Introduction to the Gaussian Free Field and Liouville Quantum Gravity

Nathanaël Berestycki

[Draft lecture notes: July 19, 2016]

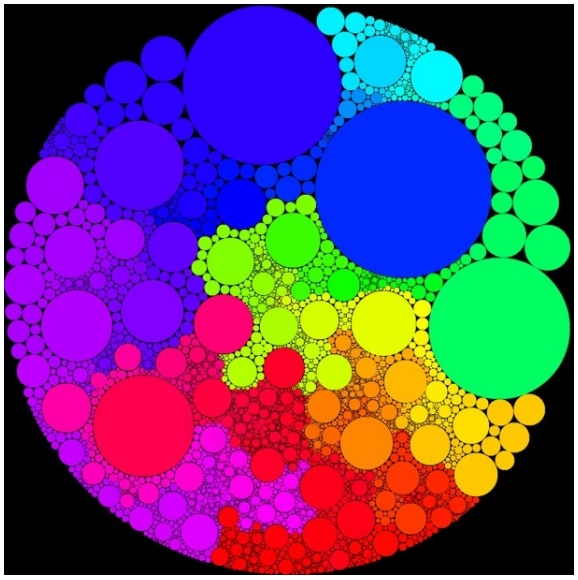

Contents

Introduction

These lecture notes are intended to be an introduction to the two-dimensional continuum Gaussian Free Field and Liouville Quantum Gravity. Topics covered include the definition and main properties of the GFF, the construction of the Liouville measure, its non degeneracy and conformal covariance, and applications to the KPZ formula. An extra chapter covers Sheffield's quantum zipper. This section is quite technical and readers are advised that this will be of most use to people who are actively working in this area. (While the arguments in that section follow the main ideas of [She10a], some are new and others have been simplified; and the structure of the proof has been rethought. The result is, I hope, that some of the main ideas are more transparent and will be helpful to others as it was to me.)

The theory is in full blossom and attempting to make a complete survey of the field would be foolish, so quickly is the theory developing. Nevertheless, as the theory grows in complexity and applicability, it has appeared useful to me to summarise some of its basic and foundational aspects in one place, especially since complete proofs of some of these facts can be spread over a multitude of papers.

Clearly, the main drawback of this approach is that many of the important subsequent developments and alternative points of view are not included. For instance the theory is deeply intertwined with Kahane's theory of Gaussian multiplicative chaos, but this will not be apparent in these notes. Likewise, topics such as random planar maps, the Brownian map, Liouville Brownian motion, critical Gaussian multiplicative chaos, as well as the beautiful way in which Duplantier, Miller and Sheffield have made sense of and analysed the "peanosphere" (which is really the scaling limit of random planar maps equipped with a collection of loops coming from a critical FK model), imaginary geometry and QLE, are not covered. The list is endless.

To help beginners, I have tried to make a distinction between what is an important calculation and what can be skipped in a first reading without too much loss. Such sections are indicated with a star. In a first reading, the reader is encouraged to skip these sections.

Acknowledgements:

A first draft was written in preparation for the LMS / Clay institute Research school on "Modern developments in probability" taking place in Oxford, July 2015. The draft was revised in April 2016 on the occasion of the Spring School on "Geometric models in probability" in Darmstadt. In both cases I thank the organisers (respectively, Christina Goldschmidt and Dmitry Beliaev; and Volker Betz and Matthias Meiners) for their invitations and superb organisation. Thanks also to Benoit Laslier for agreeing to run the exercise sessions accompanying the lectures at the initial Clay Research school.

The notes were written towards the end of my stay at the Newton institute during the semester on *Random Geometry*. Thanks to the INI for its hospitality and to some of the programme participants for enlightening discussions related to aspects discussed in these

notes: especially, Omer Angel, Juhan Aru, Stéphane Benoît, Bertrand Duplantier, Ewain Gwynne, Nina Holden, Henry Jackson, Benoit Laslier, Jason Miller, James Norris, Ellen Powell, Gourab Ray, Scott Sheffield, Xin Sun, Wendelin Werner.

I received comments on versions of this draft at various stages by Henry Jackson, Benoit Laslier, Ellen Powell, Gourab Ray, Mark Sellke, Huy Tran, Fredrik Viklund. I am grateful for their input which helped me correct some typos and other minor problems, and emphasise some of the subtle aspects of the arguments.

Thanks also to Jason Miller and to Jérémie Bettinelli and Benoit Laslier for letting me use some of their beautiful simulations which can be seen on the cover and in Section **??**.

1 Definition and properties of the GFF

1.1 Discrete case*

The discrete case is included here only for the purpose of helping the intuituion when we come to the continuous case. Proofs will therefore be left as exercises.

Consider a finite graph $G = (V, E)$ (possibly with weights on the edges, but we'll ignore that for simplicity of the computations). Let ∂ be a distinguished set of vertices, called the boundary of the graph, and set $\hat{V} = V \setminus \partial$. Let X_n be a random walk on G. Let P denote the transition matrix and $d(x) = \deg(x)$ be the degree, which is a reversible invariant measure for X; let τ be the first hitting time of ∂.

Definition 1.1 (Green function)**.** *The Green function $G(x, y)$ is defined for $x, y \in V$ by putting*

$$G(x, y) = \frac{1}{d(y)} \mathbb{E}_x \left(\sum_{n=0}^{\infty} \mathbf{1}_{\{X_n = y; \tau > n\}} \right).$$

In other words this is the expected time spent at y, started from x, until hitting the boundary, but normalised by $1/d(y)$, where $d(y) = \deg(y)$ The following properties of the Green function are left as an exercise:

Proposition 1.2. *The following hold:*

1. *Let $\hat{P}$ denote the restriction of P to $\hat{V}$. Then $(I - \hat{P})^{-1}(x, y) = G(x, y)d(y)$ for all $x, y \in \hat{V}$.*
2. *G is a symmetric nonnegative semidefinite function. That is, one has $G(x, y) = G(y, x)$ and if $(\lambda_x)_{x \in V}$ is a vector then $\sum_{x,y \in V} \lambda_x \lambda_y G(x, y) \geq 0$. Equivalently, all its eigenvalues are nonnegative.*
3. *$G(x, \cdot)$ is discrete harmonic in $\hat{V} \setminus \{x\}$, and more precisely $\Delta G(x, \cdot) = -\delta_x(\cdot)$.*

Here, as usual, $\Delta f(x) = \sum_{y \sim x} (1/d(x))(f(y) - f(x))$ is the discrete Laplacian.

Definition 1.3 (Discrete GFF)**.** *The Discrete GFF is the centered Gaussian vector $(h(x))_{x \in V}$ with covariance given by the Green function G.*

Remark 1.4. This definition is justified. Indeed suppose $C(x, y)$ is a given function. Then there exists a centered Gaussian field X having covariance matrix C if and only if C is symmetric and a nonnegative semidefinite function (in the sense of property 2 above).

Note that if $x \in \partial$, then $G(x, y) = 0$ for all $y \in V$ and hence $h(x) = 0$ almost surely.

Usually, for Gaussian fields it is better to use the covariance for getting some intuition. However, as it turns out here the joint probability density function of the n vectors is actually more transparent for the intuition.

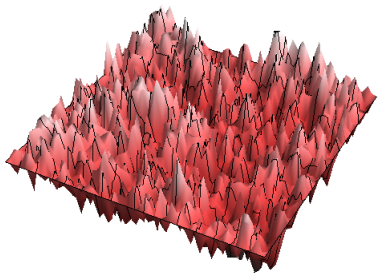

Figure 1: A discrete Gaussian free field

Theorem 1.5 (Law of GFF is Dirichlet energy)**.** *The law of $(h(x))_{x\in V}$ is absolutely continuous with respect to $d\mathbf{x} = \prod_{1\le i\le n} dx_i$, and the joint pdf is proportional to*

$$\mathbb{P}(h(\mathbf{x})) = \frac{1}{Z}\exp\left(-\frac{1}{4}\sum_{x,y\in V:x\sim y}(h(x)-h(y))^2\right)dh(\mathbf{x})$$

where Z is a normalising constant (called the partition function).

Remark 1.6. The way to read this formula is as follows: if $V = \{x_1,\dots,x_n\}$ and $Y_i = h(x_i)$, then $\mathbb{P}((Y_1,\dots,Y_n)\in A) = \int_A \frac{1}{Z}\exp(-\frac{1}{4}\sum(y_i-y_j)^2)\prod_i dy_i$ where the sum is over all i,j such that $x_i\sim x_j$, for any Borel set A such that whenever $\mathbf{y}\in A$ then $y_i = 0$ for indices i corresponding to $x_i\in\partial$.

For a given function $f: V\to\mathbb{R}$, the sum $\sum_{x\sim y}(f(x)-f(y))^2$ is known as the Dirichlet energy, and is the discrete analogue of $\int_D |\nabla f|^2$.

Proof. For a centred Gaussian vector $(Y_1,\dots,Y_n)$ with invertible covariance matrix V, the joint pdf is proportional to $f(y_1,\dots,y_n) = (1/Z)\exp(-\frac{1}{2}y^TV^{-1}y)$ where Z is a normalisation constant.

For us $x\in V$ plays the role of $1\le i\le n$ in the above formula, and $h(x)$ plays the role of y_i. To get a nondegenerate covariance matrix we restrict ourselves to vertices in $\hat{V}$, in which case G is invertible by Proposition 1.2. Hence we wish to compute $h(\hat{\mathbf{x}})^TG^{-1}h(\hat{\mathbf{x}})$ where $h(\hat{\mathbf{x}})$ stands for the vector of values of the field on vertices in $\hat{V}$. Recall that $(I-\hat{P})^{-1}(x,y) = G(x,y)d(y)$ for $x,y\in\hat{V}$. In matrix notations, $(I-\hat{P})^{-1} = GD$ where D is the diagonal matrix (of dimension $|\hat{V}|$) with entries $d(x)$. Therefore, inverting this relation and multiplying by D, $G^{-1} = D(I-\hat{P})$, or $G^{-1}(x,y) = d(x)(I-\hat{P})(x,y)$ for $x,y\in\hat{V}$.

Hence

$$\begin{aligned} h(\hat{\mathbf{x}})^TG^{-1}h(\hat{\mathbf{x}}) &= \sum_{x,y\in\hat{V}} G^{-1}(x,y)h(x)h(y) \\ &= \sum_{x,y\in\hat{V}:x\sim y} -d(x)\hat{P}(x,y)h(x)h(y) + \sum_{x\in\hat{V}} h(x)^2 \end{aligned}$$

$$= - \sum_{x,y \in V : x \sim y} h(x)h(y) + \sum_{x,y \in V : x \sim y} \frac{1}{2}(h(x)^2 + h(y)^2)$$
$$= \sum_{x,y \in V : x \sim y} \frac{1}{2}(h(x) - h(y))^2,$$

where in going from the second to the third line we used the fact that $h(x) = 0$ when $x \in \partial$. This concludes the proof. □

Now, the Dirichlet energy of functions is minimised for harmonic functions. This means that the Gaussian free field is a Gaussian perturbation of harmonic functions: as much as possible, it "tries" to be harmonic (which is a little ironic, given that in the continuum it is not even a function).

This property is at the heart of the Markov property, which is without a doubt the most useful property of the GFF. We will state this without proof, as we will later prove the continuous counterpart (which is very similar).

Theorem 1.7 (Markov property of the discrete GFF)**.** *Fix $U \subset V$. The discrete GFF $h(x)$ can be decomposed as follows:*

$$h = h_0 + \varphi$$

where: h_0 is Gaussian Free Field on U, and ϕ is harmonic in U. Moreover, h_0 and φ are independent.

By a Gaussian Free Field in U we mean that we have set $\partial = V \setminus U$. In other words, what this means is the following: conditional on the values of h outside U, the field can be written as the sum of two terms, one which is an independent, zero boundary GFF, and another which is just the harmonic extension inside U of the boundary values outside U.

1.2 Green function

We will follow a route which is similar to the previous, discrete case. First we need to recall the definition of the Green function. We will only cover the basics here, and readers who want to know more are advised to consult for instance Lawler's book [Law05] which reviews important facts in a very accessible way.

For $D \subset \mathbb{R}^d$ a bounded domain, we define $p_t^D(x, y)$ to be the transition probability of Brownian motion killed when leaving D. In other words, $p_t^D(x, y) = p_t(x, y)\pi_t^D(x, y)$ where $p_t(x, y) = (2\pi t)^{-d/2} \exp(-(|x - y|^2/2t)$, the Gaussian transition probability, and $\pi_t^D(x, y)$ is the probability that a Brownian bridge of duration t remains in D.

Definition 1.8 (Continuous Green function)**.** *The Green function $G(x, y) = G_D(x, y)$ is given by*

$$G(x, y) = \pi \int_0^\infty p_t^D(x, y)dt.$$

[The factor π in front is purely so as to make some later computations in LQG slightly easier]. Note that $G(x,x) = \infty$ for all $x \in D$, since $\pi_t^D(x,x) \to 1$ as $t \to 0$. However $G(x,y) < \infty$ as soon as $x \neq y$ and for instance D is bounded. The Green function can be finite on some unbounded domains.

Example. Suppose $D = \mathbb{H}$ is the upper half plane. Then it is not hard to see that $p_t^{\mathbb{H}}(x,y) = p_t(x,y) - p_t(x,\bar{y})$ by a reflection argument, and hence one can deduce:

$$G_{\mathbb{H}}(x,y) = \log \left| \frac{x-\bar{y}}{x-y} \right|. \tag{1.1}$$

(see Exercise 1 for a hint on the proof).

From now on, we consider only $d = 2$. Then, as shown in this proposition G inherits conformal invariance from that of Brownian motion.

Proposition 1.9. *If $T : D \to D'$ is a conformal map (i.e., holomorphic and one-to-one), $G_{T(D)}(T(x), T(y)) = G_D(x,y)$.*

Proof. Essentially this is just change of variables. The Jacobian term $|T'(y)|^2$ arising from the change of variable just cancels the term $|T'(B_s)|^2$ arising from Itô's formula in the conformal invariance of Brownian motion. More precisely, let ϕ be a test function. Let $x' = T(x)$. Then

$$\int_{D'} G_{D'}(x',y')\phi(y')dy' = \mathbb{E}_{x'}[\int_0^{\tau'} \phi(B'_{t'})dt']$$

where B' is a Brownian motion and τ' is its exit time from D'. On the other hand, the change of variable formula applied to the left hand side gives us, letting $y' = T(y)$ and $dy' = |T'(y)|^2 dy$:

$$\int_{D'} G_{D'}(x',y')\phi(y')dy' = \int_D G_{D'}(T(x),T(y))\phi(T(y))|T'(y)|^2 dy.$$

Now we apply Itô's formula to the right hand side. This allows us to write $B'_{t'} = T(B_{F^{-1}(t)})$ where $F(t) = \int_0^t |T'(B_s)|^2 ds$ for $s \leq \tau$, the first exit time of D by B, and F^{-1} is the cadlag inverse of F. Moreover $\tau' = F^{-1}(\tau)$. Therefore,

$$\begin{aligned}
\mathbb{E}_{x'}[\int_0^{\tau'} \phi(B'_{t'})dt'] &= \mathbb{E}_x[\int_0^{F^{-1}(\tau)} \phi(T(B_{F^{-1}(t)}))dt \\
&= \mathbb{E}_x[\int_0^{\tau} \phi(T(B_s))F'(s)ds] \\
&= \mathbb{E}_x[\int_0^{\tau} \phi(T(B_s))|T'(B_s)|^2 ds] \\
&= \int_D G_D(x,y)\phi(T(y))|T'(y)|^2 dy.
\end{aligned}$$

Hence identifying the left and right hand sides, since the test function ϕ is arbitrary, we conclude that $G_{D'}(T(x), T(\cdot)) = G_D(x,\cdot)$ as distributions. Hence the result follows. □

Together with (1.1) and the Riemann mapping theorem, this allows us to determine G_D in any simply connected proper domain $D \subset \mathbb{C}$. In particular we deduce:

Proposition 1.10. *The following properties hold:*

1. *$G(x, \cdot)$ is harmonic in $D \setminus \{x\}$; and as a distribution $\Delta G(x, \cdot) = -2\pi\delta_x(\cdot)$*
2. *$G(x, y) = -\log(|x - y|) + O(1)$ as $y \to x$.*

Proof. For the first point, observe that harmonic functions stay harmonic under composition by a conformal map. For the second point, use the explicit form of G on $\mathbb{H}$ and the conformal invariance. □

In fact, one can be a bit more precise in the above argumentation:

$$G(x, y) = -\log(|x - y|) + \log R(x; D) + o(1) \tag{1.2}$$

as $y \to x$, where $R(x; D)$ is the **conformal radius** of x in D: that is, $R(x; D) = |f'(0)|$ where f is any conformal map taking $\mathbb{D}$ to D and satisfying $f(0) = x$. (Note that this unambiguously defined). To see this, first note (using the explicit Möbius transform $\varphi(z) = (i - z)/(i + z)$ which maps the upper-half plane to the unit disc) that if $D = \mathbb{D}$ is the unit disc, then

$$G_{\mathbb{D}}(0, z) = \log |z|$$

whence (1.2) is obvious for $\mathbb{D}$ and $x = 0$. Then (1.2) follows immediately in the general case by conformal invariance and definition of the conformal radius.

The conformal radius appears in Liouville quantum gravity in various formulae which will be discussed later on in the course. The reason it shows up in these formulae is usually because of (1.2).

1.3 GFF as a stochastic process

Essentially, as in the discrete case, the GFF is a Gaussian "random function" with mean zero and covariance given by the Green function. However the logarithmic divergence of the Green function on the diagonal means that the GFF cannot be defined pointwise (as the variance would have to be infinite). Instead it is defined as a random distribution or generalised function in the sense of Schwartz[1]. More precisely, we will take the point of view that it assigns values to certain measures with finite Green energy. In doing so we follow the point of view in the two sets of lecture notes [BN11] and [Wer]. The latter in particular contains a great deal more about the relationship between the GFF, SLE, Brownian loop soups and conformally invariant random processes in the plane, which will not be discussed in these notes. The paper by Dubédat [Dub09] is also an excellent source of information regarding basic properties of the Gaussian free field.

[1]This conflicts with the usage of distribution to mean the law of a random variable but is standard and should not cause confusion.

Let D be a bounded domain (or more generally a domain in which the Green function is finite – such a domain is called **Greenian**). Let $\mathcal{M}_+$ denote the set of (nonnegative) measures with compact support in D such that $\int \rho(dx)\rho(dy)G(x,y) < \infty$. Note that due to the logarithmic nature of the Green function blowup on the diagonal, this includes the case where $\rho(x) = f(x)dx$ and f is continuous, but does not include Dirac point masses. Denote $\mathcal{M}$ the set of signed measures, which can be written of the form $\rho = \rho_+ - \rho_-$, where $\rho_\pm \in \mathcal{M}_+$. For test functions $\rho_1, \rho_2 \in \mathcal{M}$, we set

$$\Gamma(\rho_1, \rho_2) := \int_{D^2} G_D(x,y)\rho_1(dx)\rho_2(dy).$$

We also define $\Gamma(\rho) = \Gamma(\rho, \rho)$.

Recall the following elementary facts from measure theory. Let I be an index set. A **stochastic process indexed by** I is just a collection of random variables $(X_i, i \in I)$, defined on some given probability space. The law of the process is a measure on $\mathbb{R}^I$, endowed with the product topology. It is uniquely characterised by its finite-dimensional marginals, via Kolmogorov's extension theorem. Note that in such a setting, it might not be possible to 'observe' simultaneously more than a countable number of random variables. In order to do so we will have to rely on the existence of suitable versions with nice continuity properties (where recall that a version is another stochastic process indexed by the same set and whose finite–dimensional marginals a.s. coincide). The following is both an easy theorem and a definition of the standard GFF on a domain.

Theorem 1.11 (Zero boundary GFF)**.** *There exists a unique stochastic process $(h_\rho)_{\rho\in\mathcal{M}}$, indexed by $\mathcal{M}$, such that for every choice of $\rho_1, \dots, \rho_n$, $(h_{\rho_1}, \dots, h_{\rho_n})$ is a centered Gaussian vector with covariance structure* $\mathrm{Cov}(h_{\rho_i}, h_{\rho_j}) = \Gamma(\rho_i, \rho_j)$.

Definition 1.12. *The process $(h_\rho)_{\rho\in\mathcal{M}}$ is called the Gaussian free field in D (with Dirichlet or zero boundary conditions).*

Proof. We need to check several things:

- the finite-dimensional distributions exist and are uniquely specified.
- they are consistent.

The consistency is an immediate consequence of the Gaussianity of the finite dimensional marginals: indeed, the restriction of a Gaussian vector to a subset of coordinates is still a Gaussian vector with the same covariance structure.

For the first point, symmetry is a consequence of the symmetry of the Green function (which itself follows from the fact that $p_t^D(x,y) = p_t^D(y,x)$.) To check the nonnegative semidefinite character, we need to check that for every $\rho_1, \dots, \rho_n \in \mathcal{M}$, for every $\lambda_1, \dots, \lambda_n \in \mathbb{R}$, we have that

$$\sum_{i,j} \lambda_i\lambda_j\Gamma(\rho_i, \rho_j) \geq 0.$$

However, by linearity the above sum is nothing but $\Gamma(\rho)$ where $\rho = \sum_i \lambda_i \rho_i$. Hence it suffices to prove

$$\Gamma(\rho, \rho) \geq 0; \qquad \rho \in \mathcal{M}. \tag{1.3}$$

This will require a few arguments.

We assume first that ρ is a smooth, compactly supported function. We will need the following very useful fact, known as the **Gauss-Green formula** (integration by parts):

Lemma 1.13. *If f, g are smooth functions, then*

$$\int_D \nabla f \cdot \nabla g = -\int_D f \Delta g + \int_{\partial D} f \frac{\partial g}{\partial n}$$

(Note that this identity extends immediately to the case where say $f \in \mathcal{D}'(D)$ and $g \in \mathcal{D}(D)$. In the absence of Dirichlet boundary conditions, an extra boundary term must be added.)

With this lemma we turn to the proof of (1.3) in the case where ρ is a smooth compactly supported function. First, set

$$f(x) = -\int G_D(x, y)\rho(y)dy.$$

Then $\Delta f = 2\pi\rho$. Hence by the Gauss–Green formula (Lemma 1.13), noting that there are no boundary terms arising in each application of the Gauss–Green formula,

$$\begin{aligned}
\Gamma(\rho) &= \frac{1}{2\pi}\int_x \rho(x) \int_y G(x, y)\Delta_y f(y)dydx && \text{since } \rho(y) = \frac{1}{2\pi}\Delta_y f(y), \\
&= \frac{1}{2\pi}\int_x \rho(x) \int_y \Delta_y G(x, y) f(y)dydx && \text{by integration by parts}, \\
&= -\int_x f(x)\Delta f(x)dx && \text{since } \Delta G(x, \cdot) = -2\pi\delta_x(\cdot) \\
&= \int_D |\nabla f|^2 && \text{by integration by parts}. \qquad (1.4)
\end{aligned}$$

Hence $\Gamma(\rho, \rho) \geq 0$ if $\rho \in \mathcal{D}(D)$. In the general case, observe that if $\rho \in \mathcal{M}$, and if $\rho_\varepsilon(x) = \int_D f_\varepsilon(x - x')\rho(dx')$ with f_ε a smooth approximation of identity (such as the heat kernel $f_\varepsilon(z) = p_\varepsilon(0, z)$), then $\Gamma(\rho_\varepsilon) \to \Gamma(\rho)$. This itself can be seen from uniform integrability assuming $\rho \in \mathcal{M}_+$, which is relatively easy to establish from the assumption. Hence $\Gamma(\rho, \rho) \geq 0$ for all $\rho \in \mathcal{M}$ and so (1.3) is proved. This finishes the proof of the theorem. □

In the rest of this text, we will write (h, ρ) for h_ρ, and we will think of (h, ρ) integrated against ρ. Other boundary conditions than the zero boundary conditions considered here are of great interest in practice. In that case, we make the following definitions. Suppose f is a (possibly random) continuous function on the *conformal boundary* of the domain (equivalent to the Martin boundary of the domain for Brownian motion). Then the GFF with boundary data given by f is the random variable $h = h_0 + \varphi$, where h_0 is an independent Dirichlet GFF, and φ is the harmonic extension of f to D.

Remark 1.14. Note that (h, ρ) is linear in ρ: if $\rho, \rho' \in \mathcal{M}$, and if $\alpha, \beta \in \mathbb{R}$ $(h, \alpha\rho + \beta\rho') = \alpha(h, \rho) + \beta(h, \rho')$, almost surely (which can be seen by checking that the variance and mean of the difference are both zero).

1.4 GFF as a random distribution: Dirichlet energy*

Let $\mathcal{D}(D)$ denote the set of compactly supported, C^∞ functions in D, also known as test functions. The set $\mathcal{D}(D)$ is equipped with a topology in which convergence is characterised as follows. A sequence $f_n \to 0$ in $\mathcal{D}(D)$ if and only if there is a compact set $K \subset D$ such that $\mathrm{supp} f_n \subset K$ for all n and f_n and all its derivatives converge to 0 uniformly on K. A continuous linear map $u : \mathcal{D}(D) \to \mathbb{R}$ is called a *distribution* on D. Thus, the set of distributions on D is the dual space of $\mathcal{D}(D)$. It is denoted by $\mathcal{D}'(D)$ and is given the weak-$*$ topology. Thus $u_n \to u$ in $\mathcal{D}'(D)$ if and only if $u_n(\rho) \to u(\rho)$ for all $\rho \in \mathcal{D}(D)$.

Definition 1.15 (Dirichlet energy, Sobolev space)**.** *For $f, g \in \mathcal{D}(D)$, introduce the Dirichlet inner product:*

$$(f, g)_\nabla := \frac{1}{2\pi} \int_D \nabla f \cdot \nabla g.$$

By definition, the Sobolev space $H_0^1(D)$ is the completion of $\mathcal{D}(D)$ with respect to this inner product; this is the Sobolev space *of index 1, and consists of all distributions which are $L^2(D)$ functions and whose gradient is also an $L^2(D)$ function.*

At this stage we do not know yet that a GFF may be viewed as a random distribution (a random variable in $\mathcal{D}'(D)$). But suppose we did know it. Then for $f \in \mathcal{D}(D)$, the expression $(h, f)_\nabla$ would make sense. Set $\rho = -\Delta f$. Then observe that, by integration by parts (the Gauss–Green formula),

$$(h, f)_\nabla = -\frac{1}{2\pi}(h, \Delta f) = \frac{1}{2\pi}(h, \rho)$$

So the random variable $(h, f)_\nabla$ would be a centered Gaussian random variable with a variance $(2\pi)^{-2}\Gamma(\rho)$, with $\rho = -\Delta f$. Using the calculation (1.4) in the argument for the positive definite character of the Green function, we deduce

$$\mathrm{Var}(h, f)_\nabla = \|f\|_\nabla^2.$$

By polarisation, the covariance between $(h, f)_\nabla$ and $(h, g)_\nabla$ would hence be given by $(f, g)_\nabla$. (This property is sometimes taken as the definition of the GFF, as in [She07]). This suggests the following construction of a GFF. Let f_n denote an orthonormal basis of $H_0^1(D)$. Since f_n, f_m are orthogonal by definition, the random variables $(h, f_n)_\nabla$ and $(h, f_m)_\nabla$ will be uncorrelated (and hence independent) Gaussian random variables, with unit variances. Thus set X_n to be i.i.d. standard Gaussian random variables, and set

$$h = \sum_n X_n f_n. \tag{1.5}$$

It is not clear a priori that the series converges in $\mathcal{D}'(D)$. (Note that it does not converge a.s. in $H_0^1(D)$, as the square norm would be infinite a.s.) But if $f \in \mathcal{D}(D)$, and in fact more generally if $f \in H_0^1(D)$ and $h_N = \sum_{n=1}^N X_n f_n$, then the series

$$(h_N, f)_\nabla := \sum_{n=1}^N X_n (f_n, f)_\nabla \tag{1.6}$$

converges in $L^2(\mathbb{P})$ and almost surely by the martingale convergence theorem, as $N \to \infty$. Its limit is a Gaussian random variable with variance $\sum_n (f_n, f)_\nabla^2 = \|f\|_\nabla^2$ by Parseval's identity. This defines a random variable which we call $(h, f)_\nabla$ which has the law of a mean zero Gaussian random variable with variance $\|f\|_\nabla^2$, as desired. Hence while the series (1.5) does *not* converge in H_0^1, when we take the inner product with a given $f \in H_0^1$ then this converges almost surely. This is similar to the following familiar fact: if B is a Brownian motion, then dB_t is not a function but nevertheless $\int f(s) dB_s$ makes sense almost surely whenever $f \in L^2$: this is **Itô's isometry**. See Exercise 9 for an isometry associated to the GFF which is the exact analogue of Itô's isometry.

An eigenvalue calculation can be used to show that the series (1.5) converges in the space of distributions (and in fact, in a nice Sobolev space, known as $H_0^{-\varepsilon}(D)$, for any $\varepsilon > 0$. See Exercise 9). Therefore, for any given smooth test function $\rho \in \mathcal{D}(D)$, $\sum_{n=1}^\infty X_n (f_n, \rho)$ converges almost surely and defines a stochastic process (indexed by $\mathcal{D}(D)$) whose law coincides with the restriction of the GFF to $\mathcal{D}(D)$. Hence the restriction of the stochastic process h to $\mathcal{D}(D)$ has a modification which is a random distribution.

Furthermore, a density argument (see below for details) can be used to show that this incudes a stochastic process on all of $\mathcal{M}$, where for a given $\rho \in \mathcal{M}$, (h, ρ) is defined as the limit in probability of (h, ρ_ε) for smooth approximations ρ_ε of ρ. In fact, the series (1.5) can be used directly to concretely evaluate the Gaussian free field (h, ρ), as follows.

Theorem 1.16 (GFF as a random Fourier series)**.** *Let D be a Greenian domain. Set $h_N = \sum_{n=1}^N X_n f_n$ to be truncated series in (1.5). For any $\rho \in \mathcal{M}$,*

$$(h_N, \rho) \to (h, \rho)$$

in $L^2(\mathbb{P})$ (and hence in probability as well) as $N \to \infty$.

Proof. First we spell out the density argument. Let $\rho \in \mathcal{M}$, and let $\rho_\varepsilon = \rho \star \theta_\varepsilon$, where θ_ε is an approximation of the identity (e.g., θ_ε is the heat kernel at time ε). Then it is easy to see that if $\nu_\varepsilon = \rho - \rho_\varepsilon$, and if h is a GFF, then $\mathrm{Var}(h, \nu_\varepsilon) = \int G(x,y) \nu_\varepsilon(dx) \nu_\varepsilon(dy) \to 0$ and hence $(h, \rho_\varepsilon) \to (h, \rho)$ in L^2 and in probability, as desired. Hence if we start from the series $\tilde{h} = \sum X_n f_n$ (which hence coincides in law with h on $\mathcal{D}(D)$), we see that the $(\tilde{h}, \rho_\varepsilon)$ form a Cauchy sequence in L^2 as $\varepsilon \to 0$. Hence $\tilde{h}$ defines a stochastic process on $\mathcal{M}$ and it is easy to check that it is a Gaussian free field too in the sense of Theorem/Definition 1.11.

To finish the proof, we claim that for every $\nu \in \mathcal{M}$, and for every $N \geq 1$,

$$\mathrm{Var}(h_N, \nu) \leq \mathrm{Var}(h, \nu). \tag{1.7}$$

To see this, observe that

$$\mathrm{Var}(h_N, \nu) = \sum_{n=1}^{N} (f_n, \nu)^2$$

When ν is smooth the sum is equal to $\sum_{n=1}^{N} (f_n, f)_\nabla^2$ where $\nu = -\Delta f$ and so converges to $\|f\|_\nabla^2$ which, as we have already seen, is equal to $\iint G(x,y)\nu(x)\nu(y)dxdy$ and hence equal to $\mathrm{Var}(h, \nu)$. Hence (1.7) is proved for ν smooth, and a density argument similar to above shows that (1.7) holds for all $\nu \in \mathcal{M}$.

Applying (1.7) to $\nu = \nu_\varepsilon$ we deduce that (h_N, ν_ε) converges to 0 in L^2 and in probability as $\varepsilon \to 0$, uniformly in N. The result follows. □

Consequently, the restriction of a Gaussian free field h with zero boundary conditions to $\mathcal{D}(D)$ is the unique stochastic process such that for all $f \in \mathcal{D}(D)$, $(h, f)_\nabla$ is a Gaussian centered variable with variance $(f, f)_\nabla$.

1.5 Markov property

We are now ready to state one of the main properties of the GFF, which is the (domain) Markov property. As before, informally speaking it states that, conditionally on the values of h outside of a given subset U, the free field inside U is obtained by adding an independent GFF with Dirichlet boundary conditions, to a function which is the harmonic extension inside U of the values of h outside. Note that it is not obvious even that this harmonic extension is well defined!

Theorem 1.17 (Markov property)**.** *Fix $U \subset D$, open, and take h a GFF (with zero boundary condition on D). Then we may write*

$$h = h_0 + \varphi,$$

where:

1. *h_0 is a zero boundary condition GFF in U, and is zero outside of U.*
2. *φ is harmonic in U.*
3. *h_0 and φ are independent.*

This makes sense whether we view h as a random distribution or a stochastic process indexed by $\mathcal{M}$.

Proof. The key point is the Hilbertian decomposition: $H_0^1(D) = \mathrm{Supp}(U) \oplus \mathrm{Harm}(U)$, where $\mathrm{Harm}(U)$ consists of harmonic functions in U, and $\mathrm{Supp}(U) \equiv H_0^1(U)$ consists of functions supported in U (the closure of compactly supported smooth functions in U). To see why this is true, observe that any function $f \in \mathrm{Supp}(U)$ is orthogonal to any function in $\mathrm{Harm}(U)$, by the Gauss–Green formula. Now let us check that the sum spans the entire space. Let

$f \in H_0^1(D)$, and let f_0 denote the orthogonal projection of f onto $\mathrm{Supp}(U)$. Let $\varphi = f - f_0$, our aim is to show that φ is harmonic in U.

Note that φ is (by definition) orthogonal to $\mathrm{Supp}(U)$. Hence for any test function $\psi \in \mathcal{D}(U)$, we have that $(\varphi, \psi)_\nabla = 0$. By the Gauss–Green formula, we deduce that

$$\int_D (\Delta\varphi)\psi = \int_U (\Delta\varphi)\psi = 0$$

and hence $\Delta\varphi = 0$ as a distribution in U. Elliptic regularity arguments (going beyond the scope of these notes) show that a distribution which is harmonic in the sense of distributions must in fact be a smooth function, harmonic in the usual sense. Therefore $\varphi \in \mathrm{Harm}(U)$ and we are done.

