Hyperbolic random maps

by

Gourab Ray

B.Stat., Indian Statistical Institute, 2008
M.Stat., Indian Statistical Institute, 2010

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate and Postdoctoral Studies

(Mathematics)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

July 2014

Abstract

Random planar maps have been an object of utmost interest over the last decade and half since the pioneering works of Benjamini and Schramm, Angel and Schramm and Chassaing and Schaeffer. These maps serve as models of random surfaces, the study of which is very important with motivations from physics, combinatorics and random geometry.

Uniform infinite planar maps, introduced by Angel and Schramm, which are obtained as local limits of uniform finite maps embedded in the sphere, serve as a very important discrete model of infinite random surfaces. Recently, there has been growing interest to create and understand hyperbolic versions of such uniform infinite maps and several conjectures and proposed models have been around for some time. In this thesis, we mainly address these questions from several viewpoints and gather evidence of their existence and nature.

The thesis can be broadly divided into two parts. The first part is concerned with half planar maps (maps embedded in the upper half plane) which enjoy a certain domain Markov property. This is reminiscent of that of the SLE curves. Chapters 2 and 3 are mainly concerned with classification of such maps and their study, with a special focus on triangulations. The second part concerns investigating unicellular maps or maps with one face embedded in a high genus surface. Unicellular maps are generalizations of trees in higher genera. The main motivation is that investigating such maps will shed some light into understanding the local limit of general maps via some well-known bijective techniques. We obtain certain information about the large scale geometry of such maps in Chapter 4 and about the local limit of such maps in Chapter 5.

Preface

The new results of this thesis are based primarily on four research articles: [9, 12, 81, 82]. Article [12] is joint work with Omer Angel. Paper [9] is joint work with Omer Angel, Guillaume Chapuy and Nicolas Curien. Articles [81, 82] are independent works of the author of this thesis.

The results of [12] are presented in Chapter 2, that of [82] in Chapter 3. The material of [81] is collected in Chapter 4. Finally the results in [9] form the basis of Chapter 5.

Table of Contents

Appendices

List of Figures

List of symbols and acronyms

A list of symbols and acronyms along with the page number where they appear for the first time is given below chapter wise.

Chapter 1

$B_r(G)$	The combinatorial ball of radius r around the root of a rooted graph G....page 4.
$\mathrm{Cat}(n)$	The nth catalan number: $\frac{1}{n+1}\binom{2n}{n}$....page 14.
$E(M)$	The set of edges of a map M....page 2.
$F(M)$	The set of faces of a map M....page 2.
HUIPT	Half-plane uniform infinite planar triangulations....page 6.
$\phi_{n,m}$	Number of triangulations without loops of an m-gon with n internal vertices....page 7.
SHIQ	Stochastic hyperbolic infinite quadrangulations....15.
$\mathcal{T}_{g,n}$	Set of triangulations of genus g with n vertices....page 16
UIPT	Uniform infinite planar triangulations....page 6.
UIPQ	Uniform infinite planar quadrangulations....page 6
$\mathcal{U}_{g,n}$	The set of unicellular maps of genus g with n edges....page 82.
$V(M)$	The set of vertices of a map M....page 2.
$\|X\|$	Cardinality of a set X...page 2.

Chapter 2

α	The probability of the event that the triangle incident to the root edge is incident to an internal vertex....page 21.
$A_{Q,k}$	The event defined in page 20.

β^{p-2}	The probability of the event that the face incident to the root edge consists of a single contiguous segment of length $p-2$ from the infinite boundary and no internal vertex..see page 30 (for $p=3$) and page 43.
γ^{p-2}	The probability of the event that the face incident to the root edge consists of a single edge from the infinite boundary and $p-2$ internal vertices...page 43.
$\mathcal{H}$	Generic class of half plane maps....page 19.
$\mathcal{H}'$	Generic class of simple faced half plane maps....page 19.
$\mathcal{H}_p$	The class of infinite one-ended half plane p-angulations....19
$\mathcal{H}'_p$	The class of simple faced infinite one-ended half plane p-angulations....19
$\mathbb{H}_\alpha$	Translation invariant and domain Markov measures supported on half-plane loopless triangulations with parameter α....page 21.
$\mathbb{H}_\alpha^{(p)}$	Translation invariant and domain Markov measures supported on half-plane simple faced p-angulations with parameter α....page 21.
$\mathbb{H}_{\alpha,q,\nu}$	Translation invariant and domain Markov measures supported on half-plane triangulations with loops allowed....page 38.
$\mu_{m,n}$	Uniform measure on the set of triangulations of an m-gon with n internal vertices....page 23.
$p_{i,k}$	The probability that the triangle incident to a fixed boundary edge e of a domain Markov half planar triangulation has its third vertex on the boundary at a distance i along the boundary to the right (or left) of e and that this triangle separates k internal vertices of the triangulation from infinity....page 29.
p_i	The probability that the triangle incident to a fixed boundary edge e of a domain Markov half planar triangulation has its third vertex on the boundary at a distance i along the boundary from e....page 36
$Z_m(x)$	The generating function of the number of p-angulations of an m-gon with weight x for every internal vertex....page 28 (for $p=3$) and page 44.

Chapter 3

$B_r(\alpha)$	The hull of the ball of radius r around the root of a map with law $\mathbb{H}_\alpha$....page 56.
$\mathsf{END}(\mathcal{C})$	The space of ends of a percolation cluster $\mathcal{C}$....page 61.
$I_m(q)$	Number of internal vertices of a freely distributed triangulation of an m-gon with parameter q.....page 62.
(L, i)	The event that the triangle incident to the root edge is attached to a vertex on the boundary which is at a distance i along the boundary to the left of the root edge....page 54.
$\partial B_r(\alpha)$	Internal boundary of the hull of the ball of radius r around the root of a map with law $\mathbb{H}_\alpha$....page 56.
$\partial_E S$	Set of edges with one vertex in S another in the complement of S....page 57.
P_n	The sub-map revealed after n steps of peeling....page 63.
$\mathbb{P}_p$	The overall law of the map along with Bernoulli(p) site percolation on it....page 60.
p_c	Critical percolation probability...page 60.
p_u	Infimum over p such that there exists a unique infinite percolation cluster for Bernoulli (p) percolation....page 60.
(R, i)	The event that the triangle incident to the root edge has its third vertex on the boundary which is at a distance i along the boundary to the right of the root edge....page 54.
$\|S\|_E$	Sum of the degrees of the vertices in S....page 57.
T_n	The complement of the revealed map after n steps of peeling....page 63.
S_n	$= V_n - X_n$....page 68.
τ_r	Number of peeling steps required to reveal the hull of the ball of radius r in the peeling algorithm....page 64.
V_n	Volume of the map revealed after n steps of peeling....page 68.
X_n	Internal boundary of the revealed map after n steps of peeling....page 65.

Chapter 4

$\mathcal{C}_k$	The event that collision has not occurred up to step k....page 112.
$\mathcal{C}_{g,n}$	The set of C-decorated trees of genus g with n edges....page 86.
$d^G(.,.)$	The graph distance metric in a graph G....page 82.
d_λ	The graph distance metric in the underlying graph of $T_\lambda(n)$...page 93.
δ	The step number when we stop the exploration process I (resp. exploration process II, stage 1)....page 95 (resp. 107).
$\mathrm{diam}(G)$	The diameter of a graph G....page 83.
$d_{TV}(X,Y)$	The total variation distance between X and Y....page 98.
E_r	Vertices revealed after round r of exploration....pages 94, 106.
GW	The Galton-Watson tree defined in page 110.
$I_{g,n}$	The local injectivity radius of $U_{g,n}$....page 85.
λ	An element from the set $\mathcal{P}$....page 87.
$m(v)$	The mark of v....page 87.
$mark(v)$	The set of vertices with the same mark as v....page 94.
M_r	$\max_{v\in V(U_{0,n})} \lvert B_r(v)\rvert$ where $B_r(v)$ is the combinatorial ball of radius r around v in $U_{0,n}$....page 91.
N_r	$= E_r \setminus E_{r-1}$: The set of active vertices in E_r....page 94.
$\mathcal{P}$	The set of ordered N-tuples of odd positive integers which add up to $n+1$....page 87.
$\Phi_{\sigma,e}$	Number of rooted embedded forests consisting of σ trees and e edges....page 100.
$\mathbb{P}_\lambda$	The conditional measure induced by T_λ.... page 89.
Q_k	The set of dead vertices revealed after k steps of exploration....page 106.
R_k	The subgraph revealed at the kth step of exploration....pages 94, 106.
S	$= (S_0, S_1, \dots S_{\delta'})$: The set of seeds revealed up to step δ....page 96.
$\widetilde{S}$	$= (\widetilde{S_0}, \widetilde{S_1}, \dots \widetilde{S_{\delta'}})$: An i.i.d. sample of uniformly picked vertices from $U_{0,n}$....page 98.

$\mathcal{T}_\lambda(n)$	The set of marked trees corresponding to the N-tuple λ....page 88.
$T_\lambda(n)$	A uniformly picked element from $\mathcal{T}_\lambda(n)$....page 88.
(t, m)	A marked tree with plane tree t and a marking function m....page 93.
T_C	The auxiliary tree defined in page 108.
τ_r	Number of steps needed to complete round r of exploration....pages 94, 106.
$\mathcal{U}_{g,n}$	The set of unicellular maps of genus g with n edges....page 82.
$U_{g,n}$	Uniformly picked element from $\mathcal{U}_{g,n}$.....page 82.
v_-	The v_*-ancestor which is not dead and nearest to v in the tree t....page 107.
Z_r	The number of vertices at a distance r from the root vertex of T_C...page 108.
Z_j^{GW}	The number of vertices at a distance j from the root vertex of GW..page 111.

Chapter 5

β_θ	The solution of β for (5.1.1)....page 115.
$\mathrm{Geom}(\xi)$	Geometric distribution with parameter ξ....page 115.
T_ξ^∞	Galton-Watson tree with geometric(ξ) -1 offspring distribution conditioned to be infinite....page 115.
$\mathcal{U}_{g,n}$	The set of unicellular maps of genus g with n edges....page 82.
$U_{g,n}$	Uniformly picked element from $\mathcal{U}_{g,n}$.....page 82.
X_β	Random Variable defined in page 116.

Acknowledgements

First of all, I would like to thank my supervisor Prof. Omer Angel for providing me invaluable support throughout my stay in UBC. He always answers my questions with utmost patience and care and is very generous about spending time with me discussing mathematics. Among many other things, he was the first to introduce me to the intriguing and beautiful world of random maps and I have been hooked ever since. Apart from discussing mathematics, he also introduced me to many wonderful mathematicians which also played a very important role in shaping my research career so far. He was also very generous in providing funding for my stay in UBC as well as for traveling to many interesting meetings and conferences. I would like to specially thank him for providing support for my trip to the summer school in Saint-Flour in 2011 where I had the privilege to meet Prof. Itai Benjamini for the first time and for the visit to Paris in 2013 where I had the chance to collaborate with Guillaume Chapuy and Nicolas Curien along with Omer himself.

Being part of the supremely active and dynamic probability group in UBC is an experience on its own. I have been part of a number of research level courses, reading seminars and projects which helped me immensely to grow as a researcher. Many thanks to my friends in the probability group: Roland Bauerschmidt, Yu Ting Chen, Tyler Helmuth , Tim Hulshof, Bálázs Rath, Dennis Timmers, Yuri Mejia Miranda, Alex Tomberg, Daniel Valesin, Tom Hutchcroft, Brett Kolesnik, Benjamin Wallace, Zichun Ye, Jing Yu and Qingsan Zhu for making my stay in UBC memorable. I would also like to thank professors Martin Barlow, David Brydges, Ed Perkins, Gordon Slade, Juan Souto and Asaf Nachmias for running some wonderful courses and reading seminars and also for patiently answering all my questions. I gained a lot from them. A special thanks should go out to Asaf Nachmias for listening to all my blabbering for the past three years. Thanks also to Jim Richardson and Asif Zaman for all the wonderful time we spent together in my first year here in UBC.

Simply put, life is meaningless without friends. I do not make close friends very easily, but nevertheless I now have two very close friends I first met in UBC: Nishant Chandgotia and Vasu Tewari. Thanks for all the wonderful time we spent together. Special thanks to Vasu for enduring my cooking skills for two years and helping me learn just enough French so that I could make my way to Saint-Flour. Special thanks to Nishant for forcing me to climb all those mountains. Thanks also to Subhajit Jana for all the interesting discussions we had time and again over the past year or so.

Thanks to all my friends and professors from ISI for providing me support and writing me recommendation letters which helped me come to UBC. Special thanks to Professors Krishanu Maulik and Amitesh Dasgupta for all the support during my phd applications. I would not have made anywhere near this far unless I was inspired by three of my closest friends from ISI: Sayan Banerjee, Ritabrata Das and Abhirup Datta. Thanks also to Pramita Bagchi, Somak Dutta, Subhajit Mandal, Swarnali Banerjee, Debashree Ray, Adityadeb Mukherjee and Tamoghna Roy for being such wonderful friends and bearing with me for so many years. I would also like to thank Sir Debashish Burman for creating the stepping stone of my mathematics career during my high school days.

I met Poulami for the first time at the beginning of my phd which should count as one of the significant events of my life. Thanks Poulami for always being with me and helping me through the tough times.

Finally, the biggest acknowledgement should go out to my parents from whom I first learnt about numbers. Without their constant support, sacrifice and love, nothing would have been possible. Dhonnobad Baba aar Maa.

Dedication

To my parents and all my friends.

Chapter 1

Introduction

1.1 Definitions

We begin with some basic definitions of the objects we shall work with.

Definition 1.1.1. *A* ***map*** *is a proper embedding of a connected (multi)graph on a compact orientable surface viewed up to orientation preserving homeomorphisms from the surface to itself. Further, the embedding is such that the complement of the embedding is a union of disjoint topological discs.*

Let us clarify few terms in Definition 1.1.1. By a proper embedding, we mean the embedding is such that the edges do not cross each other. From the classification of surfaces theorem (see [73], Theorem 1 or [32]), we know that every compact orientable surface can be viewed as a connected sum of a finite number of tori. Thus the reader can imagine the surfaces on which we embed the graphs to be just a connected sum of finite number of tori (which we sometimes refer to as handles) and the number of handles is called the **genus** of the underlying surface of the map. When the genus is 0, that is the underlying surface of the map is a sphere, the map is called a **planar map**. We shall think of a planar map as being embedded in the plane instead of the sphere. A major portion of this thesis is concerned with maps embedded in the half-plane, that is in $\mathbb{R} \times \mathbb{R}^+$ (more in Chapter 2).

Definition 1.1.1 is more topological in nature and there exist other equivalent ways to define maps. We refer the interested reader to [78] for a more fundamental approach. We also mention here that maps on non-orientable surfaces have also been studied (see [26, 55, 78]), but such maps are not the focus of this thesis and hence is not part of Definition 1.1.1.

For infinite graphs, there are certain extra assumptions about the embedding. A map is **one-ended** if the complement of any finite subset of the map has precisely one infinite connected component. A map is **locally finite** if all the vertex degrees are finite. We shall only consider maps which are **locally finite** and **one-ended** in this thesis. We shall also mark one oriented edge of the map as the **root edge**. Sometimes in the literature,

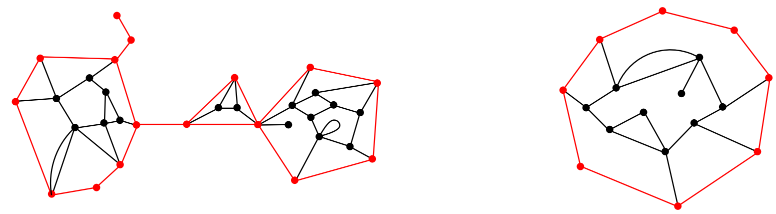

Figure 1.1: Left: A map with a boundary. Right: A map with a simple boundary. The boundary edges and vertices are marked red.

maps are defined to be face rooted or corner rooted (see [13, 34]). However, we shall stick to edge rooted maps in this thesis.

The connected components of the complement of the embedding are called **faces**. So as per Definition 1.1.1, the faces are homeomorphic to discs. Further, the graphs can have multiple edges or self-loops. The number of edges incident to a face is called the degree of a face. A face of degree 3 is called a triangle, that of degree 4 a quadrangle and so on. A map is called a **triangulation** if all its faces are triangles except possibly a face which is identified to be external to the map. A triangulation with no external face will be referred to as the triangulation of the surface. The collection of edges incident to the external face is called the **boundary** of the map (see Figure 1.1). Notice that the boundary of a map *a priori* might not be a simple cycle. If a map has m boundary vertices, n non-boundary vertices and the boundary forms a simple cycle then we shall refer to it as a triangulation of an m-gon with n internal vertices.

For a finite map M, $V(M), E(M)$ and $F(M)$ will denote the set of its vertices, edges and faces. For any finite set S, let $|S|$ denote its cardinality. For a finite map M with an underlying surface of genus g, recall Euler's formula for orientable surfaces:

$$|V(M)| - |E(M)| + |F(M)| = 2 - 2g$$

This formula can impose certain restriction of the genus on which a certain class of maps can be embedded. For example, if we consider a triangulation with n faces on a genus g surface, Euler's formula gives that the maximal genus surface on which it can be embedded is $\lfloor n/4 \rfloor$ (notice that for a triangulation T, we have $3|F(T)| = 2|E(T)|$). Similarly for a map with n edges and one face, the maximal genus possible is $\lfloor n/2 \rfloor$.

One can view suitable classes of maps as metric spaces with a suitable metric on them. There are two possible ways to do this. The first one is the local metric introduced by Benjamini and Schramm ([22]) which roughly says that two maps are close if the combinatorial balls of radius r are isomorphic for large r. The second is the Gromov-Hausdorff distance which views two finite maps as compact metric spaces induced by their appropriately rescaled combinatorial graph distances and the distance between the maps is computed by computing the Hausdorff distance between two isometric embeddings of the maps in a common metric space and then taking infimum over all such embeddings. After adding a suitable topology, we can also define probability measures on such measure spaces and consider weak limits of such measures as the class of maps we consider becomes large. Very broadly speaking, proving existence of such limits and studying their geometric properties have been a major theme of work in this field. The weak limit in the local metric is called the local limit and that of the properly rescaled metric in the Gromov-Hausdorff topology is called the scaling limit.

One of the simplest examples of scaling limits is perhaps the classical fact that the simple random walk in $\mathbb{Z}$ when properly rescaled converges to one dimensional standard Brownian motion. To be more precise let $\{X_n\}_{n\geq 1}$ be the simple random walk in $\mathbb{Z}$. Then $\frac{1}{\sqrt{n}}(X_{\lfloor nt \rfloor})_{0\leq t\leq 1}$ converges weakly to $(B_t)_{0\leq t\leq 1}$ in the appropriate topology where B_t is the standard Brownian motion at time t (see [29] for more.) Scaling limits will not be the focus of this thesis but still let us mention that proving existence of scaling limits for natural models of random graphs is an area of extensive ongoing research. The theme of such works is that the graphs are viewed as properly rescaled metric spaces and the goal is to find a continuum object as a limit of such metric spaces in the above mentioned Gromov-Hausdorff topology. For example, the scaling limits of Erdos-Renyi random graphs are obtained in [1], the random tree in [4], minimal spanning trees in [2], random dissections in [40] etc. For random triangulations and $2p$-angulations for $p \geq 1$, scaling limits are obtained by LeGall ([70]) and the limiting object (which is universal for all these class of maps and is conjectured to hold for odd p-angulations as well) is called the Brownian map which forms the continuum analogue of the uniform infinite maps. Independently, Miermont in [76] has obtained the scaling limit results for quadrangulations using a different approach. A fine exposition about the scaling limits of random maps can be found in [75]. Many basic properties of the Brownian map are still unknown and this is an area of growing interest and extensive research.

1.2 The local topology

Let $\mathcal{G}^*$ denote the space of all connected, locally finite rooted graphs. Then $\mathcal{G}^*$ is endowed with the **local topology**, where two graphs are close if large balls around their corresponding roots are isomorphic. The local topology is generated by the following metric: for $G, G' \in \mathcal{G}^*$, we define

$$d(G,G') = (R+1)^{-1} \qquad \text{where} \qquad R = \sup\{r : B_r(G) \cong B_r(G')\}.$$

Here B_r denotes the ball of radius r around the corresponding roots, and $\cong$ denotes isomorphism of rooted maps. This metric on $\mathcal{G}^*$ is non-Archimedian. Finite graphs are isolated points, and infinite graphs are the accumulation points. Sometimes e^{-R} is used instead of $(R+1)^{-1}$ in the above definition of the metric, but the actual metric is of little importance and we mainly care about the topology induced by it.

The local topology on maps induces a weak topology on measures on $\mathcal{G}^*$. One such natural choice of measures is the Benjamini-Schramm limit of a sequence of finite planar graphs [22], which is the weak limit of the laws of these graphs with a uniformly chosen root edge. The work of Aldous and Steele in the context of trees in [5] also deserves mention here. Let us give a few examples.

- The $n \times n$ triangular (or square) lattice with the root chosen uniformly converges to the infinite triangular (or square) lattice. This follows from the observation that the cardinality of the set of boundary edges is small compared to the total volume. Hence with high probability, the root is not chosen from an edge close to the boundary in the finite graphs

- The binary tree of level n with the root chosen uniformly converges to what is popularly known as the *canopy tree*. This is because with high probability, the root is near the highest level rather than at the lower levels (see Figure 1.2).

- The Erdos-Renyi random graph with n vertices where two vertices are connected by an edge independently with probability c/n converges to the Galton-Watson tree with a Poisson(c) offspring distribution except the root vertex which has size biased Poisson(c) offspring distribution. Notice that the degree of the root in the limit is obtained by degree biasing the limiting degree of a vertex in the Erdos-Renyi graph because we choose the root vertex with probability proportional to its degree.

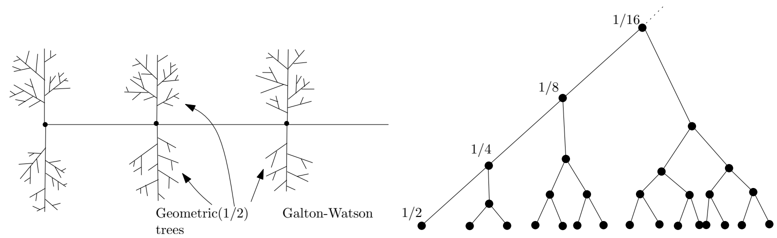

Figure 1.2: Left: The critical Galton-Watson tree with Geometric$(1/2)$ offspring distribution conditioned to survive. Right: A canopy tree which is the local limit of binary trees cut off at the nth level. The numbers beside the vertices represent the probability of them being the root. For example the probability of the leftmost vertex in the bottom level has probability $1/2$ of being the root and so on.

- Consider a uniform rooted (embedded) tree with n edges (see [69] for the precise formalism). The local limit as $n \to \infty$ is an object known as the **critical Galton-Watson tree conditioned to survive** with Geometric(1/2) offspring distribution (see [63]). The degree of the root is given by a size biased Geometric (1/2) offspring distribution. It has an infinite *spine* with critical Galton-Watson trees attached on both sides (see Figure 1.2)

We can endow any class of planar maps with the above local topology (with the graph isomorphism replaced by map isomorphism). Physicists were interested in random surfaces and their geometry. More about this background motivation is in Section 1.3.

Let T_n be uniformly distributed among all triangulations of the sphere with n faces. Then the weak limit of T_n as $n \to \infty$ forms a natural candidate for a discretized version of random surfaces. Its existence was proved by Angel and Schramm.

Theorem 1.2.1 (Angel, Schramm [13])**.** *The weak limit of T_n exists. The limiting object is an infinite, one-ended and planar triangulation.*

Notice that the completion of the space of all triangulations of the sphere with respect to the local topology *a priori* might not be one-ended: a limit of a sequence of finite triangulations on the sphere may contain more than one accumulation points. However, Theorem 1.2.1 ensures that the limiting measure is supported on one-ended infinite triangulations of the sphere. The limit is popularly known as the **uniform infinite planar triangulation** or the UIPT in short. Similarly, existence of the local limit for quadrangulations (called UIPQ) was proved by Krikun (see [65]).

Frequently we shall encounter measures of the following type. Let $T_{m,n}$ be a uniform distribution on triangulations of an m-gon with n internal vertices. Then it is known that

$$T_{m,n} \xrightarrow[n\to\infty]{} T_{m,\infty} \xrightarrow[m\to\infty]{} T_{\infty,\infty}.$$

The first limit is an infinite triangulation in an m-gon, and the second limit is known as the **half-plane uniform infinite planar triangulation** (HUIPT). The same limits exist for quadrangulations (yielding the half-plane UIPQ, see e.g. [44]) and many other classes of maps. These half-plane maps have the interesting property that if we remove any simply connected sub-map Q connected to the boundary, the distribution of the remaining map after removing Q is the same as the original distribution of the half plane map we started with. We call this property **domain Markov** and define this precisely in Chapter 2. The name is chosen in analogy with the related conformal domain Markov property that Schramm-Loewner evolution (SLE) curves have (a property which was central to the discovery of SLE [85]). This property appears in some forms also in the physics literature [6], and more recently played a central role in several works on planar maps [7, 10, 20].

1.3 The big picture

The motivation for studying random planar maps comes from several directions. This section is devoted to give a rough sketch of the bigger picture in the background and the history associated with this area of research.

1.3.1 Enumeration

Enumeration results of planar maps go back to the work of Tutte in the 60's. The main aim was to prove the four color theorem and the strategy was to show that the number of planar maps and the number of four colorable maps are the same. An array of results in this regard (see [31, 87–89]) were obtained by solving functional equations involving generating functions. Generalization to maps on general surfaces were also obtained (see [54, 55]). We roughly present below the general idea for this technique of enumeration. Suppose we wish to enumerate the number of triangulations with n faces and an external face of degree m edges rooted on the boundary. We define the following generating function:

$$F(x,y) = \sum x^{\#\text{faces}} y^{\#\text{edges on the boundary}}$$

Now the triangle incident to the root edge has the third vertex (the vertex which do not belong to the root edge) either on the boundary or the third vertex is an internal vertex of the triangulation. If the third vertex is a boundary vertex, it breaks up the triangulation into two independent triangulations with a smaller number of faces and smaller boundary sizes. Similarly, if the third vertex is an internal vertex, then the remaining triangulation has one less face and one more boundary edge. This gives us the required functional equation for $F(x,y)$. Solving such an equation requires a technique called the *quadratic method* (see [87–89] for details). The following proposition is an example of many results of similar flavor obtained using this technique.

Proposition 1.3.1. *For $n, m \geq 0$, the number of rooted triangulations of a disc with $m+2$ boundary vertices and n internal vertices is*

$$\phi_{n,m+2} = \frac{2^{n+1}(2m+1)!(2m+3n)!}{m!^2 n!(2m+2n+2)!} \tag{1.3.1}$$

Note that this formula is for triangulations with multiple edges allowed, but no self-loops (type II in the notations of [13]). The case of $\phi_{0,2}$ requires special attention. A triangulation of a 2-gon must have at least one internal vertex so there are no triangulations with $n = 0$, yet the above formula gives $\phi_{0,2} = 1$. This is reconciled by the convention that if a 2-gon has no internal vertices then the two edges are identified, and there are no internal faces.

This makes additional sense for the following reason: Frequently a triangulation of an m-gon is of interest not on its own, but as part of a larger triangulation. Typically, it may be used to fill an external face of size m of some other triangulation by gluing it along the boundary. When the external face is a 2-gon, there is a further possibility of filling the hole by gluing the two edges to each other with no additional vertices. Setting $\phi_{0,2} = 1$ takes this possibility into account.

Using Stirling's formula, the asymptotics of $\phi_{n,m}$ as $n \to \infty$ are easily found to be

$$\phi_{n,m} \sim C_m n^{-5/2} \left(\frac{27}{2}\right)^n. \tag{1.3.2}$$

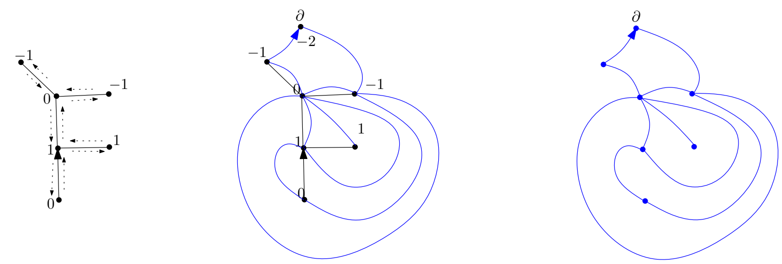

Figure 1.3: An illustration of the Schaeffer bijection. Left: Figure shows the way we move along the contour of the tree with 6 vertices starting from the root vertex. This also introduces a natural contour order between the vertices. Middle: Observe that the minimal label in the tree is -1. We add a new vertex ∂ with label -2. For each vertex encountered while moving around the contour, draw an arc between the vertex and the next vertex in the contour order with label one less. Draw an arc between vertices with label -1 and the vertex ∂. Orient the edge between the root vertex and ∂ in one of the two ways. Right: The rooted quadrangulation, pointed at ∂ with 7 vertices obtained by erasing the tree.

where C_m is given by

$$C_{m+2} = \frac{\sqrt{3}(2m+1)!}{4\sqrt{\pi}m!^2}\left(\frac{9}{4}\right)^m \sim Cm^{1/2}9^m.$$

The asymptotics in the above equation is obtained by using Stirling's formula. The power terms $n^{-5/2}$ and $m^{1/2}$ are common to many classes of planar structures. They arise from the common observation that a cycle partitions the plane into two parts (Jordan's curve Theorem) and that the two parts may generally be triangulated (or for other classes, filled) independently of each other. Similar enumeration results for other classes of planar maps were also obtained.

Later, bijective techniques were introduced due to Schaeffer *et al* (see [37, 84]) to find different interpretations of enumeration formulas like Proposition 1.3.1. The main theme of the bijections is to find a correspondence between classes of planar maps and tree-like objects maybe with some additional decorations on the vertices and edges. For example, the simplest of such results is a bijection initially worked out by Cori and Vauquelin ([37]) and developed later by Schaeffer ([84]). The bijection popularly used is that between rooted quadrangulations with n faces and a distinguished vertex (sometimes called a *rooted pointed quadrangulation*) and labelled rooted trees with n edges. A labelled tree is a rooted tree where each vertex is assigned a label such that the absolute value of the difference between the labels in two adjacent vertices is at most 1 and the root vertex is assigned label 0. The prescription to obtain the quadrangulation from the tree is illustrated in Figure 1.3.

These beautiful bijections provide us with a method to analyze planar maps via analyzing the corresponding labelled trees. For example the bijection illustrated in Figure 1.3 immediately gives that the number of rooted quadrangulations of the sphere with n faces is given by

$$\frac{2}{(n+2)}\frac{3^n}{(n+1)}\binom{2n}{n}$$

where the factor $\binom{2n}{n}\frac{1}{n+1}$ counts the number of trees with n edges, 3^n counts the number of labellings, the factor $(n+2)^{-1}$ comes from the fact that the quadrangulation can be pointed in one of its $n+2$ vertices and the factor 2 comes from orienting the root in one of the two directions. Furthermore, the labels on the vertices are the distances between the vertices and the pointed vertex up to an added constant. This enables us to analyze the distances in the map by looking at the labelling function. See [75] for a comprehensive survey of this approach.

1.3.2 Universality

The polynomial correction term $n^{-5/2}$ in the asymptotic formula in (1.3.2) also appears if we replace triangulations by quadrangulations or any reasonable class of maps. This is no coincidence and suggests that the large scale properties of random maps are universal and do not depend upon the local properties. The results in [70, 76] which depicts the universality of the Brownian Map illustrates this fact. This sort of universality results and conjectures are similar flavor to convergence of random walk to the Brownian motion for a wide class of walks.

Simple random walk is also conjectured to behave similarly in the local limit of random maps. Here is a conjecture which has been open for quite some time. Let $X_0, X_1, \ldots$ be the simple random walk on the UIPT or UIPQ.

Conjecture 1.3.2. *Show that almost surely, the graph distance between X_n and X_0 grows like $n^{1/4}$ up to poly-logarithmic fluctuations.*

Note that the exponent of 1/4 is not yet proved and the best known upper bound is 1/3 for the UIPQ obtained in [20]. Recently, the work of Gurel-Gurevich and Nachmias (see [59]) shows that if we consider a sequence of finite planar graphs (might be random) where the root is chosen uniformly, and if the distribution of the root vertex has exponential tail, then the local limit is recurrent. Previously, the same result was proved for bounded degree graphs by Benjamini and Schramm in [22].

If we enter the world of conformal embeddings of random maps then things are much less clear. A finite triangulation can be embedded in the sphere in many ways such that the embedding is conformally invariant. One popular way is to use the circle packing theory. For precise questions, we refer to [18], Section 3.2. Very recently, some information about the conformal structure of the boundary of the half planar maps have been obtained in [38]. The very recent work in [77] also deserves mention here.

1.3.3 Quantum gravity

The motivation to understand random surfaces, keeping aside the aesthetic interests, also comes from the physics literature. Developing the theory of two dimensional quantum gravity requires one to extend the Feynman path integrals to integrals over surfaces. This motivates the study of random surfaces and allows one to wonder: what is the geometry of a large typical surface? If we look into the discrete world of lattices, there is *a priori* no particular reason to prefer the $\mathbb{Z}^2$ lattice over say the triangular lattice or the hyperbolic plane.

One way to make sense of the above question is to consider uniform measures on all possible finite surfaces and then take weak limits of such measures. For example, all possible triangulations of a surface can be viewed as a discretized version of a 2-dimensional manifold and the local limit of such objects can be viewed as a discretized model of an infinite random surface. Hence, Theorem 1.2.1 suggests that the UIPT or UIPQ are natural candidates for such a model.

One can consider many models of statistical mechanics (for example the Ising model or the Bernoulli percolation model) on a random surface. A connection between dimensions of random fractal objects (for example the critical Bernoulli percolation cluster or the critical Ising model clusters) in the fixed surface (for example, $\mathbb{Z}^2$) and a random surface has been conjectured in [64]. The formula is

$$\Delta - \Delta_0 = \frac{\Delta(1-\Delta)}{k+2}$$

where $2 - 2\Delta$ is the dimension of the random fractal object in a random surface and $2 - 2\Delta_0$ is the dimension of the random fractal object in the deterministic surface. The parameter k depends upon the model of statistical mechanics in question. Some progress in this area has been made using the multiplicative cascades (see [23]) and the Gaussian free field (see [47]). On models of random surfaces such as the UIPQ, such a formula has been verified for pioneer points of random walk (see [20]) and that of the Bernoulli percolation critical cluster (see [10]). However, much is not known in this area as of now.

1.4 Percolation on random maps

Bernoulli percolation is one of the simplest models of statistical mechanics that one can consider. This is the main motivation behind understanding behavior of percolation clusters on random maps. Recall that the Bernoulli(p) site percolation model on a graph is defined as follows: for every vertex in the graph is colored black with probability p or white with probability $1-p$. We consider the percolation model on random graphs and we are interested in quenched statements about the percolation model. This means we are interested in statements about the percolation clusters conditioned on the map. As usual, the critical percolation probability p_c is the infimum over p such that for almost all triangulations there exists an infinite cluster with probability 1 for Bernoulli(p) site percolation on the triangulation.

Recall that the dual graph of a planar map is a graph with a vertex corresponding to each face and two vertices are joined by an edge if the corresponding faces in the primal graph have a common adjacent edge (see Figure 1.4). Site percolation on triangulations are nicer to analyze in the sense that we can consider the interface between two clusters. The interface is a path in the dual graph, that uses exactly the dual edges of edges with differently colored endpoints. Since faces are triangles, dual vertices have degree 3, and either two or none of the three dual edges are part of an interface. Thus the interfaces can not intersect each other. In a triangulation, the interfaces form cycles, simple paths or infinite lines. The simple paths must start at the boundary and end at the boundary.

Theorem 1.4.1 (Angel [7, 8])**.** *Almost surely, the critical site percolation probability on the UIPT is* $1/2$. *Almost surely, the critical site percolation probability for the half plane UIPT is also* $1/2$.

Thus, the site percolation probability matches with that of the triangular grid. The structure of the critical percolation cluster around the root vertex

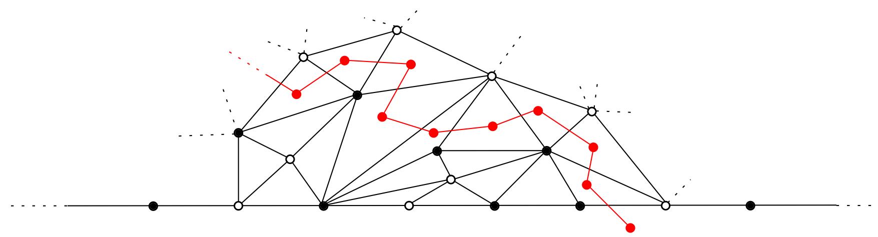

Figure 1.4: An illustration of an interface between a black and a white cluster for percolation in a half planar triangulation.

for half plane models as well as bond and face percolation models on the half plane UIPT is studied in [10]. The scaling limit of the critical interface in the full plane UIPT is studied in [41]. Some similar results for uniform infinite planar maps were obtained in [74].

1.5 Unicellular maps

Unicellular maps are maps with a single face. In genus 0, maps with a single face are nothing but plane embedded trees. So unicellular maps can be thought of as a generalization of plane trees in higher genera. The set of all unicellular maps of genus g and n edges will be denoted by $\mathcal{U}_{g,n}$.

In this thesis, random permutations will play a crucial role in Chapter 4. So there will be two notions of cycles floating around: one for cycle decomposition of permutations and the other for maps and graphs. To avoid confusion, we shall refer to a cycle in the context of graphs as a **circuit**. A circuit in a planar map is a subset of its vertices and edges whose image under the embedding is topologically a closed loop. A circuit is called contractible if its image under the embedding on the surface can be contracted to a point and not contractible otherwise. Recall that since we wanted the complement of the embedding to be disjoint union of discs, in higher genus the underlying graph of unicellular maps must contain circuits. Also notice that every circuit in a unicellular map is non-contractible.

There are alternate equivalent ways of defining an unicellular map. One way is to take a polygon with $2n$ edges and then glue together pairs of edges in the boundary of the polygon (to ensure orientability, one must glue together edges in the opposite directions in any cyclic orientation of the polygon.)

One can iteratively delete all the leaves of a unicellular map and then *erase* the degree 2 vertices as in Figure 1.6 and obtain a unicellular map

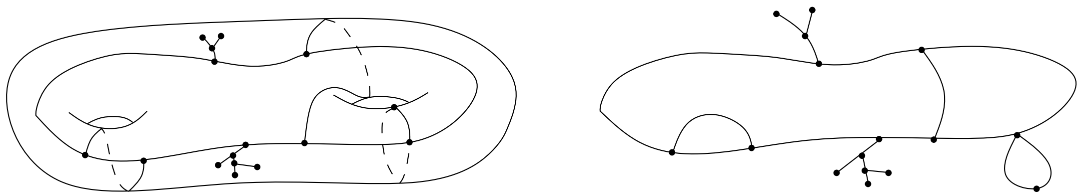

Figure 1.5: A unicellular map and its underlying graph.

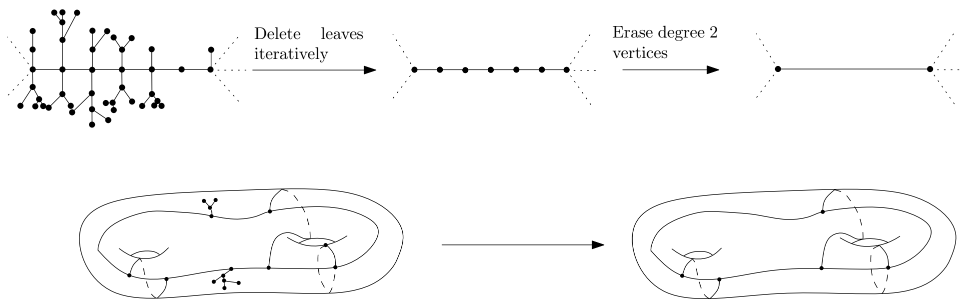

Figure 1.6: The process of obtaining a scheme of a unicellular map.

with each vertex degree at least 3. Such a map is called the **scheme** or the skeleton of the unicellular map. Clearly, the scheme carries all the information about the topology of the unicellular map. Thus, a unicellular map can be decomposed into its scheme and a forest corresponding to every edge in the scheme (see Figure 1.6).

We can explore the contour of a unicellular map in the same way as we do in a plane tree: we start from any corner, and continue exploring the corners of the face until we come back to the corner we started with (see Figure 1.7). This gives a cyclic order (called the face order) to the half-edges in the unicellular map. However observe that at every vertex, the half edges adjacent to it can be ordered in an anti clockwise manner which we call the vertex order. Notice that for a plane tree, the vertex order of edges around a vertex always coincides with their face order. However, in a higher genus unicellular map, the face order of edges around a vertex may not coincide with their vertex order giving rise to *intertwinings*. These intertwinings in some sense carry the information of the underlying topology of the unicellular maps.

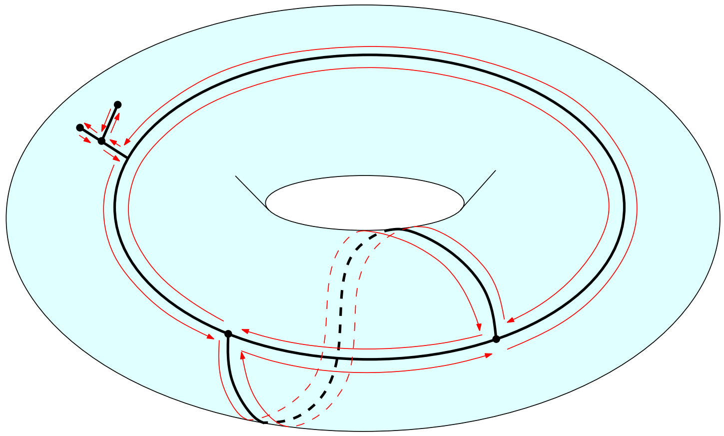

Figure 1.7: The way of exploring the face in unicellular map.

