Universal Randomness in 2D

Jason Miller and Scott Sheffield

Abstract

We begin by reviewing one-dimensional stochastic processes that are *universal* in the sense that they arise in many contexts — in particular as scaling limits of large families of discrete models — and *canonical* in the sense that they are uniquely characterized by scale invariance and other natural symmetries. Examples include Brownian motion, Bessel processes, and stable Lévy processes.

We then introduce several universal and canonical random objects that are *planar* in the sense that they can be either embedded in or parameterized by a two dimensional surface. These objects include trees, distributions, curves, loop ensembles, surfaces, and growth trajectories.

Finally, we discuss the intricate and surprising relationships *between* these universal objects. We explain how to use generalized functions to construct curves and vice versa; how to conformally *weld* a pair of surfaces to produce a surface decorated by a simple curve; how to conformally *mate* a pair of trees to obtain a surface decorated by a non-simple curve; and how to *shuffle* certain mating and welding operations to produce random growth trajectories on random surfaces. We present both discrete and continuum analogs of these constructions.

Several topics in these notes are inspired and motivated by physics, especially string theory, conformal field theory, gauge theory, and statistical mechanics. The mathematics can nonetheless be understood independently of the physical motivation.

Keywords include continuum random tree, stable Lévy tree, stable looptree, Gaussian free field, Schramm-Loewner evolution, percolation, uniform spanning tree, loop-erased random walk, Ising model, FK cluster model, conformal loop ensemble, Brownian loop soup, random planar map, Liouville quantum gravity, Laplacian determinant, Brownian map, Brownian snake, diffusion limited aggregation, first passage percolation, dielectric breakdown model, imaginary geometry, quantum zipper, peanosphere, and quantum Loewner evolution.

Contents

Preface

The goals of these notes are very simple. We will

1. introduce fundamental random objects, and
2. explain how they are related to one another.

The fundamental random objects include processes, trees, distributions (a.k.a. generalized functions), curves, loop ensembles, surfaces, and growth trajectories.

All of these objects are in some sense *universal*. That is, they arise as *macrosopic limits* of many different kinds of random systems, which may have very different *microscopic* behavior. This usage of the term "universal" comes from statistical physics. Physicists tell us that many phenomena (such as phase transitions) are surprisingly similar from one material to another. Informally speaking, physical systems — and mathematical models — that look very different on the microscopic level (different atoms, molecules, etc.) are declared to belong the same *universality class* if they "behave the same way" in some macroscopic limit. The convergence of arbitrary random walks with i.i.d. increments to Brownian motion (when the increment law has zero mean and finite variance) is an example of mathematical universality. We will encounter many other examples during the course of this text, some proven and some conjectural.

The random objects introduced in these notes are also all in some sense *canonical*. Many fundamental objects in mathematics are singled out by special symmetries. For example, in a universe full of roughly round-ish shapes, the sphere is special: it is uniquely determined by rotational invariance, equidistance of points from a center, etc. Similarly, among all random variables taking values in the space of continuous paths, Brownian motion is (up to multiplicative constant) the only one with reflection invariance, stationarity, and independence of increments. It has a strong claim to be *the* canonical continuous random path. These notes will survey objects that can claim with equal justification to be the canonical random planar tree, the canonical random non-self-crossing planar curve, the canonical random surface, and so forth. Like Brownian motion itself, many of these objects can be constructed from Gaussian noise in some way. We will see that they are also closely related to Poisson point processes and stable Lévy processes.

Among the various symmetries that make these objects special, many involve some sort of *conformal invariance*. Recall that the Riemann uniformization theorem implies the existence of a conformal map between any two sphere-homeomorphic surfaces. When the sphere is replaced by a multi-handled torus or a disk with holes, the space of conformal equivalence classes (a.k.a. the *moduli space*) remains finite dimensional. This remarkable fact is a peculiar feature of two dimensions and seems to be a large part of what makes the two dimensional theory interesting. In the 1980's and 1990's a branch of physics called *conformal field theory*, motivated by both string theory and two dimensional statistical mechanics, began to discover and explore some surprisingly far reaching consequences of conformal symmetry assumptions in physical models. Mathematicians have more recently expanded these ideas further, building in particular on the introduction of the so-called Schramm-Loewner evolution in 1999.

The focus of this text is on the mathematics, and in particular on a few of the most fundamental discrete and continuum mathematical objects in one and two dimensions. However we will also provide some cursory discussion of the motivating problems that link them to physics and to other fields.

Sections 1 through 6 introduce both discrete and continuum analogs of several universal random objects: processes, trees, distributions, curves, loop ensembles, surfaces, and growth trajectories. Sections 7 through 10 then explore some of the the intricate and often surprising relationships between these objects. To put this another way, the first half of this text introduces a certain cast of characters, and the second half explores the drama that takes place when these characters interact.

These notes are intended as a broad introductory overview of this field and as such they cover a good deal of material. With additional detail, each individual chapter could be (and in most cases already has been) expanded into an entire book or book-length monograph of its own. Some parts of these notes have been lifted (and adapted) from the more expository portions of some of the longer works by the current authors. We do not provide fully detailed proofs of every result cited in this text. However, we aim to provide enough rigor and detail to enable the reader to appreciate the overall narrative and to begin further research in this field.

Acknowledgments. J.M.'s work was partially supported by DMS-1204894 and S.S.'s work was partially supported by DMS-1209044, a fellowship from the Simons Foundation, and EPSRC grants EP/L018896/1 and EP/I03372X/1. S.S. presented this material as a graduate topics course at MIT in Fall 2015 and would like to thank the students there for feedback and support. We also thank the participants in the 2016 summer school at St. Flour where this text is used as lecture notes.

1 Random processes

1.1 Brownian motion

We trust that most readers are familiar on some level with Brownian motion and the Gaussian (a.k.a. normal) distribution, which form the core of many of the more sophisticated constructions that appear in this book.

Concepts to keep in mind for the current text include Itô's formula, the martingale representation theorem, local martingales, quadratic variation, and Girsanov's theorem.

More detailed accounts of this material can be found in basic probability texts like [Dur10], the book on Brownian motion by Mörters and Peres [MP10], and stochastic calculus texts such as [KS91a, RY99a].

1.2 Bessel processes

An introduction to Bessel process can be founded for example in [RY99a, Chapter 11]. We recall here that when δ is a real constant, the Bessel process of dimension δ, also written BES^δ, is a (non-negative) solution to the SDE

$$dX_t = dB_t + \frac{\delta - 1}{2X_t}dt, \quad X_0 \geq 0 \tag{1.1}$$

where B is a standard Brownian motion. When δ is a positive integer, X_t agrees in law with the distance from 0 of a standard Brownian motion in $\mathbf{R}^\delta$ (i.e., each of the components is an independent standard Brownian motion).

Suppose that $\alpha > 0$ is a positive constant. Then Iˆto's formul implies

$$d(X_t^\alpha) = \alpha X_t^{\alpha-1}\big(dB_t + \frac{\delta-1}{2}X_t^{-1}dt\big) + \frac{\alpha(\alpha-1)}{2}X_t^{\alpha-2}dt.$$

The dt terms cancel when $\alpha(\delta-1)/2 = -\alpha(\alpha-1)/2$, which holds when $(\delta-1) = -(\alpha-1)$ so that $\alpha = \delta - 2$. This implies that $X_t^{\delta-2}$ evolves as a local martingale (at least up until the first time at which $X_t = 0$). Using this, it is not hard to see that if X_0 is set to a positive fixed value, then X_t almost surely reaches 0 before ∞ when $\delta < 2$, and almost surely tends to ∞ without hitting zero when $\delta > 2$, and almost surely oscillates between values arbitrarily close to zero and arbitrarily large when $\delta = 2$.

Another application of Iˆto's formula shows that the sum of a δ_1 Bessel process and an independent δ_2 Bessel process is a $\delta_1 + \delta_2$ Bessel process, even when δ_1 and δ_2 are not positive integers.

In the case that $\delta < 2$, one interesting question is how to extend the Bessel process definition beyond times at which X_t reaches zero. One approach to constructing such an extension is to first define a process that jumps up by ϵ each time it hits 0, and take a limit as $\epsilon \to 0$. Another standard approach (adopted for example in [RY99a, Chapter 11]) is to first construct the *squared Bessel process* $Y_t = X_t^2$, which turns out to fit more neatly into the framework of some general existence and uniqueness theorems in SDE theory, and then take $X_t = \sqrt{Y_t}$.

When $\delta > 1$, (1.1) holds in the sense that X is a.s. *instantaneously reflecting* at 0 (i.e., the set of times for which $X_t = 0$ has Lebesgue measure zero) and a.s. satisfies

$$X_t = X_0 + B_t + \int_0^t \frac{\delta-1}{2X_s}ds, \quad X_0 \geq 0. \tag{1.2}$$

In particular, assuming $\delta > 1$, the integral in (1.2) is finite a.s. so that X_t is a semi-martingale. The solution is a strong solution in the sense of [RY99a], which means that X is adapted to the filtration generated by the Brownian motion B. The law of X is determined by the fact that it is a solution to (1.1) away from times where $X_t = 0$, instantaneously reflecting where $X_t = 0$, and adapted to the filtration generated by B.

Regardless of δ, standard SDE results imply that (1.2) has a unique solution up until the first time t that $X_t = 0$. When $\delta < 1$, however, (1.2) cannot hold beyond times at which $X_t = 0$ without a so-called principal value correction, because the integral in (1.2) is almost surely infinite beyond such times (see [She09a, Section 3.1] for additional discussion of this point). Bessel processes can be defined for all time whenever $\delta > 0$ but they are not semi-martingales when $\delta \in (0, 1)$.

1.3 Brownian excursions, meanders, and bridges

One may define a Brownian excursion indexed by $[0, 1]$ by conditioning a Brownian motion, started at ϵ, to end in $[0, \epsilon]$, and then taking the $\epsilon \to 0$ limit. Brownian motion conditioned to stay in a cone (starting from the apex) is explained in [Shi85] along with the relationship to Bessel processes.

1.4 Stable Lévy processes

We assume the reader has (or is able to quickly acquire) a Wikipedia-level understanding of stable distributions and the corresponding stable Lévy processes. See also the textbooks on Lévy processes by Sato, by Bertoin and by Barndoff-Nielson, Mikosch, and Resnick [Sat99, Ber96, BNMR01].

We are mainly interested in *strictly stable* process, whose laws remain unchanged when space is rescaled by C and time is rescaled by C^α, where $\alpha \in (0, 2]$ is a fixed parameter. These processes have jumps whose magnitudes can be understood as a Poisson point process on the product of the real numbers and $u^{-\alpha-1}du$ (where u denotes jump magnitude). An additional parameter $\beta \in [-1, 1]$ is chosen so that a $(\beta + 1)/2 \in [0, 1]$ fraction of the jumps are positive, with the others being negative.

The characteristic function of a standard stable random variable with parmeters (α, β) is given by

$$\exp\bigl[-|ct|^\alpha(1 - i\beta \mathrm{sgn}(t)\Phi)\bigr],$$

where $\Phi = \tan \frac{\pi\alpha}{2} 4$ if $\alpha \neq 1$ and $\frac{2}{\pi} \log |t|$ if $\alpha = 1$.

1.5 Continuous state branching processes

We next recall some basic facts about *continuous state branching processes*, which were introduced by Jiřina and Lamperti several decades ago [Jiř58, Lam67a, Lam67c] (see also the more recent overview in [LG99] as well as [Kyp06, Chapter 10]). A Markov process $(Y_t, t \geq 0)$ with values in $\mathbf{R}_+$, whose sample paths are càdlàg (right continuous with left limits) is said to be a continuous state branching process (CSBP for short) if the transition kernels $P_t(x, dy)$ of Y satisfy the additivity property:

$$P_t(x + x', \cdot) = P_t(x, \cdot) * P_t(x', \cdot). \tag{1.3}$$

Remark 1.1. Note that (1.3) implies that the law of a CSBP at a *fixed* time is infinitely divisible. In particular, this implies that for each fixed t there exists a subordinator (i.e., a non-decreasing process with stationary, independent increments) A^t with $A^t_0 = 0$ such that $A^t_t \stackrel{d}{=} Y_t$. (We emphasize though that Y does not *evolve* as a subordinator in t.) We will make use of this fact several times.

The Lamperti representation theorem states that there is a simple time-change procedure that gives a one-to-one correspondence between CSBPs and non-negative Lévy processes without negative jumps (stopped when they reach zero), where each is a time-change of the other. The statement of the theorem we present below is lifted from a recent expository treatment of this result [CLUB09].

Consider the space $\mathcal{D}$ of càdlàg functions $f\colon [0, \infty] \to [0, \infty]$ such that $\lim_{t\to\infty} f(t)$ exists in $[0, \infty]$ and $f(t) = 0$ (resp. $f(t) = \infty$) implies $f(t + s) = 0$ (resp. $f(t + s) = \infty$) for all $s \geq 0$. For any $f \in \mathcal{D}$, let $\theta_t := \int_0^t f(s)ds \in [0, \infty]$, and let κ denote the right-continuous inverse of θ, so $\kappa_t := \inf\{u \geq 0 : \theta_u > t\} \in [0, \infty]$, using the convention $\inf \emptyset = \infty$. The *Lamperti transformation* is given by $L(f) = f \circ \kappa$. The following is the Lamperti representation theorem, which applies to $[0, \infty]$-valued processes indexed by $[0, \infty]$.