Having this decomposition, we deduce the Markov property in a rather straightforward way. Indeed, let f_n^0 be an orthonormal basis of $\mathrm{Supp}(U)$, and let ϕ_n be an orthonormal basis of $\mathrm{Harm}(U)$. If (X_n, Y_n) is an i.i.d. sequence of independent standard Gaussian random variables, set $h_0 = \sum_n X_n f_n^0$ and $\varphi = \sum_n Y_n \phi_n$. Then the first series converges in $\mathcal{D}'(D)$ since it is a series of a GFF in U. The sum of the two series gives h by construction, and so the second series also converges in the space of distributions. In the space of distributions, the limit of harmonic distributions must be harmonic as a distribution, and hence (by the same elliptic regularity arguments as above) a true harmonic function. This proves the theorem. □

Remark 1.18. It is worth pointing out an important message from the proof above: any orthogonal decomposition of $H_0^1(D)$ gives rise to a decomposition into independent summands of the GFF.

1.6 Conformal invariance

A straightforward change of variable formula shows that the Dirichlet inner product is conformally invariant: if $\varphi : D \to D'$ is a conformal map, then

$$\int_{D'} \nabla(f \circ \varphi^{-1}) \cdot \nabla(g \circ \varphi^{-1}) = \int_D \nabla f \cdot \nabla g$$

Consequently, if (f_n) is an orthonormal basis of $H_0^1(D)$, then $f_n \circ \varphi^{-1}$ defines an orthonormal basis of $H_0^1(D')$. We deduce from the construction of h in the previous theorem the following important property of conformal invariance of the Gaussian Free Field:

Theorem 1.19 (Conformal invariance of the GFF)**.** *If h is a random distribution on $\mathcal{D}'(D)$ with the law of the Gaussian Free Field on D, then the distribution $h \circ \varphi^{-1}$ is a GFF on D'.*

1.7 Circle averages

An important tool in the study of the GFF is the process which gives the average values of the field at small circles centered around a point $x \in D$. Fix $z \in D$ and let $0 < \varepsilon < \mathrm{dist}(z, \partial D)$.

Let $\rho_{z,\varepsilon}$ denote the uniform distribution on the circle of radius ε around z, and note that $\rho_{z,\varepsilon} \in \mathcal{M}$. We set $h_\varepsilon(z) = (h, \rho_{z,\varepsilon})$ The following theorem, which is a consequence of the Kolmogorov-Čentsov continuity theorem (a multidimensional generalisation of the more classical Kolmogorov continuity criterion), will not be proved (see Proposition 3.1 of [DS11] for a proof).

Proposition 1.20 (Circle average is jointly Hölder)**.** *There exists a modification of h such that $(h_\varepsilon(z), z \in D, 0 < \varepsilon < \operatorname{dist}(z, \partial D))$ is jointly Hölder continuous of order $\gamma < 1/2$ on all compacts of $(z \in D, 0 < \varepsilon < \operatorname{dist}(z, \partial D))$.*

In fact it can be shown that this version of the GFF is the same as the version which turns h into a random distribution in Theorem 1.16.

The reason circle averages are so useful is because of the following result.

Theorem 1.21 (Circle average is a Brownian motion)**.** *Let h be a GFF on D. Fix $z \in D$ and let $0 < \varepsilon_0 < \operatorname{dist}(z, \partial D)$. For $t \geq t_0 = \log(1/\varepsilon_0)$, set*

$$B_t = h_{e^{-t}}(z),$$

then $(B_t, t \geq t_0)$ has the law of a Brownian motion started from B_{t_0}.

Proof. Various proofs can be given. For instance, the covariance function can be computed explicitly (check it as an exercise). Alternatively, we can use the Markov property of the GFF to see that B_t must have independent increments. Indeed, suppose $\varepsilon_1 > \varepsilon_2$, and we condition on h outside $B(z, \varepsilon_1)$. Ie, we can write $h = h^0 + \varphi$, where φ is harmonic in $U = B(z, \varepsilon_1)$ and h^0 is a GFF in U. Then $h_{\varepsilon_2}(z)$ is the sum of two terms: $h^0_{\varepsilon_2}(z)$, and the second which is the circle average of the harmonic extension φ on $\partial B(z, \varepsilon_2)$. By harmonicity of φ the latter is nothing else than $h_{\varepsilon_1}(z)$. This gives that the increment can be expressed as

$$h_{\varepsilon_2}(z) - h_{\varepsilon_1}(z) = h^0_{\varepsilon_2}(z)$$

and hence, since h^0 is independent of φ, the increments are independent. Moreover, by applying a the change of scale $w \mapsto (w - z)/\varepsilon_1$ so that the outer circle is mapped to the unit circle, we see that the distribution of the random variable on the right hand side depends only on $r = \varepsilon_2/\varepsilon_1$. This gives stationarity of the increments. Since $h_\varepsilon(z)$ is mean zero and has finite variance, it follows that necessarily B_t is a Brownian motion of some diffusivity.

We can compute the variance explicitly as follows: by the above it suffices to check that if h is a GFF in the unit disc $\mathbb{D}$ and $r < 1$ then $h_r(0)$ has variance $-\log r$. Let ρ denote the uniform distribution on the circle at distance r from the origin. Then

$$\operatorname{Var}(h_r(0)) = \int_{\mathbb{D}^2} G_{\mathbb{D}}(x, y)\rho(dx)\rho(dy).$$

By harmonicity of $G_{\mathbb{D}}(x, \cdot)$ in $\mathbb{D} \setminus \{x\}$ and the mean value property, the above integral is simply

$$\operatorname{Var}(h_r(0)) = \int_{\mathbb{D}} G_{\mathbb{D}}(x, 0)\rho(dx).$$

On the other hand $G_{\mathbb{D}}(x, 0) = -\log|x| = -\log r$ on this circle. This concludes the proof. □

As we zoom in towards a point, the average values of the field oscillate like those of a Brownian motion, giving us a very precise sense in which the field can not be defined pointwise.

1.8 Thick points

An important notion in the study of Liouville Quantum Gravity is that of thick points of the Gaussian Free Field. Indeed, although these points are atypical from the point of view of Euclidean geometry, we will see that they are typical from the point of the associated quantum geometry.

Definition 1.22. *Let h be a GFF in D and let $\alpha > 0$. We say a point $z \in D$ is α-thick if*

$$\liminf_{\varepsilon \to 0} \frac{h_\varepsilon(z)}{\log(1/\varepsilon)} = \alpha.$$

In fact, the lim inf in the definition could be replaced with a lim sup or lim. Note that a given point $z \in D$ is almost surely not thick: the typical value of $h_\varepsilon(z)$ is of order $\sqrt{\log 1/\varepsilon}$ since $h_\varepsilon(z)$ is a Brownian motion at scale $\log 1/\varepsilon$. At this stage, the most relevant result is the following fact, proved by Hu, Miller and Peres [HMP10] (though this was independently and earlier proved by Kahane in the context of his work on Gaussian multiplicative chaos).

Theorem 1.23. *Let $\mathcal{T}_\alpha$ denote the set of α-thick points. Almost surely,*

$$\dim(\mathcal{T}_\alpha) = (2 - \frac{\alpha^2}{2})_+$$

and $\mathcal{T}_\alpha$ is empty if $\alpha > 2$.

Sketch of proof. The upper bound is easy to understand. Indeed, for a given $\varepsilon > 0$,

$$\begin{aligned}\mathbb{P}(h_\varepsilon(z) \geq \alpha \log(1/\varepsilon)) &= \mathbb{P}(\mathcal{N}(0, \log(1/\varepsilon) + O(1)) \geq \alpha \log(1/\varepsilon)) \\ &= \mathbb{P}(\mathcal{N}(0,1) \geq \alpha \sqrt{\log(1/\varepsilon) + O(1)}) \leq \varepsilon^{\alpha^2/2}\end{aligned}$$

using scaling and the standard bound $\mathbb{P}(X > t) \leq \text{const} \times t^{-1} e^{-t^2/2}$ for $X \sim \mathcal{N}(0,1)$. Suppose without loss of generality that $D = (0,1)^2$ is the unit square. Then the expected number of squares of size ε such that the centre z satisfies $h_\varepsilon(z) \geq \alpha \log 1/\varepsilon$ is bounded by $\varepsilon^{-2+\alpha^2/2}$. A slightly more elaborate argument shows that in fact the Minkowski dimension of $\mathcal{T}_\alpha$ is a.s. less than $2 - \alpha^2/2$ and therefore so is, a.s., the Hausdorff dimension. To prove the lower bound, one uses a fairly standard second moment argument, see [HMP10] for details. A simple alternative new proof of this fact is also proposed in Exercise 5 of Chapter 2. □

The value $\alpha = 2$ corresponds informally to the maximum of the free field, and the study of the set $\mathcal{T}_2$ is, informally at least, related to the study of extremes in a branching Brownian motion (see [ABBS13, ABK13]).

1.9 Exercises

Discrete GFF

1. Describe the GFF on a binary tree of depth n, where ∂ is the root of the tree.

2. Using an orthonormal basis of eigenfunctions for $\hat{P}$, show that the partition function Z in Theorem 1.5 is given by

$$Z = \det(I - \hat{P})^{-1/2}$$

3. Prove that the minimiser of the discrete Dirichlet energy is discrete harmonic.

4. Prove the spatial Markov property of the discrete GFF.

Continuous GFF

1. Show that on the upper half pane,

$$G_{\mathbb{H}}(x, y) = \log \left| \frac{x - \bar{y}}{x - y} \right|.$$

   Hint: use that $p_t^{\mathbb{H}}(x, y) = p_t(x, y) - p_t(x, \bar{y})$ by symmetry, and use the formula $e^{-a/t} - e^{-b/t} = t^{-1} \int_a^b e^{-x/t} dx$.

   Deduce the value of G on the unit disc.

2. Let $p_t(x, y)$ be the transition function of Brownian motion on the whole plane. Show that $\pi \int_0^1 p_t(x, y) dt = -\log|x - y| + O(1)$ as $x \to y$. Then use this to argue that $G(x, y) = -\log|x - y| + O(1)$ as $x \to y$, thereby recovering the property 2 of Proposition 1.9.

3. Show that eigenfunctions of $-\Delta$ corresponding to different eigenvalues are orthogonal on $H_0^1(D)$.

4. Compute the Green function in one dimension on $D = (0, \infty)$ and $D = (0, 1)$ respectively. What are the corresponding Gaussian free fields? What is the spatial Markov property in this context?

5. Consider $\tilde{h}_\varepsilon(z)$, the average value of the GFF on a square distance ε from z. Is this a Brownian motion as a function of $t = \log 1/\varepsilon$? If not, how can you modify it so that it becomes a Brownian motion? More generally, what about the average of the field on a scaled contour $\varepsilon\lambda$, where λ is a piecewise smooth loop?

6. Let $\rho \in \mathcal{M}^*$ (recall that by definition this means $\rho \in \mathcal{M}$ and if $f(x) = -\int G(x,y)\rho(dy)$ then $f \in H_0^1(D)$). Show that the series

$$(h,\rho) = \sum_{n=1}^{\infty} X_n(f_n,\rho)$$

converges almost surely and the limit agrees with the value h_ρ coming from the stochastic process definition of h. (Here, as in (1.5), f_n is an orthonormal basis of $H_0^1(D)$ and X_n are i.i.d. standard Gaussian random variables).

Hence deduce that if f_n^ε is the circle average of the smooth function f_n, we have that for all fixed $z \in D$, and for all fixed $\varepsilon > 0$, almost surely:

$$\sum_{n=1}^{N} X_n f_n^\varepsilon(z) \to h_\varepsilon(z).$$

(Hint: when $|x-y| \to 0$, $|\nabla_y G(x,y)| = 1/|x-y| + O(1)$.)

7. Let $D = (0,1)^2$ and let $\phi_{j,k}(x,y) = \sin(j\pi x)\sin(k\pi y)$ for $j,k \geq 1$. Compute the H_0^1 norm of $\phi_{j,k}$ and deduce that an expression for the GFF is

$$h = \sum_{j,k\geq 1} \frac{2}{\pi} \frac{X_{j,k}}{\sqrt{j^2+k^2}} \phi_{j,k}.$$

(This is convenient for computer simulations). How does this relate to Fourier series?

8. Give a proof of Theorem 1.21 based on computing the covariances. [If you are stuck, look in Lemma 2.2 for inspiration].

9. **Itô's isometry for the GFF.** Let e_n denote a set of eigenfunctions of $-\Delta$ on D, orthonormal for the L^2 product, with eigenvalue $\lambda_n \geq 0$. For $s \in \mathbb{R}$, set

$$H^s(D) = \{f \in \mathcal{D}'(D) : \sum_n (f,e_n)^2 \lambda_n^s < \infty\}.$$

This is the Sobolev space of index $s \in \mathbb{R}$. We equip H^s with the inner product: $(f,g)_s = \sum_n (f,f_n)(g,f_n)\lambda_n^2$ which turns H^s into a Hilbert space.
(a) Show that $H^1(D)$ corresponds to our previous definition of $H_0^1(D)$.
(b) Show that the series expansion of h given by (1.5) $h = \sum_n X_n f_n$ converges a.s. in $H^{-\varepsilon}(D)$ for any $\varepsilon > 0$, where f_n is an orthonormal basis of $H_0^1(D)$. (Hint: first check that $\lambda_n \asymp n$ as $n \to \infty$).
(c) For $f \in H^s$, define an operator (Laplacian to a fractional power) $(-\Delta)^s f = \sum_n (f,e_n)\lambda_n^s e_n$. Check that

$$(f,g)_s = ((-\Delta)^{s/2} f, (-\Delta)^{s/2} g).$$

(d) Check that the map $f \in \mathcal{D}(D) \mapsto (h, f)$ extends uniquely to an isometry from $H^{-1}(D)$ onto $L^2(\mathbb{P})$. In particular, the value of h integrated a test function is a.s. well defined as soon as $f \in H^{-1}$. (This is the GFF analogue of Itô's isometry, allowing one to integrate an L^2 function against dB). Explain how that fits into the context of defining circle averages.

10. **Radial decomposition.** Suppose $D = \mathbb{D}$ is the unit disc and h is a GFF. Then show that h can be written as the sum

$$h = h_{\text{rad}} + h_{\text{circ}}$$

where h_{rad} is a radially symmetric function, h_{circ} is a distribution with zero average on each disc, and the two parts are independent. Specify the law of each of these two parts.

2 Liouville measure

In this chapter we begin the study of Liouville Quantum Gravity per se. Informally, this is the random surface whose "Riemann metric tensor" can be expressed as

$$e^{\gamma h(z)}dz.$$

This should be interpreted as follows. Some abstract Riemann surface has been parametrised, after Riemann uniformisation, by a domain of our choice – perhaps the disc, assuming that it has a boundary, or perhaps the unit sphere in three dimensions if it doesn't. In this parametrisation, the conformal structure is preserved: i.e., curves crossing at an angle θ at some point in the domain would also correspond to curves crossing at an angle θ in the original surface.

However, in this parametrisation, the metric and the volume have been distorted. Namely, a small element of volume dz in the domain really corresponds to a small element of volume $e^{\gamma h(z)}dz$ in the original surface. Hence points where h is *very big* (e.g., thick points) correspond in reality to relatively big portions of the surface; while points where h is very low are points which correspond to small portions of the surface. The first points will tend to be typical from the point of view of sampling from the volume measure, while the second points will be where geodesics tend to travel.

Naturally, we will want to take h a Gaussian Free Field, but there is a big problem: the exponential of a distribution is not *a priori* defined. This corresponds to the fact that while h is regular enough to be a *distribution*, so small oscillations do not matter when we average h over regions of space, these oscillations become highly magnified when we take the exponential and they do no longer cancel out.

Instead, we will require a fair amount of work to make sense of a measure $\mu(dz)$ which can be interpreted as $e^{\gamma h(z)}dz$. This will be the Liouville measure. It is defined via an approximation procedure, as follows:

$$\mu_\varepsilon(dz) := e^{\gamma h_\varepsilon(z)}\varepsilon^{\gamma^2/2}dz, \tag{2.1}$$

where $h_\varepsilon(z)$ is a jointly continuous version of the circle average. It is straightforward to see that μ_ε is a (random) Radon measure on D. Our goal will be to prove the following theorem:

Theorem 2.1. *Suppose $\gamma < 2$. Then the random measure μ_ε converges in a.s. weakly to a random measure μ, the* (bulk) Liouville measure*, along the subsequence $\varepsilon = 2^{-k}$. μ has a.s. no atoms, and for any $A \subset D$ open, we have $\mu(A) > 0$ a.s. In fact, $\mathbb{E}\mu(A) = \int_A R(z,D)^{\gamma^2/2}dz \in (0,\infty)$.*

(Recall that in this theorem, $R(z,D)$ is the conformal radius of D seen from z: that is, $R(z,D) = |f'(0)|$ where f is any conformal map taking $\mathbb{D}$ to D and 0 to z.)

In this form, the result is due to Duplantier and Sheffield [DS11]. It could also have been deduced from earlier work of Kahane [Kah85] who used a different approximation procedure, together with results of Robert and Vargas [RV+10] showing universality of the limit with respect to the approximating procedure. (In fact, these two results would have

given convergence in distribution of μ_ε rather than in probability; and hence would not show that the limiting measure μ depends solely on the free field h. However, a strengthening of the arguments of Robert and Vargas due to Shamov [Sha] has recently yielded convergence in probability.) Earlier, Høegh-Krohn [HK71] had introduced a similar model in the context of quantum field theory, and analysed it in the relatively easy L^2 phase when $\gamma < \sqrt{2}$. Here we will follow the elementary approach developed in [Ber15], which works in a more general context. However we will restrict ourselves to the case of interest for the proof of this theorem, which simplifies further some arguments.

2.1 Preliminaries

Before we start the proof of Theorem 2.1 we first observe that this is the right normalisation.

Lemma 2.2. *We have that* $\operatorname{Var} h_\varepsilon(x) = \log(1/\varepsilon) + \log R(x, D)$. *As a consequence,* $\mathbb{E}\mu_\varepsilon(A) = \int_A R(z, D)^{\gamma^2/2} dz \in (0, \infty)$.

Proof. Fix $x \in D$. By definition,

$$\operatorname{Var} h_\varepsilon(x) = \Gamma(\rho_{x,\varepsilon}) = \int \rho_{x,\varepsilon}(dz)\rho_{x,\varepsilon}(dw)G(z, w).$$

For a fixed z, $G(z, \cdot)$ is harmonic on $D \setminus \{z\}$ and so $\int \rho_{x,\varepsilon}(dw)G(z, w) = G(z, x)$ by the mean value property. Therefore,

$$\operatorname{Var} h_\varepsilon(x) = \Gamma(\rho_{x,\varepsilon}) = \int \rho_{x,\varepsilon}(dz)G(z, x).$$

Now, observe that $G(x, \cdot) = -\log|x - \cdot| + \xi(\cdot)$ where $\xi(\cdot)$ is the harmonic extension of $\log(|x-\cdot|)$ from the boundary: indeed the difference has zero boundary condition, is bounded and harmonic in $D \setminus \{z\}$ so must be zero in all of D by uniqueness of solution to the Dirichlet problem among bounded functions. Therefore, by harmonicity of ξ,

$$\operatorname{Var} h_\varepsilon(x) = G_\varepsilon(x) = 2\pi \int_y G(\cdot, y)\rho_{x,\varepsilon}(dy) = \log(1/\varepsilon) + \xi(x)$$

Now it remains to recall that $\xi(x) = \log R(x, D)$, which follows from the fact $G(x, y) = \log(1/|x - y|) + \xi(x) + o(1)$ as $y \to x$ and (1.2). □

We now make a couple of remarks:

1. Not only is the expectation constant, but we have that for each fixed z, $e^{\gamma h_\varepsilon(z)}\varepsilon^{\gamma^2/2}$ forms a martingale as a function of ε. This is nothing but the exponential martingale of a Brownian motion.

2. However, the integral $\mu_\varepsilon(A)$ is NOT a martingale: this is because the underlying filtration in which $e^{\gamma h_\varepsilon(z)}\varepsilon^{\gamma^2/2}$ is a martingale depends on z. If we try to condition on all the information $h_\varepsilon(z), z \in D$, then this is too much information and we lose the martingale property.

2.2 Convergence and uniform integrability in the L^2 phase

We assume without loss of generality that S is bounded and open (for instance, $S = (0,1)^2 \subset D$ for concreteness) and we set $I_\varepsilon = \mu_\varepsilon(S)$. Fix $\gamma \in (0,2)$. We first check that I_ε is integrable. This is easy to check when $\gamma < \sqrt{2}$, but difficulties arise when $\gamma \in [\sqrt{2}, 2)$. (As luck would have it this coincides precisely with the phase which is interesting from the point of view of the geometry of random planar maps). We start with the easy case when $\gamma < \sqrt{2}$; this is the so-called L^2 phase.

Let $\varepsilon > 0$, and let $\delta = \varepsilon/2$. Then we claim that

Proposition 2.3. *We have the estimate $\mathbb{E}((I_\varepsilon - I_\delta)^2) \leq C\varepsilon^{2-\gamma^2}$. In particular, I_ε is a Cauchy sequence in $L^2(\mathbb{P})$ and so converges to a limit in probability. Along $\varepsilon = 2^{-k}$, this convergence occurs almost surely.*

Proof. For ease of notations, let $\bar{h}_\varepsilon(z) = \gamma h_\varepsilon(z) - (\gamma^2/2)\operatorname{Var}(h_\varepsilon(z))$, and let $\sigma(dz) = R(z,D)^{\gamma^2/2}$.

The idea is to say that if we consider the Brownian motions $h_\varepsilon(x)$ and $h_\varepsilon(y)$ (viewed as a function of $\varepsilon = e^{-t}$, then they are (approximately) identical until $\varepsilon \leq |x-y|$, after which they evolve (exactly) independently.

Observe that by Fubini's theorem,

$$\mathbb{E}((I_\varepsilon - I_\delta)^2) = \int_{S^2} \mathbb{E}\left[(e^{\bar{h}_\varepsilon(x)} - e^{\bar{h}_\delta(x)})(e^{\bar{h}_\varepsilon(y)} - e^{\bar{h}_\delta(y)})\right] \sigma(dx)\sigma(dy).$$

By the Markov property, $h_\varepsilon(x) - h_\delta(x)$ and $h_\varepsilon(y) - h_\delta(y)$ are independent as soon as $|x-y| \geq 2\varepsilon$: indeed, we can apply the Markov property in $U = B(x,\varepsilon) \cup B(y,\varepsilon)$, which allows us to write $h = \tilde{h} + \varphi$ where φ is harmonic in U and $\tilde{h}$ is an independent GFF in U. Since U is a disjoint union of two balls in U, the restriction of $\tilde{h}$ to each of these balls is independent. Hence

$$(e^{\bar{h}_\varepsilon(x)} - e^{\bar{h}_\delta(x)})(e^{\bar{h}_\varepsilon(y)} - e^{\bar{h}_\delta(y)}) = e^{\bar{h}_\varepsilon(x)+h_\varepsilon(y)}(1 - e^{\bar{h}_\delta(x)-\bar{h}_\varepsilon(x)})(1 - e^{\bar{h}_\delta(y)-\bar{h}_\varepsilon(y)}).$$

Note that all three terms in this product are independent: indeed, the first is measurable with respect to h outside of U (i.e., depends only on φ), the second term depends only on $\tilde{h}$ in the ball $B(x,\varepsilon)$, and the third term depends only on $\tilde{h}$ in $B(y,\varepsilon)$.

Hence if $|x-y| \geq 2\varepsilon$, then the expectation in the above integral is

$$\begin{aligned} &= \mathbb{E}\left[(e^{\bar{h}_\varepsilon(x)} - e^{\bar{h}_\delta(x)})(e^{\bar{h}_\varepsilon(x)} - e^{\bar{h}_\delta(x)})\right] \\ &= \mathbb{E}(e^{\bar{h}_\varepsilon(x)+\bar{h}_\varepsilon(y)})\mathbb{E}(1 - e^{\bar{h}_\delta(x)-\bar{h}_\varepsilon(x)})\mathbb{E}(1 - e^{\bar{h}_\delta(y)-\bar{h}_\varepsilon(y)}) \end{aligned}$$

But both second and third terms are equal to zero, essentially because of the pointwise martingale property. Therefore the expectation is just 0 as soon as $|x-y| > 2\varepsilon$.

Hence using Cauchy–Schwarz in the case where $|x-y| \leq 2\varepsilon$,

$$\mathbb{E}((I_\varepsilon - I_\delta)^2) \leq \int_{|x-y|\leq 2\varepsilon} \sqrt{\mathbb{E}((e^{\bar{h}_\varepsilon(x)} - e^{\bar{h}_\delta(x)})^2)\mathbb{E}((e^{\bar{h}_\varepsilon(y)} - e^{\bar{h}_\delta(y)})^2)}\sigma(dx)\sigma(dy)$$

$$= C \int_{|x-y|\le 2\varepsilon} \sqrt{\mathbb{E}(e^{2\bar{h}_\varepsilon(x)})\mathbb{E}(e^{2\bar{h}_\varepsilon(y)})}\sigma(dx)\sigma(dy) \tag{2.2}$$
$$\le C \int_{|x-y|\le 2\varepsilon} \varepsilon^{\gamma^2} e^{\frac{1}{2}(2\gamma)^2 \log(1/\varepsilon)}\sigma(dx)\sigma(dy)$$
$$\le C\varepsilon^{2+\gamma^2-2\gamma^2} = C\varepsilon^{2-\gamma^2}.$$

This proves the proposition. □

The moral of this proof is the following: while I_ε is not a martingale (because there is filtration common to all points x such that $e^{\bar{h}_\varepsilon(x)}$ forms a martingale, though it is pointwise a martingale), when it comes to second moment we can use the pointwise martingale to estimate the second moment of the increment $I_\varepsilon - I_\delta$. Only for points x, y which are very close (of order ε) do we get a nontrivial contribution.

We defer the proof of the general case until a bit later (see Section 2.4).

2.3 The GFF viewed from a Liouville typical point

Let h be a Gaussian Free Field on a domain D, let $\gamma < 2$. Let μ be the associated Liouville measure. An interesting question is the following: if z is a random point sampled according to the Liouville measure, normalised to be a probability distribution (this is possible when D is bounded), then what does h look like near the point z? This gives rise to the *rooted measure* in the terminology of [DS11] or to the Peyrière measure in the terminology of Gaussian multiplicative chaos.

We expect some atypical behaviour: after all, for any given fixed $z \in D$, $e^{\gamma h_\varepsilon(z)}\varepsilon^{\gamma^2/2}$ converges a.s. to 0, so the only reason μ could be nontrivial is if there are enough points on which h is atypically big. Of course this leads us to suspect that μ is in some sense carried by certain thick points of the GFF. It remains to identify the level of thickness. As mentioned before, simple back-of-then-envelope calculation (made slightly more rigorous in the next result) suggests that these points should be γ-thick. As we will see, this in fact a simple consequence of Girsanov's lemma: essentially, when we bias h by $e^{\gamma h(z)}$, we shift the mean value of the field by $\gamma G(\cdot, z) = \gamma \log 1/|\cdot - z| + O(1)$, thereby resulting in a γ-thick point.

Theorem 2.4. *Suppose D is bounded. Let z be a point sampled according to the Liouville measure μ, normalised to be a probability distribution. Then, a.s.,*

$$\lim_{\varepsilon\to 0} \frac{h_\varepsilon(z)}{\log(1/\varepsilon)} = \gamma.$$

In other words, z is almost surely a γ-thick point ($z \in \mathcal{T}_\gamma$).

When D is not bounded we can simply say that $\mu(\mathcal{T}_\gamma^c) = 0$, almost surely. In particular, μ is singular with respect to Lebesgue measure, a.s.

Proof. The proof is elegant and simple, but the first time one sees it, it is somewhat perturbing. We require the following important but elementary lemma, which can be seen as a (completely elementary) version of Girsanov's theorem.

Lemma 2.5 (Tilting lemma / Girsanov)**.** *Let $X = (X_1, \dots, X_n)$ be a Gaussian vector under the law $\mathbb{P}$, with mean μ and covariance matrix V. Let $\alpha \in \mathbb{R}^n$ and define a new probability measure by*

$$\frac{d\mathbb{Q}}{d\mathbb{P}} = \frac{e^{\langle \alpha, X\rangle}}{Z},$$

where $Z = \mathbb{E}(e^{\langle \alpha, X\rangle})$ is a normalising constant. Then under $\mathbb{Q}$, X is still a Gaussian vector, with covariance matrix V and mean $\mu + V\alpha$.

It is worth rephrasing this lemma in plain words. Suppose we weigh the law of a Gaussian vector by some linear functional. Then the process remains Gaussian, with unchanged covariances, however the mean is shifted, and the new mean of the variable X_i say, is

$$\mu_i' = \mu_i + \mathrm{Cov}(X_i, \langle \alpha, X\rangle).$$

In other words, the mean is shifted by an amount which is simply the covariance of the quantity we are considering and what we are weighting by.

Proof. Assume for simplicity (and in fact without loss of generality) that $\mu = 0$. It is simple to check it with Laplace transforms: indeed if $\lambda \in \mathbb{R}^n$, then

$$\begin{aligned}\mathbb{Q}(e^{\langle \lambda, X\rangle}) &= \frac{1}{Z}\mathbb{E}(e^{\langle \lambda+\alpha, X\rangle}) \\ &= \frac{1}{e^{\frac{1}{2}\langle \alpha, V\alpha\rangle}} e^{\frac{1}{2}\langle \alpha+\lambda, V(\alpha+\lambda)\rangle} \\ &= e^{\frac{1}{2}\langle \lambda, V\lambda\rangle + \langle \lambda, V\alpha\rangle}\end{aligned}$$

The first term in the exponent $\langle \lambda, V\lambda\rangle$ is the Gaussian term with variance V, while the second term $\langle \lambda, V\alpha\rangle$ shows that the mean is now $V\alpha$, as desired. □

Let $\mathbb{P}$ be the law of the GFF, and let Q_ε denote the joint law of (z, h) where h has the law $\mathbb{P}$ and given h, z is sampled proportionally to μ_ε. That is,

$$Q_\varepsilon(dz, dh) = \frac{1}{Z} e^{\gamma h_\varepsilon(z)} \varepsilon^{\gamma^2/2} dz \mathbb{P}(dh).$$

Here Z is a normalising (nonrandom) constant depending solely on ε. Define also $Q(dz, dh) = \mu_h(dz)\mathbb{P}(dh)$ where by μ_h we mean the Liouville measure which is a.s. defined by h. Note that Q_ε converges to Q weakly.

Note that under the law Q_ε, the marginal law of h is simply

$$Q_\varepsilon(dh) = \frac{1}{Z}\mu_\varepsilon(D)\mathbb{P}(dh) \tag{2.3}$$

so it has the law of a GFF biased by its total mass, and we deduce that $Z = \mathbb{E}(\mu_\varepsilon(D)) = \int_D R(z,D)^{\gamma^2/2}dz$ does not even depend on ε (in fact there are small effects from the boundary which we freely ignore).

Furthermore, the marginal law of z is

$$Q_\varepsilon(dz) = \frac{1}{Z}dz\mathbb{E}(e^{\gamma h_\varepsilon(z)}\varepsilon^{\gamma^2/2}) = \frac{dz}{Z}R(z,D)^{\gamma^2/2}.$$

Here again, the law does not depend on ε and is nice, absolutely continuous with respect to Lebesgue measure. Finally, it is clear that under Q_ε, given h, the conditional law of z is just given by a sample from the Liouville measure.

We will simply reverse the procedure, and focus instead on the *conditional distribution of* h given z. Then by definition,

$$Q_\varepsilon(dh|z) = \frac{1}{Z(z)}e^{\gamma h_\varepsilon(z)}\varepsilon^{\gamma^2/2}\mathbb{P}(dh).$$

In other words, the law of the Gaussian field h has been reweighted by an exponential linear functional. By Girsanov's lemma, we deduce that under $Q_\varepsilon(dh|z)$, h is a field with same covariances and a nonzero mean at point w given by $\gamma\,\mathrm{Cov}(h(w),\gamma h_\varepsilon(z)) = \gamma\log(1/|w-z|) + O(1)$. In other words, a logarithmic singularity of strength γ has been introduced at the point z in the mean.

Now, taking the limit as $\varepsilon \to 0$, we find that under $Q(dh|z)$, a.s.,

$$\lim_{\delta\to 0}\frac{h_\delta(z)}{\log(1/\delta)} = \gamma,$$

so $z \in \mathcal{T}_\gamma$, a.s. as desired. We conclude the proof of the theorem by observing that the marginal laws $Q(dh)$ and $\mathbb{P}(dh)$ are mutually absolutely continuous with respect to one another, so any property which holds a.s. under Q holds also a.s. under $\mathbb{P}$. (This absolute continuity follows simply from the fact that $\mu(S) \in (0,\infty), \mathbb{P}-$a.s.) □

2.4 General case*

To address the difficulties that arise when $\gamma \geq \sqrt{2}$, we proceed as follows. Roughly, we claim that the second moment of I_ε blows up because of rare points which are *too thick* and which do not contribute to the integral in an a.s. sense, but inflate the value of the second moment. So we will remove these points by hand. To see which points to remove, a small back-of-the-envelope calculation with Girsanov's lemma suggests that typical points will be γ-thick.

Let $\alpha > 0$ be fixed (it will be chosen $> \gamma$ and very close to γ soon). We define a good event $G_\varepsilon^\alpha(x) = \{h_\varepsilon(x) \leq \alpha\log(1/\varepsilon)\}$. This is the good event that the point x is not too thick at scale ε.

Lemma 2.6 (Liouville points are no more than γ-thick)**.** *For $\alpha > \gamma$ we have*

$$\mathbb{E}(e^{\bar{h}_\varepsilon(x)}1_{G_\varepsilon^\alpha(x)}) \geq 1 - p(\varepsilon)$$

where the function p may depend on α and for a fixed $\alpha > \gamma$, $p(\varepsilon) \to 0$ as $\varepsilon \to 0$, polynomially fast. The same estimate holds if $\bar{h}_\varepsilon(x)$ is replaced with $\bar{h}_{\varepsilon/2}(x)$.

Proof. Note that

$$\mathbb{E}(e^{\bar{h}_\varepsilon(x)}\mathbf{1}_{\{G_\varepsilon^\alpha(x)\}}) = \tilde{\mathbb{P}}(G_\varepsilon^\alpha(x)), \text{ where } \frac{d\tilde{\mathbb{P}}}{d\mathbb{P}} = e^{\bar{h}_\varepsilon(x)}.$$

By Girsanov's lemma, under $\tilde{\mathbb{P}}$, the process $X_s = h_{e^{-s}}(x)$ has the same covariance structure as under $\mathbb{P}$ and its mean is now $\gamma \operatorname{Cov}(X_s, X_t) = \gamma s + O(1)$ for $s \le t$. Hence it is a Brownian motion with drift γ, and the lemma follows from the fact that such a process does not exceed αt at time t with high probability when t is large (and the error probability is exponential in t, or polynomial in ε, as desired).