One can think of starting from a plane tree then iteratively gluing together vertices in the correct fashion so that the number of faces remain 1 at any step. Using Euler's formula it is easy to see that this gluing gives rise to a higher genus unicellular map. Using this heuristic, Chapuy in [33] obtained a recursive formula for the number of unicellular maps of genus g which is recursive only in the genus.

$$2g|\mathcal{U}_{g,n}| = \sum_{p=0}^{g-1} \binom{n+1-2p}{2g-2p+1} |\mathcal{U}_{p,n}|$$

The idea of iterative gluing can also be used to recover the enumeration formula (see [33])

$$|\mathcal{U}_{g,n}| = \mathrm{Cat}(n) \frac{1}{2^{2g}} \sum_{\lambda \in \mathcal{L}_{n+1-2g}(n+1)} \frac{(n+1)!}{\prod_i m_i! i^{m_i}}, \tag{1.5.1}$$

where $\mathrm{Cat}(n) = \frac{1}{n+1}\binom{2n}{n}$ is the nth Catalan number and the sum is over all the partitions λ of $n+1$ of size $n+1-2g$ such that each part has odd number of elements and m_i is the number of parts with i elements. The enumeration formula in eq. (1.5.1) is known as the Lehmann-Walsh formula [91].

This idea of obtaining a unicellular map of higher genus by gluing together vertices in a lower genus unicellular map was further simplified in [34] and a very simple bijection between unicellular maps and special classes of objects called C-decorated trees were obtained. This bijection in [34] forms the main tool for obtaining the results in Chapters 4 and 5.

We finish this subsection by mentioning that unicellular maps have appeared frequently in the field of combinatorics in the past few decades. They are related to representation theory of symmetric group, permutation factorization, matrix integrals computation and also the general theory of enumeration of maps. See the introduction section of [25, 33] for a nice overview and see [68] for connections to other areas of mathematics and references therein.

1.6 Hyperbolic random maps

There has been a growing interest in constructing hyperbolic versions of the uniform infinite planar maps. The uniform infinite planar maps are parabolic in the sense that they can be conformally embedded in the full plane (see [22, 61]). Is there a way to naturally construct a random map which looks more like the hyperbolic plane?

One construction of such a possible measure supported on infinite hyperbolic maps which can be found in [19]. Consider the supercritical Galton-Watson tree and condition it to survive to obtain an infinite tree. Now for every edge in the tree, attach independent edge weights uniform on $\{-1, 0, 1\}$. Assign label 0 to the root vertex and the label of any vertex v can be obtained by summing the edge weights along the geodesic joining the root vertex and v.

Notice that there are infinitely many faces in the tree and in every face, the corners can be given a natural order corresponding to exploring the corners in a clockwise direction similarly as in Figure 1.3. For every corner c with label l, draw an edge between the vertex corresponding to c and the first corner in the clockwise order whose associated vertex has label $l-1$. Now consider the quadrangulation obtained only by the edges added between vertices in the tree. Such a quadrangulation is named **Stochastic hyperbolic infinite quadrangulation** or SHIQ. The fact that the edges are non-crossing and the resulting map is a quadrangulation is a consequence of the results in [36].

It is easy to see that the SHIQ is almost surely a locally finite, one ended quadrangulation with exponential volume growth. Many question about the SHIQ remain open, see [19] Section 6.3 for details.

A half planar version of hyperbolic triangulations was constructed in [12] which forms the basis of Chapter 2. The geometry of these maps were further studied in [82] and it was shown that these maps possess exponential volume growth and anchored expansion. Recently in [11], the speed of simple random walk and return probabilities are also studied on such maps.

Another recent construction of planar hyperbolic triangulation called **stochastic hyperbolic planar triangulation** was obtained by Curien in [39]. He obtained a one-parameter family $\boldsymbol{T}_\kappa$ where the parameter $\kappa \in [0, 2/27)$ with $\kappa = 2/27$ corresponding to the UIPT. These objects are constructed using some of the ideas in Chapter 2. Results about volume growth and random walk speed are also obtained for these maps in [39]. The natural question that arises from these constructions is to obtain finite approximations for such maps analogous to the uniform infinite planar maps.

The idea is to look into random maps in a high genus surface. Gamburd and Makover showed (see [52, 53]) that if we choose a random triangulation with n faces and no restriction on the genus, typically, the genus is the maximal possible which is $\lfloor n/4 \rfloor$. The conjecture is that if we restrict to a genus Cn where $0 < C < 1/4$, then the local limit obtained is topologically planar and is hyperbolic. The intuition behind such a conjecture is that in higher genus triangulations, the average degree is higher than 6, which gives rise to negative curvature in the limiting maps, provided the distributional limit is planar.

Conjecture 1.6.1 below appears in this form in [39]. Fix $\theta \geq 0$ and suppose we look at $\mathcal{T}_{\lfloor \theta n \rfloor, n}$ which is the set of rooted triangulation of genus $\lfloor \theta n \rfloor$ and n vertices and let $T_{\lfloor \theta n \rfloor, n}$ be a uniformly picked element from it with root ρ_n. Now observe that the degree of the root vertex is continuous in the local topology since we pick the root with probability proportional to its degree. Now if $T_{\lfloor \theta n \rfloor, n}$ converges to $\boldsymbol{T}_\kappa$ then

$$\begin{aligned} \mathbb{E}((\deg(\rho_n))^{-1}) &= \frac{1}{6(n+2g-2)|\mathcal{T}_{\lfloor \theta n \rfloor, n}|} \sum_{t \in \mathcal{T}_{\lfloor \theta n \rfloor, n}} \sum_{x \in V(t)} \deg(x) \frac{1}{\deg(x)} \\ &= \frac{n}{6(n+2g-2)} \xrightarrow[n \to \infty]{} \frac{1}{6(1+2\theta)} = \mathbb{E}(\deg(\rho_\kappa)^{-1}) \end{aligned} \quad (1.6.1)$$

where ρ_k is the root of $\boldsymbol{T}_\kappa$.

Conjecture 1.6.1. *(Benjamini and Curien [39]) For any $\theta \geq 0$, let $\kappa \in (0, 2/27]$ be such that $\mathbb{E}(deg(\rho_\kappa)^{-1}) = (6(1+2\theta))^{-1}$ where ρ_κ is the root of $\boldsymbol{T}_\kappa$. Then*

$$T_{\lfloor \theta n \rfloor, n} \xrightarrow[n \to \infty]{(d)} \boldsymbol{T}_\kappa$$

Of course, similar conjectures can be coined for quadrangulations or other reasonable classes of planar maps. To prove that the local limit is

$$\begin{tikzcd} Q_n \arrow[r] \arrow[d, leftrightarrow, "\Phi_n"] & \mathrm{UIPQ} \arrow[d, leftrightarrow, "\Phi_\infty"] \\ \mathrm{LT}_n \arrow[r] & \mathrm{LT}_\infty \end{tikzcd}$$

Figure 1.8: Illustration of the construction of UIPQ in [43].

topologically a plane, we have to prove with high probability we do not see non-contractible cycles near the root.

One standard way to attack such a problem is via the bijective techniques due to Schaeffer *et al.* For example, the UIPQ can be constructed using the bijective techniques as illustrated in [43]. Roughly, the idea is as follows: we know that finite quadrangulations Q_n with n faces are in bijection Φ_n with a uniform labelled tree LT_n with n edges. Then one can take the weak limit over n for the trees to easily obtain a uniform infinite labelled tree LT_∞. The local limit of a uniformly picked tree with n edges is the critical Galton-Watson tree conditioned to survive (see Figure 1.2). The label of a vertex v of the uniform infinite labelled tree defined by assigning i.i.d. uniform $\{-1, 0, 1\}$ edge weights to every edge and then summing these weights along the (unique) path from the root to v with the root vertex assigned label 0. Now one can obtain a quadrangulation from LT_∞ by using a similar recipe Φ_∞ used for the SHIQ: join each corner c with label l to the immediately next corner in the clockwise contour of the tree with label $l-1$. Thus we have the commutative diagram in Figure 1.8. Note that the arrow connecting UIPQ and LT_∞ in Figure 1.8 is correct provided the recipe Φ_∞ is continuous in the local topology which can be seen easily and is proved in [43].

One can hope to undertake a similar strategy for the high genus case as well. There exists a bijection between quadrangulations of genus g and labelled unicellular maps due to Chapuy, Marcus and Schaeffer (see [35]), but obtaining the local limit of labelled unicellular maps is not so straightforward as the genus 0 case. As a first step to solve this problem using this approach, we find the local limits of unicellular maps in high genus in Chapter 5. But it is no longer clear whether the labels are obtained just by adding i.i.d. edge weights as was the situation in the genus 0 case since the presence of cycles in unicellular maps makes understanding the distribution of the labels non-trivial. Thus these results alone do not suffice and we need new ideas to pursue this apporach.

Chapter 2

Classification of half planar maps

[1] The primary goal of this work is to classify all probability measures on half-planar maps which are domain Markov, and which additionally satisfy the simpler condition of translation invariance. As we shall see, these measures form a natural one (continuous) parameter family of measures.

We shall consider many different classes of planar maps in this chapter. We focus on triangulations, where all faces *except* possibly the external face are triangles, and on p-angulations where all faces are p-gons (except possibly the external face). We denote by $\mathcal{H}_p$ the class of all infinite, one-ended, half-planar p-angulations. However, it so transpires that $\mathcal{H}_p$ is not the best class of maps for studying the domain Markov property, for reasons that will be made clear later. At the moment, to state our results let us also define $\mathcal{H}'_p$ to be the subset of $\mathcal{H}_p$ of **simple maps**, where all faces are simple p-gons (meaning that each p-gon consists of p distinct vertices). Note that — as usual in the context of planar maps — multiple edges between vertices are allowed. However, multiple edges between two vertices cannot be part of any single simple face. We shall use $\mathcal{H}$ and $\mathcal{H}'$ to denote generic classes of half-planar maps, and simple half-planar maps, without specifying which. For example, this could also refer to the class of all half-planar maps, or maps with mixed face valencies.

2.1 Translation invariant and domain Markov measures

The **translation** operator $\theta : \mathcal{H} \to \mathcal{H}$ is the operator translating the root of a map to the right along the boundary. Formally, $\theta(M) = M'$ means that M and M' are the same map, except that the root edge of M' is the the edge immediately to the right of the root edge of M. Note that θ is a

[1] The results in this section are taken from the paper [12] and is joint work with Omer Angel.

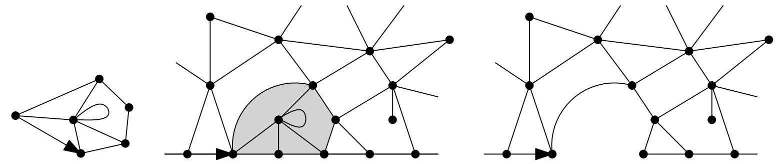

Figure 2.1: Left: A finite map Q. Centre: part of a map M containing Q with 2 edges along the boundary. Right: the resulting map $\tilde{M}$. The domain Markov property states that $\tilde{M}$ has the same law as M.

bijection. A measure μ on $\mathcal{H}$ is called **translation invariant** if $\mu \circ \theta = \mu$. Abusing language, we will also say that a random map M with law μ is translation invariant, even though typically moving the root of M yields a different (rooted) map.

The domain Markov property is more delicate, and may be informally described as follows: if we condition on the event that M contains some finite configuration Q and remove the sub-map Q from M, then the distribution of the remaining map is the same as that of the original map (see Figure 2.1).

We now make this precise. Let Q be a finite map in an m-gon for some finite m, and suppose the boundary of Q is simple (i.e. is a simple cycle in the graph of Q), and let $0 < k < m$ be some integer. Define the event $A_{Q,k} \subset \mathcal{H}$ that the map M contains a sub-map which is isomorphic to Q, and which contains the k boundary edges immediately to the right of the root edge of M, and *no other* boundary edges or vertices. Moreover, we require that the root edge of Q corresponds to the edge immediately to the right of the root of M. On this event, we can think of Q as being a subset of M, and define the map $\tilde{M} = M \setminus Q$, with the understanding that we keep vertices and edges in Q if they are part of a face not in Q (see Figure 2.1). Note that $\tilde{M}$ is again a half-planar infinite map.

Definition 2.1.1. *A probability measure μ on $\mathcal{H}$ is said to be* **domain Markov**, *if for any finite map Q and k as above, the law of $\tilde{M}$ constructed from a sample M of μ conditioned on the event $A_{Q,k}$ is equal to μ.*

Note that for translation invariant measures, the choice of the k edges to the right of the root edge is rather arbitrary: any k edges will result in $\tilde{M}$ with the same law. Similarly, we can re-root $\tilde{M}$ at any other deterministically chosen edge. Thus it is also possible to consider k edges that include the root edge, and mark a new edge as the root of $\tilde{M}$.

This definition is a relatively restrictive form of the domain Markov property. There are several other natural definitions, which we shall discuss below. While some of these definitions are superficially stronger, it turns out that several of them are equivalent to Definition 2.1.1.

2.2 Main results

Our main result is a complete classification and description of all probability measures on $\mathcal{H}'_p$ which are translation invariant and have the domain Markov property.

Theorem 2.2.1. *Fix $p \geq 3$. The set of domain Markov, translation invariant probability measures on $\mathcal{H}'_p$ forms a one parameter family $\{\mathbb{H}_\alpha^{(p)}\}$ with $\alpha \in \mathcal{I}_p \subset [0,1)$. The parameter α is the probability of the event that the p-gon incident to any fixed boundary edge is also incident to $p-2$ internal vertices.*

Moreover, for $p = 3$, $\mathcal{I}_3 = [0,1)$, and for $p > 3$ we have $(\alpha_0(p), 1) \subset \mathcal{I}_p$ for some $\alpha_0(p) < 1$.

We believe that $\mathcal{I}_p = [0,1)$ for all p although we have been able to prove this fact only for $p = 3$. We emphasize here that our approach would work for any p provided we have certain enumeration results. See Section 2.10 for more on this.

We shall normally omit the superscript (p), as p is thought of as any fixed integer. The measures $\mathbb{H}_\alpha^{(p)}$ are all mixing with respect to the translation θ and in particular are ergodic. This actually follows from a much more general proposition which is well known among experts for the standard half planar random maps, but we could not locate a reference. We include it here for future reference.

Proposition 2.2.2. *Let μ be domain Markov and translation invariant on $\mathcal{H}$. Then the translation operator is mixing on $(\mathcal{H}, \mu)$, and in particular is ergodic.*

Proof. Let Q, Q' and $A_{Q,k}, A_{Q',k'}$ be as in Definition 2.1.1. Since events of the form $A_{Q,k}$ are simple events in the local topology (recall discussion in Section 1.2), it suffices to prove that

$$\mu(A_{Q,k} \cap \theta^n(A_{Q',k'})) \to \mu(A_{Q,k})\mu(A_{Q',k'})$$

as $n \to \infty$ where θ^n is the n-fold composition of the operator θ. However, since on $A_{Q,k}$ the remaining map $\widetilde{M} = M \setminus Q$ has the same law as M, and since $\theta^n(A_{Q',k'})$ is just $\theta^{n'}(A_{Q',k'})$ in $\widetilde{M}$, for some n', we find from the domain Markov property that for large enough integer n, the equality $\mu(A_{Q,k} \cap \theta^n(A_{Q',k'})) = \mu(A_{Q,k})\mu(A_{Q',k'})$ holds. □

An application of Proposition 2.2.2 shows that the measures in the set $\{\mathbb{H}_\alpha^{(p)} : \alpha \in \mathcal{I}\}$ are all singular with respect to each other. This is because the density of the edges on the boundary for which the p-gon containing it is incident to $p-2$ internal vertices is precisely α by translation invariance. Note that the domain Markov property is not preserved by convex combinations of measures, so the measures $\mathbb{H}_\alpha$ are not merely the extremal points in the set of domain Markov measures.

Note also that the case $\alpha = 1$ is excluded. It is possible to take a limit $\alpha \to 1$, and in a suitable topology we even get a deterministic map. However, this map is not locally finite and so this can only hold in a topology strictly weaker than the local topology on rooted graphs. Indeed, this map is the plane dual of a tree with one vertex of infinite degree (corresponding to the external face) and all other vertices of degree p. As this case is rather degenerate we shall not go into any further details.

In the case of triangulations we get a more explicit description of the measures $\mathbb{H}_\alpha^{(3)}$, which we use in a future paper [82] to analyze their geometry. This can be done more easily for triangulations because of readily available and very explicit enumeration results. We believe deriving similar explicit descriptions for other p-angulations, at least for even p is possible with a more careful treatment of the associated generating functions, but leave this for future work. This deserves some comment, since in most works on planar maps the case of quadrangulations $q = 4$ yields the most elegant enumerative results. The reason the present work differs is the aforementioned necessity of working with simple maps. In the case of triangulations this precludes having any self loops, but any triangle with no self loop is simple, so there is no other requirement. For any larger p (including 4), the simplicity does impose further conditions. For example, a quadrangulation may contain a face consisting of two double edges.

We remark also that forbidding multiple edges in maps does not lead to any interesting domain Markov measures. The reason is that in a finite map Q it is possible that there exists an edge between any two boundary vertices. Thus on the event $A_{Q,k}$, it is impossible that $\tilde{M}$ contains any edge between boundary edges. This reduces one to the degenerate case of $\alpha = 1$, which is not a locally finite graph and hence excluded.

Our second main result is concerned with limits of uniform measures on finite maps. Let $\mu_{m,n}$ be the uniform measure on all simple triangulations of an m-gon containing n internal (non-boundary) vertices (or equivalently, $2n+m-2$ faces, excluding the external face). Recall we assume that the root edge is one of the boundary edges. The limits as $n \to \infty$ of $\mu_{m,n}$ w.r.t. the local topology on rooted graphs (formally defined in Section 1.2) have been studied in [13], and lead to the well-known UIPT. Similar limits exist for other classes of planar maps, see e.g. [65] for the case of quadrangulations. It is possible to take a second limit as $m \to \infty$, and the result is the half-plane UIPT measure (see also [44] for the case of quadrangulations). A second motivation for the present work is to identify other possible accumulation points of $\mu_{m,n}$. These measures would be the limits as $m, n \to \infty$ jointly with a suitable relation between them.

Theorem 2.2.3. *Consider sequences of non-negative integers m_l and n_l such that $m_l, n_l \to \infty$, and $m_l/n_l \to a$ for some $a \in [0, \infty]$. Then μ_{m_l,n_l} converges weakly to $\mathbb{H}_\alpha^{(3)}$ where $\alpha = \frac{2}{2a+3}$.*

The main thing to note is that the limiting measure does not depend on the sequences $\{m_l, n_l\}$, except through the limit of m_l/n_l. A special case is the measure $\mathbb{H}_{2/3}$ which correspond to the half-planar UIPT measure. Note that in this case, $a = 0$, that is the number of internal vertices grows faster than the boundary. Note that the only requirement to get this limit is $m_l = o(n_l)$. This extends the definition of the half-planar UIPT, where we first took the limit as $n_l \to \infty$ and only then let $m_l \to \infty$.

The other extreme case $\alpha = 0$ (or $a = \infty$) is also of special interest. To look into this case it is useful to consider the dual map. Recall that the **dual map** M^* of a planar map M is the map with a vertex corresponding to each face of M and an edge joining two colours faces (that is faces which share at least an edge), or more precisely a dual edge crossing every edge of M. Note that for a half-planar map M, there will be a vertex of infinite degree corresponding to the face of infinite degree. All other vertices shall have a finite degree (p in the case of p-angulations). To fit into the setting of locally finite planar maps, we can simply delete this one vertex, though a nicer modification is to break it up instead into infinitely many vertices of degree 1, so that the degrees of all other vertices are not changed. For half planar triangulations this gives a locally finite map which is 3-regular except for an infinite set of degree 1 vertices, each of which corresponds to a boundary edge. We can similarly define the duals of triangulations of an $m-$gon, where each vertex is of degree 3 except for m degree 1 vertices.

For a triangulation of an $m-$gon with no internal vertices ($n = 0$), the dual is a 3 regular tree with m leaves. Let T be the critical Galton-Watson tree where a vertex has 0 or 2 offspring with probability 1/2 each. We add a leaf to the root vertex, so that all internal vertices of T have degree 3. Then the law of M^* under $\mu_{m,0}$ is exactly T conditioned to have m leaves. Hence, this measure has a weak limit known as the *critical Galton-Watson tree conditioned to survive* (see Section 1.2). This is the law of the dual map M^* under $\mathbb{H}_0$. Observe that in $\mathbb{H}_0$, $\alpha = 0$, hence the probability that the triangle incident to any boundary edge has the third vertex also on the boundary is 1. As before, note that the only condition on m_l, n_l in Theorem 2.2.3 to get this limiting measure is that $n_l = o(m_l)$. For $p > 3$ the measure $\mathbb{H}_0^{(p)}$ has a similar description using trees with $p-1$ or 0 offspring.

Note that Theorem 2.2.3 gives finite approximations of $\mathbb{H}_\alpha$ for $\alpha \in [0, 2/3]$, so it is natural to ask for finite approximations to $\mathbb{H}_\alpha$ for $\alpha \in (2/3, 1)$? In this regime, the maps behave differently than those in the regime $\alpha < 2/3$ or $\alpha = 2/3$. Maps with law $\mathbb{H}_\alpha^{(3)}$ are hyperbolic in nature, and for example have exponential growth (we elaborate on the difference in Section 2.8 and investigate this further in [82]). Following the discussion in Section 1.6, we make the following conjecture similar in lines to Conjecture 1.6.1:

Conjecture 2.2.4. *Let T_{n_l,g_l,p_l} be a uniformly picked triangulation of genus g_l, n_l vertices and p_l boundary vertices rooted at the boundary. Also assume $g_l/n_l \to \theta$ for some $\theta \geq 0$ and $p_l = o(n_l)$. Then T_{n_l,g_l,p_l} converges weakly to $\mathbb{H}_\alpha$ where $\alpha \geq 2/3$ and is only a function of θ. Further if $\theta = 0$ then $\alpha = 2/3$.*

As indicated in Section 2.13, a similar phase transition is expected for p-angulations as well. Thus, we expect a similar conjecture about finite approximation to hold for any p, and not only triangulations.

Organization We prove the classification theorem for triangulations in Sections 2.6 and 2.7 and for p-angulations in Section 2.10 and also discuss the variation of maps with non-simple faces. In Section 2.12 we examine limits of uniform measures on finite maps, and prove Theorem 2.2.3.

2.3 Other approaches to the domain Markov property

In this section we discuss alternative possible definitions of the domain Markov property, and their relation to Definition 2.1.1. The common theme

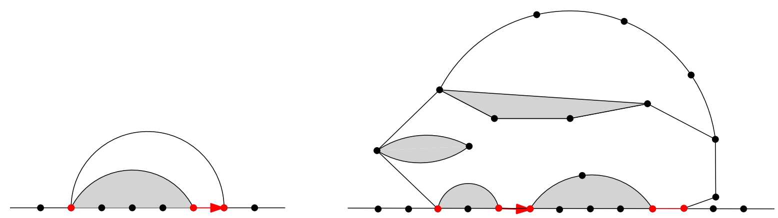

Figure 2.2: Possibilities when removing a sub-map Q connected to the boundary. The red part is C_M which is identified with C_Q. Left: Q consists of a single triangle. Right: Q consists of two faces in a general map. The shaded areas are the holes — finite components of the complement of Q.

is that a map M is conditioned to contain a certain finite sub-map Q, connected to the boundary at specified locations. We then remove Q to get a new map $\tilde{M}$. The difficulty arises because it is possible in general for $\tilde{M}$ to contain several connected components. See Figure 2.2 for some ways in which this could happen, even when the map Q consists of a single face.

To make this precise, we first introduce some topological notions. A sub-map of a planar map M is a subset of the faces of M along with the edges and vertices contained in them. We shall consider a map as a subset of the sphere on which it is embedded.

Definition 2.3.1. *A sub-map of a planar map is said to be* connected *if it is connected as a subset of the sphere. A connected sub-map E of a half planar map M is said to be simply connected if its union with the external face of M is a simply connected set in the sphere.*

Let Q denote a finite planar map, and let some (but at least one) of its faces be marked as external, and the rest as internal. We assume that the internal faces of Q are a connected set in the dual graph Q^*. One of the external faces of Q is singled out, and a non-empty subset C_Q containing at least one edge of the boundary of that external face is marked (in place of the k edges we had before). Note that C_Q need not be a single segment now. Fix also along the boundary of M a set C_M of the same size as C_Q, consisting of segments of the same length as those of C_Q and in the same order. We consider the event

$$A_Q = \{Q \subset M, \partial M \cap \partial Q = C_M\},$$

that Q is a sub-map of M, with C_Q corresponding to C_M. Figure 2.2 shows an example of this where Q has a single face.

On the event A_Q, the complement $M \setminus Q$ consists of one component with infinite boundary in the special external face of Q, and a number of components with finite boundary, one in each additional external face of Q. Let us refer to the components with finite boundary sizes as *holes*. Note that because M is assumed to be one-ended, the component with infinite boundary size, which is denoted by $\tilde{M}$ is the only infinite component of $M \setminus Q$. All versions of the domain Markov property for a measure μ state that

conditioned on A_Q, the infinite component of $M \setminus Q$ has law μ.

However, there are several possible assumptions about the distribution of the components of $M \setminus Q$ in the holes. We list some of these below.

1. No additional information is given about the distribution of the finite components.
2. The finite components are independent of the distribution of the infinite component.
3. The finite components are independent of the distribution of the infinite component and of each other.
4. The law of the finite components depends only on the sizes of their respective boundaries (i.e. two maps Q with holes of the same size give rise to the same joint distribution for the finite components).

It may seem at first that these are all stronger than Definition 2.1.1, since our definition of the domain Markov property only applies if Q is simply connected, in which case there are no finite components to $M \setminus Q$. This turns out to be misleading. Consider any Q as above, and condition on the finite components of $M \setminus Q$. Together with Q these form some simply connected map $\bar{Q}$ to which we may apply Definition 2.1.1. Thus for any set of finite maps that fill the holes of Q, $\tilde{M}$ has law μ. Since the conditional distribution of $\tilde{M}$ does not depend on our choice for the finite components, the finite components are independent of $\tilde{M}$. Thus options 1 and 2 are both equivalent to Definition 2.1.1, and the simple-connectivity condition for Q may be dropped.

In the case of p-angulations with simple faces, we have a complete classification of domain Markov measures. Along the proof, it will become clear that those in fact also satisfy the stronger forms 3 and 4 of the domain Markov property. This shows that for simple faced maps, every definition of the domain Markov property gives the same set of measures. If we allow non-simple faces, however, then different choices might yield smaller classes.

For example, if a non-simple face surrounds two finite components of the map, then under the domain Markov property as defined above, the parts of the map inside these components need not be independent of each other.

2.4 Peeling

Let us briefly describe the concept of peeling which has its roots in the physics literature [6, 92], and was used in the present form in [7]. It is a useful tool for analyzing planar maps, see e.g. applications to percolation and random walks on planar maps in [8, 10, 20]. While there is a version of this in full planar maps, it takes its most elegant form in the half plane case.

Consider a probability measure μ supported on a subset of $\mathcal{H}$ and consider a sample M from this measure. The peeling process constructs a growing sequence of finite simply-connected sub-maps (P_i) in M with complements $M_i = M \setminus P_i$ as follows. (The complement of a sub-map P contains every face not in P and every edge and vertex incident to them.) Initially $P_0 = \emptyset$ and $M_0 = M$. Pick an edge a_i in the boundary of M_i. Next, remove from M_i the face incident on a_i, as well as all finite components of the complement. This leaves a single infinite component $M_{i+1} = \tilde{M}_i$, and we set $P_{i+1} = M \setminus M_{i+1}$.

If μ is domain Markov and the choice of a_i depends only on P_i and an independent source of randomness, but not on M_i, then the domain Markov property implies by induction that M_n has law μ for every n, and moreover, M_n is independent of P_n. We will see that this leads to yet another interesting viewpoint on the domain Markov property.

In general, it need not be the case that $\bigcup P_i = M$ (for example, if the distance from the peeling edge a_i to the root grows very quickly). However, there are choices of edges a_i for which we do have $\bigcup P_i = M$ a.s. One way of achieving this is to pick a_i to be the edge of ∂M_i nearest to the root of M in the sub-map M_i, taking e.g. the left-most in case of ties. Note that this choice of a_i only depends on P_i and this strategy will exhaust any locally finite map M.

Let $Q_i = M_{i-1} \setminus M_i = P_i \setminus P_{i-1}$ for $i \geq 1$. This is the finite, simply connected map that is removed from M at step i. We also mark Q_i with information on its intersection with the boundary of M_{i-1} and the peeling edge a_{i-1}. This allows us to reconstruct P_i by gluing $Q_1, \dots, Q_i$. In this way, the peeling procedure encodes an infinite half planar map by an infinite sequence (Q_i) of marked finite maps. If the set of possible finite maps is denoted by $\mathcal{S}$, then we have a bijection $\Phi : \mathcal{H} \to \mathcal{S}^{\mathbb{N}}$. It is straightforward to see that this bijection is even a homeomorphism, where $\mathcal{H}$ is endowed with the local topology on rooted graphs (see Section 1.2), and $\mathcal{S}^{\mathbb{N}}$ with the product topology (based on the trivial topology on $\mathcal{S}$).

Now, if μ is a domain Markov measure on $\mathcal{H}$, then the pull-back measure $\mu^* = \mu \circ \Phi^{-1}$ on $\mathcal{S}^{\mathbb{N}}$ is an i.i.d. product measure, since the maps M_i all have the same law, and each is independent of all the Q_js for $j < i$. However, translation invariance of the original measure does not have a simple description in this encoding.

2.5 Preliminaries

2.5.1 Boltzmann distribution

We will also sometimes be interested in triangulations of discs where the number of internal vertices is not fixed, but is also random. The following measure is of particular interest:

Definition 2.5.1. *The* Boltzmann *distribution on rooted triangulations of an m-gon with weight $q \leq \frac{2}{27}$, is the probability measure on the set of finite triangulations with a finite simple boundary that assigns weight $q^n/Z_m(q)$ to each rooted triangulation of the m-gon having n internal vertices, where*

$$Z_m(q) = \sum_n \phi_{n,m} q^n.$$

where $\phi_{n,m}$ is given by eq. (1.3.1). From the asymptotics of ϕ as $n \to \infty$ we see that $Z_m(q)$ converges for any $q \leq \frac{2}{27}$ and for no larger q. The precise value of the partition function will be useful, and we record it here:

Proposition 2.5.2. *If $q = \theta(1-2\theta)^2$ with $\theta \in [0, 1/6]$, then*

$$Z_{m+2}(q) = \big((1-6\theta)(m+1)+1\big)\frac{(2m)!}{m!(m+2)!}\big(1-2\theta\big)^{-(2m+2)}.$$

In particular, at the critical point $q = 2/27$ we have $\theta = 1/6$ and Z takes the values

$$Z_{m+2} = Z_{m+2}\left(\frac{2}{27}\right) = \frac{(2m)!}{m!(m+2)!}\left(\frac{9}{4}\right)^{m+1}.$$

The proof can be found as intermediate steps in the derivation of $\phi_{n,m}$ in [56]. The above form may be deduced after a suitable reparametrization of the form given there.

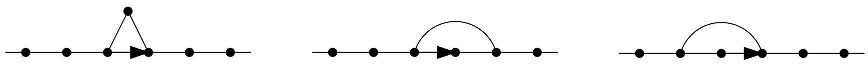

Figure 2.3: Basic building blocks for triangulations. Left: the event A_α. Centre and right: the two events of type A_β.

2.6 Half planar triangulations

For the sake of clarity, we begin by proving the special case $p = 3$ of Theorem 2.2.1 of half planar triangulations. In the case of triangulations, the number of simple maps and corresponding generating functions are known explicitly, making certain computations simpler. Somewhat surprisingly, the case of quadrangulations is more complex here, and the generating function is not explicitly known. Apart from the lack of explicit formulae, the case of general p presents a number of additional difficulties, and is treated in Section 2.10.

Theorem 2.6.1. *All translation invariant, domain Markov probability measures on $\mathcal{H}'_3$ form a one parameter family of measures $\mathbb{H}_\alpha$ for $\alpha \in [0, 1)$. Moreover, in $\mathbb{H}_\alpha$ the probability that the triangle containing any given boundary edge is incident to an internal vertex is α.*

In what follows, let μ be a measure supported on $\mathcal{H}'_3$, that is translation invariant and satisfies the domain Markov property. We shall first define a certain family of events and show that their measures can be calculated by repeatedly using the domain Markov property. Let $T \in \mathcal{H}'_3$ denote a triangulation with law μ. Let α be the μ-measure of the event that the triangle incident to a fixed boundary edge e is also incident to an interior vertex (call this event A_α, see Figure 2.3). The event depends on the boundary edge chosen, but by translation invariance its probability does not depend on the choice of e. As stated, our main goal is to show that α fully determines the measure μ.

For $i \geq 1$ define $p^{(r)}_{i,k}$ (resp. $p^{(l)}_{i,k}$) to be the μ-measure of the event that the triangle incident to a fixed boundary edge e of T is also incident to a vertex on the boundary to the right (resp. left) at a distance i along the boundary from the edge e and that this triangle separates k vertices of T that are not on the boundary from infinity. Note that because of translation invariance, these probabilities only depends on i and k and hence we need not specify e in the notation. It is not immediately clear, but we shall see later that $p^{(l)}_{i,k} = p^{(r)}_{i,k}$ (see Corollary 2.6.4 below). In light of this, we shall later drop the superscript.

The case $i = 1$, $k = 0$ is of special importance. Since there is no triangulation of a 2-gon with no internal vertex, if the triangle containing e is incident to a boundary vertex adjacent to e, then it must contain also the boundary edge next to e. (See also the discussion in Section 2.5.1.) We call such an event A_β, shown in Figure 2.3. By translation invariance, we now see that $p_{1,0}^{(r)} = p_{1,0}^{(l)}$. We shall denote $\beta = p_{1,0}^{(r)} = p_{1,0}^{(l)}$.

In what follows, fix α and β. Of course, not every choice of α and β is associated with a domain Markov measure, and so there are some constraints on their values. We compute below these constraints, and derive β as an explicit function of α for any $\alpha \in [0, 1)$.

Let Q be a finite simply connected triangulation with a simple boundary, and let $B \subsetneq \partial Q$ be a marked, nonempty, connected segment in the boundary ∂Q. Fix a segment in ∂T of the same length as B, and let A_Q be the event that Q is isomorphic to a sub-triangulation of $T \in \mathcal{H}_3'$ with B being mapped to the fixed segment in ∂T, and no other vertex of Q being mapped to ∂T. Let $F(Q)$ be the set of faces of Q, $V(Q)$ the set of vertices of Q (including those in ∂Q), and $V(B)$ the set of vertices in B, including the endpoints.

Lemma 2.6.2. *Let μ be a translation invariant domain Markov measure on $\mathcal{H}_3'$. Then for an event A_Q as above we have*

$$\mu(A_Q) = \alpha^{|V(Q)|-|V(B)|}\beta^{|F(Q)|-|V(Q)|+|V(B)|} \tag{2.6.1}$$

Furthermore, if a measure μ satisfies (2.6.1) *for any such Q, then μ is translation invariant and domain Markov.*

Remark 2.6.3. *$|V(Q)| - |V(B)|$ is the number of vertices of Q not on the boundary of T. This shows that the probability of the event A_Q depends only on the number of vertices not on the boundary of T and the number of faces of Q, but nothing else.*

The proof of Lemma 2.6.2 is based on the idea that the events A_α and A_β form basic "building blocks" for triangulations. More precisely, there exists some ordering of the faces of Q such that if we reveal triangles of Q in that order and use the domain Markov property, we only encounter events of type A_α, A_β. Moreover, in any such ordering the number of times we encounter the events A_α and A_β are the same as for any other ordering. Also observe that, for every event of type A_α encountered, we add a new vertex while for every event of type A_β encountered, we add a new face. Thus the exponent of A_α counts the number of "new" vertices added while the exponent of A_β counts the number of "remaining" faces.

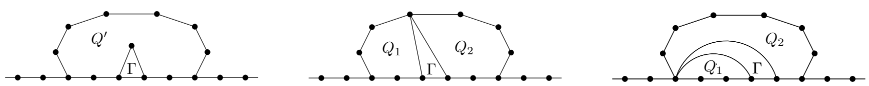

Figure 2.4: Cases in the inductive step in the proof of Lemma 2.6.2.

Proof of Lemma 2.6.2. We prove (2.6.1) by induction on $|F(Q)|$: the number of faces of Q. If $|F(Q)| = 1$, then Q is a single triangle, and B contains either one edge or two adjacent edges. If it has one edge, then the triangle incident to it must have the third vertex not on the boundary of T. By definition, in this case $\mu(A_Q) = \alpha$ and we are done since $|V(Q)| = 3$ and $|V(B)| = 2$. Similarly, if B contains two edges, then $|V(B)| = 3$ and A_Q is just the event A_β, with probability β, consistent with (2.6.1).

Next, call the vertices of Q that are not in B *new* vertices. Suppose $|F(Q)| = n$, and that we have proved the lemma for all Q' with $|F(Q)| < n$. Pick an edge e_0 from B (there exists one by hypothesis), and let Γ be the face of Q incident to this edge. There are three options, depending on where the third vertex of Γ lies in Q (see Figure 2.4):

- the third vertex of Γ is internal in Q,
- the third vertex of Γ is in $\partial Q \setminus B$,
- the third vertex of Γ is in B.

We treat each of these cases separately.

In the first case, we have that $Q' = Q - \Gamma$ is also a simply connected triangulation, if we let B' include the remaining edges from B as well as the two new edges from Γ, we can apply the induction hypothesis to Q'. By the domain Markov property, we have that

$$\mu(A_Q) = \mu(A_\Gamma)\mu(A_Q|A_\Gamma) = \alpha\mu(A_{Q'}).$$

This implies the claimed identity for Q, since Q' has one less face and one less new vertex than Q.

In the case where the third vertex of Γ is in $\partial Q \setminus B$, we have a decomposition $Q = \Gamma \cup Q_1 \cup Q_2$, where Q_1 and Q_2 are the two connected components of $Q \setminus \Gamma$ (see Figure 2.4). We define B_i, to contain the edges of B in Q_i, and one edge of Γ that is in Q_i. We have that $|F(Q)| = |F(Q_1)| + |F(Q_2)| + 1$, and that the new vertices in Q_1 and Q_2 except for the third vertex of $F(Q)$ together are the new vertices of Q. By the domain Markov property, conditioned on A_Γ, the inclusion of Q_1 and of Q_2 in T are independent events with corresponding probabilities $\mu(A_{Q_i})$. Thus

$$\begin{aligned}\mu(A_Q) &= \alpha\mu(A_{Q_1}|A_\Gamma)\mu(A_{Q_2}|A_\Gamma) \\ &= \alpha\mu(A_{Q_1})\mu(A_{Q_2}) = \alpha^{|V(Q)|-|V(B)|}\beta^{|F(Q)|-|V(Q)|+|V(B)|},\end{aligned}$$

as claimed.

Finally, consider the case that the third vertex of Γ is in B. As in the previous case, we have a decomposition $Q = \Gamma \cup Q_1 \cup Q_2$, where Q_1 is the triangulation separated from infinity by Γ, and Q_2 is the part adjacent to the rest of T (see Figure 2.4.) We let B_1 consist of the edges of B in Q_1 and let B_2 be the edges of B in Q_2 with the additional edge of Γ. We then have

$$\mu(A_Q) = \mu(A_{Q_1})\mu(A_{Q_1\cup\Gamma}|A_{Q_1})\mu(A_Q|A_{Q_1\cup\Gamma}).$$

By the induction hypothesis, the first term is

$$\alpha^{|V(Q_1)|-|V(B_1)|}\beta^{|F(Q_1)|-|V(Q_1)|+|V(B_1)|}.$$

By the domain Markov property, the second term is just β. Similarly, the third term is $\alpha^{|V(Q_2)|-|V(B_2)|}\beta^{|F(Q_2)|-|V(Q_2)|+|V(B_2)|}$. As before we have that $|F(Q)| = |F(Q_1)|+|F(Q_2)|+1$, and this time $|V(Q)|-|V(B)| = (|V(Q_1)|-|V(B_1)|) + (|V(Q_2)| - |V(B_2)|)$, since the new vertices of Q are the new vertices of Q_1 together with the new vertices of Q_2. The claim again follows.

Note that in the last case it is possible that Q_1 is empty, in which case Γ contains two edges from ∂Q. All formulae above hold in this case with no change.

For the converse, note first that since the events A_Q are a basis for the local topology on rooted graphs, they uniquely determine the measure μ. Moreover, the measure of the events of the form A_Q do not depend on the location of the root and so μ is translation invariant. Now observe from Remark 2.6.3 that the measure of any event of the form A_Q only depends on the number of new vertices and the number of faces in Q. Now suppose we remove any simple connected sub-map Q_1 from Q. Then the union of new vertices in Q_1 and $Q \setminus Q_1$ gives the new vertices of Q. Also clearly, the union of the faces of Q_1 and $Q \setminus Q_1$ gives the faces of Q. Hence it follows that $\mu(A_Q|A_{Q_1}) = \mu(A_{Q\setminus Q_1})$, and thus μ is domain Markov. $\square$

Corollary 2.6.4. *For any i, k we have*

$$p_{i,k} := p_{i,k}^{(r)} = p_{i,k}^{(l)} = \phi_{k,i+1}\alpha^k\beta^{i+k} \tag{2.6.2}$$

Proof. This is immediate because the event with probability $p_{i,k}$ is a union of $\phi_{k,i+1}$ disjoint events of the form A_Q, corresponding to all possible triangulations of an $i+1$-gon with k internal vertices. A triangulation contributing to $\phi_{k,i+1}$ has k internal vertices by the Euler characteristic formula, $2k+i-1$ faces. The triangle that separates it from the rest of the map is responsible for the extra factor of β. □

Since the probability of any finite event in $\mathcal{H}'_3$ can be computed in terms of the peeling probabilities $p_{i,k}$'s, we see that for any given α and β we have at most a unique measure μ supported on $\mathcal{H}'_3$ which is translation invariant and satisfies the domain Markov property. The next step is to reduce the number of parameters to one, thereby proving the first part of Theorem 2.6.1. This is done in the following lemma.