Theorem 1.2. *The Lamperti transformation is a bijection between CSBPs and Lévy processes with no negative jumps stopped when reaching zero. In other words, for any CSBP Y, $L(Y)$ is a Lévy process with no negative jumps stopped whenever reaching zero; and for any Lévy process X with no negative jumps stopped when reaching zero, $L^{-1}(X)$ is a CSBP.*

Informally, the CSBP is just like the Lévy process it corresponds to except that its speed (the rate at which jumps appear) is given by a constant times its current value (instead of being independent of its current value). The following is now immediate from Theorem 1.2 and the definitions above:

Proposition 1.3. *Suppose that X_t is a Lévy process with non-negative jumps that is strictly α-*stable *in the sense that for each $C > 0$, the rescaled process $X_{C^\alpha t}$ agrees in law with CX_t (up to a change of starting point). Let $Y = L^{-1}(X)$. Then Y is a CSBP with the property that $Y_{C^{\alpha-1}t}$ agrees in law with CY_t (up to a change of starting point). The converse is also true. Namely, if Y is a CSBP with the property that $Y_{C^{\alpha-1}t}$ agrees in law with CY_t (up to a change of starting point) then Y is the CSBP obtained as a time-change of the α-stable Lévy process with non-negative jumps.*

Proposition 1.3 will be useful on occasions when we want to prove that a given process Y is the CSBP obtained as a time change of the α-stable Lévy process with non-negative jumps. (We refer to this CSBP as the *α-stable CSBP* for short.[1]) It shows that it suffices in those settings to prove that Y is a CSBP and that it has the scaling symmetry mentioned in the proposition statement. To avoid dealing with uncountably many points, we will actually often use the following slight strengthening of Proposition 1.3:

Proposition 1.4. *Suppose that Y is a Markovian process indexed by the dyadic rationals that satisfies the CSBP property* (1.3) *and that $Y_{C^{\alpha-1}t}$ agrees in law with CY_t (up to a change of starting point) when $C^{\alpha-1}$ is a power of 2. Assume that Y is not trivially equal to 0 for all positive time, or equal to ∞ for all positive time. Then Y is the restriction (to the dyadic rationals) of an α-stable CSBP.*

Proof. By the CSBP property 1.3, the law of Y_1, assuming $Y_0 = a > 0$, is infinitely divisible and equivalent to the law of the value A_a where A is a subordinator and $A_0 = 0$ (recall Remark 1.1). Fix $k \in \mathbf{N}$ and pick $C > 0$ such that $C^{1-\alpha} = 2^{-k}$. Similarly, by scaling, we have that $Y_{C^{1-\alpha}} \overset{d}{=} C^{-1}A_{Ca}$. By the law of large numbers, this law is concentrated on $a\mathbf{E}[A_1]$ when k is large; we observe that $\mathbf{E}[A_1] = 1$ since otherwise (by taking the $k \to \infty$ limit) one could show that Y is equal to 0 (if $\mathbf{E}[A_1] < 1$) or ∞ (if $\mathbf{E}[A_1] > 1$) for all positive time.

From this we deduce that Y is a martingale, and the standard upcrossing lemma allows us to conclude that almost surely Y has only finitely many upcrossings across the interval $(x, x+\epsilon)$ for any x and ϵ, and that Y a.s. is bounded above. This in turn guarantees, for all $t \geq 0$, the existence of left and right limits of Y_{t+s} as $s \to 0$. It implies that Y is a.s. the restriction to the dyadic rationals of a càdlàg process; and there is a unique way to extend Y to a càdlàg process defined for all $t \geq 0$. Since left limits exist almost surely at any fixed time, it is straightforward to verify that the hypotheses of Proposition 1.3 apply to Y. □

CSBPs are often introduced in terms of their Laplace transform [LG99], [Kyp06, Chapter 10] and Proposition 1.3 is also immediate from this perspective. We will give a brief review of this here, since this perspective will also be useful in this article. In the case of an α-stable CSBP Y_t, this Laplace transform is explicitly given by

$$\mathbf{E}[\exp(-\lambda Y_t) \,|\, Y_s] = \exp(-Y_s u_{t-s}(\lambda)) \quad \text{for all} \quad t > s \geq 0 \tag{1.4}$$

where

$$u_t(\lambda) = \left(\lambda^{1-\alpha} + (\alpha-1)t\right)^{1/(1-\alpha)}. \tag{1.5}$$

More generally, CSBPs are characterized by the property that they are Markov processes on $\mathbf{R}_+$ such that their Laplace transform has the form given in (1.4) where $u_t(\lambda)$, $t \geq 0$,

[1]This process is also referred to as a ψ-CSBP with "branching mechanism" $\psi(u) = u^\alpha$ in other work in the literature, for example [DLG05].

is the non-negative solution to the differential equation

$$\frac{\partial u_t}{\partial t}(\lambda) = -\psi(u_t(\lambda)) \quad \text{for} \quad u_0(\lambda) = \lambda. \tag{1.6}$$

The function ψ is the so-called *branching mechanism* for the CSBP and corresponds to the Laplace exponent of the Lévy process associated with the CSBP via the Lamperti transform (Theorem 1.2). In this language, an α-stable CSBP is a called a "CSBP with branching mechanism $\psi(u) = u^\alpha$."

One of the uses of (1.4) is that it provides an easy derivation of the law of the extinction time of a CSBP, which we record in the following lemma.

Lemma 1.5. *Suppose that Y is an α-stable CSBP and let $\zeta = \inf\{t \geq 0 : Y_t = 0\}$ be the extinction time of Y. Then we have that*

$$\mathbf{P}[\zeta > t] = 1 - \exp\left(c_\alpha t^{1/(1-\alpha)} Y_0\right) \quad \textit{where} \quad c_\alpha = (\alpha - 1)^{1/(1-\alpha)}. \tag{1.7}$$

Proof. Note that $\{\zeta > t\} = \{Y_t > 0\}$. Consequently,

$$\mathbf{P}[\zeta > t] = \mathbf{P}[Y_t > 0] = 1 - \lim_{\lambda \to \infty} \mathbf{E}[e^{-\lambda Y_t}] = 1 - \exp(c_\alpha t^{1/(1-\alpha)} Y_0),$$

which proves (1.7). □

1.6 Ranges of stable subordinators

The range of a stable subordinator is a random closed subset of $\mathbf{R}_+$. which can also be understood as the zero set of a Bessel process. If we condition the endpoints of the Bessel process to be zero at 1, we can also define a random closed subset of $[0, 1]$. These random sets can be characterized by renewal and scale invariance properties, which are similar to the properties we will later use to characterize conformal loop ensembles (the complement of the union of the interiors of these loops will turn out to be a random closed subset of $\mathbf{R}^2$).

2 Random trees

2.1 Galton-Watson trees

Galton-Watson trees and their scaling limits are described by Duquesne and Le Gall in [DLG05]. See also [LGLJ98, DLG06, DLG09]. One of the interesting features of Galton-Watson trees is the phase transition: when the expected number of children is less than one, the tree is easily seen to be finite almost surely. (The expected number of children at level k decays exponentially in k.) When the expected number of children is greater than one, the tree has a positive probability of being infinite.

When the expected number of children is equal to 1, one may observe offspring sets of vertices one at a time, exploring tree boundary in a clockwise way, so that the number of live vertices is a martingale. This martingale is closely related to the contour function of the tree (but not exactly the same; see Lévy tree story below).

2.2 Aldous's continuum random tree

The continuum random tree was introduced in a series of papers by Aldous in 1991 [Ald91a, Ald91b, Ald93]. It can be understood as a scaling limit of Galton-Watson trees.

2.3 Lévy trees and stable looptrees

There are some very simple analogs of the CRT in which stable Lévy excursions play the role of the Brownian excursion [DLG05]. These can also be understood as scaling limits of Galton-Watson trees, when the number of children has a power law tail (finite mean but infinite variance).

There is a closely related construction in which each of the countably many big branch points is replaced with a loop; the resulting "tree of loops" called a looptree. See the work by Curien and Korchemski on *stable looptrees* [CK13], as well as the exposition in [DMS14].

2.4 Brownian snakes

A Brownian snake is essentially a Brownian motion indexed by a CRT. It will play a role later in the construction of a certain canonical random surface called the Brownian map, but it was actually studied independently before its relationship to random surfaces was discovered [DLG05].

3 Random generalized functions

3.1 Tempered distributions and Fourier transforms

The *Schwartz space* on $\mathbb{R}^d$ is the space of C^∞ functions ϕ such that for any multi-indices α and β the seminorm $\sup D^\alpha \phi(x) x^\beta$ is bounded. These seminorms induce a topology on the Schwartz space; continuous linear functionals on the Schwartz space are called *tempered distributions*. The space of tempered distributions is the smallest space which includes the bounded continuous functions and is closed under both differentiation and the Fourier transform.

In the exposition on Gaussian free fields, we will often find it convenient to limit attention to compactly supported test functions (instead of test functions in the Schwartz space) as this will allow us to more easily isolate the effects of boundary conditions.

3.2 Gaussian free fields

Gaussian Hilbert spaces are introduced in [Jan97]. Surveys of the Gaussian free field can be found in [She07, Ber].

3.2.1 Dirichlet inner product

Fix a simply connected planar domain $D \subset \mathbf{C}$ (with $D \neq \mathbf{C}$). Let $H_s(D)$ be the space of smooth, compactly supported functions on D, and let $H(D)$ (sometimes denoted by $\mathbb{H}_0^1(D)$ or $W^{1,2}(D)$) be its Hilbert space closure under the Dirichlet inner product

$$(f_1, f_2)_\nabla := (2\pi)^{-1} \int_D \nabla f_1(z) \cdot \nabla f_2(z) dz.$$

Let ψ be a conformal map from another domain $\widetilde{D}$ to D. Then an elementary change of variables calculation shows that

$$\int_{\widetilde{D}} \nabla(f_1 \circ \psi) \cdot \nabla(f_2 \circ \psi)\, dx = \int_D (\nabla f_1 \cdot \nabla f_2)\, dx.$$

In other words, the Dirichlet inner product is invariant under conformal transformations.

We write $(f_1, f_2) = \int_D f_1(x) f_2(x) dx$ for the L^2 inner product on D. We write $\|f\| := (f,f)^{1/2}$ and $\|f\|_\nabla := (f,f)_\nabla^{1/2}$. If $f_1, f_2 \in H_s(D)$, then integration by parts gives

$$(f_1, f_2)_\nabla = \frac{1}{2\pi}(f_1, -\Delta f_2). \tag{3.1}$$

3.2.2 Distributions and the Laplacian

It is conventional to use $H_s(D)$ as a space of test functions. This space is a topological vector space in which the topology is defined so that $\phi_k \to 0$ in $H_s(D)$ if and only if there is a compact set on which all of the ϕ_k are supported and the mth derivative of ϕ_k converges uniformly to zero for each integer $m \geq 1$.

A **distribution** on D is a continuous linear functional on $H_s(D)$. Since $H_s(D) \subset L^2(D)$, we may view every $h \in L^2(D)$ as a distribution $\rho \mapsto (h, \rho)$. A **modulo-additive-constant** distribution on D is a continuous linear functional on the subspace of $H_s(D)$ consisting of ρ for which $\int_D \rho(z)dz = 0$. We will frequently abuse notation and use h — or more precisely the map denoted by $\rho \to (h, \rho)$ — to represent a general distribution (which is a functional of ρ), even though h may not correspond to an element of $L^2(D)$. (Later, we will further abuse notation and use ρ to represent a non-smooth function or a measure; in the latter case (h, ρ), when defined, will represent the integral of h against that measure.)

We define partial derivatives and integrals of distributions in the usual way (via integration by parts), i.e., for $\rho \in H_s(D)$,

$$\big(\frac{\partial}{\partial x} h, \rho\big) := -\big(h, \frac{\partial}{\partial x}\rho\big).$$

In particular, if h is a distribution then Δh is a distribution defined by $(\Delta h, \rho) := (h, \Delta \rho)$. When h is a distribution and $\rho \in H_s(D)$, we also write

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

When $x \in D$ is fixed, we let $\widetilde{G}_x(y)$ be the harmonic extension to $y \in D$ of the function of y on ∂D given by $-\log|y - x|$. Then **Green's function in the domain** D is defined by

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

When $x \in D$ is fixed, Green's function may be viewed as a distributional solution of $\Delta G(x, \cdot) = -2\pi\delta_x(\cdot)$ with zero boundary conditions [She07]. It is non-negative for all $x, y \in D$ and $G(x, y) = G(y, x)$.

For any $\rho \in H_s(D)$, we write

$$-\Delta^{-1}\rho := \frac{1}{2\pi}\int_D G(\cdot, y)\,\rho(y)\,dy.$$

This is a C^∞ function in D whose Laplacian is $-\rho$. Indeed, a similar definition can be made if ρ is any signed measure (with finite positive and finite negative mass) rather than a smooth function. Recalling (3.1), if $f_1 = -2\pi\Delta^{-1}\rho_1$ then $(h, f_1)_\nabla = (h, \rho_1)$, and similarly if $f_2 = -2\pi\Delta^{-1}\rho_2$. Now $(f_1, f_2)_\nabla = (\rho_1, -2\pi\Delta^{-1}\rho_2)$ describes a covariance that can (by the definition of $-\Delta^{-1}\rho_2$ above) be rewritten as

$$\mathrm{Cov}\big((h, \rho_1), (h, \rho_2)\big) = \int_{D\times D} \rho_1(x)\,G(x, y)\,\rho_2(y)\,dx\,dy. \tag{3.2}$$

If $\rho \in H_s(D)$, may define the map $(h, \cdot)$ by $(h, \rho) := (h, -2\pi\Delta^{-1}\rho)_\nabla$, and this definition describes a distribution [She07]. (It is not hard to see that $-2\pi\Delta^{-1}\rho \in H(D)$, since its Dirichlet energy is given explicitly by (3.2).)