Changing ε into $\varepsilon/2$ means that the drift of X_s is $\gamma s + O(1)$ over a slightly larger interval of time, namely until time $t + \log 2$. In particular the same argument as above shows that the same estimate holds for $\bar{h}_{\varepsilon/2}(x)$ as well. □

We therefore see that points which are more than γ-thick do not contribute significantly to I_ε in expectation and can therefore be safely removed. We therefore fix $\alpha > \gamma$ and introduce:

$$J_\varepsilon = \int_S e^{\bar{h}_\varepsilon(x)}\mathbf{1}_{G_\varepsilon(x)}\sigma(dx); \quad J'_{\varepsilon/2}(x) = \int_S e^{\bar{h}_{\varepsilon/2}(x)}\mathbf{1}_{G_\varepsilon(x)}\sigma(dx) \tag{2.4}$$

with $G_\varepsilon(x) = G_\varepsilon^\alpha(x)$. Note that a consequence of Lemma 2.6 is that $\mathbb{E}(|I_\varepsilon - J_\varepsilon|) \le p(\varepsilon)|S|$ tends to zero and $\mathbb{E}(|I_{\varepsilon/2} - J'_{\varepsilon/2}|) \le p(\varepsilon)|S|$ also tends to zero.

Lemma 2.7. *We have the estimate $\mathbb{E}((J_\varepsilon - J'_{\varepsilon/2})^2) \le \varepsilon^r$ for some $r > 0$. In particular, I_ε is a Cauchy sequence in L^1 and so converges to a limit in probability. Along $\varepsilon = 2^{-k}$, this convergence occurs almost surely.*

Proof. The proof of this lemma is virtually identical to the case of the L^2 phase (see Proposition 2.3. The key observation was that if $|x - y| \ge 2\varepsilon$, then the increments $h_\varepsilon(x) - h_{\varepsilon/2}(x)$ and $h_\varepsilon(y) - h_{\varepsilon/2}(y)$ were independent of each other, but they are in fact also independent of $\mathcal{F}$, the σ-algebra generated by the values of h outside balls of radius ε around both x and y. Hence, since the events $G_\varepsilon(x)$ and $G_\varepsilon(y)$ are both measurable with respect to $\mathcal{F}$, we deduce from that proof (see (2.2)) that

$$\mathbb{E}((J_\varepsilon - J'_{\varepsilon/2})^2) \le C \int_{|x-y|\le 2\varepsilon} \sqrt{\mathbb{E}(e^{2\bar{h}_\varepsilon(x)}\mathbf{1}_{G_\varepsilon(x)})\mathbb{E}(e^{2\bar{h}_\varepsilon(y)}\mathbf{1}_{G_\varepsilon(y)})}\sigma(dx)\sigma(dy).$$

Now,

$$\begin{aligned}\mathbb{E}(e^{2\bar{h}_\varepsilon(x)}\mathbf{1}_{G_\varepsilon(x)}) &\le \mathbb{E}(e^{2\bar{h}_\varepsilon(x)}\mathbf{1}_{\{h_\varepsilon(x)\le \alpha\log(1/\varepsilon)\}}) \\ &\le O(1)\varepsilon^{-\gamma^2}\mathbb{Q}(h_\varepsilon(x) \le \alpha \log 1/\varepsilon)\end{aligned}$$

where by Girsanov's lemma, under $\mathbb{Q}$, $h_\varepsilon(x)$ is a normal random variable with mean $2\gamma \log(1/\varepsilon)+O(1)$ and variance $\log 1/\varepsilon + O(1)$. Hence

$$\mathbb{Q}(h_\varepsilon(x) \leq \alpha \log 1/\varepsilon) \leq O(1) \exp(-\frac{1}{2}(2\gamma-\alpha)^2 \log 1/\varepsilon).$$

Hence

$$\mathbb{E}((J_\varepsilon - J'_{\varepsilon/2})^2) \leq O(1)\varepsilon^{2-\gamma^2}\varepsilon^{\frac{1}{2}(2\gamma-\alpha)^2}.$$

Again, choosing $\alpha > \gamma$ sufficiently close to γ ensures that the bound on the right hand side is at most $O(1)\varepsilon^r$ for some $r > 0$, as desired. This finishes the proof of the lemma, hence also of Theorem 2.1 in the general case $\gamma < 2$. □

2.5 The phase transition in Gaussian multiplicative chaos

The fact that the Liouville measure $\mu = \mu_\gamma$ is supported on the γ-thick points, $\mathcal{T}_\gamma$, is very helpful to get a clearer picture of Gaussian multiplicative chaos (Kahane's general theory of measures of the form $e^{\gamma X(z)}dz$ where X is a log-correlated Gaussian field).

Indeed recall that $\dim(\mathcal{T}_\gamma) = (2-\gamma^2/2)_+$, and $\mathcal{T}_\gamma$ is empty if $\gamma > 2$. The point is that $\mu = \mu_\gamma$ does not degenerate *because* there are thick points to support it. Once $\gamma > 2$ there are no longer any thick points, and therefore it is in some sense "clear" that μ_γ must degenerate to the zero measure. When $\gamma = 2$ however, $\mathcal{T}_\gamma$ is not empty, and there is therefore a hope to construct a meaningful measure μ corresponding to the *critical Gaussian multiplicative chaos*. Such a construction has indeed been done in two separate papers by Duplantier, Rhodes, Sheffield, and Vargas [DRSV12, DRS$^+$14]. However the normalisation must be done more carefully – see these two papers for details, as well as the more recent preprint by Junnila and Saksman [JS15].

2.6 Conformal covariance

Of course, it is natural to wonder in what way the conformal invariance of the GFF manifests itself at the level of the Liouville measure. As it turns out these measures are not simply conformally invariant - this is easy to believe intuitively, since the total mass of the Liouville measure has to do with total surface area (measured in quantum terms) enclosed in a domain eg via circle packing, and so this must grow when the domain grows.

However, the measure turns out to be conformally covariant: that is, one must include a correction term accounting for the inflation of the domain under conformal map (and hence there will be a derivative term appearing). To formulate the result, it is convenient to make the following notation. Suppose h is a given distribution – perhaps a realisation of a GFF, but also perhaps one of its close relative (eg maybe it has been shifted by a deterministic function), and suppose that its circle average is well defined. Then we define μ_h to be the measure, if it exists, given by $\mu_h(dz) = \lim_{\varepsilon\to 0} e^{\gamma h_\varepsilon(z)}\varepsilon^{\gamma^2/2}dz$. Of course, if h is just a GFF, then μ_h is nothing else but the measure we have constructed in the previous part. If h can be written as $h = h_0 + \varphi$ where φ is deterministic and h_0 is a GFF, then $\mu_h(dz) = e^{\gamma\varphi(z)} \cdot \mu_{h_0}(dz)$ is absolutely continuous with respect to the Liouville measure μ_{h_0}.

Theorem 2.8 (Conformal covariance of Liouville measure)**.** *Let $f : D \to D'$ be a conformal map, and let h be a GFF in D. Then $h' = h \circ f^{-1}$ is a GFF in D' and we have*

$$\begin{aligned}\mu_h \circ f^{-1} &= \mu_{h\circ f^{-1}+Q\log|(f^{-1})'|} \\ &= e^{\gamma Q \log|(f^{-1})'|}\mu_{h'},\end{aligned}$$

where

$$Q = \frac{\gamma}{2} + \frac{2}{\gamma}.$$

In other words, pushing forward the Liouville measure μ_h by the map f, we get a measure which is absolutely continuous with respect to Liouville measure on D', having density $e^{\gamma\psi}$ with $\psi(z) = Q\log|(f^{-1})'(z)|$, $z \in D'$.

Informal proof. The reason for this formula may be understood quite easily. Indeed, note that $\gamma Q = \gamma^2/2 + 2$. When we use the map f, a small circle of radius ε is mapped *approximately* into a small circle of radius $\varepsilon' = |f'(z)|\varepsilon$ around $f(z)$. So $e^{\gamma h_\varepsilon(z)}\varepsilon^{\gamma^2/2}dz$ approximately corresponds to

$$e^{\gamma h'_{|f'(z)|\varepsilon}(z')}\varepsilon^{\gamma^2/2}\frac{dz'}{|f'(z)|^2}$$

by the usual change of variable formula. This can be rewritten as

$$e^{\gamma h'_{\varepsilon'}(z')}(\varepsilon')^{\gamma^2/2}\frac{dz'}{|f'(z)|^{2+\gamma^2/2}}$$

Letting $\varepsilon \to 0$ we get, at least heursitically speaking, the desired result. □

Proof of Theorem 2.8. Of course, the above heuristics is far from a proof, and the main reason is that $h_\varepsilon(z)$ is not a very well-behaved approximation of h under conformal maps. Instead, one uses a different approximation of the GFF, using the orthonormal basis of $H_0^1(D)$ (which is conformally invariant). In view of this, we make the following definition: suppose $h = \sum_n X_n f_n$, where X_n are i.i.d. standard normal random variables, and f_n is an orthonormal basis of $H_0^1(D)$. Set $h^N(z) = \sum_{i=1}^N X_i f_i$ to be the truncated series, and define

$$\mu^N(S) = \int_S \exp\left(\gamma h^N(z) - \frac{\gamma^2}{2}\operatorname{Var}(h^N(z))\right)\sigma(dz)$$

where recall that $\sigma(dz) = R(z,D)^{\gamma^2/2}dz$. Note that $\mu^N(S)$ has the same expected value as $\mu(S)$. Furthermore, $\mu^N(S)$ is a nonnegative martingale with respect to the filtration $(\mathcal{F}_N)$ generated by (X_N), so has an almost sure limit which we will call $\mu^*(S)$.

Lemma 2.9. *Almost surely, $\mu^*(S) = \mu(S)$.*

Proof. When we take the circle averages of the series we obtain

$$h_\varepsilon = h_\varepsilon^N + h'_\varepsilon$$

where h'_ε is independent from h^N, and h^N_ε denotes the circle average of the function h^N. Hence

$$\varepsilon^{\gamma^2/2} e^{\gamma h_\varepsilon(z)} = e^{\gamma h^N_\varepsilon(z)} \varepsilon^{\gamma^2/2} e^{\gamma h'_\varepsilon(z)}.$$

Consequently, integrating over S and taking the conditional expectation given $\mathcal{F}_n$, we obtain

$$\mathbb{E}(\mu_\varepsilon(S)|\mathcal{F}_n) = \mu^N_\varepsilon(S) := \int_S \exp\left(\gamma h^N_\varepsilon(z) - \frac{\gamma^2}{2}\operatorname{Var}(h^N(z))\right)\sigma(dz).$$

When $\varepsilon \to 0$, the right hand side converges to $\mu^N(S)$ as h^N is a nice smooth function. Consequently,

$$\mu^N(S) = \lim_{\varepsilon\to 0} \mathbb{E}(\mu_\varepsilon(S)|\mathcal{F}_N).$$

By Fatou's lemma, the right hand side is greater or equal to $\mathbb{E}(\mu(S)|\mathcal{F}_n)$. Hence

$$\mu^N(S) \geq \mathbb{E}(\mu(S)|\mathcal{F}_N).$$

We can take $N \to \infty$ and deduce from the martingale convergence theorem (and the fact that $\mu(S)$ is measurable with respect to $\mathcal{F}_\infty$) that

$$\mu^*(S) \geq \mu(S).$$

However, from Fatou's lemma again, we know that

$$\mathbb{E}(\mu^*(S)) \leq \lim_{N\to\infty} \mathbb{E}(\mu^N(S)) = \int_D R(z,D)^{\gamma^2/2} dz.$$

We also know (by uniform integrability of $\mu_\varepsilon(S)$) that

$$\mathbb{E}(\mu(S)) = \int_D R(z,D)^{\gamma^2/2} dz.$$

Consequently we see that $\mathbb{E}(\mu^*(S)) \leq \mathbb{E}(\mu(S))$. The only possibility is that $\mathbb{E}(\mu(S)) = \mathbb{E}(\mu^*(S))$ and therefore also $\mu(S) = \mu^*(S)$, as desired. □

To finish the proof of conformal covariance (Theorem 2.8) we now simply recall that if f_n is an orthonormal basis of $H^1_0(D)$ then $f_n \circ f^{-1}$ gives an orthonormal basis of $H^1_0(D')$. Hence if $h' = h \circ f^{-1}$, then its truncated series h'_N can also simply be written as $h'_N = h^N \circ f^{-1}$. Thus, consider the measure μ^N and apply the map f. We obtain a measure $\tilde{\mu}'_N$ in D' such that

$$\begin{aligned}\tilde{\mu}'_N(D') &= \int_{D'} \exp\{\gamma h^N(f^{-1}(z')) - \frac{\gamma^2}{2}\operatorname{Var}(h^N(f^{-1}(z')))\} R(f^{-1}(z'),D)^{\gamma^2/2} \frac{dz'}{|f'(f^{-1}(z'))|^2} \\ &= \int_{D'} d\mu'_N(z') e^{(2+\gamma^2/2)|(f^{-1})'(z')|},\end{aligned}$$

where $d\mu'_N$ is the approximating measure to $\mu_{h'}$ in D'. (The second identity is justified by properties of the conformal radius). Letting $N \to \infty$, and recalling that $d\mu'_N$ converges to $d\mu_{h'}$ by the previous lemma, we obtain the desired statement of conformal covariance. This finishes the proof of Theorem 2.8. □

2.7 Random surfaces

Essentially, we want to consider the surfaces enoded by μ_h and by $\mu_h \circ f^{-1}$ to be the same. This idea, and the covariance formula, allowed Duplantier and Sheffield to formulate a mathematically rigorous notion of random surfaces. Define an equivalence relation between (D_1, h_1) and (D_2, h_2) if there exists $f : D_1 \to D_2$ a conformal map such that $h_2 = h_1 \circ f^{-1} + Q \log |(f^{-1})'|$. It is easy to see that this is an equivalence relation.

Definition 2.10. *A (random) surface is a pair (D, h) consisting of a domain and a (random) distribution $h \in \mathcal{D}'(D)$, where the pair is considered modulo the above equivalence relation.*

Interesting random surfaces arise, among other things, when we sample a point according to Liouville measure (either in the bulk, or on the boundary when the free field has a nontrivial boundary behaviour), and we 'zoom in' near this point. Roughly speaking, these are the *quantum cones* and *quantum wedges* introduced by Sheffield in [She10a]. A particular kind of wedges will be studied in a fair amount of details in these notes later on (see Theorem 5.20).

2.8 Exercises

1. Explain why Lemma 2.7 and 2.6 imply uniform integrability of $\mu_\varepsilon(S)$.

2. Let μ_n be a random sequence of Borel measures an open set $D \subset \mathbb{R}^d$. Suppose that for each open $S \subseteq D$, $\mu_n(S)$ converges in probability to a finite limit $\ell(S)$. Explain why ℓ defines uniquely a Borel measure μ and why μ_n converges in probability weakly to μ.

3. Let μ be the Liouville measure with parameter $\gamma < 2$. Use uniform integrability and the Markov property of the GFF to show that $\mu(S) > 0$ a.s.

4. How would you normalise $e^{\gamma h_\varepsilon(z)}$ if you are just aiming to define the Liouville measure on some segment contained in D? Show that with this normalisation you get a nondegenerate limit. What is the conformal covariance in this case?

5. Recall the events $G_\varepsilon(z) = \{h_r(z) \leq \alpha \log 1/r, \text{ for all } r \in [\varepsilon_0, \varepsilon]\}$ from the proof of uniform integrability of the Liouville measure in the general case. Show that if $d < 2 - \gamma^2/2$, then

$$\mathbb{E}\left[\int_{S^2} \frac{1}{|x-y|^d} e^{\bar{h}_\varepsilon(x)} \mathbf{1}_{G_\varepsilon(x)} \sigma(dx)\ e^{\bar{h}_\varepsilon(y)} \mathbf{1}_{G_\varepsilon(y)} \sigma(dy)\right] \leq C < \infty$$

   where C does not depend on ε. Deduce that

$$\dim(\mathcal{T}_\gamma) \geq 2 - \gamma^2/2.$$

   Conclude with a proof of Theorem 1.23.

6. Write carefully the continuity argument of Q_ε in the proof of Theorem 2.4. Under the rooted measure $Q(dh|z)$, what are the fluctuations of $h_\varepsilon(z)$ as $\varepsilon \to 0$?

   Show that if h is a GFF and z is sampled according to the Liouville measure, then given z, h is absolutely continuous with respect to $\tilde{h}(\cdot) + \gamma \log(1/|\cdot - z|)$, where $\tilde{h}$ is a Gaussian free field.

3 The KPZ relation

The KPZ formula, named after Knizhnik–Polyakov–Zamolodchikov [KPZ88], is a far-reaching identity, relating the 'size' (dimension) of a given set $A \subset \mathbb{R}^2$ from the point of view of standard Euclidean geometry, to its counterpart from the point of view of the random geometry induced by a Gaussian free field h, via the formal definition of the metric $e^{\gamma h(z)}dz$.

As I will discuss, this is fundamentally connected with the question of computing critical exponents of models of statistical physics for which it is believed that there is a conformally invariant scaling limit – we will look for instance at the computation of percolation exponents in this manner.

3.1 Scaling exponents; KPZ theorem

To formulate it, we need the notion of *scaling exponent* of a set A. Fix A deterministic for the moment, or possibly random but (and this is crucial) independent of the GFF. Suppose $A \subset \mathbb{D}$ for instance. The scaling exponent of A, tells us how likely it is, if we pick a point z uniformly in $\mathbb{D}$, that z falls within distance ε of the set A.

Definition 3.1. *The (Euclidean) scaling exponent of A is the limit, if it exists, of the quantity*

$$x = \lim_{\varepsilon \to 0} \frac{\log \mathbb{P}(B(z,\varepsilon) \cap A \neq \emptyset)}{\log(\varepsilon^2)}$$

We need to make a few comments about this definition.

1. First, this is equivalent to saying that the volume of A_ε decays like ε^{2x}. In other words, A can be covered with $\varepsilon^{-(2-2x)}$ balls of radius ε, and hence typically the Hausdorff dimension of A is simply

$$\dim(A) = 2 - 2x = 2(1-x).$$

   In particular, note that $x \in [0,1]$ always; $x = 0$ means that A is practically the full space, $x = 1$ means it is practically empty.

2. In the definition we chose to divide by $\log(\varepsilon^2)$ because ε^2 is the volume, in Euclidean geometry on $\mathbb{R}^2$, of a ball of radius ε. In the quantum world, we would need to replace this by the Liouville area of a ball of radius ε.

We therefore make the following definition – to be taken with a grain of salt, as we shall allow ourselves to consider slightly different variants. Since we do not have a metric to speak about, we resort to the following expedient: for $z \in \mathbb{D}$, we call $B^\delta(z)$ the quantum ball of mass δ, the Euclidean ball centered at z whose radius is chosen so that its Liouville area is precisely δ. (In [DS11], this is called the *isothermal* ball of mass δ at z).

Definition 3.2. *Suppose z is sampled from the Liouville measure μ, normalised to be a probability distribution. The (quantum) scaling exponent of the (possibly random subset) $A \subset D$ is the limit, if it exists, of the quantity*

$$\Delta = \lim_{\delta \to 0} \frac{\log \mathbb{P}(B^\delta(z) \cap A \neq \emptyset)}{\log(\delta)}.$$

Again, this calls for a few parallel comments.

1. The metric dimension (i.e., the Hausdorff dimension) of $\mathbb{D}$ equipped with the random geometry induced by $e^{\gamma h(z)}$ is not currently well defined, and even if there was a metric, it would be unclear what its value is. Of course there is one exception, which is the case when $\gamma = \sqrt{8/3}$, corresponding to the limit of uniform random planar maps, in which case the dimension is known explicitly and is equal to 4, due to properties of the Brownian map.

2. If we call D this dimension, then as before there is a relation between quantum scaling exponent and Hausdorff dimension of A: namely,

$$\dim_\gamma(A) = D(1 - \Delta).$$

Again, $\Delta \in [0, 1]$ always.

3. As explained above, there is no consensus (even in the physics literature) about the value of D. However, the following formula, proposed by Watabiki, seems to have a reasonable chance of being correct:

$$D(\gamma) = 1 + \frac{\gamma^2}{4} + \sqrt{(1 + \frac{\gamma^2}{4})^2 + \gamma^2}.$$

Simulations are notoriously difficult because of large fluctuations, but this formula matches reasonably well with what is observed numerically, and it has the desirable property that $D(\sqrt{8/3}) = 4$.

4. In the definition of Δ, the probability $\mathbb{P}(B^\delta(z) \cap A \neq \emptyset)$ is averaged over everything: z, h and A itself if it is random. So it is really an annealed probability.

Theorem 3.3 (Hausdorff KPZ formula)**.** *Suppose A is independent of the GFF. If A has Euclidean scaling exponent x, then it has quantum scaling exponent Δ where x and Δ are related by the formula*

$$x = \frac{\gamma^2}{4}\Delta^2 + (1 - \frac{\gamma^2}{4})\Delta. \tag{3.1}$$

A few observations:

1. $x = 0, 1$ if and only if $\Delta = 0, 1$.

2. This is a quadratic relation with positive discriminant so can be inverted.

3. In the particular but important case of uniform random planar maps, $\gamma = \sqrt{8/3}$ so the relation is
$$x = \frac{2}{3}\Delta^2 + \frac{1}{3}\Delta. \tag{3.2}$$

Various versions of this result have now been proved. The above version deals with a notion of dimension which is in expectation, and another one, due to Aru [Aru14] will soon be stated where the notion of dimension is more closely related to Minkowski dimension and is also in expectation. However, almost sure versions also exist – let us mention, in particular, the work of Rhodes and Vargas [RV11] who proved a theorem which is in some ways stronger than the above (it is more robust, and deals with an a.s. notion of dimension) and which appeared simultaneously to the paper of Duplantier and Sheffield [DS11] from which the above theorem is taken. More recently, a version of the KPZ formula was formulated and proved using the so-called Liouville heat kernel [BGRV14], thereby eliminating the need to reference the underlying Euclidean structure in the notion of dimension.

3.2 Applications of KPZ to exponents*

At the discrete level, the KPZ formula can be interpreted as follows. Suppose that a certain subset A within a random map of size N has a size $|A| \approx N^{1-\Delta}$. Then its Euclidean analogue within a box of area N (and thus of side length $n = \sqrt{N}$) occupies a size $|A'| \approx N^{1-x} = n^{2-2x}$, where x, Δ are related by the KPZ formula:
$$x = \frac{\gamma^2}{4}\Delta^2 + (1 - \frac{\gamma^2}{4})\Delta.$$
In particular, observe that the approximate (Euclidean) Hausdorff dimension of A' is then $2 - 2x$, consistent with our definitions.

We now explain how this can be applied to determine the value of certain exponents. Suppose we know the critical exponent of a critical FK loop configuration in random geometry, by which I mean the exponent β such that
$$\mathbb{P}(|\partial L| \geq k) \approx k^{-\beta}$$
where ∂L denotes the interface of boundary of the smallest loop containing a given point, say the origin.

In order to apply KPZ, we consider a planar map with say N faces (or N edges...), equipped with a critical FK loop configuration. We could take for our set A the boundary of the largest cluster. The largest cluster will have a volume of order N, and its boundary will have a length (number of triangles it goes through) of order $N^{1/\beta}$. Then for this set A, we have $\Delta = 1 - 1/\beta$, and hence the value x is known also in the Euclidean world via the KPZ formula. We deduce that the largest cluster in a box of volume N (and of sidelength $n = \sqrt{N}$) will be, for the Euclidean models, $N^{1-x} = n^{2-2x}$. This can be used to "deduce"

the Hausdorff dimension of interfaces in the Euclidean world. Recall that these interfaces are believed to be given (in the scaling limit) by SLE_κ curves, for which the dimension is $1 + \kappa/8$ by a result of Beffara.

Example: percolation interface. The paper [BLR15] gives

$$\mathbb{P}(\partial L \geq k) \approx k^{-4/3}$$

as $k \to \infty$, where k is a uniformly chosen loop. Hence the size of the largest cluster interface in a map of size N will be $N^{3/4}$, so $\Delta = 1 - 3/4 = 1/4$. Applying KPZ, we find $x = \Delta(2\Delta + 1)/3 = 1/8$, so the Hausdorff dimension of the percolation interface should be $2 - 2x = 7/4$. This is consistent with Beffara's result on the Hausdorff dimension of SLE_6, which is $1 + \kappa/8 = 7/4$ as well!

Why is the independence assumption fulfilled?* In applying KPZ it is important to check that A is independent of the underlying free field / metric. In the above cases this might not seem so obvious: after all, the map and the configuration are chosen together. But note that:

– it is the case that given the map, one can sample the FK configuration of loops independently.

– An alternative justification is as follows, and relies just on conformal invariance of the interfaces. We are applying the KPZ relation to these interfaces, after embedding the map conformally. But at the scaling limit, even given the surface – or equivalently, given the realisation of the free field – the law of the conformal embedding of the interfaces must be CLE_κ by conformal invariance. Since this law does not depend on the actual realisation of the free field, we conclude that it is independent of it, at least in the scaling limit.

Hence the use of KPZ is justified.

3.3 Proof in the case of expected Minkowski dimension

We will sketch two proofs of the KPZ relation (Theorem 3.3). The first one is probably the simplest, and is due to Aru [Aru14]. The second is close to part of the argument by Duplantier and Sheffield [KPZ88]. Both are instructive in their own right.

Aru's proof relies on the formula giving the power law spectrum of Liouville measures.

Proposition 3.4 (Scaling relation for Liouville measure)**.** *Let $\gamma \in (0, 2)$. Let $B(r)$ be a ball of radius r that is at distance at least ε_0 from the boundary of the domain D, for some ε_0. Then uniformly over $r \in (0, \varepsilon_0)$, and uniformly over the centre of the ball, and for any $q \in [0, 1]$,*

$$\mathbb{E}(\mu(B(r))^q) \asymp r^{(2+\gamma^2/2)q - \gamma^2 q^2/2},$$

where the implied constants depend only on q, ε_0 and γ.

We defer the proof of this proposition until the end of the section, and content ourselves in saying it is simply the result of approximate scaling relation between $h_r(z)$ and $h_{r/2}(z)$

say – boundary conditions of the GFF make this scaling relations only approximative and hence introduce a few tedious complications.

Also we point out that the fact the exponent in the right hand side is not linear in q is an indication that the geometry has an interesting *multifractal structure*: by definition this means that the geometry is not defined by a single fractal exponent (compare for instance with $\mathbb{E}(|B_t|^q)$ which is proportional to $t^{q/2}$).

Armed with this proposition we will prove a slightly different version of the KPZ theorem, using a different notion than scaling exponents - rather we will use Minkowski dimension. Hence we make the following definitions. Let $\mathcal{S}_n$ denote the nth level dyadic covering of the domain D by squares $S_i, i \in \mathcal{S}_n$. The (Euclidean) 2^{-n}-Minkowski content of A is, by definintion:

$$M_\delta(A; 2^{-n}) = \sum_{i \in \mathcal{S}_n} \mathbf{1}_{\{S_i \cap A \neq \emptyset\}} \operatorname{Leb}(S_i)^\delta$$

The (Euclidean) Minkowsi dimension of A is then defined by

$$d_M(A) = \inf\{\delta : \limsup_{n \to \infty} M_\delta(A, 2^{-n}) < \infty\}$$

and the Minkowski scaling exponent is

$$x_M = 1 - d_M.$$

On the quantum side,

$$M_\delta^\gamma(A, 2^{-n}) = \sum_{i \in \mathcal{S}_n} \mathbf{1}_{\{S_i \cap A \neq \emptyset\}} \ \mu(S_i)^\delta$$

The quantum expected Minkowski dimension is

$$q_M = \inf\{\delta : \limsup_{n \to \infty} \mathbb{E} M_\delta^\gamma(A, 2^{-n}) < \infty\}$$

and the quantum Minkowski scaling exponent is

$$\Delta_M = 1 - q_M.$$

The KPZ relation for the Minkowski scaling exponents is then $x_M = (\gamma^2/4)\Delta_M^2 + (1 - \gamma^2/4)\Delta_M$ (formally this is the same as the relation in Theorem 3.3). Equivalently, this can be rephrased as follows.

Proposition 3.5 (Expected Minkowski KPZ, [Aru14])**.**

$$d_M = (1 + \gamma^2/4)q_M - \gamma^2 q_M^2/4. \tag{3.3}$$

Proof. Fix $d \in (0, 1)$ and let q be such that $d = (1 + \gamma^2/4)q - q^2\gamma^2/4$. Then

$$\mathbb{E}(\sum_{i \in \mathcal{S}_n} \mathbf{1}_{\{S_i \cap A \neq \emptyset\}} \mu(S_i)^q) \asymp \sum_{i \in \mathcal{S}_n} \mathbf{1}_{\{S_i \cap A \neq \emptyset\}} \operatorname{Leb}(S_i)^d$$

and consequently the limsup of the left hand side is infinite if and only if the limsup of the right hand side is infinite. In other words, d_M and q_M satisfy (3.3). (Note here, we are ignoring boundary effects). □

3.4 Sketch of proof of Hausdorff KPZ using circle averages

We sketch the argument used by Duplantier and Sheffield to prove Theorem 3.3. The proof is longer, but the result is also somewhat stronger, as the expected Minkowski dimension is quite a weak notion of dimension.

Informal description of the idea of the proof. We wish to evaluate the probability $\mathbb{P}(B^\delta(z) \cap A \neq \emptyset)$, where z is a point sampled from the Liouville measure, and B^δ is the Euclidean ball of Liouville mass δ around z. Of course the event that this ball intersects A is rather unlikely, since the ball is small. But it can happen because of two reasons. The first one is simply that z lands very close (in the Euclidean sense) to A – this has a cost governed by the Euclidean scaling exponent of A, by definition, since we may think of z as being sampled from the Lebesgue measure and then sampling the Gaussian free field given z, as in the description of the rooted measure. However, it is more economical for z to land relatively further away from z, and instead require that the ball of quantum mass δ have a bigger than usual radius. As the quantum mass of the ball of radius r around z is essentially governed by the size of the circle average $h_r(z)$, which behaves like a Brownian motion plus some drift, we find ourselves computing a large deviation probability for a Brownian motion. The KPZ formula is hence nothing else but the large deviation function for Brownian motion. □

Sketch of proof of Theorem 3.3. Now we turn the informal idea above into more concrete mathematics, except for two approximations which we will not justify. Suppose z is sampled according to the Liouville measure μ. Then we know from Theorem 2.4 (see also the Exercise 6 in that chapter) that the free field is absolutely continuous with respect to $h^0(z) + \gamma \log|\cdot - z| + O(1)$ where h^0 is an independent GFF (see the section on the GFF viewed from a typical Liouville point). Hence the mass of a ball of radius ε is approximately given by

$$\begin{aligned}\mu(B_\varepsilon(z)) &\approx \varepsilon^{\gamma^2/2} e^{\gamma h_\varepsilon(z)} \times \varepsilon^2 \\ &\asymp e^{\gamma(h^0_\varepsilon(z) + \gamma \log 1/\varepsilon)} \varepsilon^{2+\gamma^2/2} \\ &= \varepsilon^{2-\gamma^2/2} e^{\gamma h^0_\varepsilon(z)}. \end{aligned} \tag{3.4}$$

It takes some time to justify rigorously the approximation in (3.4), but the idea is that the field h_ε fluctuates on a spatial scale of size roughly ε. Hence we are not making a big error by pretending that h_ε is constant on $B_\varepsilon(z)$, equal to $h_\varepsilon(z)$. In a way making this precise is the most technical part of the paper [DS11]. We will not go through the arguments which do so, and instead we will see how, assuming it, one is naturally led to the KPZ relation.

Going on an exponential scale which is more natural for circle averages, and calling $B_t = h^0_{e^{-t}}(z)$, we find

$$\log \mu(B_{e^{-t}}(z)) \approx \gamma B_t - (2 - \gamma^2/2)t.$$

We are interested in the maximum radius ε such that $\mu(B_\varepsilon(z))$ will be approximately δ: this will give us the Euclidean radius of the quantum ball of mass δ around z. So let

$$T_\delta = \inf\{t \geq 0 : \gamma B_t - (2 - \gamma^2/2)t \leq \log \delta\}$$

$$= \inf\{t \geq 0 : B_t + (\frac{2}{\gamma} - \frac{\gamma}{2})t \geq \frac{\log(1/\delta)}{\gamma}\}.$$

where the second equality is in distribution. Note that since $\gamma < 2$ the drift is positive, and hence $T_\delta < \infty$ a.s.

Now, recall that if $\varepsilon > 0$ is fixed, the probability that z will fall within (Euclidean) distance ε of A is approximately ε^{2x}. Hence, applying this with $\varepsilon = e^{-T_\delta}$ the probability that $B^\delta(z)$ intersects A is, approximately, given by

$$\mathbb{P}(B^\delta(z) \cap A \neq \emptyset) \approx \mathbb{E}(\exp(-2xT_\delta)),$$

Consequently, we deduce that

$$\Delta = \lim_{\delta \to 0} \frac{\log \mathbb{E}(\exp(-2xT_\delta))}{\log \delta}.$$

For $\beta > 0$, consider the martingale

$$M_t = \exp(\beta B_t - \beta^2 t/2),$$

and apply the optional stopping theorem at the time T_δ (note that this is justified). Then we get, letting $a = 2/\gamma - \gamma/2$,

$$1 = \exp(\beta \frac{\log(1/\delta)}{\gamma}) \mathbb{E}(\exp(-(a\beta + \beta^2/2)T_\delta)).$$

Set $2x = a\beta + \beta^2/2$. Then $\mathbb{E}(\exp(-2xT_A)) = \delta^{\beta/\gamma}$. In other words, $\Delta = \beta/\gamma$, where β is such that $2x = a\beta + \beta^2/2$. Equivalently, $\beta = \gamma\Delta$, and

$$2x = (\frac{2}{\gamma} - \frac{\gamma}{2})\gamma\Delta + \frac{\gamma^2}{2}\Delta^2.$$

This is exactly the KPZ relation. □

3.5 Proof of multifractal spectrum of Liouville measure*

We now explain where the arguments for the multifractal spectrum of μ come from. Recall that we seek to establish: for $q \in [0, 1]$,

$$\mathbb{E}(\mu(B(r))^q) \asymp r^{(2+\gamma^2/2)q - q^2\gamma^2/2}. \tag{3.5}$$

As mentioned before, this is a relation which follows in a fairly straightforward way from scaling arguments. Unfortunately $h_\varepsilon(z)$ is only approximately scale invariant (because of boundary conditions), and hence we need a tool which allows us to do the comparison with a field which enjoys exact scale invariance property. In order to do so, Kahane discovered the following beautiful comparison lemma, which is crucial to his theory of Gaussian multiplicative chaos (but note for instance that it is not needed in the approach in [Ber15]).