Lemma 2.6.5. *Let μ be a domain Markov, translation invariant measure on $\mathcal{H}'_3$, and let α,β be as above. Then*

$$\beta = \begin{cases} \frac{1}{16}(2-\alpha)^2 & \alpha \leq 2/3, \\ \frac{1}{2}\alpha(1-\alpha) & \alpha \geq 2/3. \end{cases}$$

Proof. The key is that since the face incident to the root edge is either of type α, or of the type with probability $p_{i,k}$ for some i,k, (with $i=1, k=0$ corresponding to type β) we have the identity

$$\alpha + \sum_{i\geq 1}\sum_{k\geq 0}\left(p_{i,k}^{(r)} + p_{i,k}^{(l)}\right) = 1.$$

In light of Corollary 2.6.4 we may write this as

$$1 = \alpha + 2\sum_i \beta^i \sum_k \phi_{k,i+1}(\alpha\beta)^k = \alpha + 2\sum_i \beta^i Z_{i+1}(\alpha\beta).$$

From Proposition 2.5.2 we see that the sum above converges if and only if $\alpha\beta \leq \frac{2}{27}$. In that case, there is a $\theta \in [0,1/6]$ with $\alpha\beta = \theta(1-2\theta)^2$. Using the generating function for ϕ (see e.g. [56]) and simplifying gives the explicit identity

$$(2\theta + \alpha - 1)\sqrt{1 - \frac{4\theta}{\alpha}} = 0. \tag{2.6.3}$$

Thus $\theta \in \{\frac{1-\alpha}{2}, \frac{\alpha}{4}\}$. Of these, only one solution satisfies $\theta \in [0,1/6]$ for any value of α. If $\alpha \leq 2/3$, then we must have $\theta = \alpha/4$ which yields

$$\beta = \frac{1}{4}\left(1 - \frac{\alpha}{2}\right)^2 = \frac{1}{16}(2-\alpha)^2$$

If $\alpha \geq 2/3$ one can see from (2.6.3) that the solution satisfying $\theta \in [0, 1/6]$ is $\theta = (1-\alpha)/2$ which in turn gives

$$\beta = \frac{\alpha(1-\alpha)}{2}. \qquad \square$$

2.7 Existence

As we have determined β in terms of α, and since Lemma 2.6.2 gives all other probabilities $p_{i,k}$ in terms of α and β, we have at this point proved uniqueness of the translation invariant domain Markov measure with a given $\alpha < 1$. However we still need to prove that such a measure exists. We proceed now to give a construction for these measures, via a version of the peeling procedure (see Section 2.4). For $\alpha \leq 2/3$, we shall see with Theorem 2.2.3 that the measures $\mathbb{H}_\alpha$ can also be constructed as local limits of uniform measures on finite triangulations.

In light of Lemma 2.6.2, all we need is to construct a probability measure μ such that the measure of the events of the form A_Q (as defined in Lemma 2.6.2) is given by (2.6.1).

If we reveal a face incident to any fixed edge in a half planar triangulation along with all the finite components of its complement, then the revealed faces form some sub-map Q. The events A_Q for such Q are disjoint, and form a set we denote by $\mathcal{A}$. If we choose α and β according to Lemma 2.6.5, then the prescribed measure of the union of the events in $\mathcal{A}$ is 1.

Let α and let β be given by Lemma 2.6.5. We construct a distribution μ_r on the hull of the ball of radius r in the triangulation (which consists of all faces with a corner at distance less than r from the root, and with the holes added to make the hull).

Repeatedly pick an edge on the boundary which has at least an endpoint at a distance strictly less than r from the root edge in the map revealed so far. Note that as more faces are added to the map, distances may become smaller, but not larger. Reveal the face incident to the chosen edge and all the finite components of its complement. Given α and β we pick which event in $\mathcal{A}$ occurs by (2.6.1), independently for different steps. We continue the process as long as any vertex on the exposed boundary is at distance less than r from the root. Note that this is possible since the revealed triangulation is always simply connected with at least one vertex on the boundary, the complement must be the upper half plane.

Proposition 2.7.1. *The above described process almost surely ends after finitely many steps. The law of the resulting map does not depend on the* *order in which we choose the edges.*

Proof. We first show that the process terminates for some order of exploration. The following argument for termination is essentially taken from [7]. Assume that at each step we pick a boundary vertex at minimal distance (say, k) from the root (w.r.t. the revealed part of the map), and explore along an edge containing that vertex. At any step with probability $\beta > 0$ we add a triangle such that the vertex is no longer on the boundary. Any new revealed vertex must have distance at least $k+1$ from the root. Moreover, any vertex that before the exploration step had distance greater than k to the root, still has distance greater than k, since the shortest path to any vertex must first exit the part of the map revealed before the exploration step. Thus the number of vertices at distance k to the root cannot increase, and has probability $\beta > 0$ of decreasing at each step. Thus almost surely after a finite number of steps all vertices at distance k are removed from the boundary. Once we reach distance r, we are done.

The probability of getting any possible map T is a monomial in α and β, and is the same regardless of the order in which the exploration takes place (with one α for each non-boundary vertex of the map, and a β term for the difference between faces and vertices). It remains to show that the process terminates for any other order of exploration. For some order of exploration, let $\nu_i(T)$ be the probability that the process terminated after at most i steps and revealed T as the ball of radius r. For i large enough (larger than the number of faces in T) we have that $\nu_i(T) = \mu_r(T)$. Summing over T and taking the limit as $i \to \infty$, Fatou's lemma implies that $\lim_i \sum_T \nu_i(T) \geq \sum \mu_r(T)$. However, the last sum must equal 1, since for some order of exploration the process terminates a.s. $\square$

It is clear from Proposition 2.7.1 that μ_r is a well-defined probability measure. Since we can first create the hull of radius r and then go on to create the hull of radius $r+1$, (μ_r) forms a consistent sequence of measures. By Kolmogorov's extension Theorem, $(\mu_r)_{r\in\mathbb{N}}$ can be extended to a measure $\mathbb{H}_\alpha$ on $\mathcal{H}_3'$. Also, we have the following characterization of $\mathbb{H}_\alpha$ for any simple event of the form A_Q as defined in Lemma 2.6.2.

Lemma 2.7.2. *For any A_Q and B as defined in Lemma 2.6.2,*

$$\mathbb{H}_\alpha(A_Q) = \alpha^{|V(Q)|-|V(B)|}\beta^{|F(Q)|-|V(Q)|+|V(B)|} \tag{2.7.1}$$

We alert the reader that such a characterization is not obvious from the fact that the events of the form $\{\overline{B_r} = T\}$ have the $\mathbb{H}_\alpha$ measure exactly as asserted by Lemma 2.7.2 where $\overline{B_r}$ denotes the hull of the ball of radius r around the root vertex. Any finite event like A_Q can be written in terms of the measures of $\mathbb{H}_\alpha(\overline{B_r} = T)$ for different $T \in \mathcal{T}$ by appropriate summation. However it is not clear *a priori* that the result will be as given by (2.7.1).

Proof of Lemma 2.7.2. Since Q is finite, there exists a large enough r such that Q is a subset of $\overline{B_r}$. Now we claim that $\mu_r(A_Q)$ is given by the right hand side of (2.7.1). This is because crucially, μ_r is independent of the choice of the sequence of edges, and hence we can reveal the faces of Q first and then the rest of $\overline{B_r}$. However the measure of such an event is given by the right hand side of (2.7.1) by the same logic as Proposition 2.7.1. Now the lemma is proved because $\mathbb{H}_\alpha(A_Q) = \mu_r(A_Q)$ since $\mathbb{H}_\alpha$ is an extension of μ_r. □

We now have all the ingredients for the proof of Theorem 2.6.1.

Proof of Theorem 2.6.1. We have the measures $\mathbb{H}_\alpha$ constructed above which are translation invariant and domain Markov (this follows from the second part of Lemma 2.6.2). If μ is a translation invariant domain Markov measure, then by Lemmas 2.6.2, 2.6.5 and 2.7.2, μ agrees with $\mathbb{H}_\alpha$ on every event of the form A_Q, and thus $\mu = \mathbb{H}_\alpha$ for some α. □

2.8 The phase transition

In the case of triangulations, we call the measures $\mathbb{H}_\alpha$ *subcritical, critical* and *supercritical* when $\alpha < \frac{2}{3}$, $\alpha = \frac{2}{3}$, and $\alpha > \frac{2}{3}$ respectively. We summarize here for future reference the peeling probabilities $p_{i,k}$ and $p_i = 2\sum_{k\geq 0} p_{i,k}$ for every $\alpha \in [0,1)$. Recall that θ is defined by $\alpha\beta = \theta(1-2\theta)^2$ and $\theta \in [0, \frac{1}{6}]$.

Critical case: $\alpha = \frac{2}{3}$ This case is the well-known half plane UIPT (see [7]) Here $\beta = \frac{1}{9}$ and $\theta = \frac{1}{6}$. The two possible values of θ coincide at $\frac{1}{6}$ and hence $\beta = \frac{1}{9}$. Using Corollary 2.6.4 and Proposition 2.5.2, we recover the probabilities

$$\begin{aligned} p_{i,k} &= \phi_{k,i+1} \left(\frac{1}{9}\right)^i \left(\frac{2}{27}\right)^k \\ p_i &= \frac{2}{4^i} \frac{(2i-2)!}{(i-1)!(i+1)!} \end{aligned} \tag{2.8.1}$$

Note that in $\mathbb{H}_{2/3}$ we have the asymptotics $p_i \sim ci^{-5/2}$ for some $c > 0$.

Sub-critical case: $\alpha < \frac{2}{3}$ Here $\theta = \alpha/4$ and hence $\beta = \frac{(2-\alpha)^2}{16}$. Using Corollary 2.6.4 and Proposition 2.5.2, we get

$$\begin{aligned} p_{i,k} &= \phi_{k,i+1}\left(\frac{2-\alpha}{4}\right)^{2i}\left(\frac{\alpha}{4}\left(1-\frac{\alpha}{2}\right)^2\right)^k \\ p_i &= \frac{2}{4^i}\frac{(2i-2)!}{(i-1)!(i+1)!}\cdot\left(\left(1-\frac{3\alpha}{2}\right)i+1\right) \end{aligned} \tag{2.8.2}$$

As before, we get the asymptotics $p_i \sim ci^{-3/2}$ for some $c = c(\alpha) > 0$. Note that p_i is closely related to a linearly biased version of p_i for the critical case.

Super-critical case: $\alpha > 2/3$ Here $\theta = \frac{1-\alpha}{2}$ and hence $\beta = \frac{\alpha(1-\alpha)}{2}$. Using Corollary 2.6.4 and Proposition 2.5.2, we get

$$\begin{aligned} p_{i,k} &= \phi_{k,i+1}\alpha^{i+2k}\left(\frac{1-\alpha}{2}\right)^{i+k} \\ p_i &= \frac{2}{4^i}\frac{(2i-2)!}{(i-1)!(i+1)!}\cdot\left(\frac{2}{\alpha}-2\right)^i((3\alpha-2)i+1) \end{aligned} \tag{2.8.3}$$

Here, the asymptotics of p_i are quite different, and p_i has an exponential tail: $p_i \sim c\gamma^i i^{-3/2}$ for some c and $\gamma = \frac{2}{\alpha} - 2$. The differing asymptotics of the connection probabilities p_i indicate very different geometries for these three types of half plane maps. These are almost (though not quite) the probabilities of edges between boundary vertices at distance i.

2.9 Non-simple triangulations

So far in this chapter, we have only considered one type of map: triangulations with multiple edges allowed, but no self loops. Forbidding double edges combined with the domain Markov property, leads to a very constrained set of measures. The reason is that a step of type α followed by a step of type β can lead to a double edge. If μ is supported on measures with no multiple edges, this is only possible if $\alpha\beta = 0$. As seen from the discussion above, this gives the unique measure $\mathbb{H}_0$ which has no internal vertices at all. A similar phenomenon occurs for p-angulations for any $p \geq 3$, and we leave the details to the reader.

In contrast, the reason one might wish to forbid self-loops is less clear. We now show that on the one hand, allowing self-loops in a triangulation leads to a very large family of translation invariant measures with the domain Markov property. On the other hand, these measures are all in an essential way very close to one of the $\mathbb{H}_\alpha$ measures already encountered. The reason that uniqueness breaks as thoroughly as it does, is that here it is possible for removal of a single face to separate the map into two components, one of which is only connected to the infinite part of the boundary through the removed face. We remark that for triangulations with self loops, the stronger forms of the domain Markov property discussed in Section 2.3 are no longer equivalent to the weaker ones that we use.

Let us construct a large family of domain Markov measures as promised. Our translation invariant measures on triangulations with self-loops are made up of three ingredients. The first is the parameter $\alpha \in [0, 1)$ which corresponds to a measure $\mathbb{H}_\alpha$ as above. Next, we have a parameter $\gamma \in [0, 1)$ which represents the density of self loops. Taking $\gamma = 0$ will result in no self-loops and the measure will be simply $\mathbb{H}_\alpha$. Finally, we have an arbitrary measure ν supported on triangulations of the 1-gon (i.e. finite triangulations whose boundary is a self-loop, possibly with additional self-loops inside). From α, γ and ν we construct a measure denoted $\mathbb{H}_{\alpha,\gamma,\nu}$. More precisely, we describe a construction for a triangulation with law $\mathbb{H}_{\alpha,\gamma,\nu}$.

Given α, take a sample triangulation T from $\mathbb{H}_\alpha$. For each edge e of T, including the boundary edges, take an independent geometric variable G_e with $\mathbb{H}_{\alpha,q,\nu}(G_e = k) = (1 - q)q^{k-1}$. Next, replace the edge e by G_e parallel edges, thereby creating $G_e - 1$ faces which are all 2-gons. In each of the 2-gons formed, add a self-loop at one of the two vertices, chosen with equal probability and independently of the choices at all other 2-gons. This has the effect of splitting the 2-gon into a triangle and a 1-gon. Finally, fill each self-loop created in this way with an independent triangulation with law ν (see Figure 2.5).

Proposition 2.9.1. *The measures $\mathbb{H}_{\alpha,q,\nu}$ defined above are translation invariant and satisfy the domain Markov property. For $\alpha > 0$, these are all the measures on half planar triangulations with these properties.*

Recall that we use α to denote the probability of the event of type α that the triangle incident on any boundary edge also contains an internal vertex. The case of triangulations with $\alpha = 0$ is special for reasons that will be clearer after the proof, and is the topic of Proposition 2.9.2. In that case we shall require another parameter, and another measure ν'. This will be the only place where we shall demonstrate domain Markov measures that are not symmetric w.r.t. left-right reflection.

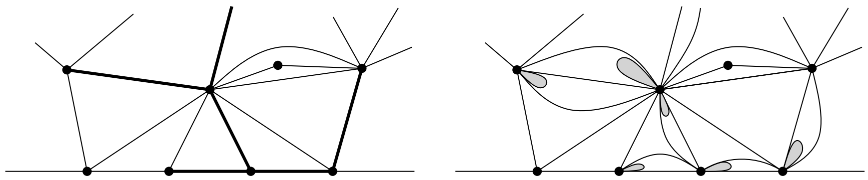

Figure 2.5: Non-uniqueness for triangulation with self-loops. Starting with a triangulation with simple faces (left), each edge is replaced by a geometric number of parallel edges with a self-loop at one of the two vertices between any pair (greater than 1 at the bold edges). Independent maps with arbitrary distribution are added inside the self-loops (shaded). Note that multiple edges may occur on the left (but not self-loops).

Coming back to the case $\alpha > 0$, note that since ν is arbitrary, the structure of domain Markov triangulations with self-loops is much less restricted than without the self-loops. For example, ν could have a very heavy tail for the size of the maps, or for the degree of the vertex in the self-loop, which will affect the degree distribution of vertices in the map. However, the measures $\mathbb{H}_{\alpha,q,\nu}$ are closely related to $\mathbb{H}_\alpha$, since the procedure described above for generating a sample of $\mathbb{H}_{\alpha,q,\nu}$ from a sample of $\mathbb{H}_\alpha$ is reversible. Indeed, if we take a sample from $\mathbb{H}_{\alpha,q,\nu}$ and remove each loop and the triangulation inside it, we are left with a map whose faces are triangles or 2-gons. If we then glue the edges of each 2-gon into a single edge, we are left with a simple triangulation. We refer to this operation as *taking the 2-connected core* of the triangulation, since the dual of the triangulation contains a unique infinite maximal 2-connected component, which is a subdivision of the dual of the triangulation resulting from this operation. Clearly the push-forward of the measures $\mathbb{H}_{\alpha,q,\nu}$ via this operation has law $\mathbb{H}_\alpha$. Thus $\mathbb{H}_\alpha$ does determine in some ways the large scale structure of $\mathbb{H}_{\alpha,q,\nu}$.

Proof of Proposition 2.9.1. Translation invariance is clear as $\mathbb{H}_\alpha$ is translation invariant, the variables G_e and triangulations in the self-loops do not depend upon the location of the root.

To see that $\mathbb{H}_{\alpha,q,\nu}$ is domain Markov, let T be a half planar triangulation with law $\mathbb{H}_{\alpha,q,\nu}$. Let core$(\cdot)$ denote the 2-connected core of a map, and observe that core(T) is a map with law $\mathbb{H}_\alpha$ from which T was constructed. Let Q be a finite simply connected triangulation (which may contain non-simple faces), and let A_Q be the event as defined in Lemma 2.6.2. To establish the domain Markov property for $\mathbb{H}_{\alpha,q,\nu}$, we need to show that conditionally on A_Q, $\tilde{T} = T \backslash Q$ (as defined in Section 2.1) has the same law as T. On the event A_Q, a corresponding event $A_{\mathrm{core}(Q)}$ that $\mathrm{core}(Q) \subset \mathrm{core}(T)$ also holds. Moreover, on these events, $\mathrm{core}(\tilde{T}) = \mathrm{core}(T) \setminus \mathrm{core}(Q)$ has law $\mathbb{H}_\alpha$, since $\mathbb{H}_\alpha$ is domain Markov. We therefore need to show that to get from $\mathrm{core}(T) \backslash \mathrm{core}(Q)$ to $\tilde{T}$ each edge is replaced by a $\mathrm{Geom}(q)$ number of parallel non-simple triangles with ν-distributed triangulations inside the self-loops. Any edge of $\mathrm{core}(T) \setminus \mathrm{core}(Q)$ is split in $\tilde{T}$ into an independent $\mathrm{Geom}(q)$ number of parallel edges. Indeed, for edges not in $\mathrm{core}(Q)$ this number is the same as in T, and for edges in the boundary of Q, the number is reduced by those non-simple triangles that are in Q, but is still $\mathrm{Geom}(q)$ due to the memory-less property of the geometric variables. The triangulations inside the self-loops are i.i.d. samples of ν, since they are just a subset of the ones in T which are i.i.d. and ν-distributed.

For the second part of the proposition, note first that if μ is domain Markov, then the push-forward of μ w.r.t. taking the core is also domain Markov, hence must be $\mathbb{H}_\alpha$ for some $\alpha \in [0, 1)$ by Theorem 2.6.1.

Fix an edge along the boundary, let q be the probability that the face containing it is not simple. By the domain Markov property, conditioned on having such a non-simple face and removing it leaves the map unchanged in law, and so this is repeated $\mathrm{Geom}(q)$ times before a simple face is found. Removing all of these faces also does not change the rest of the map, and so this number is independent of the multiplicity at any other edge of the map. Similarly, the triangulation inside the self-loop within each such non simple face is independent of all others, and we may denote its law by ν. Since any edge inside the map may be turned into a boundary edge by removing a suitable finite sub-map, the same holds for all edges.

To see that $\mu = \mathbb{H}_{\alpha,q,\nu}$, it remains to show that the self-loops are equally likely to appear at each end-point of the 2-gons and are all independent. The independence follows as for the triangulations inside the self-loops. To see that the two end-points are equally likely (and only to this end) we require $\alpha > 0$. The configuration shown in Figure 2.6 demonstrates this. After removing the face on the right, the self-loop is at the right end-point of a 2-gon on the boundary. Removing the triangle on the left leaves the self-loop on the left end-point, and so the two are equally likely. $\square$

As noted above, the case $\alpha = 0$ is special. In this case, no boundary edge has its third vertex internal to the triangulation. Note that this is not the same as saying that the triangulation has no internal vertices - they could all be inside self-loops, which are attached to the boundary vertices.

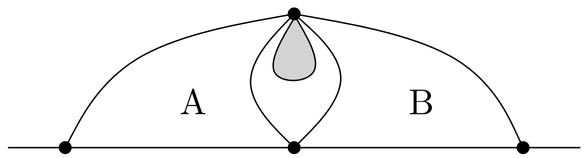

Figure 2.6: Exploring in different orders shows that self-loops are equally likely to be at each end-point of a 2-gon. Conditioning on face A and removing it leaves a non-simple face along the boundary with the self-loop at the left vertex. Removing instead face B leaves the self-loop on the right vertex.

The contraction operation described above still necessarily yields a sample T of $\mathbb{H}_0$. Similarly, each edge of T must correspond to an independent, geometric number of edges in the full map, and the triangulations inside the corresponding self-loops must be independent.

However, without steps of type α we cannot show that the the two choices for the location of the self-loop in 2-gons are equally likely. Indeed, since all 2-gons connect a pair of boundary vertices, it is possible to tell them apart. Adding the self-loop always on the left vertex will not be the same as adding it always on the right. This reasoning leads to a complete characterization also in the case $\alpha = 0$. In each 2-gon the self-loop is on the left vertex with some probability $\gamma \in [0, 1]$, and these must be independent of all other 2-gons. The triangulations inside the self-loops are all independent, but their laws may depend on whether the self-loop is on the left or right vertex in the 2-gon, so we need to specify two measures ν_L, ν_R on triangulations of the 1-gon. Thus we get the following:

Proposition 2.9.2. *A domain Markov, translation invariant triangulation with $\alpha = 0$ is determined by the intensity of multiple edges q, the probability $\gamma \in [0, 1]$ that the self-loop is attached to the left vertex in each 2-gon, and probability measures ν_L, ν_R on triangulations of the 1-gon.*

2.10 Simple and general p-angulations

Here we prove the general case of Theorem 2.2.1. The proof is similar to the proof of Theorem 2.6.1, with some additional complications: There are more than the two types of steps α and β, and the generating function for simple p-angulations is not explicitly known. There are implicit formulae relating it to the generating function for general maps with suitably chosen

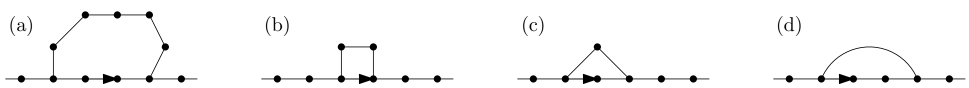

Figure 2.7: Building blocks for quadrangulations and general p-angulations. Shown: an event of type A_5 for $p = 9$ and the three building blocks for $p = 4$.

weights for various face sizes, which are fairly well understood in the case of even p. For quadrangulations, even more is known. In [83], the problem of enumerating 2-connected loopless near 4-regular planar maps (see [83] for exact definitions) is considered. This is easily equivalent to our problem of enumerating simple faced quadrangulations with a simple boundary. The generating function is computed there in a non-closed form. With careful analysis, this might lead to explicit expressions analogous to the ones we have for the triangulation case at least for the case of quadrangulations. We have not been able to obtain such expressions, and thus our description of the corresponding $\mathbb{H}_\alpha$'s still depends on an undetermined parameter $\beta = \beta(\alpha)$. Instead, uniqueness is proved by a softer argument based on monotonicity. The proof of existence used for triangulations goes through with no significant changes, but is now conditional on the existence of a solution to a certain equation.

Proof of Theorem 2.2.1. As before, let μ be a probability measure supported on the set $\mathcal{H}'_p$ of half planar simple p-angulations which is translation invariant and satisfies the domain Markov property. The building blocks for simple p-angulations, taking the place of A_α and A_β, will be the events where the face incident to the root edge consists of a single contiguous segment from the infinite boundary, together with a simple path in the interior of the map closing the cycle, with the path in any fixed position relative to the root (see Figure 2.7(a)). The number of internal vertices can be anything from 0 to $p-2$. Let the μ-measure of such an event with i internal vertices (call the event A_i) be α_i for $i = 0, \dots, p-2$. For example, in the case of $p = 3$ we have $\alpha_1 = \alpha$ and $\alpha_0 = \beta$. We shall continue to use α for α_{p-2}, i.e. the μ-probability that the face on the root edge contains no other boundary vertices. Note that there are several such events of type A_i, which differ only in the location of the root. However because of translation invariance, each such event has the same probability α_i. For quadrangulations ($p = 4$), there are three possible building blocks, shown in Figure 2.7(b–d).

We have a generalization of Lemma 2.6.2, that shows that the measure μ is determined by $\alpha_0 \dots, \alpha_{p-2}$, leaving us with $p-1$ degrees of freedom.

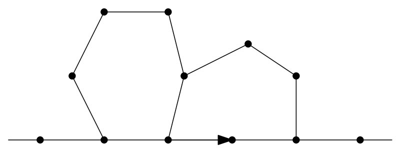

Figure 2.8: The event B_4 for $p = 6$. Depending on the order of exploration, its probability is found to be α_3^2 or $\alpha_4\alpha_2$.

However, before doing that, let us reduce these to two degrees of freedom. For any $i = 1, \dots, p-2$, consider the event B_i defined as follows (see e.g. Figure 2.8):

(i) The face incident to the root edge has $i-1$ internal vertices and its intersection with the boundary is a contiguous segment of length $p-i+1$ with the leftmost of those vertices being the root.

(ii) The face incident to the edge to the left of the root edge has i internal vertices, its intersection with the boundary is a contiguous segment of length $p-i$, with the root vertex being the right end-point.

(iii) The two faces above share precisely one common edge between them which is also incident to the root vertex.

The probability $\mu(B_i)$ can be computed by exploring the faces incident to the root edge, and with the edge to its left in the two possible orders. We find that $\alpha_{i-1}^2 = \alpha_i\alpha_{i-2}$, and hence the numbers $\{\alpha_0, \dots, \alpha_{p-2}\}$ form a geometric series, leaving two degrees of freedom. In order to simplify subsequent formulae we reparametrize these as follows. Denote

$$\beta^{p-2} = \alpha_0, \qquad\qquad \gamma^{p-2} = \alpha_{p-2}$$

so that the geometric series is given by $\alpha_i = \gamma^i\beta^{p-2-i}$. This is consistent with the previous definition of β in the case $p = 3$.

Lemma 2.10.1. *Let μ be a measure supported on $\mathcal{H}'_p$ which is translation invariant and domain Markov. Let Q be a finite simply connected simple p-angulation and $2 \le k < |\partial Q|$. As before, $A_{Q,k}$ is the event that Q is isomorphic to a sub-map of M with k consecutive vertices being mapped to the boundary of M. Then*

$$\mu(A_{Q,k}) = \alpha_1^{|V(Q)|-k}\alpha_0^{|F(Q)|-|V(Q)|+k} = \beta^{(p-2)|F(Q)|-|V(Q)|+k}\gamma^{|V(Q)|-k}. \tag{2.10.1}$$

Furthermore, if μ satisfies (2.10.1) *for any such Q and k, then μ is translation invariant and domain Markov.*

The proof is almost the same as in the case of triangulations, and we omit some of the repeated details, concentrating only on the differences.

Proof. We proceed by induction on the number of faces of Q. If Q has a single face, then we are looking at one of the events A_i. Then the face connected to the root sees i new vertices. The measure of such an event is α_i which is equal to $\alpha_0(\alpha_1/\alpha_0)^i$ since $\{\alpha_0, \ldots, \alpha_{p-2}\}$ form a geometric series. Hence (2.10.1) holds.

In general, the face Γ connected to the root can be connected to the boundary of Q and to the interior of Q in several possible ways. $Q \setminus \Gamma$ has several components some of which are connected to the infinite component of $M \setminus Q$ and some are not. We shall explore the components not connected to the infinite component of $M \setminus Q$ first, then the face Γ and finally the rest of the components. Note that in every step of exploration if we encounter an event of type A_i, we get a factor of $\alpha_1^v \alpha_0^{f-v}$ for the probability, where v is the number of *new* vertices added and f is the number of new faces added since $\{\alpha_0, \alpha_1, \ldots, \alpha_p\}$ are in geometric progression. Notice that the number of new vertices in all the components and Γ add up to that of Q and similarly the number of faces in all the components and Γ also add up to that of Q. Also in each component the number of faces is strictly smaller than Q. Hence we use induction hypothesis to finish the proof of the claim similarly as in Lemma 2.6.2. $\square$

Returning to the proof of Theorem 2.2.1, let $Z_m(x) = \sum_{i \geq 0} \psi_{m,i}^{(p)} x^i$ be the generating function for p-angulations of an m-gon with weight x for each internal vertex. The probability of any particular configuration for the face containing the root is found by summing (2.10.1) over all possible ways of filling the holes created by removal of the face. A hole which includes $k \geq 2$ vertices from the boundary of the half planar p-angulation and has a total boundary of size m can be filled in $\psi_{m,n}^{(p)}$ ways with n additional vertices. A p-angulation of an m-gon with n internal vertices has $\frac{m+2n-2}{p-2}$ faces, and so each of these contributes a factor of

$$\beta^{(p-2)|F(Q)|-|V(Q)|+k} \gamma^{|V(Q)|-k} = \beta^{n+k-2} \gamma^{n+m-k}.$$

to the product in (2.10.1). Summing over p-angulations, these weights add up to

$$\beta^{k-2} \gamma^{m-k} Z_m(\beta\gamma).$$

Now, suppose there are a number of holes with boundary sizes given by a sequence (m_i) involving (k_i) boundary vertices respectively (see Figure 2.9).

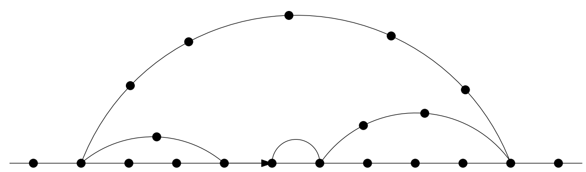

Figure 2.9: A possible configuration for the root face in a 13-angulation. The hole parameters (k_i, m_i) from left to right are $(4,5)$, $(2,2)$, $(5,7)$. There are $j = 5$ vertices exposed to infinity, so the probability of this configuration is $\alpha_5 \cdot (\beta^2\gamma Z_5) \cdot (Z_2) \cdot (\beta^3\gamma^2 Z_7)$.

Since any p-angulation can be placed in each of the holes and the weights are multiplicative, the total combined probability of all ways of filling the holes is

$$\prod_i \beta^{k_i-2}\gamma^{m_i-k_i} Z_{m_i}(\beta\gamma).$$

This must still be multiplied by a probability α_j of seeing the face containing the root conditioned on any compatible filling of the holes (see Figure 2.9). Thus we have the final identity $R(\beta,\gamma) = 1$, where we denote

$$R(\beta,\gamma) = \sum \alpha_j \prod_i \beta^{k_i-2}\gamma^{m_i-k_i} Z_{m_i}(\beta\gamma), \tag{2.10.2}$$

where the sum is over all possible configurations for the face containing the root edge, and (m_i, k_i) and j are as above.

For any possible configuration for the face at the root, and each hole it creates we have $k_i \geq 2$ (since $k = 1$ would imply a self-loop) and $m_i \geq k_i$ (since k counts a subset of the vertices at the boundary of the hole). We also have $\alpha_j = \gamma^j \beta^{p-2-j}$, and so each term in R is a power series in β, γ with all non-negative coefficients. In particular, R is strictly monotone in β and γ, and consequently for any γ there exists at most a single β so that $R(\beta,\gamma) = 1$. ☐

As an example of (2.10.2), consider the next simplest case after $p = 3$, namely $p = 4$. Here, there are 8 topologically different configurations for the face attached to the root, shown in Figure 2.10. Of those, in the leftmost shown and its reflection the hole must have a boundary of size at least 4. In all others, the hole or holes can be of any even size. summing over the possible even sizes, we get the total

$$R = \gamma^2 + \frac{4\gamma}{\beta} Z - 2\gamma\beta Z_2(\beta\gamma) + \frac{3}{\beta^2} Z^2,$$

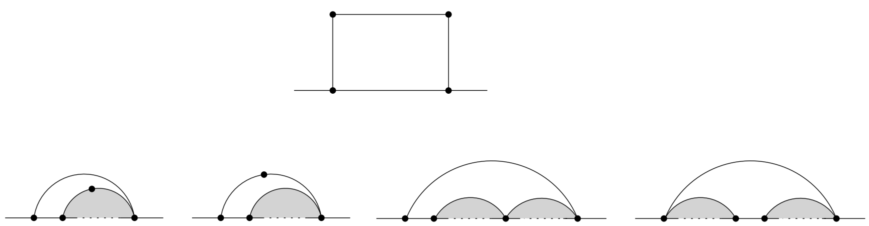

Figure 2.10: Possible faces incident to a boundary edge for quadrangulations. The first three may also be reflected to give the 8 topologically distinct possibilities. The holes (shaded) can have boundary of any even length.

where $Z = \sum_{k\geq 2} \beta^k Z_k(\beta\gamma)$ is the complete generating function for simple-faced quadrangulations with a simple boundary.

To get existence of the measures $\mathbb{H}_\alpha^{(p)}$, we need to show that for any $\gamma = \alpha^{1/(p-2)}$ there exists a β so that $R(\beta, \gamma)$ as defined in (2.10.2) equals 1. By monotonicity, and since $R(0, \gamma) = \gamma^{p-2} < 1$ (the only term with no power β corresponds to the event A_{p-2} with probability α), it suffices to show that some β satisfies $1 \leq R(\beta, \gamma) < \infty$. Note that just from steps of type A_0 and A_{p-2} we get $R(\beta, \gamma) \geq \beta^{p-2} + \gamma^{p-2}$. Thus for β close to 1 we have $R(\beta, \gamma) > 1$, provided it is finite. We prove this holds at least for α sufficiently close to 1:

Proposition 2.10.2. *For any $p \geq 4$, and any $\alpha \in (\alpha_0(p), 1)$ there is some β so that $R(\beta, \alpha^{1/(p-2)}) > 1$, and so the measure $\mathbb{H}_\alpha^{(p)}$ exists for $\alpha > \alpha_0(p)$.*

Proof. To see that $Z_m(q) < \infty$ for small enough q we need that the number of p-angulations grows at most exponentially. For triangulations or even p this is known from exact enumerative formulae. For any p-angulation we can partition each face into triangles to get a triangulation of the m-gon. The number of those is at most exponential in the number of vertices. The number of p-angulations corresponding to a triangulation is at most 2 to the number of edges, as each edge is either in the p-angulation or not. Thus we get a (crude) exponential bound also for odd p.

It is easy to see that there exists a $0 < q_c < 1$ such that $Z_m(q) < \infty$ for $q < q_c \neq 0$. We expect $Z_m(q_c) < \infty$ as well, though that is not necessary for the rest of the argument. Now we need some general estimate giving exponential growth of Z_m. Fix any $q < q_c$. Note that $\psi_{m,n} \geq \psi_{m+p-2,n-p+2}$ by just counting maps where the face containing the root is incident to no other boundary vertices. Thus $Z_m(q) \geq q^{p-2} Z_{m+p-2}(q)$, and so $Z_m(q) \leq$

Cq^{-m} for some constant $C > 0$, provided it is finite. Of course, this crude bound does not give the correct rate of increase for Z as $m \to \infty$.

In each term of (2.10.2), the $m_i - k_i$ are bounded, but while keeping them fixed, the k_i's could take any value (subject to parity constraints for even p). Fixing $m_i - k_i$ and summing over the possibilities for the k_i's we see that $R(\beta, \gamma) < \infty$ provided that $\sum_m \beta^m Z_m(\beta\gamma) < \infty$. Now $Z_m(q)$ is an increasing function of q as long as it is finite since all the coefficients of Z_m are non-negative integers. Thus we have for $\beta = q_c/(4\gamma)$, any choice of $\gamma > 1/2$ and the estimate on Z_m found above,

$$\sum_m \beta^m Z_m(\beta\gamma) = \sum_m \beta^m Z_m\left(\frac{q_c}{4}\right) < \sum_m \left(\frac{q_c}{4\gamma}\right)^m Z_m\left(\frac{q_c}{2}\right) < \sum_m (2\gamma)^{-m} < \infty \tag{2.10.3}$$

Thus for a choice of γ close to 1 and $\beta = q_c/4\gamma$ we have $R(\beta, \gamma) < \infty$ and $R(\beta, \gamma) \geq \beta^{p-2} + \gamma^{p-2} > 1$.

Having found a γ so that $R(\beta, \gamma) = 1$, we know the probability that the map contains any given finite neighborhood of the root. The rest of the construction is similar to the triangulation case as described in Section 2.7 with no significant changes. $\square$

Based on the behavior in the case of $p = 3$, we expect the measures $\mathbb{H}_\alpha$ to exist for all $\alpha < 1$. Moreover, we expect that $R(q_c/\gamma, \gamma) > 1$ when $\gamma^{p-2} = \alpha > \alpha_c$ and that for smaller γ the maximal finite value taken by R is exactly 1 where α_c will be a critical value of α at which a phase transition occurs analogous to the triangulation case. We see below that $\mathbb{H}_\alpha^{(4)}$ exists for $\alpha \leq \frac{3}{8}$, and a similar argument holds for other even p (when there are explicit enumeration results).

2.11 Non-simple p-angulations

Finally, let us address the situation with p-angulations with non-simple faces. In the case of p-angulations for $p > 3$, uniqueness breaks down thoroughly, and a construction similar to Section 2.9 applies. For even p self-loops are impossible since a p-angulation is bi-partite. However, inspection of the construction of $\mathbb{H}_{\alpha,q,\nu}$ shows that it works not because of the self-loop, but because it is possible for a single face to completely surround other faces of the map.

Consider first the case $p = 4$, and suppose we are given a measure μ supported on $\mathcal{H}_4$ satisfying translation invariance and the domain Markov property. Take a sample from μ, and replace each edge by an independent

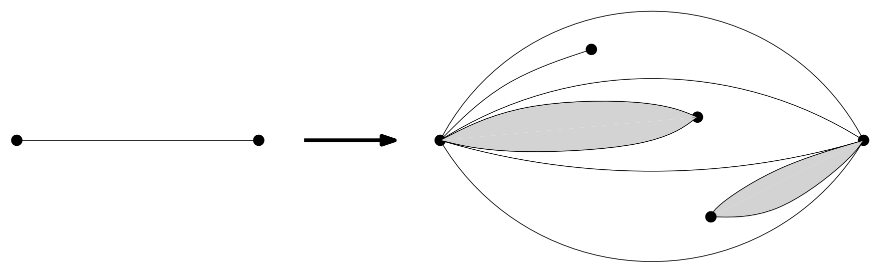

Figure 2.11: Non-uniqueness for quadrangulations: each edge is replaced with a geometric number of parallel edges. In each 2-gon an internal 2–gon is added at a uniformly chosen endpoint, and filled with an independent finite (possibly empty) quadrangulation.

geometric number of parallel edges. In each of the 2-gons created, add another 2-gon attached to one of the two vertices with equal probability, thereby creating a quadrangle. Fill the smaller 2-gons with i.i.d. samples from an arbitrary distribution supported on quadrangulations of 2-gons (see Figure 2.11). As with triangulations, this results in a measure which is domain Markov and translation invariant.

Hence we see that faces which completely surround other faces of the map prevent us from getting only a one-parameter family of domain Markov measures. For triangulations and quadrangulations, the external boundary of such a face can only consist of 2 edges (i.e. there are precisely two edges connecting the face to the infinite component of the complement). Removing such faces and identifying the two edges results in a domain Markov map with simple faces, which falls into our classification. Similarly to Proposition 2.9.1, it is possible to get a complete characterization of all domain Markov maps on quadrangulations in terms of α, the density γ of non-simple faces, and a measure ν on quadrangulations in a 2-gon.

For $p \geq 5$, things get messier. Similar constructions work for any $p > 3$, with inserted 2-gons for even p, and any combination of 2-gons and self-loops for p odd. However, here this no longer gives all domain Markov p-angulations. A non-simple face can have external boundary of any size from 2 up to $p - 1$ (with parity constraint for even p). Thus it is not generally possible to get a p-angulation with simple faces from a general one. Removing the non-simple faces leaves a domain Markov map with simple faces of unequal sizes. It is possible to classify such maps, and these are naturally parametrized by a finite number of parameters, since we must also allow for the relative frequency of different face sizes. Much of such a classification is similar to the proofs of Theorems 2.2.1 and 2.6.1, and we do not pursue this here.

2.12 Approximation by finite maps

We prove Theorem 2.2.3, identifying the local limits of uniform measures on finite triangulations in this section. Here, we are concerned only with the measures $\mathbb{H}_\alpha$ on triangulations for critical and sub-critical $\alpha \le 2/3$. Recall from the statement of the theorem, that we have sequences $(m_l)_{l\in\mathbb{N}}$, $(n_l)_{l\in\mathbb{N}}$ of integers such that $m_l/n_l \to a$ for some $a \in [0,\infty]$ and $m_l, n_l \to \infty$. We show that μ_{m_l,n_l} — the uniform measure on triangulations of an m-gon with n internal vertices — converges weakly to $\mathbb{H}_\alpha$ where $\alpha = \frac{2}{2a+3}$. To simplify the notation, we drop the index l from the sequences m_l and n_l and assume that m is implicitly a function of n. Note that since $[0,\infty]$ is compact, it follows that $\{\mathbb{H}_\alpha\}_{\alpha\le 2/3}$ are all the possible local limits of the $\mu_{m,n}$s.

Here is an outline of the proof: A direct computation shows that the $\mu_{m,n}$ measure of the event that the hull of the ball of radius r is a particular finite triangulation T converges to the $\mathbb{H}_\alpha$ measure of the same event (for any T), as given by Lemma 2.7.2. While *a priori* this only gives convergence in the vague topology, since the limit $\mathbb{H}_\alpha$ is a probability measure, it actually follows that $\mu_{m,n}$ is a tight family of measures and hence converges weakly. Thus we show the convergence of the hulls of balls. Note that the hulls of balls around the root always have a simple boundary.

We start with a simple estimate on relative enumerations on the number of triangulations of a polygon.