3.2.3 Zero boundary GFF

An instance of the GFF with zero boundary conditions on D is a random sum of the form $h = \sum_{j=1}^{\infty} \alpha_j f_j$ where the α_j are i.i.d. one-dimensional standard (unit variance, zero mean) real Gaussians and the f_j are an orthonormal basis for $H(D)$. This sum almost surely does not converge within $H(D)$ (since $\sum_{j=1}^{\infty} |\alpha_j|^2$ is a.s. infinite). However, it does converge almost surely within the space of distributions — that is, the limit $(\sum_{j=1}^{\infty} \alpha_j f_j, \rho)$ almost surely exists for all $\rho \in H_s(D)$, and the limiting value as a function of ρ is almost surely a continuous functional on $H_s(D)$ [She07]. We may view h as a sample from the measure space $(\Omega, \mathcal{F})$ where $\Omega = \Omega_D$ is the set of distributions on D and $\mathcal{F}$ is the smallest σ-algebra that makes (h, ρ) measurable for each $\rho \in H_s(D)$, and we sometimes denote by dh the probability measure which is the law of h. If f_j are chosen in $H_s(D)$, then the values α_j are clearly $\mathcal{F}$-measurable. In fact, for any $f \in H(D)$ with $f = \sum_j \beta_j f_j$ the sum $(h, f)_\nabla := \sum_j \alpha_j \beta_j$ is a.s. well defined and is a Gaussian random variable with mean zero and variance $(f, f)_\nabla$.

3.2.4 Green's functions on C and $\mathbb{H}$: free boundary GFF

The GFF with free boundary conditions is defined the same way as the GFF with zero boundary conditions except that we replace $H_s(D)$ with the space of all smooth functions with gradients in $L^2(D)$ (i.e., we remove the requirement that the functions be compactly supported). However, to make the correspondingly defined $H(D)$ a Hilbert space, we have to consider functions only modulo additive constants (since all constant functions have norm zero). On the whole plane $\mathbf{C}$, we may define the Dirichlet inner product on the Hilbert space closure $H(\mathbf{C})$ of the space of such functions defined modulo additive constants.

Generally, given a compactly supported ρ (or more generally, a signed measure), we can write

$$- \Delta^{-1}\rho(\cdot) := \frac{1}{2\pi} \int_{\mathbf{C}} G(\cdot, y)\rho(y)dy, \tag{3.3}$$

with $G(x, y) = -\log|x - y|$.

As before, for compactly supported f and g, we have $(f, g)_\nabla = \frac{1}{2\pi}(f, -\Delta g)$ by integration by parts, and moreover $(f, -\Delta^{-1}\rho)_\nabla = \frac{1}{2\pi}(\rho, f)$. The same holds for bounded and not necessarily compactly supported smooth functions f and g if the gradient of $-\Delta^{-1}\rho$ tends to zero at infinity, which in turn happens if and only if $\int_{\mathbf{C}} \rho(z)dz = 0$.

If $\int_{\mathbf{C}} \rho(z)dz \neq 0$ then the Dirichlet energy of $-\Delta^{-1}\rho$ will be infinite and moreover (h, ρ) will not be independent of the additive constant chosen for h. (If we view $\mathbf{C}$ as a Riemann sphere, then $\int_{\mathbf{C}} \rho(z)dz \neq 0$ can also be interpreted as the statement that the Laplacian of $-\Delta^{-1}\rho$ has a point mass at ∞.) When h is the free boundary GFF on $\mathbf{C}$, we will thus define the random variables (h, ρ) only if the integral of ρ over $\mathbf{C}$ is zero. If ρ_1 and ρ_2 each have total integral zero, we may write

$$\operatorname{Cov}((h,\rho_1),(h,\rho_2)) = \int_{\mathbf{C}\times\mathbf{C}} \rho_1(x)G(x,y)\rho_2(y)dxdy. \tag{3.4}$$

Using $z \to \bar{z}$ to denote complex conjugation, we define, for smooth functions $h \in H(\mathbf{C})$, the pair of projections

$$\begin{aligned} h^{\mathrm{O}}(z) &:= \frac{1}{\sqrt{2}}(h(z) - h(\bar{z})), \\ h^{\mathrm{E}}(z) &:= \frac{1}{\sqrt{2}}(h(z) + h(\bar{z})). \end{aligned}$$

If h is an instance of the free boundary GFF on $\mathbf{C}$, we may still define h^{O} and h^{E} as projections of h onto complementary orthogonal subspaces. Their restrictions to $\mathbb{H}$ are instances of the zero boundary GFF and free boundary GFF, respectively on $\mathbb{H}$. For ρ supported on $\mathbb{H}$ we write (for $z \in \mathbb{C}$) $\rho^*(z) := \rho(\overline{z})$. Then we have by definition

$$\begin{aligned} (h^{\mathrm{O}},\rho) &= \frac{1}{\sqrt{2}}(h, \rho - \rho^*) \\ (h^{\mathrm{E}},\rho) &= \frac{1}{\sqrt{2}}(h, \rho + \rho^*). \end{aligned}$$

Note that (h^{E},ρ) is only defined if the total integral of ρ is zero, while (h^{O},ρ) is defined without that restriction (since in any case the total integral of $\rho - \rho^*$ will be zero).

For ρ_1 and ρ_2 supported on $\mathbb{H}$ we now compute the following (first integral taken over $\mathbf{C}\times\mathbf{C}$, second over $\mathbb{H}\times\mathbb{H}$):

$$\begin{aligned} \operatorname{Cov}\big((h^{\mathrm{O}},\rho_1),(h^{\mathrm{O}},\rho_2)\big) &= \frac{1}{2}\int (\rho_1(x) - \rho_1^*(x)) \log|x-y| (\rho_2(y) - \rho_2^*(y)) dxdy \\ &= \int \rho_1(x) G^{\mathbb{H}_0}(x,y)\rho_2(y)dxdy, \end{aligned} \tag{3.5}$$

where $G^{\mathbb{H}_0}(x,y) := \log|x-\bar{y}| - \log|x-y|$. Similarly (first integral over $\mathbf{C}\times\mathbf{C}$, second over $\mathbb{H}\times\mathbb{H}$),

$$\begin{aligned} \operatorname{Cov}\big((h^{\mathrm{E}},\rho_1),(h^{\mathrm{E}},\rho_2)\big) &= \frac{1}{2}\int (\rho_1(x) + \rho_1^*(x)) \log|x-y| (\rho_2(y) + \rho_2^*(y)) dxdy \\ &= \int \rho_1(x) G^{\mathbb{H}_F}(x,y)\rho_2(y)dxdy, \end{aligned} \tag{3.6}$$

where $G^{\mathbb{H}_F}(x,y) := -\log|x-\bar{y}| - \log|x-y|$.

3.2.5 GFF as a continuous functional

Note that we could have used (3.5) and (3.6) to give an alternate and more direct definition of the zero and free boundary Gaussian free fields on $\mathbb{H}$. Here (3.5) and (3.6) define inner products on the space of functions ρ on $\mathbb{H}$. They are well defined when ρ_1 and ρ_2 are smooth and compactly supported functions on $\mathbb{H}$ (each with total integral zero in the case of (3.6)). By taking the Hilbert space closure of functions of this type, we get a larger space of ρ, which correspond to Laplacians of elements of $H(\mathbb{H})$, and which cannot all be interpreted as functions on $\mathbb{H}$. For example, the ρ for which (h, ρ) is $h_\varepsilon(z)$, the mean value of h on $\partial B_\varepsilon(z)$, is not a function, though it can be interpreted as a measure — a uniform measure on $\partial B_\varepsilon(z)$ — and the inner products (3.5) and (3.6) still make sense when $\rho_1(z)dz$ and $\rho_2(z)dz$ are replaced with more general measures, as do the definitions of $-\Delta^{-1}\rho_1$ and $-\Delta^{-1}\rho_2$.

The (h, ρ) are centered jointly Gaussian random variables, defined for each ρ in this Hilbert space, with covariances given by the inner products (3.5) and (3.6) (which can be defined on the entire Hilbert space). For each particular ρ in this Hilbert space, (h, ρ) is a.s. well defined and finite; however, $\rho \to (h, \rho)$ is almost surely not a continuous linear functional defined on the entire Hilbert space, since a.s. $h \notin H(\mathbb{H})$.

In addition to the description of h as a distribution above, there are various ways to construct a space of ρ values — a subset of the complete Hilbert space — endowed with a topology that makes $\rho \to (h, \rho)$ almost surely continuous. For example, the map $h \to h_\varepsilon(z)$ is an a.s. Hölder continuous function of ε and z [DS11a]. Also, the zero boundary GFF can be defined as a random element of $(-\Delta)^{-\varepsilon}L^2(D)$ for any $\varepsilon > 0$, and is hence a continuous linear function on $(-\Delta)^{\varepsilon}L^2(D)$, if D is bounded. (See [She07] for definitions and further discussion of fractional powers of the Laplacian in this context.) Also, as mentioned earlier, both the free and zero boundary GFFs can be understood as random distributions [She07].

The issues that come up when defining $\rho \to (h, \rho)$ as a continuous function on some topological space of ρ values are the same ones that come up when rigorously constructing a Brownian motion B_t: one can give the joint law of B_t for any finite set of t values explicitly by specifying covariances, and this determines the law for any fixed countable set of t values, but one needs to overcome some (mild) technicalities in order to say "B_t is almost surely a continuous function." Indeed, if one uses the smallest σ-algebra in which B_t is measurable for each fixed t, then the event that B_t is continuous is not even in the σ-algebra.

On the other hand, if we are given a construction that produces a random continuous function with the right finite dimensional marginals, then it must be a Brownian motion. A standard fact (proved using characteristic functions and Fourier transforms) states that a random variable on a finite dimensional space is a centered Gaussian with a given covariance structure if and only if all of its one dimensional projections are centered Gaussians with the appropriate variance. Thus, to establish that B_t is a Brownian motion, it is enough to show that each finite linear combination of B_t values is a (one-dimensional) centered Gaussian with the right variance. The following proposition formalizes the analogous notion in the GFF context. It is a standard and straightforward result about Gaussian processes (see [She07] for a proof in the zero boundary case; the free boundary case is identical):

Proposition 3.1. *The zero boundary GFF on $\mathbb{H}$ is the only random distribution h on $\mathbb{H}$ with the property that for each $\rho \in H_s(\mathbb{H})$ the random variable (h, ρ) is a mean-zero Gaussian with variance given by* (3.5) *(with $\rho_1 = \rho_2 = \rho$). Similarly, the free boundary GFF is the only random modulo-additive-constant distribution on $\mathbb{H}$ with the property that for each $\rho \in H_s(\mathbb{H})$ with $\int_{\mathbb{H}} \rho(z)dz = 0$ the random variable (h, ρ) is a mean-zero Gaussian with variance given by* (3.6)*.*

In our proofs of Theorem 7.1 and Theorem 7.2, we will first construct a random distribution in the manner prescribed by the theorem statement and then check the laws of the one dimensional projections (which determine the laws of the finite and countably infinite dimensional projections) to conclude by Proposition 3.1 that it must be the GFF.

We remark that knowing h as a distribution determines the values of α_j in a basis expansion $h = \sum_j \alpha_j f_j$, as long as the $-\Delta f_j$ are sufficiently smooth. This in turn determines the value of $h_\varepsilon(z)$ almost surely for a countable dense set of ε and z values, which determines the values for all ε and z by the almost sure continuity of $h_\varepsilon(z)$ [DS11a]. This is convenient because it means that h, as a distribution, a.s. determines $(z, \varepsilon) \to h_\varepsilon(z)$ as a function, which in turn determines μ_h and ν_h. (We could alternatively have defined $h_\varepsilon(z)$ — and hence μ_h and ν_h — using weighted averages of h defined by integrating against smooth bump functions on $B_\varepsilon(z)$ instead of averages on $\partial B_\varepsilon(z)$. Though we won't do this here, one can easily construct measures this way that are almost surely equivalent to μ_h and ν_h.)

3.2.6 Dirichlet inner product

Let D be a domain in $\mathbf{C}$ with smooth, **harmonically non-trivial** boundary. The latter means that the harmonic measure of ∂D is positive as seen from any point in D. Let $C_0^\infty(D)$ denote the set of C^∞ functions compactly supported in D. The **Dirichlet inner product** is defined by

$$(f, g)_\nabla = \frac{1}{2\pi} \int_D \nabla f(x) \cdot \nabla g(x) dx \quad \text{for} \quad f, g \in C_0^\infty(D). \tag{3.7}$$

More generally, (3.7) makes sense for $f, g \in C^\infty(D)$ with L^2 gradients.

3.2.7 Distributions

We view $C_0^\infty(D)$ as a space of test functions and equip it with the topology where a sequence (ϕ_k) in $C_0^\infty(D)$ satisfies $\phi_k \to 0$ if and only if there exists a compact set $K \subseteq D$ such that the support of ϕ_k is contained in K for every $k \in \mathbf{N}$ and ϕ_k as well as all of its derivatives converge uniformly to zero as $k \to \infty$. A **distribution** on D is a continuous linear functional on $C_0^\infty(D)$ with respect to the aforementioned topology. A **modulo additive constant distribution** on D is a continuous linear functional on the subspace of functions f of $C_0^\infty(D)$ with $\int_D f(x)dx = 0$ with the same topology.