Theorem 3.6 (Kahane's convexity inequality)**.** *Let $(X_i), (Y_i)$ be two Gaussian vectors such that:*

$$\mathbb{E}(X_i X_j) \leq \mathbb{E}(Y_i Y_j). \tag{3.6}$$

Then for all nonnegative vectors (p_i), and all real convex function $F : (0, \infty) \to \mathbb{R}$, with at most polynomial growth at infinity, we have

$$\mathbb{E}\left[F\left(\sum_{i=1}^n p_i e^{X_i - \frac{1}{2}\mathbb{E}(X_i^2)}\right)\right] \leq \mathbb{E}\left[F\left(\sum_{i=1}^n p_i e^{Y_i - \frac{1}{2}\mathbb{E}(Y_i^2)}\right)\right]$$

The proof of this theorem can be found in [Kah85]. The proof is not easy and not particularly illuminative (also, it is in French which makes it intimidating for some people). If one is willing to assume that F is also increasing (but this excludes application to concave functions, which will be needed here), then a short and simple proof can be found in [Kah86], which is also sketeched in [RV+10] under slightly different assumptions. However, even then the proof is nothing short of miraculous.

Now, fix a bounded continuous function $\phi : [0, \infty) \to [0, \infty)$ which is a bounded positive definite function and such that $\phi(x) = 0$ if $x \geq 1$. For instance, we choose $\phi(x) = \sqrt{(1 - |x|)_+}$, see the discussion in Example 2.3 in [RV+10]. Define an auxiliary centered Gaussian random field $(X_\varepsilon(x))_{x \in \mathbb{R}^d}$ by specifying its covariance

$$c_\varepsilon(x, y) := \mathbb{E}(X_\varepsilon(x) X_\varepsilon(y)) = \log_+\left(\frac{1}{|x - y| \vee \varepsilon}\right) + \phi\left(\frac{|y - x|}{\varepsilon}\right).$$

Then

$$c_\varepsilon(x, y) - a \leq \mathbb{E}(\gamma h_\varepsilon(x) \gamma h_\varepsilon(y)) \leq c_\varepsilon(x, y) + b \tag{3.7}$$

As a result it will be possible to estimate the moments of $\mu(B(r))$ up to constants by computing those of $\tilde{\mu}(B(r))$, where $\tilde{\mu}$ is obtained from X in a similar way as μ is obtained from h.

The crucial observation about X_ε (and the reason why we introduce it) is that it enjoys an exact scaling relation, as follows:

Lemma 3.7. *For $\lambda < 1$,*

$$(X_{\lambda\varepsilon}(\lambda x))_{x \in B(0,1)} =_d (\Omega_\lambda + X_\varepsilon(x))_{x \in B(0,1)},$$

where Ω_λ is an independent centered Gaussian random variable with variance $\log(1/\lambda)$.

Proof. One easily checks that for all $x, y \in \mathbb{R}^2$ such that $\|x - y\| \leq 1$, $c_{\lambda\varepsilon}(\lambda x, \lambda y) = \log(1/\lambda) + c_\varepsilon(x, y)$. □

We now turn to the proof of (3.5). Define a measure

$$\tilde{\mu}_\varepsilon(z) = \exp(\gamma X_\varepsilon(z) - (\gamma^2/2)\mathbb{E}(X_\varepsilon(z)^2))dz.$$

and note that $\mathbb{E}(X_\varepsilon(x)^2) = \log(1/\varepsilon) + C$ for some constant C. Fix $\varepsilon > 0$, and $\lambda = r < 1$. Then

$$\begin{aligned}\tilde{\mu}_{r\varepsilon}(B(r)) &= \int_{B(r)} Ce^{\gamma X_{r\varepsilon}(z)}(r\varepsilon)^{\gamma^2/2}dz \\ &= Cr^{2+\gamma^2/2}\int_{B(1)} e^{\gamma X_{r\varepsilon}(rw)}\varepsilon^{\gamma^2/2}dw\end{aligned}$$

by changing variables $z = rw$. By Lemma 3.7,

$$\tilde{\mu}_{r\varepsilon}(B(r)) = r^{2+\gamma^2/2}e^{\gamma\Omega_r}\tilde{\mu}'(B(1))$$

where $\tilde{\mu}'$ is a copy of $\tilde{\mu}$ and Ω_r is an independent $\mathcal{N}(0, \log(1/r))$. Raising to the power q, and taking expectations, we get:

$$\begin{aligned}\mathbb{E}(\mu_{r\varepsilon}(B(r)^q)) &\asymp \mathbb{E}(\tilde{\mu}_{r\varepsilon}(B(r)^q)) \\ &= r^{q(2+\gamma^2/2)}\mathbb{E}(e^{\gamma q\Omega_r})\mathbb{E}(\tilde{\mu}_\varepsilon(B(1))^q) \\ &\asymp r^{\xi(q)}\mathbb{E}(\mu_\varepsilon(B(1)^q))\end{aligned}$$

where $\xi(q) = q(2 + \gamma^2/2) + \gamma^2q^2/2$. Now, since $q < 1$ and $\mu_\varepsilon(B(r))$ is uniformly integrable, there is no problem in showing that the expectations converge to their limit as $\varepsilon \to 0$. Hence

$$\mathbb{E}(\mu(B(r))^q) \asymp Cr^{\xi(q)},$$

as desired.

3.6 Exercises

1. By considering the set of thick points or otherwise, show that the KPZ relation does not need to hold if the set A is allowed to depend on the free field.

2. Consider simple random walk on the (infinite) uniform random planar map. If $n \geq 1$, a pioneer point for the walk $(X_1, \dots X_n)$ is a point x such that x is visited at some time $m \leq n$ and is on the boundary of the unbounded component of $G \setminus \{X_1, \dots, X_m\}$. A beautiful theorem of Benjamini and Curien [BC13] shows that by the time simple random walk on an infinite uniform random planar map exits a ball of radius R, it has had $\approx R^3$ pioneer points. Deduce the value of Brownian pioneer points (the set of points B_s such that B_s is on the frontier at time s, i.e., the boundary of the unbounded component of the complement of $B[0, s]$). (The answer is $7/4$, see [LSW00]).

3. Consider a simple random walk (X_n) on the infinite local limit of FK weighted planar map. Try to argue (without being fully rigorous!), using the KPZ relation, that $\mathrm{dist}(o, X_n)$ must be approximately equal to $n^{1/D}$ where D is the dimension of the space. (Hint: the range of Brownian motion must satisfy $\Delta = 0$). In particular, on the UIPT, one conjectures that $\mathrm{dist}(o, X_n) \approx n^{1/4}$.

4. Let $A \subset D$, and let $q \in [0,1]$ say. Show that $\mathbb{E}(\mu(A)^q)$ is a nondecreasing function of $\gamma \in [0,2)$.

5. Use the scaling invariance properties developed in the proof of the multifractal spectrum to show that μ has a.s. no atoms. (You should admit that $\mathbb{E}(\mu(S)^q) < \infty$ for some $q > 1$.) Observe that this also follows from the energy estimate in Exercise 5 in Chapter 2.

4 Statistical physics on random planar maps

4.1 Fortuin–Kesteleyn weighted random planar maps

In this chapter we change our focus from the continuous to the discrete, and describe the model of random planar maps weighted by critical Fortuin–Kasteleyn percolation. These maps can be thought of as canonical discretisations of the surface of interest. First, we make a few definitions.

Planar map, dual map. Recall that a **planar map** is a proper embedding of a (multi) graph with a finite number of edges in the plane $\mathbb{C} \cup \{\infty\}$ (viewed as the Riemann sphere), which is viewed up to orientation preserving homeomorphisms from the sphere to itself. Let $\boldsymbol{m}_n$ be a map with n edges and $\boldsymbol{t}_n$ be the subgraph induced by a subset of its edges and *all* of its vertices. We call the pair $(\boldsymbol{m}_n, \boldsymbol{t}_n)$ a **decorated map**. Let $\boldsymbol{m}_n^\dagger$ denote the *dual map* of $\boldsymbol{m}_n$. Recall that the vertices of the dual map correspond to the faces of $\boldsymbol{m}_n$ and two vertices in the dual map are adjacent if and only if their corresponding faces are adjacent to a common edge in the primal map. Every edge e in the primal map corresponds to an edge $e^\dagger$ in the dual map which joins the vertices corresponding to the two faces adjacent to e. The dual map $\boldsymbol{t}_n^\dagger$ is the graph formed by the subset of edges $\{e^\dagger : e \notin \boldsymbol{t}_n\}$. We fix an oriented edge in the map $\boldsymbol{m}_n$ and define it to be the root edge. With an abuse of notation, we will still write $\boldsymbol{m}_n$ for the rooted map; and we let $\mathcal{M}_n$ be the set of maps with n edged together with one distinguished edge called the root.

Canonical triangulation. Given a subgraph decorated map $(\boldsymbol{m}_n, \boldsymbol{t}_n)$, one can associate to it a set of loops which form the interface between the two clusters. To define it precisely, we consider a refinement of the map $\boldsymbol{m}_n$ which is formed by joining the dual vertices in every face of $\boldsymbol{m}_n$ with the primal vertices incident to that face. We call these edges **refinement edges**. Every edge in $\boldsymbol{m}_n$ corresponds to a quadrangle in its refinement formed by the union of the two triangles incident to its two sides. In fact the refinement of $\boldsymbol{m}_n$ is a quadrangulation and this construction defines a bijection between maps with n edges and quadrangulations with n faces, sometimes called the **Tutte bijetcion**.

Given a subgraph decorated map $(\boldsymbol{m}_n, \boldsymbol{t}_n)$ define the map $(\bar{\boldsymbol{m}}_n, \bar{\boldsymbol{t}}_n)$ to be formed by the union of $\boldsymbol{t}_n, \boldsymbol{t}_n^\dagger$ and the refinement edges. The root edge of $(\bar{\boldsymbol{m}}_n, \bar{\boldsymbol{t}}_n)$ is the refinement edge corresponding to the root edge in $\boldsymbol{m}_n$ oriented towards the dual vertex. The **root triangle** is the triangle to the right of the root edge. It is easy to see that such a map is a triangulation: every face in the refinement of $\boldsymbol{m}_n$ is divided into two triangles either by a primal edge in $\boldsymbol{t}_n$ or a dual edge in $\boldsymbol{t}_n^\dagger$. Thus every triangle in $(\bar{\boldsymbol{m}}_n, \bar{\boldsymbol{t}}_n)$ is formed either by a primal edge and two refinement edges or by a dual edge and two refinement edges. For future reference, we call a triangle in $(\bar{\boldsymbol{m}}_n, \bar{\boldsymbol{t}}_n)$ with a primal edge to be a **primal triangle** and that with a dual edge to be a **dual triangle** (Figure 5).

Loops. Finally, given $\boldsymbol{m}_n, \boldsymbol{t}_n$ we can define the loops induced by $\boldsymbol{t}_n$ as a collection of subgraphs of the refined map $(\bar{\boldsymbol{m}}_n, \bar{\boldsymbol{t}}_n)$. Each loop is simply a collection of triangles "separating"

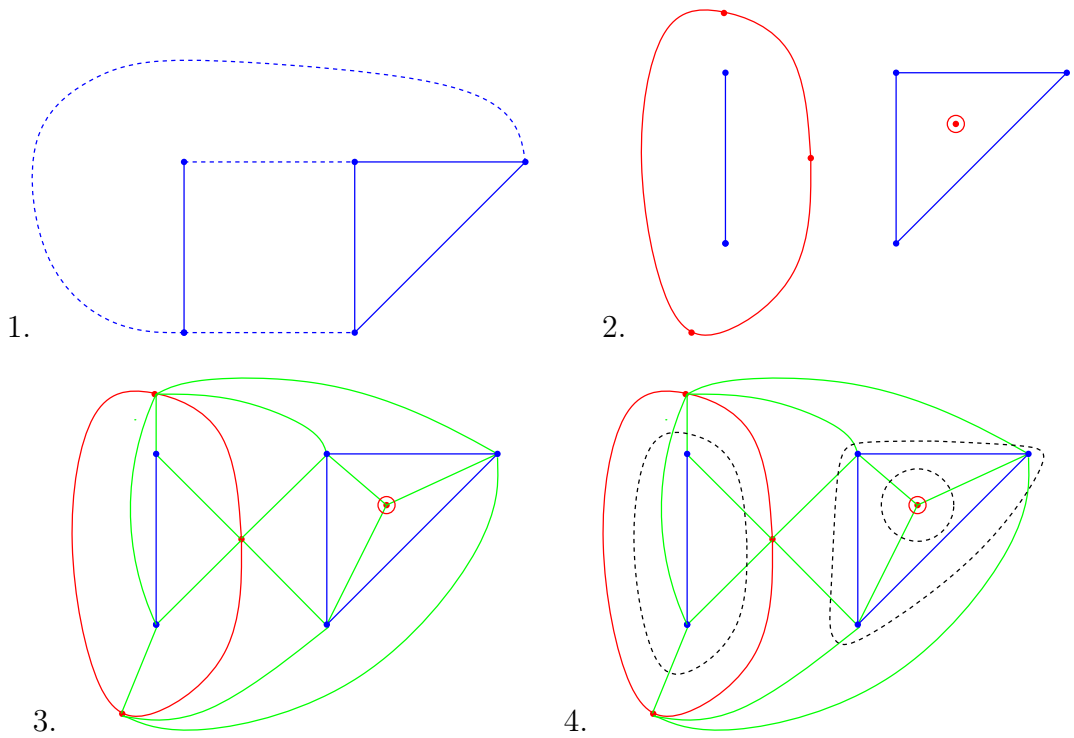

Figure 2: A map m decorated with loops associated with a set of open edges t. *Top left:* the map is in blue, with solid open edges and dashed closed edges. *Top right:* Open clusters and corresponding open dual clusters in blue and red. *Bottom left:* every dual vertex is joined to its adjacent primal vertices by a green edge. This results in an augmented map $\bar{m}$ which is a triangulation. *Bottom right:* the primal and dual open clusters are separated by loops, which are drawn in black and are dashed. Each loop crosses every triangle once, and so can be identified with the set of triangles it crosses.

a primal connected component of $\boldsymbol{t}_n$ from a dual connected component in $\boldsymbol{t}_n^\dagger$, or vice-versa. Note that the set of loops is "space filling" in the sense that every triangle of the refined map is contained in a loop. Also every loop configuration corresponds to a unique configuration $\boldsymbol{t}_n$ and vice versa. Let $\ell(\boldsymbol{m}_n, \boldsymbol{t}_n)$ denote the number of loops corresponding to a configuration $(\boldsymbol{m}_n, \boldsymbol{t}_n)$.

Fortuin–Kasteleyn model. The particular distribution on planar maps which we consider was introduced in [She10b] and is roughly the following. Let $q \geq 0$ and let $n \geq 1$. The random map M_n that we consider is decorated with a (random) subset T_n of edges. The map T_n induces a dual collection of edges $T_n^\dagger$ on the dual map of M (see Figure 2). Let $\mathbf{m}$ be a planar map with n edges, and $\mathbf{t}$ a given subset of edges of $\mathbf{m}$. Then the probability to pick a particular $(\mathbf{m}, \mathbf{t})$ is, by definition, proportional to

$$\mathbb{P}(M_n = \mathbf{m}, T_n = \mathbf{t}) \propto \sqrt{q}^{\ell}, \tag{4.1}$$

where ℓ is the (total) number of loops in between both primal and dual vertex clusters in $\mathbf{t}$ which is equal to the combined number of cluster in T_n and $T_n^\dagger$ minus 1. Equivalently given the map $M_n = \mathbf{m}$, the collection of edges T_n follows the distribution of the self-dual Fortuin–Kasteleyn model, which is in turn closely related to the critical q-state Potts model, see [BKW76]. Accordingly, the map M_n is chosen with probability proportional to the partition function of the Fortuin-Kasteleyn model on it.

One reason for this particular choice, is the belief that for $q < 4$, after Riemann uniformisation, a large sample of such a map closely approximates a *Liouville quantum gravity* surface, described by an area measure of the form $e^{\gamma h(z)} dz$, where the parameter $\gamma \in (0, 2)$ is believed to be related to the parameter q of (4.1) by the relation

$$q = 2 + 2\cos\left(\frac{8\pi}{\kappa'}\right); \quad \gamma = \sqrt{\frac{16}{\kappa'}}. \tag{4.2}$$

Note that when $q \in (0, 4)$ we have that $\kappa' \in (4, 8)$ so that it is necessary to generate the Liouville quantum gravity with the associated dual parameter $\kappa = 16/\kappa' \in (0, 4)$. This ensures that $\gamma = \sqrt{\kappa} \in (0, 2)$.

4.2 Conjectured connection with Liouville Quantum Gravity

The distribution (4.1) gives us a natural family of sitributions on planar maps (indexed by the parameter $q \geq 0$). In this model, the weight of a particular map $\boldsymbol{m}$ is proportional to the *partition function* $Z(\boldsymbol{m}, q)$ of the critical FK model on the map $\boldsymbol{m} \in \mathcal{M}_n$ (the set of planar maps with n edges). Now, in the Euclidean world, scaling limits of critical FK models are believed to be related to CLE_κ collection of loops, where κ is related to the parameter q via the relation

$$q = 2 + 2\cos\left(\frac{8\pi}{\kappa'}\right). \tag{4.3}$$

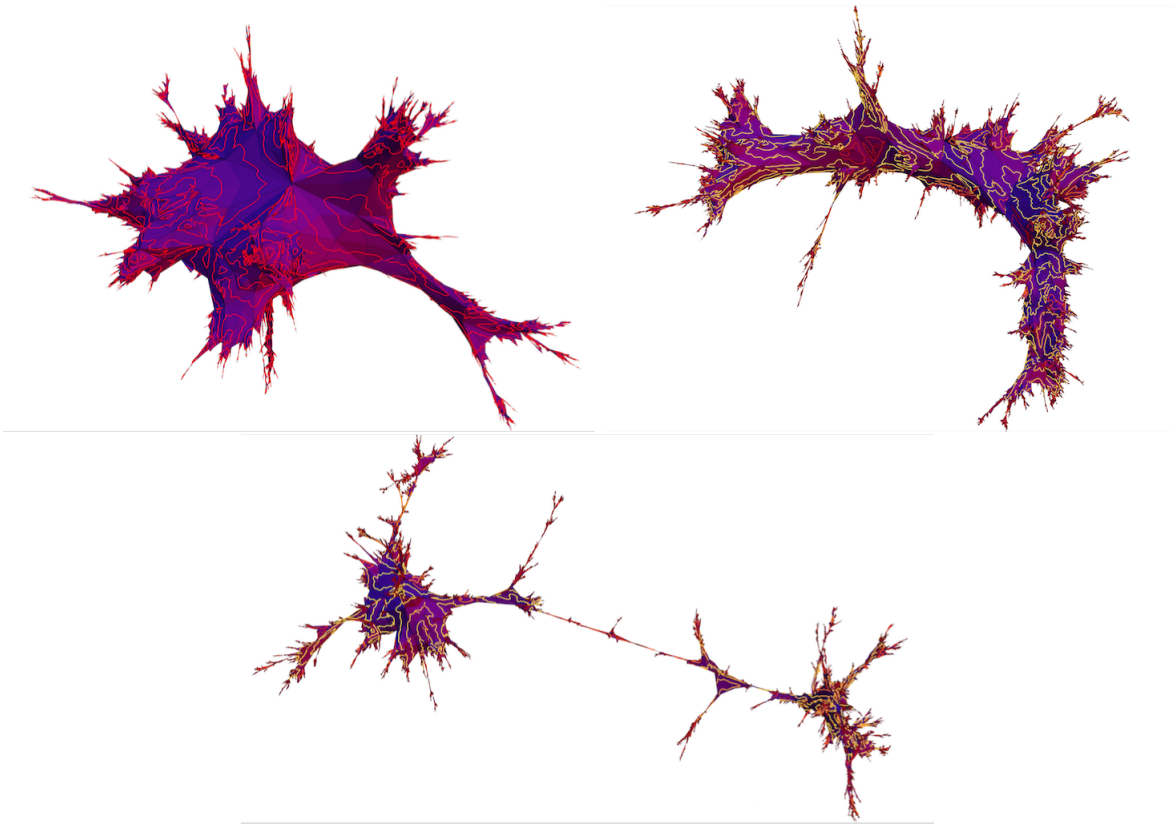

Figure 3: A map weighted by the FK model with $q = 0.5$, $q = 2$ (corresponding to the Ising model) and $q = 9$ respectively, together with some of their loops. Simulation by J. Bettinelli and B. Laslier. When $q > 4$ it is believed that the maps become tree-like, and the limiting metric space should be Aldous' continuum random tree.

For instance, when $q = 1$, the FK model reduces to percolation and one finds $\kappa = 6$, which is of course now rigorously established thanks to Stas Smirnov's proof of conformal invariance of critical percolation.

Uniform random planar maps. Observe that when $q = 1$, the FK model reduces to ordinary bond percolation. Hence this corresponds to the case where $\boldsymbol{m}_n$ is chosen according to the uniform probability distribution on planar maps with n edges. This is a situation in which remarkably detailed information is known about the structure of the planar map. In particular, a landmark result due to Miermont [Mie11] and Le Gall [LG13] is that, viewed as a metric space, and rescaling edge lengths to be $n^{-1/4}$, the random map converges to a multiple of a certain universal random metric space, known as the *Brownian map*. (In fact, the results of Miermont and Le Gall apply respectively to uniform quadrangulations with n faces and to p-angulation for $p = 3$ or p even, whereas the convergence result concerning uniform planar maps with n edges was established a bit later by Bettinelli, Jacob and Miermont [BJM14]). Critical percolation on a related half-plane version of the maps has been analysed in a recent work of Angel and Curien [AC15a], while information on the full plane percolation model was more recently obtained by Curien and Kortchemski [CK15]. Related works on loop models (sometimes rigorous, sometimes not) appear in [GJSZJ12, BBG12c, EK95, BBM11, BBG12a, BBG12b].

Likewise, when $q = 2$ corresponding to the FK-Ising model, we find $\kappa = 16/3$. Here again, this is known rigorously by results of Chelkak and Smirnov.

Conformal Embedding. Now, it is strongly believed that in the limit where $n \to \infty$, the geometry of such maps are related to Liouville quantum gravity where the parameter $\gamma^2 = 16/\kappa$. (In SLE theory, the value $16/\kappa$ is special as it transforms an SLE into a dual SLE - roughly, the exterior boundary of an SLE_κ is given by a form of $\mathrm{SLE}_{16/\kappa}$ when $\kappa \geq 4$. The self dual point is $\kappa = 4$, corresponding to $\gamma = 2$, which is also known to be critical for Liouville quantum gravity!) To relate the world of planar maps to the world of Liouville quatum gravity, one must specify a "natural" embedding of the planar maps into the plane. There are various ways to do this: the simplest ways to do so are:

– via the circle packing theorem. Each triangulation can be represented as a circle packing, meaning that the vertices of the triangulation are given by the centres of the circles, and the edges correspond to tangent circles. See Figure above.

– via the Riemann mapping theorem: indeed, each map can be viewed as a Riemann surface, by declaring that each a face of degree p is a regular p-gon of unit area, endowed with the standard metric, and specifying the charts near a vertex in the natural manner. Then such a Riemann surface can be embedded into the disc say by general theory of Riemann mapping.

These embeddings are essentially unique up to Möbius transforms, and this can be made unique in a number of natural ways. Either way, it is believed that in the limit, the empirical distribution on the vertices of the triangulation converge, after appropriate rescaling, to a

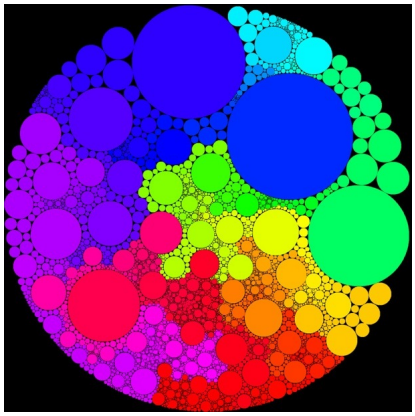

Figure 4: Circle packing of a uniform random planar map. Simulation by Jason Miller.

version of Liouville quantum gravity measure μ where:

$$q = 2 + 2\cos\left(\frac{8\pi}{\kappa'}\right); \quad \gamma = \sqrt{\frac{16}{\kappa'}}. \tag{4.4}$$

There are now several results making this connection precise at various levels: see [BLR15, DMS14a, GMS15].

Remark 4.1. Since the partition function of percolation is constant, a map weighted by the percolation model ($\kappa = 6$) is the same as a uniformly chosen random map. Consequently, the limit of a uniformly chosen map should be related to Liouville quantum gravity with parameter $\sqrt{16/\kappa} = \sqrt{16/6} = \sqrt{8/3}$.

4.3 Mullin–Bernardi–Sheffield's bijection in the case of spanning trees

We will now discuss the case where the map $\boldsymbol{m}_n$ is chosen with probability proportional to the number of spanning trees. In other words,

$$\mathbb{P}(\boldsymbol{m}_n, \boldsymbol{t}_n) \propto \mathbf{1}_{\{\boldsymbol{t}_n \text{ is a spanning tree}\}}. \tag{4.5}$$

This can be understood as the case $q \to 0^+$ of the Fortuin–Kasteleyn model above, since in that case the model concentrates on configurations such as $\ell = 0$ or, equivalently, $\boldsymbol{t}_n$ is a tree. In fact, it is immediate that in this case, given $\boldsymbol{m}_n$, $\boldsymbol{t}_n$ is a Uniform Spanning Tree (UST) on $\boldsymbol{m}_n$. We will discuss a powerful bijection due to Mullin [Mul67] and Bernardi [Ber07, Ber08] which is key to the study of such planar maps. This bijection is in fact a particular case of a bijection due to Sheffield which is sometimes called the "hamburger–cheeseburger" bijection, and which can be used in the general case where $q \geq 0$ is arbitrary. However, the case $q = 0$ of trees is considerably simpler, and so we discuss it (in the language of Sheffield, in order to prepare for the more general case which we will discuss later).

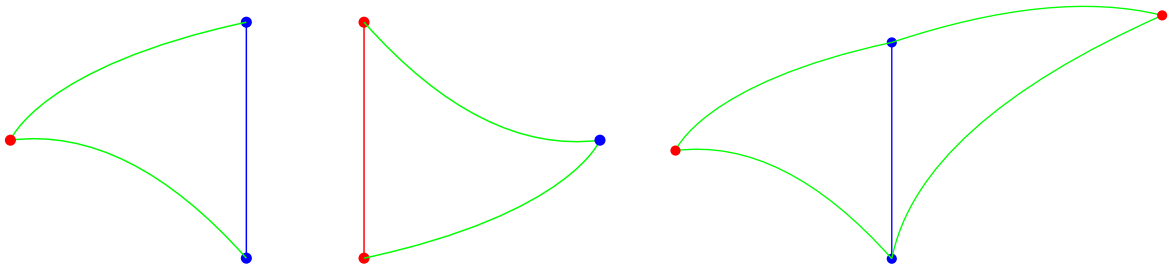

Figure 5: Refined or green edges split the map and its dual into primal and dual triangles. Each primal triangle sits opposite another primal triangle, resulting in a primal quadrangle as above.

Recall that the refinement edges split the map into triangles which can be of only two types: a primal triangle (meaning two green or refined edges and one primal edge) or a dual triangle (meaning two green or refined edges and one dual edge). For ease of reference primal triangles will be associated to hamburgers, and dual triangles to cheeseburgers. Now consider the primal edge in a primal triangle; the triangle opposite that edge is then obviously a primal triangle too. Hence it is better to think of the map as being split into quadrangles where one diagonal is primal or dual (see Figure 5).

We will reveal the map, triangle by triangle, by exploring it with a path which visits every triangle once (hence the word "space-filling"). We will keep track of the first time that the path enters a given quadrangle by saying that either a hamburger or a cheeseburger is produced, depending on whether the quadrangle is primal or dual. Later on, when the path comes back to the quadrangle for the second and last time, we will say that the burger has been eaten. We will use the letters h, c to indicate that a hamburger or cheeseburger has been produced and we will use the letters H, C to indicate that a burger has been eaten (or *ordered* and eaten immediately). So in this description we will have one letter for every triangle.

It remains to specify in what order are the triangles visited, or equivalently to describe the space-filling path. In the case where the decoration $\boldsymbol{t}_n$ consists of a single spanning tree (corresponding to $q = 0$ as we will see later) the path is simply the contour path going around the tree. Hence in that case the map is entirely described by a sequence w (or **word**) of a number N of letters in the alphabet $\Theta = \{\mathsf{h}, \mathsf{c}, \mathsf{H}, \mathsf{C}\}$. We also claim that it is always the case that $N = 2n$. To see why, recall that there is one letter for every triangle, so N is the total number of triangles. A triangle can also be identified with an edge (or in fact, half an edge, since each edge is visited exactly twice, the first time when the burger is produced, and the second time when it is eaten). Hence

$$N = 2(E(\boldsymbol{t}_n) + E(\boldsymbol{t}_n^\dagger)) = 2(V(\boldsymbol{t}_n) - 1 + V(\boldsymbol{t}_n^\dagger) - 1)$$

Now, the number of vertices of $\boldsymbol{t}_n$ is the same as $\boldsymbol{m}_n$ since $\boldsymbol{t}_n$ is spanning, and the number of vertices of $\boldsymbol{t}_n^\dagger$ is the same as the number of vertices in the dual map, which is equal to the

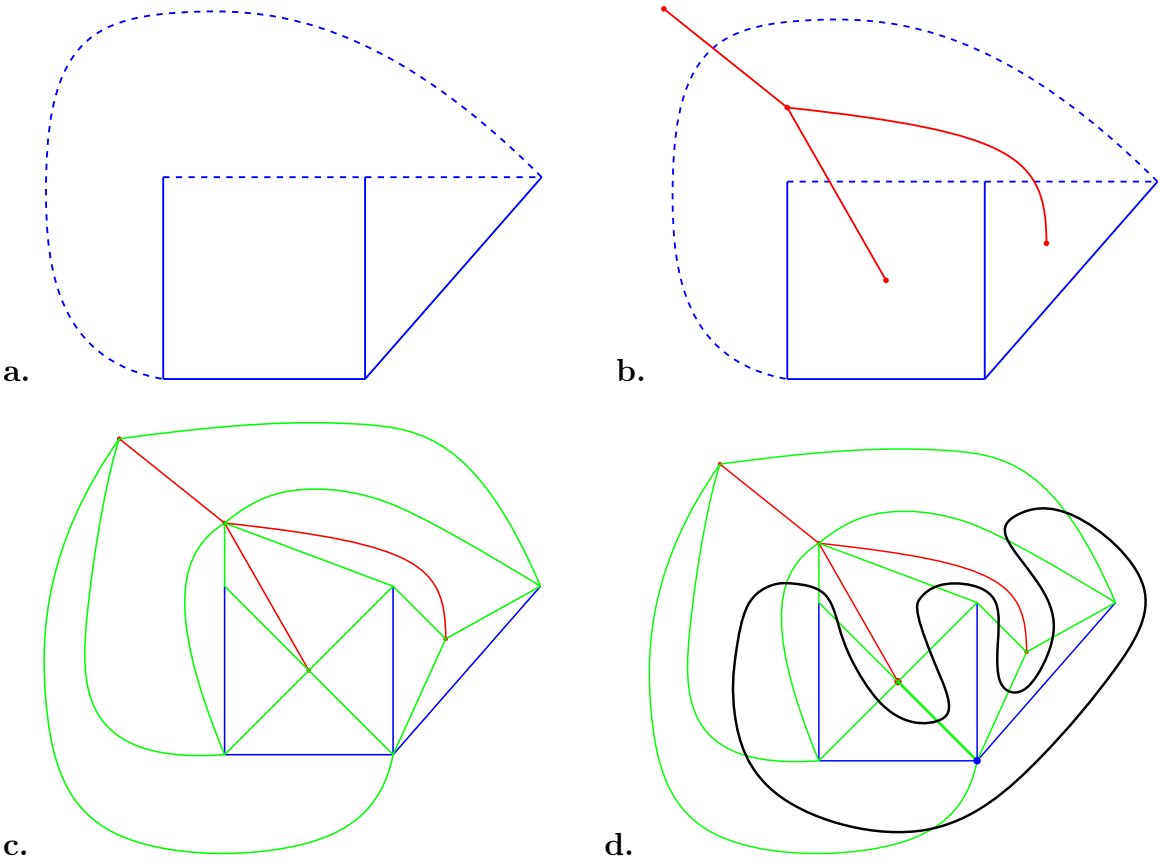

Figure 6: **a:** a map with a spanning tree. **b:** Spanning tree and dual tree. **c:** Refinement edges. **d:** Loop separating the primal and dual spanning trees, to which a root (refined) edge has been added in bold.

number of faces of $\boldsymbol{m}_n$. Hence

$$N = 2(V + F - 2). \tag{4.6}$$

Applying Euler's formula, we find $N = 2n$ as desired.

Together with this letter, we introduce the quantities $(X_k, Y_k)_{1\leq k\leq 2n}$ which count the number of hamburgers and cheeseburgers respectively in the stack at time k (i.e., the number of hamburgers or cheeseburgers which have been produced prior to time k but whose match come after time k). Then note that (X, Y) is a process which starts from the origin at time $k = 0$, and ends at the origin at time $k = 2n$. Moreover, X and Y both stay non-negative throughout. We call a process $(X_k, Y_k)_{0\leq k\leq 2n}$ satisfying these properties a **discrete excursion** (in the quarter plane).

It is obvious that $(X, Y)_{0\leq k\leq 2n}$ identify uniquely the word w encoding the decorated map $(\boldsymbol{m}_n, \boldsymbol{t}_n)$. Conversely, given such a process (X, Y) we can associate to it a unique word w in the letters of Θ such that (X, Y) is the net burger count of w.