Lemma 2.12.1. *Suppose $m, n \to \infty$ so that $m/n \to a$ for some $a \in [0,\infty]$. Then for any fixed $j, k \in \mathbb{Z}$,*

$$\lim_{n,m\to\infty} \frac{\phi_{n-k,m-j}}{\phi_{n,m}} = \left(\frac{(a+1)^2}{(2a+3)^2}\right)^j \left(\frac{2(a+1)^2}{(2a+3)^3}\right)^k$$

Proof. By applying Stirling's approximation to (1.3.1), we have for m, n

large

$$\begin{aligned}\phi_{n,m+2} &= \frac{2^{n+1}(2m+1)!(2m+3n)!}{(m!)^2 n!(2m+2n+2)!} \\ &\sim c_1 \frac{2^{n+1}(2m+1)!}{(m!)^2}\left(\frac{(2m+3n)^{2m+3n+1/2}}{(2m+2n+2)^{2m+2n+5/2} n^{n+1/2}}\right) \\ &\sim c_2 2^n 4^m \sqrt{m}\left(\frac{27}{4}\right)^n \left(\frac{9}{4}\right)^m n^{-5/2}\left(1+\frac{2m}{3n}\right)^{2m+3n}\left(1+\frac{m}{n}\right)^{-2m-2n}\end{aligned} \tag{2.12.1}$$

Taking the ratio, we have

$$\frac{\phi_{n-k,m+2-j}}{\phi_{n,m+2}} \sim \left(\frac{2}{27}\right)^k \left(\frac{1}{9}\right)^j \frac{(1+\frac{m}{n})^{2j+2k}}{(1+\frac{2m}{3n})^{2j+3k}} \quad \times \\ \left(\frac{1+\frac{2m-2j}{3n-3k}}{1+\frac{2m}{3n}}\right)^{2m+3n} \left(\frac{1+\frac{m-j}{n-k}}{1+\frac{m}{n}}\right)^{-2m-2n}. \tag{2.12.2}$$

An easy calculation shows that the product of the last two terms in the right hand side of (2.12.2) converges to 1. Indeed, if a is finite then the first tends to $e^{-2j+2ak}$ and the second to e^{2j-2ak}. If $a=\infty$ then after shifting a factor of $\left(\frac{n}{n-k}\right)^{2m}$ from the first to the second, the limits are e^{-2j} and e^{2j}.

The result follows by taking the limit and using the fact that m/n converges to a. $\square$

Let $A_Q, V(Q), F(Q), V(B)$ be as in Lemma 2.6.2, and note that A_Q makes sense also when looking for Q as a sub-map of a finite map.

Lemma 2.12.2. *Suppose $m,n\to\infty$ with $m/n\to a$ for some $a\in[0,\infty]$. Then*

$$\lim_{m,n} \mu_{m,n}(A_Q) = \left(\frac{2}{2a+3}\right)^{|V(Q)|-|V(B)|}\left(\frac{a+1}{2a+3}\right)^{2(|F(Q)|-|V(Q)|+|V(B)|)}.$$

Remark 2.12.3. *If we make the change of variable $\alpha = 2(2a+3)^{-1}$, then Lemma 2.12.2 gives us*

$$\lim_{m,n} \mu_{m,n}(A_Q) = \alpha^{|V(Q)|-|V(B)|}\left(\frac{(2-\alpha)^2}{16}\right)^{(|F(Q)|-|V(Q)|+|V(B)|)}.$$

From Lemma 2.12.2 we can immediately conclude that the $\mu_{m,n}$-measure of A_Q converges to the $\mathbb{H}_\alpha$ measure of the corresponding event.

Corollary 2.12.4. *Suppose $m, n \to \infty$ with $m/n \to a$ for some $a \in [0, \infty]$. Then we have*

$$\lim_{m,n} \mu_{m,n}(A_Q) = \mathbb{H}_\alpha(A_Q)$$

where $\alpha = \frac{2}{2a+3}$.

Proof of Lemma 2.12.2. It is clear that the number of simple triangulations of an $m+2$-gon with n internal vertices where A_Q occurs is $\phi_{n-k,m+2-j}$ where $j = 2|V(B)| - |\partial Q| - 2$ where $|\partial Q|$ is the number of vertices in the boundary of Q, and $k = |V(Q)| - |V(B)|$. Then from Lemma 2.12.1, we have

$$\lim_{m,n} \mu_{m,n}(A_Q) = \lim_{n,m} \frac{\phi_{n-k,m+2-j}}{\phi_{n,m+2}} = \left(\frac{(1+a)^2}{(2a+3)^2}\right)^j \left(\frac{2(a+1)^2}{(2a+3)^3}\right)^k$$

From Euler's formula, it is easy to see that $|F(Q)| = 2|V(B)| - |\partial Q| - 2$. This shows $j + k = |F(Q)| - |V(Q)| + |V(B)|$. Using all this, we have the Lemma. □

Proof of Theorem 2.2.3. Corollary 2.12.4 gives convergence for cylinder events. Since $\mathbb{H}_\alpha$ is a probability measure, the result follows by Fatou's lemma. □

2.13 Quadrangulations and beyond

Can we get similar finite approximations for $\mathbb{H}_\alpha^{(p)}$ for $p > 3$? We think it is possible to prove such results based on enumeration of general p-angulations with a boundary, which is available for p even. We believe that similar results should hold for any p, though do not see a way to prove them. Let us present here a recipe for quadrangulations. For higher even p there are additional complications as the core is no longer a p-angulation and results on maps with mixed face sizes are needed.

Let us first consider quadrangulations with a simple boundary. Denote by $\mathcal{Q}_{2m,n}$ the space of quadrangulations with simple boundary size $2m$ and number of internal vertices n (note that since the quadrangulation is bipartite, the boundary size is always even). Let $q_{2m,n} = |\mathcal{Q}_{2m,n}|$ be its cardinality. Enumerative results are available in this situation (see [30]). We alert the reader that our notation is slightly different from [30]: they use $\tilde{q}_{2m,n}$ for quadrangulations with a simple boundary and n denotes the number of faces, not the number of internal vertices. Using Euler's formula one can easily change from one variable to the other. Doing that, we get:

$$q_{2m,n} = 3^{n-1} \frac{(3m)!}{m!(2m-1)!} \frac{(2n+3m-3)!}{n!(n+3m-1)!} \tag{2.13.1}$$

Now suppose $m/n \to a$ for some $a \in [0, \infty]$ where m and n are sequences such that $m \to \infty$ and $n \to \infty$. Let $\nu_{2m,n}$ be the uniform measure on all quadrangulations of boundary size $2m$ and n internal vertices. A straightforward computation similar to Lemmas 2.12.1 and 2.12.2 gives us for any finite Q,

$$\lim_{m,n} \nu_{2m,n}(A_Q) = \left(\frac{4(1+3a)^3}{27(2+3a)^3}\right)^{|F(Q)|} \cdot \left(\frac{9(2+3a)}{4(1+3a)^2}\right)^{|V(Q)|-|V(B)|} \tag{2.13.2}$$

where $V(Q)$ is the set of vertices in Q, $V(B)$ is the set of vertices of Q on the boundary of M, and $F(Q)$ is the set of faces in Q (by Euler's characteristic, the "change" in the boundary length when removing Q is $2(|V(Q)|-|V(B)|-|F(Q)|)$).

The limit (2.13.2) in itself is not enough to give us distributional convergence of $\nu_{2m,n}$, as we are missing tightness. It is possible to get tightness for $\nu_{2m,n}$ using the same ideas presented for example in [13, 65] or the general approach found in [21]. The key is that it suffices to show the tightness of the root degree. The interested reader can work out the details and we shall not go into them here. Instead, throughout the remaining part of this section, we shall assume that the distributional limits of $\nu_{2m,n}$ exist. We remark here that when $a = 0$ the limiting measures of the events described by (2.13.2) matches exactly with that of the half planar UIPQ measure (see [44]) and that for $a = \infty$ we get the dual of a critical Galton-Watson tree conditioned to survive. Thus in these two extreme cases, the distributional limit has already been established.

To handle all a, we define the operator core : $\mathcal{H}_4 \to \mathcal{H}_4'$, which is the reverse of the process used to define the measures $\mathbb{H}_{\alpha,q,\nu}$ in Section 2.9, and acts on the dual by taking the 2-connected core. Formally, any face which is not simple must have an external double edge connecting it to the rest of the map (and a 2-gon inside it). The core operator removes every such face, and identifies the two edges connecting it to the outside. This operation is defined in the same way on quadrangulations of an m-gon. As discussed in Section 2.9, if μ is domain Markov on $\mathcal{H}_4$ then $\mu \circ \mathrm{core}^{-1}$ is domain Markov on $\mathcal{H}_4'$.

Let $\mu = \lim \nu_{2m,n}$ as $m, n \to \infty$ with $m/n \to a \in [0, \infty]$. We first observe that μ is domain Markov and translation invariant. This follows from (2.13.2) and the converse part of Lemma 2.10.1.

Next, observe that the events A_i for $i = 0, 1, 2$ are not affected by core. This is because in each of them, the face containing the root is a simple face, and so is not contained in any non-simple face. At this point from (2.13.2), we obtain $\beta^2 = (4(1+3a)^3)/(27(2+3a)^3)$ and $\gamma/\beta = (9(2+3a))(4(1+3a)^2)$. Thus,

$$\mu(A_2) = \frac{3}{4(1+3a)(2+3a)} \qquad\qquad \mu(A_0) = \frac{4}{27}\left(\frac{1+3a}{2+3a}\right)^3.$$

From the first we see that as a goes from 0 to ∞ we get $\alpha \in [0, \frac{3}{8}]$. Solving for a in terms of $\alpha = \mu(A_2)$ and plugging in we find $\beta = \sqrt{\mu(A_0)} = \frac{2}{27}(\sqrt{3+\alpha} - \sqrt{\alpha})^3$, which decreases from $\sqrt{4/27}$ to $\sqrt{1/54}$ as α increases from 0 to 3/8.

This gives the measures $\mathbb{H}_\alpha^{(4)}$ as the the core of the limit of uniform measures on non-simple quadrangulations. Since the core operation is continuous in the local topology, this is also the limit of the core of uniform quadrangulations. This does not give $\mathbb{H}_\alpha^{(4)}$ as a limit of uniform measures on non-simple quadrangulations, since the number of internal vertices in the core of a uniform map from $\mathcal{Q}_{2m,n}$ is not fixed. Thus the above only proves the limit when n is taken to be random with a certain distribution (though concentrated and tending to infinity in proportion to m.) It should be possible to deduce (though we have not proved it) that uniform simple quadrangulations converge to $\mathbb{H}_\alpha^{(4)}$ by using a local limit theorem for the distribution of the size of the core (see [15, 16]).

The above indicates that a phase transition for the family $\mathbb{H}_\alpha^{(4)}$ occurs at $\alpha = 3/8$, similar to the case $p = 3$. We can similarly compute the asymptotics of p_k as in Section 2.8 and see that $p_k \sim ck^{-5/2}$ for $\alpha = 3/8$ and $p_k \sim ck^{-3/2}$ for $\alpha < 3/8$. This indicates different geometry of the maps. All these hints encourage us to conjecture that a similar picture of phase transition do exist for the measures $\mathbb{H}_\alpha^{(p)}$ for all $p > 3$.

Chapter 3

Half planar maps: geometry

[2] As a continuation of Chapter 2, we shall study the geometry of domain Markov half planar random maps and also analyze percolation on them. We shall focus mainly on half planar triangulation and the main tool will be the peeling process. Recall that Theorem 2.2.1 for $p = 3$ gives the classification result for domain Markov half planar triangulations. Also recall the parameter α which denotes the probability that the triangle incident to the root edge has the other vertex not on the boundary. Recall the phase transition as described in Section 2.8. In this chapter, we focus on the subcritical and supercritical phases of domain Markov half planar triangulations, and analyze this phase-transition in more detail. In particular, we obtain results for volume growth, isoperimetry and geometry of percolation clusters in the supercritical and subcritical phases of these maps. Finally, we extract some information about the behavior of random walk on these maps from these geometric informations.

We introduce some new notations which will be used throughout this chapter. A step of the form (L, i) (resp. (R, i)) is the event that the triangle incident to the root edge is attached to a vertex on the boundary which is at a distance i to the left (resp. right) of the root edge along the boundary (see Figure 3.2). We shall also talk about such events with the root edge replaced by any fixed edge on the boundary of the map. Because of translation invariance, the measures of such events do not depend on the edge we want to consider and it was also shown in Chapter 2 that for any fixed $i \geq 1$, the measures of (L, i) and (R, i) are the same. Let $p_{i,k}$ denote the measure of the event that a step of the form (L, i) or (R, i) occurs and the triangle incident to the root edge separates k internal vertices of the map from infinity. Let $p_i = \sum_k p_{i,k}$. We will use some of the asymptotics of p_i described in Section 2.8. In particular recall that $p_i \sim ci^{-3/2}$ in the subcritical phase and on the other hand, $p_i \sim \exp(-c'i)$ in the supercritical phase for some positive constants c, c'.

[2] The results of this chapter are from the paper [82]

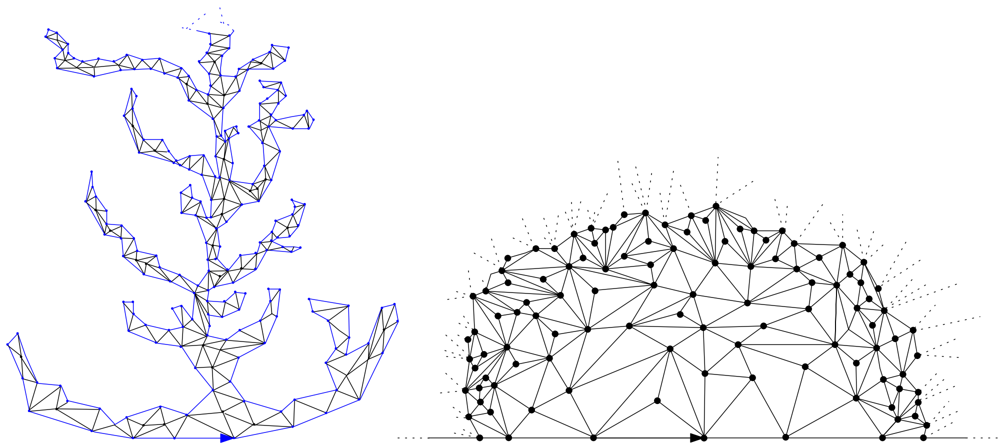

Figure 3.1: An illustration (artistic) of the geometry of a subcritical half planar triangulation to the left and that of supercritical to the right. The blue edges in the subcritical map is the boundary of the map.

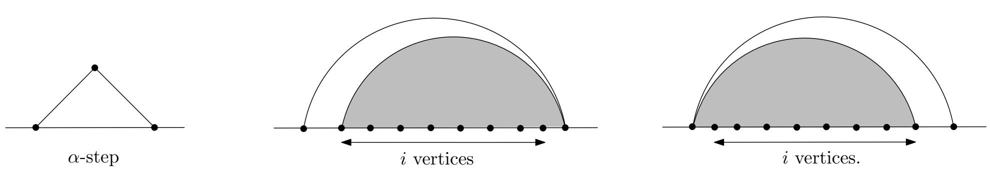

Figure 3.2: Left: An α-step. Centre: A step of the form (R, i). Right: A step of the form (L, i). The gray area denotes some unspecified triangulation.

3.1 Main results

3.1.1 Geometry

We present below the results obtained in this chapter first for supercritical and then for subcritical maps. Roughly, the behavior of supercritical maps are hyperbolic: they have exponential volume growth and anchored expansion which will imply transience of simple random walk. The subcritical maps behave, in view of their geometry, roughly like the critical Galton-Watson tree conditioned to survive (see [63]). They have quadratic volume growth, and infinitely many cut-sets of finite size. All the terms stated in this paragraph will be defined rigorously below.

We remark here that the geometric properties are certainly very different from the critical uniform infinite half planar triangulation (UIHPT). For results of similar nature regarding the UIHPT, see [7, 8, 10]. Predictably, the asymptotics of p_i stated in Section 2.8 play a crucial role in proving the results stated below for different regimes of the parameter α.

Supercritical

Roughly, the geometry of maps in the supercritical regime can be viewed as a collection of supercritical trees one attached to each of the boundary vertices with some horizontal segments being added (see Figure 3.3). Hence, we can expect exponential volume growth, large cut-sets and positive speed of random walk on these maps. In this chapter, we confirm some of these heuristics.

Throughout this subsection, we assume $\alpha \in (2/3, 1)$. For a set X, $|X|$ denotes its cardinality, and by an abuse of notation, for any finite graph or map G, $|G|$ will denote its number of vertices. The **ball** of radius r in a map denotes the sub-map formed by all the faces which have at least one vertex incident to it which is at a distance strictly less than r from the root vertex along with all the edges and vertices incident to them. Recall that the **hull** of radius r is the ball of radius r along with all the finite components of its complement. Note that since the half planar maps are one-ended, there will be at most one infinite component in the complement of the ball and the hull is always a simply connected sub-map. The **internal boundary** of a simply connected sub-map is the set of vertices and edges in the sub-map which is incident to at least one finite face which is not in the sub-map. Clearly, the internal boundary of a hull is a connected simple path in the map. We denote the hull of radius r around the root of a map with law $\mathbb{H}_\alpha$ by $B_r(\alpha)$ and its internal boundary by $\partial B_r(\alpha)$.

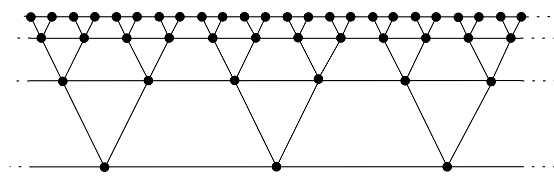

Figure 3.3: An very rough intuition of the geometry of supercritical maps.

We first show exponential volume growth of the hull and the boundary of the hull.

Theorem 3.1.1. *There exists some positive constants* $1 < c < C$ *depending only on* α *such that almost surely,*

$$\limsup \frac{|\partial B_r(\alpha)|}{C^r} = 0 \quad \textit{and} \qquad \liminf \frac{|\partial B_r(\alpha)|}{c^r} = \infty \tag{3.1.1}$$

and also,

$$\limsup \frac{|B_r(\alpha)|}{C^r} = 0 \quad \textit{and} \qquad \liminf \frac{|B_r(\alpha)|}{c^r} = \infty \tag{3.1.2}$$

Having established the exponential volume growth, we ask if there are small cut-sets in the map. The usual parameter to look for in this situation is the Cheeger constant but since our maps are random and any finite configuration does occur almost surely somewhere in the map, the correct parameter to consider is the anchored expansion (see [72] Chapter 6).

For a graph G, let $V(G)$ denote its set of vertices. For any graph G, and a subset of vertices $S \subset V(G)$, let $|\partial_E S|$ denote the number of edges in G with one vertex in S and another in $V(G) \setminus S$. Also let $|S|_E$ denote the sum of the degrees of the vertices in S. The **anchored expansion constant** $i^*_E(G)$ of a graph G is defined as follows:

$$\liminf_{n\to\infty} \left\{ \frac{|\partial_E S|}{|S|_E}; S \subset V(G) \text{ is connected, } v \in S, |S|_E \geq n \right\} = i^*_E(G)$$

Although we specify a vertex v in the above definition, the definition is independent of the choice of v. If $i^*_E(G) > 0$, we say G has anchored expansion.

Theorem 3.1.2. *A half planar triangulation with law* $\mathbb{H}_\alpha$ *for* $\alpha \in (2/3, 1)$ *has anchored expansion almost surely.*

We remark here that the exponential lower bound for the volume growth can be concluded from anchored expansion, but we prove it using a different procedure involving an exploration process because we use the same exploration process to study the subcritical maps and also we get an upper bound on the volume growth using this method.

It is shown in [90] that for bounded degree graphs having anchored expansion, the random walk has positive speed, that is, the graph distance between the position of the walker after the nth step and the starting position grows linearly. Unfortunately our maps are not bounded degree maps, so we cannot directly translate the result. However we can conclude using Theorem 3.5 of [86] and Theorem 3.1.2 that

Corollary 3.1.3. *Simple random walk on a map with law $\mathbb{H}_\alpha$ is transient almost surely if $\alpha \in (2/3, 1)$.*

We believe that the random walk does have positive speed almost surely for supercritical maps and in fact the walker should move away from the boundary at linear speed. What we need to do to prove this is to control the slowing down of the walk due to the presence of high degree "traps" in the map.

Subcritical

Throughout this subsection, $\alpha \in [0, 2/3)$. The journey of understanding subcritical triangulations begins with a result about their cut-sets. A **cut-set** of an infinite rooted graph G is a connected subgraph of G which when removed breaks up G into two or more connected components, the root being in the finite component.

Proposition 3.1.4. *In a half planar triangulation with law $\mathbb{H}_\alpha$ where $\alpha \in [0, 2/3)$, there exists infinitely many cutsets each of which consists of a single edge almost surely.*

Because of the domain Markov property, the above proposition means that a triangulation distributed as a subcritical $\mathbb{H}_\alpha$ has i.i.d. copies of finite triangulation each of which share only a single edge with each other (see Figure 3.1). Proposition 3.1.4 and the Nash-Williams criterion for recurrence (see [71] Proposition 9.15) immediately implies

Corollary 3.1.5. *Simple random walk on a half planar triangulation with law $\mathbb{H}_\alpha$ is recurrent almost surely for $\alpha \in (0, 2/3)$.*

It is interesting to study the dual maps of these maps. As described in Chapter 2, we make such a dual map locally finite by breaking the infinite degree vertex corresponding to the infinite face into infinitely many leaves. Hence the dual almost surely consists of i.i.d. sequence of 3-regular graphs (with leaves) connected to each other by a single edge. For $\alpha = 0$, it can actually be seen that the dual is a critical Galton-Watson tree conditioned to survive (see Chapter 2) and hence, the maps for different values of α can be seen as an "interpolation" between the UIHPT and critical trees. In fact, we believe that the scaling limit of such maps in the sense of local Gromov-Hausdorff topology exists and is the infinite non-compact continuum random tree (CRT). This can be viewed as the tangent cone at the root of the compact Aldous CRT (see [42]). In fact from this heuristic, we expect that the random walk behaves more or less similarly as a random walk in the critical tree conditioned to survive. The spectral dimension of the subcritical maps should be almost surely $4/3$ and should follows from much more general results for strongly recurrent graphs studied in [17, 66].

As can be expected from Proposition 3.1.4, the boundary size of the hull of radius r is a tight sequence. We prove a stronger result: the boundary sizes of the hull has exponential tail.

Theorem 3.1.6. *Let $\alpha \in [0, 2/3)$. Then there exists some positive constant $c > 0$ (depending only on α) such that*

$$\mathbb{H}_\alpha(|\partial B_r(\alpha)|) > n) < e^{-cn}$$

for all $n \geq 1$.

The following central limit theorem shows that the volume growth is quadratic: another tree-like behavior.

Theorem 3.1.7. *Suppose $\alpha \in [0, 2/3)$. Then*

$$\frac{|B_r|}{r^2} \to S_{1/2}(\alpha)$$

in distribution where $S_{1/2}(\alpha)$ is a stable random variable with parameter $1/2$ and its other parameters depend only upon α and nothing else.

3.1.2 Percolation

We are mainly interested in quenched statements about Bernoulli site percolation on random triangulations: take a half planar triangulation T with law $\mathbb{H}_\alpha$ and color each vertex independently black with probability p or white with probability $1-p$. A black (resp. white) cluster is a connected component induced by the black (resp. white) vertices on the map. Given a half planar map, denote the percolation measure on it by P_p and the expectation by E_p. Let $\mathbb{P}_p$ denote the overall measure induced by the percolation configuration on such random maps and let $\mathbb{E}_p$ denote the expectation with respect to the measure $\mathbb{P}_p$. It is understood that in these notations there is a hidden parameter α which we shall drop to lighten notation. As usual, define p_c to be the infimum over p such that there exists an infinite black cluster P_p-almost surely. Further, we will also be interested in

$$p_u = \inf\{p \in (0,1] : \text{ there exists an unique infinite cluster } P_p\text{-almost surely}\}$$

Further history of the study of percolation of random maps is discussed in Section 1.4.

Notice that it is immediate to see via Proposition 3.1.4 that percolation is uninteresting in the subcritical maps (in this case $p_c = 1$ almost surely.) It was conjectured (see [24]) by Benjamini and Schramm that on non-amenable quasitransitive graphs, $p_c < p_u$. For supercritical maps, because of anchored expansion as depicted by Theorem 3.1.2, we would expect a similar behavior.

Theorem 3.1.8. *Fix $\alpha \in (2/3, 1)$. Then $\mathbb{H}_\alpha$-almost surely,*

(i) $p_c = \frac{1}{2}\left(1 - \sqrt{3 - \frac{2}{\alpha}}\right)$

(ii) $p_u = \frac{1}{2}\left(1 + \sqrt{3 - \frac{2}{\alpha}}\right)$

Also, $\mathbb{H}_\alpha$-almost surely, there is no infinite black cluster P_{p_c}-almost surely and there is an unique infinite black cluster P_{p_u}-almost surely.

Note that $p_c < p_u$ almost surely in the regime $\alpha \in (2/3, 1)$. It is interesting to note that as $\alpha \to 2/3$, both $p_c \to 1/2$ and $p_u \to 1/2$. But in the regime (p_c, p_u) we have more than one infinite cluster. One can easily conclude via ergodicity of these maps with respect to translation of the root along the boundary (see Chapter 2) that the number of infinite black or white clusters is actually infinite almost surely. The next Theorem shows that the number of black or white clusters in fact have positive density along the boundary. Let W_k^∞, B_k^∞ be the number of infinite white and black clusters respectively which share at least vertex with a vertex which is within distance k from the root along the boundary.

Theorem 3.1.9. *Fix $\alpha \in (2/3, 1)$ and suppose $p \in (p_c, p_u)$ where p_c,p_u are as in Theorem 3.1.8. There exists a positive constant $\rho > 0$ such that almost surely,*

$$\frac{W_k^\infty}{k} \to \rho, \qquad \frac{B_k^\infty}{k} \to \rho \tag{3.1.3}$$

The constant ρ is in fact half of the probability of the event of having an infinite interface starting from a boundary edge (see Section 3.1.2 for more details.)

A **ray** in an infinite percolation cluster is a semi-infinite simple path in the cluster starting from a vertex closest to the root with ties are broken arbitrarily. Two rays r_1 and r_2 are equivalent if there is another ray r_3 which intersect both these rays infinitely many times. An **end** of a cluster is an equivalence class of rays. Let $\mathsf{END}(\mathcal{C})$ denote the space of ends of a percolation cluster $\mathcal{C}$. We shall define a metric on $\mathsf{END}(\mathcal{C})$ as follows: for any two rays ξ and η on $\mathcal{C}$, define the distance between them as

$$\begin{aligned} d(\xi, \eta) = \inf\{1/n, n = 1 \text{ or } \forall X \in \xi, \forall Z \in \eta, \exists \text{ a component } K \text{ of} \\ \mathcal{C} \setminus B_n, |X \setminus K| + |Z \setminus K| < \infty\} \end{aligned}$$

It is easy to deduce that $\mathsf{END}(\mathcal{C})$ does not depend on the choice of the vertex around which we consider the graph-distance balls and that $\mathsf{END}(\mathcal{C})$ equipped with this metric is compact.

Theorem 3.1.10. *Fix $\alpha \in (2/3, 1)$. Assume p_c, p_u are as in Theorem 3.1.8 and fix $p \in (p_c, p_u)$. Then $\mathbb{H}_\alpha$-almost surely, the subgraph formed by each infinite cluster has no isolated end and has continuum many ends P_p-almost surely.*

Organization: The chapter is organized as follows. In Section 3.2 some preliminary background material is provided. In Section 3.3 results about geometry are proved. For the supercritical case: Theorem 3.1.1 is proved in Section 3.3.2, Theorem 3.1.2 is proved in Section 3.3.2. For the subcritical case, Theorems 3.1.6 and 3.1.7 are proved in Section 3.3.2. For percolation, Theorems 3.1.8–3.1.10 are proved in Section 3.4. Some open problems are discussed in Section 3.5.

3.2 Preliminaries

In this Section we gather some results which we are going to need. The reader might skip this Section in the first reading and come back to it when it is referred to in the subsequent Sections.

3.2.1 Free triangulations

Recall the Boltzmann distribution defined in Definition 2.5.1. Also recall a freely distributed triangulation with parameter q of an m-gon will be referred to as a **free triangulation** with parameter q of an m-gon. We will need few estimates on free triangulations in this chapter.

Let $I_m(q)$ denote the number of internal vertices of a freely distributed triangulation of an m-gon with parameter q.

Proposition 3.2.1. *Fix $\theta \in [0, 1/6)$ and let $q = \theta(1-2\theta)^2$. Fix an integer $m \geq 2$.*

(i) $\mathbb{E}(I_m(q)) = \frac{(m-1)(2m-3)2\theta}{(1-6\theta)m+6\theta} = \frac{4\theta}{(1-6\theta)}m + O(1)$

(ii) $Var(I_m(q)) = \frac{(m-1)(2m-3)m(1-2\theta)}{((1-6\theta)m+6\theta)^2(1-6\theta)} = \frac{2(1-2\theta)}{(1-6\theta)^3}m + O(1)$

Proof. Note the following identity

$$\mathbb{E}(I_m(q)) = \frac{q(Z'_m(q))}{Z_m(q)} = q(\log Z_m)'(q)$$

Putting $q = \theta(1-2\theta)^2$ and using Proposition 2.5.2, we obtain after an easy computation

$$\mathbb{E}(I_m(q)) = \frac{(m-1)(2m-3)2\theta}{(1-6\theta)m+6\theta}$$

The proof of (ii) is a similar computation and is left to the reader to verify. □

We would need the following estimates whose proof is postponed to appendix A.

Lemma 3.2.2. *Fix $\theta \in [0, 1/6)$ and suppose $q = \theta(1-2\theta)^2$. Suppose Y is a variable supported on $\mathbb{N} \setminus \{0\}$ such that $\mathbb{P}(Y = k) \sim ck^{-3/2}$ for some constant $c > 0$. Then there exists positive constants c_θ, c'_θ depending upon θ such that*

(i) $\mathbb{P}(Y + I_{Y+1} > x) \sim \frac{c_\theta}{\sqrt{x}}$

(ii) $\mathbb{E}(Y + I_{Y+1})\mathbb{1}_{\{Y+I_{Y+1}<x\}} \sim c'_\theta\sqrt{x}$

as $x \to \infty$.

3.2.2 Stable random variables

The theory of stable random variables play a vital role in our subsequent analysis. Fix $\alpha \in (0, 2]$. An independent sequence $X_1, X_2, \dots$ is said to follow a stable distribution of type α if $S_n = X_1 + \ldots + X_n$ satisfies

$$S_n \stackrel{(d)}{=} n^{1/\alpha} X_n + \gamma_n$$

for some sequence γ_n and the distribution of X_1 is not concentrated around 0. See for example [49] Chapter VI or [48] for more details.

We shall be needing the following classical result. This can be found in [48].

Theorem 3.2.3. *Suppose $X_1, X_2, \dots$ are i.i.d. with a distribution that satisfies*

1. $\lim_{x\to\infty} \mathbb{P}(X_1 > x)/\mathbb{P}(|X_1| > x) = \theta \in [0,1]$
2. $\mathbb{P}(|X_1| > x) = x^{-\alpha} L(x)$

where $\alpha < 2$ and L is slowly varying. Let $S_n = X_1 + \ldots X_n$. $a_n = \inf\{x : \mathbb{P}(|X_1| > x) \le n^{-1}\}$ and $b_n = n\mathbb{E}(X_1 1_{|X_1| \le a_n})$. As $n \to \infty$ $(S_n - b_n)/a_n \to Y$ in distribution where Y is a stable random variable of type α.

We are specially interested in the case $\alpha = 1/2$. It turns out that we can add a constant to a variable following a stable distribution of type α where $\alpha \neq 1$ such that $\gamma_n = 0$ for all n in its definition (see [49]). After such a centering, its density can be explicitly written as

$$(2\pi x^3)^{-1/2} \exp(-1/2x) \mathbb{1}_{\{x>0\}} \tag{3.2.1}$$

This is known as the *Lévy distribution*.

3.2.3 Peeling

Let us recall the concept of peeling introduced in Section 2.4. We only focus on triangulations in this section and to that end let us investigate what happens when we perform peeling on triangulations.

Suppose we have a sample T from $\mathbb{H}_\alpha$ for some $\alpha \in (0, 1]$. We will construct a growing sequence of simply connected sub-maps P_n containing the root with the complement T_n defined as the set of faces in T not in P_n along with the edges and vertices incident to them (note that T_n is also a half planar triangulation because P_n is simply connected.) Next we choose an edge e_n on the boundary of T_n and reveal the triangle incident to it along with all the finite components of the complement.

Start with P_0 to be empty and $T_0 = T$. At the nth peeling step, we pick an edge e_n on the boundary of T_n and add the triangle in T_n incident to e_n along with the finite components of the complement (if there is any) to P_{n+1}. Define T_{n+1} to be the complement of P_{n+1}. For all n, T_n is distributed as $\mathbb{H}_\alpha$ via the domain Markov property. By abuse of notation, sometimes we shall re-root T_n on some other edge on the boundary of T_n and the distribution of T_n does not change by translation invariance. Further note that the triangle revealed in the peeling step is either of the form described by an α-step or is of the form (L, i) or (R, i) for $i \geq 1$ (see Figure 3.2.) Recall that if the peeling step is of the form (L, i) or (R, i), the finite triangulation of the $(i+1)$ gon separated from infinity by the triangle incident to e_n is distributed as a free triangulation of the $i+1$-gon with parameter given $\alpha\beta$ where β is given by eqs. (2.8.2) and (2.8.3).

Recall also that the choice of e_n depends only on P_n and we are free to choose any edge on the boundary of P_n for the next peeling step. We shall exploit this freedom of choice of e_n to prove the different results stated in Section 3.1.

3.3 Geometry

3.3.1 Peeling algorithm

Recall from the discussion in Section 3.2.3 that we are free to choose the edge e_n on which we apply the nth peeling step. We now describe an algorithmic procedure to choose the edges in such a way that at a certain (random) step, we reveal the hull of the ball of radius r around the root vertex. The algorithm follows the idea developed in [7] for analyzing the volume growth of the full plane UIPT, but we modify it appropriately for the half plane versions. We take up the notations of Section 3.2.3. Further recall that the hull of the ball of radius r around the root is denoted by B_r.

Recall that the internal boundary of a simply connected sub-map is the set of vertices and edges in the sub-map which is incident to at least one finite face which is not in the sub-map. Let $\tau_0 = 0$ and let P_0 be the root vertex. Suppose we have defined a (random) time τ_r so that $P_{\tau_r} = B_r$ for some $r \geq 1$. In particular, the internal boundary of P_{τ_r} is ∂B_r. The idea is to iteratively peel the edges in ∂B_r till none of the vertices in ∂B_r remain in the boundary of T_n.

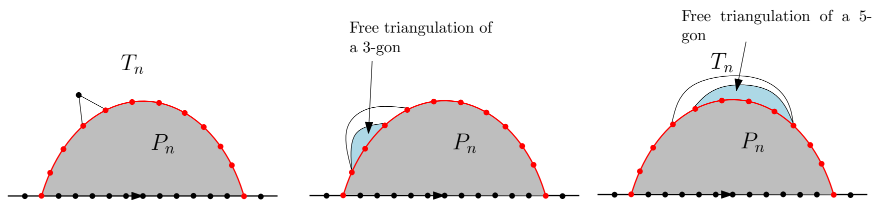

Figure 3.4: An illustration of peeling at the nth step. The gray area denote the peeled part P_n. The red vertices and edges denote the internal boundary of P_n.

Algorithm: Suppose we have described the process up to step n such that $\tau_r \le n < \tau_{r+1}$. Now look for the left most vertex v of ∂B_r which remains in the internal boundary of P_n at step n and perform a peeling step on the edge to the right of v in the boundary of T_n. If there is no vertex v of ∂B_r left in the boundary of T_{n+1}, define $n+1 = \tau_{r+1}$ and $P_{\tau_{r+1}} = B_{r+1}$.

The algorithm proceeds in such a way that for every vertex of ∂B_r, we keep on peeling at an edge incident to that vertex until it goes inside the revealed map. Hence at step τ_r, we reveal nothing but the hull of the ball of radius r for every $r \ge 1$. Recall that internal boundary of P_n is the set of edges and vertices which are incident to at least one face not in P_n. Let X_n denote the number of vertices in the internal boundary of P_n at the nth step. It is easy to see that X_n itself is not a Markov chain because the transition probabilities very much depend on the position of the edge on which we are peeling. However, a bit of thought reveals that X_{τ_r} is in fact an irreducible aperiodic Markov chain. We record the above observations in the following Proposition.

Proposition 3.3.1. *For $r \ge 1$, P_{τ_r} described in the algorithm above is the same as B_r and $X_{\tau_r} = |\partial B_r|$. Also, the sequence $\{X_{\tau_r}\}_{r\ge1}$ is an irreducible aperiodic Markov chain.*

Following the idea of [7], we estimate the size of the boundary by analyzing X_n separately for α in subcritical and supercritical regimes. Observe that $\Delta X_n = X_{n+1} - X_n \le 1$ for any n. Note also that ΔX_n has different behavior of tails in the subcritical and supercritical regimes:

$$\mathbb{H}_\alpha(\Delta X_n < -i) \approx \begin{cases} i^{-1/2} & \alpha < 2/3 \\ \exp(-ci) & \alpha > 2/3 \end{cases} \tag{3.3.1}$$

for some constant $c > 0$ and a large non negative integer $i < X_n$. Thus, conditioned on X_n, ΔX_n has negative expectation if X_n is not too small in the subcritical regime. This tells us that X_n has a drift towards 0 as soon as it gets large which implies it should be a tight sequence. However in the supercritical regime Lemma 3.3.2 below shows that ΔX_n has positive conditional expectation, which implies that X_n grows linearly. This constitutes the key point of difference between the two regimes which is made rigorous in the following Sections 3.3.2 and 3.3.3.

3.3.2 Supercritical

In this subsection, we prove Theorem 3.1.1 and hence we assume $\alpha > 2/3$ throughout this subsection. Further, we shall also borrow the notations from Sections 3.2.3 and 3.3.1.

As mentioned before, we will perform the peeling algorithm performed in Section 3.3.1 and analyze the quantity ΔX_n. To that end, we shall approximate ΔX_n by a sequence of auxiliary variables $\tilde{X_n}$ such that the variables $\Delta \tilde{X_n} = \tilde{X_{n+1}} - \tilde{X_n}$ for $n \geq 1$ form an i.i.d. sequence with $\Delta \tilde{X_n} = -i$ if a step of the form (L, i) or (R, i) occurs in the $(n+1)$th peeling step and $\Delta \tilde{X_n} = 1$ if an α step occurs in the $(n+1)th$ peeling step. Clearly, from definition, $X_n > \tilde{X_n}$ since if in a peeling step the triangle revealed has the third vertex not on the internal boundary of P_n, $\Delta X_n > \Delta \tilde{X_n}$.

Because of the exponential tail, the variables $\Delta \tilde{X_n}$ in the supercritical regime has finite variance. Further its expectation turns out to be positive.

Lemma 3.3.2.

$$\mathbb{E}(\Delta(\tilde{X_n})) = \sqrt{3\alpha - 2}\sqrt{\alpha} \tag{3.3.2}$$

In particular, $\mathbb{E}(\Delta(\tilde{X_n})) > 0$ *for* $\alpha \in (2/3, 1)$.

Proof of Lemma 3.3.2 is an easy computation which is provided in appendix A.

Lemma 3.3.3. *There exists a constant* $c > 0$ *such that almost surely*

$$c < \liminf \frac{X_n}{n} \leq \limsup \frac{X_n}{n} \leq 1 \tag{3.3.3}$$

Proof. $\limsup X_n/n \leq 1$ follows trivially because $\Delta X_n \leq 1$. Since the steps in $\Delta(\tilde{X_i})$ are i.i.d. with finite mean, strong law of large numbers imply that $\tilde{X_n}/n \to \sqrt{3\alpha - 2}\sqrt{\alpha} > 0$ almost surely as $n \to \infty$. The required lower bound now follows from the fact that $X_n > \tilde{X_n}$ by definition. □

Recall that step τ_r in the peeling algorithm marks the step when the hull of the ball of radius r is revealed. We now state some estimates on τ_r. The first part of the following Lemma 3.3.4 is essentially rephrasing Lemma 4.2 of [7]. Further, we remark that Lemma 3.3.4 is valid for any $\alpha \in [0, 1)$ and we shall use it again when dealing with the subcritical case in Section 3.3.3.

Lemma 3.3.4. *For any* $r \geq 0$*,*

(i) *There exists some constants* $A > 1$ *and* $A' > 0$ *such that for any integer* $n \geq 1$*,*

$$\mathbb{P}(\Delta\tau_r > An||X_{\tau_r}| = n) < \exp(-A'n) \tag{3.3.4}$$

(ii) *For any integer* $k \geq 1$ *and integers* $1 \leq l \leq l'$

$$\mathbb{P}(\Delta\tau_r > k||X_{\tau_r}| = l) \leq \mathbb{P}(\Delta\tau_r > k||X_{\tau_r}| = l').$$

Proof. The number of steps required for a vertex on ∂B_r to go inside the revealed map is a geometric random variable (we wait till a step of the form (L, i) occurs for some $i \geq 1$.) Thus $\Delta\tau_r$ is a sum of at most n i.i.d. geometric variables. Thus part (i) follows from a suitable large deviations estimate.

An easy coupling argument can be used to prove part (ii). To see this, let us consider two marked contiguous segments S with l vertices and S' with l' vertices on the boundary with the left most vertex being the root vertex. We can now perform the peeling algorithm described in Section 3.3.1 until all the vertices in S is inside the revealed map. Clearly, if at some step, some vertices of S are still not swallowed by the revealed map, then some vertices of S' are also not swallowed. ☐

Lemma 3.3.4 along with Lemma 3.3.3 shows that almost surely for some positive constants a, a' and for all but finitely many r

$$a'\tau_{r+1} < X_{\tau_{r+1}} < \Delta(\tau_r) < aX_{\tau_r} < a\tau_r. \tag{3.3.5}$$

For the first and last inequality in the above display, we used Lemma 3.3.3, for the third inequality, we used Lemma 3.3.4 and for the second inequality we observe that the vertices of ∂B_{r+1} are added only one at a time. This in turn shows that

Lemma 3.3.5. *There exists constants* $1 < c < C$ *such that almost surely*

$$\begin{aligned} \liminf c^{-r}\tau_r &= \infty \\ \limsup C^{-r}\tau_r &< \infty \end{aligned}$$

Let V_n denote the number of vertices in the revealed map P_n in the nth step of the peeling algorithm. Our main goal is to estimate V_{τ_r} in order to prove Theorem 3.1.1. Now suppose $S_n = V_n - X_n$. Then it is easy to see just from the description of the algorithm that S_n is a sum of n i.i.d. random variables each of which is distributed as $Y + I_{Y+1}$ where $Y = -\Delta\tilde{X}_n \mathbb{1}_{\Delta\tilde{X}_n \neq 1}$ and I_{Y+1} is distributed as a the number of internal vertices of a free triangulation of a $(Y+1)$-gon with parameter $\alpha\beta$. Notice that this definition makes sense for all values of α, not for just the supercritical regime. However in the supercritical regime, the exponential tail of Y entails that Y has finite expectation. Further conditioned on Y, the expectation of I_{Y+1} is $4\theta Y/(1-6\theta) + O(1)$ via Proposition 3.2.1 where θ is given by the relation $\theta(1-2\theta)^2 = \alpha\beta$. Thus $Y + I_{Y+1}$ has finite expectation.