3.2.8 GFF with Dirichlet and mixed boundary conditions

We let $H_0(D)$ be the Hilbert-space closure of $C_0^\infty(D)$ with respect to the Dirichlet inner product (3.7). The GFF h on D with zero Dirichlet boundary conditions can be expressed as a random linear combination of an $(\cdot,\cdot)_\nabla$-orthonormal basis (f_n) of $H_0(D)$:

$$h = \sum_n \alpha_n f_n, \quad (\alpha_n) \quad \text{i.i.d.} \quad N(0,1). \tag{3.8}$$

Although this expansion of h does not converge in $H_0(D)$, it does converge almost surely in the space of distributions or (when D is bounded) in the fractional Sobolev space $H^{-\epsilon}(D)$ for each $\epsilon > 0$ (see [She07, Proposition 2.7] and the discussion thereafter). If $f, g \in C_0^\infty(D)$ then an integration by parts gives $(f,g)_\nabla = -(2\pi)^{-1}(f, \Delta g)$. Using this, we define

$$(h,f)_\nabla = -\frac{1}{2\pi}(h, \Delta f) \quad \text{for} \quad f \in C_0^\infty(D).$$

Observe that $(h,f)_\nabla$ is a Gaussian random variable with mean zero and variance $(f,f)_\nabla$. Hence h induces a map $C_0^\infty(D) \to \mathcal{G}$, $\mathcal{G}$ a Gaussian Hilbert space, that preserves the Dirichlet inner product. This map extends uniquely to $H_0(D)$ and allows us to make sense of $(h,f)_\nabla$ for all $f \in H_0(D)$ and, moreover,

$$\mathrm{Cov}((h,f)_\nabla, (h,g)_\nabla) = (f,g)_\nabla \quad \text{for all} \quad f,g \in H_0(D).$$

For fixed $x \in D$ we let $\widetilde{G}_x(y)$ be the harmonic extension of $y \mapsto -\log|x-y|$ from ∂D to D. The **Dirichlet Green's function** on D is defined by

$$G^{\mathrm{D}}(x,y) = -\log|y-x| - \widetilde{G}_x(y).$$

When $x \in D$ is fixed, $G^{\mathrm{D}}(x,\cdot)$ may be viewed as the distributional solution to $\Delta G^{\mathrm{D}}(x,\cdot) = -2\pi\delta_x(\cdot)$ with zero boundary conditions. When $D = \mathbf{D}$, we have that

$$G^{\mathrm{D}}(x,y) = \log\left|\frac{1-x\overline{y}}{y-x}\right|. \tag{3.9}$$

Repeated applications of integration by parts also imply that

$$\begin{aligned}\mathrm{Cov}((h,f),(h,g)) &= (2\pi)^2\mathrm{Cov}((h,\Delta^{-1}f)_\nabla,(h,\Delta^{-1}g)_\nabla)\\ &= \iint_{D\times D} f(x)G^{\mathrm{D}}(x,y)g(y)dxdy\end{aligned}$$

where G^{D} is the Dirichlet Green's function on D. If h is a zero-boundary GFF on D and $F\colon D\to\mathbf{R}$ is harmonic, then $h+F$ is the GFF with Dirichlet boundary conditions given by those of F.

More generally, suppose that $D\subseteq\mathbf{C}$ is a domain and $\partial D=\partial^{\mathrm{D}}\cup\partial^{\mathrm{F}}$ where $\partial^{\mathrm{D}}\cap\partial^{\mathrm{F}}=\emptyset$. We also assume that the harmonic measure of ∂^{D} is positive as seen from any point $z\in D$. The GFF on D with Dirichlet (resp. free) boundary conditions on ∂^{D} (resp. ∂^{F}) is constructed using a series expansion as in (3.8) except the space $H_0(D)$ is replaced with the Hilbert space closure with respect to $(\cdot,\cdot)_\nabla$ of the subspace of functions in $C^\infty(D)$ which have an L^2 gradient and vanish on ∂^{D}. The aforementioned facts for the GFF with only Dirichlet boundary conditions also hold verbatim for the GFF with mixed Dirichlet/free boundary conditions. In the case that D is a smooth Jordan domain and ∂^{D}, ∂^{F} are each non-degenerate intervals of ∂D, the Green's function G is taken to solve $\Delta G(x,\cdot)=-2\pi\delta_x(\cdot)$ with $n\cdot\nabla G(x,\cdot)=0$ on ∂^{F} and $G(x,\cdot)=0$ on ∂^{D} for $x\in D$. See also the discussion in [DS11a, Section 6.2] for the GFF with mixed boundary conditions.

3.2.9 GFF with free boundary conditions

The GFF with free boundary conditions on $D\subseteq\mathbf{C}$ is constructed using a series expansion as in (3.8) except we replace $H_0(D)$ with the Hilbert space closure $H(D)$ of the subspace of functions $f\in C^\infty(D)$ with $\|f\|_\nabla^2:=(f,f)_\nabla<\infty$ with respect to the Dirichlet inner product (3.7). Since the constant functions are elements of $C^\infty(D)$ but have $\|\cdot\|_\nabla$-norm zero, in order to make sense of this object, we will work in the space of distributions modulo additive constant. As in the case of the ordinary GFF, it is not difficult to see that the series converges almost surely in this space. As in Section 3.2.8, we can view $(h,f)_\nabla$ for $f\in H(D)$ as a Gaussian Hilbert space where

$$\mathrm{Cov}((h,f)_\nabla,(h,g)_\nabla)=(f,g)_\nabla\quad\text{for all}\quad f,g\in H(D).$$

Note that we do not need to restrict to mean zero test functions here due to the presence of gradients.

The **Neumann Green's function** on D is defined by

$$G^{\mathrm{F}}(x,y)=-\log|y-x|-\widehat{G}_x(y)$$

where for $x\in D$ fixed, $y\mapsto\widehat{G}_x(y)$ is the function on D such that the normal derivative of $G^{\mathrm{F}}(x,y)$ along ∂D is equal to 1. (The reason for the superscript "F" is that, as explained below, G^{F} gives the covariance function for the GFF with free boundary conditions.) When $x \in D$ is fixed, G^{F} may be viewed as the distributional solution to $\Delta G^{\mathrm{F}}(x,\cdot) = -2\pi\delta_x(\cdot)$ where the normal derivative of $G^{\mathrm{F}}(x,\cdot)$ is equal to 1 at each $y \in \partial D$. When $D = \mathbf{D}$, we have that

$$G^{\mathrm{F}}(x,y) = -\log|(x-y)(1-x\overline{y})|. \tag{3.10}$$

Assuming that f, g have mean zero, repeated applications of integration by parts yield that

$$\begin{aligned} \mathrm{Cov}((h,f),(h,g)) &= (2\pi)^2 \mathrm{Cov}((h,\Delta^{-1}f)_\nabla,(h,\Delta^{-1}g)_\nabla) \\ &= \iint_{D\times D} f(x) G^{\mathrm{F}}(x,y) g(y) dx dy \end{aligned} \tag{3.11}$$

where G^{F} is the Neumann Green's function on D.

We note that the choice of taking the normal derivative of $G^{\mathrm{F}}(x,\cdot)$ being equal to 1 at every point on ∂D is somewhat arbitrary and it is in fact not invariant under conformal transformations. This may lead one to worry that the law of the free boundary GFF is not invariant under conformal transformations. Let us now explain why the definition given above is in fact conformally invariant. Suppose that D has smooth boundary, $\widetilde{D}$ is another domain with smooth boundary, and that $\phi\colon \widetilde{D} \to D$ is a conformal transformation. Then $\widetilde{G}^{\mathrm{F}}(x,y) = G^{\mathrm{F}}(\phi(x),\phi(y))$ solves $\Delta\widetilde{G}^{\mathrm{F}}(x,y) = -2\pi\delta_x(\cdot)$ in the distributional sense and the normal derivative of $\widetilde{G}^{\mathrm{F}}(x,\cdot)$ evaluated at $y \in \partial\widetilde{D}$ is equal to the normal derivative of ϕ at y. The important point here, however, is that this normal derivative is constant when y is fixed and x is allowed to vary. In particular, the normal derivative terms still cancel when one performs the integration by parts in the calculation in (3.11) as both f, g have mean zero.

We note that our definition of the GFF with free boundary conditions is equivalent to the definition given in [DS11a, Section 6.1] and in [She10, Section 3.3].[2]

3.2.10 Markov property of the GFF

We are now going to explain the Markov property enjoyed by the GFF with Dirichlet, free, or mixed boundary conditions. For simplicity, for the present discussion we are going to assume that h is a GFF with zero boundary conditions (though the proposition stated below is general and so is the following argument). Suppose that $W \subseteq D$ with $W \neq D$ is open. There is a natural inclusion ι of $H_0(W)$ into $H_0(D)$ where

$$\iota(f)(x) = \begin{cases} f(x) \text{ if } x \in W, \\ 0 \text{ otherwise.} \end{cases}$$

[2] Our choice of Neumann Green's function differs from that which is implicit in [DS11a, Section 6.1] because in [DS11a] the Neumann Green's function is used to describe the covariance function for the Gaussian process given by the average of the field on $\partial B(z,\epsilon)$ *minus* its average on all of D.

If $f \in C_0^\infty(W)$ and $g \in C_0^\infty(D)$, then as $(f,g)_\nabla = -(2\pi)^{-1}(f, \Delta g)$ it is easy to see that $H_0(D)$ admits the $(\cdot,\cdot)_\nabla$-orthogonal decomposition $H_0(W) \oplus H_0^\perp(W)$ where $H_0^\perp(W)$ is the subspace of functions in $H_0(D)$ which are harmonic in W. Thus we can write

$$h = h_W + h_{W^c} = \sum_n \alpha_n^W f_n^W + \sum_n \alpha_n^{W^c} f_n^{W^c}$$

where $(\alpha_n^W), (\alpha_n^{W^c})$ are independent i.i.d. sequences of standard Gaussians and (f_n^W), $(f_n^{W^c})$ are orthonormal bases of $H_0(W)$ and $H_0^\perp(W)$, respectively. Observe that h_W is a zero-boundary GFF on W, h_{W^c} is almost surely harmonic in W, and h_W and h_{W^c} are independent. We interpret h_{W^c} as the harmonic extension of the values of $h|_{\partial W}$ from ∂W to W (of course this does not make literal sense because h does not value "values" on W as it is only a distribution valued random variable). We arrive at the following proposition:

Proposition 3.2 (Markov Property)**.** *Suppose that h is a GFF with Dirichlet, free, or mixed boundary conditions. The conditional law of $h|_W$ given $h|_{D\setminus W}$ is that of the sum of a zero boundary GFF on W plus the harmonic extension of $h|_{\partial W}$ from ∂W to W. (In the case that h has free boundary conditions, this harmonic extension is only defined modulo additive constant.)*

The orthogonality of $H_0(W)$ and the set of functions in $H_0(D)$ which are harmonic in W is also proved in [She07, Theorem 2.17] and it is explained thereafter how this is related to the Markov property of the GFF.

There are other equivalent ways of defining the harmonic extension of the values of $h|_{\partial W}$ to W. Let us pause to mention one other important such possibility and why this construction is equivalent to h_{W^c} defined just above. Assume that W is simply connected and has smooth boundary and let $\mathcal{P}_W$ be the Poisson kernel for W. We cannot naively define $\int_{\partial W} h(w)\mathcal{P}_W(z,w)dw$ where dw denotes Lebesgue measure on ∂W using only that h defines a distribution on D because it is not possible to represent the harmonic measure $\mathcal{P}_W(z,w)dw$ of ∂W as seen from z as a $C_0^\infty(D)$ function. We can, however, make sense of this integral indirectly as follows. (See [DS11a, Section 3] for an analogous construction in the context of the circle average process.) We note that, for each fixed $z \in D$, $\zeta_W^z(w) = -2\pi\Delta^{-1}\mathcal{P}_W(z,w)dw \in H_0(D)$. Indeed, this can be seen because we can write

$$\zeta_W^z(w) = \int_{\partial W} G^{\mathrm{D}}(w,u)\mathcal{P}_W(z,u)du$$

where G^{D} is the Green's function for Δ on D and we extend $\mathcal{P}_W$ from W to D by value 0. Consequently, $\widetilde{h}_{W^c}(z) := -2\pi(h, \zeta_W^z)_\nabla$ is defined, at least almost surely for each fixed z (off a possibly z-dependent set of measure of 0). Since (formally) integrating by parts, we have that

$$\widetilde{h}_{W^c}(z) = (h, \zeta_W^z)_\nabla = \int_{\partial W} h(w)\mathcal{P}_W(z,w)dw,$$

we see that $\widetilde{h}_{W^c}$ is a natural definition of the harmonic extension of the values of h from ∂W to W. One can see $\widetilde{h}_{W^c}$ is almost surely defined for all z simultaneously and continuous in z by bounding the moments of its increments and using the Kolmogorov-Čentsov theorem. (See [DS11a, Section 3] for similar discussion in the context of the circle average process for the GFF.) One can also see that $\widetilde{h}_{W^c} = h_{W^c}$ as follows (using the orthogonal decomposition introduced just before the statement of Proposition 3.2):

$$\begin{aligned}
\widetilde{h}_{W^c}(z) &= \lim_{n\to\infty} \left(\sum_{j=1}^{n} \alpha_j^W f_j^W + \sum_{j=1}^{n} \alpha_j^{W^c} f_j^{W^c}, \zeta_W^z \right)_\nabla \\
&= \lim_{n\to\infty} \sum_{j=1}^{n} \left(\alpha_j^{W^c} f_j^{W^c}, \zeta_W^z \right)_\nabla \quad \text{(recall that } f_j^W \text{ is supported in } W\text{)} \\
&= \lim_{n\to\infty} \sum_{j=1}^{n} \alpha_j^{W^c} \int_{\partial W} f_j^{W^c}(w) \mathcal{P}_W(z,w) dw \\
&= \lim_{n\to\infty} \sum_{j=1}^{n} \alpha_j^{W^c} f_j^{W^c}(z) \quad \text{(recall that } f_j^{W^c} \text{ is harmonic in } W\text{)} \\
&= h_{W^c}(z).
\end{aligned}$$

Remark 3.3. Proposition 3.2 implies that if h is an instance of the free boundary GFF on D then we can write h as the sum of the harmonic extension of its boundary values from ∂D to D and an independent zero boundary GFF in D.

Remark 3.4. Proposition 3.2 implies that for each fixed $W \subseteq D$ open we can almost surely define the orthogonal projection of a GFF h onto the subspaces of functions which are harmonic in and supported in W. We will indicate these by $P_{\text{harm}}(h;W)$ and $P_{\text{supp}}(h;W)$, respectively. If W is clear from the context, we will simply write $P_{\text{harm}}(h)$ and $P_{\text{supp}}(h)$.