Another property which is easy to check (and easily seen on the picture) is that the excursions X and Y encode the spanning trees $\boldsymbol{t}_n$ and dual spanning trees $\boldsymbol{t}_n^*$ in the sense

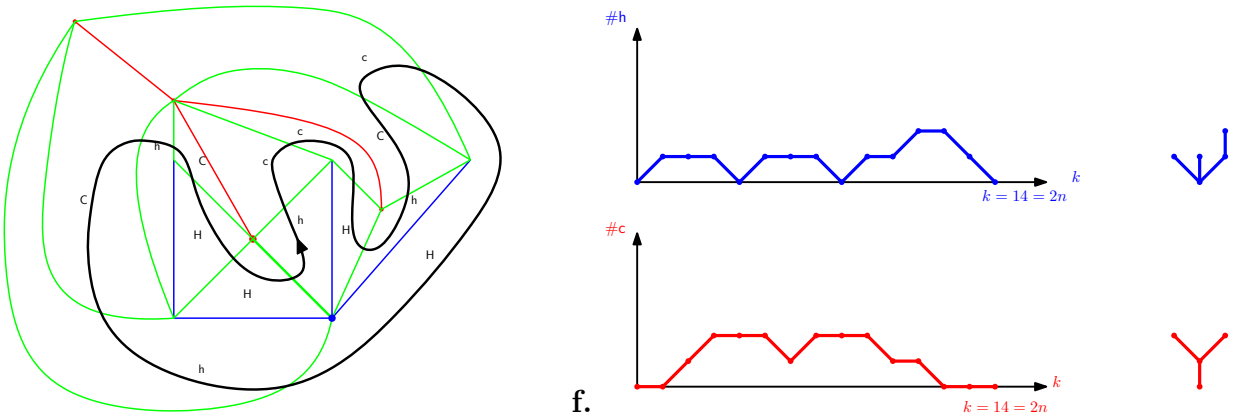

Figure 7: **e:** the word associated to $(\boldsymbol{m}_n, \boldsymbol{t}_n)$ is: $w = \mathsf{hccHhCcHhChCHH}$. **f:** The hamburger and cheeseburger counts, as well as the trees encoded by these excursions (which are identical to the primal and dual spanning trees, respectively).

that they are the contour functions: at a given time k, X_k denotes the height in the tree (distance to the root) of the last vertex discovered prior to time k. It follows from classical results of Durrett and Ingleheart that, as $n \to \infty$,

$$\frac{1}{\sqrt{n}}(X_{\lfloor 2nt \rfloor}, Y_{\lfloor 2nt \rfloor})_{0 \le t \le 1} \to (e_t, e'_t)_{0 \le t \le 1}$$

where e, e' are independent Brownian (one-dimensional) excursions (i.e., the pair (e, e') is Brownian excursion in the quarter plane).

We summarise our finding in the case of UST weighted random planar maps in the following result.

Theorem 4.2. *The set of spanning-tree decorated maps $(\boldsymbol{m}_n, \boldsymbol{t}_n)$ are in bijection with excursions $(X_k, Y_k)_{0 \le k \le 2n}$ in the quarter plane. Scaled by $\sqrt{n}$, the pair of trees $(\boldsymbol{t}_n, \boldsymbol{t}_n^*)$ converges to a pair of independent Continuous Random Trees (CRT).*

Note that the map $\boldsymbol{m}_n$ itself can then be thought of as a gluing of two independent CRTs, where the gluing is "uniform" among all possible ways to glue these two trees in a planar way. As this turns out, this procedure has a continuum analogue which has been studied in great detail by Duplantier, Miller and Sheffield [DMS14b].

4.4 The Loop-Erased Random Walk Exponent

A loop-erased random walk (or LERW for short) is the process one obtains when erasing the loops chronologically as they appear on a simple random walk trajectory. There is a well-known connection between Uniform Spanning Trees and LERW: here we will only need the following fact:

Theorem 4.3. *If x and y are two points of an arbitrary finite (connected, undirected) graph G, and if T is a UST on G, then the branch of T from x to y has the same distribution as a loop-erased random walk from x run until it hits y.*

This was originally proved by Pemantle [Pem91], and can also be seen as a straightforward consequence of generating Wilson's algorithm [Wil96]. However, since the proof is very short and beautiful, we provide the argument here.

Proof. For a possibly infinite path $\gamma = (\gamma_0, \gamma_1, \ldots)$ on the vertices V of the graph, let $T(\gamma)$ be the rooted tree obtained by retaining only the edges (γ_j, γ_{j+1}) which do not close a loop: i.e., keep this edge if and only if there is no $i < j$ such that $\gamma_i = \gamma_j$ (the root is the starting point $gamma_0$ of the path). It is obvious that this generates an acyclic graph; if the path visits every vertex then $T(\gamma)$ is a spanning tree. Now suppose that $(X_n)_{n\geq 0}$ is a stationary random walk on G, so that X_0 is distributed according to its equilibrium measure π, and let $\gamma_n = (X_n, X_{n+1}, \ldots)$ be the path started from X_n. Then the claim is that $(T(\gamma_n), X_n)$ defines a Markov chain on rooted spanning trees. This is best seen by viewing the tree $T(\gamma_n)$ as rooted at the directed edge $e_n = (X_n, X_{n+1})$. Then one step of the Markov chain is as follows. First choose an e' uniformly at random from the neighbours of the head e^+ of the edge e. Add it to the tree T, this now creates a cycle (possibly a double edge). Then remove form the tree T the unique edge touching the tail e^- in this cycle. See [LP, Chapter 4.4] for an illuminating alternative description of this argument.

When X is stationary then so is $(T(\gamma_n), X_n)$. But the unique invariant (reversible) measure for such a chain is clearly $\tilde{\pi}(\boldsymbol{t}, v) = \deg v$, the degree of the root of T. This is known as the **Markov chain tree theorem**. This particular algorithm for generating a uniform spanning tree is known as the **Aldous–Broder algorithm**. . Hence, when we condition on $X_0 = v$, we get the conditional distribution of the unrooted tree T obtained from ignoring the root in the rooted tree $T(\gamma_0)$ as

$$\begin{aligned}\mathbb{P}(T = \boldsymbol{t} | X_0 = v) &= \frac{\mathbb{P}((T(\gamma_0), X_0) = (\boldsymbol{t}, v))}{\pi(v)} \\ &\propto \frac{\tilde{\pi}(\boldsymbol{t}, v)}{\deg(v_0)} \propto 1.\end{aligned}$$

So T is a uniform (unrooted) spanning tree, and note that (remarkably) the law of T hence does not depend on the starting point v.

Now suppose the random walk X is started at $X_0 = x$, and consider the above construction of $T(X)$ which yields (as shown above) a uniform spanning tree. It is easy to check that the branch β of $T(X)$ connecting x to y to which we have removed all the loops but in a *reverse-chronological* order: that is, let γ be the path X run until it hits y, and let γ^* denote its time reversal. Then β is the loop-erasure $\Lambda(\gamma^*)$. The result is therefore proved if we show that $\Lambda(\gamma) = \Lambda(\gamma^*)$ in distribution, when viewed as a set of edges. The reason why this is true is because the ways the cycle are traversed can always be rearranged so that the backward erasure becomes the forward loop-erasure, and the probability of the whole path is unchanged by this reordering.

Another (perhaps more direct) way of seeing this is by considering a bi-infinite stationary version of the random walk $(\gamma_n, n \in \mathbb{Z})$ and looking at the tree associated with the time reversal $\hat{\gamma}_n = \gamma_{-n}$. Again, we refer to [LP] for more details on this argument.

□

Therefore, we deduce from Theorem 4.2 the following result about Loop-Erased Random Walk:

Theorem 4.4. *Let $(\boldsymbol{m}_n, \boldsymbol{t}_n)$ be chosen as in (4.5). Let x, y be two vertices chosen uniformly at random on $\boldsymbol{m}_n$, and let $(\Lambda_k)_{0\leq k\leq \xi_n}$ be a LERW starting from x, run until the random time ξ_n where it hits y. Then*

$$\frac{\xi_n}{\sqrt{n}} \to \xi_\infty$$

in distribution, where ξ_∞ is a random variable has a nondegenerate distribution on $(0, \infty)$ a.s.

A consequence of Theorem 4.4 is that for LERW the quantum scaling exponent satisfies:

$$\Delta = 1/2,$$

where Δ is the quantum scaling exponent. One can check that the associated parameter γ in this case is $\gamma = \sqrt{2}$ and hence the Euclidean scaling exponent is, by KPZ,

$$x = \frac{\gamma^2}{4}\Delta^2 + (1 - \frac{\gamma^2}{4})\Delta = 3/8.$$

So we deduce that the Loop-Erased Random Walk has dimension

$$d_{\text{Hausdorff}} = 2 - 2x = 5/4.$$

This is in accordance with Beffara's formula [Bef08] for the dimension of SLE, since LERW is in the scaling limit described by SLE_κ with $\kappa = 2$ (by a famous result of Lawler, Schramm and Werner [LSW04]). The Hausdorff dimension is therefore $1 + \kappa/8 = 5/4$, as above. In fact, the exponent for LERW had earlier been derived by Kenyon in a remarkable paper [Ken00] building on his earlier work on the dimer model and the Gaussian free field [Ken01].

4.5 Sheffield's bijection in the general case

We now describe the situation when the collection of edges $\boldsymbol{t}_n$ is arbitrary (i.e., not necessarily a tree), which is more delicate. The idea is that the space-filing path starts to go around the component of the root edge, i.e. explores the loop of the root edge, call it L_0. However, we also need to explore the rest of the map. To do this, consider the last time that L_0 is adjacent to some triangle in the complement of L_0, where by complement we mean the set of triangles which do not intersect L_0. (Typically, this time will be the time when we are about to close the loop L_0). At that time we continue the exploration as if we had **flipped**

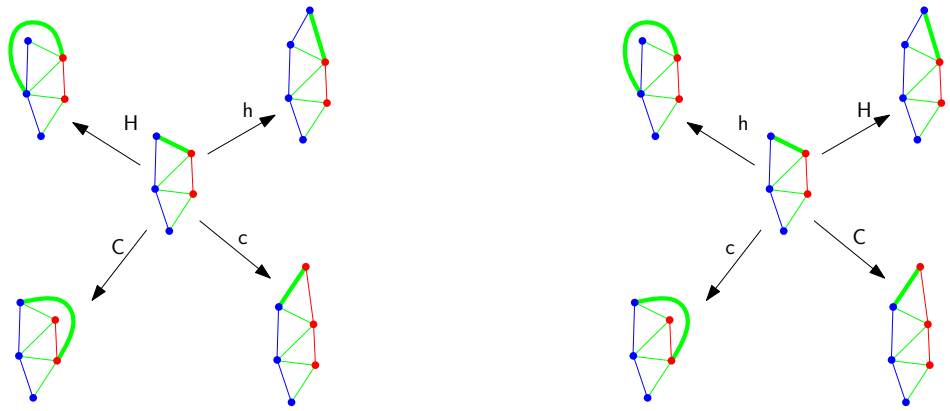

Figure 8: From symbols to map. The current position of the interface (or last discovered refined edge) is indicated with a bold line. Left: reading the word sequence from left to right or *into the future*. The map in the center is formed from the symbol sequence chc. Right: The corresponding operation when we go from right to left (or into the *past*). The map in the center now corresponds to the symbol sequence CHC.

the diagonal of the corresponding quadrangle. This has the effect the exploration path now visits two loops. We can now iterate this procedure. A moment of thought shows that this results in a space-filling path which visit every quadrangle exactly twice, going around some virtual tree which is not needed for what follows. We record a flipping event by the symbol F. More precisely, we associate to the decorated map $(\boldsymbol{m}_n, \boldsymbol{t}_n)$ a list of $2n$ symbols $(X_i)_{1\leq i\leq 2n}$ taking values in the alphabet $\Theta = \{\mathsf{h}, \mathsf{c}, \mathsf{H}, \mathsf{C}, \mathsf{F}\}$. For each triangle visited by the space-filling exploration path we get a symbol in Θ defined as before if there was no flipping, and we use the symbol F the second time the path visit a flipped quadrangle.

It is not obvious but true that this list of symbols completely characterises the decorated map $(\boldsymbol{m}_n, \boldsymbol{t}_n)$. Moreover, observe that each loop corresponds to a symbol F (except the loop through the root).

Inventory accumulation We now explain how to reverse the bijection. One can interpret an element in $\{\mathsf{h}, \mathsf{c}, \mathsf{H}, \mathsf{C}\}^{2n}$ as a last-in, first-out inventory accumulation process in a burger factory with two types of products: hamburgers and cheeseburgers. Think of a sequence of events occurring per unit time in which either a burger is produced (either ham or cheese) or there is an order of a burger (either ham or cheese). The burgers are put in a single **stack** and every time there is an order of a certain type of burger, the freshest burger in the stack of the corresponding type is removed. The symbol h (resp. c) corresponds to a ham (resp. cheese) burger production and the symbol H (resp. C) corresponds to a ham (resp. cheese) burger order.

Reversing the procedure when there is no F symbol is pretty obvious (see e.g. Figure 8). So we discuss the general case straight away. The inventory interpretation of the symbol F is the following: this corresponds to a customer demanding the freshest or the topmost burger in the stack irrespective of the type. In particular, whether an F symbol corresponds to a

hamburger or a cheeseburger order depends on the topmost burger type at the time of the order. Thus overall, we can think of the inventory process as a sequence of symbols in Θ with the following reduction rules

- $\mathsf{cC} = \mathsf{cF} = \mathsf{hH} = \mathsf{hF} = \emptyset$,
- $\mathsf{cH} = \mathsf{Hc}$ and $\mathsf{hC} = \mathsf{Ch}$.

Given a sequence of symbols X, we denote by $\bar{X}$ the reduced word formed via the above reduction rule.

Given a sequence X of symbols from Θ, such that $\bar{X} = \emptyset$, we can construct a decorated map $(\boldsymbol{m}_n, \boldsymbol{t}_n)$ as follows. First convert all the F symbols to either a H or an C symbol depending on its order type. Then construct a spanning tree decorated map as is described above (Figure 8). The condition $\bar{X} = \emptyset$ ensures that we can do this. To obtain the loops, simply switch the type of every quadrangle which has one of the triangles corresponding to an F symbol. That is, if a quadrangle formed by primal triangles has one of its triangles coming from an F symbol, then replace the primal map edge in that quadrangle by the corresponding dual edge and vice versa. The interface is now divided into several loops and the number of loops is exactly one more than the number of F symbols.

Generating FK-weighted maps. Fix $p \in [0, 1/2)$. Let $X_1, \ldots, X_n$ be i.i.d. with the following law

$$\mathbb{P}(\mathsf{c}) = \mathbb{P}(\mathsf{h}) = \frac{1}{4}, \mathbb{P}(\mathsf{C}) = \mathbb{P}(\mathsf{H}) = \frac{1-p}{4}, \mathbb{P}(\mathsf{F}) = \frac{p}{2}. \tag{4.7}$$

conditioned on $\overline{X_1, \ldots, X_n} = \emptyset$.

Let $(\boldsymbol{m}_n, \boldsymbol{t}_n)$ be the random associated decorated map as above. Then observe that since n hamburgers and cheeseburgers must be produced, and since $\#\mathsf{H} + \#\mathsf{C} = n - \#\mathsf{F}$,

$$\begin{aligned} \mathbb{P}((\boldsymbol{m}_n, \boldsymbol{t}_n)) &= \left(\frac{1}{4}\right)^n \left(\frac{1-p}{4}\right)^{\#\mathsf{H}+\#\mathsf{C}} \left(\frac{p}{2}\right)^{\#\mathsf{F}} \\ &\propto \left(\frac{2p}{1-p}\right)^{\#\mathsf{F}} = \left(\frac{2p}{1-p}\right)^{\#\ell(\boldsymbol{m}_n, \boldsymbol{t}_n)-1} \end{aligned} \tag{4.8}$$

Thus we see that $(\boldsymbol{m}_n, \boldsymbol{t}_n)$ is a realisation of the critical FK-weighted cluster random map model with $\sqrt{q} = 2p/(1-p)$. Notice that $p \in [0, 1/2)$ corresponds to $q = [0, 4)$. From now on we fix the value of p and q in this regime. (Recall that $q = 4$ is believed to be a critical value for many properties of the map).

4.6 Local limits and the geometry of loops

The following theorem due to Sheffield and made more precise later by Chen [She10b, Che15] shows that the decorated map (M_n, T_n) has a local limit as $n \to \infty$ in the local topology. Roughly two maps are close in the local topology if the finite maps near a large neighbourhood of the root are isomorphic as maps.

Theorem 4.5 ([She10b, Che15]). *Fix $p \in [0,1)$. We have*

$$(M_n, T_n) \xrightarrow[n\to\infty]{(d)} (M, T)$$

in the local topology.

Furthermore, (M, T) can be described by applying the obvious infinite version of Sheffield's bijection to the bi-infinite i.i.d. sequence of symbols with law given by (4.7).

The idea behind the proof of Theorem 4.5 is the following. Let $X_1, \dots, X_{2n}$ be i.i.d. with law given by (4.7) conditioned on $\overline{X_1 \dots X_{2n}} = \emptyset$. It is shown in [She10b, Che15] that the probability of $\overline{X_1 \dots X_{2n}} = \emptyset$ decays sub exponentially. Using Cramer's rule one can deduce that locally the symbols around a uniformly selected symbol from $\{X_i\}_{1 \le i \le n}$ converge to a bi-infinite i.i.d. sequence $\{X_i\}_{i \in \mathbb{Z}}$ in law. The proof is now completed by arguing that the correspondence between the finite maps and the symbols is continuous in the local topology.

Notice that uniformly selecting a symbol corresponds to selecting a uniform triangle in $(\bar{M}_n, \bar{T}_n)$ which in turn corresponds to a unique refinement edge which in turn corresponds to a unique oriented edge in M_n. Because of the above interpretation and the invariance under re-rooting, one can think of the triangle corresponding to X_0 as the root triangle in (M, T).

One important thing that comes out of the proof is that every symbol in the i.i.d. sequence $\{X_i\}_{i \in \mathbb{Z}}$ has an almost sure unique **match**, meaning that every order is fulfilled and every burger is consumed with probability 1. Let $\varphi(i)$ denote the match of the ith symbol. Notice that $\varphi : \mathbb{Z} \mapsto \mathbb{Z}$ defines an involution on the integers.

Let L_n denote a *typical loop*, that is, a loop chosen uniformly at random from the set of loops induced by (M_n, T_n) which follow the law given by (4.1); here we identify a loop with the set of triangles that it passes through. Let $\mathsf{Len}(L_n)$ denote the number of triangles in the loop and let $\mathsf{Area}(L_n)$ denote the number of triangles inside it (where the "inside" of the loop is the connected component of the loop which does not contain the root triangle). Let

$$p_0 = \frac{\pi}{4 \arccos\left(\frac{\sqrt{2-\sqrt{q}}}{2}\right)} = \frac{\kappa'}{8} \tag{4.9}$$

where q and κ' are related as in (4.2).

Theorem 4.6. *We have that $\mathsf{Len}(L_n) \to \mathsf{L}$ and $\mathsf{Area}(L_n) \to \mathsf{A}$ in law. Further, the random variables L and A satisfy the following.*

$$\mathbb{P}(\mathsf{L} > k) = k^{-1/p_0 + o(1)}, \tag{4.10}$$

and

$$\mathbb{P}(\mathsf{A} > k) = k^{-1+o(1)}. \tag{4.11}$$

4.7 Exercises

1. By a theorem of [BLR15] (see also [GMS15] for related statements), the loop critical exponent of FK-weighted maps is given by

$$\mathbb{P}(\partial L \geq k) \approx k^{-1/p_0}; p_0 = \frac{\pi}{4 \arccos\left(\frac{\sqrt{2-\sqrt{q}}}{2}\right)}$$

   Use KPZ and the relation between $q = 2 + 2\cos(8\pi/\kappa)$ to recover the dimension of SLE_κ.

2. Another proof of uniformity.

5 Scale-Invariant Random Surfaces*

Here we introduce some variants of the Gaussian free field which define random surfaces that are scale-invariant in some precise sense: zooming in near a given point of the surface produces the same surface in distribution.

We first need to set up the scene correctly. This will require working with a slightly different version of the GFF with nonzero boundary conditions: the so-called Neumann (or free boundary) GFF. As this section requires familiarity with SLE, proofs will be less elementary than in previous chapters.

5.1 Definition of Neumann boundary GFF

We start with the notion of Neumann boundary Gaussian Free Field (also known as free boundary GFF). In general if we wish to add boundary data to a GFF it is natural to simply add a function which is harmonic in the domain (though it can have relatively wild behaviour on the boundary). We will seek to impose **Neumann boundary conditions**: recall that for a smooth function this means that the normal derivative vanishes along the boundary (if the domain is smooth). Of course for an object as rough as the GFF it is a priori unclear what that condition means, and we will see that indeed the resulting object is the same as when we don't impose any condition (which is why the field can also be called a free boundary GFF, as is the case in the papers of Sheffield [She10a] and the subsequent papers by Duplantier, Miller and Sheffield [DMS14a]).

One annoying problem with this object is that it really is only defined up to a global additive constant, meaning that it can either be viewed as a **distribution modulo constants** (two distributions are equivalent if their difference is a constant function) or it can be defined as a distribution after fixing some normalisation (eg by requiring that the integral of h be zero, or if in the upper half plane $\mathbb{H}$ that its average on the unit upper half disc or unit upper half circle be zero). The latter is a concrete object and a bit easier to define and to work with in practice, but is not conformally invariant or indeed scale invariant. The former on the other hand is a better point of view as it enjoys conformal invariance, but in this setup things like the Liouville measure will only be defined up to a global multiplicative constant. This too can have serious drawbacks for proofs. Therefore, depending on the context, one point of view may be more appropriate than the other, and we will make use of both in these notes.

Let D be a domain and let $\tilde{\mathcal{D}}(D)$ be the set of smooth functions on $\bar{D}$ (meaning all derivatives extend continuously to the boundary), with bounded support, and such that $\int_D f = 0$. Let us assume for now that $\mathcal{D}$ has smooth boundary. Let $\tilde{\mathcal{D}}_+(D)$ be the set of smooth functions in $\bar{D}$ with bounded support such that

$$\int_{\partial D} \frac{\partial f}{\partial n} = 0; \quad \int f(x)dx = 0$$

On $\tilde{\mathcal{D}}_+(D)$, observe that the Dirichlet product $(f,g)_\nabla$ still defines an inner product even though the functions do not have zero boundary conditions. The completion of $\tilde{\mathcal{D}}_+(D)$ is

a Hilbert space which we call $\tilde{H}^1(D)$. The $+$ in the index is for the boundary condition (Neumann boundary conditions have something to do with reflected Brownian motion) and the $\sim$ on top is for the normalisation whereby we require the mean of the function to be zero.

Definition 5.1. *The GFF with Neumann boundary conditions (also known as **free boundary conditions**) and zero mean is the random distribution $\tilde{h} = \sum_n X_n \tilde{f}_n$ where $\tilde{f}_n$ is an orthonormal basis of $\tilde{H}^1(D)$.*

Remark 5.2. It is easy to check that smooth functions with Neumann boundary conditions, $\tilde{\mathcal{D}}_+(D)$, is dense in $\tilde{\mathcal{D}}(D)$ for the Dirichlet inner product. By consequence the two sets have the same Hilbert space completion, and so the GFF associated to Neumann boundary conditions and to free boundary conditions are deterministically the same object.

Remark 5.3. Another reason why the term *free boundary conditions* is employed by the authors of [DMS14a], [She10a], is that in the discrete, a box where we don't impose any particular boundary conditions to vertices on the boundary ("free boundary condition", similar to the terminology in use for e.g. the definition of uniform spanning trees) give rise to a random walk that is reflected on the boundary of the box, and hence, in the scaling limit, to reflecting Brownian motion, whose generator is associated with Neumann boundary conditions.

It is easy to see that the sum defining the Neumann GFF in Definition 5.1 converges a.s. in the space of distributions, and indeed in the Sobolev space $H^{-\varepsilon}$ as before. Note that if h is defined by Definition 5.1, then (h, f) is a.s. well defined even if f itself doesn't have zero mean; rather this definition should be understood as saying that it is h which has an integral equal to zero (or more precisely the truncated series has an integral which is always zero). The choice of having mean zero is arbitrary: we could alternatively have used mean zero on the unit disc or the unit circle, etc. All such choices correspond to a Neumann GFF with what we call **a different normalisation**.

Distribution modulo constants. Alternatively, we can define the Neumann GFF as a random variable $\bar{h}$ in the space of distribution modulo additive constants. First let $\bar{\mathcal{D}}_+(D)$ be the space of smooth functions on $\bar{D}$ with Neumann boundary conditions, modulo additive constants: i.e., two functions are considered to be equivalent if their difference is a constant function. Then we view $(f, g)_\nabla$ as an inner product on the space $\bar{\mathcal{D}}_+(D)$ and make the following definition:

Definition 5.4. *The Neumann GFF, viewed as a distribution modulo constants, is given by $\bar{h} = \sum_n X_n \bar{f}_n$ where $\bar{f}_n$ is an orthonormal basis of $\bar{H}^1(D)$, the Hilbert space completion of $\bar{\mathcal{D}}_+(D)$ induced by $(f, g)_\nabla$.*

Note that the distribution $\tilde{h}$ of Definition 5.1 induces a "distribution modulo constant" $\bar{h}$ (simply consider the equivalence class of $\tilde{h}$) which has the same law as the random variable from Definition 5.4. This follows from the fact that if $\tilde{f}_n$ is an orthonormal basis of $\tilde{H}^1(D)$ then their equivalence classes form an orthonormal basis of $\bar{H}^1(D)$.

The following proposition is then straightforward to prove due to the conformal invariance of $(f,g)_\nabla$ in $\bar{\mathcal{D}}(D)$:

Proposition 5.5. *If $\bar{h}$ is a GFF with Neumann boundary conditions (viewed as a random distribution modulo constant), then $\bar{h}$ is conformally invariant: namely, if $\varphi : D \to D'$ is a conformal map then $\bar{h} \circ \varphi^{-1}$ is a GFF with free boundary conditions in D'.*

This conformal invariance allows us to define the Neumann GFF for domains which don't necessarily have a smooth boundary. Note that when we integrate a Neumann GFF with a fixed normalisation (say $\tilde{h}$) against a smooth function ρ with mean zero, the choice of the normalisation does not affect the value of $(\tilde{h}, \rho)$. In other words, $(\tilde{h}, \rho)$ is constant over the equivalence class of $\tilde{h}$, and hence we may unambiguously refer to the value of $(\bar{h}, \rho)$. In other words, the space of distributions modulo constants is simply the space of linear functionals on smooth functions with zero mean.

Even / odd decomposition. The whole plane GFF can be constructed more or less in the same manner, by considering the Hilbert space completion $H_{\mathbb{C}}$ of the set of smooth functions in $\mathbb{C}$ with bounded support and zero mean, with respect to the Dirichlet inner product. The whole plane GFF is then the Gaussian process associated to this Hilbert space: i.e., $h_{\mathbb{C}} = \sum_n X_n f_n$, where X_n are i.i.d. standard normal random variables, and f_n is an orthonormal basis of $H_{\mathbb{C}}$. Let

$$H_{\text{even}} = \{h \in H_{\mathbb{C}} : h(z) = h(\bar{z}), z \in \mathbb{C}\}$$

and likewise let

$$H_{\text{odd}} = \{h \in H_{\mathbb{C}} : h(z) = -h(\bar{z}), z \in \mathbb{C}\}.$$

Then it is straightforward to check that $H_{\mathbb{C}} = H_{\text{even}} \oplus H_{\text{odd}}$ (orthogonality follows from the change of variable $z \mapsto \bar{z}$). Therefore, we can write

$$h_{\mathbb{C}} = h_{\text{even}} + h_{\text{odd}}$$

where $h_{\text{even}}, h_{\text{odd}}$ are the standard Gaussian fields in $H_{\text{even}}, H_{\text{odd}}$ respectively. But observe that odd functions in H_{odd} are exactly obtained by considering Dirichlet functions on the upper half-plane in $\mathbb{H}$ and reflecting with a minus sign. Likewise, even functions in H_{even} on $\mathbb{C}$ are exactly obtained by consider Neumann boundary conditions functions on $\mathbb{H}$ and reflecting. We deduce that the restriction of h_{odd} to the upper half plane is a Dirichlet boundary condition GFF on $\mathbb{H}$. Likewise, the restriction of h_{even} to $\mathbb{H}$ is a Neumann GFF on $\mathbb{H}$. We summarise this with the following proposition:

Proposition 5.6. *The even part of a whole plane GFF is a Neumann GFF on the upper half plane.*

5.2 Green function in the Neumann case

We will focus on the case of the upper half plane, where calculations are a bit easier, as we already observed in the Dirichlet case. Recall that by substracting the trajectories ending

in the negative half plane, we obtained the formula that the Green function with Dirichlet boundary conditions in the upper half plane was given by $G_0(x,y) = \log(|x-\bar{y}|) - \log(|x-y|)$. It is reasonable to guess (because of the relationships between Neumann boundary conditions and refelecting Brownian motion) that the corresponding Green function in the Neumann case, meaning the "covariance" of the Neumann GFF, will be given by $G(x,y) = -\log(|x-\bar{y}|) - \log(|x-y|)$. Note the change of sign. Indeed, this is what we will find – with one caveat however: the fact that the field is really only defined *up to a constant* means that fixing this constant by some mean (e.g. by imposing zero mean in some region or other) leads to different exact values for the covariance. However, as explained before, when a test function ρ has zero mean, the value of (h,ρ) is constant over the whole equivalence class of h, and so its distribution is independent of the normalisation of h. We state the following result:

Theorem 5.7. *Suppose $\rho \in \tilde{\mathcal{D}}(\mathbb{H})$: i.e., ρ is smooth over $\bar{\mathbb{H}}$, has bounded support, and $\int \rho(z)dz = 0$. Let h be a Neumann GFF on $\mathbb{H}$ with some arbitrary normalisation. Then (h,ρ) is a centered Gaussian random variable with variance*

$$\mathrm{Var}(h,\rho) = \int_{\mathbb{H}} \rho(x)\rho(y)G(x,y)dxdy$$

where

$$G(x,y) = -\log(|x-\bar{y}|) - \log(|x-y|).$$

Remark 5.8. Note that since ρ has mean zero, the distribution of $(h.\rho)$ does *not* depend on the normalisation of the Neumann GFF. So Theorem 5.7 can be used as a characterisation of the Neumann GFF viewed as a distribution modulo constants, a fact we will rely on for later arguments.

Remark 5.9. As in the Dirichlet case, this formula extends readily to the case of signed measures ρ with zero mass, satisfying the condition that $\int G(x,y)|\rho(dx)||\rho(dy)| < \infty$, allowing us to turn $\bar{h}$ or $\tilde{h}$ into a stochastic process indexed by signed measures of zero mass and finite Green energy. In turn this can be used as the definition of the object, along the same lines as the definition in the case of Dirichlet boundary conditions. The only property that is nontrivial to check is the nonnegative semidefinite character of G. This follows from an integration by parts formula similar to the one used below in the proof, or the fact that G is the inverse of $(-1/2)$ times the Laplacian with Neumann boundary conditions, and hence has only nonnegative eigenvalues.

Sketch of proof of Theorem 5.7. Let $\rho \in \mathcal{D}(\bar{\mathbb{H}})$ be a test function with compact support in $\bar{\mathbb{H}}$ and we wish to compute the variance of (h,ρ). Suppose we could find f such that $\Delta f = -2\pi\rho$ and f has Neumann boundary condition. The problem:

$$\text{Find } f \text{ such that: } \begin{cases} \Delta f = -2\pi\rho & \text{in } D \\ \dfrac{\partial f}{\partial n} = v & \text{on } \partial D \end{cases} \tag{5.1}$$

is called the Neumann problem. It is known that it has a solution if and only if $\int_D -2\pi\rho = \int_{\partial D} v$ (this condition comes from the divergence theorem: the integral of v along the boundary measures the total flux across the boundary, while the integral the Laplacian of f inside the domain measures the pointwise flow). This solution is then unique, up to a global constant. That is, this solution is unique in $\bar{H}^1(D)$. When $D = \mathbb{H}$ and $v = 0$, one such solution is then given by

$$f(x) = \int_{\mathbb{H}} G(x,y)\rho(y)dy. \tag{5.2}$$

(The Neumann boundary conditions of G give rise to the Neumann boundary conditions of f, as is easily checked). We call $\bar{f}$ the equivalence class of f in $\bar{H}^1(D)$.

Hence assume ρ has mean zero, so the Neumann problem with $\partial f/\partial n = 0$ has a solution $f \in H^1(\mathbb{H})$ as above (in fact, see Remark 5.10 after the proof for a subtlety here). Consider the truncated sum $h_N = \sum_{n=1}^{N} h_n f_n$. Doing an integration by parts, taking into account the boundary terms (this is the general form of the Gauss–Green formula), we know that

$$\begin{aligned} 2\pi(h_N, f)_\nabla &= -\int_{\mathbb{H}} h_N \Delta f + \int_{\mathbb{R}} h_N \frac{\partial f}{\partial n} \\ &= -\int_{\mathbb{H}} h_N \Delta f = 2\pi(h_N, \rho) \end{aligned} \tag{5.3}$$

since f satisfies Neumann boundary condition. On the other hand, it is clear by linearity that

$$(h_N, f)_\nabla = \sum_{n=1}^{N} X_n (f_n, f)_\nabla.$$

Letting $n \to \infty$, and dividing by 2π, we obtain that

$$(h, \rho) = \sum_{n=1}^{\infty} X_n (f_n, f)_\nabla. \tag{5.4}$$

(Convergence of the left hand side is because the series converges in the sense of distribution, while convergence of the right hand side is guaranteed by the martingale convergence theorem). By (5.4) we deduce, since adding constants do not change the value of either terms (as ρ has mean zero),

$$(\bar{h}, \rho) = \sum_{n=1}^{\infty} X_n (\bar{f}_n, \bar{f})_\nabla. \tag{5.5}$$

Thus

$$\mathrm{Var}(\bar{h}, \rho) = \sum_{n=1}^{\infty} (f_n, \bar{f})_\nabla^2 = \|\bar{f}\|_\nabla^2.$$

since $(\cdot,\cdot)_\nabla$ is an inner product on $\bar{H}^1_+$. By the Gauss–Green formula with boundary terms,

$$\|\bar{f}\|_\nabla^2 = -\frac{1}{2\pi}\int \bar{f}\Delta\bar{f}$$

since $\bar{f}$ has Neumann boundary conditions (note that the integral on the right is also defined unambiguously across the equivalence class $\bar{f}$) and so

$$\mathrm{Var}(\bar{h}, \rho) = \int f(x)\rho(x)dx = \int \rho(x)\rho(y)G(x,y)dxdy$$

by definition of f in (5.2). □

Remark 5.10. The proof above hides one additional subtlety. Recall the choice of $f(x) = \int G(x,y)\rho(y)dy$, and suppose for a moment that ρ does not have zero mean. While it is true that f has Neumann boundary condition across all of $\mathbb{R}$, this becomes wrong when we think of the boundary of $\mathbb{H}$ as being $\mathbb{R}$ together with an extra point at infinity. This becomes apparent when we replace $\mathbb{H}$ by the unit disc $\mathbb{D}$, by applying the Möbius map $z \mapsto m(z) = (i-z)/(i+z)$, which maps ∞ to -1. Then, in the sense of distributions, the Neumann boundary condition of G_+ is

$$\nabla G(x,\cdot)\cdot n = 2\pi\delta_{-1}(\cdot),$$

see Exercise 4. This will result in the fact that

$$\nabla f \cdot n = 2\pi|\rho|\delta_{-1}$$

as a distribution too, where $|\rho| = \int \rho(z)dz$ is the total mass of ρ. Hence it is only when ρ has mass zero that we can really say that f has Neumann boundary condition across all of its boundary (i.e., including the point at infinity). This is in fact why the identity (5.3) (and the rest of the proof) is justified: when we apply the Gauss–Green formula in $\mathbb{H}$, the boundary should really include a point at infinity (as is apparent when we apply the formula not to $\mathbb{H}$ but to a large bounded domain) and hence it is important to know that the Neumann boundary condition is zero also at infinity. A clean proof of Theorem 5.7 would therefore require working in the unit disc throughout the proof, but we have preferred not to bother with the calculations.