Proof of Theorem 3.1.1. Recall that $|B_r| = V_{\tau_r}$. Now since S_n is a sum of i.i.d. random variables with finite mean, S_n/n converges almost surely. This fact along with (3.3.5) and Lemma 3.3.5 completes the proof. $\square$

Anchored expansion

Now we turn to the proof of Theorem 3.1.2.

Recall that internal boundary of a simply connected sub-map with a simple boundary is the set of vertices and edges in the sub-map which is incident to at least one face which do not belong to the sub-map. Also recall that distance between two vertices on the boundary along the boundary is the number of edges on the boundary between the vertices. Clearly, distance along the boundary is at most the graph distance in the whole map. We show in the following lemma that the graph distance between vertices on the boundary in the whole map is at least linear in the distance between them along the boundary.

Lemma 3.3.6. *Let v be a vertex at distance $n \geq 1$ along the boundary from the root vertex on a map with law $\mathbb{H}_\alpha$ where $\alpha \in (2/3, 1)$. There exists a constant $t(\alpha) > 0$ depending only on α such that the probability of the distance between v and the root being smaller than $t(\alpha)n$ is at most* $\exp(-cn)$ *for some $c > 0$.*

Proof. Let us assume without loss of generality that v is to the right of the root vertex. We use the peeling algorithm described in Section 3.3.1 and reveal the hulls of radius r for $r \geq 1$ around the root vertex. Recall the notations P_n which denotes the revealed map after n peeling steps and τ_r which denotes the step in which we finish exploring the hull of radius r. Now

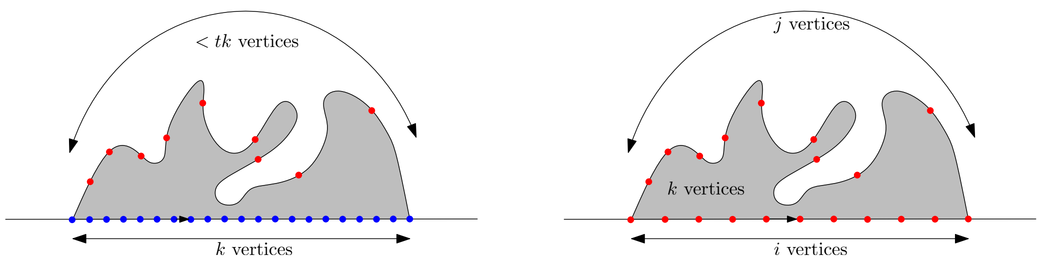

Figure 3.5: Left: An illustration of a t-bad segment. The segment consisting of blue vertices is a t-bad segment. The gray area is some fixed finite triangulation. Right: The red vertices form a $(k, i + j)$-separating loop

the vertices of the boundary to the right of the root vertex which goes inside the peeled map is entirely determined by the last step and is easily seen to have exponential tail and a finite expectation depending only on α. Hence the probability that v is in the hull of radius at most tn around the root vertex is at most the probability of the event that the sum of tk independent variables with finite expectation and exponential tail is larger than k. The latter event has probability $\exp(-ck)$ for some $c > 0$ if t is small enough (depending only on α) by a suitable large deviations estimate. □

A connected segment X on the boundary of the map containing the root edge is said to be a t-**bad segment** for some $t > 0$ if there exists a simply connected sub-map Q with a simple boundary whose intersection with the boundary of the map is X and the internal boundary has at most $t|X|$ vertices (see Figure 3.5).

Lemma 3.3.7. *For small enough t (depending on α), there exists finitely many t-bad segments almost surely.*

Proof. Let us fix a connected segment X of length k containing the root edge. The event that X is t-bad is contained in the event that the distance (in the whole map) between the leftmost and the rightmost vertices in X is at most tk. If $t > 0$ is small enough this event has probability at most $\exp(-ck)$ for some $c > 0$ using Lemma 3.3.6 and translation invariance. Since there are at most k connected segments of length k containing the root, the rest of the proof follows from Borel-Cantelli. □

We restate a Lemma which was proved in Chapter 2.

Lemma 3.3.8. *Fix $\alpha \in [0,1)$ and β is given by eqs. (2.8.2) and (2.8.3). Let Q be a simply connected triangulation with $k+i+j$ vertices and boundary size $i+j$. Let T be a sample from $\mathbb{H}_\alpha$. The event that Q is a sub-map of T with a marked connected segment with i vertices on the boundary of Q being mapped to a marked connected segment on the boundary of T with i vertices and no other vertex of Q being mapped to the boundary of T has probability*

$$\alpha^{k+j}\beta^{i+k-2}$$

We call a simple cycle in a half planar map a (k,l) **separating loop** if it has l vertices, its intersection with the boundary forms a connected segment containing the root edge and it separates k internal vertices of the map from infinity.

Lemma 3.3.9. *Fix $\alpha \in (2/3,1)$. There exists a constant $c(\alpha)$ depending upon α such that $\mathbb{H}_\alpha$-almost surely there are finitely many (k,l)-separating loops with $l < c(\alpha)k$.*

Proof. Recall that $\phi_{n,m}$ denotes the number of triangulations of an m-gon with n internal vertices. From Lemma 3.3.8 and union bound, the probability that there exists a (k,l)-separating loop with i vertices on the boundary of the map is at most

$$i\phi_{k,l}\alpha^{k+j}\left(\frac{\alpha(1-\alpha)}{2}\right)^{k+i-2} \tag{3.3.6}$$

where the factor i comes from the fact that the root can be any one of the edges of the intersection of the separating loop with the boundary. Let $j = l - i$. Now it is easy to see from eq. (2.12.1) that

$$\phi_{k,l} < (27/2)^k 9^l \left(1+\frac{2l}{3k}\right)^{2l+3k} k^{-5/2}\sqrt{l} \tag{3.3.7}$$

Combining (3.3.6) and (3.3.7) and summing over $i < l$ where $l < tk$, we get that the probability of existence of a (k,l)-separating loop is at most

$$l^{5/2}k^{-5/2}\cdot\left(\frac{27\alpha^2(1-\alpha)}{4}\right)^k\cdot\left(\frac{9\alpha(1-\alpha)}{2}\right)^l\cdot\left(1+\frac{2t}{3}\right)^{(2t+3)k}\cdot\left(\frac{2}{1-\alpha}\right)^{tk} < \exp(-ck) \tag{3.3.8}$$

for some constant $c > 0$ if t is small enough. To see this, observe that $\alpha^2(1-\alpha) < 4/27$ and $\alpha(1-\alpha) < 2/9$ if $\alpha \in (2/3,1)$. The sum of the bound in (3.3.8) over $k > l/t$ and then over l is finite. The rest of the proof follows from Borel-Cantelli. □

For $S \subset V(G)$, recall that the notation $|S|_E$ denotes the sum of the degrees of the vertices in S and $\partial_E S$ denotes the number of edges which are incident to one vertex in S and another in $G \setminus S$.

Proof of Theorem 3.1.2. Consider a connected set of vertices S containing the root vertex such that $|S|_E > n$ and suppose $\partial_E S < t^2|S|_E$ for some $t > 0$. By an abuse of notation, denote by S the finite map induced by S and without loss of generality assume it contains the root edge. Add to S all the faces which share at least one vertex with S along with the edges and vertices incident to it. Then add all the connected finite components of the complement and call the resulting finite triangulation $\overline{S}$. Note that $\overline{S}$ is simply connected with a simple boundary with $|\overline{S}| > n$ and the vertices and edges in the boundary of $\overline{S}$ form a separating loop. Suppose the internal boundary of $\overline{S}$ has j vertices. From the definition of $\partial_E S$: $j < \partial_E S < t^2|S|_E$. Let i be the number of vertices of S on the boundary of the map and suppose $k = |\overline{S}| - i - j$. Now the assumption $\partial_E S < t^2|S|_E$ and Euler's formula for $\overline{S}$ yields

$$k > \frac{1-5t^2}{6}|S|_E - \frac{2i}{3} \tag{3.3.9}$$

If $i < t|S|_E$ then $i+j < tCk$ for some universal constant $C > 0$ using (3.3.9) which can occur for finitely many n almost surely via Lemma 3.3.9 if t is small enough. If $i > t|S|_E$ then $j < ti$ and this can occur for finitely many n almost surely via Lemma 3.3.7. □

Proofs of Lemmas 3.3.7 and 3.3.9 and Theorem 3.1.2 in fact says that the probability of the existence of a set with small boundary containing the root vertex is exponentially small. We record it here for future reference.

Proposition 3.3.10. *There exists a $t > 0$ depending upon α such that the probability that there exists a connected set of vertices S containing the root vertex such that $|S|_E > n$ and $\partial_E S < t|S|_E$ is at most* $\exp(-cn)$ *for some $c > 0$.*

3.3.3 Subcritical

In this section we prove Theorem 3.1.6. We shall use the notations of Sections 3.3.1 and 3.3.2 and $\alpha \in [0, 2/3)$ throughout this section. Recall that in this regime, probability that a peeling step of the form (L, i) or (R, i) occurs for $i \geq k$ is roughly $k^{-1/2}$.

Boundary size estimates

To understand the boundary sizes, we need to understand the variables X_{τ_r} for $r \geq 1$. As a warm up we prove Proposition 3.1.4 first.

Proof of Proposition 3.1.4. From the definition, $X_n \leq n+2$ for all $n \in \mathbb{N}$. If at step numbers $2k$ and $2k+1$, events: $\cup_{j>2k+2}\{(R,j)\}$ and $\cup_{j>2k+2}\{(L,j)\}$ happen respectively then $X_{2k+2} = 2$. But this event has probability at least c/k for some $c > 0$ and for different k's these events are independent by the domain Markov property. The proof follows by Borel-Cantelli. $\square$

We know via Proposition 3.3.1 that X_{τ_r} is an irreducible aperiodic Markov chain with state space a $\mathbb{N} \setminus \{0,1\}$. We now show that $\{X_{\tau_r}\}_{r\geq 1}$ is a tight sequence with exponential tail. Suppose N_k for $k \geq 0$ denote the number of vertices in the internal boundary of P_{τ_r+k} which do not belong to ∂B_r.

Lemma 3.3.11. *For any $r \geq 0, k \geq 1, n \geq 1$*

$$\mathbb{P}(N_k > n | X_{\tau_r}) < \exp(-Bn)$$

for some positive constant B. In particular, this bound is independent of the conditioning.

Proof. First fix an n_0 large enough such that

$$\alpha - \frac{1}{2}\sum_{i=1}^{\lfloor n_0/2 \rfloor} i p_i < -\varepsilon$$

for some $\varepsilon > 0$ where p_i is given by eq. (2.8.2) (observe that such a choice of n_0 exists due to the heavy tail of p_i.) The above choice of n_0 depends only on α. Now choose an integer $n > n_0$. Let $\Delta N_k := N_{k+1} - N_k$ for $k \geq 1$. Observe that N_k increases by at most 1 in any step because of the evolution of X_k and $N_0 = 0$. This has several implications. Firstly, this implies that it is enough to consider $k > n$ or otherwise the requested probability is 0. Secondly, if $N_k > n$, then for some integer $1 \leq j \leq k$, N_j is equal to n_0. Let $M = \max\{1 \leq j \leq k : N_j = n_0\}$. Finally, we must have $M \leq k - n + n_0$. Now note that

$$\mathbb{P}(N_k > n, M = j) < \mathbb{P}(N_i \geq n_0 \text{ for all } j \leq i \leq k) \tag{3.3.10}$$

Now for any $i > j$, conditioned on $N_i \geq n_0$, there are at least $n_0/2$ vertices of the internal boundary of P_{τ_r+k} which do not belong to ∂B_r either to the left or right of the edge we perform the $(i+1)$th peeling step because of the way the exploration process evolves. Hence it is clear that conditioned on $N_i \geq n_0$, $\Delta(N_i)$ is dominated by a variable D with $\mathbb{E}(D) < -\varepsilon$ because of the choice of n_0. Thus,

$$\mathbb{P}(N_i \geq n_0 \text{ for all } j \leq i \leq k) < \mathbb{P}\left(\sum_{i=1}^{k-j} D_i > 0\right) < \gamma^{k-j} \tag{3.3.11}$$

for some $0 < \gamma < 1$ depending only on n_0 where $\{D_i\}_{i \geq 1}$ are i.i.d. copies of D and the last inequality of (3.3.11) follows from suitable large deviations estimate. Now using (3.3.10) and (3.3.11),

$$\mathbb{P}(N_k > n | X_{\tau_r}) = \sum_{j=1}^{k-n+n_0} \mathbb{P}(N_k > n, M = j) < \sum_{j=1}^{k-n+n_0} \gamma^{k-j} < \exp(-Bn) \tag{3.3.12}$$

for some $B > 0$ for large enough n. Decrease B suitably so that the requested bound is true even for smaller values of n. $\square$

We remarked before that Lemma 3.3.4 is true for any value of α. We shall now use this fact and induction to prove Theorem 3.1.6.

Proof of Theorem 3.1.6. First, get hold of the constants $A > 1, A' > 0, B > 0$ such that Lemma 3.3.4, part (i) and Lemma 3.3.11 are true for $n \geq 1$. Now fix a C such that $0 < C < B$. Then choose a large N to ensure that for all $n > N$,

$$\max\{\exp(-A'n^2), An^2\exp(-Bn), \exp(-Cn^2)\} < \frac{1}{3}\exp(-Cn).$$

We shall prove that for all $n > N$, the Theorem is true for the above choice of $C > 0$ by induction on r. Note that for $r = 0$, the Theorem is true trivially since $X_{\tau_0} = X_0 = 1$. Now assume, the Theorem is true for $r' = r - 1$ for any $n > N$ for above choice of C, N. Now recall the notation N_j from Lemma 3.3.11 and observe that $N_j = X_j$ for $j \geq \Delta\tau_r$. Clearly, for $n > N$, using Lemma 3.3.11

$$\mathbb{P}(X_{\tau_r} > n, \Delta\tau_{r-1} = j | X_{\tau_{r-1}}) < \mathbb{P}(N_j > n | X_{\tau_{r-1}}) < \exp(-Bn) \tag{3.3.13}$$

Now for any choice of $n > N$, using (3.3.13),

$$\begin{aligned}
\mathbb{P}(X_{\tau_r} > n) &< \mathbb{P}(X_{\tau_{r-1}} > n^2) + \mathbb{P}(\Delta\tau_{r-1} > An^2 | X_{\tau_{r-1}} \leq n^2) + \sum_{j=1}^{An^2} \exp(-Bn) \\
&< \exp(-Cn^2) + \exp(-A'n^2) + An^2\exp(-Bn) && (3.3.14) \\
&< \exp(-Cn) && (3.3.15)
\end{aligned}$$

where (3.3.14) follows from induction step, Lemmas 3.3.4 and 3.3.11. Also, (3.3.15) follows from the choice of N. The proof is completed by induction. □

Hull Volumes

First, we wish to estimate the growth rate of τ_r. Note that conditioned on X_{τ_r} the distribution of $\Delta(\tau_r)$ depends only on X_{τ_r} and not r. It is easy to see that $Z_r := (X_{\tau_r}, \Delta\tau_r)$ is an irreducible aperiodic Markov chain. Using Theorem 3.1.6 and Lemma 3.3.4 it is not difficult to see that the sequence $\{Z_r\}_{r\geq 1}$ forms a tight sequence. Hence, Z_r has a stationary probability distribution. Let us denote the marginal of the second coordinate of this stationary distribution by π. It is also easy to see using Theorem 3.1.6 and Lemma 3.3.4 that π has exponential tail and hence finite expectation. If we start the Markov chain $\{Z_r\}_{r\geq 1}$ from stationarity, ergodic theorem gives us that τ_r/r converges almost surely to $\sum_{i\geq 0} i\pi(i)$. However if we start the Markov chain $\{Z_r\}_{r\geq 1}$ from any fixed number, it is still absolutely continuous with respect to the corresponding chain starting from stationarity. This argument proves

Lemma 3.3.12. *Almost surely,*

$$\frac{\tau_r}{r} \to \sum_{i\geq 0} i\pi(i)$$

Recall the notation $Y, Z, \{S_n\}_{n\geq 1}$ from Section 3.3.2. Recall that the volume of the triangulation revealed at the n-th step of peeling is given by $V_n = S_n + X_n$. We wish to estimate $V_{\tau_r} = |B_r|$. Recall that S_n is a sum of n i.i.d. copies of W where $W = Y + I_{Y+1}$. From Lemma 3.2.2 part (i), we conclude $\mathbb{P}(W > x) \sim c_\alpha x^{-1/2}$ as $x \to \infty$ for some constant c_α depending on α.

Lemma 3.3.13. *For some sequence of real numbers a_n and b_n*

$$\frac{V_n - b_n}{a_n} \Rightarrow S \tag{3.3.16}$$

where S follows a stable distribution of type $1/2$. Also $a_n \sim cn^2$ and $b_n \sim c'n^2$ for some positive constants c, c'.

Proof. Note that since $X_n \leq n$ and since $V_n = S_n + X_n$, it is enough to prove the result with V_n replaced by S_n. Since S_n is a sum of i.i.d. sequence of variables, each of which is distributed as W, we apply Theorem 3.2.3. Recall from Theorem 3.2.3, the centering sequence $a_n = \inf\{t : \mathbb{P}(W > t) \leq 1/n\}$.

Recall that we also obtained the tail estimate of W, $\mathbb{P}(W > x) \sim c_\alpha x^{-1/2}$. It is easy to see from this tail estimate of W that $a_n \sim c_\alpha^2 n^2$. The asymptotics of b_n is provided in Lemma 3.2.2 part (ii). $\square$

We need one final lemma before we prove the theorem. Recall the distribution π from Lemma 3.3.12.

Lemma 3.3.14. *$V_{\tau_r}/V_{\lfloor \gamma r \rfloor}$ converges in probability to 1 where $\gamma = \sum_{i\geq 0} i\pi(i)$.*

Proof. Observe that it is enough to prove $S_{\tau_r}/S_{\lfloor \gamma r \rfloor}$ converges to 1 in probability. Notice that since S_r is nondecreasing in r, for any $\eta > 0$ and $\varepsilon > 0$, we have

$$\begin{aligned} &\mathbb{P}(|S_{\tau_r}/S_{\lfloor \gamma r \rfloor} - 1| > \eta, (1-\varepsilon)\gamma r < \tau_r < (1+\varepsilon)\gamma r) \\ &\quad < \mathbb{P}\left(\frac{S_{\lfloor (1+\varepsilon)\gamma r \rfloor} - S_{\lfloor (1-\varepsilon)\gamma r \rfloor}}{S_{\lfloor \gamma r \rfloor}} > \eta, (1-\varepsilon)\gamma r < \tau_r < (1+\varepsilon)\gamma r\right) \quad (3.3.17) \end{aligned}$$

Recall that a stable law is absolutely continuous (see [49], Chapter VI.1, Lemma 1) and hence via Lemma 3.3.13 we can conclude both $\{S_r/r^2\}_{r\geq 1}$ and $\{r^2/S_r\}_{r\geq 1}$ form a tight sequence in r. Further notice that $\tau_r/\gamma r \to 1$ almost surely via Lemma 3.3.12. Combining all these pieces, it is easy to see that for any $\eta > 0$, there exists an $\varepsilon > 0$ such that the right hand side of eq. (3.3.17) can be made smaller than any prescribed $\delta > 0$ for large enough r. The details are left to the reader. $\square$

Proof of Theorem 3.1.7. Notice $V_{\tau_r} = |B_r|$. Also observe

$$\frac{V_{\tau_r} - b_{\lfloor \gamma r \rfloor}}{a_{\lfloor \gamma r \rfloor}} = \frac{V_{\lfloor \gamma r \rfloor} - b_{\lfloor \gamma r \rfloor}}{a_{\lfloor \gamma r \rfloor}} + \frac{V_{\lfloor \gamma r \rfloor}}{a_{\lfloor \gamma r \rfloor}}\left(\frac{V_{\tau_r}}{V_{\lfloor \gamma r \rfloor}} - 1\right) \quad (3.3.18)$$

The first term of eq. (3.3.18) converges to a stable random variable of type 1/2 via Lemma 3.3.13. The second term in eq. (3.3.18) converges to 0 in probability via Lemma 3.3.14. The proof follows combining these two facts. $\square$

3.4 Percolation

In this Section, we prove Theorems 3.1.8–3.1.10. We will use the peeling procedure and use the notations P_n, T_n introduced in Section 3.2.3. Along with revealing the face on the edge we peel, we might also reveal the color of the new vertex (if any) revealed. It will be useful to consider several boundary conditions, which specifies the colors of the boundary vertices. If we consider a percolation configuration on the whole graph including the boundary vertices, we say it is a random i.i.d. boundary condition.

Algorithm: We start with the root vertex black and every other vertex on the boundary white. At the $n+1$th step, we perform a peeling step at the edge on the boundary of T_n with a black vertex to the right and a white vertex to the left. We stop at the nth step if there is no black vertex left on the boundary of T_n.

A simple topological argument shows that the event that the above algorithm stops is the same as the event that the black cluster containing the root vertex is finite. Now consider the following variable B. If the peeling step is an α-step and a black vertex is revealed set $B = 1$. If the peeling step is of the form (R, i), set $B = -i$. Otherwise set $B = 0$. The following Lemma is a simple computation which essentially follows from Lemma 3.3.2.

Lemma 3.4.1.

$$\mathbb{E}_p(B) = \alpha p - \frac{1}{2}(\alpha - \sqrt{\alpha}\sqrt{3\alpha - 2}) \tag{3.4.1}$$

In particular, $\mathbb{E}_p(B) > 0$ *if and only if* $p > 1/2(1 - \sqrt{3 - 2/\alpha})$.

The following proof is essentially follows the idea of [10]. We add it for completeness. Recall the notation T_n from Section 3.2.3.

Proof of Theorem 3.1.8(for p_c). Assume the boundary condition: the root vertex is black and the rest of the vertices on the boundary are white. Start with $B_0 = 1$ and suppose B_k is the number of black vertices left in the boundary of T_k. Clearly $B_{k+1} - B_k$ are i.i.d. with the same distribution as B. Lemma 3.4.1 shows that B_k eventually goes to 0 almost surely if and only if $p \leq \frac{1}{2}(\alpha - \sqrt{\alpha}\sqrt{3\alpha - 2})$. Modifying the proof to a random i.i.d. boundary is an easy exercise of imitating Proposition 9 of [10] and is left to the reader. The almost sure existence of a black cluster if $p > \frac{1}{2}(\alpha - \sqrt{\alpha}\sqrt{3\alpha - 2})$ follows from ergodicity of the map with respect to translation of the root as proved in Proposition 2.2.2. □

Corollary 3.4.2. *With random i.i.d. boundary condition, $\mathbb{H}_\alpha$-almost surely,*

$$p_u \leq 1/2(1 + \sqrt{3 - 2/\alpha}) \tag{3.4.2}$$

Proof. Assume $p \geq 1/2(1+\sqrt{3-2/\alpha})$. Consider the event $\mathcal{E}$ that there are two infinite black clusters. Then one of the components of the complement of one of them must be infinite. Then the vertices in this component which connect to the infinite black cluster must be white. This means that there is also an infinite white cluster since the map is locally finite and one ended almost surely. Since white clusters are finite almost surely in the given regime of p (using Theorem 3.1.8 (i) and symmetry), $\mathcal{E}$ has probability 0. □

One can define an interface between the black and white clusters of a percolation configuration. More precisely, there is a well defined path in the dual configuration which separates the white and black clusters (see Figure 1.4). We are interested in the interfaces which are connected to the boundary. These are the interfaces which start on those edges on the boundary which have a black and a white vertex incident to it. Interfaces mark the boundary between a white and a black cluster on both its side. An interface might be finite or infinite. Finite interface separate finite clusters from infinity while infinite interfaces correspond to an infinite black cluster on one side and an infinite white cluster on the other. So in particular, if $p \in [0, p_c] \cup [p_u, 1]$, every interface is finite almost surely.

In the following exploration procedure The vertices whose colors have not been revealed yet will be called **free** vertices.

Algorithm 2: We start with the root vertex colored white, the vertex incident to the right of the root edge colored black and every other vertex on the boundary free. Now we start performing peeling on the root edge.

Suppose after n steps of peeling, the boundary of T_n consists of free vertices except for a finite contiguous white segment followed by a finite contiguous black segment to the right of the white segment. We now peel on the unique boundary edge of T_n connecting the black and white segments. If after a peeling step the third vertex of the face revealed is free, we reveal its color. If the triangle revealed *swallows* all the black vertices to the right (resp. white vertices to the left) and the revealed third free vertex is white (resp. black), then we reveal the colors of the vertices along the boundary to the right (resp. left) of revealed third vertex until we find a black (resp. white) vertex. Notice that after such a step, we are again left with a boundary which consists of free vertices except for a finite contiguous white segment followed by a finite contiguous black segment to its right. We can now continue this procedure.

Suppose we have a random i.i.d. boundary condition. Let $\mathcal{I}$ be the event that there is an infinite interface starting from the root edge.

Lemma 3.4.3. *If* $p \in (1/2(1-\sqrt{3-2/\alpha}), 1/2(1+\sqrt{3-2/\alpha}))$, *then*

$$\mathbb{P}_p(\mathcal{I}) > 0.$$

Proof. Suppose we are on the event that the root vertex is colored white, the vertex incident to the right of the root edge colored black. Now we perform algorithm 2. Let B_k be the size of the black connected segment and W_k be that of the white connected segment at the kth step of the algorithm. Recall the definition of the variable B defined in Lemma 3.4.1. Conditioned on B_k, B_{k+1} stochastically dominates a variable which has the same distribution as $(B_k + B)^+ + \mathbb{1}_{\{B_k+B\le 0\}}$ and $B_{k+1} - B_k$ are independent for every k. The domination comes from the fact that if a white segment is *swallowed*, we add a geometric p number of black vertices to B_k which we ignore in the prescribed expression. Now for p in the given range, $\mathbb{E}(B) > 0$, hence B_k forms a random walk with a positive drift. This implies $B_k \to \infty$ almost surely. Similarly by symmetry, $W_k \to \infty$ almost surely for p in the given range. All this implies the event $\{B_k > 1, W_k > 1 \text{ for all } k \ge 0\}$ has positive probability. But $B_k > 1$ and $W_k > 1$ for all $k \ge 0$ implies that the interface we started with is infinite. This completes the proof. □

Recall the notations W_k^∞, B_k^∞ the number of black and white infinite clusters respectively which has least one vertex on the boundary within distance k along the boundary from the root vertex.

Proof of Theorem 3.1.8 (for p_u) and Theorem 3.1.9. Fix a number p in the following range: $p \in (1/2(1-\sqrt{3-2/\alpha}), 1/2(1+\sqrt{3-2/\alpha}))$. Let E_k be the number of edges within distance k from the root edge along the boundary such that there is an infinite interface starting from that edge. Now note that the measure $\mathbb{P}_p$ is ergodic with respect to translation of the root. Hence Birkhoff's ergodic theorem implies that almost surely,

$$\frac{E_k}{k} \to \mathbb{P}_p(\mathcal{I}). \tag{3.4.3}$$

Note that $W_k^\infty + B_k^\infty = E_k + 1$ and also $|W_k^\infty - B_k^\infty| \le 1$. Hence,

$$W_k^\infty/k \to \rho \qquad \text{and} \qquad B_k^\infty/k \to \rho$$

where $\rho = \mathbb{P}_p(\mathcal{I})/2 > 0$ from Lemma 3.4.3. This proves Theorem 3.1.9 as well as shows that $p_u \ge 1/2(1+\sqrt{3-2/\alpha})$. □

Now we turn to the proof of Theorem 3.1.10. We will need the following technical Lemma which can be easily shown using optional stopping Theorem. For details, we refer the reader to [51] Corollary 9.4.1 and Exercise 9.13.

Lemma 3.4.4. *Let $X_1, X_2, \ldots$ be an i.i.d. sequence of random variables such that $E(X_1) > 0$ and $\mathbb{E}(\exp(\lambda X_1))$ exists for values of λ in a neighborhood around 0. Let $S_n = \sum_{i=1}^n X_i$. Then for any $k > 0$ there exists some constant $c > 0$ such that*

$$\mathbb{P}(\cup_{n\geq 1}\{S_n \leq -k\}) < \exp(-ck)$$

Lemma 3.4.5. *Fix $p \in (p_c, p_u)$. The $\mathbb{P}_p$ probability that the root vertex is contained in an infinite black cluster with one end or an infinite white cluster with one end is 0.*

Proof. Suppose without loss of generality the color of the root vertex is black and we shall prove that the probability that this vertex is contained in an infinite black cluster with one end is 0. Reveal vertices to the left and right of this vertex along the boundary until we find a white vertex on both sides. In the exploration we describe now, there will be a contiguous finite white segment followed by a contiguous finite black segment followed by a contiguous finite white segment on the boundary and the rest of the vertices on the boundary are free. We shall peel alternately at the two edges connecting the black and the white segments to the left and to the right. If at any step we swallow all the black vertices we stop. If we swallow all the white vertices to the left (resp. to the right), we reveal black vertices to the left (resp. to the right) along the boundary until we find a white vertex. Consider the sequence of maps T_n. Define the root edge of this map to be the same root edge as in the previous step if it has not been swallowed in that step. If it is swallowed, define the edge in the middle of the black segment in the boundary of T_n oriented from left to right to be the new root edge.

Lemma 3.4.6. *The root edge is swallowed finitely many times almost surely in the above described exploration.*

Proof. Notice that on the event we stop the exploration, the Lemma is true by definition. Let L_n (resp. R_n) be the distance between the root vertex and edge to the left (resp. right) on which we perform the nth peeling step in the above described exploration and let $B_n = L_n + R_n$ be the length of the black segment. Clearly, the sequence $\{\Delta B_n\}_{n\geq 1}$ is an i.i.d. sequence of

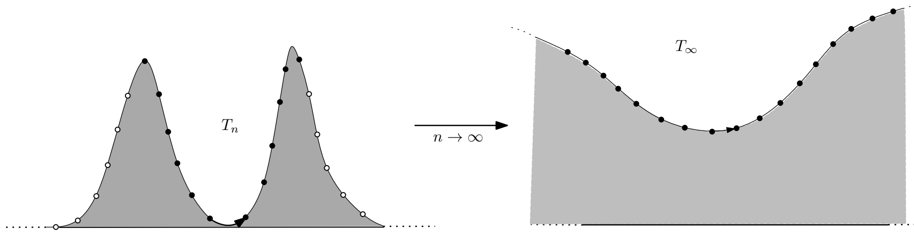

Figure 3.6: An illustration of the proof of Lemma 3.4.5. The gray area to the left denotes the revealed part in the exploration.

variables with each of which is distributed as B. Recall that B has positive expectation in the given regime of p (using Lemma 3.4.1). Hence using standard large deviation estimates, on the event that we do not stop the exploration, the probability of $B_n \leq tn$ for small enough t has probability at most $\exp(-cn)$ for some constant $c > 0$.

Now consider the event $\mathcal{E}_n$ that the root edge is swallowed in the nth step and is swallowed again in some step after the nth step. On the event $B_n > tn$ if the root edge is swallowed in the nth step, then by description of the exploration both L_n and R_n are at least $tn/2 - 1$. If the root edge is swallowed again, then either $\{L_k\}_{k \geq n}$ or $\{R_k\}_{k \geq n}$ has to reach 0 starting from at least $tn/2-1$. This event has probability at most $\exp(-c'n)$ for some $c' > 0$ via Lemma 3.4.4 since L_n as well as R_n has i.i.d. increments with positive expectation in every alternate step until the root edge is swallowed. Combining the pieces, we see that $\mathcal{E}_n$ has probability at most $\exp(-c''n)$ for some $c'' > 0$ which means $\mathcal{E}_n$ occurs for finitely many n by Borel-Cantelli lemma. This completes the proof. $\square$

Let T be a map with law $\mathbb{P}_p$ and we perform the above exploration. Let B_n be the number of black vertices on the boundary of T_n. On the event that $B_n \to \infty$, T_n converges almost surely to a sub map T_∞ of T since the root edge is swallowed finitely often almost surely via Lemma 3.4.6 (see Figure 3.6). However, on the event $B_n \to \infty$, from the domain Markov property, the distribution of the map T_n converges to a map with law $\mathbb{P}_p$ with all boundary vertices black and the rest of the vertices are not revealed. This map almost surely contains an infinite white cluster with the given range of p via Theorem 3.1.8. However this means that the cluster containing the root has at least two ends almost surely on the event that the cluster is infinite. □

Corollary 3.4.7. *Every infinite cluster does not contain an isolated end almost surely.*

Proof. We prove the corollary for an infinite cluster containing the root vertex. The proof for any infinite cluster is an easy exercise using the domain Markov property, and is left to the reader. Suppose with positive probability there is an infinite cluster in T containing the root vertex which has an isolated end. This implies that with positive probability there exists an r such that $T \setminus B_r(\alpha)$ has an infinite cluster incident to the boundary with one end. This is a contradiction because of Lemma 3.4.5 and domain Markov property. □

Proof of Theorem 3.1.10. Corollary 3.4.7 shows that each infinite cluster do not contain an isolated end. Since END is compact, the non-isolated points form a perfect subset via the Cantor Bendixson Theorem. Hence this implies that the set of ends has cardinality of the continuum. (see [67]). □

3.5 Open questions

We conclude with several open problems for possible future research. In Theorem 3.1.1, it is shown that the volume growth is exponential. A natural question is: what is the exact rate of growth of the volume? We expect similar behavior as exhibited by a supercritical Galton-Watson tree.

Question 3.5.1. *Suppose $\alpha \in (2/3, 1)$. Show that almost surely,*

$$\frac{\log |B_r(\alpha)|}{r} \to c$$

for some constant c depending only on α. Show further that $|B_r(\alpha)|/c^r$ converges to some non-degenerate random variable.

In Theorem 3.1.10 it is shown that the supercritical percolation clusters in the regime $p \in (p_c, p_u)$ have uncountably many ends. It would be interesting to know how a supercritical percolation cluster behaves.

Question 3.5.2. *Fix $\alpha \in (2/3, 1)$ and $p \in (p_c, p_u)$. Does the supercritical percolation cluster have exponential volume growth? Anchored expansion? Is the simple random walk on it transient?*

The key to understand the supercritical cluster in this regime is to understand if the supercritical clusters have long thin cutsets which kills anchored expansion.

Chapter 4

Unicellular maps: large scale structure

[3] This chapter is concerned with looking into the large scale structure of unicellular maps with genus proportional to the number of edges as the number of edges becomes large. In particular, we prove in Theorem 4.1.1 and Corollary 4.1.2 that the typical distances and diameter are roughly logarithmic in the number of edges. We also show that the map is locally planar and prove Theorem 4.1.4 quantifying this fact. The motivation for studying these maps is discussed in Section 1.6.

Let us first perform a small calculation. Suppose v is the number of vertices in a unicellular map of genus g with n edges. Then Euler's formula yields

$$v - n = 1 - 2g \tag{4.0.1}$$

Observe from eq. (4.0.1) that the genus of a unicellular map with n edges can be at most $n/2$. We are concerned in this chapter with unicellular maps whose genus grows like θn for some constant $0 < \theta < 1/2$. Specifically, we are interested in the geometry of a typical element among such maps as n becomes large.

4.1 Main results

Recall $\mathcal{U}_{g,n}$ denotes the set of unicellular maps of genus g with n edges and let $U_{g,n}$ denote a uniformly picked element from $\mathcal{U}_{g,n}$ for integers $g \geq 0$ and $n \geq 1$. For a graph G, let $d^G(.,.)$ denote its graph distance metric. Our first main result shows that the distance between two uniformly and independently picked vertices from $U_{g,n}$ is of logarithmic order if g grows like θn for some constant $0 < \theta < 1/2$.

[3] The results of this chapter are from [81].

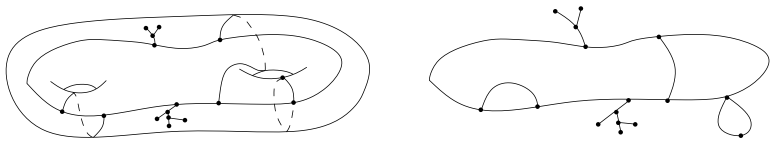

Figure 4.1: On the left: a unicellular map of genus 2. On the right: its underlying graph.

Theorem 4.1.1. *Let $\{g_l, n_l\}_l$ be a sequence in $\mathbb{N}^2$ such that $\{g_l, n_l\} \to \{\infty, \infty\}$ and $g_l/n_l \to \theta$ for some constant $0 < \theta < 1/2$. Suppose V_1 and V_2 are two uniformly and independently picked vertices from U_{g_l,n_l}. Then there exists constants $0 < \varepsilon < C$ (depending only on θ) such that*

(i) $\mathbb{P}(d^{U_{g_l,n_l}}(V_1, V_2) > \varepsilon \log n_l) \to 1$ *as* $l \to \infty$.

(ii) $\mathbb{P}(d^{U_{g_l,n_l}}(V_1, V_2) > C \log n_l) < c(n_l)^{-3}$ *for some* $c > 0$.

We remark here that in the course of the proof of part (i) of Theorem 4.1.1, a polynomial lower bound on the rate of convergence will be obtained. But since it is far from being sharp and is not much more enlightening, we exclude it from the statement of the Theorem. For part (ii) however, we do provide an upper bound on the rate. Notice that part (ii) enables us to immediately conclude that the diameter of U_{g_l,n_l} is also of order $\log n$ with high probability. For any finite map G, let $\operatorname{diam}(G)$ denote the diameter of its underlying graph.

Corollary 4.1.2. *Let $\{g_l, n_l\}_l$ be a sequence in $\mathbb{N}^2$ such that $\{g_l, n_l\} \to \{\infty, \infty\}$ and $g_l/n_l \to \theta$ for some constant $0 < \theta < 1/2$. Then there exists constants $\varepsilon > 0, C > 0$ such that*

$$\mathbb{P}(\varepsilon \log n < \mathit{diam}(U_{g_l,n_l}) < C \log n) \to 1$$

as $n \to \infty$.

Proof. The existence of $\varepsilon > 0$ such that $\mathbb{P}(\operatorname{diam}(U_{g_l,n_l}) > \varepsilon \log n) \to 1$ follows directly from Theorem 4.1.1 part (i). For the other direction, pick the same constant C as in Theorem 4.1.1. Let N be the number of pairs of vertices (v, w) in U_{g_l,n_l} where the distance between them is least $C \log n$. From part (ii) of Theorem 4.1.1, $\mathbb{E}(N) < c n_l^{-1}$ for some $c > 0$. Hence $\mathbb{E}N$ converges to 0 as $l \to \infty$. Consequently, $\mathbb{P}(N > 0)$ also converges to 0 which completes the proof. $\square$

If the genus is fixed to be 0, that is in the case of plane trees, the geometry is well understood (see [69] for a nice exposition on this topic.) In particular, it can be shown that the typical distance between two uniformly and independently picked vertices of a uniform random plane tree with n edges is of order $\sqrt{n}$. The diameter of such plane trees is also of order $\sqrt{n}$. These variables when properly rescaled, converge in distribution to appropriate functionals of the Brownian excursion. This characterization stems from the fact that a plane tree can be viewed as a metric space and the metric if rescaled by $\sqrt{n}$ (up to constants) converges in the Gromov-Hausdorff topology (see [58] for precise definitions) to the Brownian continuum random tree (see [4] for more on this.) The Benjamini-Schramm limit in the local topology (see [13, 22] for definitions), of the plane tree as the number of edges grow to infinity is also well understood: the limit is a tree with an infinite spine with critical Galton-Watson trees of geometric(1/2) offspring distribution attached on both sides (see Figure 1.2 and see [63] for details.)

Thus Theorem 4.1.1 depicts that the picture is starkly different if the genus of unicellular maps grow linearly in the number of vertices. The main idea behind the proof of Theorem 4.1.1 is that locally, $U_{g,n}$ behaves like a supercritical Galton-Watson tree, hence the logarithmic order. We believe that the quantity $d^{U_{g_l,n_l}}(V_1, V_2)$ of Theorem 4.1.1 when rescaled by $\log n$ should converge to a deterministic constant. Further, we also believe that the diameter of U_{g_l,n_l} when rescaled by $\log n$ should also converge to another deterministic constant. This constant obtained from the rescaled limit of the diameter should be different from the constant obtained as a rescaled limit of typical distances. The heuristic behind this extra length of the diameter is the existence of large "bushes" of order $\log n$ on the scheme of the unicellular map (recall the definition of scheme in Chapter 1 and Figure 1.6), a behavior reminiscent of Erdos-Renyi random graphs.

Tools developed for proving Theorem 4.1.1 also helps us conclude that locally $U_{g,n}$ is in fact planar with high probability which is our next main result. In fact, we are also able to quantify up to what distance from the root does $U_{g,n}$ remain planar. This will be made precise in the next theorem. A natural question at this point is what is the planar distributional limit of $U_{g,n}$ in the local topology. This is investigated in Chapter 5.

We now introduce the notion of **local injectivity radius** of a map. Recall that a circuit in a planar map is a subset of its vertices and edges whose image under the embedding is topologically a loop. A circuit is called contractible if its image under the embedding on the surface can be contracted to a point. A circuit is called non-contractible if it is not contractible.