3.3 Local sets of the GFF

The theory of local sets, developed in [SS13], extends the Markovian structure of the field (Proposition 3.2) to the setting of conditioning on the values it takes on a *random set* $A \subseteq D$. More precisely, suppose that (A,h) is a coupling of a GFF (with either Dirichlet, free, or mixed boundary conditions) h on D and a random variable A taking values in the space Γ of closed subsets of $\overline{D}$, equipped with the Hausdorff metric. Then A is said to be a **local set** of h [SS13, Lemma 3.9, part (4)] if there exists a law on pairs (A,h_1) where h_1 takes values in the space of distributions on D with $h_1|_{D\setminus A}$ harmonic is such that a sample with the law (A,h) can be produced by

1. choosing the pair (A,h_1),

2. then sampling an instance h_2 of the zero boundary GFF on $D \setminus A$ and setting $h = h_1 + h_2$.

There are several other characterizations of local sets which are discussed in [SS13, Lemma 3.9]. These are stated and proved for the GFF with Dirichlet boundary conditions, however the argument goes through verbatim for the GFF with either free or mixed boundary conditions. For the convenience of the reader, we restate this result here:

Lemma 3.5. *([SS13, Lemma 3.9]) Suppose that (A, h) is a random variable which is a coupling of an instance h of the GFF on D (with either Dirichlet, mixed, or free boundary conditions) with a random element A of Γ. Then the following are equivalent:*

(i) *For each deterministic open $U \subseteq D$, we have that* given *the projection of h onto $H^\perp(U)$, the event $A \cap U = \emptyset$ is independent of the projection of h onto $H(U)$. In other words, the conditional probability that $A \cap U = \emptyset$ given h is a measurable function of the projection of h onto $H^\perp(U)$.*

(ii) *For each deterministic open $U \subseteq D$, we let S be the event that A intersects U and let*

$$\widetilde{A} = \begin{cases} A & \text{on} \quad S^c, \\ \emptyset & \text{otherwise}. \end{cases}$$

Then we have that given *the projection of h onto $H^\perp(U)$, the pair $(S, \widetilde{A})$ is independent of the projection of h onto $H(U)$.*

(iii) *Conditioned on $\mathcal{A}$, (a regular version of) the conditional law of h is that of $h_1 + h_2$ where h_2 is the GFF with zero boundary values on $D \setminus A$ (extended to all of D) and h_1 is an $\mathcal{A}$-measurable random distribution (i.e., as a distribution-valued function on the space of distribution-set pairs (A, h), h_1 is $\mathcal{A}$-measurable) which is almost surely harmonic on $D \setminus A$.*

(iv) *A is a local set for h. That is, a sample with the law of (A, h) can be produced as follows. First choose the pair (A, h_1) according to some law where h_1 is almost surely harmonic on $D \setminus A$. Then sample an instance h_2 of the GFF on $D \setminus A$ and set $h = h_1 + h_2$.*

For a given local set A, we will write $\mathcal{C}_A$ for h_1 as above. We can think of $\mathcal{C}_A$ as being given by $P_{\text{harm}}(h; D \setminus A)$. We can also interpret $\mathcal{C}_A$ as the conditional expectation of h given A and $h|_A$. In the case that h is a GFF with free boundary conditions, $\mathcal{C}_A$ is defined modulo additive constant.

Throughout this article, we will often work with increasing families of closed subsets $(K_t)_{t \geq 0}$ each of which is local for a GFF h. The following is a restatement of [MS12a, Proposition 6.5] and describes the manner in which $\mathcal{C}_{K_t}$ evolves with t. In the following statement, for a domain $U \subseteq \mathbf{C}$ with simply-connected components and $z \in U$, we write $\mathrm{C}R(z;U)$ for the conformal radius of the component U_z of U containing z as seen from z. That is, $\mathrm{C}R(z;U) = \phi'(0)$ where ϕ is the unique conformal map which takes $\mathbf{D}$ to U_z with $\phi(0) = z$ and $\phi'(0) > 0$.

Proposition 3.6. *Suppose that $D \subseteq \mathbf{C}$ is a non-trivial simply connected domain. Let h be a GFF on D with either Dirichlet, free, or mixed boundary conditions. Suppose that $(K_t)_{t\geq 0}$ is an increasing family of closed sets such that K_τ is local for h for every (K_t) stopping time τ and $z \in D$ is such that $\mathrm{C}R(z; D \setminus K_t)$ is almost surely continuous and monotonic in t. Then $\mathcal{C}_{K_t}(z) - \mathcal{C}_{K_0}(z)$ has a modification which is a Brownian motion when parameterized by $\log \mathrm{C}R(z; D \setminus K_0) - \log \mathrm{C}R(z; D \setminus K_t)$ up until the first time $\tau(z)$ that K_t accumulates at z. In particular, $\mathcal{C}_{K_t}(z)$ has a modification which is almost surely continuous in $t \geq 0$. (In the case that h has free boundary conditions, we use the normalization $\mathcal{C}_{K_0}(z) = 0$.)*

3.4 Fractional and log-correlated Gaussian fields

The GFF can be generalized in several ways. See the survey articles [DRSV14b, LSSW14] for more on fractional Gaussian fields and log correlated Gaussian fields in general d-dimensional spaces. These are obtained by applying powers of the Laplacian to white noise. The Gaussian free field can be understood as the restriction to two dimensions of log correlated fields defined in higher dimensions.

3.5 Dimer models and uniform spanning trees

The UST height function is arguably the simplest discrete analog of the GFF. See Kenyon's scaling limit proof [Ken00b, Ken01a], which makes use of the equivalent formulation of the model in terms of dimers.

4 Random curves and loop ensembles

4.1 Schramm-Loewner evolution: basic definitions and phases

Much of the work on Schramm-Loewner evolution is prefigured in the physics literature on conformal field theory [DFMS97]. Schramm's original paper [Sch00a] has been followed by many excellent survey articles and textbooks [Wer03, KN04, Car06, Law09a, BN11]. The so-called natural parameterization is described in [LS11, LR12, LZ13].

The SLE_κ are a one-parameter family of conformally invariant random curves, indexed by a parameter κ that roughly encodes how "windy" the curves are. These curves were introduced by Oded Schramm in [Sch00a] as a candidates for (and later proved to be) the scaling limits of loop erased random walk [LSW04a] and the interfaces in critical percolation [Smi01a, CN06a]. Schramm's curves have been shown so far also to arise as the scaling limit of the macroscopic interfaces in several other models from statistical physics: [Smi10a, CS, SS05a, SS09a, Mil10b]. More detailed introductions to SLE can be found in many excellent survey articles of the subject, e.g., [Wer04a, Law05a].

An SLE_κ in $\mathbf{H}$ from 0 to ∞ is defined by the random family of conformal maps g_t obtained by solving the Loewner ODE (7.14) with $W = \sqrt{\kappa}B$ and B a standard Brownian motion. Write $K_t := \{z \in \mathbf{H} : \tau(z) \le t\}$ where $\tau(z) = \sup\{t \ge 0 : \mathrm{Im}(g_t(z)) > 0\}$. Then g_t is the unique conformal map from $\mathbf{H}_t := \mathbf{H} \setminus K_t$ to $\mathbf{H}$ satisfying $\lim_{|z|\to\infty} |g_t(z) - z| = 0$.

Rohde and Schramm showed that there a.s. exists a curve η (the so-called SLE *trace*) such that for each $t \ge 0$ the domain $\mathbf{H}_t$ of g_t is the unbounded connected component of $\mathbf{H} \setminus \eta([0,t])$, in which case the (necessarily simply connected and closed) set K_t is called the "filling" of $\eta([0,t])$ [RS05a]. An SLE_κ connecting boundary points x and y of an arbitrary simply connected Jordan domain can be constructed as the image of an SLE_κ on $\mathbf{H}$ under a conformal transformation $\varphi\colon \mathbf{H} \to D$ sending 0 to x and ∞ to y. (The choice of φ does not affect the law of this image path, since the law of SLE_κ on $\mathbf{H}$ is scale invariant.)

4.2 Definition of $\mathrm{SLE}_\kappa(\rho)$

The so-called $\mathrm{SLE}_\kappa(\underline{\rho})$ processes are an important variant of SLE_κ in which one keeps track of additional marked points. Just as with regular SLE_κ, one constructs $\mathrm{SLE}_\kappa(\underline{\rho})$ using the Loewner equation except that the driving function W is replaced with a solution to the SDE (7.18). The purpose of this section is to construct solutions to (7.18) in a careful and canonical way. We will not actually need to think about the Loewner evolution on the half plane for any of the discussion in this subsection. It will be enough for now to think about the Loewner evolution restricted to the real line.

Fix a value $\rho > -2$ and write

$$\delta = 1 + \frac{2(\rho+2)}{\kappa},$$

noting that $\delta > 1$. Let X be an instantaneously reflecting solution to (1.1) for some $X_0 = x_0 \ge 0$. We would like to define a pair W and V^R that solves the SDE (7.18) with $W_0 = 0$ and some fixed initial value $V_0^R = x_0^R \ge 0$. To motivate the definition, note that (7.18) formally implies that the difference $V^R - W$ solves the same SDE as $\sqrt{\kappa}X$, away from times where it is equal to zero. Thus it is natural to write

$$\begin{aligned} V_t^R &= x_0^R + \int_0^t \frac{2}{\sqrt{\kappa}X_s} ds, \\ W_t &= V_t^R - \sqrt{\kappa}X_t. \end{aligned} \tag{4.1}$$

The standard definition of (single-force-point) $\mathrm{SLE}_\kappa(\rho)$ is the Loewner evolution driven by the process W defined in (4.1).

4.3 Loop erased random walk and uniform spanning tree

See Wilson's algorithm [Wil96, PW98] and the original UST/LERW convergence paper [LSW04b].

4.4 Critical percolation interfaces

Percolation interface scaling limits are tractable thanks to a fundamental discovery by Stanislav Smirnov [Smi01a].

4.5 Gaussian free field level lines

See [SS05b, SS09b, SS13] and the universality theorem in [Mil10c].

4.6 Ising, Potts, and FK-cluster models

Some of these "next simplest after percolation" models are also tractable [CS]

4.7 Bipolar orientations

This is another simple model conjectured to scale to SLE_{12}. The conjecture is easy to state, but the motivation behind the conjecture will not be explained until the sections on imaginary geometry and the peanosphere.

4.8 Restriction measures, self-avoiding walk, and loop soups

The relationship between $\mathrm{SLE}_{8/3}$ and Brownian motion is especially beautiful and has an especially beautiful history. See the account in the early work by Lawler, Schramm, and Werner [LSW03a]. Loop soups provide a natural construction of CLE_κ for $\kappa \in (8/3, 4]$.

4.9 Conformal loop ensembles

Given that the discrete interfaces that scale to SLE have "loop ensemble" variants, one would expect there to be a natural "loop ensemble" variant of SLE itself. See the introduction in [She09a, SW12].

4.10 Forward and reverse radial SLE_κ

For $u \in \partial\mathbf{D}$ and $z \in \mathbf{D}$, let

$$\Psi(u,z) = \frac{u+z}{u-z} \quad \text{and} \quad \Phi(u,z) = z\Psi(u,z). \tag{4.2}$$

A radial SLE_κ in $\mathbf{D}$ starting from 1 and targeted at 0 is described by the random family of conformal maps obtained by solving the radial Loewner ODE

$$\dot{g}_t(z) = \Phi(e^{iW_t}, g_t(z)), \quad g_0(z) = z \tag{4.3}$$

where $W_t = \sqrt{\kappa}B_t$ and B is a standard Brownian motion. We refer to e^{iW_t} as the **driving function** for (g_t). For each $z \in \mathbf{D}$ let $\tau(z) = \sup\{t \geq 0 : |g_t(z)| < 1\}$ and write $K_t := \{z \in \mathbf{D} : \tau(z) \leq t\}$. For each $t \geq 0$, g_t is the unique conformal map which takes $\mathbf{D}_t := \mathbf{D} \setminus K_t$ to $\mathbf{D}$ with $g_t(0) = 0$ and $g_t'(0) > 0$. Time is parameterized by log conformal radius so that $g_t'(0) = e^t$ for each $t \geq 0$. Rohde and Schramm showed that there almost surely exists a curve η (the so-called SLE **trace**) such that for each $t \geq 0$ the domain $\mathbf{D}_t$ of g_t is equal to the connected component of $\mathbf{D} \setminus \eta([0,t])$ which contains 0. The (necessarily simply connected and closed) set K_t is called the "filling" of $\eta([0,t])$ [RS05a].

In our construction of QLE, it will sometimes be more convenient to work with *reverse* radial SLE_κ rather than *forward* radial SLE_κ (as defined in (4.3)). A reverse SLE_κ in $\mathbf{D}$ starting from 1 and targeted at 0 is the random family of conformal maps obtained by solving the reverse radial Loewner ODE

$$\dot{g}_t(z) = -\Phi(e^{iW_t}, g_t(z)), \quad g_0(z) = z \tag{4.4}$$

where $W_t = \sqrt{\kappa}B_t$ and B is a standard Brownian motion. As in the forward case, we refer to e^{iW_t} as the **driving function** for (g_t).

Remark 4.1. Forward and reversal radial SLE_κ are related in the following manner. Suppose that (g_t) solves the reverse radial Loewner equation (4.4) with driving function $W_t = \sqrt{\kappa}B_t$ and B a standard Brownian motion. Fix $T > 0$ and let $f_t = g_{T-t}$ for $t \in [0,T]$. Then (f_t) solves the forward radial Loewner equation with driving function $t \mapsto W_{T-t}$ and with initial condition $f_0(z) = g_T(z)$. Equivalently, we can let q_t for $t \in [0,T]$ solve (4.3) with driving function $t \mapsto W_{T-t}$ and $q_0(z) = z$ and then take $f_t = q_t \circ g_T$. Indeed, this follows from standard uniqueness results for ODEs. Then

$$z = g_0(z) = f_T(z) = q_T(g_T(z)).$$

That is, q_T is the inverse of g_T. This implies that the image of g_T can be expressed as the complementary component containing zero of an SLE_κ curve η_T in $\mathbf{D}$ drawn up to log conformal radius time T. We emphasize here that the path η_T changes with T. (See the end of the proof of [Law05b, Theorem 4.14] for a similar discussion.)

