[66] N. Curien and C. Marzouk. How fast planar maps get swallowed by a peeling process. *Electron. Commun. Probab.*, 23:11 pp., 2018.

[67] N. Curien and L. Ménard. The skeleton of the UIPT, seen from infinity. *arXiv:1803.05249*, Mar. 2018.

[68] N. Curien, L. Ménard, and G. Miermont. A view from infinity of the uniform infinite planar quadrangulation. *Lat. Am. J. Probab. Math. Stat.*, 10(1):45–88, 2013.

[69] N. Curien and G. Miermont. Uniform infinite planar quadrangulations with a boundary. *Random Structures & Algorithms*, 47(1):30–58.

[70] F. den Hollander. Probability theory: The coupling method. *http://websites.math.leidenuniv.nl/probability/lecturenotes/CouplingLectures.pdf*.

[71] R. Doney. The Martin boundary and ratio limit theorems for killed random walks. *Journal of the London Mathematical Society*, 58(3):761–768, 1998.

[72] P. G. Doyle and J. L. Snell. *Random walks and electric networks*, volume 22 of *Carus Mathematical Monographs*. Mathematical Association of America, Washington, DC, 1984.

[73] H. Duminil-Copin, C. Garban, and G. Pete. The near-critical planar FK-Ising model. *Comm. Math. Phys.*, 326(1):1–35, 2014.

[74] B. Duplantier and S. Sheffield. Duality and the Knizhnik-Polyakov-Zamolodchikov relation in Liouville quantum gravity. *Phys. Rev. Lett.*, 102(15):150603, 4, 2009.

[75] B. Durhuus, T. Jonsson, and J. Wheater. On the spectral dimension of causal triangulations. *Journal of Statistical Physics*, 139:859–881, 2010.

[76] W. Esty. The reverse Galton-Watson process. *J. of Applied Prob.*, 12:574–580, 1975.

[77] C. Garban, G. Pete, and O. Schramm. The scaling limits of near-critical and dynamical percolation. *J. of European Math. Society*, 20:1195–1268, 2018.

[78] M. Gorny, É. Maurel-Segala, and A. Singh. The geometry of a critical percolation cluster on the UIPT. *Ann. Inst. H. Poincaré Probab. Statist.*, to appear.

[79] I. P. Goulden and D. M. Jackson. *Combinatorial enumeration*. A Wiley-Interscience Publication. John Wiley & Sons Inc., New York, 1983. With a foreword by Gian-Carlo Rota, Wiley-Interscience Series in Discrete Mathematics.

[80] I. P. Goulden and D. M. Jackson. The KP hierarchy, branched covers, and triangulations. *Adv. Math.*, 219(3):932–951, 2008.

[81] M. Gromov. Hyperbolic manifolds, groups and actions. In *Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978)*, volume 97 of *Ann. of Math. Stud.*, pages 183–213. Princeton Univ. Press, Princeton, N.J., 1981.

[82] O. Gurel-Gurevich and A. Nachmias. Recurrence of planar graph limits. *Ann. Maths*, 177(2):761–781, 2013.

[83] T. Hutchcroft and Y. Peres. Boundaries of planar graphs: a unified approach. *Electron. J. Probab.*, 22, 2017.

[84] J. Jurkiewicz, A. Krzywicki, and B. Petersson. A numerical study of discrete euclidean Polyakov surfaces. *Physics Letters B*, 168(3):273–278, 1986.

[85] V. A. Kaimanovich and A. M. Vershik. Random walks on discrete groups: boundary and entropy. *Ann. Probab.*, 11(3):457–490, 1983.

[86] V. Kazakov, I. Kostov, and A. Migdal. Critical properties of randomly triangulated planar random surfaces. *Physics Letters B*, 157(4):295 – 300, 1985.

[87] H. Kesten. Aspects of first passage percolation. In *École d'été de probabilités de Saint-Flour, XIV—1984*, volume 1180 of *Lecture Notes in Math.*, pages 125–264. Springer, Berlin, 1986.

[88] H. Kesten. Subdiffusive behavior of random walk on a random cluster. *Ann. Inst. H. Poincaré Probab. Statist.*, 22(4):425–487, 1986.

[89] H. Kesten. Scaling relations for 2D-percolation. *Comm. Math. Phys.*, 109(1):109–156, 1987.

[90] V. G. Knizhnik, A. M. Polyakov, and A. B. Zamolodchikov. Fractal structure of 2D-quantum gravity. *Modern Phys. Lett. A*, 3(8):819–826, 1988.

[91] H. Komuro. The diagonal flips of triangulations on the sphere. *Yokohama Math. Journal*, 44(2):115–122, 1997.

[92] M. Krikun. Local structure of random quadrangulations. *arXiv:0512304*.

[93] M. Krikun. On one property of distances in the infinite random quadrangulation. *arXiv:0805.1907*.