Definition 4.1.3. *The local injectivity radius of a planar map with root vertex v^* is the largest r such that the sub-map formed by all the vertices within graph distance r from v^* does not contain any non-contractible circuit.*

In the world of Riemannian geometry, injectivity radius around a point p on a Riemannian manifold refers to the largest r such that the ball of radius r around p is diffeomorphic to an Euclidean ball via the exponential map. This notion is similar in spirit to what we are seeking in our situation. Notice however that a circuit in a unicellular map is always non-contractible because it has a single face. Hence looking for circuits and looking for non-contractible circuits are equivalent in our situation.

Theorem 4.1.4. *Let $\{g_l, n_l\} \to \{\infty, \infty\}$ and $g_l/n_l \to \theta$ for some constant $0 < \theta < 1/2$ as $l \to \infty$. Let I_{g_l,n_l} denote the local injectivity radius of U_{g_l,n_l}. Then there exists a constant $\varepsilon > 0$ such that*

$$\mathbb{P}\left(I_{g_l,n_l} > \varepsilon \log n_l\right) \to 1$$

as $l \to \infty$.

The girth of a map is defined as the minimum of the circuit sizes in it. The girth of $U_{g,n}$ also deserves some comment. It is possible to conclude via second moment method that the girth of U_{g_l,n_l} form a tight sequence. This shows that there are small circuits somewhere in the unicellular map, but they are far away from the root with high probability.

The main tool for the proofs is a bijection due to Chapuy, Feray and Fusy ([34]) which gives us a connection between unicellular maps and certain objects called C-decorated trees which preserve the underlying graph properties (details in Section 4.2.1.) This bijection provides us a clear roadway for analyzing the underlying graph of such maps.

From now on throughout this chapter, for simplicity, we shall drop the suffix l in $\{g_l, n_l\}$, and assume g as a function of n such that $g \to \infty$ as $n \to \infty$ and $g/n \to \theta$ where $0 < \theta < 1/2$.

Overview of the chapter: In Section 4.2 we gather some useful preliminary results we need. Proofs and references of some of the results in Section 4.2 are provided in appendices B and C. An overview of the strategy of the proofs of Theorems 4.1.1 and 4.1.4 is given in Section 4.3. Part (ii) of Theorem 4.1.1 along with Theorem 4.1.4 is proved in Section 4.4. Part (i) of Theorem 4.1.1 is proved in Section 4.5.

4.2 Preliminaries

In this section, we gather some useful results which we shall need.

4.2.1 The bijection

Chapuy, Féray and Fusy in ([34]) describes a bijection between unicellular maps and certain objects called C-decorated trees. The bijection describes a way to obtain the underlying graph of $U_{g,n}$ by simply gluing together vertices of a plane tree in an appropriate way. This description gives us a simple model to analyze because plane trees are well understood. In this section we describe the bijection in [34] and define an even simpler model called marked trees. The model of marked trees will contain all the information about the underlying graph of $U_{g,n}$.

For a graph G, let $V(G)$ denote the collection of vertices and $E(G)$ denote the collection of edges of G. The subgraph induced by a subset $V' \subseteq V(G)$ of vertices is a graph (V', E') where $E' \subseteq E(G)$ and for every edge $e \in E'$, both the vertices incident to e is in V'.

A permutation of order n is a bijective map σ from the set $\{1, 2, \ldots, n\}$ to itself. As is classically known, σ can be written as a composition of disjoint cycles. The length of a cycle is the number of elements in the cycle. The cycle type of a permutation is an unordered list of the lengths of the cycles in the cycle decomposition of the permutation. A **cycle-signed** permutation of order n is a permutation of order n where each cycle in its cycle decomposition carries a sign, either $+$ or $-$.

Definition 4.2.1 ([34])**.** *A C-**permutation** of order n is a cycle-signed permutation σ of order n such that each cycle of σ in its cycle decomposition has odd length. The genus of σ is defined to be $(n - N)/2$ where N is the number of cycles in the cycle decomposition of σ.*

Definition 4.2.2 ([34])**.** *A C-**decorated tree** on n edges is the pair (t, σ) where t is a rooted plane tree with n edges and σ is a C-permutation of order $n + 1$. The genus of (t, σ) is the genus of σ.*

The set of all C-decorated trees of genus g is denoted by $\mathcal{C}_{g,n}$. One can canonically order and number the vertices of t from 1 to $n+1$. Hence in a C-decorated tree (t, σ), the permutation σ can be seen as a permutation on the vertices of the tree t. To obtain the **underlying graph of a C-decorated tree** (t, σ), any pair of vertices x, y whose numbers are in the same cycle of σ are glued together (note that this might create loops and multiple edges.)

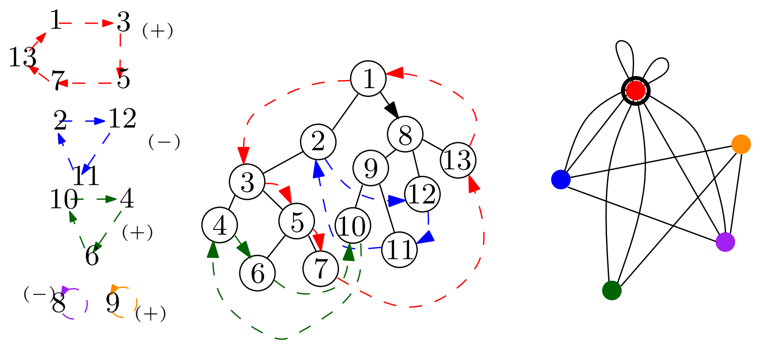

(a) (b) (c)

Figure 4.2: An illustration of a C-decorated tree. (a) A C-permutation σ where each cycle is marked with a different color. (b) A plane tree t with the vertices in the same cycle of σ joined by an arrow of the same color as the cycle. Note that vertices numbered 8 and 9 are fixed points in the C-permutation. (c) The underlying graph of the C-decorated tree (t, σ). The root vertex is circled.

The underlying graph of (t, σ) is the vertex rooted graph obtained from (t, σ) after this gluing procedure. So there are N vertices of the underlying graph of (t, σ), each correspond to a cycle of σ (see Figure 4.2). By Euler's formula, if the underlying graph of (t, σ) is embedded in a surface such that there is only one face, then the underlying surface must have genus g given by $N = n + 1 - 2g$.

For a set $\mathcal{A}$, let $k\mathcal{A}$ denote k distinct copies of $\mathcal{A}$. Recall that underlying graph of a unicellular map is the vertex rooted graph whose embedding is the map.

Theorem 4.2.3. *(Chapuy, Féray, Fusy [34]) There exists a bijection*

$$2^{n+1}\mathcal{U}_{g,n} \longleftrightarrow \mathcal{C}_{g,n}.$$

Moreover, the bijection preserves the underlying graph.

As promised, we shall now introduce a further simplified model which we call **marked tree** to analyze the underlying graph of C-decorated trees. Let $\mathcal{P}$ denote the set of ordered N-tuple of odd positive integers which add up to $n + 1$.

Definition 4.2.4. *A* ***marked tree*** *with n edges corresponding to an N-tuple $\lambda = (\lambda_1, \ldots, \lambda_N) \in \mathcal{P}$ is a pair (t, m) such that $t \in \mathcal{U}_{0,n}$ and $m : V(t) \to \mathbb{N}$*

is a function which takes the value i for exactly λ_i vertices of t for all $i = 1, \ldots, N$. The underlying graph of (t, m) is the rooted graph obtained when we merge together all the vertices of t with the same mark.

Given a $\lambda \in \mathcal{P}$, let $\mathcal{T}_\lambda(n)$ be the set of marked trees corresponding to λ and let $T_\lambda(n)$ be a uniformly picked element from it. Now pick λ from $\mathcal{P}$ according to the following distribution

$$\mathbb{P}(\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_N)) = \frac{\prod_{i=1}^N \lambda_i^{-1}}{Z} \tag{4.2.1}$$

where $Z = \sum_{\lambda \in \mathcal{P}} (\prod_{i=1}^N \lambda_i^{-1})$.

Proposition 4.2.5. *Choose λ according to the distribution given by* (4.2.1)*. Then the underlying graph of $U_{g,n}$ and $T_\lambda(n)$ has the same distribution.*

Proof. First observe that it is enough to show the following sequence of bijections

$$2^N \bigcup_{\lambda = (\lambda_1, \ldots, \lambda_N) \in \mathcal{P}} \prod_{i=1}^N (\lambda_i - 1)! \mathcal{T}_\lambda(n) \overset{\Psi}{\longleftrightarrow} N! \mathcal{C}_{g,n} \overset{\Phi}{\longleftrightarrow} 2^{n+1} N! \mathcal{U}_{g,n}$$

where Φ and Ψ are bijections which preserve the underlying graph. This is because for each $\lambda \in \mathcal{P}$, it is easy to see that the number of elements in $\prod_i (\lambda_i - 1)! \mathcal{T}_\lambda(n)$ is $(n+1)! \prod_{i=1}^N \lambda_i^{-1}$ and given a λ, the underlying graph of an uniform element of $\prod_i (\lambda_i - 1)! \mathcal{T}_\lambda(n)$ and $T_\lambda(n)$ has the same distribution.

Now the existence of bijection Φ which also preserves the underlying graph is guaranteed from Theorem 4.2.3. For Ψ, observe that the factor $\prod_{i=1}^N (\lambda_i - 1)!$ comes from the ordering of the elements within the cycle of C-permutations and the factor 2^N comes from the signs associated with each cycle of the C-permutations. The factor $N!$ comes from all possible ordering each cycle type of a C-permutation which is taken into account in the marked trees but not C-permutations. The proof is now complete by putting the pieces together. □

Because of Lemma 4.2.5 it is enough to look at the underlying graph of $T_\lambda(n)$ to prove the Theorems stated in Section 2.2 where λ is chosen according to the distribution given by (4.2.1). Our strategy is to show that a typical λ satisfies some "nice" conditions (which we will call condition (A) later), condition on such a λ satisfying those conditions and then work with $T_\lambda(n)$.

Recall $N = n + 1 - 2g$. Since $g/n \to \theta$ where $0 < \theta < 1/2$, $n/N \to (1-2\theta)^{-1}$. Denote $\alpha = (1-2\theta)^{-1}$. Clearly $\alpha > 1$. The reader should bear in mind that α will remain in the background throughout the rest of the chapter. We also warn the readers not to confuse this α with the α in Chapter 2.

4.2.2 Typical λ

Recall the definition of $\mathcal{P}$ from Section 4.2.1. Suppose C_0, C_1, C_2, d_1, d_2 are some positive constants which we will fix later. We say that an element in $\lambda = (\lambda_1, \lambda_2, \ldots \lambda_N) \in \mathcal{P}$ satisfies **condition** (A) if it satisfies

(i) $\lambda_{\max} < C_0 \log n$ where $\lambda_{\max}$ is the maximum in the set $\{\lambda_1, \lambda_2, \ldots \lambda_N\}$.

(ii) $C_1 n < \sum_{i=1}^N \lambda_i^2 < \sum_{i=1}^N \lambda_i^3 < C_2 n$.

(iii) $d_1 n < |i : \lambda_i = 1| < d_2 n$

The following Lemma ensures that λ satisfies condition (A) with high probability for appropriate choice of the constants. The proof is provided in appendix B

Lemma 4.2.6. *Suppose λ is chosen according to the distribution given by* (4.2.1)*. Then there exists constants C_0, C_1, C_2, d_1, d_2 depending only upon α such that condition (A) holds with probability at least $1 - cn^{-3}$ for some constant $c > 0$.*

Now we state a Lemma which will be useful later. Given a λ, we shall denote by $\mathbb{P}_\lambda$ the conditional measure induced by $T_\lambda(n)$.

Lemma 4.2.7. *Fix a tree $t \in \mathcal{U}_{0,n}$ and a $\lambda \in \mathcal{P}$ satisfying condition (A). Fix $\mathcal{I} \subset \{1, 2, \ldots, N\}$ such that $|\mathcal{I}| < n^{3/4}$. Condition on the event $\mathcal{E}$ that the plane tree of $T_\lambda(n)$ is t and S is the set of all the vertices in t whose mark belong to $\mathcal{I}$ where S is some fixed subset of $V(t)$ (S is chosen so that $\mathcal{E}$ has non-zero probability.) Let $\{v, w, z\} \subset V(t) \setminus S$ be any set of three distinct vertices in t and $i \notin \mathcal{I}$. Then*

$$\mathbb{P}_\lambda(m(v) = i|\mathcal{E}) \sim \lambda_i/n \tag{4.2.2}$$

$$\mathbb{P}_\lambda(m(v) = m(w)|\mathcal{E}) \asymp n^{-1} \tag{4.2.3}$$

$$\mathbb{P}_\lambda(m(v) = m(w) = m(z)|\mathcal{E}) \asymp n^{-2} \tag{4.2.4}$$

Proof. Notice that $|S| < C_0 n^{3/4} \log n$ because of part (i) of condition (A). The proof of (4.2.2) follows from the fact that

$$\mathbb{P}_\lambda(m(v) = i|\mathcal{E}) = \frac{(n - |S| - 1)!\lambda_i!}{(n - |S|)!(\lambda_i - 1)!} = \frac{\lambda_i}{n - |S|} \sim \lambda_i/n$$

since $|S| < C_0 n^{3/4} \log n$.

Now we move on to prove (4.2.3). Conditioned on S, t the probability that v and w have the same mark $j \notin \mathcal{I}$ with $\lambda_j \geq 3$ is

$$\frac{(n-|S|-2)!\lambda_j!}{(n-|S|)!(\lambda_j-2)!} \sim \frac{\lambda_j(\lambda_j-1)}{n^2}$$

All we need to prove is $\sum_{j \notin \mathcal{I}} \lambda_j(\lambda_j - 1) \asymp n$ which is clear from part (ii) of condition (A) and the fact that $|\mathcal{I}| < n^{3/4}$.

Proof of eq. (4.2.4) is very similar to that of eq. (4.2.3) and is left to the reader. □

4.2.3 Large deviation estimates on random trees

Galton-Watson trees

A Galton-Watson tree, roughly speaking, is the family tree of a Galton-Watson process which is also sometimes referred to as a branching process in the literature. These are well studied in the past and goes far back to the work of Harris ([60]). A fine comprehensive coverage about branching processes can be found in [14]. Given a Galton-Watson tree, we denote by ξ the offspring distribution. Let $\mathbb{P}(\xi = k) = p_k$ for $k \geq 1$. Let Z_r be the number of vertices at generation r of the tree. We shall also assume

- $p_0 + p_1 < 1$
- $\mathbb{E}(e^{\lambda\xi}) < \infty$ for small enough $\lambda > 0$.

We need the following lower deviation estimate. The proof essentially follows from a result in [14] and is provided in appendix C.

Lemma 4.2.8. *Suppose $\mathbb{E}\xi = \mu > 1$ and the distribution of ξ satisfies the assumptions as above. For any constant γ such that $1 < \gamma < \mu$, for all $r \geq 1$*

$$\mathbb{P}(Z_r \leq \gamma^r) < c\exp(-c'r) + \mathbb{P}(Z_r = 0)$$

for some positive constants c, c'.

Random plane trees

A random plane tree with n edges is a uniformly picked ordered tree with n edges (see [69] for a formal treatment.) In other words a random plane tree with n edges is nothing but $U_{0,n}$ as per our notation. We shall need the following large deviation result for the lower bounds and upper bounds on the diameter of $U_{0,n}$. This follows from Theorem 1.2 of [3] and the discussion in Section 1.1 of [3].

Lemma 4.2.9. *For any $x > 0$,*

(i) $P(Diam(U_{0,n}) \leq x) < c\exp(-c_1(n-2)/x^2)$

(ii) $P(Diam(U_{0,n}) > x) < c\exp(-c_1 x^2/n)$

where $c > 0$ and $c_1 > 0$ are constants.

We shall also need some estimate of local volume growth in random plane trees. For this purpose, let us define for an integer $r \geq 1$,

$$M_r = \max_{v \in V(U_{0,n})} |B_r(v)|$$

where $B_r(v)$ denotes the ball of radius r around v in the graph distance metric of $U_{0,n}$. In other words, M_r is the maximum over v of the volume of the ball of radius r around a vertex v in $U_{0,n}$. It is well known that typically, the ball of radius r in $U_{0,n}$ grows like r^2. The following Lemma states that M_r is not much larger than r^2 with high probability. Proof is provided in appendix C.

Lemma 4.2.10. *Fix $j \geq 1$ and $r = r(n)$ is a sequence of integers such that $1 \leq r(n) \leq n$. Then there exists a constant $c > 0$ such that*

$$\mathbb{P}(M_r > r^2 \log^2 n) < \exp(-c \log^2 n)$$

4.3 Proof outline

In this section we describe the heuristics of the proofs of Theorems 4.1.1 and 4.1.4.

Let us describe an exploration process on a given marked tree starting from any vertex v in the plane tree. This process will describe an increasing sequence of subsets of vertices which we will call the set of revealed vertices. In the first step, we reveal all the vertices with the same mark as v. Then we explore the set of revealed vertices one by one. At each step when we explore a vertex, we reveal all its neighbors and also reveal all the vertices which share a mark with one of the neighbors. If a neighbor has already been revealed, we ignore it. We then explore the unexplored vertices and continue.

We can associate a branching process with this exploration process where the number of vertices revealed while exploring a vertex can be thought of as the offsprings of the vertex. It is well known that the degree of any uniformly picked vertex in $U_{0,n}$ is roughly distributed as a geometric(1/2) variable and we can expect such behavior of the degree as long as the number of vertices revealed by the exploration is small compared to the size of the tree. Now the expected number of vertices with the same mark as a vertex is roughly a constant strictly larger than 1 because of part (ii) of condition (A). Hence the associated branching process will have expected number of offsprings a constant which is strictly larger than 1. Thus we can stochastically dominate this branching process both from above and below by supercritical Galton-Watson processes which will account for the logarithmic order of typical distances.

Once we have such a domination, observe that the vertices at distance at most r from the root in the underlying graph of the marked tree is approximately the vertices in the ball of radius r around the root in a supercritical Galton-Watson tree. Hence by virtue of the fact that supercritical Galton-Watson trees have roughly exponential growth, we can conclude that the number of vertices at a distance at most $\varepsilon \log n$ from the root in the underlying graph of the marked tree is $\ll \sqrt{n}$ if $\varepsilon > 0$ is small enough. Hence note that to have a circuit within distance $\varepsilon \log n$ in the underlying graph of the marked tree, two of the vertices which are revealed within $\ll \sqrt{n}$ many steps must be close in the plane tree. But observe that the distribution of the revealed vertices is roughly a uniform sample from the set of vertices in the tree up to the step when at most roughly $\sqrt{n}$ many vertices are revealed. It is not hard to see from this that the probability of revealing two vertices which are close in the plane tree up to roughly $\sqrt{n}$ many steps is small. This argument shows that the local injectivity radius is at least $\varepsilon \log n$ for some small enough $\varepsilon > 0$.

The rest of the chapter is devoted to make these heuristics precise.

4.4 Lower bound and injectivity radius

Recall condition (A) as described in the behavior of Section 4.2.2. Pick a λ satisfying condition (A). Recall that $T_\lambda(n)$ denotes a uniformly picked element from $\mathcal{T}_\lambda(n)$. Throughout this section we shall fix a λ satisfying condition (A) and work with $T_\lambda(n)$. Also recall that $T_\lambda(n) = (U_{0,n}, M)$ where $U_{0,n}$ is a uniformly picked plane tree with n edges and M is a uniformly picked marking function corresponding to λ which is independent of $U_{0,n}$.

Let $d_\lambda(.,.)$ denote the graph distance metric in the underlying graph of $T_\lambda(n)$. In this section we prove the following Theorem.

Theorem 4.4.1. *Fix a λ satisfying condition (A). Suppose x and y are two uniformly and independently picked numbers from $\{1, 2, \ldots, N\}$ and V_x and V_y are the vertices in the underlying graph of $T_\lambda(n)$ corresponding to the marks x and y respectively. Then there exists a constant $\varepsilon > 0$ such that*

$$\mathbb{P}_\lambda(d_\lambda(V_x, V_y) < \varepsilon \log n) \to 0$$

as $n \to \infty$.

Proof of Theorem 4.1.1 part (i). Follows from Theorem 4.4.1 along with Proposition 4.2.5 and Lemma 4.2.6. □

As a by-product of the proof of Theorem 4.4.1, we also obtain the proof of Theorem 4.1.4 in this section.

Note that for any finite graph, if the volume growth around a typical vertex is small, then the distance between two typical vertices is large. Thus to prove Theorem 4.4.1, we aim to prove an upper bound on volume growth around a typical vertex. Note that with high probability the maximum degree in $U_{0,n}$ is logarithmic and $\lambda_{\max}$ is also logarithmic (via condition (A) part (i) and Lemma 4.2.9.) Hence it is easy to see using the idea described in Section 4.3 that the typical distance is at least $\varepsilon \log n / \log\log n$ with high probability if $\varepsilon > 0$ is small enough. This is enough, as is heuristically explained in Section 4.3, to ensure that the injectivity radius of $U_g(n)$ is at least $\varepsilon \log n / \log\log n$ with high probability for small enough constant $\varepsilon > 0$. The rest of this section is devoted to the task of getting rid of the $\log\log n$ factor. This is done by ensuring that while performing the exploration process for reasonably small number of steps, we do not reveal vertices of high degree with high probability.

Given a marked tree (t, m), we shall define a nested sequence $R_0 \subseteq R_1 \subseteq R_2 \subseteq \ldots$ of subgraphs of (t, m) where R_k will be the called the **subgraph revealed** and the vertices in R_k will be called the **vertices revealed** at the kth step of the exploration process. We will also think of the number of steps as the amount of time the exploration process has evolved. There will be two states of the vertices of R_k: **active** and **neutral**. Along with $\{R_k\}$, we will define another nested sequence $E_0 \subseteq E_1 \subseteq E_2 \subseteq \ldots$. In the first step, $R_0 = E_0$ will be a set of vertices with the same mark and hence E_0 will correspond to a single vertex in the underlying graph of (t, m). The subgraph of the underlying graph of (t, m) formed by gluing together vertices with the same mark in E_r will be the ball of radius r around the vertex corresponding to E_0 in the underlying graph of (t, m). The process will have rounds and during round i, we shall reveal the vertices which correspond to vertices at distance exactly i from the vertex corresponding to E_0 in the underlying graph of (t, m). Define $\tau_0 = 0$ and we now define τ_r which will denote the time of completion of the rth round for $r \geq 1$. Let $N_r = E_r \setminus E_{r-1}$. Inductively, having defined N_r, we continue to explore every vertex in N_r in some predetermined order and τ_{r+1} is the step when we finish exploring N_r. For a vertex v, $mark(v)$ denotes the set of marked vertices with the same mark as that of v. For a vertex set S, $mark(S) = \cup_{v \in S} mark(v)$. We now give a rigorous algorithm for the exploration process.

Exploration process I

(i) **Starting rule:** Pick a number x uniformly at random from the set of marks $\{1, 2, \ldots, N\}$ and let $E_0 = R_0 = mark(x)$. Declare all the vertices in $mark(x)$ to be active. Also set $\tau_0 = 0$.

(ii) **Growth rule:**

1. For some $r \geq 1$, suppose we have defined the nested subset of vertices of $E_0 \subseteq \ldots \subseteq E_r$ such that $N_r := E_r \setminus E_{r-1}$ is the set of active vertices in E_r. Suppose we have defined the increasing sequence of times $\tau_0 \leq \ldots \leq \tau_r$ and the nested sequence of subgraphs $R_0 \subseteq R_1 \subseteq \ldots \subseteq R_{\tau_r}$ such that $R_{\tau_r} = E_r$. The number r denotes the number of **rounds** completed in the exploration process at time τ_r.
2. Order the vertices of N_r in some arbitrary order. Now we explore the first vertex v in the ordering of N_r. Let S_v denote all the neighbors of v in t which do not belong to R_{τ_r}. Suppose S_v has l vertices $\{v_1, v_2, \ldots, v_l\}$ which are ordered in an arbitrary way. For $1 \leq j \leq l$, at step $\tau_r + j$, define $R_{\tau_r + j}$ to be the subgraph induced by $V(R_{\tau_r + j - 1}) \cup mark(v_j)$. At step $\tau_r + l$ we finish exploring v. Define all the vertices in $R_{\tau_r + l} \setminus R_{\tau_r}$ to be active and declare v to be neutral. Then we move on to the next vertex in N_r and continue.
3. Suppose we have finished exploring a vertex of N_r in step k and obtained R_k. If there are no more vertices left in N_r, define $k = \tau_{r+1}$ and $E_{r+1} = R_{\tau_{r+1}}$. Declare round $r+1$ is completed and go to step 1.
4. Otherwise, we move on to the next vertex v' in N_r according to the pre-described order. Let $S_{v'} = \{v_1, v_2, \ldots, v_{l'}\}$ be the

neighbors of v' which do not belong to R_k. For $1 \leq j \leq l'$, at step $k+j$, define R_{k+j} to be the subgraph induced by $V(R_{k+j-1}) \cup mark(v_j)$. Define all the vertices in $R_{k+l'} \setminus R_k$ to be active and declare v' to be neutral. Now go back to step 3.

(iii) **Threshold rule:** We stop if the number of steps exceeds $n^{1/10}$ or the number of rounds exceeds $\log n$. Let δ be the step number when we stop the exploration process.

Recall that V_x denotes the vertex in the underlying graph of $T_\lambda(n)$ corresponding to the mark x. The following proposition is clear from the description of the exploration process and is left to the reader to verify.

Proposition 4.4.2. *For every $j \geq 1$, all the vertices with the same mark in $E_j \setminus E_{j-1}$ when glued together form all the vertices at a distance exactly j from V_x in the underlying graph of (t, m).*

In step 0, define $mark(x)$ to be the **seeds** revealed in step 0. At any step, if we reveal $mark(z)$ for some vertex z, then $mark(z) \setminus z$ is called the seeds revealed at that step. The nomenclature seed comes from the fact that a seed gives rise to a new connected component in the revealed subgraph unless it is a neighbor of one of the revealed subgraph components. However we shall see that the probability of the latter event is small and typically every connected component has one unique seed from which it "starts to grow".

Now suppose we perform the exploration process on $T_\lambda(n) = (U_{0,n}, M)$ where recall that M is a uniformly random marking function which is compatible with λ on the set of vertices of $U_{0,n}$ and is independent of the tree $U_{0,n}$. Let $\mathcal{F}_k$ be the sigma field generated by $R_0, R_1, R_2, \ldots, R_k$.

The aim is to control the growth of R_k and to that end, we need to control the size of $mark(S_v)$ while exploring the vertex v conditioned up to what we have revealed up to the previous step. It turns out that it will be more convenient to condition on a subtree which is closely related to the connected tree spanned by the vertices revealed.

Definition 4.4.3. *The **web** corresponding to R_k is defined to be the union of the unique paths joining the root vertex v^* and the vertices closest to v^* in each of the connected components of R_k including the vertices at which the paths intersect R_k. The web corresponding to R_k is denoted by P_{R_k}.*

As mentioned before, the idea is to condition on the web. Observe that after removing the web from $U_{0,n}$ at any step, we are left with a uniformly

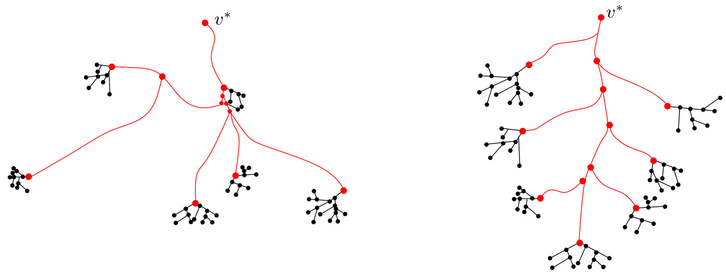

Figure 4.3: The web is denoted by the red paths. On the left: a general web structure. *A priori* the web structure might be very complicated. Many paths in the web might pass through the same vertex as is depicted here. On the right: A typical web structure

distributed forest with appropriate number of edges and trees. What stands in our way is that in general the web corresponding to a subtree might be very complicated (see Figure 4.3). The paths joining the root and several components might "go through" the same component. Hence conditioned on the web, a vertex might *a priori* have arbitrarily many of its neighbors belonging to the web. To show that this does not happen with high probability we need the following definitions.

For any vertex u in t, the **ancestors** of u are the vertices in t along the unique path joining u and the root vertex v^*. For any two vertices u, v in t let $u \wedge v$ denote the common ancestor of u and v which is farthest from the root vertex v^* in t. Let

$$C(u,v) = d^t(u \wedge v, \{u, v, v^*\})$$

A pair of vertices (u, v) is called a **bad pair** if $C(u,v) < \log^2 n$ (see Figure 4.4.)

Recall that we reveal some set of seeds (possibly empty) at each step of the exploration process. Suppose we uniformly order the seeds revealed at each step and then concatenate them in the order in which they are revealed. More formally, Let $(s_{i_0}, s_{i_1}, \ldots, s_{i_{k_i}})$ be the set of seeds revealed in step i ordered in uniform random order. Let $S = (s_{1_0}, s_{1_1}, \ldots, s_{1_{k_1}}, \ldots, s_{\delta_1}, \ldots, s_{\delta_{k_\delta}})$. To simplify notation, let us denote $S = (S_0, S_1, \ldots, S_{\delta'})$ where $\delta' + 1$ counts the number of seeds revealed up to step δ. The reason for such ordering is technical and will be clearer later in the proof of Lemma 4.4.8.

Lemma 4.4.4. *If S does not contain a bad pair then each connected component of R_δ contains an unique seed and the web P_{R_δ} intersects each connected component of R_δ at most at one vertex.*

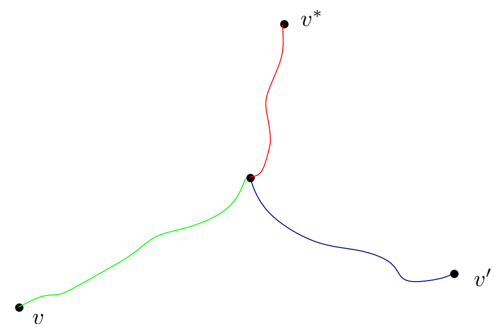

Figure 4.4: v^* denotes the root vertex. (v, v') is a bad pair if either of red, green or blue part has at most $\log^2 n$ many vertices.

Remark 4.4.5. *In Lemma 4.4.10, we shall prove that the probability of S containing a bad pair goes to 0 as $n \to \infty$. This and Lemma 4.4.4 shows that for large n, the typical structure of the web is like the right hand figure of Figure 4.3.*

Proof. Clearly, every connected component of R_δ must contain at least one seed. Also note that every connected component of R_δ has diameter at most $2 \log n$ because of the threshold rule. Since the distance between any pair of seeds in R_δ is at least $\log^2 n$ if S do not contain a bad pair, each component must contain a unique seed.

Suppose at any arbitrary step there is a connected component C which intersects the web in more than two vertices. Then there must exist a component C' such that the path of the web joining the root and C' intersects C in more than one vertex. This implies that the (unique) seeds of C and C' form a bad pair since the diameter of both C and C' are at most $2 \log n$. ☐

Now we want to prove that with high probability, S do not contain a bad pair. Observe that the distribution of the set of seeds revealed is very close to a uniformly sampled set of vertices without replacement from the set of vertices of the tree as long as $R_\delta \ll \sqrt{n}$, because of the same effect as the birthday paradox. We quantify this statement and further show that an i.i.d. sample of size δ' from the set of vertices do not contain a bad pair with high probability.

We first show that the cardinality of the set R_δ cannot be too large with high probability.

Lemma 4.4.6. $R_\delta = O(n^{1/10} \log n)$

Proof. At each step at most $\lambda_{\max}$ many vertices are revealed and $\lambda_{\max} = O(\log n)$ via condition (A). $\square$

Given S, let $\widetilde{S} = \{\widetilde{S}_0, \widetilde{S}_1, \ldots, \widetilde{S}_{\delta'}\}$ be an i.i.d. sample of uniformly picked vertices from $U_{0,n}$. First we need the following technical lemma.

Lemma 4.4.7. *Suppose $a = a(n)$ and $b = b(n)$ are sequences of positive integers such that $(a+b)^2 = o(n)$. Then for large enough n,*

$$n^b \left| \frac{1}{b!}\binom{n-a}{b}^{-1} - \frac{1}{n^b} \right| < 4\left(\frac{(a+b)b}{n}\right)$$

Proof. Observe that

$$\begin{aligned}
\frac{1}{b!}\binom{n-a}{b}^{-1} &= \frac{1}{(n-a-b+1)\ldots(n-a)} \\
&= \frac{1}{n^b}\prod_{j=1}^{b}\left(1 + \frac{a+b-j}{n-(a+b-j)}\right) \qquad (4.4.1)\\
&< \frac{1}{n^b}\prod_{j=1}^{b}(1 + 2(a+b)/n)) \\
&< \frac{1}{n^b}\exp\left(\frac{2(a+b)b}{n}\right) \\
&= \frac{1}{n^b}\left(1 + \frac{2(a+b)b}{n} + o\left(\frac{(a+b)b}{n}\right)\right) \\
&< \frac{1}{n^b}\left(1 + 4\left(\frac{(a+b)b}{n}\right)\right)
\end{aligned}$$

where the third inequality follows because $n-(a+b) > n/2$ for large enough n and $a+b-j < a+b$. The second last equality follows since $b(a+b) = o(n)$ via the hypothesis. The other direction follows from the fact that the expression in the right hand side of eq. (4.4.1) is larger than $1/n^b$. $\square$

For random vectors X, Y let $d_{TV}(X,Y)$ denote the total variation distance between the measures induced by X and Y.

Lemma 4.4.8.

$$d_{TV}(S, \widetilde{S}) < 4n^{-2/3}$$

Proof. First note that $|S| < |R_\delta| < n^{1/9}$ from Lemma 4.4.6. Let us denote by $(S_1, S_2, \dots S_d)$ the ordered set of seeds revealed in the first step after uniform ordering. Then

$$d_{TV}((S_1, \dots, S_d), (\widetilde{S}_1, \dots, \widetilde{S}_d)) < n^d \left| \frac{1}{d!} \binom{n}{d}^{-1} - \frac{1}{n^d} \right| < 4n^{-7/9} \qquad (4.4.2)$$

where the factor n^d in the first inequality of (4.4.2) comes from the definition of total variation distance and the fact that there are n^d many d-tuple of vertices and the second inequality of (4.4.2) follows from Lemma 4.4.7 and the fact that $d < |S| < n^{1/9}$. We will now proceed by induction on the number of steps. Suppose up to step t, $(S_1, \dots, S_m)$ is the ordered set of seeds revealed. Assume

$$d_{TV}((S_1, \dots, S_m), (\widetilde{S}_1, \dots, \widetilde{S}_m)) < 4mn^{-7/9} \qquad (4.4.3)$$

Recall $\mathcal{F}_t = \sigma(R_0, \dots, R_t)$. Now suppose we reveal $S_{m+1}, \dots, S_{m+L}$ in the $t+1$th step where L is random depending upon the number of seeds revealed in the $t+1$th step. Observe that to finish the proof of the lemma, it is enough to prove that the total variation distance between the measure induced by $(S_{m+1}, \dots, S_{m+L})$ conditioned on $\mathcal{F}_t$ and $(\widetilde{S}_{m+1}, \dots, \widetilde{S}_{m+L})$ (call this distance Δ) is at most $4n^{-7/9}$. This is because using induction hypothesis and $\Delta < 4n^{-7/9}$, we have the following inequality

$$\begin{aligned} &d_{TV}((S_1, \dots, S_{m+L}), (\widetilde{S}_1, \dots, \widetilde{S}_{m+L})) \\ &\quad < d_{TV}((S_1, \dots, S_m), (\widetilde{S}_1, \dots, \widetilde{S}_m)) + 4n^{-7/9} < 4(m+1)n^{-7/9}. \qquad (4.4.4) \end{aligned}$$

Thus (4.4.4) along with induction implies $d_{TV}(S, \widetilde{S}) < 4n^{1/9}n^{-7/9} < 4n^{-2/3}$ since $\delta' < n^{1/9}$.

Let $\mathcal{F}'_t$ be the sigma field induced by $\mathcal{F}_t$ and the mark revealed in step $t+1$. To prove $\Delta < 4n^{-7/9}$, note that it is enough to prove that the total variation distance between the measure induced by $S_{m+1}, \dots, S_{m+L}$ conditioned on $\mathcal{F}'_t$ and $(\widetilde{S}_{m+1}, \dots, \widetilde{S}_{m+L})$ (call it Δ') is at most $4n^{-7/9}$. But if l many seeds are revealed in step $t+1$ (note l only depends on the mark revealed) then a calculation similar to (4.4.2) shows that

$$\Delta' < n^l \left| \frac{1}{l!} \binom{n - |R_t| - 1}{l}^{-1} - \frac{1}{n^l} \right| < 4n^{-7/9}$$

where the last inequality above again follows from Lemma 4.4.7. The proof is now complete. $\square$

We next show that the probability of obtaining a bad pair of vertices in the collection of vertices $\widetilde{S}$ is small.

Lemma 4.4.9.

$$\mathbb{P}_\lambda(\widetilde{S} \text{ contains a bad pair }) = O(n^{-1/10})$$

Proof. Let (V, W) denote a pair of vertices uniformly and independently picked from the set of vertices of $U_{0,n}$. Let P be the path joining the root vertex and V. Let A be the event that the unique path joining W and P intersects P at a vertex which is within distance $\log^2 n$ from the root vertex or V. Since V and W have the same distribution and since there are at most $n^{2/9}$ pairs of vertices in $\widetilde{S}$, it is enough to prove $\mathbb{P}_\lambda(A) = O(n^{-1/3}\log^2 n)$

Recall the notation M_r of Lemma 4.2.10: M_r is the maximum over all vertices v in $U_{0,n}$ of the volume of the ball of radius r around v. Let $|P|$ denote the number of vertices in P. Consider the event $E = \{M_{\lfloor n^{1/3}\rfloor} < n^{2/3}\log^2 n\}$. On E, the probability of $\{|P| < n^{1/3}\}$ is $O(n^{-1/3}\log^2 n)$. Since the probability of the complement of E is $O(\exp(-c\log^2 n))$ for some constant $c > 0$ because of Lemma 4.2.10, it is enough to prove the bound for the probability of A on $|P| > n^{1/3}$.

Condition on P to have k edges where $k > n^{1/3}$. Observe that the distribution of $U_{0,n} \setminus P$ is given by an uniformly picked of rooted forests with $\sigma = 2k + 1$ trees and $n - k$ edges. Hence if we pick another uniformly distributed vertex W independent of everything else, the unique path joining W and P intersects P at each vertex with equal probability. Hence the probability that the unique path joining W and P intersects P at a vertex which is at a distance within $\log^2 n$ from the root or V is $O(n^{-1/3}\log^2 n)$ by union bound. This completes the proof. □

Lemma 4.4.10.

$$\mathbb{P}_\lambda(S \text{ contains a bad pair}) = O(n^{-1/10})$$

Proof. Using Lemmas 4.4.8 and 4.4.9, the proof follows. □

We will now exploit the special structure of the web on the event that S do not contain a bad pair to dominate the degree of the explored vertex by a suitable random variable of finite expectation for all large n. To this end, we need some enumeration results for forests. Note that the forests we consider here are rooted and ordered. Let $\Phi_{\sigma,e}$ denote the number of forests with σ trees and e edges. It is well known (see for example, Lemma 3 in [27]) that

$$\Phi_{\sigma,e} = \frac{\sigma}{2e+\sigma}\binom{2e+\sigma}{e} \tag{4.4.5}$$

We shall need the following estimate. The proof is postponed for later.

Lemma 4.4.11. *Suppose e is a positive integer such that $e < n$. Suppose d_0, d_1 denote the degree of the roots of two trees of a uniformly picked forest with $n-e$ edges and σ trees. Let $j \leq n-e$. Then*

$$\max\{\mathbb{P}(d_0+d_1=j), \mathbb{P}(d_0=j)\} < \frac{4j(j+1)}{2^j}$$

We shall now show the degree of an explored vertex at any step of the exploration process can be dominated by a suitable variable of finite expectation which do not depend upon n or the step number. Recall that while exploring v we spend several steps of the exploration process which depends on the number of neighbors of v which have not been revealed before.

In the following Lemmas 4.4.12 and 4.4.14, we assume v_{k+1} is the vertex we *start exploring* in the $(k+1)$th step of the exploration process.

Lemma 4.4.12. *The distribution of the degree of v_{k+1} conditioned on R_k such that R_k do not contain a bad pair is stochastically dominated by a variable X where $\mathbb{E}X < \infty$ and the distribution of X do not depend on n or k.*

Proof. Consider the conditional distribution of the degree of v_{k+1} conditioned on R_k as well as P_{R_k}. Without loss of generality assume P_{R_k} do not contain n edges for then the Lemma is trivial. Note that P_{R_k} cannot intersect a connected component of R_k at more than one vertex because of Lemma 4.4.4. Suppose $e < n$ is the number of edges in the subgraph $P_{R_k} \cup R_k$. It is easy to see that the distribution of $U_{0,n} \setminus (P_{R_k} \cup R_k)$ is a uniformly picked element from the set of forests with σ trees and $n-e$ edges for some number σ. If v_{k+1} is not an isolated vertex in R_k (that is there is an edge in R_k incident to v_{k+1}), the degree of v is at most 2 plus the sum of the degrees of the root vertices of two trees in a uniformly distributed forest of σ trees and $n-e$ edges. If v_{k+1} is an isolated vertex, the degree of v_{k+1} is 1 plus the degree of the root of a tree in a uniform forest of σ trees and $n-e$ edges. Now we can use the bound obtained in Lemma 4.4.11 and observe that the bound do not depend on the conditioning of the web P_{R_k}. It is easy now to choose a suitable variable X. The remaining details are left to the reader. □

Now we stochastically dominate the number of seeds revealed at a step conditioned on the subgraph revealed up to the previous step by a variable Y with finite expectation which is independent of the step number or n.