5 Random surfaces

5.1 Planar maps

A planar map is a planar map together with an embedding in the plane (defined up to topological equivalence). Enumeration work was done by Tutte in the 1960's [Tut62, Tut68a].

5.2 Decorated surfaces and Laplacian determinants

The Laplacian determinant and its inverse are related to partition functions for the GFF and UST models in surprisingly simple ways. See Kenyon's work on scaling limits of determinant Laplacians on grids [Ken00a] and the broad survey by Merris [Mer94] which describes Kirchhoff's matrix tree theorem, among other things. See also Sarnak's work on height functions and the Polyakov-Alvarez formula in [Sar90], and works by Dubédat and by Lawler on partition functions and the Polyakov-Alvarez formula in the SLE context [Law09b, Dub09b].

Given any finite connected graph (V, E) the Laplacian on the graph can be defined as a linear operator Δ from $\mathbf{R}^V$ itself. Its matrix is given by

$$M_{i,j} = \begin{cases} 1 & i \neq j, (v_i, v_j) \in E \\ 0 & i \neq j, (v_i, v_j) \notin E \\ -\deg(v_i) & i = j. \end{cases}.$$

Let $R \subset \mathbf{R}^V$ be the set of functions with mean zero. Then $-\Delta : R \to R$ is invertible, and Kirchhoff's matrix tree theorem states that if α is the determinant of this invertible operator on R then α is the number of spanning trees of V. The quantity α is also the product of all of the non-zero eigenvalues of the matrix M.

The DGFF partition function can be be written $\int_R (2\pi)^{-|V-1|/2} e^{-(f,-\Delta f)/2} df$. Expanding over eigenbases, and using the fact the $\frac{1}{\sqrt{2\pi}} \int e^{-tx^2/2} dt = t^{-1/2}$, we find that quantity is $\alpha^{-1/2}$.

Two ways to measure size of a connected graph: number of edges (the log of the number of edge subsets) and the log of the number of spanning trees. For now, let A be first number, B second. Then $A \geq B$ with equality only if the graph is a tree. If we choose a random planar map from the Boltzmann measure, these two size measures are coupled and random, then we expect the pair $(A/n, B/n)$ to satisfy a large deviations principle as $n \to \infty$, with a rate function that is linear on lines through the origin. If we weight the original by $e^{aA-cB/2}$ for appropriately chosen a and c then we expect the measure to have a power law decay.

In addition to weighting by determinant Laplacian powers, another way to interpolate involves the Tutte polynomial: namely, the FK cluster model partition function.

5.3 Mullin-Bernardi bijection

There is a very simple bijection between discrete lattice walks in $\mathbb{Z}_+^2$ starting and ending at zero and rooted planar maps with distinguished spanning trees. See [Mul67a, Ber07] as well as the exposition in [She11a].

5.4 Cori-Vauquelin-Schaeffer bijection

The Cori-Vaquelin-Schaeffer bijection gives a way to bijectively count *undecorated* planar maps [CV81, JS98, Sch99]. Every quadrangulation with a root can be decorated by a directed breadth first spanning tree spannign all of the edges. When multiple incoming edges come into the same vertex, each outgoing edge is connected to only one of them in this tree namely, the next one over in clockwise ordering.

5.5 Hamburger-cheeseburger bijection

There is a generalization of the Mullin-Bernardi bijection in which the rooted planar map comes with an arbitrary distinguished edge subset, instead of a distinguished spanning tree [She11a].

5.6 Bipolar bijection

The scaling limit of the pair of trees can be easily described in this case, as it can in each of the other cases described above.

5.7 Brownian map

The idea behind the Cori-Vaquelin-Schaeffer bijection can be used to define a continuum random metric space [MM06a, LG13a, Mie13a, LG14], which has a natural infinite volume analog [CL12]. See Le Gall's ICM notes [Le 14] or the survey by Miermont and Le Gall [LGM$^+$12]. An axiomatic characterization of the Brownian map in terms of it symmetries appears in [MS15a].

5.8 Liouville quantum gravity

Polaykov conceived of a random surface model based on an action closely related to the Gaussian free field [Pol81a]. If h is an instance of the Gaussian free field, one attempts to define a meausre of the form $e^{\gamma h(z)}dz$, which in turn encodes the volume form of a random surface after a conformal map back to a fixed parameter space (say, a disk in the plane). The rigorous construction of this random measure was given by Høegh-Krohn in 1971 [HK71], for the range $\gamma \in [0, \sqrt{2})$, and the full range $[0, 2)$ was treated by Kahane (who used the term *multiplicative chaos*) in 1985 [Kah85], see also the survey [RV14]. The construction of the measure as a measure-valued function on the space of instances h of the GFF was done in [DS11a]. The case $\gamma = 2$ is different but one can make sense of the measure by different means [DRSV14a, DRSV14c].

5.9 KPZ (Knizhnik-Polyakov-Zamolodchikov) scaling relations

A relationship between scaling dimensions was discovered by Knizhnik, Polyakov, and Zamolodchikov in [KPZ88a]. In a recent memoir [Pol08a] Polyakov explains how the discover of this relationship cemented the belief that the discrete planar map models were (in some sense) equivalent to Liouville quantum gravity. See [DS11a] and the references therein. See the Hausdorff variant in [RV08a].

See the derivation of the $d = 26$ value for the bosonic string by Lovelace in 1971 in [Lov71].

5.10 Quantum wedges, cones, spheres and disks

There are natural ways to define quantum surfaces using Bessel process excursions. There are some natural probability measures on the space of infinite volume surfaces. There are also some natural infinite measures on the space of finite volume surfaces. See the introduction to [DMS14].

5.11 Quantum boundary length measures

We now summarize a few important facts which are based on the discussion in [DS11a, Section 6] regarding the Liouville quantum gravity boundary length measure. Suppose that h is a GFF with mixed Dirichlet/free boundary conditions on a Jordan domain $D \subseteq \mathbf{C}$ where both the Dirichlet and free parts ∂^{D} and ∂^{F}, respectively, of ∂D are non-degenerate boundary arcs. In [DS11a, Theorem 6.1], it is shown how to construct the measure $\nu_h^\gamma = \exp(\frac{\gamma}{2}h(u))du$ on ∂^{F} for $\gamma \in (-2, 2)$ fixed. Formally, this means that the Radon-Nikodym derivative of ν_h^γ with respect to Lebesgue measure on ∂^{F} is given by $\exp(\frac{\gamma}{2}h)$. This does not make literal sense because h does not take values in the space of functions and, in particular, does not take on a specific value at a given point in ∂^{F}.

One can construct ν_h^γ rigorously as follows. First, suppose that ∂^{F} consists of a single linear segment. For each $z \in \partial^{\mathrm{F}}$ and $\epsilon > 0$, let $h_\epsilon(z)$ be the average of h on the semi-circle $\partial B(z,\epsilon) \cap D$. For completeness, let us briefly recall the construction and basic properties of $h_\epsilon(z)$ (we refer the reader to [DS11a, Section 3 and Section 6] for background on the circle average process).

We will start with the case that h has Dirichlet boundary conditions for simplicity. For each $z \in D$ and $\epsilon > 0$, we let ρ_ϵ^z denote the uniform measure on $D \cap \partial B(z,\epsilon)$. Then $h_\epsilon(z)$ is formally given by $\int h d\rho_\epsilon^z$. To make sense of this rigorously, when $B(z,\epsilon) \subseteq D$ we let

$$\xi_\epsilon^z(y) = -\log \max(\epsilon, |y-z|) - \widetilde{G}_{z,\epsilon}(y), \tag{5.1}$$

where $\widetilde{G}_{z,\epsilon}(y)$ is the harmonic extension of $y \mapsto -\log\max(\epsilon,|y-z|)$ from ∂D to D. Then $\Delta \xi_\epsilon^z = -2\pi \rho_\epsilon^z$ (in the distributional sense), so we can define $h_\epsilon(z)$ by setting $h_\epsilon(z) = (h, \xi_\epsilon^z)_\nabla$. This defines $h_\epsilon(z)$ as a Gaussian random variable which is almost surely defined off a set of measure zero which depends on ϵ and z. Consequently, for a fixed *countable* collection of ϵ, z pairs, we have that the Gaussian variables $h_\epsilon(z)$ are defined off a common set of measure zero. A continuity argument is needed in order to define $h_\epsilon(z)$ for *all* ϵ, z pairs off a common set of measure zero. This can be accomplished using the Kolmogorov-Čentsov theorem. In particular, by bounding the moments of the increments of $h_\epsilon(z)$ and using the Kolmogorov-Čentsov theorem, one can see that $(\epsilon, z) \mapsto h_\epsilon(z)$ has a Hölder continuous modification [DS11a, Proposition 3.1]. In the case that $B(z,\epsilon), B(w,\delta) \subseteq D$ are disjoint, $\mathrm{Cov}(h_\epsilon(z), h_\delta(w))$ is given by the Green's function for Δ on D with Dirichlet boundary conditions evaluated at (z,w). In the case that $z = w$ so that $B(z,\epsilon), B(w,\delta)$ are concentric, and $\epsilon \leq \delta$ then $\mathrm{Cov}(h_\epsilon(z), h_\delta(z)) = -\log\epsilon + \log \mathrm{C}R(z;D)$.

One can construct and analyze the circle average process in the case that h has mixed boundary conditions by relating the GFF in this case to a GFF with Dirichlet boundary conditions using the so-called odd/even decomposition (see [DS11a, Section 6] and [She10, Section 3.2]). For this purpose, we suppose that $D \subseteq \mathbf{H}$ and that ∂D contains an interval $[a,b] \in \mathbf{R}$ with $a < b$. Then the law of the GFF on D with free (resp. Dirichlet) boundary conditions on $\partial^{\mathrm{F}} = [a,b]$ (resp. $\partial^{\mathrm{D}} = \partial D \setminus \partial^{\mathrm{F}}$) can be sampled from as follows:

1. Sample $h^\dagger$ as a GFF on $D^\dagger = D \cup \overline{D}$ where $\overline{D} = \{\overline{z} : z \in D\}$ with Dirichlet boundary conditions.
2. For $\rho \in C_0^\infty(D)$, set $(h,\rho) = \frac{1}{\sqrt{2}}(h^\dagger, \rho + \overline{\rho})$ where $\overline{\rho}(z) = \rho(\overline{z})$.

Using this representation of h, we can define the average of h on $D \cap \partial B(z,\epsilon)$ for $z \in [a,b]$ by setting it equal to the average of $h^\dagger$ on $\partial B(z,\epsilon)$.

Finally, in the case that h has purely free boundary conditions one can still construct the circle average process at points on ∂D along which ∂D is linear by using absolute continuity to compare the law of the GFF with free boundary conditions to the law of the GFF with mixed boundary conditions.

For each $\gamma \in (-2,2)$, the measure $\exp(\frac{\gamma}{2}h(u))du$ is defined as the almost sure limit

$$\nu_h^\gamma = \lim_{\epsilon \to 0} \epsilon^{\gamma^2/4} \exp(\tfrac{\gamma}{2}h_\epsilon(u))du \tag{5.2}$$

along negative powers of two as $\epsilon \to 0$ with respect to the weak topology. Upon showing that the limit in (5.2) exists for linear ∂^{F}, the boundary measure for other domains is defined via conformal mapping and applying the change of coordinates rule for quantum surfaces.

One can similarly make sense of the limit in (5.2) in the case that h has free boundary conditions, i.e. $\partial^{\mathrm{D}} = \emptyset$. If we consider h as a distribution defined modulo additive constant, then the measure ν_h^γ will only be defined up to a multiplicative constant. This means that if $A, B \subseteq \partial D$ with $\nu_h^\gamma(B) \in (0,\infty)$ then the value of $\nu_h^\gamma(A)/\nu_h^\gamma(B)$ is well-defined but the values of $\nu_h^\gamma(A)$ and $\nu_h^\gamma(B)$ are not (however, the event $\nu_h^\gamma(B) \in (0,\infty)$ does make sense). We can "fix" the additive constant in various ways, for example by:

1. Taking the average of h on an open set to be equal to 0 (or any other fixed real number),
2. Taking the integral of h against a given test function ρ with $\int \rho \neq 0$ to be equal to 0 (or any other fixed real number),
3. Setting $\nu_h^\gamma(A) = 1$ (or any other fixed value in $(0,\infty)$) for some $A \in \partial D$ such that $\nu_h^\gamma(A) \in (0,\infty)$ almost surely.

With any of these choices, we get that ν_h^γ is an actual measure. In article, we will have $D = \mathbf{D}$ and typically normalize so that $\nu_h^\gamma(\partial \mathbf{D}) = 1$. This normalization is convenient because we will use ν_h^γ to sample points from $\partial \mathbf{D}$ in the construction of QLE, so we would like to think of ν_h^γ as a probability measure on $\partial \mathbf{D}$.

One also has the following analog of [DS11a, Proposition 1.2] for the boundary measures associated with the free boundary GFF on $\mathbf{D}$. Suppose that (f_n) is any orthonormal basis consisting of smooth functions for the Hilbert space used to define h. For each $n \in \mathbf{N}$, let h^n be the orthogonal projection of h onto the subspace spanned by $\{f_1, \ldots, f_n\}$. One can similarly define ν_h^γ as the almost sure limit

$$\nu_h^\gamma = \lim_{n \to \infty} \exp\left(\tfrac{\gamma}{2}h^n(u) - \tfrac{\gamma^2}{4}\mathrm{Var}(h^n(u))\right) du. \tag{5.3}$$

That these two definitions for ν_h^γ almost surely agree is not explicitly stated in [DS11a] for boundary measures however its proof is exactly the same as in the case of bulk measures which is given in [DS11a, Proposition 1.2].