[94] M. Krikun. A uniformly distributed infinite planar triangulation and a related branching process. *Zap. Nauchn. Sem. S.-Peterburg. Otdel. Mat. Inst. Steklov. (POMI)*, 307(Teor. Predst. Din. Sist. Komb. i Algoritm. Metody. 10):141–174, 282–283, 2004.

[95] M. Krikun. Explicit enumeration of triangulations with multiple boundaries. *Electron. J. Combin.*, 14(1):Research Paper 61, 14 pp. (electronic), 2007.

[96] J.-F. Le Gall. Random geometry on the sphere. *Proceedings of the ICM 2014*.

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

[98] J.-F. Le Gall. Geodesics in large planar maps and in the Brownian map. *Acta Math.*, 205:287–360, 2010.

[99] J.-F. Le Gall. Uniqueness and universality of the Brownian map. *Ann. Probab.*, 41:2880–2960, 2013.

[100] J.-F. Le Gall. The Brownian cactus II: upcrossings and local times of super-Brownian motion. *Probability Theory and Related Fields*, 162(1):199–231, Jun 2015.

[101] J.-F. Le Gall and T. Lehéricy. Separating cycles and isoperimetric inequalities in the uniform infinite planar quadrangulation. *Ann. of Probab.*, to appear.

[102] J.-F. Le Gall and G. Miermont. Scaling limits of random planar maps with large faces. *Ann. Probab.*, 39(1):1–69, 2011.

[103] J.-F. Le Gall and F. Paulin. Scaling limits of bipartite planar maps are homeomorphic to the 2-sphere. *Geom. Funct. Anal.*, 18(3):893–918, 2008.

[104] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times*. Providence, R.I. American Mathematical Society, 2009.

[105] S. Lin. Harmonic measure for biased random walk in a supercritical Galton-Watson tree. *arXiv:1707.01811*, July 2017.

[106] R. Lyons. Random walks and percolation on trees. *Ann. Probab.*, 18(3):931–958, 1990.

[107] R. Lyons. Random walks, capacity and percolation on trees. *The Annals of Probability*, 20(4):2043–2088, 1992.

[108] R. Lyons, R. Pemantle, and Y. Peres. Ergodic theory on Galton-Watson trees: speed of random walk and dimension of harmonic measure. *Ergodic Theory Dynam. Systems*, 15(3):593–619, 1995.

[109] R. Lyons, R. Pemantle, and Y. Peres. Biased random walks on Galton-Watson trees. *Probab. Theory Related Fields*, 106(2):249–264, 1996.

[110] R. Lyons and Y. Peres. *Probability on Trees and Networks*. Cambridge University Press, 2017.

[111] J.-F. Marckert and G. Miermont. Invariance principles for random bipartite planar maps. *Ann. Probab.*, 35(5):1642–1705, 2007.

[112] C. Marzouk. Scaling limits of random bipartite planar maps with a prescribed degree sequence. *Random Structures and Algorithms*, 2018.

[113] L. McShine and P. Tetali. On the mixing time of the triangulation walk and other Catalan structures. In *Randomization Methods in Algorithm Design, Proceedings of a DIMACS Workshop, Princeton, New Jersey, USA, December 12-14, 1997*, pages 147–160, 1997.

[114] G. Miermont. Aspects of random maps. *preprint, http://perso.ens-lyon.fr/gregory.miermont/coursSaint-Flour.pdf*.

[115] G. Miermont. On the sphericity of scaling limits of random planar quadrangulations. *Electron. Commun. Probab.*, 13:248–257, 2008.

[116] G. Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.*, 210(2):319–401, 2013.

[117] J. Miller and S. Sheffield. An axiomatic characterization of the Brownian map. *arXiv:1506.03806*, 2015.

[118] J. Miller and S. Sheffield. Liouville quantum gravity and the Brownian map I: The QLE(8/3,0) metric. *ArXiv e-prints*, July 2015.

[119] J. Miller and S. Sheffield. Liouville quantum gravity and the Brownian map II: geodesics and continuity of the embedding. *ArXiv e-prints*, May 2016.

[120] J. Miller and S. Sheffield. Liouville quantum gravity and the Brownian map III: the conformal structure is determined. *ArXiv e-prints*, Aug. 2016.

[121] M. S. Molloy, B. Reed, and W. Steiger. On the mixing rate of the triangulation walk. In *DIMACS Series in Discrete Math. Theor. Comput. Sci*, volume 43, pages 179–190, 1998.

[122] L. Ménard. Volumes in the uniform infinite planar triangulation: From skeletons to generating functions. *Combinatorics, Probability and Computing*, page 1–28, 2018.

[123] P. Nolin and W. Werner. Asymmetry of near-critical percolation interfaces. *J. Amer. Math. Soc.*, 22(3):797–819, 2009.