Remark 5.11 (**Non-uniqueness; Green function in the unit disc**)**.** The Green function G in the statement of Theorem 5.7 and in the proof is not unique. For instance, if $\tilde{G}$ is chosen so that the boundary condition is $\nabla G(\cdot, y)\cdot n = 1$ along all the boundary then $f(x) = \int \tilde{G}(x,y)\rho(y)dy$ also satisfies

$$\nabla f \cdot n = 2\pi|\rho| = 0$$

and hence the arguments of the proof above go through equally well. We end up with a different expression for the covariance of the Neumann GFF. However, it can be checked that we would have $\tilde{G}(x,y) = G(x,y) + \phi_1(x) + \phi_2(y)$ for some functions f, g. Therefore, when we integrate against a test function ρ which has mean zero, the resulting variance is

$$\int \tilde{G}(x,y)\rho(x)\rho(y)dxdy = \int G(x,y)\rho(x)\rho(y)dxdy$$

which is the same as the variance appearing in the statement of Theorem 5.7. This is another expression of the fact that the field is only defined up to a constant, so that the pointwise covariance might be different – yet the integrated variance (tested against a mean zero function) is well defined independently of the normalisation.

Depending on which domain one works with, certain calculations are easier with a particular choice of the Green function. For instance, in the upper half plane, it is more natural to put the singularity in $\nabla G \cdot n$ at infinity and thus have $G(x,y) = -\log|x-y| - \log|x-\bar{y}|$ as in the statement of the theorem. However, in the unit disc $\mathbb{D}$, it is more natural to spread the singularity uniformly on the boundary of the unit disc and hence require $\nabla G \cdot n = 1$ along $\partial\mathbb{D}$. This leads to the choice of the Green function

$$\tilde{G}(x,y) = -\log|(x-y)(1-x\bar{y})|; x, y \in \mathbb{D}.$$

This is indeed the choice being made in the paper by Miller and Sheffield on QLE [MS13], see e.g. (4.7) in that paper, as well as the work of Huang, Rhodes and Vargas which constructs (true) Liouville quantum gravity on the unit disc [HRV15]. Again, it is worth emphasising that this choice is irrelevant: though the expressions for the Green functions G and $\tilde{G}$ are pointwise different, they nevertheless serve to define the same GFF, up to an additive (and possibly random) constant.

5.3 Semicircle averages and boundary Liouville measure

Let h be a Neumann GFF on $\mathbb{H}$ with some arbitrary normalisation. An immediate consequence of the above calculations is the following fact. For $x \in \mathbb{R}$, let $h_\varepsilon(x)$ denote the average of h on the upper semicircle $\partial B(x,\varepsilon) \cap \mathbb{H}$ of radius ε about x: hence if $\rho^+_{x,\varepsilon}$ is the uniform probability distribution on this set, then $h_\varepsilon(x) = (h, \rho^+_{x,\varepsilon})$.

Theorem 5.12. *For any $x \in \mathbb{R}$, the process $(h_{e^{-t}}(x) - h_1(x); t \in \mathbb{R})$ is a two-sided Brownian motion with variance 2 (so $\mathrm{Var}((h_{e^{-t}}(x) - h_1(x)) = 2|t|$).*

Note that the statement of the theorem makes sense: the increments $h_{e^{-t}}(x) - h_1(x)$ are well defined independently of the normalisation of h. (Indeed, $\rho = \rho^+_{x,\varepsilon} - \rho^+_{x,\delta}$ has mean zero).

Proof. Without loss of generality we may take $x = 0$ and assume that h is normalised so that $h_1(0) = 0$. Then it is clear that $X_t = h_{e^{-t}}(0)$ has stationary increments by scaling properties of h (with this normalisation) and these increments are independent by an easily proved Markov property (see exercises). Moreover the mean of these increments is clearly zero and the variance is finite, so $X_t = B_{\kappa t}$ for some $\kappa > 0$, where B is a standard Brownian motion. It remains to check that $\kappa = 2$. This follows from the fact that $G(0,y) = 2\log(1/\varepsilon)$ if $|y| = \varepsilon$. The result follows. □

Having identified the boundary behaviour of the Neumann GFF, we can now construct an interesting measure supported on the boundary of $\mathbb{H}$. As it turns out, the measure of interest to us is again an exponential of the Neumann GFF, but the multiplicative factor in the exponential is $\gamma/2$ rather than γ. The reason for this choice is rather deep, and has to

do with the fact that we plan to use it to measure the "quantum length of an SLE" (another justification comes from the fact that it satisfies the same KPZ equation as in the bulk case).

Theorem 5.13. *Let h be a Neumann GFF with some normalisation. Define a measure ν_ε on $\mathbb{R}$ by setting $\nu_\varepsilon(dx) = \varepsilon^{\gamma^2/4} e^{(\gamma/2)h_\varepsilon(x)} dx$. Then for $\gamma < 2$, the measure ν_ε converges a.s. along the dyadic subsequence $\varepsilon = 2^{-k}$ to a nontrivial, nonatomic measure ν called the boundary Liouville measure.*

Note the normalisation which is by $\varepsilon^{\gamma^2/4}$. This is because, as proved in Theorem 5.12, when $x \in \mathbb{R}$ and h is a Neumann GFF, $\operatorname{Var} h_\varepsilon(x) = 2\log(1/\varepsilon) + O(1)$. Once this is observed the proof of the construction of the bulk Liouville measure μ adapts without any trouble (the use of the Markov property is justified with appropriate modifications – though in fact, as remarked before, this property is an unnecessary assumption, see e.g. [Ber15]).

The boundary measure is not conformally covariant: i.e., the analogue of Theorem 2.8 does not hold in a straightforward way. This is because the behaviour of conformal maps near the boundary of a domain can be very wild. For instance if the conformal map is only Hölder with a certain exponent, then the boundary Liouville measure can be transformed into a boundary Liouville measure corresponding to a different exponent. In a way this is what happens in the case of SLE. (On the other hand, if the domain is smooth enough that the conformal map extends across the boundary say, then the conformal covariance will clearly hold).

5.4 Convergence of random surfaces

The Neumann GFF is conformally invariant (and in particular scale invariant) when viewed as a distribution modulo constants. However when we fix a normalisation this is no longer true, and so the abstract surface described by $(\mathbb{H}, h)$ is not scale invariant (in the sense that applying the conformal map $z \mapsto rz$ and applying the conformal change of coordinates formula does not yield the same surface in distribution). In order to construct a surface $(\mathbb{H}, h)$ which is invariant under scaling, Sheffield introduced the notion of *quantum wedges* which will play an important role in our study of the zipper. Roughly, this is the surface that one obtains by "zooming in" close to a point at the boundary. Since this surface is obtained as a scaling limit it is then automatic that it will be invariant under scalings.

In order to make sense of the above we first need to discuss a notion of convergence for random surfaces – and more precisely, for surfaces with two marked points on its boundary (which will be 0 and ∞ in the case of the upper half plane). Recall first our definition of an abstract surface – a pair (D, h) modulo the equivalence class $(D, h) \sim (D', h')$ if there exists a conformal map $\phi : D \to D'$ such that $h' = h \circ \phi^{-1} + Q \log |(\phi^{-1})'|$, where $Q = \gamma/2 + 2/\gamma$.

Definition 5.14. *A random surface (D, h) with two marked points z_0, z_1 on its (conformal) boundary is called **half-plane like** if some neighbourhood of z_0 has finite Liouville mass, and any neighbourhood of z_1 has infinite Liouville mass.*

Such an abstract surface may be embedded in the domain $D = \mathbb{H}$ with the two marked points being $z_0 = 0$, $z_1 = \infty$. Such an embedding is not unique: indeed, the set of conformal

maps mapping $\mathbb{H}$ to itself and mapping 0 and ∞ to themselves consists precisely of the dilations $z \mapsto rz$.

Definition 5.15. *The **canonical description** of a half-plane like surface (D, h, z_0, z_1) is the unique equivalent surface $(\mathbb{H}, \tilde{h}, 0, \infty)$ such that $\mu_{\tilde{h}}(\mathbb{D}) = 1$.*

In other words, for a half-plane like surface, we fix its embedding in $\mathbb{H}$ by choosing the dilation parameter r so that $\mathbb{D}$ has Liouville mass 1 (such a choice is always possible by our definition of half-plane like surface).

Example: zooming in (important!) Let h be a Neumann GFF, say normalised to have average zero in the unit ball. Then the canonical description of h and of $h + 100$ (say) are *very different* – and this can be confusing at first since h is in some sense defined "up to a constant". Indeed to find the canonical description of h we just need to find the (random) r such that $\mu_h(B(0, r)) = 1$, and apply the conformal map $z \mapsto z/r$; the resulting field

$$\tilde{h}(z) = h(rz) + Q \log(r)$$

is the canonical description of the surface $(\mathbb{H}, h)$.

On the other hand, in order to find the canonical description of $h + 100$, we need to find $s > 0$ such that $\mu_{h+100}(B(0, s)) = 1$. That is, we need to find $s > 0$ such that $\mu_h(B(0, s)) = e^{-100\gamma}$. The resulting field

$$h^*(z) = h(sz) + Q \log(s) + 100$$

is the canonical description of $(\mathbb{H}, h + 100)$.

Note that in this example, the ball of radius s is much smaller than the ball of radius r. Yet in $\tilde{h}$, the ball of radius r has been scaled to become the unit disc, while in h^* it is the ball of radius s which has been scaled to become the unit disc. In other words, and since s is much smaller than r, the canonical description of $h + 100$ corresponds to **zooming in** the surface $(\mathbb{H}, h)$. This property has nothing to do with the Neumann GFF per se – it is a completely deterministic observation (in particular, this has nothing to do with the fact that the Neumann GFF is only really defined up to a constant, since we had to start by fixing a normalisation).

Having fixed a canonical description (or canonical embedding), we can formulate a notion of convergence of surfaces.

Definition 5.16. *We say that a sequence of half-plane like surfaces (D_n, h_n, z_0^n, z_1^n) converges to a half-plane like surface (D, h, z_0, z_1) if in the canonical description of these surfaces $(\mathbb{H}, h_n, 0, \infty)$ and $(\mathbb{H}, h, 0, \infty)$ we have that the measures μ_{h_n} converge weakly towards μ_h.*

5.5 Quantum wedges

As we will see later on, quantum wedges are abstract random surfaces arising as limits (in the sense discussed above) of surfaces of the form $h + C$, as $C \to \infty$, where h is a Neumann GFF with a certain logarithmic singularity at the origin. Thus, as explained in the example above, they correspond to zooming in near the origin of such a surface.

In practice however we prefer to work with a concrete definition and then prove that this surface can indeed be seen as a limit as above. It will turn out to be more convenient to define it in the infinite strip $S = \mathbb{R} \times (0, \pi)$ rather than the upper half plane, with the two marked boundary points being $+\infty$ and $-\infty$ respectively. The conformal map transforming $(S, \infty, -\infty)$ into $(\mathbb{H}, 0, \infty)$ is $z \mapsto -e^{-z}$. Then vertical segments are mapped to semicircles; and more precisely the segment $\{z : \Re(z) = s\}$ is mapped into $\partial B(0, e^{-s})$.

Lemma 5.17. *Let S be the infinite strip $S = \mathbb{R} \times (0, \pi)$. Let $\mathcal{H}_{\text{rad}}$ be the subspace of $H^1_+(S)$ obtained as the closure of smooth functions which are constant on each vertical segment. Let $\mathcal{H}_{\text{circ}}$ be the subspace obtained as the closure of smooth functions which have mean zero on all vertical segments. Then*

$$H^1_+ = \mathcal{H}_{\text{rad}} \oplus \mathcal{H}_{\text{circ}}.$$

Consequently, a Neumann GFF on S, normalised so that its average on $(0, i\pi)$ is zero, can be written as $h = h_{\text{rad}} + h_{\text{circ}}$ where:

- $h_{\text{rad}}, h_{\text{circ}}$ are independent,
- $h_{\text{rad}}(z) = B_{2s}$ if $\Re(z) = s$ and B is an independent standard Brownian motion (by Theorem 5.12)
- $h_{\text{circ}}(z)$ has mean zero on each vertical segment.

Note that all the roughness of h is contained in the h_{circ} part, as the h_{rad} is a nice continuous functions (which is even constant on vertical segments). Also, on the upper half plane, this would correspond to a radial decomposition of h into a part which is a radially symmetric continuous function, and one which has zero average on every semicircle, explaining our notation.

Let $\alpha < Q$. We will define an α-quantum wedge by specifying separately its radially symmetric behaviour on $\mathcal{H}_{\text{rad}}$ and its mean zero behaviour on $\mathcal{H}_{\text{circ}}$. The behaviour on $\mathcal{H}_{\text{circ}}$ is exactly the same as in the standard Neumann GFF case, it is only the radially symmetric part which is different.

Definition 5.18. *Let*

$$h_{\text{rad}}(z) = \begin{cases} B_{2s} + (\alpha - Q)s & \text{if } \Re(z) = s \text{ and } s \geq 0 \\ \widehat{B}_{-2s} + (\alpha - Q)s & \text{if } \Re(z) = s \text{ and } s < 0 \end{cases} \tag{5.6}$$

where B is a standard Brownian motion, and $\widehat{B}$ is an independent Brownian motion conditioned so that $\widehat{B}_{2t} + (Q - \alpha)t > 0$ for all $t > 0$.

*Let $h_{\text{circ}}(z)$ be the zero semicircle average part of a Neumann GFF. Then $h = h_{\text{rad}} + h_{\text{circ}}$ is called an α-**quantum wedge** in $\mathbb{H}$.*

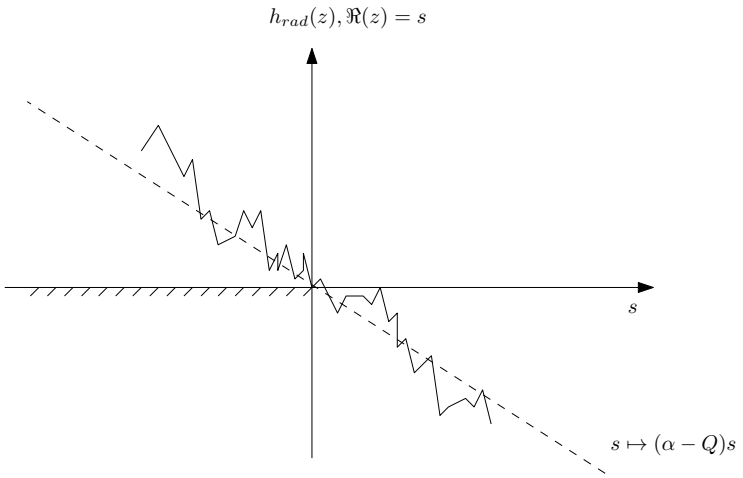

Figure 9: Schematic representation of the radially symmetric part of a quantum wedge in a strip. When $s < 0$, the function is conditioned to be positive.

Remark 5.19. We have defined the quantum wedge as an abstract random surface (S, h) in the sense of Definition 2.10. Using the change of coordinate formula (Theorem 2.8) we could thus also view it as parametrised by the upper half plane, and we would obtain a field $\tilde{h}$ defined on $\mathbb{H}$. However the expression for $\tilde{h}$ is not particularly nice, and makes the following proofs more difficult to follow (which is why we take the strip S as our domain of reference, for once).

The conditioning in the case where $s < 0$ (corresponding to radii greater than 1) takes some work to define, especially if $\alpha = Q$, and it will not be the purpose of these notes to specify its meaning. Note that when $s > 0$, $h_{\mathrm{rad}}(s)$ has a drift of coefficient $\alpha - Q < 0$ as $s \to \infty$. This means that, embedding in the upper half plane using $z \mapsto -e^{-z}$ and taking into account the conformal change of variables, we obtain a quantum wedge field h defined in the upper half plane, which has a logarithmic singularity of coefficient α near zero (i.e., it looks like $\alpha \log(1/z)$ near zero). Note also that if h is a quantum wedge (embedded in $\mathbb{H}$ as above), then $\tilde{h}(z) = h(z) - \alpha \log 1/|z|$, restricted to $\mathbb{D} \cap \mathbb{H}$, has the same law as a Neumann GFF restricted to $\mathbb{D} \cap \mathbb{H}$, normalised so that it has zero average on the circle of radius 1.

We emphasise that the description of the quantum wedge given above is in no way the *canonical description* of this wedge. In the context of strips the canonical description is obtained by making a horizontal translation by an amount s chosen so that $\mu_h(\{z : \Re(z) \geq s\}) = 1$. This amount is a random variable with very good tails when $\alpha < Q$.

We can now state the result about the quantum wedge being the scaling limit of a Neumann GFF with a logarithmic singularity near the origin.

Theorem 5.20. *Fix $\alpha \leq Q$. Then the following hold:*
(i) *Let $\tilde{h}$ be a Neumann GFF in $\mathbb{H}$ with some normalisation and let $h(z) = \tilde{h}(z) + \alpha \log 1/|z|$. Then the quantum surface described by $(\mathbb{H}, h + C/\gamma, 0, \infty)$ converges in law, as $C \to \infty$, to an α-quantum wedge.*

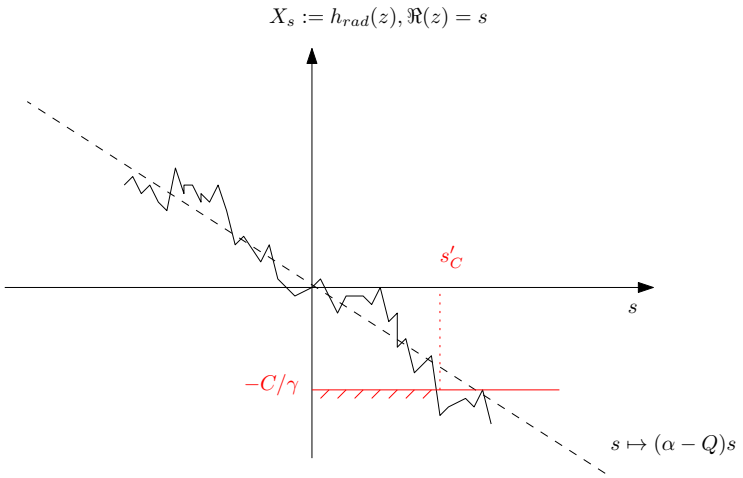

Figure 10: Proof of the theorem.

(ii) *If* $(\mathbb{H}, h, 0, \infty)$ *is the canonical description of an* α*-quantum wedge, and if* h_C *is the canonical description in* $(\mathbb{H}, 0, \infty)$ *of* $h + C/\gamma$*, then* $h_C = h$ *in distribution.*

Recall that (ii) says that a quantum wedge is invariant under rescaling, while (i) says that a quantum wedge is the limit, zooming in near zero, of the surface described by $\tilde{h}(z) + \alpha \log 1/|z|$.

Proof. We will assume for simplicity that the normalisation of $\tilde{h}$ is such that the average value of h on the unit upper semicircle is zero (all other normalisations can be treated with identical arguments). We first embed the field h into the strip S via the conformal map $z \in S \mapsto \phi(z) = -e^{-z} \in \mathbb{H}$. Then applying the change of coordinates rule, we see that the radially symmetric part of h becomes

$$h_{\text{rad}}(z) = B_{2s} + (\alpha - Q)s; \quad z \in S, \Re(z) = s \in \mathbb{R}, \tag{5.7}$$

where B is a standard two-sided Brownian motion. The term $+\alpha$ is the drift corresponding to the logarithmic singularity, while the term $-Q$ comes from the change of variables.

In order to prove the convergence in part (i) of the theorem, it is sufficient to prove convergence of a modified canonical description: namely, we can choose the translation parameter s by requiring $s = \inf\{t \in \mathbb{R} : h_{\text{rad}}(z) = 0\}$ (rather than requiring $\mu_h(\{\Re(z) \geq s\}) = 1$). Once again this identifies uniquely an embedding of a given surface within the strip. Then convergence of $h + C/\gamma$ in the sense of Definition 5.16 will follow from convergence of the *modified canonical description* of $h + C/\gamma$ to the process of Definition 5.18, in total variation say. Since the h_{circ} part of h is independent from h_{rad} and is a stationary process, it suffices to prove convergence in total variation of $h_{\text{rad}}(\cdot + s_C) + c/\gamma$ towards the process $h_{\text{rad}}(z)$ from Definition 5.18 (see (5.6)).

Therefore, let $X_s = B_{2s} + (\alpha - Q)s$, as in (5.7), so $h_{\text{rad}}(z) = X_s$ whenever $\Re(z) = s$. Set $s'_C = \inf\{s > 0 : X_s = -C/\gamma\}$ (see Figure 10). Let

$$Y_s^C = X_{s'_C + s} + C/\gamma$$

Then we make the following observations about Y^C.

1. By the strong Markov property of X, $(Y^C_s)_{s\geq 0}$ is a Brownian motion with variance 2 and drift $Q-\alpha$, independent of s'_C.

2. For $s \in [0, s'_C]$, it is also not hard to see that conditionally on s'_C, $(Y^C_{-s})_{s\in[0,s'_C]}$ is a Brownian motion with variance 2, conditioned on being positive during that interval (see picture), and conditioned to take the value C/γ at time s'_C.

3. For $s \geq s'_C$, and conditionally on s'_C, $(Y^C_{-s})_{s\geq s'_C}$ is just a Brownian motion with drift $Q-\alpha$.

As $C \to \infty$, we have that $s'_C \to \infty$ almost surely. Hence it is not hard to see that the conditioning in 2. just becomes a conditioning on being positive, while the behaviour described in 3. is irrelevant in the limit. In other words, Y^C converges (in total variation) towards the process in (5.6), and part (i) of the theorem is proved.

Point (ii) follows immediately since scaling limits must be invariant under scaling. □

5.6 Exercises

1. Prove that smooth functions in $\mathbb{H}$ satisfying the Neumann boundary condition are dense in smooth functions for the Dirichlet inner product. Deduce that the Neumann GFF and the free boundary conditions GFF are identical.

2. Prove that the Neumann GFF in $\mathbb{H}$ can be written as the sum of an independent harmonic function in $\mathbb{H}$ and a Dirichlet GFF. State and prove a Markov property for the Neumann GFF.

3. What can be said about the law of the restriction of a Neumann GFF on $\mathbb{H}$ to $\mathbb{R}$? And what about the restriction of a full plane GFF to $\mathbb{R}$? Given this restriction, what is the law of the GFF on the rest of its domain in either of these cases?

4. Check that if $G_{\mathbb{D}}$ is the Green function on the unit disc obtained by the Green function G on $\mathbb{H}$ of Theorem 5.7 by the Möbius map $z \mapsto m(z) = (i-z)/(i+z)$, then $\nabla G_{\mathbb{D}}(x,\cdot)\cdot n = 2\pi\delta_{-1}(\cdot)$. (Hint: first check the Neumann boundary conditions on $\mathbb{R}$ for G, then apply the Gauss Green formula on $\mathbb{D}$).

5. Check that the boundary Liouville measure ν satisfies the same KPZ relation as the bulk Liouville measure.

6. Suppose we sample a point x from the boundary Liouville measure ν (restricted to (0,1) and renormalised so that it is a probability distribution). Is the point x thick for the field? If so, how thick?

7. The boundary Liouville measure is not conformally covariant. But can you construct a measure on the segment $(0, i)$ by exponentiating a Neumann GFF in $\mathbb{H}$ which enjoys a property of conformal covariance?

8. Show that Theorem 5.20 (i) remains true if we replace h by $h = \tilde{h} + \alpha \log(1/|\cdot|) + \varphi$, where $\tilde{h}$ is a Neumann GFF on $\mathbb{H}$ with some normalisation and φ is a (possibly random) function which is continuous at 0: that is, show that if h is as above then as $C \to \infty$, the surfaces $(\mathbb{H}, h + C/\gamma)$ converge to an α-thick wedge.

9. Let $D = \{z : \arg(z) \in [0, \theta]\}$ be the (Euclidean) wedge of angle θ, and suppose that $\theta \in (0, 2\pi)$. Let h be a Neumann GFF in D. Show that by zooming in (D, h) near the tip of the wedge, we obtain a thick quantum wedge with $\alpha = Q(\theta/\pi - 1)$ (which satisfies $\alpha < Q$ if $\theta < 2\pi$).

10. (Proposition 4.7 in [DMS14a]): show the following characterisation of quantum wedges. Fix $\alpha < Q$ and suppose that $(\mathbb{H}, h, 0, \infty)$ is a canonical description of a quantum surface parameterised by $\mathbb{H}$ such that the following hold:

    (i) The law of $(\mathbb{H}, h, 0, \infty)$ (as a quantum surface) is invariant under the operation of multiplying its area by a constant. That is, if we fix $C \in \mathbb{R}$, and then let $(\mathbb{H}, h_C, 0, \infty)$ be a canonical description of $(\mathbb{H}, h + C/\gamma, 0, \infty)$, then $(\mathbb{H}, h_C, 0, \infty) = (\mathbb{H}, h, 0, \infty)$

    (ii) The total variation distance between the law of $(\mathbb{H}, h, 0, \infty)$ restricted to $B(0, r)$ and the law of an α-quantum wedge restricted to $B(0, r)$ tends to 0 as $r \to 0$.

    Then $(\mathbb{H}, h, 0, \infty)$ has the law of an α-quantum wedge.

11. (**Quantum cones.**) By replacing the strip $S = \mathbb{R} \times (0, \pi)$ with the infinite cylinder $\bar{S} = \mathbb{R} \times (0, \pi) \mod \pi$ (so the top and bottom boundaries of the strip are identified, define for all $\alpha < Q$ a quantum surface $(\bar{S}, h)$ such that Theorem 5.20 holds, with the Neumann GFF h in that theorem replaced by a whole plane GFF h on $\mathbb{C}$. Such a surface is called a (thick) quantum cone [DMS14a].

    State and prove the analogue of Exercise 9.

6 SLE and the quantum zipper*

In this section we discuss some fundamental results due to Sheffield [She10a], which have the following flavour.

1. Theorem 6.1: An SLE_κ curve has a 'nice' coupling with $e^{\gamma h}$. This coupling can be formulated as a Markov property analogous to the domain Markov properties inherent to discrete random maps. This makes the conjectures about convergence of random maps toward Liouville quantum gravity plausible, and in particular this justifies that the 'right' relation between κ and γ is $\kappa = \gamma^2$.

2. Theorem 6.9: An SLE_κ curve can be endowed with a random measure which can be roughly interpreted as $e^{\gamma h} d\lambda$ where $d\lambda$ is a natural length measure on the curve. (This is in itself hard to define, so this is not the route which we will take – moreover, the exponent γ would need to be changed to a slightly different value to take into account the quantum scaling exponent of the SLE curve: see [BSS14] for a discussion).

3. Theorem 6.22: An SLE_κ curve slits the upper half plane into two independent random surfaces, glued according to boundary length.

6.1 SLE and GFF coupling: Markov property for random surfaces

Here we state one of the two couplings between the GFF and SLE. This was first stated in the context of Liouville quantum gravity (but in a slightly different way from here) in [She10a]. However, ideas for a related coupling go back to two seminal papers by Schramm and Sheffield [SS13] on the one hand an Dubédat [Dub09] on the other.

Let h be a Neumann GFF on $\mathbb{H}$, let $\kappa = \gamma^2 \in (0,4)$, and set

$$h_0 = h + \varphi \qquad \text{where } \varphi(z) = \frac{2}{\gamma} \log|z|; \quad z \in \mathbb{H}. \tag{6.1}$$

Hence h_0 is a Neumann GFF to which we have *substracted* (rather than added) a logarithmic singularity at zero, of a slightly unusual multiple.

Let $\eta = (\eta_t)_{t \geq 0}$ be an independent chordal SLE_κ curve in $\mathbb{H}$, going from 0 to ∞, where recall that $\kappa = \gamma^2$. Hence η is paramaterised by its half-plane capacity. Let g_t be the unique conformal map $g_t : \mathbb{H} \setminus \{\eta_s\}_{s \leq t} \to \mathbb{H}$ such that $g_t(z) = z + 2t/z + o(1/z)$ as $z \to \infty$ (we will call g_t the Loewner map). Then

$$\frac{dg_t(z)}{dt} = \frac{2}{g_t(z) - \xi_t}; \qquad z \notin \{\eta_s\}_{s \leq t}$$

where $(\xi_t)_{t \geq 0}$ is the Loewner driving function and has the law of a linear Brownian motion of variance κ. Set $\tilde{g}_t(z) = g_t(z) - \xi_t$ to be the *centered* Loewner map.

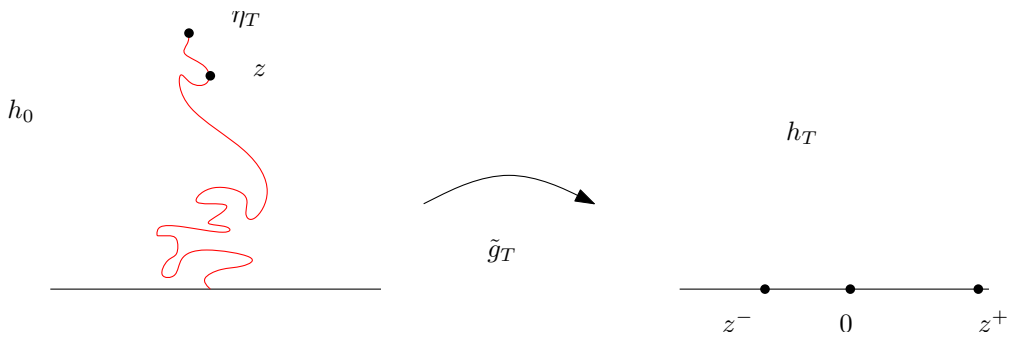

Figure 11: Start with the field h_0 and an independent SLE_κ curve run up to some time T. After mapping h_0, restricted to the complement of the curve H_T, by the Loewner map g_T and applying the change of coordinate formula, we obtain a distribution h_T in $\mathbb{H}$ which by the theorem has the same law as h_0. This is the Markov property for random surfaces.

Theorem 6.1. *Let $T > 0$, and set*

$$h_T = h_0 \circ \tilde{g}_T^{-1} + Q \log |(\tilde{g}_T^{-1})'|, \text{ where } Q = \frac{2}{\gamma} + \frac{\gamma}{2}.$$

Then h_T defines a distribution in $\mathbb{H}$ modulo constants which has the same law as h_0 for all fixed T.

Remark 6.2. Here we have started with a field h_0 with a certain law (described in (6.1)) and a curve η which is *independent* of h_0. However, η is *not* independent of h_T. In fact, we will see later on that h_T entirely *determines* the curve η up to time T. More precisely, we will see in Theorem 6.9 that when we apply the map $\tilde{g}_T$ to the curve $(\eta_s)_{0\leq s\leq T}$ the boundary lengths (measured with h_T) of the two intervals to which η is mapped by g_T must agree: that is, on Figure 11, the lengths of $[z^-, 0]$ and $[0, z^+]$ agree). We will later (in Theorem 6.21) be able to show that given h_T, the curve $(\eta_s)_{0\leq s\leq T}$ is in fact determined by the requirement that g_T^{-1} maps intervals of equal quantum length to identical pieces of the curve η. This is the idea of **conformal welding** (we are welding $\mathbb{H}$ to itself by welding together pieces of its boundary that have the same quantum length).

Remark 6.3. The proof of the theorem (and the statement which can be found in Sheffield's paper [She10a, Theorem 1.2], involves the reverse Loewner flow (f_t) rather than, for a fixed t, the map $\tilde{g}_t^{-1}$. In this context, the theorem is equivalent to saying that

$$h_T = h_0 \circ \tilde{f}_T + Q \log |f_T'|, \text{ where } Q = \frac{2}{\gamma} + \frac{\gamma}{2}.$$

Moreover, in this case the theorem is also true if T is say a bounded stopping time (for the underlying reverse Loewner flow). I have chosen this formulation however, because the usual forward Loewner flow is a simpler object and more natural in the context of the Markovian interpretation discussed below. Nevertheless the formulation in terms of the reverse flow will be the most useful in the rest of the section.

Discussion and interpretation. Let $H_T = \mathbb{H} \setminus \{\eta_s\}_{s \leq t}$. In the language of random surfaces, the theorem states that the random surface (H_T, h_0) has the same distribution $(\mathbb{H}, h_0)$, because h_T is precisely obtained from h_0 by mapping its restriction to H_T through the centered Loewner map $\tilde{g}_T$ and applying the change of coordinate formula (the conformal covariance of Theorem 2.8). In other words, suppose we start with a surface described by $(\mathbb{H}, h_0)$. Then we explore a small portion of it using an independent SLE_κ started where the logarithmic singularity is located (here it is important to assume that γ and κ are related by $\kappa = \gamma^2$). What is the law of the surface which remains to be discovered after some time T? The theorem states that this law is the same as the original one. Hence this theorem can be seen as simply stating a **Markov property for random surfaces**.

Connection with the discrete picture. This Markov property is to be expected from the discrete side of the story. To see this, consider for instance the Uniform Infinite Half Plane Triangulation (UIHPT) constructed by Angel [Ang03]. This is obtained as the local limit of a uniform planar map with a large number of faces and a large boundary, rooted at a uniform edge along the boundary. One can further add a critical site percolation process on this process by colouring vertices black or white independently with probability $1/2$ (as shown by Angel, this is indeed the critical value). We make an exception for vertices along the boundary, where those to the left of the root edge are coloured in black, and those to the right in white. This generates an interface and it is possible to use that interface to discover the map. Such a procedure is called *peeling* and was used with great efficacy by Angel and Curien [AC15b] to study critical percolation on the UIHPT. The important point for us is that conditionally on the map begin discovered up to a certain point using this peeling procedure, it is straightforward to see that the rest of the surface that remains to be discovered has again the law of the UIHPT. This has an analogue for all FK models, $q \in (0, 4)$.

Of course, the above discussion only suggests that a nice coupling between the GFF and SLE must exist, when taking the scaling limit. But the strength of the logarithmic singularity at zero (the coefficient $2/\gamma = 2/\sqrt{\kappa}$) cannot be guessed using the above arguments. Instead, this comes out of the calculations in the proof. In turn, this suggests that the uniform measure on the vertices of the UIHPT, embedded conformally (say using circle packing) into $\mathbb{H}$, converge in the scaling limit to $e^{\gamma h_0}$ rather than $e^{\gamma h}$ - in other words, this logarithmic singularity must be present in the scaling limit of discrete map models.