Proposition 5.1. *Fix $\gamma \in (-2,2)$. Consider a random pair (h,u) where u is sampled uniformly from $\partial \mathbf{D}$ using Lebesgue measure and, given u, the conditional law of h is that of a free boundary GFF on $\mathbf{D}$ plus $-\gamma \log|\cdot - u|$ viewed as a distribution defined* *modulo additive constant. Let ν_h^γ denote the γ boundary measure associated with h. Then given h, the conditional law of u is that of a point uniformly sampled from ν_h^γ (as explained above, ν_h^γ is only defined up to a multiplicative constant, but can be normalized to be a probability measure).*

Proof. Let A_r for $0 < r < 1$ be the annulus $\mathbf{D} \setminus \overline{B(0,r)}$. Let $\widetilde{A}_r$ be the larger annulus $B(0,1/r) \setminus \overline{B(0,r)}$. Let dh be the law of an instance h be the GFF on A_r with zero boundary conditions on the inner boundary circle $\partial B(0,r)$ and free boundary conditions on $\partial\mathbf{D}$. Let ν_h^γ denote the boundary γ-LQG measure associated with h on $\partial\mathbf{D}$. Since the h that we are considering here has mixed boundary conditions (and not purely free), we note that ν_h^γ is a well-defined measure (i.e., there is no need to fix the multiplicative constant). Then it is not hard to see that the following ways to produce a random pair u, h are equivalent (and very similar statements are proved in [DS11a, Section 6]):

1. First sample u uniformly on $\partial\mathbf{D}$ and the let h be a sample from the law described above *plus* the deterministic function $f_{u,r}(\cdot) = \gamma G_{\widetilde{A}_r}(u, \cdot)$ where $\gamma G_{\widetilde{A}_r}$ is the Green's function on $\widetilde{A}_r$.
2. First sample h from the measure $\nu_h^\gamma(\partial\mathbf{D})dh$ and then, conditioned on h, sample u from the boundary measure ν_h (normalized to be a probability measure).

Indeed, to see the equivalence of these two methods of sampling, we first recall that by the odd/even decomposition of the GFF, the law of h can also be sampled from by:

1. Sampling a GFF $h^\dagger$ on $\widetilde{A}_r$ with zero-boundary conditions and then
2. For $\rho \in C_0^\infty(A_r)$, setting $(h, \rho) = \frac{1}{\sqrt{2}}(h^\dagger, \rho + \overline{\rho})$ where $\overline{\rho}(z) = \rho(1/\overline{z})$.

Let ξ_ϵ^u be as in (5.1). The boundary measure ν_h^γ is then given by the almost sure weak limit as $\epsilon \to 0$ along negative powers of two of

$$\epsilon^{\gamma^2/4} \exp(h_\epsilon^\dagger(u))du = \epsilon^{\gamma^2/4} \exp((h^\dagger, \xi_\epsilon^u)_\nabla)du.$$

Consider the measure on pairs $(u, h^\dagger)$ given by

$$\epsilon^{\gamma^2/4} \exp(h_\epsilon^\dagger(u))dudh^\dagger = \epsilon^{\gamma^2/4} \exp((h^\dagger, \xi_\epsilon^u)_\nabla)dudh^\dagger.$$

Recall that if $Z \sim N(0,1)$ then the law of Z weighted by $e^{\mu x}$ is $N(\mu, 1)$. Applying this to the coordinates in the expansion of $h^\dagger$ using an orthonormal basis of $H_0(\widetilde{A}_r)$, it is easy to see from the form of the law above that we can sample from it using both of the following two methods of sampling.

1. First sample u uniformly on $\partial\mathbf{D}$ and the let $h^\dagger$ be a sample from the law described above *plus* the deterministic function $\gamma\xi_\epsilon^u$.

2. First sample $h^\dagger$ from the measure $\nu_h^{\gamma,\epsilon}(\partial\mathbf{D})dh^\dagger$ where

$$\nu_h^{\gamma,\epsilon}(\partial\mathbf{D}) = \int_{\partial\mathbf{D}} \epsilon^{\gamma^2/4}\exp(h_\epsilon^\dagger(w))dw$$

and then, conditioned on $h^\dagger$, sample u from the boundary measure $\nu_h^{\gamma,\epsilon}$ (normalized to be a probability measure).

The claim follows because as $\epsilon \to 0$ we have that ξ_ϵ^z converges to $G_{\widetilde{A}_r}(z,\cdot)$ and $\nu_h^{\gamma,\epsilon}(\partial\mathbf{D}) \to \nu_h^\gamma(\partial\mathbf{D})$.

The lemma is proved by taking the limit $r \to 0$ (with the corresponding h being considered modulo additive constant). Note that on the set $\partial\mathbf{D}$, the functions $f_{u,r}(\cdot)$ (treated modulo additive constant) converge uniformly to $-\gamma\log|\cdot - u|$ as $r \to 0$. $\square$

6 Random growth trajectories

6.1 Eden model and first passage percolation

There are a number of natural growth trajectories. The Eden model, introduced by Edein in 1961 [Ede61], is the simplest to describe. Here, every edge has an exponential clock, and when it rings the edge is added to the growing edge cluster (if it is incident to the existing cluster). A generalization of this story known as first-passage percolation was introduced by Hammersley and Welsh in 1965 [HW65].

6.2 Diffusion limited aggregation and the dieelectric breakdown model

Diffusion limited aggregation, as introduced by Witten and Sander in 1981 [WJS81, WS83], is a model for growth in which new particle locations on the boundary are chosen from harmonic measure instead of uniform measure. See early conjectures in [Mea86] and the theorem of Kesten [Kes87].

Diffusion limited aggregation (DLA) was introduced by Witten and Sander in 1981 and has been used to explain the especially irregular "dendritic" growth patterns found in coral, lichen, mineral deposits, and crystals [WJS81, WS83].[3] Sander himself wrote a general overview of the subject in 2000 [San00]; see also the review [Hal00].

[3] Note that here (and throughout the remainder of this paper) we use the term DLA alone to refer to *external DLA*. The so-called *internal DLA* is a growth process introduced by Meakin and Deutch in 1986 [MD86] to explain the especially smooth growth/decay patterns associated with electropolishing, etching, and corrosion. Internal DLA clusters grow spherically with very small (log order) fluctuations, much smaller than the fluctuations observed for the Eden model on a grid. Although external DLA

The most famous and substantial result about planar external DLA to date is Kesten's theorem [Kes87], which states that the diameter of the DLA cluster obtained after n particles have been added almost surely grows asymptotically at most as fast as $n^{2/3}$. Another way to say this is that by the time DLA reaches radius n (for all n sufficiently large), there are least $n^{3/2}$ particles in the cluster. This seems to suggest (though it does not imply) that any scaling limit of DLA should have dimension at least $3/2$.[4] Although there is an enormous body of research on the behavior of DLA simulations, even the most basic questions about the scaling limit of DLA (such as whether the scaling limit is space-filling, or whether the scaling limit has dimension greater than 1) remain unanswered mathematically.

The effects of lattice anisotropy on DLA growth also remain mysterious. We mentioned above that limit shapes for FPP and Eden clusters need not be exactly round — the anisotropy of the lattice can persist in the limit. Intuitively, this makes sense: there is no particular reason, on a grid, to expect the rate of growth in the vertical direction to be exactly the same as the rate of growth in a diagonal direction. In the case of DLA, effects of anisotropy can be rather subtle, and it is hard to detect anything anisotropic from a glance at a DLA cluster like the one in Figure 10.5. Nonetheless, simulations suggest that anisotropy may also affect scaling limits for DLA (perhaps by decreasing the overall scaling limit dimension from about 1.7 to about 1.6). One recent overview of the scaling question (with many additional references) appears in [Men12], and effects of anisotropy are studied in [MS11]. There is some simulation-based evidence for universality among different isotropic "off-lattice" formulations of DLA (which involve differently-shaped dust particles performing Brownian motion until they attach themselves to a growing cluster) [LYTZC12]. There is also some evidence that different types of isotropic models (such as DLA and the so-called viscous fingering) have common scaling limits [MPST06]. Meakin proposed already in 1986 that off-lattice DLA and DLA on systems with five-fold or higher symmetry belong to one universality class, while DLA on systems with lower symmetry belong to one or more different universality classes [Mea86].

Closely related to DLA is the so-called Hele-Shaw flow, which is itself an active area of research. See, e.g., [LTW09] and the references therein as well as [Hal00] for a more expository account of the relationship between DLA and Hele-Shaw.

In the DLA simulations generated in this paper, the square to add to a cluster is essentially chosen by running a Brownian motion from far away and choosing the first

---

has had more attention in the physics literature, there has been much more mathematical progress on internal DLA, beginning with works by Diaconis and Fulton and by Lawler, Bramson, and Griffeath from the early 1990's [DF91, LBG92]. More recently, the second author was part of an IDLA paper series with Levine and Jerison that describes the size and nature of internal DLA fluctuations in great detail and relates these fluctuations to a variant of the GFF [JLS12a, JLS10, JLS11], see also [AG13a, AG13b].

[4] In his 2006 ICM paper, Schramm discussed the problem of understanding DLA on $\mathbf{Z}^2$ and wrote that Kesten's theorem "appears to be essentially the only theorem concerning two-dimensional DLA, though several very simplified variants of DLA have been successfully analysed" [Sch07].

cluster-adjacent square the Brownian motion hits. This is a little different from doing a simple random walk on the graph of squares started at a far away target vertex (and it was actually a little easier to code efficiently). It is possible that our approach is somehow "isotropic enough" to ensure that the growth models in the simulations converge to a universal isotropic scaling limit as δ tends to zero, but we do not know how to prove this. We stress that the QLE evolutions that we construct in this paper are rotationally invariant, and can thus only be scaling limits of growth models that have isotropic scaling limits.

6.3 KPZ (Kardar-Parisi-Zhang) growth

The KPZ growth model is the logarithm of the stochastic heat equation with geometric noise. It was introduced in a slightly different form, and without a rigorous construction, Karder, Parisi, and Zhang in [KPZ86]. It does not itself describe the conjectural scaling limit of Eden model fluctuations; rather, it describes what amounts to a sort of "off critical" variant, which is believed to converge to a fixed point as a certain parameter tends to zero. These models can be viewed as interesting in their right, or interesting as approximations to the (still conjectural) KPZ fixed point, which is in turn the conjectural scaling limit of Eden model fluctuations. The fixed point conjecture is described by Corwin and Quastel in [CQ11]. See also Corwin's survey article [Cor12a].

6.4 Hastings-Levitov

The Hastings-Levitov model was designed as an approximation of what should be a continuum DLA theory. The hope was that one could prove the existence of an isotropic scaling limits of this model, and that would be easier than establishing the analogous result for (an isotropic form of) ordinary DLA. While this goal has not yet been achieved, there has been some recent progress in understanding Hastings-Levitov; see, e.g., [JVST12].

6.5 Internal DLA

Internal DLA is a growth model introduced by Meakin and Deuthch in 1986 [MD86]. Internal DLA growth seems to be much smoother than Eden model, with logarithmic fluctuations [LBG92, JLS12b, JLS13, AG13b, AG13a]. Unlike ordinary (external) DLA and most of the other growth models presented in this section, fluctuations of internal DLA on the grid have a well understood scaling limit, which can be described by a variant of the Gaussian free field [JLS$^+$14].

7 Imaginary geometry

All readers are familiar with two dimensional Riemannian geometries whose Gaussian curvature is purely positive (the sphere), purely negative (hyperbolic space), or zero (the plane). In this paper, we study "geometries" whose Gaussian curvature is purely *imaginary*. We call them *imaginary geometries*.

Imaginary geometries have zero real curvature, which means (informally) that when a small bug slides without twisting around a closed loop, the bug's angle of rotation is unchanged. However, the bug's *size* may change (an *Alice in Wonderland* phenomenon that further justifies the term "imaginary").[5] "Straight lines" and "angles" are well-defined in imaginary geometry, and the angles of a triangle always sum to π, but "distance" is not defined.

A simply connected imaginary geometry can be described by a simply connected subdomain D of the complex plane $\mathbf{C}$ and a function $h\colon D \to \mathbf{R}$.[6] The angle-θ ray beginning at a point $z \in D$ is the flow line of $e^{i(h+\theta)}$ beginning at z, i.e., the solution to the ODE

$$\eta'(t) = e^{i(h(\eta(t))+\theta)} \quad \text{for} \quad t > 0, \quad \eta(0) = z \tag{7.1}$$

as in Figure 7.8.[7] In this paper we concern ourselves only with these rays, which we view as a simple and complete description of the imaginary geometry.[8] Our goal is to make sense of and study the properties of these flow lines when h is a constant multiple of a random *generalized* function called the Gaussian free field.

Many of the results in this section are developed in a series of *imaginary geometry* papers by the current authors [MS12a, MS12b, MS12c, MS13a]. The idea is to try to define flow lines of $e^{ih(z)/\chi}$ where $\chi > 0$ is a fixed parameter and h is an instance of the Gaussian free field. We begin by discussing paths that originate at the boundary, and are related to forms of chordal SLE.

[5]In both real and imaginary geometries, parallel transport about a simple loop multiplies a $\mathbf{C}$-identified tangent space by e^{iC} where C is the integral of the enclosed curvature; these transformations are rotations when C is real, dilations when C is imaginary.

[6]In the language of differential geometry, an *imaginary geometry* is a two dimensional manifold endowed with a torsion-free affine connection whose holonomy group consists entirely of dilations (c.f. ordinary Riemannian surfaces, whose Levi-Civita holonomy groups consist entirely of rotations), and straight lines are geodesic flows of the connection. The connection endows the manifold with a conformal structure, and by the uniformization theorem one can conformally map the geometry to a planar domain on which the geodesics are determined by some function h in the manner described here [She15].