Lemma 4.4.13. *Number of vertices added to R_{j-1} the jth step of the exploration process conditioned on R_{j-1} is stochastically dominated by a variable Y with $\mathbb{E}Y < C$ where C is a constant which do not depend upon j or n.*

Proof. Recall r_i denotes the cardinality of the set $\{j : \lambda_j = i\}$. Now note that because of the condition (A), we can choose $\vartheta > 1$ such that $\sum_{i\geq 3} \vartheta i r_i < d_3 n$ for some number $0 < d_3 < 1$. Since $|R_k| < n^{1/9}$, the probability that the number of vertices added to R_{j-1} in the jth step is i for $i \geq 3$ is at most $ir_i/(n - n^{1/9}) < \vartheta i r_i/n$ for large enough n using eq. (4.2.2). Now define Y as follows:

$$\mathbb{P}(Y = i) = \begin{cases} \vartheta \frac{ir_i}{n} & \text{if } i \geq 3 \\ 1 - \sum_{i\geq 3} \vartheta \frac{ir_i}{n} := p_2 & \text{if } i = 2 \end{cases}$$

Note further that

$$\mathbb{E}(Y) = 2p_2 + \vartheta \sum_{i\geq 3} i^2 r_i/n < 2p_2 + \sum_{i=1}^{N} \lambda_i^2/n < 2 + C_2$$

from condition (A). Thus clearly Y satisfies the conditions of the Lemma. $\square$

Again, recall the definition of v_{k+1} from Lemma 4.4.12. The following lemma is clear now.

Lemma 4.4.14. *Let X, Y be distributed as in Lemmas 4.4.12 and 4.4.13 and suppose they are mutually independent. Conditioned on R_k such that R_k do not have any bad pair, the number of vertices added to R_k when we finish exploring v_{k+1} is stochastically dominated by a variable Z where Z is the sum of X independent copies of the variable Y. Consequently $\mathbb{E}Z < C$ where C is a constant which do not depend upon k or n.*

Proof of Theorem 4.4.1. We perform exploration process I. Let r_δ be the maximum integer r such that $\tau_r < \delta$. Let $B_r^\lambda(V_x)$ denote the ball of radius r around the vertex V_x in the underlying graph of $T_\lambda(n)$. Recall that because of Proposition 4.4.2, $B_r^\lambda(V_x)$ is obtained by gluing together vertices with the same mark in $R_{\tau_r} = E_r$. Note that if $|B^\lambda_{\lfloor \varepsilon \log n \rfloor}(V_x)| \leq n^{1/9}$ then the probability that V_y lies in $B^\lambda_{\lfloor \varepsilon \log n \rfloor}(V_x)$ is $O(n^{-8/9} \log n)$ because of condition

(A) part (i). Hence it is enough to prove $\mathbb{P}_\lambda(r_\delta < \varepsilon \log n) \to 0$. Further, because of Lemma 4.4.10, it is enough to prove $\mathbb{P}_\lambda(r_\delta < \varepsilon \log n \cap \mathcal{B}) \to 0$ where $\mathcal{B}$ is the event that S do not contain a bad pair.

Consider a Galton-Watson tree with offspring distribution Z as specified in Lemma 4.4.14 and suppose Z_r is the number of offsprings in generation r for $r \geq 1$. Then from Lemma 4.4.14, we get

$$\mathbb{P}_\lambda(r_\delta < \varepsilon \log n \cap \mathcal{B}) < \mathbb{P}_\lambda\left(\sum_{k=1}^{\lfloor \varepsilon \log n \rfloor} Z_k > n^{1/9}\right) \to 0 \tag{4.4.6}$$

if $\varepsilon > 0$ is small enough which follows from the fact that $\mathbb{E}(Z_r) < C^r$ where C is the constant in Lemma 4.4.14 and Markov's inequality. □

Now we finish the proof of Theorem 4.1.4.

Proof of Theorem 4.1.4. We shall use the notations used in the proof of Theorem 4.4.1. Observe that if the ball of radius r_δ in the underlying graph of $T_\lambda(n)$ contains a circuit, then two connected components must coalesce to form a single component at some step $k < \delta$. However this means that there exists a bad pair. Thus on the event $\mathcal{B}$, the underlying graph of R_δ do not contain a circuit. Hence on the event $\mathcal{B}$, the ball of radius $\varepsilon \log n$ contains a circuit in the underlying graph of $T_\lambda(n)$ implies $r_\delta < \varepsilon \log n$. However from eq. (4.4.6), we see that the probability of $\{r_\delta < \varepsilon \log n \cap \mathcal{B}\} \to 0$ for small enough $\varepsilon > 0$. The rest of the proof follows easily from Lemmas 4.2.6 and 4.4.10 and Proposition 4.2.5. □

Now we finish off by providing the proof of Lemma 4.4.11.

Proof of Lemma 4.4.11. It is easy to see that

$$\mathbb{P}(d_0 = j) = \frac{\Phi_{\sigma+j-1,n-e-j}}{\Phi_{\sigma,n-e}}$$

where $\Phi_{\sigma,n}$ is given by eq. (4.4.5). A simple computation shows that

$$\begin{aligned}\frac{\Phi_{\sigma+j-1,n-e-j}}{\Phi_{\sigma,n-e}} &= \frac{\sigma+j-1}{\sigma}\frac{1}{2^j} \\ &\times \left(\frac{(n-e+\sigma)\prod_{i=1}^{j-1}(1-i/(n-e))}{(2(n-e)+\sigma-1)\prod_{i=2}^{j+1}(1+(\sigma-i)/2(n-e))}\right)\end{aligned} \tag{4.4.7}$$

Now we can assume $(n-e+\sigma)/(2(n-e)+\sigma-1) \leq 1$ (since $e \neq n$ by assumption). Also notice

$$1 - \frac{i}{n-e} < 1 + \frac{\sigma - i}{2(n-e)}$$

for $i \geq 1$. Hence eq. (4.4.7) yields

$$\begin{aligned} \frac{\Phi_{\sigma+j-1,n-e-j}}{\Phi_{\sigma,n-e}} &\leq \frac{\sigma+j-1}{\sigma}\frac{1}{2^j}\left(\frac{(1-1/(n-e))}{\prod_{i=j}^{j+1}(1+(\sigma-i)/2(n-e))}\right) \\ &\leq \frac{\sigma+j-1}{\sigma}\frac{4}{2^j} \leq \frac{4j}{2^j} \end{aligned} \tag{4.4.8}$$

which follows because $\prod_{i=j}^{j+1}(1+(\sigma-i)/2(n-e)) \geq 1/4$ since $n-e \geq j$ and for the second inequality of (4.4.8), we use the trivial bound $(\sigma+j-1)/\sigma \leq j$.

Further note that $\mathbb{P}(d_0 = k, d_1 = j-k)$ for any $0 \leq k \leq j$ is given by $\Phi_{\sigma+j-2,n-e-j}/\Phi_{\sigma,n-e}$. Hence summing over k,

$$\mathbb{P}(d_0 + d_1 = j) = (j+1)\frac{\Phi_{\sigma+j-2,n-e-j}}{\Phi_{\sigma,n-e}}.$$

Now keeping n fixed, $\Phi_{\sigma,n}$ is an increasing function of σ, hence using the bound obtained in (4.4.8), the proof is complete. □

4.5 Upper bound

Throughout this section, we again fix a λ satisfying condition (A) as described in Section 4.2.2. Recall $d_\lambda(.,.)$ denotes the graph distance metric in the underlying graph of $T_\lambda(n)$. In this section we prove the following Theorem.

Theorem 4.5.1. *Fix a λ satisfying condition (A). Suppose V_1 and V_2 be vertices corresponding to the marks* 1 *and* 2 *in $T_\lambda(n)$. Then there exists a constant $C > 0$ such that*

$$\mathbb{P}_\lambda(d_\lambda(V_1, V_2) > C \log n) = O(n^{-3})$$

Note that the distribution of $T_\lambda(n)$ is invariant under permutation of the marks. Hence the choice of marks 1 and 2 in Theorem 4.5.1 plays the same role as an arbitrary pair of marks.

Proof of Theorem 4.1.1 part (ii). Proof follows from Theorem 4.5.1, Proposition 4.2.5, and Lemma 4.2.6. □

To prove Theorem 4.5.1, we plan to use an exploration process similar to that in Section 4.4 albeit with certain modification to overcome technical hurdles. We start the exploration process from a vertex v_1 with mark 1 and continue to explore for roughly $n^{3/4}$ steps. Then we start from the vertex v_2 with mark 2 and explore for another $n^{3/4}$ steps. Since the sets of vertices revealed are approximately uniformly and randomly selected from the set of vertices of the tree, the distance between these sets of vertices should be small with high probability, because of the same reasoning as the birthday paradox problem. Then we show that the distance in the underlying graph of $T_\lambda(n)$ from the set of vertices revealed and 1 or 2 is roughly $\log n$ to complete the proof. To this end, we shall find a supercritical Galton-Watson tree whose offspring distribution will be dominated by the vertices revealed in every step of the process.

However, if we proceed as the exploration process described in Section 4.4, since an unexplored vertex has a reasonable chance of being a leaf, the corresponding Galton-Watson tree will also have a reasonable chance of dying out. However, we need the dominated tree to survive for a long time with high probability. To overcome this difficulty, we shall invoke the following trick. Condition on the tree $U_{0,n}$ to have diameter $\gg \log^2 n$. Consider the vertex v_* which is farthest from $\{v_1, v_2\}$. For each vertex we explore, we reveal its unique neighbor which lie on the path joining the vertex and v_* instead of revealing all the neighbors which do not lie in the set of revealed vertices. Note that the revealed vertices by the exploration process now will mostly be disjoint paths increasing towards v_* and we shall always have at least one child if the paths do not intersect. However the chance of paths intersecting is small. Since expected size of $mark(v)$ for any non-revealed vertex v is larger than 1 throughout the process, we have exponential growth accounting for the logarithmic distance. The rest of the Section is devoted to rigorously prove the above described heuristic.

We shall now give a brief description of the exploration process we shall use in this section which is a modified version of exploration process described in Section 4.4. Hence, we shall not write down details of the process again to avoid repetition, and concentrate on the differences with exploration process I as described in Section 4.4.

Conditioning on the tree: For the proof of Theorem 4.5.1, we only need randomness of the marking function M and not that of the tree $U_{0,n}$. Hence, throughout this section, we shall condition on a plane tree $U_{0,n} = t$ where $t \in \mathcal{U}_{0,n}$ such that

(i) $diam(t) > \sqrt{n}/\log n$.

(ii) $M_{\lfloor \log^3 n \rfloor} \leq \log^8 n$.

where recall that M_r is as defined in Lemma 4.2.10: maximum over all vertices v in $U_{0,n}$ of the volume of the ball of radius r around v. Let us call this condition, **condition** (B). Although apparently it should only help if the diameter of t is small, the present proof fails to work if the diameter is too small and requires a different argument which we do not need. Note that by Lemmas 4.2.9 and 4.2.10, the probability that $U_{0,n}$ satisfies condition (B) is at least $1 - \exp(-c \log^2 n)$ for some constant $c > 0$. Hence it is enough to prove Theorem 4.5.1 for the conditional measure which we shall also call $\mathbb{P}_\lambda$ by an abuse of notation.

We start with a marked tree (t, m) where t satisfies condition (B). As planned, the exploration process will proceed in two stages, in the first stage, we start exploring from a vertex with mark 1 and in the second stage from a vertex with mark 2.

Exploration process II, stage 1: There will be three states of vertices **active, neutral** or **dead**. We shall again define a nested sequence of subgraphs $R_0 \subseteq R_1 \subseteq R_2 \subseteq \ldots$ which will denote the subgraph revealed. Alongside $\{R_k\}_{k=0,1,\ldots}$, we will define another nested sequence $Q_0 \subseteq Q_1 \subseteq Q_2 \subseteq \ldots$ which will denote **dead vertices revealed**.

We shall similarly define the sequences $\{N_r\}$, $\{E_r\}$ and $\{\tau_r\}$ as in exploration process I. We call v to be a v_*-ancestor of another vertex v' if v lies on the unique path joining v' and v_*. The v_*-ancestor which is also the neighbor of v is called the v_*-parent of v.

Starting Rule: We start from a vertex v_1 with mark 1 and v_2 with mark 2 (if there are more than one, select arbitrarily.) Let v_* be a vertex farthest from $\{v_1, v_2\}$ in t (break ties arbitrarily.) Note that because of the lower bound on the diameter via condition (B), $d^t(v_1, v_*)$ and $d^t(v_2, v_*)$ are at least $\sqrt{n}(3\log n)^{-1}$. Declare v_1 to be active and let $R_0 = \{v_1\}$. Declare all the vertices in $mark(v_1) \setminus v_1$ to be dead and let $Q_0 = mark(v_1) \setminus v_1$. Set $\tau_0 = 0$ and $E_0 = R_0$.

Growth rule: Suppose we have defined $E_0 \subseteq \ldots \subseteq E_r$, $\tau_0 \leq \ldots \leq \tau_r$ and also $R_0 \subseteq \ldots \subseteq R_{\tau_r}$ such that $R_{\tau_r} = E_r$ and $N_r := E_r \setminus E_{r-1}$ is the set of active vertices in E_r. Now we explore vertices in N_r in some predetermined order and suppose we have determined R_k for some $k \geq \tau_r$. We now move on to the next vertex in N_r. If there is no such vertex, declare $k = \tau_{r+1}$ and $E_{r+1} = R_{\tau_{r+1}}$.

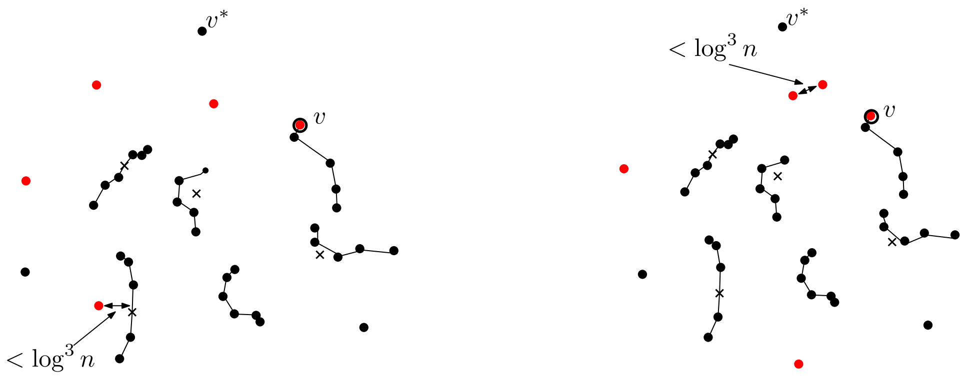

Figure 4.5: Illustration of the death rule in exploration process II. A snapshot of the revealed vertices when we are exploring the circled vertex v is given. The black vertices and edges correspond to neutral and active vertices, while the crosses correspond to dead vertices. We are exploring v and $mark(v)$ is denoted by the red vertices. On the left: a red vertex comes within distance $\log^3 n$ of one the revealed vertices, hence death rule is satisfied. On the right: two of the revealed vertices are within distance $\log^3 n$. Hence death rule is satisfied.

Otherwise suppose v is the vertex to be explored in the $k+1$th step. Let v_- denote the v_*-ancestor which is not dead and is nearest to v in the tree t. If v_- is already in R_k then we terminate the process.

Death rule: Otherwise, declare v_- to be active, v to be neutral and let $\Lambda = mark(v_-) \setminus v_-$. If any vertex $u \in \Lambda$ is within distance $\log^3 n$ from $R_k \cup Q_k \cup v_*$ or another $u' \in \Lambda$, we say *death rule is satisfied* (see Figure 4.5.) If death rule is satisfied declare all the vertices in Λ to be dead and set $Q_{k+1} = Q_k \cup \Lambda$, $R_{k+1} = R_k \cup v_-$. Otherwise declare all the vertices in Λ to be active, set $R_{k+1} = R_k \cup mark(v_-)$ and $Q_{k+1} = Q_k$.

Threshold rule: We stop if the number of steps exceed $n^{3/4}$ or r exceeds $\log^2 n$. Let δ denote the step when we stop stage 1 of the exploration procedure.

Exploration process II, stage 2: Similarly as in stage 1, we start with $R'_0 = v_2$ being active and $Q'_0 = mark(v_2) \setminus v_2$ being dead. We proceed exactly as in stage 1, except for the following change: if v_2 is a neighbor of $R_\delta \cup Q_\delta$, or while exploring v, if any of the vertices in $mark(v_-)$ is a neighbor of $R_\delta \cup Q_\delta$, we say a **collision** has occurred and terminate the procedure.

We shall see later (see Lemma 4.5.9 and Corollary 4.5.10) that with high probability, we perform the exploration for $n^{3/4}$ steps and the number of rounds is approximately $\log n$ in stage 1. Also in stage 2, collision occurs with high probability and the number of rounds is at most $\log n$ with high probability.

In what follows, we shall denote by X' in stage 2 the set or variable corresponding to that denoted by X in stage 1 (for example, R'_k, Q'_k will denote the set of revealed subgraphs and dead vertices respectively up to stage k in stage 2 etc.)

Now we shall define a new tree T_C which is defined on a subset of vertices of t. The tree T_C will capture the growth process associated with the exploration process. We start with the tree t and remove all its edges so we are left with only its vertices. The root vertex of T_C is v_1. For every step of exploring v, we add an edge neighboring v and every vertex of $mark(v_-)$ (where v_- is defined as in growth rule) in T_C if death rule is not satisfied. Otherwise we add an edge between v and v_- in T_C. The vertices we connect by an edge to v while exploring v is called the offsprings of v in T_C similar in spirit to a Galton-Watson tree. It is clear that T_C is a tree (since we terminate the procedure if $v_- \in R_k$). Let Z_r denotes the number of vertices at distance r from v_1 in T_C. Clearly, if we glue together vertices with the same mark which are at a distance at most r in T_C, we obtain a subgraph of the ball of radius r in the underlying graph of $T_\lambda(n)$. We similarly define another tree corresponding to stage 2 of the process which we call T'_C which starts from the root vertex v_2.

Lemma 4.5.2. *The volume of $R_\delta \cup Q_\delta$ is at most $C_0 n^{3/4} \log n$. Also the volume of $R_{\delta'} \cup Q_{\delta'}$ is at most $C_0 n^{3/4} \log n$ where C_0 is as in part (i) of condition (A).*

Proof. In every step, at most $C_0 \log n$ vertices are revealed by condition (A). □

Define v_1 and v_2 to be the **seeds** revealed in the first step. While exploring vertex v, we call the vertices in $mark(v_-) \setminus v_-$ to be the seeds revealed at that step if the death rule is not satisfied. Note that because of the prescription of the death rule, seeds are necessarily isolated vertices (not a neighbor of any other revealed neutral or dead vertex up to that step.)

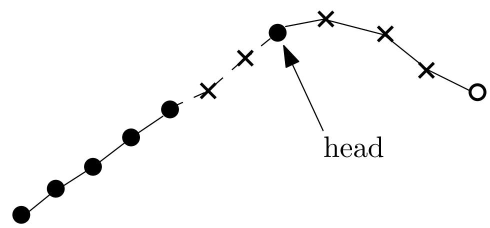

Figure 4.6: An illustration of a worm at a certain step in the exploration process. The black vertices denote the vertices of the worm, the crosses are the dead vertices and the head of the worm is as shown. The circle is the v_*-ancestor of the head which is not dead. This worm has faced death 5 times so far.

Definition 4.5.3. *A **worm** corresponding to a seed s denotes a sequence of vertices $\{w_0, w_1, w_2, \ldots, w_d\}$ such that w_0 is s and w_{i+1} is the vertex w_{i-} for $i \geq 0$.*

Note that in the above definition, w_{i-} depends on the vertices revealed up to the time we explore w_i in exploration process II. Note that in a worm, w_{i+1} is a neighbor of w_i if the v_*-parent of w_i is not a dead vertex. If it is a dead vertex we move on to the next nearest ancestor of w_i which is not dead. Note that the ancestors of w_i which lie on the path joining w_{i+1} and w_i are necessarily dead. If there are p dead vertices on the path between w_0 and the nearest v_*-ancestor of w_d which is not dead, we say that the worm has **faced death p times** so far (see Figure 4.6.) Call w_d the **head** of the worm. The **length** of the worm is the distance in $U_{0,n}$ between w_d and w_0.

We want the worms to remain disjoint so that conditioned up to the previous step, the number of children of a vertex in the tree T_C or T'_C remain independent of the conditioning. Now for any worm, if $w_i \in N_r$ (resp. $w_i \in N'_r$), then it is easy to see that $w_{i+1} \in N_{r+1}$ (resp. $w_{i+1} \in N'_{r+1}$) because of the way the exploration process evolves. Hence if none of the worms revealed during the exploration process face a dead vertex, then the length of each worm is at most $\log^2 n$ from threshold rule. Since every seed is at a distance at least $\log^3 n$ from any other seed via the death rule, the worms will remain disjoint from each other if death does not occur.

Unfortunately, many worms will face a dead vertex with reasonable chance. But fortunately, none of them will face many dead vertices with high probability. We say that a **disaster** has occurred at step k if after performing step k, there is a worm which has faced death at least 16 times. The following proposition is immediate from the threshold rule for exploration process II and the discussion above.

Proposition 4.5.4. *If disaster does not occur, then the length of each worm is at most* $\log^2 n + 16$ *and hence no two worms intersect during the exploration process for large enough* n.

We will now provide a series of Lemmas using which we will prove Theorem 4.5.1. The proofs of Lemmas 4.5.5 and 4.5.7 are postponed to Section 4.5.1 for clarity.

We start with a Lemma that shows that disaster does not happen with high probability and consequently the length of each worm is at most $\log^2 n + 16$ with high probability.

Lemma 4.5.5. *With* $\mathbb{P}_\lambda$*-probability at least* $1 - cn^{-3}$ *disaster does not occur where* $c > 0$ *is some constant.*

Lemma 4.5.6. *Suppose disaster does not occur and* $r \geq 0$*. Then in the tree* T_C*, for* $v \in Z_r$*,* $d^{T_C}(v, v_1) < 16r$*. Also in* T'_C*, for* $v \in Z'_r$*,* $d^{T'_C}(v, v_2) < 16r$*.*

Proof. We prove only for stage 1 as for stage 2 the proof is similar. The Lemma is trivially true for $r = 0$. Suppose now the Lemma is true for $r' = r$. Now for any vertex in Z_r and its offspring, the vertices corresponding to their marks in the underlying graph of $T_\lambda(n)$ must lie at a distance at most 16 because otherwise disaster would occur. Hence the distance of every vertex in Z_{r+1} from v_1 is at most $16r + 16 = 16(r + 1)$. We use induction to complete the proof. □

Let us denote by $\mathcal{E}$ (resp. $\mathcal{E}'$) the event that disaster has not occurred up to step δ (resp. δ'). Let $\mathcal{F}_k$ denote the sigma field generated by $R_0, \dots, R_k, Q_0, \dots, Q_k$. Let $\mathcal{F}'_k = \sigma(R'_0, \dots, R'_k, Q'_0, \dots, Q'_k) \vee \mathcal{F}_\delta$. Let v_{k+1} (resp. v'_{k+1}) be the vertex we explore in the $k + 1$th stage in the exploration process stage 1 (resp. stage 2).

Lemma 4.5.7. *Conditioned on* $\mathcal{F}_k$ *(resp.* $\mathcal{F}'_k$*) such that* $k < \delta$ *(resp.* δ'*) and disaster has not occurred up to step* k*, the* $\mathbb{P}_\lambda$*-probability that* v_{k+1} *in* T_C *(resp.* T'_C*) has*

(i) *no offsprings is* 0.

(ii) *at least* 3 *offsprings is at least* k_0 *for some constant* $k_0 > 0$.

We now aim to construct a supercritical Galton-Watson tree which will be stochastically dominated by both T_C and T'_C. Consider a Galton-Watson tree GW with offspring distribution ξ where

- $\mathbb{P}(\xi = 1) = (1 - k_0/2)$
- $\mathbb{P}(\xi = 2) = k_0/2$

and k_0 is the constant obtained in part (ii) of Lemma 4.5.7. Let Z_r^{GW} be the number of offsprings in the r-th generation of GW. Lemma 4.5.7 and the definition of GW clearly shows that Z_r stochastically dominates Z_r^{GW} for all $r \geq 1$ if disaster does not occur up to step τ_r. Let $r_\delta = \max\{r : \tau_r < \delta\}$ and similarly define $r_{\delta'} = \max\{r : \tau_r < \delta'\}$. Thus, we have

Lemma 4.5.8. *For any integer $j \leq r_\delta$ (resp. $j \leq r_{\delta'}$), Z_j stochastically dominates Z_j^{GW} on the event $\mathcal{E}$ (resp. $\mathcal{E}'$).*

It is clear that the mean offspring distribution of GW is strictly greater than 1 and hence GW is a supercritical Galton-Watson tree. Also GW is infinite with probability 1.

Now we are ready to show that the depth of T_C (resp. T'_C) when we run the exploration up to time δ (resp. δ') is of logarithmic order with high probability.

Lemma 4.5.9. *There exists a $C > 0$ such that*

(i) $\mathbb{P}_\lambda((r_\delta > C \log n) \cap \mathcal{E}) = O(n^{-3})$

(ii) $\mathbb{P}_\lambda((r_{\delta'} > C \log n) \cap \mathcal{E}') = O(n^{-3})$

Proof. We shall prove only (i) as proof of (ii) is similar. Because of Lemmas 4.5.2 and 4.5.8 we have for a large enough choice of $C > 0$,

$$\begin{aligned}\mathbb{P}_\lambda((r_\delta > C \log n) \cap \mathcal{E}) &< \mathbb{P}_\lambda\left(\sum_{i=1}^{\lfloor C \log n \rfloor} Z_i^{GW} \leq C_0 n^{3/4} \log n\right) \\ &< \mathbb{P}_\lambda\left(Z_{\lfloor C \log n \rfloor}^{GW} \leq C_0 n^{3/4} \log n\right) \\ &< n^{-3} \end{aligned} \tag{4.5.1}$$

where (4.5.1) follows by applying Lemma 4.2.8, choosing $C > 0$ large enough and observing the fact that $\mathbb{P}_\lambda(Z_r = 0) = 0$ for any r from definition of GW. □

Recall that in stage 2, we stop the process if we have revealed a vertex which is a neighbor of $R_\delta \cup Q_\delta$, and we say a **collision** has occurred. Let us denote the event that collision does not occur up to step k by $\mathcal{C}_k$. Since $\delta \leq \lfloor n^{3/4} \rfloor$ implies either disaster has occurred in stage 1 or $r_\delta > \log^2 n$ and $\delta' \leq \lfloor n^{3/4} \rfloor$ implies either disaster has occurred in stage 2 or $r'_\delta > \log^2 n$ or a collision has occurred we have the immediate corollary

Corollary 4.5.10. *On the event $\mathcal{E}$, the $\mathbb{P}_\lambda$-probability that $\delta \le \lfloor n^{3/4} \rfloor$ is $O(n^{-3})$. On the event $\mathcal{E}' \cap \mathcal{C}_{\delta'}$, the $\mathbb{P}_\lambda$-probability that $\delta' \le \lfloor n^{3/4} \rfloor$ is $O(n^{-3})$.*

Now we are ready to prove our estimate on the typical distances. We show next, that a collision will occur with high probability.

Lemma 4.5.11. *Probability that a collision occurs before step δ' is at least $1 - cn^{-3}$ for some constant $c > 0$.*

Proof. Let $\mathcal{Q}$ be the event that disaster does not occur up to step δ, $\delta = \lfloor n^{3/4} \rfloor + 1$. Let $A(R_\delta)$ be the set of v_*- parents of the heads of the worms in R_δ. Since at each step at least one vertex is revealed, $\delta = \lfloor n^{3/4} \rfloor + 1$ implies the number of vertices revealed is at least $n^{3/4}$. If disaster does not occur, then from Lemma 4.5.5 and Proposition 4.5.4, the worms are disjoint and each worm has length at most $\log^2 n + 16$. Hence the number of vertices in $A(R_\delta)$ is at least $n^{3/4}/(\log^2 n + 16)$. Also, the number of vertices in $A(R_\delta)$ is at most $C_0 n^{3/4} \log n$ from Lemma 4.5.2. For any $k < \delta'$, conditioned on the event $\mathcal{C}_k$ that no collision has occurred up to step k, the probability that collision occurs in step $k+1$ when we are exploring a vertex v is at least (using Bonferroni's inequality),

$$\sum_{w \in A(R_\delta)} \mathbb{P}_\lambda(m(v) = m(w) | \mathcal{C}_k, \mathcal{Q}) - \sum_{w,z \in A(R_\delta)} \mathbb{P}_\lambda(m(v) = m(w) = m(z) | \mathcal{C}_k, \mathcal{Q})$$

$$> \frac{c}{n^{1/4} \log^2 n} - \frac{c \log^2 n}{\sqrt{n}} \tag{4.5.2}$$

$$> \frac{c}{n^{1/4} \log^2 n} \tag{4.5.3}$$

for some constant $c > 0$. The first term of (4.5.2) follows from the lower bound of (4.2.3). The second term of (4.5.2) follows from (4.2.4) and noting that the number of terms in the sum is $O(n^{3/2} \log^2 n)$. Since the bound on the probability displayed in (4.5.3) is independent of the conditioning,

$$\begin{aligned}
&\mathbb{P}_\lambda(\mathcal{C}_{\delta'} \cap \delta' = \lfloor n^{3/4} \rfloor + 1 | \mathcal{Q}) + \mathbb{P}_\lambda(\mathcal{C}_{\delta'} \cap \delta' \le \lfloor n^{3/4} \rfloor | \mathcal{Q}) \\
< &\left(1 - \frac{c}{n^{1/4} \log^2 n}\right)^{n^{3/4}} + \mathbb{P}_\lambda(\mathcal{C}_{\delta'} \cap \mathcal{E}' \cap \delta' \le \lfloor n^{3/4} \rfloor | \mathcal{Q}) + \mathbb{P}_\lambda((\mathcal{E}')^c | \mathcal{Q}) \\
< &\exp(-c\sqrt{n}/\log^2 n) + O(n^{-3}) \qquad (4.5.4) \\
= &O(n^{-3})
\end{aligned}$$

where the bound on the second term in eq. (4.5.4) follows from Corollary 4.5.10 and Lemma 4.5.5. The Lemma now follows because the probability of the complement of $\mathcal{Q}$ is $O(n^{-3})$ again from Corollary 4.5.10 and Lemma 4.5.5. $\square$

Proof of Theorem 4.5.1. Suppose we have performed exploration process I stage 1 and 2. Let $\mathcal{G}$ be the event that $r_\delta \le C \log n, r_{\delta'} \le C \log n$, disaster does not occur before step δ or δ' and a collision occurs. On the event $\mathcal{G}$ the distance between V_1 and V_2 in the underlying graph of $T_\lambda(n)$ is at most $32C \log n + 1$ by Lemma 4.5.6. But by Lemmas 4.5.5, 4.5.9 and 4.5.11, the complement of the event $\mathcal{G}$ has probability $O(n^{-3})$. $\square$

4.5.1 Remaining proofs

The proofs of both the Lemmas in this subsection are for stage 1 of the exploration process as the proof for stage 2 is the same.

Proof of Lemma 4.5.5. Let s be a seed revealed in the kth step of exploration process II . Suppose P denotes the set of vertices at a distance at most $\log^2 n + 16$ from s along the unique path joining s and v_*. Note that none of the vertices in P are revealed yet because of the death rule. If the worm corresponding to s faces more than 16 dead vertices, then more than 16 dead vertices must be revealed in P during the exploration from step k to δ. Conditioned up to the previous step, the probability that one of the revealed vertices lie in P in a step is $O(n^{-1}(\log^2 n + 16))$ from (4.2.3) and union bound. Since this bound is independent of the conditioning, the probability that this event happens at least 16 times during the process is $O(n^{-16} \log^{32} n \cdot n^{12}) = O(n^{-4} \log^{32} n)$ where the factor n^{12} has the justification that $\binom{\lfloor n^{3/4} \rfloor}{16} = O(n^{12})$ is the number of combination of steps by which this event can happen 16 times. Observe that more than one vertex may be revealed in P in a step, but the probability of that event is even smaller. Thus taking union over all seeds, we see that the probability of disaster occurring is $O(n^{-13/4} \log^{33} n) = O(n^{-3})$ using Lemma 4.5.2 and union bound. $\square$

Proof of Lemma 4.5.7. It is clear that on the event of no disaster, every explored vertex has at least the offspring corresponding to its closest non-dead v_*-ancestor in the tree T_C. This is because on the event of no disaster, no two worms intersect. For stage 2, the closest non-dead v_*-ancestor cannot belong to $R_\delta \cup Q_\delta$ because otherwise, the process would have stopped. Hence (i) is trivial.

Now for (ii), first recall that condition (A) ensures that the number of indices i such that $\lambda_i \geq 3$ is at least $(1-d_2)n$. Now since the number of vertices revealed up to any step $k < \delta$ is $O(n^{3/4} \log n)$, the number of vertices left with mark i such that $\lambda_i \geq 3$ is at least $(1-d_2)n - O(n^{3/4} \log n) > (1-d_2)n/2$ for large enough n. Note that the number of offsprings of v in T_C is at least 3 if the number of vertices with the same mark as v_- is at least 3 and death does not occur. Hence if we can show that the probability of death rule being satisfied in a step is $o(1)$, we are done.

To satisfy the death rule in step $k+1$, a vertex in $mark(v_-) \setminus v_-$ must be within distance $\log^3 n$ in the tree $U_{0,n}$ to another vertex in $mark(v_-) \setminus v_-$ or $R_k \cup Q_k \cup v^*$. Now using of part (ii) of condition (B) and Lemma 4.5.2, the number of vertices within $\log^3 n$ of $R_\delta \cup Q_\delta \cup v^*$ is $O(n^{3/4} \log^9 n)$. Hence the probability that the death rule is satisfied is $O(n^{-1/4} \log^9 n) = o(1)$ by union bound. This completes the proof. $\square$

Chapter 5

Unicellular maps: the local limit

[4] Our goal in this section is to identify explicitly the local limit of $U_{g_n,n}$ where $g_n \sim \theta n$ for $\theta \in (0,1/2)$ as a super-critical geometric Galton-Watson tree conditioned to survive.

5.1 Main results

We write Geom(ξ) to denote a random variable which follows the geometric distribution with parameter $\xi \in (0,1)$. In other words,

$$\mathbb{P}(\text{Geom}(\xi) = k) = (1-\xi)^{k-1}\xi \qquad \text{for } k \geq 1.$$

For any $\xi \in (0,1)$ we shall use T_ξ to denote the Galton-Watson tree with offspring distribution Geom(ξ) -1. We denote by T_ξ^∞ the tree T_ξ conditioned to be infinite. For $\xi < 1/2$ this tree is super-critical and hence the conditioning is in the classical sense. We define $T_{1/2}^\infty$ to be the limit as $n \to \infty$ of the critical tree $T_{1/2}$ conditioned to have n edges. This limit is known to exist in a much more general setting, see [63], Section 1.2 and Figure 1.2.

Theorem 5.1.1. *Assume g_n is such that $g_n/n \to \theta$ with $\theta \in [0,1/2)$. Then we have the following convergence in distribution for the local topology:*

$$U_{g_n,n} \xrightarrow[n\to\infty]{(d)} T_{\xi_\theta}^\infty,$$

where $\xi_\theta = \frac{1-\beta_\theta}{2}$, and β_θ is the unique solution in $\beta \in [0,1)$ of

$$\frac{1}{2}\left(\frac{1}{\beta} - \beta\right)\log\frac{1+\beta}{1-\beta} = (1-2\theta). \tag{5.1.1}$$

[4] The results of this section are taken from the paper [9] and is joint work with Omer Angel, Guillaume Chapuy and Nicolas Curien

For $\theta = 0$, the genus is much smaller than the size of the map, so it is not surprising that the local limit is the same as that of a critical tree conditioned to survive.

Note that the mean of the geometric offspring distribution in Theorem 5.1.1 is given by $(1+\beta_\theta)/(1-\beta_\theta) > 1$ and in particular the Galton-Watson tree is supercritical.

In order to prove Theorem 5.1.1 we first determine the root degree distribution of unicellular maps using the bijection of [34]. This is done in Section 5.2, where we also obtain an asymptotic formula for $|\mathcal{U}_{g,n}|$. This enables us to compute in Section 5.4.1 the probability that the ball of radius r around the root in $U_{g_n,n}$ is equal to any given tree. In Chapter 4 it is shown that the local limit of unicellular maps is supported on trees. However, we do not rely on this result. In Section 5.5 we show that the probabilities computed below are sufficient to characterize the local limit of $U_{g,n}$.

5.2 Enumeration and root degree distribution

Recall the definition of C-decorated tree from Section 4.2.1 and also recall that the unicellular maps of genus g and n edges are in 2^{n+1} to one correspondence between C-decorated trees of genus g and n edges (Theorem 4.2.3)

Using this correspondence we will obtain the two main theorems of this section, Theorems 5.2.1 and 5.2.2. Before stating these theorems we introduce a probability distribution on the odd integers that will play an important role in the sequel. For $\beta \in (0,1)$, we let X_β be a random variable taking its values in the odd integers, whose law is given by:

$$\mathbb{P}(X_\beta = 2k+1) := \frac{1}{Z_\beta}\frac{\beta^{2k+1}}{2k+1},$$

where

$$Z_\beta = \sum_{k\geq 0}\frac{\beta^{2k+1}}{2k+1} = \frac{1}{2}\log\frac{1+\beta}{1-\beta} = \operatorname{arctanh}\beta.$$

It is easy to check that eq. (5.1.1) is equivalent to

$$\mathbb{E}[X_\beta] = \frac{1}{Z_\beta}\frac{\beta}{1-\beta^2} = \frac{1}{1-2\theta}. \tag{5.2.1}$$

Theorem 5.2.1. *Assume $g_n \sim \theta n$ where $\theta \in (0,1/2)$. Let β_n be such that $\mathbb{E}[X_{\beta_n}] = \frac{n}{s_n} + o\left(n^{-1/2}\right)$ and $s_n = n+1-2g_n$. As n tends to infinity we* *have*

$$|\mathcal{U}_{g_n,n}| \sim A_\theta \frac{(2n)!}{n! s_n! \sqrt{s_n}} \frac{(Z_{\beta_n})^{s_n}}{4^{g_n} \beta_n^{n+1}},$$

where $A_\theta = \frac{2}{\sqrt{2\pi \operatorname{Var}(X_{\beta_\theta})}}$.

Note that $\beta_n \to \beta_\theta$. If $g_n = \theta n + o(\sqrt{n})$ we may take β_n to be just β_θ and not depend on n.

Proof. For $s, n \geq 1$, let $\mathcal{L}_s(n+1)$ be the set of partitions of $n+1$ having s parts, all of odd size. Recall that if λ is a partition of $n+1$, the number of permutations having cycle-type λ is given by

$$\frac{(n+1)!}{\prod_i m_i! i^{m_i}},$$

where for $i \geq 1$, $m_i = m_i(\lambda)$ is the number of parts of λ with size equal to i. Therefore by Theorem 4.2.3, the number of unicellular maps of genus g_n with n edges is given by

$$|\mathcal{U}_{g_n,n}| = \operatorname{Cat}(n) \frac{2^{s_n}}{2^{n+1}} \sum_{\lambda \in \mathcal{L}_{s_n}(n+1)} \frac{(n+1)!}{\prod_i m_i! i^{m_i}}, \tag{5.2.2}$$

where $\operatorname{Cat}(n) = \frac{(2n)!}{n!(n+1)!}$ is the nth Catalan number, i.e. the number of rooted plane trees with n edges, the sum counts permutations, and the powers of 2 are from the signs on cycles of the permutation and since the correspondence is 2^{n+1} to 1. Recall that this is the well-known Lehman-Walsh formula ([91]) as was described in eq. (1.5.1).

Now, let $\beta \in (0,1)$ and let $X_1, X_2, \ldots, X_s$ be i.i.d. copies of X_β. By the local central limit theorem [80, Chap.7], if $n+1 = s\mathbb{E}[X_\beta] + o(\sqrt{s})$ has the same parity as s, then $\mathbb{P}(\sum_{i \leq s} X_i = n+1) \sim As^{-1/2}$ where $A = 2/\sqrt{2\pi \operatorname{Var}(X_\beta)}$. The additional factor 2 comes from the fact that the support of X_i are odd numbers. On the other hand, we have

$$\begin{aligned}
\mathbb{P}\left(\sum_{i \leq s} X_i = n+1\right) &= \sum_{\substack{k_1 + \cdots + k_s = n+1 \\ k_i \text{ odd}}} \prod_i \frac{\beta^{k_i}}{Z_\beta \cdot k_i} \\
&= \frac{\beta^{n+1}}{(Z_\beta)^s} \sum_{\lambda \in \mathcal{L}_s(n+1)} \frac{s!}{\prod_i m_i! i^{m_i}},
\end{aligned} \tag{5.2.3}$$

since $\frac{s!}{\prod_i m_i!}$ is the number of distinct ways to order of the parts of the partition λ.

Therefore if, as in the statement of the theorem, we pick β_n so that $\mathbb{E}[X_{\beta_n}] = (n+1)/s_n + o(1/\sqrt{n})$, noticing that $\beta_n \to \beta_\theta$ and $\mathrm{Var}(X_{\beta_n}) \to \mathrm{Var}(X_{\beta_\theta})$, it follows from eq. (5.2.2) and the last considerations that

$$|\mathcal{U}_{g_n,n}| \sim \frac{1}{2^{2g_n}} \mathrm{Cat}(n) \frac{(n+1)!}{s_n!} \frac{(Z_{\beta_n})_n^s}{\beta_n^{n+1}} A_\theta s_n^{-1/2}. \qquad \square$$

The following theorem gives an asymptotic enumeration of unicellular maps of high genus with a prescribed root degree.

Theorem 5.2.2 (Root degree distribution)**.** *Assume $g_n \sim \theta n$ with $\theta \in (0,1/2)$, and let β_θ be the solution of eq.* (5.1.1)*. Then for every $d \in \mathbb{N}$ we have*

$$\mathbb{P}\left(U_{g_n,n} \text{ has root degree } d\right) \xrightarrow[n\to\infty]{} \left(\frac{1-\beta_\theta^2}{4}\right) \frac{(1+\beta_\theta)^d - (1-\beta_\theta)^d}{2^d \beta_\theta}.$$

Equivalently, the degree of the root of $U_{g_n,n}$ converges in distribution to an independent sum $\mathrm{Geom}(\frac{1+\beta_\theta}{2}) + \mathrm{Geom}(\frac{1-\beta_\theta}{2}) - 1$.