Proof of Theorem 6.1. First, the idea is to use the *reverse Loewner flow* rather than the ordinary Loewner flow (g_t) and its centered version $\tilde{g}_t(z) = g_t(z) - \xi_t$. Recall that while $\tilde{g}_t(z) : H_t \to \mathbb{H}$ satisfies the SDE:

$$d\tilde{g}_t(z) = \frac{2}{\tilde{g}_t(z)} dt - d\xi_t.$$

In contrast, the reverse Loewner flow is the map $f_t : \mathbb{H} \to H_t := f_t(\mathbb{H})$ defined by the SDE:

$$df_t(z) = -\frac{2}{f_t(z)} dt - d\xi_t.$$

Note the change of signs in the dt term, which corresponds to a change in the direction of time. This Loewner flow is building the curve from the ground up rather than from the tip. More precisely, in the ordinary (forward) Loewner flow, an unusual increment for $d\xi_t$ will be reflected in an unusual behaviour of the curve near its tip at time t. But in the reverse Loewner flow, this increment is reflected in an unusual behaviour near the origin. Furthermore, by using the fact that for any fixed time $T > 0$, the process $(\xi_T - \xi_{T-t})_{0 \le t \le T}$ is a Brownian motion with variance κ run for time $T > 0$, the reader can check that $f_T = \tilde{g}_T^{-1}$ in distribution.

Lemma 6.4. *Let*

$$M_t = M_t(z) := \frac{2}{\sqrt{\kappa}} \log |f_t(z)| + Q \log |f_t'(z)|.$$

Then $M(z)$ is a local martingale if and only if $\kappa = \gamma^2$. Furthermore, if $z, w \in \mathbb{H}$, then the quadratic cross variation between $M(z)$ and $M(w)$ satisfies

$$d[M(z), M(w)]_t = 4\Re(\frac{1}{f_t(z)})\Re(\frac{1}{f_t(w)})dt$$

Proof. Set $Z_t = f_t(z)$. Then $dZ_t = -2/Z_t dt - d\xi_t$. Set $M_t^* = \frac{2}{\sqrt{\kappa}} \log f_t(z) + Q \log f_t'(z)$, so $M_t = \Re(M_t^*)$. Then applying Itô's formula we get

$$\begin{aligned} d \log Z_t &= \frac{dZ_t}{Z_t} - \frac{1}{2}\frac{d[\xi]_t}{Z_t^2} \\ &= -\frac{d\xi_t}{Z_t} + \frac{1}{Z_t^2}(-2 - \kappa/2)dt. \end{aligned}$$

and

$$df_t'(z) = 2\frac{f_t'(z)}{Z_t^2}dt$$

so

$$d \log f_t'(z) = \frac{df_t'(z)}{f_t'(z) - \xi_t} = \frac{2}{Z_t^2}dt.$$

Adding the two pieces together we find:

$$dM_t^* = -\frac{2d\xi_t}{\sqrt{\kappa} Z_t} + \frac{1}{Z_t^2}\left(\frac{2}{\sqrt{\kappa}}(-2 - \kappa/2) + 2Q\right)dt. \tag{6.2}$$

The dt term vanishes if and only if $2/\sqrt{\kappa} + \sqrt{\kappa}/2 = Q$. Clearly this happens if and only if $\gamma = \sqrt{\kappa}$ given the range of these two parameters.

Furthermore, taking the real part in (6.2), if z, w are two points in the upper half plane $\mathbb{H}$, then the quadratic cross variation between $M(z)$ and $M(w)$ is a process which can be identified as

$$d[M(z), M(w)]_t = 4\Re(\frac{1}{f_t(z)})\Re(\frac{1}{f_t(w)})dt$$

so Lemma 6.4 follows. □

One elementary but tedious calculation shows that if $G_t(z,w) = G(f_t(z), f_t(w))$ then $G_t(z,w)$ is a finite variation process and furthermore:

Lemma 6.5. *We have that*

$$dG_t(z,w) = -4\Re(\frac{1}{f_t(z)})\Re(\frac{1}{f_t(w)})dt.$$

In particular, $d[M(z), M(w)]_t = -dG_t(z,w)$.

Proof. We encourage the reader to skip this proof, which is included only for completeness. (However, the result itself will be quite important in what follows.)

Set $X_t = f_t(z)$ and $Y_t = f_t(w)$. From the definition of the Neumann Green function,

$$\begin{aligned} dG_t(x,y) &= -d\log(|X_t - \bar{Y}_t|) - d\log(|X_t - Y_t|) \\ &= -\Re(d\log(X_t - \bar{Y}_t)) - \Re(d\log(X_t - Y_t)). \end{aligned}$$

Now, $dX_t = (2/X_t)dt - d\xi_t$ and $dY_t = (2/Y_t) - d\xi_t$ so taking the difference

$$d(X_t - Y_t) = \frac{2}{X_t}dt - \frac{2}{Y_t}dt = 2\frac{Y_t - X_t}{X_t Y_t}dt$$

and so

$$d\log(X_t - Y_t) = -\frac{2}{X_t Y_t}dt; \quad d\log(X_t - \bar{Y}_t) = -\frac{2}{X_t \bar{Y}_t}dt.$$

Thus we get

$$dG_t(x,y) = -2\Re(\frac{1}{X_t Y_t} + \frac{1}{X_t \bar{Y}_t})dt. \tag{6.3}$$

Now, observe that for all $x, y \in \mathbb{C}$,

$$\frac{1}{xy} + \frac{1}{x\bar{y}} = \frac{\bar{x}\bar{y} + \bar{x}y}{|xy|^2} = \frac{\bar{x}(\bar{y} + y)}{|xy|^2} = \frac{2\Re(y)}{|xy|^2}\bar{x}.$$

Therefore, plugging into (6.3)

$$dG_t(x,y) = -4\frac{\Re(X_t)\Re(Y_t)}{|X_t Y_t|^2} = -4\Re(\frac{1}{X_t})\Re(\frac{1}{Y_t})$$

as desired. □

Equipped with the above two Lemmas, we prove Theorem 6.1. Set $h_0 = \tilde{h} + \frac{2}{\gamma}\log|z|$, and let f_t be an independent reverse Loewner flow as above. Define

$$h_t = h_0 \circ f_t + Q\log|f_t'|.$$

Then, viewed as a distribution modulo constants, we claim that:

$$h_t \text{ has the same distribution as } h_0. \tag{6.4}$$

Let ρ be a test function with zero average, so $\rho \in \tilde{\mathcal{D}}(\mathbb{H})$. To prove (6.4), it suffices to check that (h_t, ρ) is a centered Gaussian with variance as in Theorem 5.7, i.e., $\sigma^2 = \int \rho(dz)\rho(dw)G(z,w)$ (see Remark 5.8).

Now, taking conditional expectations given $\mathcal{F}_t = \sigma(\xi_s, s \le t)$,

$$\begin{aligned}\mathbb{E}[e^{i(h_t,\rho)}] &= \mathbb{E}[e^{i(\frac{2}{\gamma}\log|f_t|+Q\log|f_t'|,\rho)} \times \mathbb{E}[e^{i(h_0\circ f_t,\rho)}|\mathcal{F}_t] \\ &= \mathbb{E}[e^{iM_t(\rho)} \times e^{-\frac{1}{2}\int \rho(dz)\rho(dw)G_t(z,w)}]\end{aligned}$$

where $M_t(\rho) = \int M_t(z)\rho(z)dz$, and we have used that given $\mathcal{F}_t$, $h \circ f_t$ is (by conformal invariance) a Neumann GFF in H_t, with covariance given precisely by G_t (also by conformal invariance). Now, by an application of the conditional Fubini theorem (which would require a bit of justification), $M_t(\rho)$ is also a martingale and furthermore its quadratic variation is

$$d[M(\rho)]_t = \int \rho(dz)\rho(dw)d[M(z), M(w)]_t.$$

Hence by Lemma 6.5,

$$\int \rho(dx)\rho(dy)G_t(x,y) = \int \rho(dz)\rho(dw)G(z,w) - [M(\rho)]_t.$$

Therefore,

$$\mathbb{E}[e^{i(h_t,\rho)}] = e^{-\frac{1}{2}\int \rho(dz)\rho(dw)G(z,w)}\mathbb{E}[e^{iM_t(\rho)+\frac{1}{2}[M(\rho)]_t}].$$

By Itô's formula, $e^{iM_t(\rho)+\frac{1}{2}[M(\rho)]_t}$ is an exponential local martingale, and it is not hard to see that it is a true martingale. We deduce that the expectation in the right hand side above is equal to 1, and therefore

$$\mathbb{E}[e^{i(h_t,\rho)}] = e^{-\frac{1}{2}\int \rho(dz)\rho(dw)G(z,w)}.$$

This proves (6.4). Then arguing that f_t and $\tilde{g}_t^{-1}$ have the same distribution finishes the proof of the theorem. □

Remark 6.6. The reason why we expect M_t to be a martingale for each z is that it is the expected height of the field h_t at z given the filtration generated by the Loewner flow (or equivalently the driving Brownian motion) up to time t.

Remark 6.7. As mentioned earlier, since the proof relies on martingale computation and the optional stopping theorem, the theorem remains true if T is a (bounded) stopping time of the reverse Loewner flow.

Remark 6.8. This martingale is obtained by taking the real part of a certain complex martingale. Taking its imaginary part (in the case of the forward flow) gives rise to the imaginary geometry developed by Miller and Sheffield in a striking series of papers.

6.2 Quantum length of SLE

We start with one of the main theorems of this section, which allows us, given a chordal SLE_κ curve and an independent Neumann GFF, to define a notion of quantum length of the curve unambiguously. The way this is done is by mapping the curve away using the centered Loewner map $\tilde{g}_t$, and using the boundary measure ν introduced above. However, when we map away the curve using the map g_t, each point of the curve corresponds to two points on the boundary (except for the tip of the curve, which is sent to the origin since we consider the centered map). Hence, to measure the length of the curve, we need to know that measuring on one side of 0 gives almost surely the same answer as measuring on the other side of 0.

This is basically the content of the next theorem. For ease of proof, the theorem is stated in the case where h is not a Neumann GFF but rather a certain wedge. However, since these two fields are mutually absolutely continuous with respect to one another, we will see that this is no loss of generality.

Theorem 6.9. *Let $0 < \gamma < 2$. Let $(\mathbb{H}, h, 0, \infty)$ be an α-quantum wedge where $\alpha = \gamma - 2/\gamma$. Let η be an independent SLE_κ with $\kappa = \gamma^2$. Let $\tilde{g}_t$ be the Loewner flow, and let $h_t = h \circ g_t^{-1} + Q \log |(g_t^{-1})'|$ as before, and let ν_{h_t} be the boundary Liouville measure in $\mathbb{H}$ associated with h_t. Given a point $z \in \eta[0,t]$ let $z^- < z^+$ be the two images of z under $\tilde{g}_t$. Then*

$$\nu_{h_t}([z^-, 0]) = \nu_{h_t}([0, z^+]),$$

almost surely for all $z \in \eta[0,t]$.

Corollary 6.10. *The same result is true with h replaced by a Neumann GFF on $\mathbb{H}$ with an arbitrary normalisation.*

Definition 6.11. *The quantity is $\nu_{h_t}([z^-, 0]) = \nu_{h_t}([0, z^+])$ is called the **quantum length** of $\eta([s,t])$ where $z = \eta(s)$.*

False proof of Theorem 6.9. The following argument does not work but helps explain the idea and why wedges are a useful notion. Since $\kappa < 4$ we can consider η to be the infinite curve. Let $L(t)$ be the quantum length of left hand side of the curve η up to time t (measured by computing the boundary quantum length on the left of zero after applying the map $\tilde{g}_t$) and likewise, let $R(t)$ be the quantum length of the right-hand side of η. Then it is *tempting* (but wrong) to think that, because SLE is stationary via the domain Markov property, and the Neumann GFF is invariant by Theorem 6.1, $L(t)$ and $R(t)$ form processes with stationary increments. If that were the case, we would conclude from Birkhoff's ergodic theorem for stationary increments processes that $L(t)/t$ converges almost surely to a possibly random constant, and $R(t)/t$ converges also to a random constant. We would deduce that $L(t)/R(t)$ converges to a possibly random constant. This constant can in fact not be random because of tail triviality of SLE (i.e., of driving Brownian motion) and must be one because of left-right symmetry. On the other hand by scale invariance, the distribution of $L(t)/R(t)$ is constant. Hence we would deduce $L(t) = R(t)$.

This proof is wrong on at least two counts: first of all, it is not true that $L(t)$ and $R(t)$ have stationary increments. This does not hold for instance because h loses its stationarity (i.e., the relation $h_T = h_0$ in distribution does not hold) as soon as a normalisation is fixed for the Neumann GFF. Likewise the scale invariance does not hold in this case. This explains the importance of the concept of wedges, for which a certain form of stationarity holds (see Proposition 6.19) and so does scale invariance by definition, allowing us to make the above proof rigorous.

Proof of Theorem 6.9. The proof is quite tricky, and consists of several steps. Some of these will only be sketched.

Step 0: orientation (*readers who prefer to cut straight to the point can look at Figure 12 and skip this step*). Theorem 6.1 considered a field h_0 with law (6.1) and an independent curve η, and gave a description of the surface obtained by removing a piece of the curve η. We then end up with a new surface h_t which, when parametrised by $\mathbb{H}$ (using the Loewner flow $\tilde{g}_t$), had the same law as the original surface. We will call this "cutting" the surface $(\mathbb{H}, h)$, or zipping down.

Now we want to define the reverse procedure: we want to start with the surface $(\mathbb{H}, h_t)$ and glue pieces of its boundary back together (using a reverse Loewner flow f_t). This will form a surface equipped with a curve: this curve indicates where the two pieces of the boundary of $(\mathbb{H}, h_t)$ that have been glued together can be found. We call this process "welding" the surface $(\mathbb{H}, h_t)$ to itself, or zipping up: see the accompanying figure. ("Up" and "down" refer to the fact that the imaginary part of the Loewner flow which achieves the desired operation is increasing or decreasing).

It is a priori unclear how to define this procedure. If we already knew the theorem we are trying to prove, we could decide that this surface is obtained from $(\mathbb{H}, h_T)$ by "conformally welding" $\mathbb{H}$ to itself by welding parts of its boundary whose quantum length agree (see the discussion in Remark 6.2). Of course we are not allowed to do this yet. The problem might therefore look intractable. However, we can proceed in distribution for now: there is a well defined joint distribution of (h_T, h_0) and hence a well-defined conditional distribution of h_0 given h_T. More precisely, we will do the following. First of all, because of what we are trying to do, it is more natural to call h_T, $\hat{h}_0$, and to call h_0, $\hat{h}_T$. The reason being that, in this perspective, T is arbitrary and you could try to weld what we now call $\hat{h}_0$ as much as you like, all the way up to time ∞. (The $\hat{h}$ notation indicates that we are working in this perspective, starting from $\hat{h}_0$ and trying to weld it to obtain $\hat{h}_T$.)

Step 1: the capacity zipper. To put the discussion in Step 0 on a mathematical footing, let $T > 0$, and let h_0 be as in the original Theorem 6.1 (i.e. h_0 has the distribution (6.1)), and let $\eta = \eta^0$ be an independent infinite SLE_κ curve from 0 to ∞. As in the coupling theorem, set $h_t = h_0 \circ \tilde{g}_t^{-1} + Q \log |(\tilde{g}_t^{-1})'|$, and let η^t be the image by $\tilde{g}_t$ of the initial infinite curve $\eta = \eta^0$. Then Theorem 6.1 says that $h_t = h_0$ in distribution and in fact we can also see that the joint distribution (h_t, η^t) is identical to that of (h_0, η^0).

For $0 \leq t \leq T$, let $\hat{h}_t = h_{T-t}$, and let $\hat{\eta}^t = \eta^{T-t}$. Then it is an easy consequence of Theorem 6.1 that the following lemma holds:

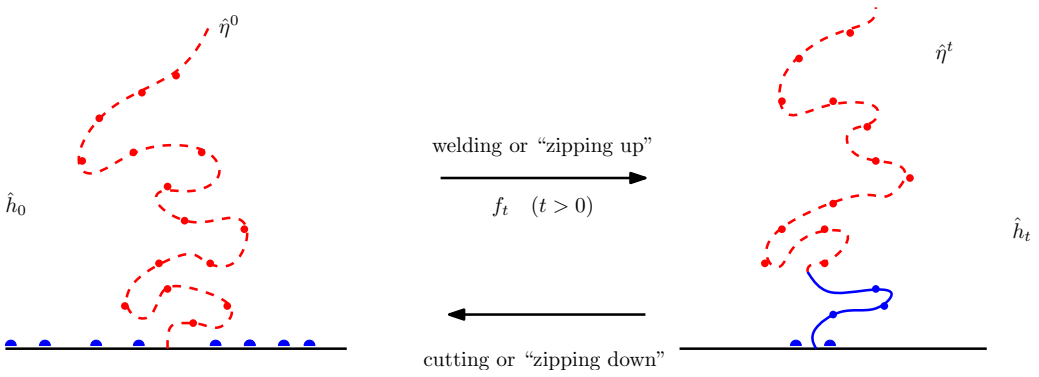

Figure 12: Zipping up and down

Lemma 6.12. *The laws of $(\hat{h}_t, \hat{\eta}^t)_{0\le t\le T}$ are consistent as T increases.*

By Lemma 6.12, and applying Kolmogrov's extension theorem, it is obvious that there is a well-defined process $(\hat{h}_t, \hat{\eta}^t)_{0\le t<\infty}$ whose restriction to $[0, T]$ agrees with the process described above. Hence starting from $\hat{h}_0$ and an infinite curve $\hat{\eta}^0$, there is a well-defined welding procedure giving rise to $\hat{h}_t$. This field is obtained by applying a reverse Loewner flow $(f_s)_{s\le t}$ and applying the change of coordinates formula to $\hat{h}_0$, but we stress that here the reverse Loewner flow is not independent of $\hat{h}_0$ (rather, it will end up being uniquely determined by $\hat{h}_0$, and will be independent of $\hat{h}_t$).

But we could also go in the other direction, cutting $\mathbb{H}$ along $\hat{\eta}^0$, as in Theorem 6.1. Indeed we could define, for $t < 0$ this time, a field $\hat{h}_t$ by considering the centered Loewner flow $(\hat{g}_{|t|})_{t<0}$ associated to the infinite curve $\hat{\eta}^0$, and setting

$$\hat{h}_t = \hat{h}_0 \circ \hat{g}_{|t|}^{-1} + Q \log |(\hat{g}_{|t|}^{-1})'| \quad (t < 0).$$

We can also, of course, get a new curve $\hat{\eta}^t$ by pushing $\hat{\eta}^0$ through the map $\hat{g}_{|t|}$. This gives rise to a two-sided process $(\hat{h}_t, \hat{\eta}^t)_{t\in\mathbb{R}}$.

Definition 6.13. *The process $(\hat{h}_t, \hat{\eta}^t)_{t\in\mathbb{R}}$ is called the (capacity) zipper. We call its law Γ.*

Remark 6.14. An equivalent way to define this process would be as follows. Start from the setup of Theorem 6.1: thus h_0 is a field distributed as in (6.1), and η^0 an independent infinite SLE_κ curve. Set $h_t = h_0 \circ \tilde{g}_t^{-1} + Q \log |(\tilde{g}_t^{-1})'|$ as before, and $\eta^t = g_t(\eta^0 \setminus \eta^0[0, t])$. Then Theorem 6.1 tells us that $(h_t, \eta^t)_{t\ge 0}$ is a stationary process, so we can consider the limit as $t_0 \to \infty$ of $(h_{t_0+t}, \eta^{t_0+t})_{t\ge -t_0}$ which defines a two-sided process $(h_t, \eta^t)_{t\in\mathbb{R}}$. The capacity zipper process is then defined as $(\hat{h}_t, \hat{\eta}^t)_{t\in\mathbb{R}}$, where $\hat{h}_t = h_{-t}$ and $\hat{\eta}^t = \eta_{-t}$.

Thus given a field $\hat{h}_0$ and an independent SLE_κ infinite curve $\hat{\eta}^0$, we can either "zip it up" (weld it to itself) to obtain the configuration $(\hat{h}_t, \hat{\eta}^t)$ for some $t > 0$, or "zip it down" (cut it open along $\hat{\eta}^0$) to obtain the configuration $(\hat{h}_t, \hat{\eta}^t)$ for some $t < 0$. (The term 'capacity'

in the definition refers to the fact that we are zipping up a curve with $2t$ units of half-plane capacity[2].)

Martingales. In practice it is only the direction $t > 0$ of the capacity zipper, where we zip up the curve, which is of interest to us for the rest of the proof. When $t > 0$, the field $\hat{h}_t$ and the curve $\hat{\eta}^t$ are obtained from $(\hat{h}_0, \hat{\eta}^0)$ by a reverse Loewner flow f_t. More precisely, we have

$$\hat{h}_t = \hat{h}_0 \circ f_t^{-1} + Q \log |(f_t^{-1})'| \quad (t > 0), \tag{6.5}$$

and $\hat{\eta}^t = f_t(\hat{\eta}^0)$. This equality holds in the sense of distribution modulo constants on H_t and therefore on $\mathbb{H}$ as well.[3]

From the proof of Theorem 6.1 we have that

$$M_t(z) = \frac{2}{\gamma} \log |f_t(z)| + Q \log |f_t'(z)|; \quad t \geq 0 \tag{6.6}$$

is a local martingale satisfying

$$d[M(z), M(w)]_t = -dG(f_t(z), f_t(w)). \tag{6.7}$$

This is immediate from Lemmas 6.4 and 6.5, once we observe that when we do the relabeling relabel $\hat{h}_t = h_{T-t}$ to go from the setup of the proof of Theorem 6.1 to the capacity zipper studied here, the reverse Loewner flow f_t remains the same. It is only the 'name' of h which changes. These martingales play a crucial role in the following, as they did in the proof of Theorem 6.1.

The stationarity of Theorem 6.1 immediately implies:

Proposition 6.15. *The law Γ of the capacity zipper is stationary. That is, if $(\hat{h}_t, \hat{\eta}^t)_{t\in\mathbb{R}}$ has the law Γ, then the law of $(\hat{h}_{t_0+t}, \hat{\eta}^{t_0+t})_{t\in\mathbb{R}}$ is also Γ, for any $t_0 \in \mathbb{R}$.*

Step 2: reweighting. Let $(\hat{h}_t, \hat{\eta}^t)_{t\in\mathbb{R}}$ be a capacity zipper as in the previous step. Given $\hat{h}_0$ we will sample a point x from the boundary measure $\nu_{\hat{h}_0}$, restricted (say) to the interval $[1, 2]$. More precisely, we consider the law

$$dQ((\hat{h}_t, \hat{\eta}^t)_{t\in\mathbb{R}}, x) = \frac{1}{Z} 1_{x\in[1,2]} d\nu_{\hat{h}_0}(x) d\Gamma((\hat{h}_t, \hat{\eta}^t)_{t\in\mathbb{R}})$$

where Γ is, as before, the law of the capacity zipper. (This reweighting of the law is analogous to the argument used to describe the GFF viewed from a Liouville typical point, see Theorem

[2]Once Theorem 6.9 is proved, it is perhaps more natural to consider the 'quantum zipper', which is the same object but parametrised by quantum length as opposed to half-plane capacity. This is the point of view taken by Sheffield in [She10a], and it will appear naturally in the proof here as well (see Definition 6.18).

[3]To be more precise, we know from Theorem 6.1 that

$$\hat{h}_0 = \hat{h}_t \circ \tilde{f}_t + Q \log |f_t'|.$$

This equality holds as distributions modulo constants on $\mathbb{H}$, which implies that (6.5) holds on H_t by applying f_t^{-1}. Once this equality holds on H_t it follows that it also holds as distributions modulo constants on $\mathbb{H}$, even though a priori the right hand side is not necessarily defined on $\mathbb{H}$.

2.4). As in this proof, we can reverse the order in which $(\hat{h}, \hat{\eta})$ and x are sampled: thus x will be chosen according to the uniform distribution on $[1, 2]$ and given x, $\hat{h}_0$ has an extra $\gamma -$ log singularity at x. That is, under $Q(d\hat{h}_0|x)$, we can write

$$\hat{h}_0 = \tilde{h} + \frac{2}{\gamma} \log |\cdot| + \gamma \log(\frac{1}{|x - \cdot|}) \tag{6.8}$$

where $\tilde{h}$ is a Neumann GFF with some normalisation. When we zip up $\hat{h}_0$ (recall that, when we do so, we obtain a curve which, informally at the moment, describes the welding of $\mathbb{H}$ to itself where pieces of the boundary with the same quantum length), the presence of the extra log singularity at x is bound to influence the geometry of the curve $(\hat{\eta}^t(s))_{0 \leq s \leq t}$ for $t > 0$. We now explain that indeed, instead of being a regular SLE_κ as is normally the case for the capacity zipper, we obtain a curve which is a reverse $\mathrm{SLE}_\kappa(\rho)$, with force point at x.

Lemma 6.16. *Let $\tau = \inf\{t \geq 0 : f_t(x) = 0\}$. The law of $(\hat{\eta}^t[0, t])_{0 \leq t \leq \tau}$ under $dQ(\cdot|x)$ is that of a reverse $SLE_\kappa(\rho)$ flow with force point at x and $\rho = \kappa$. Furthermore, under $dQ(\cdot|x)$, the law of $\hat{h}_t$ (viewed as a distribution modulo constants), conditioned on f_t and $0 \leq t \leq \tau$, is given by*

$$\hat{h}_t = \tilde{h} + \frac{2}{\gamma} \log |z| + \frac{\gamma}{2} G(f_t(x), z), \tag{6.9}$$

where $\tilde{h}$ is an independent Neumann GFF.

Proof. We start with the description of the flow as a reverse $\mathrm{SLE}_\kappa(\rho)$ Loewner flow. To see why this is the case, we make the following comments. Suppose we try to understand the law of the pair $(\hat{h}_t, \hat{\eta}_t)$ biased by $e^{(\gamma/2)(\hat{h}_0, \mu)}/Z$, where μ is some given measure (think of the uniform distribution on some small circle around x), and Z a normalisation constant. Let us call Q_μ this law. Then, from (6.5) we have that

$$\hat{h}_0 = \hat{h}_t \circ f_t + Q \log |f_t'|$$

and by the stationarity of the capacity zipper in the unbiased measure Γ (Proposition 6.15), f_t is independent from $\hat{h}_t$ (indeed this is true at time 0, and recall that this comes from the setup of Theorem 6.1, where the SLE is assumed to be independent from the Neumann GFF). Furthermore we can write $\hat{h}_t = \tilde{h} + \frac{2}{\gamma} \log |\cdot|$ where under Γ, $\tilde{h}$ is a Neumann GFF (independent of f_t). Moreover, the reverse Loewner flow satisfies

$$df_t(z) = \frac{-2}{f_t(z)} - d\xi_t$$

as before. Once we bias the law Γ by $e^{(\gamma/2)(\hat{h}_0, \mu)}/Z$, then the law of $(\xi_t)_{t \geq 0}$ is no longer that of a Brownian motion: like everything else, its law is biased by $e^{(\gamma/2)(\hat{h}_0, \mu)}/Z$. Thus if F is any bounded functional,

$$\mathbb{E}_{Q_\mu}(F(\xi)) = \mathbb{E}\left(F(\xi) \frac{e^{\frac{\gamma}{2}(\hat{h}_0, \mu)}}{Z}\right)$$

$$= \mathbb{E}\left(F(\xi)\frac{e^{\frac{\gamma}{2}(\hat{h}_t\circ f_t + Q\log|f_t'|,\mu)}}{Z}\right)$$

Taking the conditional expectation given the filtration $\mathcal{F}_t$ generated by $(\xi_s)_{s\leq t}$ (or equivalently $(f_s)_{s\leq t}$):

$$\mathbb{E}_{Q_\mu}(F(\xi)) = \mathbb{E}\left(F(\xi)\frac{e^{\frac{\gamma}{2}M_t(\mu)}}{Z}\mathbb{E}(e^{\frac{\gamma}{2}(\tilde{h}\circ f_t,\mu)}|\mathcal{F}_t)\right)$$

where $M_t(\mu) = \int M_t(z)\mu(z)dz$ and by definition, $M_t(z) = \frac{2}{\gamma}\log|f_t(z)| + Q\log|f_t'(z)|$ is the martingale from (6.6) (itself coming from Lemma 6.4).

Now, given $\mathcal{F}_t$ using Theorem 5.7 and a change of variables we have that

$$\mathbb{E}(e^{\frac{\gamma}{2}(\tilde{h}\circ f_t,\mu)}|\mathcal{F}_t) = \exp(\frac{\gamma^2}{4}\int \mu(dx)\mu(dy)G_t(x,y))$$

where $G_t(x,y) = G(f_t(x), f_t(y))$. However, we also know by (6.7) (see Lemma 6.5) that if $[M_t(x), M_t(y)]$ denotes the quadratic cross variation between the martingales $M_t(x)$ and $M_t(y)$ then

$$d[M_t(x), M_t(y)] = -dG_t(x,y).$$

Consequently, applying a Fubini argument, we see that $M_t(\mu)$ is a martingale and that

$$\begin{aligned}\mathbb{E}_{Q_\mu}(F(\xi)) &= \frac{1}{Z}\mathbb{E}\left(F(\xi)e^{\frac{\gamma}{2}M_t(\mu)-\frac{\gamma^2}{4}[M_t(\mu)]+\frac{\gamma^2}{4}\int\mu(dx)\mu(dy)G(x,y)}\right)\\ &= \frac{1}{Z'}\mathbb{E}\left(F(\xi)e^{\frac{\gamma}{2}M_t(\mu)-\frac{\gamma^2}{4}[M_t(\mu)]}\right)\end{aligned}$$

where $Z' = Z/e^{\frac{\gamma^2}{4}\int\mu(dx)\mu(dy)G(x,y)}$ is a different normalisation constant. But taking $F = 1$ and noting that the exponent is an exponential martingale, we see that $Z' = 1$.

Furthermore, we see that the law of ξ under the new measure is nothing but that of a Brownian motion weighted by an exponential martingale. Hence by Girsanov's theorem (the "sophisticated" one about Brownian motion, rather than the elementary Lemma 2.5 proved in these notes), under Q_μ, the process $\tilde{\xi}_t = \xi_t + [\frac{\gamma}{2}M_t(\mu), \xi_t]$ is a local martingale and its quadratic variation is unchanged. By Lévy's theorem, this martingale is therefore a Brownian motion with variance κ. Moreover, by Lemma 6.4

$$dM_t(\mu) = -\frac{2}{\sqrt{\kappa}}\int\mu(dz)\Re(\frac{1}{f_t(z)})d\xi_t,$$

and thus

$$d[\frac{\gamma}{2}M_t(\mu), \xi_t] = -\int\mu(dz)\Re(\frac{1}{f_t(z)})\kappa dt.$$

Consequently, under Q_μ, we may write

$$\xi_t = \tilde{\xi}_t + \int_0^t\int\mu(dz)\Re(\frac{1}{f_s(z)})\kappa ds \tag{6.10}$$

where $\tilde{\xi}$ is a Brownian motion of variance κ. Taking $\mu = \mu^\varepsilon$ the uniform distribution[4] on a semicircle of radius ε about x, and letting $\varepsilon \to 0$, we find that under the biased law,

$$\xi_t = \tilde{\xi}_t + \int_0^t \frac{\kappa}{f_s(x)} ds,$$

since $f_t(x)$ is real. This is the defining property of a (reverse) $\mathrm{SLE}_\kappa(\rho)$ with $\rho = \kappa$, and force point at x: in particular, if $X_t = f_t(x)$, then we must have

$$\begin{aligned} dX_t &= \frac{-2}{X_t} dt - d\xi_t \\ &= \sqrt{\kappa} dB_t + (-2+\rho)/X_t dt \end{aligned}$$

so that $\tilde{X}_t = X_t/\sqrt{\kappa}$ is a Bessel process of dimension

$$\delta = 1 + 2(\rho - 2)/\kappa. \tag{6.11}$$

We no turn to (6.9). We wish to understand the law of $\hat{h}_t$ biased by $e^{(\gamma/2)(\hat{h}_0, \mu)}$, where $\hat{h}_0 = \hat{h}_t \circ f_t + Q \log |f_t'|$, given f_t. But if f_t is given, then the weighting is simply proportional to $e^{(\gamma/2)(\hat{h}_t \circ f_t, \mu)}$, in which case (recalling the elementary Girsanov Lemma 2.5), the effect of this biasing is simply to add to $\hat{h}_t$ a drift term of the form $z \mapsto \int G(z,y) d\mu \circ f_t^{-1}(y)$, in other words the (Neumann) Green function integrated against the image by f_t of the measure μ. When $\mu = \mu^\varepsilon$ is the the uniform distribution on the semi-circle of radius ε about x, and we let $\varepsilon \to 0$, this drift term is simply $z \mapsto G(f_t(x), z)$ as desired. □

Remark 6.17. The conclusion of this lemma remains valid even when applied to a (bounded) stopping time T of the filtration generated by f_t. Indeed the stationarity of the capacity zipper extends to such times (because the proof of the coupling theorem, relying on a martingale computation, extends to such times), and the arguments above are unchanged if $t = T$ is a bounded stopping time.

Step 3: Zooming in to get a wedge and an independent SLE. We now examine closer the law of $(\hat{h}_\tau, \hat{\eta}^\tau)$ under $dQ(\cdot|x)$. On the one hand, thanks to (6.9) (applied to the time $T = \tau \wedge t$, and then letting $t \to \infty$), the field $\hat{h}_\tau$ can be described by

$$\hat{h}_\tau = \tilde{h} + \frac{2}{\gamma} \log |z| + \frac{\gamma}{2} G(0, z)$$

where $\tilde{h}$ is an independent Neumann GFF. In other words, as we zip up $\hat{h}_0$ using the map f_t, the logarithmic singularity which is initially at x travels to $f_t(x)$ until time τ, where

[4] Actually, the above calculation is only valid when μ has average zero, since we have used Theorem 5.7. We can nevertheless apply the conclusion to a positive measure such as μ^ε by introducing a negative unit point mass at a point tending to infinity. Then observe that the drift term in (6.10) does not "feel" the negative mass in the limit.