[124] G. Ray. Geometry and percolation on half planar triangulations. *Electron. J. Probab.*, 19(47):1–28, 2014.

[125] L. Richier. Limits of the boundary of random planar maps. *Probability Theory and Related Fields*, to appear.

[126] S. Rohde. Oded Schramm: From circle packing to SLE. *Ann. Probab.*, 39(5):1621–1667, 2011.

[127] G. Schaeffer. Conjugaison d'arbres et cartes combinatoires aléatoires. PhD thesis. 1998.

[128] S. Sheffield. Quantum gravity and inventory accumulation. *Ann. Probab.*, 44(6):3804–3848, 11 2016.

[129] R. Stephenson. Local convergence of large critical multi-type Galton–Watson trees and applications to random maps. *Journal of Theoretical Probability*, 31(1):159–205, Mar 2018.

[130] A.-S. Sznitman and M. Zerner. A law of large numbers for random walks in random environment. *Ann. Probab.*, 27(4):1851–1869, 10 1999.

[131] G. 't Hooft. A planar diagram theory for strong interactions. *Nuclear Physics B*, 72:461–473, 1974.

[132] W. T. Tutte. A census of Hamiltonian polygons. *Canad. J. Math.*, 14:402–417, 1962.

[133] W. T. Tutte. A census of planar triangulations. *Canad. J. Math.*, 14:21–38, 1962.

[134] W. T. Tutte. A census of slicings. *Canad. J. Math.*, 14:708–722, 1962.

[135] W. T. Tutte. A census of planar maps. *Canad. J. Math.*, 15:249–271, 1963.

[136] W. T. Tutte. On the enumeration of four-colored maps. *SIAM J. Appl. Math.*, 17:454–460, 1969.

[137] B. Virág. Anchored expansion and random walk. *Geom. Funct. Anal.*, 10(6):1588–1605, 2000.

[138] K. Wagner. Bemerkungen zum Vierfarbenproblem. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 46:26–32, 1936.

[139] Y. Watabiki. Construction of non-critical string field theory by transfer matrix formalism in dynamical triangulation. *Nuclear Phys. B*, 441(1-2):119–163, 1995.

[140] Y. Wen. The brownian plane with minimal neck baby universe. *Random Structures & Algorithms*, 51(4):729–752.

université PARIS-SACLAY | École doctorale de mathématiques Hadamard (EDMH)

Titre : Cartes aléatoires hyperboliques

Mots Clefs : probabilités, cartes aléatoires, hyperboliques.

Résumé : Cette thèse s'inscrit dans la théorie des cartes planaires aléatoires, et plus précisément dans l'étude de modèles de nature hyperbolique.
Dans un premier temps, nous nous intéressons à un modèle de triangulations aléatoires dynamiques de la sphère basé sur les flips d'arêtes, et nous montrons une borne inférieure sur le temps de mélange de ce modèle.
Dans la suite, l'objet d'étude principal est une famille de triangulations aléatoires hyperboliques, appelées PSHIT, qui sont des variantes de la triangulation uniforme du plan (UIPT). Nous établissons un résultat de limite d'échelle quasi-critique, et obtenons à la limite un nouvel espace métrique aléatoire appelé plan brownien hyperbolique. Nous étudions également des propriétés métriques fines des PSHIT et du plan brownien hyperbolique, et notamment la structure de leurs géodésiques infinies.
Enfin, nous nous intéressons à un autre modèle naturel de cartes aléatoires hyperboliques : les cartes causales surcritiques. Nous établissons des résultats d'hyperbolicité métrique, ainsi que des propriétés de la marche aléatoire sur ces cartes, dont un résultat de vitesse positive. Certaines des propriétés obtenues sont robustes, et peuvent se généraliser à n'importe quelle carte planaire contenant un arbre de Galton–Watson surcritique.

Title: Hyperbolic random maps

Keys words: probability, random planar maps, hyperbolic.

Abstract: This thesis deals with random planar maps, and more precisely with hyperbolic models.
We are first interested in a model of dynamical random triangulations of the sphere based on edge-flips, for which we prove a lower bound on the mixing time.
We then focus on a family of random hyperbolic triangulations called the PSHIT, which are variants of the Uniform Infinite Planar Triangulation (UIPT). We establish a near-critical scaling limit result, and obtain in the limit a new random metric space called the hyperbolic Brownian plane. We also study fine metric properties of the PSHIT and of the hyperbolic Brownian plane, and in particular the structure of their infinite geodesics.
Finally, we are interested in another natural model of hyperbolic random maps: supercritical causal maps. We establish metric hyperbolicity results, as well as properties of the simple random walk on these maps, including a result of positive speed. Some of the properties we obtain are robust, and may be generalized to any planar map containing a supercritical Galton–Watson tree.