Proof. As in the proof of Theorem 5.2.1, we see that the length of a uniformly chosen cycle in a uniform random C-decorated tree with n edges and $n+1-2g_n$ cycles is distributed as the random variable X_1 conditioned on the fact that $X_1 + \cdots + X_s = n+1$, where the X_i's are i.i.d. copies of X_β for any choice of $\beta \in (0,1)$, and $s = n+1-2g_n$. This follows by writing down the required probability distributions and using eqs. (5.2.2) and (5.2.3) and Theorem 4.2.3. Using the local central limit theorem, we see that with β_n chosen according to Theorem 5.2.1, when n tends to infinity, this random variable converges in distribution to X_{β_θ}.

Since the permutation is independent of the tree, the probability that a cycle contains the root vertex is proportional to its size. Therefore the size of the cycle containing the root vertex converges in distribution to a size-biased version of X_{β_θ}, which is a random variable K with distribution $\mathbb{P}(K = 2k+1) = (1-\beta_\theta^2)\beta_\theta^{2k}$, i.e. $K = 2\,\mathrm{Geom}(1-\beta_\theta^2) - 1$.

Now by Theorem 4.2.3, conditionally on the fact that the cycle containing the root vertex has length $2k+1$, the root degree in $U_{g_n,n}$ is distributed as $\sum_{i=0}^{2k} D_i$, where D_0 if the degree of the root of a random plane tree of size n, and $(D_i)_{i>0}$ are the degrees of $2k$ uniformly chosen distinct vertices of the tree. It is classical, and easy to see, that when n tends to infinity the variables $(D_i)_{i>0}$ converge in distribution to independent $\mathrm{Geom}(1/2)$

random variables, while D_0 converges to $Y+Y'-1$, where Y, Y' are further independent Geom(1/2) variables. All geometric variables here are also independent of K.

From this it is easy to deduce that when n tends to infinity, the root degree in $U_{g_n,n}$ converges in law to $\sum_{i=0}^{K} Y_i - 1$ where K is as above and the Y_i's are independent Geom(1/2) variables. Since the probability that the sum of ℓ i.i.d. Geom(1/2) random variables equals m is $2^{-m}\binom{m-1}{\ell-1}$, we thus obtain that for all $d \geq 1$, the probability that the root vertex has degree d tends to:

$$\frac{1-\beta_\theta^2}{\beta_\theta}\sum_{k\geq 0}\beta_\theta^{2k+1}2^{-d-1}\binom{d}{2k+1} = \frac{1-\beta_\theta^2}{4\beta_\theta}\frac{(1+\beta_\theta)^d-(1-\beta_\theta)^d}{2^d}. \qquad \square$$

Remark 5.2.3. *It may be possible to prove Theorem 5.2.2 using the enumeration results for unicellular maps by vertex degrees found in [57], although this would require some computations. Here we prefer to prove it using the bijection of [34], since the proof is quite direct and gives a good understanding of the probability distribution that arises. This is also the reason we prove Theorem 5.2.1 from the bijection, rather than starting directly from the Lehman-Walsh formula* (5.2.2)*.*

We now comment on a "paradox" that the reader may have noticed. For any rooted graph G and any $r \geq 0$ we denote by $B_r(G)$ the set of vertices which are at distance less than r from the origin of the graph. In $U_{g,n}$ the mean degree can be computed as

$$\lim_{r\to\infty}\frac{1}{|B_r(U_{g,n})|}\sum_{u\in B_r(U_{g,n})}\deg(u) = \frac{2n}{v} \xrightarrow[n\to\infty]{} 2(1-2\theta)^{-1}.$$

However, if one interchanges $\lim_{n\to\infty}$ and $\lim_{r\to\infty}$ a different larger result appears. Indeed, easy arguments about Galton-Watson processes show that in $T_{\xi_\theta}^\infty$ we have

$$\lim_{r\to\infty}\frac{1}{|B_r(T_{\xi_\theta}^\infty)|}\sum_{u\in B_r(T_{\xi_\theta}^\infty)}\deg(u) = \frac{2}{1-\beta_\theta}.$$

5.3 The low genus case

Proof of Theorem 5.1.1 for $\theta = 0$. As noted, the case $g = 0$ is well known. We argue here that the local limit for $g = o(n)$ is the same as for $g = 0$.

Indeed, the permutation on the tree contains $n+1-2g$ cycles, and so has at most $3g$ non-fixed points. (If cycles of length 2 were allowed this would be $4g$.) Since the permutation is independent of the tree, and since the ball of radius r in the tree distance is tight, the probability that any vertex in the ball is in a non-trivial cycle is $o(1)$ (with constant depending on r). In particular, the local limit of the unicellular map and of the tree are the same. $\square$

5.4 The local limit

5.4.1 Surgery

Throughout this subsection, we fix integers $n, g \geq 0$. Let t be a rooted plane tree of height $r \geq 1$ with k edges and exactly d vertices at height r.

Lemma 5.4.1. *For any $n, g, k, d, r \geq 0$ we have*

$$|\{m \in \mathcal{U}_{g,n} : B_r(m) = t\}| = |\{m \in \mathcal{U}_{g,n-k+d} \textit{ with root degree } d\}|.$$

Proof. The lemma follows from a surgical argument illustrated in Fig. 5.1: if $m \in \mathcal{U}_{g,n}$ is such that $B_r(m) = t$ we can replace the r-neighborhood of the root by a star made of d edges which diminishes the number of edges of the map by $k-d$ and turns it into a map of $\mathcal{U}_{g,n-k+d}$ having root degree d. To be precise, consider the leaf of t first reached in the contour around t. The edge to this leaf is taken to be the root of the new map.

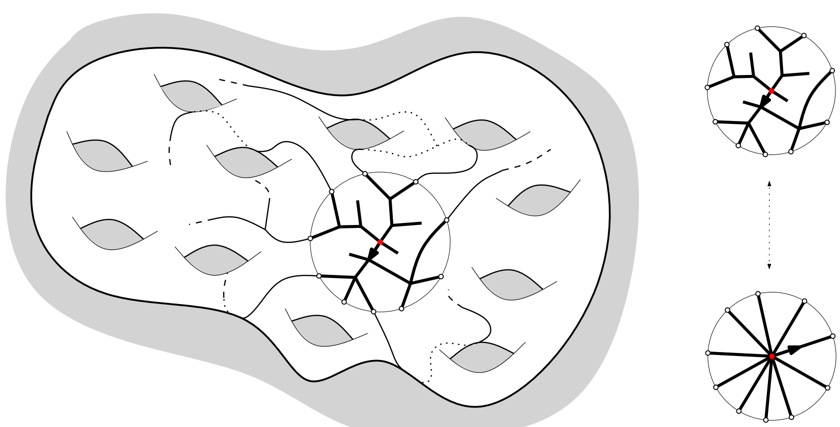

Figure 5.1: Illustration of the surgical operation

It is clear that this operation is invertible. To see that it is a bijection between the two sets in question we need to establish that it does not change the genus or number of faces in a map. One way to see this is based on an alternative description of the surgery, namely that it contracts every edge of t except those incident to the leaves, and it is easy to see that edge contraction does not change the number of faces or genus of a map.

□

5.5 Identifying the limit

Recall that for $\xi \in (0,1)$ we denote by T_ξ the law of a Galton-Watson tree with $\mathrm{Geom}(\xi) - 1$ offspring distribution. Note that when $\xi \in (0, 1/2)$ the mean offspring is strictly greater than 1 and so the process is supercritical, and recall that T_ξ^∞ is T_ξ conditioned to survive. Plane trees can be viewed as maps, rooted at the edge from the root to its first child. For every $r \geq 0$, if t is a (possibly infinite) plane tree we denote by $B_r(t)$ the rooted subtree of t made of all the vertices at height less than or equal to r.

Proposition 5.5.1. *Fix $\xi \in (0, 1/2)$. For any tree t of height exactly r having k edges and exactly d vertices at maximal height, we have*

$$\mathbb{P}\big(B_r(T_\xi^\infty) = t\big) = \frac{\big(\xi(1-\xi)\big)^{k+1-d}\big((1-\xi)^d - \xi^d\big)}{1-2\xi}.$$

Note that the probability of observing t does not depend on r, but only on the number of edges and vertices where t is connected to the rest of T_ξ.

Proof. Since $\xi \in (0, 1/2)$ the Galton-Watson process is supercritical and by standard result the extinction probability p_{die} is strictly less than 1 and is the root of $x = \sum_{k \geq 0} x^k (1-\xi)^k \xi$ in $(0,1)$. Hence

$$p_{\mathrm{die}} = \frac{\xi}{1-\xi}.$$

Next, fix a tree t of height exactly r with k edges and d vertices at height r. By the definition of T_ξ if k_u denotes the number of children of the vertex u in t we have

$$\mathbb{P}(B_r(T_\xi) = t) = \prod_u (1-\xi)^{k_u} \xi = (1-\xi)^k \xi^{k+1-d}$$

where the product is taken over all the vertices of t which are at height less than r. Conditioned on the event $\{B_r(T_\xi) = t\}$, by the branching property, the probability that the tree survives forever is $(1 - p_{\text{die}}^d)$. Combining the pieces, we get the statement of the proposition. $\square$

Proof of Theorem 5.1.1 for $\theta \in (0, 1/2)$. Under the assumptions of Theorem 5.1.1, fix r and let t be a rooted oriented tree of height exactly r having k edges and exactly d vertices at height r. By Lemma 5.4.1 we have

$$\begin{aligned}\mathbb{P}(B_r(U_{g_n,n}) = t) &= \frac{|\{m \in \mathcal{U}_{g_n,n-k+d} \text{ with root degree } d\}|}{|\mathcal{U}_{g_n,n}|} \\ &= \frac{|\mathcal{U}_{g_n,n-k+d}|}{|\mathcal{U}_{g_n,n}|} \cdot \mathbb{P}(\text{root degree of } U_{g_n,n-k+d} = d).\end{aligned}$$

Applying Theorem 5.2.2 we have

$$\mathbb{P}(\text{root degree of } U_{g_n,n-k+d} = d) \xrightarrow[n\to\infty]{} \left(\frac{1-\beta_\theta^2}{4\beta_\theta}\right) \frac{(1+\beta_\theta)^d - (1-\beta_\theta)^d}{2^d}. \tag{5.5.1}$$

On the other hand, since $n/s = (n-k+d)/(s-k+d) + o(1/\sqrt{n})$ we can apply Theorem 5.2.1 for the asymptotic of $|\mathcal{U}_{g_n,n-k+d}|$ and $|\mathcal{U}_{g_n,n}|$ with the same sequence (β_n) and get that

$$\frac{|\mathcal{U}_{g_n,n-k+d}|}{|\mathcal{U}_{g_n,n}|} \sim \frac{(2n+2d-2k)!n!s!Z_{\beta_n}^{d-k}}{(2n)!(n+d-k)!(s+d-k)!\beta_n^{d-k}}.$$

Since d, k are fixed, and using the facts that $\beta_n \to \beta_\theta$, $Z_{\beta_n} \to Z_{\beta_\theta}$ and $s/n \to (1-2\theta)$, the last display is also equivalent to

$$\frac{|\mathcal{U}_{g_n,n-k+d}|}{|\mathcal{U}_{g_n,n}|} \sim \left(\frac{\beta_\theta(1-2\theta)}{4Z_{\beta_\theta}}\right)^{k-d} = \left(\frac{1-\beta_\theta^2}{4}\right)^{k-d}, \tag{5.5.2}$$

by the definition of β_θ in eq. (5.2.1). Plugging (5.5.1) and (5.5.2) together and using Proposition 5.5.1 we find that

$$\mathbb{P}(B_r(U_{g_n,n}) = t) \xrightarrow[n\to\infty]{} \mathbb{P}(B_r(T_{\xi_\theta}^\infty) = t),$$

with $\xi_\theta = (1-\beta_\theta)/2$.

Finally, note that the law of $B_r(T_{\xi_\theta}^\infty)$ is a probability measure on the set of finite plane trees. It follows that $B_r(U_{g_n,n})$ is tight, and converges in distribution to $B_r(T_{\xi_\theta}^\infty)$. Since r is arbitrary, this completes the proof of the Theorem. $\square$

5.6 Questions and remarks

Planarity. A consequence of Theorem 5.1.1 is that $U_{g_n,n}$ is locally a tree (hence planar) near its root. More precisely, the length of a minimal non-trivial cycle containing the root edge diverges in probability as $n \to \infty$. Note that a much stronger statement has been proved in Chapter 4 where quantitative estimates on cycle lengths are obtained. As noted above, the proof presented in this section does not rely on this result and our approach is softer. Note that our method of proof only requires to prove convergences of the quantities $\mathbb{P}(B_r(U_{g_n,n}) = t)$ when t is a tree since we were able to identify these limits as coming from a probability measure on infinite trees.

Open questions. We gather here a couple of possible extensions of our work.

Question 5.6.1. *Find more precise asymptotic formulae for $|\mathcal{U}_{g,n}|$ as the sequence $\{g, n\} \to \{\infty, \infty\}$. Theorem 5.2.1 gives a first order approximation.*

Question 5.6.2. *Quantify the convergence of $U_{g_n,n}$ to T_{ξ_θ}. In particular, let $r_n = o(\log n)$. Is it possible to couple $U_{g_n,n}$ with T_{ξ_θ} so that $B_{r_n}(U_{g_n,n}) = B_{r_n}(T_{\xi_\theta})$ with high probability?*

Bibliography

[1] 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.

[2] L. Addario-Berry, N. Broutin, C. Goldschmidt, and G. Miermont. The scaling limit of the minimum spanning tree of the complete graph. 2013. arXiv:1301.1664.

[3] L. Addario-Berry, L. Devroye, and S. Janson. Sub-Gaussian tail bounds for the width and height of conditioned Galton-Watson trees. *Ann. Probab.*, 41(2):1072–1087, 2013.

[4] D. Aldous. The continuum random tree. I. *Ann. Probab.*, 19(1):1–28, 1991.

[5] 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.

[6] J. Ambjørn. Quantization of geometry. In *Géométries fluctuantes en mécanique statistique et en théorie des champs (Les Houches, 1994)*, pages 77–193. North-Holland, Amsterdam, 1996.

[7] O. Angel. Growth and percolation on the uniform infinite planar triangulation. *Geom. Funct. Anal.*, 13(5):935–974, 2003.

[8] O. Angel. Scaling of percolation on infinite planar maps, I. arXiv:math/0501006, 2005.

[9] O. Angel, G. Chapuy, N. Curien, and G. Ray. The local limit of unicellular maps in high genus. *Elec. Comm. Prob*, 18:1–8, 2013.

[10] O. Angel and N. Curien. Percolations on random maps I: half-plane models. *Ann. Inst. H. Poincaré*, 2013. To appear.

[11] O. Angel, A. Nachmias, and G. Ray. Random walks in stochastic hyperbolic planar triangulations. In preparation, 2014.

[12] O. Angel and G. Ray. Classification of half planar maps. *Ann. Probab.*, 2013. arXiv:1303.6582.

[13] O. Angel and O. Schramm. Uniform infinite planar triangulations. *Comm. Math. Phys.*, 241(2-3):191–213, 2003.

[14] K. Athreya and P. Ney. The local limit theorem and some related aspects of supercritical branching processes. *Trans. Amer. Math. Soc.*, 152(2):233251, 1970.

[15] C. Banderier, P. Flajolet, G. Schaeffer, and M. Soria. Planar maps and Airy phenomena. In *Automata, languages and programming (Geneva, 2000)*, volume 1853 of *Lecture Notes in Comput. Sci.*, pages 388–402. Springer, Berlin, 2000.

[16] C. Banderier, P. Flajolet, G. Schaeffer, and M. Soria. Random maps, coalescing saddles, singularity analysis, and Airy phenomena. *Random Structures Algorithms*, 19(3-4):194–246, 2001. Analysis of algorithms (Krynica Morska, 2000).

[17] M. T. Barlow, A. A. Járai, T. Kumagai, and G. Slade. Random walk on the incipient infinite cluster for oriented percolation in high dimensions. *Comm. Math. Phys.*, 278(2):385–431, 2008.

[18] I. Benjamini. Random planar metrics. In *Proceedings of the International Congress of Mathematicians. Volume IV*, pages 2177–2187, New Delhi, 2010. Hindustan Book Agency.

[19] I. Benjamini. Coarse geometry and randomness. *cole dt de Probabilits de Saint-Flour XLI*, 2011. http://www.wisdom.weizmann.ac.il/ itai/stflouraug24.pdf.

[20] 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.

[21] I. Benjamini, R. Lyons, and O. Schramm. Unimodular random trees. arXiv:1207.1752, 2012.

[22] I. Benjamini and O. Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6:no. 23, 1–13, 2001.

[23] I. Benjamini and O. Schramm. KPZ in one dimensional random geometry of multiplicative cascades. *Comm. Math. Phys.*, 289(2):653–662, 2009.

[24] I. Benjamini and O. Schramm. Percolation beyond $\mathbb{Z}^d$, many questions and a few answers [mr1423907]. In *Selected works of Oded Schramm. Volume 1, 2*, Sel. Works Probab. Stat., pages 679–690. Springer, New York, 2011.

[25] O. Bernardi. An analogue of the Harer-Zagier formula for unicellular maps on general surfaces. *Adv. in Appl. Math.*, 48(1):164–180, 2012.

[26] O. Bernardi and G. Chapuy. Counting unicellular maps on non-orientable surfaces. *Adv. in Appl. Math.*, 47(2):259–275, 2011.

[27] J. Bettinelli. Scaling limits for random quadrangulations of positive genus. *Electron. J. Probab.*, 15:no. 52, 1594–1644, 2010.

[28] J. D. Biggins and N. H. Bingham. Large deviations in the supercritical branching process. *Adv. in Appl. Probab.*, 25(4):757–772, 1993.

[29] P. Billingsley. *Convergence of probability measures.* Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons, Inc., New York, second edition, 1999. A Wiley-Interscience Publication.

[30] J. Bouttier and E. Guitter. Distance statistics in quadrangulations with a boundary, or with a self-avoiding loop. *J. Phys. A*, 42(46):465208, 44, 2009.

[31] W. G. Brown. Enumeration of triangulations of the disk. *Proc. London Math. Soc. (3)*, 14:746–768, 1964.

[32] C. E. Burgess. Classification of surfaces. *Amer. Math. Monthly*, 92(5):349–354, 1985.

[33] G. Chapuy. A new combinatorial identity for unicellular maps, via a direct bijective approach. *Adv. in Appl. Math.*, 47(4):874–893, 2011.

[34] G. Chapuy, V. Féray, and É. Fusy. A simple model of trees for unicellular maps. *Journal of Combinatorial Theory, Series A*, 120(8):2064–2092, 2013.

[35] G. Chapuy, M. Marcus, and G. Schaeffer. A bijection for rooted maps on orientable surfaces. *SIAM J. Discrete Math.*, 23(3):1587–1611, 2009.

[36] P. Chassaing and G. Schaeffer. Random planar lattices and integrated superBrownian excursion. *Probab. Theory Related Fields*, 128(2):161–212, 2004.

[37] R. Cori and B. Vauquelin. Planar maps are well labeled trees. *Canad. J. Math.*, 33(5):1023–1042, 1981.

[38] N. Curien. A glimpse of the conformal structure of random planar maps. arXiv:1308.1807, 2013.

[39] N. Curien. Planar stochastic hyperbolic infinite triangulations. arXiv:1401.3297, 2014.

[40] N. Curien, B. Haas, and I. Kortchemski. The CRT is the scaling limit of random dissections. 2013. arXiv:1305.3534.

[41] N. Curien and I. Kortchemski. Percolation on random triangulations and stable looptrees. 2013. arXiv:1307.6818.

[42] N. Curien and J.-F. Le Gall. The Brownian plane. *J. Theoret. Probab.*, pages 1–43, 2012.

[43] N. Curien, L. Ménard, and G. Miermont. A view from infinity of the uniform infinite planar quadrangulation. *ALEA Lat. Am. J. Probab. Math. Stat.*, 10(1):45–88, 2013.

[44] N. Curien and G. Miermont. Uniform infinite planar quadrangulations with a boundary. *Random Structures Algorithms*, 2012. to appear.

[45] B. Davis and D. McDonald. An elementary proof of the local central limit theorem. *J. Theoret. Probab.*, 8(3):693–701, 1995.

[46] L. Devroye and S. Janson. Distances between pairs of vertices and vertical profile in conditioned Galton-Watson trees. *Random Structures Algorithms*, 38(4):381–395, 2011.

[47] B. Duplantier and S. Sheffield. Liouville quantum gravity and KPZ. *Invent. Math.*, 185(2):333–393, 2011.

[48] R. Durrett. *Probability: theory and examples.* Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, fourth edition, 2010.

[49] W. Feller. *An introduction to probability theory and its applications. Vol. II.* Second edition. John Wiley & Sons Inc., New York, 1971.

[50] K. Fleischmann and V. Wachtel. Lower deviation probabilities for supercritical Galton-Watson processes. *Ann. Inst. H. Poincaré Probab. Statist.*, 43(2):233–255, 2007.

[51] R. G. Gallager. *Discrete stochastic processes*, volume 101. Kluwer Academic Publishers Boston, 1996.

[52] A. Gamburd. Poisson-Dirichlet distribution for random Belyi surfaces. *Ann. Probab.*, 34(5):1827–1848, 2006.

[53] A. Gamburd and E. Makover. On the genus of a random Riemann surface. In *Complex manifolds and hyperbolic geometry (Guanajuato, 2001)*, volume 311 of *Contemp. Math.*, pages 133–140. Amer. Math. Soc., Providence, RI, 2002.

[54] Z. Gao. The number of rooted triangular maps on a surface. *J. Combin. Theory Ser. B*, 52(2):236–249, 1991.

[55] Z. Gao. A pattern for the asymptotic number of rooted maps on surfaces. *J. Combin. Theory Ser. A*, 64(2):246–264, 1993.

[56] 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.

[57] A. Goupil and G. Schaeffer. Factoring n-cycles and counting maps of given genus. *European J. Combin.*, 19(7):819–834, 1998.

[58] M. Gromov. *Metric structures for Riemannian and non-Riemannian spaces*. Modern Birkhäuser Classics. Birkhäuser Boston Inc., Boston, MA, english edition, 2007. Based on the 1981 French original, With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by Sean Michael Bates.

[59] O. Gurel-Gurevich and A. Nachmias. Recurrence of planar graph limits. *Ann. of Math. (2)*, 177(2):761–781, 2013.

[60] T. E. Harris. *The theory of branching processes*. Dover Phoenix Editions. Dover Publications Inc., Mineola, NY, 2002. Corrected reprint of the 1963 original [Springer, Berlin; MR0163361 (29 #664)].

[61] Z.-X. He and O. Schramm. Hyperbolic and parabolic packings. *Discrete Comput. Geom.*, 14(2):123–149, 1995.

[62] N. I. Kazimirov. The occurrence of a gigantic component in a random permutation with a known number of cycles. *Diskret. Mat.*, 15(3):145–159, 2003.

[63] H. Kesten. Subdiffusive behavior of random walk on a random cluster. *Ann. Inst. H. Poincaré Probab. Statist.*, 22(4):425–487, 1986.

[64] 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.

[65] M. Krikun. Local structure of random quadrangulations. *arXiv:math/0512304*, 2005.

[66] T. Kumagai and J. Misumi. Heat kernel estimates for strongly recurrent random walk on random media. *J. Theoret. Probab.*, 21(4):910–935, 2008.

[67] K. Kuratowski. *Topology. Vol. I.* New edition, revised and augmented. Translated from the French by J. Jaworowski. Academic Press, New York, 1966.

[68] S. K. Lando and A. K. Zvonkin. *Graphs on surfaces and their applications*, volume 141 of *Encyclopaedia of Mathematical Sciences*. Springer-Verlag, Berlin, 2004. With an appendix by Don B. Zagier, Low-Dimensional Topology, II.

[69] J.-F. Le Gall. Random trees and applications. *Probab. Surv.*, 2:245–311, 2005.

[70] J.-F. Le Gall. Uniqueness and universality of the Brownian map. *Ann. Probab.*, 41(4):2880–2960, 2013.

[71] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, Providence, RI, 2009. With a chapter by James G. Propp and David B. Wilson.

[72] R. Lyons and Y. Peres. *Probability on Trees and Networks*. Cambridge University Press. In preparation. Current version available at `http://mypage.iu.edu/~rdlyons/`.

[73] W. S. Massey. *Algebraic topology: an introduction*. Springer-Verlag, New York, 1977. Reprint of the 1967 edition, Graduate Texts in Mathematics, Vol. 56.

[74] L. Ménard and P. Nolin. Percolation on uniform infinite planar maps. 2013. arXiv:1302.2851.

[75] G. Miermont. Random maps and their scaling limits. In *Fractal geometry and stochastics IV*, volume 61 of *Progr. Probab.*, pages 197–224. Birkhäuser Verlag, Basel, 2009.

[76] G. Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Mathematica, to appear*, 2011. arXiv:1104.1606.

[77] J. Miller and S. Sheffield. Quantum Loewner Evolution. arXiv:1312.5745, 2013.

[78] B. Mohar and C. Thomassen. *Graphs on surfaces.* Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, 2001.

[79] S. V. Nagaev and N. V. Vahrusev. Estimation of probabilities of large deviations for a critical Galton-Watson process. *Theory Probab. Appl.*, 20(1):179180, 1975.

[80] V. V. Petrov. *Sums of independent random variables.* Springer-Verlag, New York, 1975. Translated from the Russian by A. A. Brown, Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 82.

[81] G. Ray. Large unicellular maps in high genus. *Ann. Inst. H. Poincaré*, 2013. To appear. arXiv:1307.1224.

[82] G. Ray. Geometry and percolation on half planar triangulations. *Electron. J. Probab.*, 19(47):1–28, 2014.

[83] H. Ren, Y. Liu, and Z. Li. Enumeration of 2-connected loopless 4-regular maps on the plane. *European J. Combin.*, 23(1):93–111, 2002.

[84] G. Schaeffer. *Conjugaison d'arbres et cartes combinatoires aléatoires.* PhD thesis, 1998.

[85] O. Schramm. Scaling limits of loop-erased random walks and uniform spanning trees. *Israel J. Math.*, 118:221–288, 2000.

[86] C. Thomassen. Isoperimetric inequalities and transient random walks on graphs. *Ann. Probab.*, 20(3):1592–1600, 1992.

[87] W. T. Tutte. A census of planar triangulations. *Canad. J. Math.*, 14:21–38, 1962.

[88] W. T. Tutte. A census of slicings. *Canad. J. Math.*, 14:708–722, 1962.

[89] W. T. Tutte. A census of planar maps. *Canad. J. Math.*, 15:249–271, 1963.

[90] B. Virág. Anchored expansion and random walk. *Geom. Funct. Anal.*, 10(6):1588–1605, 2000.

[91] T. R. S. Walsh and A. B. Lehman. Counting rooted maps by genus. *J. Combin. Theory Ser. B.*, 13:192–218, 1972.

[92] 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.

Appendix A

Proof of lemma 3.2.2

Recall that $I_m(q)$ denotes the number of internal vertices of a free triangulation of an m-gon and recall the variable θ used in Section 2.5.1 where $q = \theta(1-2\theta)^2$.

Proof of Lemma 3.2.2 part (i). Without loss of generality assume x is an integer. Let $d_\theta = \frac{4\theta}{(1-6\theta)}$. For simplicity of notation let $I_m(q) = I_m$. Notice that conditioned on $Y = k$, expectation of I_{Y+1} is $d_\theta k + O(1)$ as $k \to \infty$. We want,

$$\mathbb{P}(Y + I_{Y+1} > x) = \sum_{k \geq 1} \mathbb{P}(I_{Y+1} > x - k | Y = k)\mathbb{P}(Y = k) \qquad \text{(A.0.1)}$$

The trick is to break the sum in (A.0.1) into sums over three subsets of indices:

(i) $A_1 = \{1 \leq k \leq \lfloor x/(1+d_\theta) - x^{3/4} \rfloor\}$

(ii) $A_3 = \{k > \lfloor x/(1+d_\theta) + x^{3/4} \rfloor\}$

(iii) $A_2 = \mathbb{N} \setminus (A_1 \cup A_3)$

The sum over A_2 is $O(x^{-3/4})$ by bounding $\mathbb{P}(I_{Y+1} > x - k|Y = k)$ by 1 and using $\mathbb{P}(Y = k) \sim ck^{-3/2}$. Now note

$$\sum_{A_1} \mathbb{P}(I_{Y+1} > x - k|Y = k)\mathbb{P}(Y = k) \qquad \text{(A.0.2)}$$

$$< \sum_{A_1} \mathbb{P}(I_{Y+1} - \mathbb{E}(I_{Y+1}|Y = k) > x^{3/4} + O(1)|Y = k)\mathbb{P}(Y = k) \qquad \text{(A.0.3)}$$

$$< \sum_{A_1} \frac{Var(I_{Y+1}|Y = k)}{x^{3/2}} \mathbb{P}(Y = k) = O(x^{-1}). \qquad \text{(A.0.4)}$$

where we used Proposition 3.2.1 part (i) for (A.0.3) and Chebyshev's inequality followed by Proposition 3.2.1 part (ii) for (A.0.4). Finally,

$$\sum_{k\in A_3} \mathbb{P}(I_{Y+1} > x-k|Y=k)\mathbb{P}(Y=k) \tag{A.0.5}$$

$$= \sum_{k\in A_3} \mathbb{P}(Y=k) - \sum_{k\in A_3} \mathbb{P}(I_{Y+1} \leq x-k|Y=k)\mathbb{P}(Y=k) \tag{A.0.6}$$

$$= \sum_{k\in A_3} \mathbb{P}(Y=k) - O(x^{-1}) \tag{A.0.7}$$

where the bound in the second term in the right hand side of (A.0.7) follows in the same way as (A.0.4) using Chebyshev's inequality and Proposition 3.2.1 part (ii) plus the fact that the summands are 0 when $k > x$. Finally it is easy to verify that $\sum_{k\in A_3} \mathbb{P}(Y = k) \sim c_\theta x^{-1/2}$ for some constant $c_\theta > 0$. □

Proof of Lemma 3.2.2 part (ii). Note that

$$\mathbb{E}((Y+I_{Y+1})\mathbb{1}_{\{Y+I_{Y+1}<x\}}) = \sum_{k=1}^{x-1}(\mathbb{P}(Y+I_{Y+1}\geq k) - \mathbb{P}(Y+I_{Y+1}\geq x)) \tag{A.0.8}$$

Now the asymptotics follows by using the asymptotics of part (i). □

Proof of Lemma 3.3.2. Observe that the expected change is given by

$$\alpha - \sum_{i\geq 1} ip_i = 1 - \sum_{i\geq 1}(i+1)p_i$$

where p_i is given by eq. (2.8.3) and the equality follows from the fact that $\sum_{i\geq 1} p_i = 1-\alpha$. Now from eq. (2.8.3),

$$\sum_{i\geq 1}(i+1)p_i = \sum_{i\geq 1} 2Cat(i-1)\left(\frac{2/\alpha - 2}{4}\right)^i ((3\alpha-2)i+1) \tag{A.0.9}$$

where $Cat(n) = \frac{1}{n+1}\binom{2n}{n}$ is the nth catalan number. The sum in the right hand side of eq. (A.0.9) can be easily computed using generating functions of catalan numbers. We leave this last step to the reader. □

Appendix B

Proof of lemma 4.2.6

We shall prove Lemma 4.2.6 in this section. We do the computation following the method of random allocation similar in lines of [62]. For this, we need to introduce i.i.d. random variables $\{\xi_1, \xi_2, \ldots\}$ such that for some parameter $\beta \in (0,1)$

$$P(\xi_1 = i) \begin{cases} = \frac{\beta^i}{B(\beta)(i)} & \text{if } i \text{ is odd positive integer} \\ = 0 & \text{otherwise} \end{cases} \tag{B.0.1}$$

where $B(\beta) = 1/2 \log((1+\beta)/(1-\beta))$. Recall that $\mathcal{P}$ is the set of all N-tuples of odd positive integers which sum up to $n+1$. Observe that for any $z = (z_1, \ldots, z_N) \in \mathcal{P}$,

$$P(\lambda = z) = P(\xi_1 = z_1, \ldots, \xi_N = z_N | \xi_1 + \xi_2 + \ldots \xi_N = n+1)$$

throughout this Section, we shall assume the following:

- $\{n, N\} \to \{\infty, \infty\}$ and $n/N \to \alpha$ for some constant $\alpha > 1$.
- For every n, the parameter $\beta = \beta(n)$ is chosen such that $E(\xi_1) = m = (n+1)/N$

It is easy to check using (B.0.1) that there is a unique choice of such β and β converges to some finite number β_α such that $0 < \beta_\alpha < 1$ as $(n+1)/N$ converges to α. Let $\zeta_{N,j} = \xi_1^j + \ldots \xi_N^j$ where $j \geq 1$ is an integer. It is also easy to see that for any integer $j \geq 1$, $\mathbb{E}\xi_1^j = m_j(n)$ for some function m_j which also converge to some number $m_{j\alpha}$ as $n \to \infty$. Let $\sigma_j^2 = Var(\xi_1^j)$. To simplify notation, we shall denote $\zeta_{N,1}$ by ζ_N and σ_1 by σ.

We will first prove a central limit theorem for ζ_N.

Lemma B.0.3. *We have,*

$$\frac{\zeta_N - Nm}{\sigma\sqrt{N}} \to N(0,1) \tag{B.0.2}$$

in distribution as $n \to \infty$.

Proof. Easily follows by checking the Lyapunov condition for triangular arrays of random variables (see [48]). □

We now prove a local version of the CLT asserted by Lemma B.0.3.

Lemma B.0.4. *We have*

$$P(\zeta_N = n+1) \sim \frac{2}{\sqrt{2\pi N}\sigma}$$

Proof. Let $\overline{\xi_i} = (\xi_i - 1)/2$. Apply Theorem 1.2 of [45] for the modified arrays $\{\overline{\xi_1}, \dots \overline{\xi_N}\}_{n\geq 1}$ and use Lemma B.0.3. The details are left for the readers to check. □

Lemma B.0.5. *Fix $j \geq 1$. There exists constants $C_1 > 1$ and $C_2 > 1$ (both depending only on α and j) such that*

$$P\left(C_1 n < \sum_{i=1}^{N} \lambda_i^j < C_2 n\right) > 1 - \frac{c}{n^{7/2}} \tag{B.0.3}$$

for some $c > 0$ which again depends only on α and j for large enough n.

Proof. It is easy to see that $m_j \to m_{j\alpha}$ as $n \to \infty$. Notice that

$$\begin{aligned} P\left(\sum_{i=1}^{N} \lambda_i^j > C_2 n, \sum_{i=1}^{N} \lambda_i^j < C_1 n\right) &= P\left(\zeta_{N,j} > C_2 n, \zeta_{N,j} < C_1 n | \zeta_N = n\right) \\ &< \frac{P\left(\zeta_{N,j} > C_2 n, \zeta_{N,j} < C_1 n\right)}{P(\zeta_N = n)} \end{aligned} \tag{B.0.4}$$

Choose $C_2 > m_{j\alpha}$ and $C_1 < m_{j\alpha}$. Then for some $c > 0$, for large enough n,

$$\begin{aligned} P\left(\zeta_{N,j} > C_2 n, \zeta_{N,j} < C_1 n\right) &< P\left(|\zeta_{N,j} - m_j N| > cn\right) \\ &< \mathbb{E}\left(\zeta_{N,j} - m_j N\right)^8 (cn)^{-8} \end{aligned} \tag{B.0.5}$$

It is easy to see that $\mathbb{E}\left(\zeta_{N,j} - m_j N\right)^8 = O(n^4)$ since the terms involving $\mathbb{E}(\xi_i^j - m_j)$ vanishes and all the finite moments of ξ_1 are bounded. Now plugging in this estimate into eq. (B.0.5), we get

$$P\left(\zeta_{N,j} > C_2 n, \zeta_{N,j} < C_1 n\right) = O(n^{-4}) \tag{B.0.6}$$

Now plugging in the estimate of eq. (B.0.6) into eq. (B.0.4) and observing that $P(\zeta_N = n) \asymp N^{-1/2}$ via Lemma B.0.4, the result follows. □

Lemma B.0.6. *There exists a constant $C_0 > 0$ such that*

$$P(\lambda_{\max} > C_0 \log n) = O(n^{-3})$$

Proof. Let $\xi_{\max}$ be the maximum among $\xi_1, \ldots, \xi_N$. Note that

$$\mathbb{P}(\xi_{\max} > C_0 \log n) < N\mathbb{P}(\xi_1 > C_0 \log n) = O(N\beta^{C_0 \log n}) = O(n^{-7/2}) \tag{B.0.7}$$

if $C_0 > 0$ is chosen large enough. Now the Lemma follows from the estimate of Lemma B.0.4. The details are left to the reader. □

Lemma B.0.7. *There exists constants $0 < d_1 < 1$ and $0 < d_2 < 1$ which depends only on α such that $\mathbb{P}(d_1 n < |i : \lambda_i = 1| < d_2 n) < e^{-cn}$ for some constant $c > 0$ for large enough n.*

Proof. The probability that $|i : \lambda_i = 1| < d_2 n$ for large enough n for some $0 < d_2 < 1$ follows directly from the fact that $|i : \lambda_i = 1| \leq N$. For the upper bound,

$$\mathbb{P}(|i : \lambda_i = 1| > d_1 n) = \mathbb{P}\left(\sum_{i=1}^{N} 1_{\lambda_i = 1} > d_1 n\right) < \frac{\mathbb{P}\left(\sum_{i=1}^{N} 1_{\xi_i = 1} > d_1 n\right)}{\mathbb{P}(\zeta_N = n)} \tag{B.0.8}$$

Now $\mathbb{P}(\xi_i = 1) \to \beta_\alpha / B(\beta_\alpha)$ as $n \to \infty$. The Lemma now follows by choosing d_1 small enough, applying Lemma B.0.4 to the denominator in eq. (B.0.8) and a suitable large deviation bound on Bernoulli variables. Details are standard and is left to the reader. □

Proof of Lemma 4.2.6. Follows from Lemmas B.0.5–B.0.7. □

Appendix C

Proofs of the lemmas in section 4.2.3

In this section, we shall prove Lemmas 4.2.8 and 4.2.10.

Let p_i denote $\mathbb{P}(\xi = i)$ for $i \in \mathbb{N}$ and denote the generating function by $\varphi(s) = \sum_i p_i s^i$. Let $\mu = \mathbb{E}\xi$. Let Z_n denote the number of offsprings in the n-th generation of the Galton-Watson process

C.1 Critical Galton-Watson trees

We assume ξ has geometric distribution with parameter $1/2$. Here $\mu = 1$ and we want to show that Z_r cannot be much more than r. The following large deviation result is a special case of the main theorem of [79].

Proposition C.1.1. *For all $r \geq 1$ and $k \geq 1$,*

$$\mathbb{P}(Z_r \geq k) < \frac{3}{2}\left(1 + \frac{1}{\varphi''(3/2)r/2 + 2}\right)^{-k}$$

C.2 Supercritical Galton-Watson trees

Here $\mu > 1$. Recall the assumptions

- $0 < p_0 + p_1 < 1$
- There exists a small enough $\lambda > 0$ such that $\mathbb{E}(e^{\lambda\xi}) < \infty$.

It is well known (see [60]) that Z_n/μ^n is a martingale which converges almost surely to some non-degenerate random variable W. Let $\rho := P(\lim_n Z_n = 0)$ be the extinction probability which is strictly less than 1 in the supercritical regime.

The following results may be realized as special cases of the results in [14], [50] and further necessary references can be found in these papers.

W if restricted to $(0, \infty)$ has a strictly positive continuous density which is denoted by w. In other words, we have the following limit theorem:

$$\lim_n \mathbb{P}(Z_n \geq x\mu^n) = \int_x^\infty w(t)dt, \qquad x > 0$$

Also define $\gamma := \varphi'(\rho)$ where $0 < \gamma < 1$ in our case. Define β by the relation $\gamma = \mu^{-\beta}$. It is clear that in our case $\beta \in (0, \infty)$. β is used to determine the behavior of w as $x \downarrow 0$. The following is proved in [14].

Proposition C.2.1. *Let $\eta := \mu^{\beta/(3+\beta)} > 1$. Then for all $\varepsilon \in (0, \eta)$, there exists a positive constant $C_\varepsilon > 0$ such that for all $k \geq 1$,*

$$|\mathbb{P}(Z_r = k)\mu^r - w(k/\mu^r)| \leq C_\varepsilon \frac{\eta^{-r}}{k\mu^{-r}} + (\eta - \varepsilon)^{-r} \tag{C.2.1}$$

for all $r \geq 1$.

It can be shown (see [28]) that there exists positive constants $A_1 > 0, A_2 > 0$ such that $A_1 x^{\beta-1} < w(x) < A_2 x^{\beta-1}$ as $x \downarrow 0$. Using this and eq. (C.2.1), we get

$$\mathbb{P}(Z_r = k) \leq C\frac{k^{\beta-1}}{\mu^{r\beta}} + \frac{\eta^{-r}}{k} + ((\eta - \varepsilon)\mu)^{-r} \tag{C.2.2}$$

Proof of Lemma 4.2.8. The proof is straightforward by summing k from 1 to γ^r the expression given by the right hand side of (C.2.2). □

C.3 Random plane trees

Proof of Lemma 4.2.10. Note that it is enough to prove the bound for $r \leq n$ because otherwise the probability is 0. It is well known that if we pick an oriented edge uniformly from $U_{0,n}$ and re-root the tree there then the distribution of this new re-rooted tree is the same as that of $U_{0,n}$ (see [46]). Let V denote the root vertex of the new re-rooted tree and let $\mathcal{Z}_j(V)$ denote the number of vertices at distance exactly j from V. It is well known that the probability of a critical geometric Galton-Watson tree to have n edges is $\asymp n^{-3/2}$. Using this fact and Proposition C.1.1 we get for any $k \geq 1$ and $1 \leq j \leq r$

$$\mathbb{P}(\mathcal{Z}_j(V) > k) < n^{3/2} c \exp(-c'k/j) < n^{3/2} c \exp(-c'k/r) \tag{C.3.1}$$

and some suitable positive constants c, c'. Note that if $M_r > r^2 \log^2 n$ then $\mathcal{Z}_j(v) > r \log^2 n$ for some $1 \leq j \leq r$ and some vertex $v \in U_{0,n}$. Using this and union bound to the estimate obtained in (C.3.1), we get

$$\mathbb{P}(M_r > r^2 \log^2 n) < cn^{5/2} r \exp(-c' \log^2 n) = O(\exp(-c' \log^2 n))$$

for some positive constants c and c'. This completes the proof. □