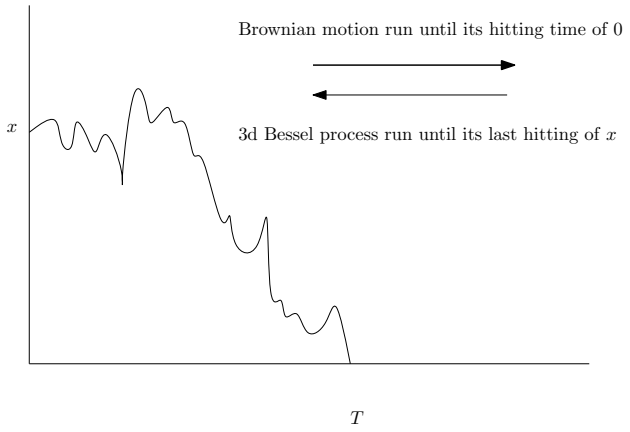

Figure 13: Illustration of Williams' path decomposition theorem

$f_\tau(x)$ itself is at zero. Then recalling the form of the Neumann Green function, we see that $G(0,z) = 2\log 1/|z|$ so that

$$h_\tau = \tilde{h} + \alpha \log 1/|z|; \quad \alpha = \gamma - \frac{2}{\gamma}. \tag{6.12}$$

Therefore, applying Theorem 5.20, we see that zooming near zero, the field converges to an α-thick wedge with $\alpha = \gamma - 2/\gamma$ as in (6.12). (Recall that this means that the law of $h + C/\gamma$, normalised to have unit Liouville mass in the unit disc, converges to that of a wedge as $C \to \infty$.)

We now turn to the curve η^τ, which by Lemma 6.16 we know is that of a reverse $\mathrm{SLE}_\kappa(\rho)$, independent of $\hat{h}_\tau$, and stopped at its continuation threshold $\tau = \inf\{t \geq 0 : X_t = 0\}$, with $X_t = f_t(x)$ the Bessel process measuring the distance between the origin and the (image by f_t of the) force point. As it turns out, this admits a nice description as an ordinary (forward) Loewner flow, which can be seen as follows.

A classical result about Brownian motion, due to Williams (see e.g. Corollary (4.6) in Chapter VII of [RY99]), says that if X is a Brownian motion started from $x > 0$ and T is its hitting time of zero, then its time-reversal $\hat{X} = (X_{T-t})_{0 \leq t \leq T}$ is distributed as three-dimensional Bessel process, run until its last visit Λ to x (see accompanying figure). More generally, if X is a Bessel process of dimension $\delta \in (0,2)$ started from $x > 0$, run until its hitting time of zero T, then its time-reversal $\hat{X} = (X_{T-t}, 0 \leq t \leq T)$ is a Bessel process of dimension $\hat{\delta} = 4 - \delta \in (2,4)$, run until its last visit Λ to x.

This means that the time reversal of X_t at time τ is now a $\hat{\delta}$-dimensional Bessel process, where because of (6.11), and since $\rho = \kappa$, we have that

$$\hat{\delta} = 4 - (1 + \frac{2(\rho - 2)}{\kappa}) = 1 + \frac{4}{\kappa}.$$

But this is precisely the dimension of the Bessel process describing $\tilde{g}_t(x)/\sqrt{\kappa}$ in an ordinary SLE_κ (forward) flow. [In general, this time reversal transforms an $\mathrm{SLE}_\kappa(\rho)$ into an $\mathrm{SLE}_\kappa(\kappa -$

ρ) flow.] Consequently, since time-reversing a reverse flow is precisely the same thing as considering a forward flow driven by the time-reversal, we see that

$$\hat{\eta}^\tau[0,\tau] \text{ has the law of an ordinary SLE}_\kappa\text{, run until } \Lambda$$

where Λ is a certain random time (which is not a stopping time in general), which is a.s. positive. Consequently, as we zoom in near zero, the image of this curve becomes an *infinite* SLE_κ curve (since $\Lambda > 0$ a.s.), which retains in the limit its independence from the wedge obtained as the limit of $\hat{h}_\tau + C/\gamma$ as $C \to \infty$. More precisely, if ψ_C is the scaling map such that $\hat{h}_\tau + C/\gamma$ has unit Liouville mass, then

$$\lim_{C\to\infty} \psi_C(\hat{h}_\tau, \eta^\tau) = (\hat{h}, \hat{\eta}), \tag{6.13}$$

where $\hat{h}$ is, as argued above, an α-thick wedge and $\hat{\eta}$ an independent SLE_κ curve.

Step 4: Stationarity of the wedge quantum zipper. A moment of thought shows that the pair $(\hat{h}, \hat{\eta})$ defined by (6.13) comes with a two-sided process $(\hat{h}_t, \hat{\eta}^t)_{t\geq 0}$ obtained by zipping up $(\hat{h}, \hat{\eta})$, such that $\hat{h}_0 = \hat{h}$ and $\hat{\eta}^0 = \hat{\eta}$. (For instance, this comes from enriching the observed structure when we take the limit $C \to \infty$ above, so that there is a whole zipping up process: see Lemma 6.20 where this idea is spelled out in more detail). As before, this zipper process is parametrised by capacity: to obtain $\hat{\eta}^t$ we have grown a curve of half-plane capacity $2t$ from zero, and mapped the original $\hat{\eta}$ on top of this new piece of curve. In fact, the only difference between this process and the one considered in Definition 6.13 is that the marginal of $\hat{h}_0$ is not a Neumann GFF with some arbitrary normalisation but an α-thick wedge. This is why we also keep the same notation: it is the same process, but with a different law (at time 0, the field is a wedge rather than a GFF with a certain log singularity).

We will now make a change of time paramaterisation, so that time corresponds to a given amount of quantum length being zipped up. Namely, for $t > 0$, let

$$R(t) = \inf\{x > 0; \nu_{\hat{h}_0}([0,x]) > t\}, \tag{6.14}$$

and let

$$\sigma_t = \inf\{s > 0 : f_s(x) = 0\} \quad \text{where} \quad x = R(t).$$

So σ_t is the first time that an interval of quantum length t is being zipped up by f.

Definition 6.18. *The process $(\hat{h}_{\sigma_t}, \hat{\eta}^{\sigma_t})_{t\geq 0}$ is called the (α-thick wedge) quantum zipper.*

Then the crucial fact on which the proof rests is the following:

Proposition 6.19. *The quantum zipper is stationary: if $t_0 \geq 0$, then $(\hat{h}_{\sigma_{t_0+t}}, \hat{\eta}^{\sigma_{t_0+t}})_{t\geq 0}$ is also a quantum zipper.*

Proof. Essentially, the idea is that when we zip up the wedge by t_0 units of quantum length, this corresponds (in the pre-limiting picture), to zipping up $(\hat{h}_\tau, \eta^\tau)$ in Step 3 by a microscopic amount of quantum length: i.e., considering the state of the capacity zipper at the first time

after τ where we have zipped up a further amount of quantum length corresponding to $t_0 e^{-C/2}$.

Alternatively, we could have replaced x by x_C, where x_C is chosen so that

$$\nu_{\hat{h}_0}([x, x_C]) = t_0 e^{-C/2},$$

and we could have just considered the time τ_C (instead of τ) where x_C (and not x) is zipped up to 0, before zooming in. Therefore the claim is that if $\tau_C = \inf\{t \geq 0 : f_t(x_C) = 0\}$ and we zoom in near zero at time τ_C, we obtain the surface $\hat{h}_{\sigma_{t_0}}$, rather than $\hat{h}$:

Lemma 6.20. *We have the convergence of the random surface equipped with a curve*

$$(\hat{h}_{\tau_C} + C/\gamma, \eta^{\tau_C}) \to (\hat{h}^{\sigma_{t_0}}, \eta^{\sigma_{t_0}})$$

(As explained above, this can be taken as a definition of the wedge quantum zipper, so we do not provide additional justification for this fact).

To conclude the proof of Proposition 6.19, we simply observe that given $\hat{h}_0$, the points x and x_C are (almost) indistinguishable: indeed, since x is itself chosen according to the boundary length measure $\nu_{\hat{h}_0}$, the total variation distance between x and x_C tends to 0 (this follows from the fact that if U is uniform in $[0, 1]$ then the total variation distance between U and $U + \varepsilon$ tends to 0 as $\varepsilon \to 0$). Therefore we may couple x and x_C with high probability given $\hat{h}_0$; in which case it is obvious that (on the event that the coupling is successful) the pairs $(\hat{h}_\tau, \hat{\eta}^\tau)$ and $(\hat{h}_{\tau_C}, \hat{\eta}^{\tau_C})$ are identical, and so are their limits when we zoom in near zero. Consequently, the distribution of $(\hat{h}^{\sigma_{t_0}}, \eta^{\sigma_{t_0}})$ must be identical to that of $(\hat{h}, \hat{\eta})$. Stationarity follows. □

To summarise the idea behind the proof of this theorem: if we want something which is invariant under zipping up by a given amount of quantum length, it suffices to sample a point according to quantum length and wait until that point is zipped up at zero. Zooming in near zero necessarily gives us a stationary process – and a scale invariant one. We are now in a position to conclude the proof of Theorem 6.9.

Step 5: conclusion of the proof. Let $(\hat{h}_t, \eta^t)_{t \geq 0}$ be a quantum zipper as in Definition 6.18. For $t \geq 0$, let

$$L(t) = x \text{ such that } f_t[x, 0] = f_t[0, R(t)],$$

where $R(t)$ is defined in (6.14) and is such that the quantum length of $[0, R(t)]$ is t (since $\nu_{\hat{h}_0}$ is nonatomic). Our goal is to show that if we call

$$\lambda_t = \nu_{\hat{h}_0}([L(t), 0])$$

then $\lambda_t = t$ almost surely for all $t > 0$. This will show that intervals to the left and to the right of zero which get zipped up together have the same quantum length.

Observe that, by Proposition 6.19, λ_t has stationary increments: the law of $R(t) - R(s)$ depends only on $t - s$ for $0 \leq s \leq t$. By Birkhoff's ergodic theorem,

$$\lim_{n \to \infty} \frac{\lambda_n}{n} = X$$

exists almost surely. (Note that the theorem is usually stated under the assumption that $\mathbb{E}(|\lambda_1|) < \infty$, but it is straightforward to see, by a truncation argument and the monotone convergence theorem, that the conclusion is also true if we only know $\lambda_1 \geq 0$ a.s., which is the case here.)

Then by scale-invariance of the α-thick wedge (part (i) in Theorem 5.20) we see that the distribution of λ_t/t is constant and hence equal to X for all $t > 0$. In particular as $t \to 0$,

$$\lim_{t\to 0} \frac{\lambda_t}{t} = X$$

almost surely. But it is obvious that the above limit is measurable with respect to the tail σ-algebra

$$\mathcal{T} = \bigcap_{\varepsilon>0} \sigma\{(h,\rho) : \rho \text{ supported in } B(0,\varepsilon)\}.$$

In turn, $\mathcal{T}$ is contained in the tail σ-algebra of the circle average process at zero of an α-thick wedge, which must be trivial (it is identical to the tail at infinity of a Brownian motion with drift by Definition 5.18, which is trivial by Blumenthal's zero-one law). Therefore X is a deterministic constant and we deduce by symmetry that $X = 1$ necessarily. Hence $\lambda_t = t$ for all $t > 0$, almost surely. In other words,

$$\nu_{\hat{h}_0}[L(t), 0] = \nu_{\hat{h}_0}([0, R(t)])$$

for all $t > 0$, almost surely. By absolute continuity between wedges modulo constants and Neumann GFF (modulo constants), away from zero, it follows immediately that the same is true for a regular Neumann GFF, and hence the theorem follows. □

6.3 Uniqueness of the welding

Consider the capacity zipper $(\hat{h}_t, \hat{\eta}^t)_{t\in\mathbb{R}}$ of Definition 6.13. The (reverse) Loewner flow associated to the curves $\hat{\eta}^t$ has the property that it zips together intervals of $\mathbb{R}$ with the same $\nu_{\hat{h}_0}$ quantum length by Theorem 6.9. Could there be any other curves with this property, or equivalently could there be any other Loewner flow with the property that intervals of identical quantum length on either side of zero are being zipped up together?

We will now show that there is a unique such flow, and hence the Loewner flow is entirely determined by $\hat{h}_0$.

Theorem 6.21. *Let $(\hat{h}_t, \hat{\eta}^t)_{t\in\mathbb{R}}$ be a capacity zipper with reverse Loewner flow $(f_t)_{t\geq 0}$. Let $t > 0$. Almost surely, the following holds: if $\tilde{f}_t : \mathbb{H} \to \tilde{H}_t = \tilde{f}_t(\mathbb{H})$ is a conformal map such that*

- *$\tilde{H}_t$ is the complement of a simple curve $\tilde{\eta}^t$,*
- *$\tilde{f}_t$ has the hydrodynamic normalisation $\lim_{z\to\infty} \tilde{f}_t(z) - z = 0$,*
- *$\tilde{f}_t$ has the property that $\tilde{f}_t(z^-) = \tilde{f}_t(z^+)$ as soon as $\nu_{\hat{h}_0}([z^-, 0]) = \nu_{\hat{h}_0}([0, z^+)$,*

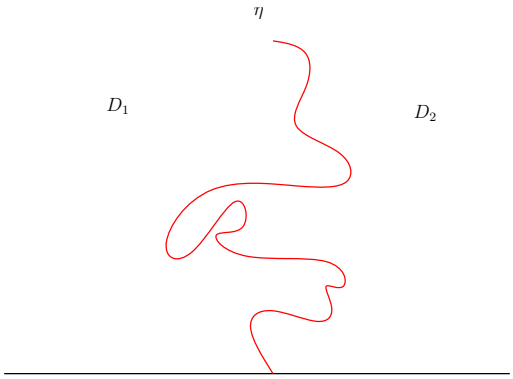

Figure 14: An independent SLE slices an α-thick wedge into two independent γ-thick wedges.

then $\tilde{f}_t = f_t$ *and* $\tilde{\eta}^t = \hat{\eta}^t$. *In particular, the reverse Loewner flow* $(f_t)_{t\geq 0}$ *(and hence* $(\hat{h}_t, \hat{\eta}^t)_{t\geq 0}$*) is uniquely determined by* $\hat{h}_0$ *only.*

Proof. Before we start the proof, we recall that from the definition of the capacity zipper in Definition 6.13, we only have defined the reverse Loewner flow as being coupled to $\hat{h}_0$ in a certain way specified by the application of Kolmogorov's theorem. Usually, proving that objects coupled to a GFF are determined by it can be quite complicated (e.g., this is the case in the setup of imaginary geometry or level lines of the GFF).

Here the proof will turn out to be quite simple given some classical results from the literature. Indeed consider

$$\phi = \tilde{f}_t \circ f_t^{-1}.$$

A priori, ϕ is a conformal map on $f_t(\mathbb{H}) = H_t$, and its image is $\phi(H_t) = \tilde{H}_t$. However, because of our assumptions on $\tilde{f}_t$ (and the properties of f_t), the definition of ϕ can be extended unambiguously to all of $\mathbb{H}$. Moreover when we do so, the extended map is a homeomorphism of $\mathbb{H}$ onto $\mathbb{H}$, which is conformal off the curve η^t. Thus the theorem will be proved if we can show that any such map must be the identity. In the terminology of complex analysis, this is equivalent to asking that the curve $\eta^t([0,t])$ is a *removable* set. Now, by a result of Rohde and Schramm [RS05], the complement H_t of the curve is a.s. a Hölder domain for $\kappa < 4$ (or $\gamma < 2$), and by a result of Jones and Smirnov [JS00] it follows that $\eta^t([0,t])$ is a removable set. Hence the theorem follows. □

6.4 Slicing a wedge with an SLE

In this section we complement our earlier results by the following remarkable result due to Sheffield [She10a]. This result is fundamental to the theory developed in [DMS14a], where the main technical tool is a generalisation of the result below.

Suppose we are given an α-quantum wedge h with $\alpha = \gamma - 2/\gamma$, and an independent SLE_κ curve η with $\kappa = \gamma^2 < 4$, then the curve η slices the wedge into two surfaces (see

picture). Then the result below says that, considered as abstract random surfaces, the two surfaces are independent and they are both γ-thick wedges.

Theorem 6.22. *Suppose we are given an α-quantum wedge h with $\alpha = \gamma - 2/\gamma$, and an independent SLE_κ curve η with $\kappa = \gamma^2 < 4$. Let D_1, D_2 be the two connected component of $\mathbb{H} \setminus \eta$, containing respectively $-\infty$ and $+\infty$. Let $h_1 = h|_{D_1}$ and $h_2 = h|_{D_2}$. Then the two surfaces (D_1, h_1) and (D_2, h_2) are independent γ-thick wedges.*

Sketch of proof. We will rely on elements of the proof of Theorem 6.9. Consider the setup of Step 2 of that proof: we have a capacity zipper $(\hat{h}_t, \hat{\eta}^t)_{t\in\mathbb{R}}$ with law Γ, and a biased law

$$dQ((\hat{h}_t, \hat{\eta}^t)_{t\in\mathbb{R}}, x) = \frac{1}{Z} 1_{x\in[1,2]} d\nu_{h_0}(x) d\Gamma((\hat{h}_t, \hat{\eta}^t)_{t\in\mathbb{R}})$$

Recall that we can reverse the order in which $(\hat{h}, \hat{\eta})$ and x are sampled: thus x will be chosen according to the uniform distribution on $[1, 2]$ and given x, $\hat{h}_0$ has an extra $\gamma-\log$ singularity at x. That is, under $Q(d\hat{h}_0|x)$, we have

$$\hat{h}_0 = \tilde{h} + \frac{2}{\gamma} \log|\cdot| + \gamma \log(\frac{1}{|x - \cdot|}) \tag{6.15}$$

Briefly, we recall that given x, if $\tau = \inf\{t > 0 : f_t(x) = 0\}$, and we zoom in near zero at time τ in $\hat{h}_\tau$, we find an α-thick wedge decorated with an independent $\mathrm{SLE}_\kappa(\rho)$ curve (see (6.13)). We will deduce the desired independence of the two surfaces in Theorem 6.22 by showing that before taking the limit, the two surfaces on either side of the curve η^τ near zero are approximately independent, and that they are each γ-thick wedges.

To make sense of this approach, we start by fixing some arbitrary $\epsilon > 0$ and let $D_2' = f_\tau(B(x, \epsilon))$. Likewise, if $y = R(x)$ denote the point to which x is welded (i.e., by Theorem 6.9, the unique $y < 0$ such that $\nu_{\hat{h}_0}([y, 0]) = \nu_{\hat{h}_0}([0, x])$), then we set $D_1' = f_\tau(B(y, \epsilon))$. Then we claim that the pair of surfaces $(D_1', \hat{h}_\tau + C/\gamma)$ and $(D_2', \hat{h}_\tau + C/\gamma)$ converges as $C \to \infty$ towards the pair of surfaces (D_1, h_1) and (D_2, h_2) in the statement of Theorem 6.22 (see Figure).

But observe that by definition, the two surfaces $(D_1', \hat{h}_\tau + C/\gamma)$ and $(D_2', \hat{h}_\tau + C/\gamma)$ are obtained by conformally mapping the surfaces $(B(y, \epsilon), \hat{h}_0 + C/\gamma)$ and $(B(x, \epsilon), \hat{h}_0 + C/\gamma)$ and applying the change of variable formula. Now, because of (6.15), it is clear that the limit of $(B(x, \epsilon), \hat{h}_0 + C/\gamma)$ (i.e., (D_2, h_2)) must be a γ-quantum wedge. By left-right symmetry, we deduce that (D_1, h_1) is also a quantum wedge. It remains to argue independence between these two surfaces. To do so, we condition on some additional amount of information: x, $L = \nu_{\hat{h}_0}([0, x])$, y (in a moment, we will even add a bit more information). When we do so, it is clear that the law of the field becomes

$$\hat{h}_0 = \tilde{h} + \frac{2}{\gamma} \log|\cdot| + \gamma \log(\frac{1}{|x - \cdot|}) + \gamma \log(\frac{1}{|y - \cdot|}) \tag{6.16}$$

where $\tilde{h}$ is a Neumann GFF conditioned so that $\hat{h}_0$ gives length L to $[0, x]$ and $[y, 0]$. We now condition further on the and the values of the field $\tilde{h}$ outside $B(x, \varepsilon)$ and $B(y, \epsilon)$. Call

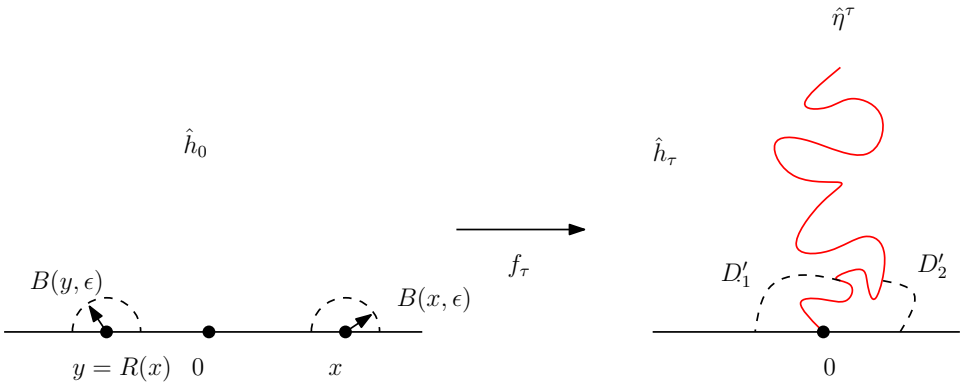

Figure 15: Summary of the idea for the proof of Theorem 6.22: as we zoom in towards zero, the two surfaces to the left and right of $\hat{\eta}^\tau$ come from very different places of the field $\hat{h}_0$ (near x and $y = R(x)$ respectively). The behaviour of the field near these two places is easily shown to be conditionally independent given x, y and the quantum length $\nu_{\hat{h}_0}([0,x])$, and the values of the field outside balls near x and y. Even given this information, the behaviour of the field when we zoom in near these points is still given by a quantum wedge.

$a = \nu_{\hat{h}_0}([0, x-\epsilon])$ and $b = \nu_{\hat{h}_0}([y+\epsilon, 0])$. Then conditionally on all this information, the law of $\tilde{h}$ in $B(x,\epsilon)$ (resp. $B(y,\epsilon)$) is that of an independent Neumann GFF, plus a harmonic function in $B(x,\epsilon)$ (resp. in $B(y,\epsilon)$), and conditioned so that the length of $[x-\varepsilon, x]$ is $L-a$ (resp. the length $[y, y+\epsilon]$ is $L-b$). Clearly, since the harmonic function in $B(x,\epsilon)$ (resp. $B(y,\varepsilon)$) is continuous in that ball and in particular continuous at x (resp. y), when we zoom in near x (resp. y), we still end up with a γ-wedge (see Exercise 8). Likewise, the conditioning that the length of $[x-\epsilon, x]$ is $L-a$ (resp. the length of $[y, y+\epsilon]$ is $L-b$), which is a macroscopic information, does not change the fact that in the limit the surface is a γ-wedge (this would require a justification). Moreover, once we have conditioned on this information recall that the free fields in the two balls are independent, and hence the limiting surfaces are conditionally independent γ-thick wedges, given all this extra information. Therefore this remains true once we remove the conditioning by taking expectations, and the result follows. □

Remark 6.23. Observe that in this proof, to obtain independence of the two surfaces it is crucial to know the conclusion of Theorem 6.9, which allows us (given $L = \nu_{\hat{h}_0}([0,x])$) to find $y = R(x)$ as the only point to the left of zero such that $[y, 0]$ also has a quantum length equal to L. If we did not know this, we could a priori only define y via as the point to which x is welded and this might a priori involve the whole curve, itself dependent a priori on the entire field. Hence it would be very difficult to argue that the local behaviour at x and y are approximately independent, which is what the argument boils down to.

6.5 Exercises

1. Let f_t be a reverse $\mathrm{SLE}_\kappa(\rho)$ Loewner flow, let

$$h_0(\cdot) = \tilde{h}(\cdot) + \frac{2}{\sqrt{\kappa}} \log|\cdot| + \frac{\rho}{2\sqrt{\kappa}} G(x, \cdot)$$

be independent of (f_t), where $G(x, y) = -\log|x - y|$ is the whole plane Green function and $\tilde{h}$ is a Neumann GFF with some normalisation.

For all $0 \leq t \leq \tau$ set

$$h_t^\rho = \tilde{h} \circ f_t + \frac{2}{\sqrt{\kappa}} \log|f_t| + Q \log|f_t'| + \frac{\rho}{2\sqrt{\kappa}} G(f_t(x), f_t(z)).$$

Show that

$$h_{t\wedge\tau}^\rho = h_0$$

in law, as distributions modulo additive constants. (This is Theorem 4.5 in [She10a]). Give two proofs of this result: one based on stochastic calculus in the manner of the proof of Theorem 6.1, and another one using Lemma 6.16.

References

[ABBS13] Elie Aïdékon, Julien Berestycki, Éric Brunet, and Zhan Shi. Branching Brownian motion seen from its tip. *Probability Theory and Related Fields*, 157(1-2):405–451, 2013.

[ABK13] Louis-Pierre Arguin, Anton Bovier, and Nicola Kistler. The extremal process of branching Brownian motion. *Probability Theory and related fields*, 157(3-4):535–574, 2013.

[AC15a] Omer Angel and Nicolas Curien. Percolations on random maps I: half-plane models. *Ann. Inst. H. Poincaré*, 51(2):405–431, 2015.

[AC15b] Omer Angel and Nicolas Curien. Percolations on random maps I: half-plane models. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 51(2):405–431, 2015.

[Ang03] Omer Angel. Growth and percolation on the uniform infinite planar triangulation. *Geometric & Functional Analysis GAFA*, 13(5):935–974, 2003.

[Aru14] Juhan Aru. KPZ relation does not hold for the level lines and SLE_κ flow lines of the Gaussian free field. *Probability Theory and Related Fields*, pages 1–62, November 2014.

[BBG12a] G Borot, J Bouttier, and E Guitter. Loop models on random maps via nested loops: the case of domain symmetry breaking and application to the potts model. *Journal of Physics A: Mathematical and Theoretical*, 45(49):494017, 2012.

[BBG12b] G Borot, Jérémie Bouttier, and E Guitter. Loop models on random maps via nested loops: the case of domain symmetry breaking and application to the Potts model. *Journal of Physics A: Mathematical and Theoretical*, 45(49):494017, 2012.

[BBG12c] Gaetan Borot, Jérémie Bouttier, and Emmanuel Guitter. A recursive approach to the $O(n)$ model on random maps via nested loops. *Journal of Physics A: Mathematical and Theoretical*, 45(4):045002, 2012.

[BBM11] Olivier Bernardi and Mireille Bousquet-Mélou. Counting colored planar maps: algebraicity results. *Journal of Combinatorial Theory, Series B*, 101(5):315–377, 2011.

[BC13] Itai Benjamini and Nicolas Curien. Simple random walk on the uniform infinite planar quadrangulation: subdiffusivity via pioneer points. *Geometric and Functional Analysis*, 23(2):501–531, 2013.

[Bef08] Vincent Beffara. The dimension of the sle curves. *The Annals of Probability*, pages 1421–1452, 2008.

[Ber07] Olivier Bernardi. Bijective counting of tree-rooted maps and shuffles of parenthesis systems. *Electron. J. Combin*, 14(1):R9, 2007.

[Ber08] Olivier Bernardi. A characterization of the tutte polynomial via combinatorial embeddings. *Annals of Combinatorics*, 12(2):139–153, 2008.

[Ber15] Nathanaël Berestycki. An elementary approach to Gaussian multiplicative chaos. *arXiv preprint arXiv:1506.09113*, 2015.

[BGRV14] Nathanael Berestycki, Christophe Garban, Rémi Rhodes, and Vincent Vargas. KPZ formula derived from Liouville heat kernel. *arXiv preprint arXiv:1406.7280*, 2014.

[BJM14] Jérémie Bettinelli, Emmanuel Jacob, and Grégory Miermont. The scaling limit of uniform random plane maps, via the ambjørn–budd bijection. *Electron. J. Probab*, 19(74):1–16, 2014.

[BKW76] Rodney J Baxter, Stewart B Kelland, and Frank Y Wu. Equivalence of the potts model or whitney polynomial with an ice-type model. *Journal of Physics A: Mathematical and General*, 9(3):397, 1976.

[BLR15] Nathanaël Berestycki, Benoît Laslier, and Gourab Ray. Critical exponents on Fortuin–Kasteleyn weighted planar maps. *arXiv preprint arXiv:1502.00450*, 2015.

[BN11] Nathanaël Berestycki and James R. Norris. Lectures on Schramm–Loewner Evolution. 2011.

[BSS14] Nathanaël Berestycki, Scott Sheffield, and Xin Sun. Equivalence of Liouville measure and Gaussian free field. *arXiv preprint arXiv:1410.5407*, 2014.

[Che15] Linxiao Chen. Basic properties of the infinite critical-FK random map. *arXiv preprint arXiv:1502.01013*, 2015.

[CK15] Nicolas Curien and Igor Kortchemski. Percolation on random triangulations and stable looptrees. *Probab. Theor. Rel. Fields*, 163(1):303–337, 2015.

[DMS14a] Bertrand Duplantier, Jason Miller, and Scott Sheffield. Liouville quantum gravity as a mating of trees. *arXiv preprint arXiv:1409.7055*, 2014.

[DMS14b] Bertrand Duplantier, Jason Miller, and Scott Sheffield. Liouville quantum gravity as a mating of trees. In *arXiv preprint arXiv:1409.7055* [DMS14a].

[DRS+14] Bertrand Duplantier, Rémi Rhodes, Scott Sheffield, Vincent Vargas, et al. Critical gaussian multiplicative chaos: convergence of the derivative martingale. *The Annals of Probability*, 42(5):1769–1808, 2014.

[DRSV12] Bertrand Duplantier, Rémi Rhodes, Scott Sheffield, and Vincent Vargas. Renormalization of critical gaussian multiplicative chaos and KPZ formula. *arXiv preprint arXiv:1212.0529*, 2012.

[DS11] Bertrand Duplantier and Scott Sheffield. Liouville quantum gravity and KPZ. *Invent. Math.*, 185(2):333–393, 2011.

[Dub09] Julien Dubédat. SLE and the free field: partition functions and couplings. *Journal of the American Mathematical Society*, 22(4):995–1054, 2009.

[EK95] B Eynard and C Kristjansen. Exact solution of the o (n) model on a random lattice. *Nuclear Physics B*, 455(3):577–618, 1995.

[GJSZJ12] A Guionnet, VFR Jones, D Shlyakhtenko, and P Zinn-Justin. Loop models, random matrices and planar algebras. *Communications in Mathematical Physics*, 316(1):45–97, 2012.

[GMS15] E. Gwynne, C. Mao, and X. Sun. Scaling limits for the critical Fortuin-Kasteleyn model on a random planar map I: cone times. *ArXiv e-prints*, February 2015.

[HK71] Raphael Høegh-Krohn. A general class of quantum fields without cut-offs in two space-time dimensions. *Comm. Math. Phys.*, 21:244–255, 1971.

[HMP10] Xiaoyu Hu, Jason Miller, and Yuval Peres. Thick points of the Gaussian free field. *Ann. Probab.*, 38(2):896–926, 2010.

[HRV15] Yichao Huang, Rémi Rhodes, and Vincent Vargas. Liouville quantum gravity on the unit disk. *arXiv preprint arXiv:1502.04343*, 2015.

[JS00] Peter W Jones and Stanislav K Smirnov. Removability theorems for Sobolev functions and quasiconformal maps. *Arkiv för Matematik*, 38(2):263–279, 2000.

[JS15] Janne Junnila and Eero Saksman. The uniqueness of the Gaussian multiplicative chaos revisited. *arXiv preprint arXiv:1506.05099*, 2015.

[Kah85] Jean-Pierre Kahane. Sur le chaos multiplicatif. *Ann. Sci. Math. Québec*, 9(2):105–150, 1985.

[Kah86] Jean-Pierre Kahane. Une inégalité du type de Slepian et Gordon sur les processus gaussiens. *Israel journal of mathematics*, 55(1):109–110, 1986.

[Ken00] Richard Kenyon. The asymptotic determinant of the discrete laplacian. *Acta Mathematica*, 185(2):239–286, 2000.

[Ken01] Richard Kenyon. Dominos and the Gaussian free field. *Ann. Probab.*, 29(3):1128–1137, 2001.

[KPZ88] Vadim G. Knizhnik, Alexander M Polyakov, and Alexander B. Zamolodchikov. Fractal structure of 2d quantum gravity. *Modern Physics Letters A*, 3(08):819–826, 1988.

[Law05] Gregory F. Lawler. *Conformally invariant processes in the plane*, volume 114 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2005.

[LG13] Jean-François Le Gall. Uniqueness and universality of the Brownian map. *Ann. Probab.*, 41(4):2880–2960, 2013.

[LP] R. Lyons and Y. Peres. *Probability on Trees and Networks*. Cambridge University Press. In preparation. Current version available at `http://mypage.iu.edu/~rdlyons/`.

[LSW00] Gregory F Lawler, Oded Schramm, and Wendelin Werner. The dimension of the planar Brownian frontier is 4/3. *arXiv preprint math/0010165*, 2000.

[LSW04] Gregory F Lawler, Oded Schramm, and Wendelin Werner. Conformal invariance of planar loop-erased random walks and uniform spanning trees. *The Annals of Probability*, 32(1B):939–995, 2004.

[Mie11] G. Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Mathematica, to appear*, 2011. arXiv:1104.1606.

[MS13] Jason Miller and Scott Sheffield. Quantum loewner evolution. *arXiv preprint arXiv:1312.5745*, 2013.

[Mul67] Ronald C Mullin. On the enumeration of tree-rooted maps. *Canad. J. Math*, 19:174–183, 1967.

[Pem91] Robin Pemantle. Choosing a spanning tree for the integer lattice uniformly. *The Annals of Probability*, pages 1559–1574, 1991.

[RS05] Steffen Rohde and Oded Schramm. Basic properties of SLE. *Annals of Mathematics*, pages 883–924, 2005.

[RV+10] Raoul Robert, Vincent Vargas, et al. Gaussian multiplicative chaos revisited. *The Annals of Probability*, 38(2):605–631, 2010.

[RV11] Rémi Rhodes and Vincent Vargas. KPZ formula for log-infinitely divisible multifractal random measures. *ESAIM Probab. Stat.*, 15:358–371, 2011.

[RY99] Daniel Revuz and Marc Yor. *Continuous martingales and Brownian motion*, volume 293 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1999.

[Sha] A. Shamov. On Gaussian multiplicative chaos. arXiv:1407.4418v2.

[She07] Scott Sheffield. Gaussian free fields for mathematicians. *Probab. Theory Related Fields*, 139(3-4):521–541, 2007.

[She10a] Scott Sheffield. Conformal weldings of random surfaces: SLE and the quantum gravity zipper. *arXiv preprint arXiv:1012.4797*, 2010.

[She10b] Scott Sheffield. Conformal weldings of random surfaces: SLE and the quantum gravity zipper. In *arXiv preprint arXiv:1012.4797* [She10a].

[SS13] Oded Schramm and Scott Sheffield. A contour line of the continuum Gaussian free field. *Probab. Theory Related Fields*, 157(1-2):47–80, 2013.

[Wer] Wendelin Werner. Topics on the two-dimensional gaussian free field. Available on the webpage of the author.

[Wil96] David Bruce Wilson. Generating random spanning trees more quickly than the cover time. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 296–303. ACM, 1996.

Index