[7]Imaginary geometries have also been called "altimeter-compass" geometries [She]. If the graph of h is viewed as a mountainous terrain, then a hiker holding an analog altimeter—with a needle indicating altitude modulo 2π—in one hand and a compass in the other can trace a ray by walking at constant speed (continuously changing direction as necessary) in such a way that the two needles always point in the same direction.

[8]This description is canonical up to conformal coordinate change, see Figure 7.13.

7.1 Forward coupling: flow lines of $e^{ih/\chi}$

Fix a planar domain D, viewed as a subset of $\mathbb{C}$, a function $h : D \to \mathbb{R}$, and a constant $\chi > 0$. An **AC ray** of h is a flow line of the complex vector field $e^{ih/\chi}$ beginning at a point $x \in \overline{D}$ — i.e., a path $\eta : [0, \infty) \to \mathbf{C}$ that is a solution to the ODE:

$$\eta'(t) := \frac{\partial}{\partial t}\eta(t) = e^{ih(\eta(t))/\chi} \text{ when } t > 0 \ , \qquad \eta(0) = x, \tag{7.2}$$

until time $T = \inf\{t > 0 : \eta(t) \notin D\}$. When h is Lipschitz, the standard Picard-Lindelöf theorem implies that if $x \in D$, then (7.2) has a unique solution up until time T (and T is itself uniquely determined). The reader can visually follow the flow lines in Figure 7.1.

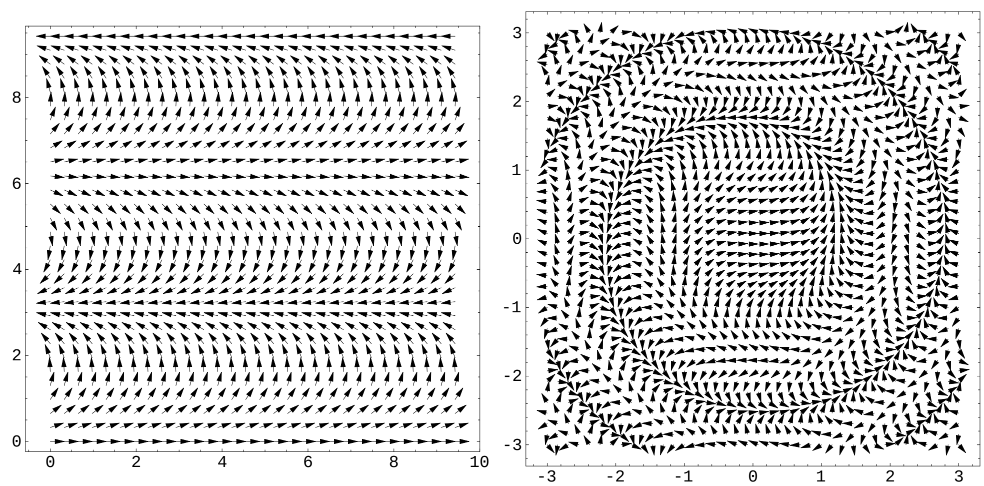

Figure 7.1: The complex vector flow e^{ih}: $h(x, y) = y$, $h(x, y) = x^2 + y^2$.

If h is continuous, then the time derivative $\eta'(t)$ moves continuously around the unit circle, and $h(\eta(t)) - h(\eta(0))$ describes the net amount of *winding* of η' around the circle between times 0 and t.

A major problem (addressed in depth in an imaginary geometry series [MS12a, MS12b, MS12c, MS13a]) is to make sense of these flow lines when h is a multiple of the Gaussian free field. We will give here just a short overview of the way these objects are constructed. Suppose that η is a smooth simple path in $\mathbb{H}$ beginning at the origin, with (forward) Loewner map $f_t = f_t^\eta$. We may assume that η starts out in the vertical direction, so that the winding number is $\pi/2$ for small times. Then when η and h are both smooth, the statement that η is a flow line of $e^{ih/\chi}$ is equivalent to the statement that for each x on $\eta\big((0, t)\big)$ we have

$$\chi \arg f_t'(z) \to -h(x) \tag{7.3}$$

as z approaches x from the left side of η and

$$\chi \arg f_t'(z) \to -h(x) + \chi\pi \tag{7.4}$$

as z approaches x from the right side of η (as Figure 7.2 illustrates). Recall that $\arg f_t'(z)$ — a priori determined only up to a multiple of 2π — is chosen to be continuous on $\mathbb{H} \setminus \eta([0,t])$ and 0 on $\mathbf{R}$. If $\chi = 0$, then (7.3) and (7.4) hold if and only if h is identically zero along the path η, i.e., η is a zero-height *contour line* of h.

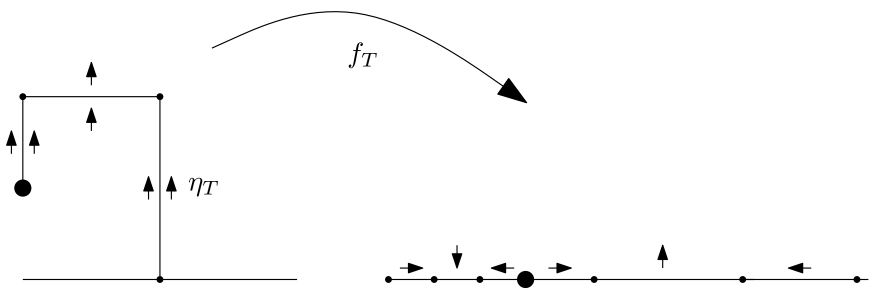

Figure 7.2: Winding number along η_T determines $\arg f_T'$, which is the amount a small arrow near η_T is rotated by f_T.

In [SS09c, SS13], it is shown that when one takes certain approximations h^ε of an instance h of the GFF that are piecewise linear on an ε-edge-length triangular mesh, then conditioned on a zero chordal contour line of h^ε there is in some $\varepsilon \to 0$ limiting sense a constant "height gap" between the expected heights immediately to one side of the contour line and those heights on the other. We might similarly expect that if one looked at the expectation of h^ε, given a chordal flow line η^ε of $e^{ih^\varepsilon/\chi}$, there would be a constant order limiting height gap between the two sides, see Figure 7.3.

This suggests the form of $\mathbf{H}_t$ given in Theorem 7.1, which comes from taking (7.3) and (7.4) and modifying the height gap between the two sides by adding a multiple of $\arg f_t$. (As in [SS10], the size of the height gap — and hence the coefficient of $\arg f_t$ in the definition of $\mathbf{H}_t$ — is determined by the requirement that $\mathbf{H}_t(z)$ be a martingale in t). Interestingly, the fact that winding may be ill-defined at a particular point on a fractal curve turns out to be immaterial. It is the harmonic extension of the boundary winding values (the $\arg f_t'$) that is needed to define $\mathbf{H}_t$, and this is defined even for non-smooth curves.

The time-reversal of a flow line of $e^{ih^\varepsilon/\chi}$ is a flow line of $e^{i(h^\varepsilon/\chi+\pi)}$, which at first glance appears to imply that there should not be a height gap between the two sides (since if the left side were consistently smaller for the forward path, then the right side would be consistently smaller for the reverse path). To counter this intuition, observe that, in the left diagram in Figure 7.1, the left-going infinite horizontal flow lines (at vertical heights of $k\pi$, k odd) are "stable" in that the flow line beginning at a generic point slightly off one of these lines will quickly converge to the line. The right-going horizontal flow lines (at heights $k\pi$, k even) are unstable. In a stable flow line, h appears to generally be larger to the right side of the flow line and smaller to the left side. It is reasonable to expect that a flow line of $e^{ih^\varepsilon/\chi}$ started from a generic point would be approximately

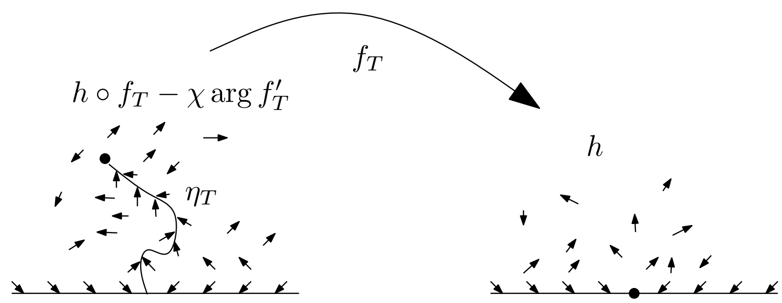

Figure 7.3: Forward coupling with arrows in $e^{ih/\chi}$ direction (sketch), illustrating the constant angle gap between the two sides of the curve η, constant angles along the positive and negative real axes, and random angles (not actually point-wise defined if h is the GFF) in $\mathbb{H} \setminus \eta$.

stable in that direction — and in particular would look qualitatively different from the time reversal of a flow line of $e^{i(h^\varepsilon/\chi+\pi)}$ started from a generic point.

7.2 Chordal SLE/GFF couplings

We will give explicit relationships between the Gaussian free field and both "forward" and "reverse" forms of SLE in Theorems 7.1 and 7.2 below. The forward couplings will be useful for the imaginary geometry discussion of this section and the reverse couplings will be useful later in the context of Liouville quantum gravity.

We will prove Theorem 7.1 in Section 7.3 using a series of calculations.[9]

The maps

$$f_t(z) := g_t(z) - W_t$$

satisfy

$$df_t(z) = \frac{2}{f_t(z)}dt - \sqrt{\kappa}dB_t,$$

and $f_t(\eta(t)) = 0$. Throughout this section, we will use f_t rather than g_t to describe the Loewner flow. If $\eta_T = \eta([0,T])$ is a segment of an SLE trace, denote by K_T the complement of the unbounded component of $\mathbb{H} \setminus \eta_T$. In the statements of Theorem 7.1

[9]The argument presented in Section 7.3 to prove Theorem 7.1, together with the relevant calculations, first appeared in lecture slides [She05]. Dubédat presented another short derivation of this statement within a long foundational paper [Dub09c]. More recent variants appear in [HBB10a, IK10], and in the imaginary geometry series [MS12a, MS12b, MS12c, MS13a]. Prior to these works, Kenyon and Schramm derived (but never published) a calculation relating SLE to the GFF in the case $\kappa = 8$. One could also have inferred the existence of such a relationship from the fact — due to Lawler, Schramm, and Werner — that SLE_8 is a continuum scaling limit of uniform spanning tree boundaries [LSW04c], and the fact — due to Kenyon — that the winding number "height functions" of uniform spanning trees have the GFF as a scaling limit [Ken00c, Ken01b, Ken08].

and Theorem 7.2 below and throughout the paper, we will discuss several kinds of random distributions on $\mathbb{H}$. To show that these objects are well defined as distributions on $\mathbb{H}$, we will make implicit use of some basic facts about distributions:

1. If h is a distribution on a domain D then its restriction to a subdomain is a distribution on that subdomain. (This follows by simply restricting the class of test functions to those supported on the subdomain.)

2. If h a distribution on a domain D and ϕ is a conformal map from D to a domain $\tilde{D}$ then $h \circ \phi^{-1}$ is a distribution on $\tilde{D}$. (Recall Footnote 14.)

3. An instance of the zero boundary GFF on a subdomain of D is also well defined as a distribution on all of D. (See Section 2.1 of [SS10].)

4. If h is an L^1 function on D, then h can be understood as a distribution on D defined by $(h, \rho) = \int_D \rho(z)h(z)dz$.

In the proof of Theorem 7.1 in Section 7.3, we will show that even though the function $\arg f_t'$ that appears in the theorem statement is a.s. unbounded, it can also a.s. be understood as a distribution on $\mathbb{H}$ (see the discussion after the theorem statement below).

Theorem 7.1. *Fix $\kappa \in (0, 4]$ and let η_T be the segment of SLE$_\kappa$ generated by the Loewner flow*

$$df_t(z) = \frac{2}{f_t(z)}dt - \sqrt{\kappa}dB_t, \qquad f_0(z) = z \tag{7.5}$$

up to a fixed time $T > 0$. Write

$$\begin{aligned} \mathbf{H}_0(z) &:= \frac{-2}{\sqrt{\kappa}} \arg z, \qquad\qquad \chi := \frac{2}{\sqrt{\kappa}} - \frac{\sqrt{\kappa}}{2}, \\ \mathbf{H}_t(z) &:= \mathbf{H}_0(f_t(z)) - \chi \arg f_t'(z). \end{aligned}$$

Here $\arg(f_t(z))$ (which is a priori defined only up to an additive multiple of 2π) is chosen to belong $(0, \pi)$ when $f_t(z) \in \mathbb{H}$; we similarly define $\arg f_t'(z)$ by requiring that (when t is fixed) it is continuous on $\mathbb{H} \setminus \eta_T$ and tends to 0 at ∞. Let $\widetilde{h}$ be an instance of the zero boundary GFF on $\mathbb{H}$, independent of B_t. Then the following two random distributions on $\mathbb{H}$ agree in law:[10]

$$\begin{aligned} h &:= \mathbf{H}_0 + \widetilde{h}. \\ h \circ f_T - \chi \arg f_T' &= \mathbf{H}_T + \widetilde{h} \circ f_T. \end{aligned}$$

The two distributions above also agree in law when $\kappa \in (4, 8)$ if we replace $\widetilde{h} \circ f_T$ with a GFF on $\mathbb{H} \setminus \eta([0, t])$ (which in this case means the sum of an independent zero boundary GFF on each component of $\mathbb{H} \setminus \eta([0, t])$) and take $\mathbf{H}_t(z) := \lim_{s \to \tau(z)_-} \mathbf{H}_s(z)$ if z is absorbed at time $\tau(z) \le t$.

[10]Note that f_T maps $\mathbb{H} \setminus K_T$ to $\mathbb{H}$, so $(\mathbb{H}, h)$ and $(\mathbb{H} \setminus K_T, h \circ f_T - \chi \arg f_T')$ describe equivalent AC surfaces by (8.4).

Alternative statement of Theorem 7.1: Using our coordinate change and AC surface definitions, we may state the theorem when $\kappa < 4$ somewhat more elegantly as follows: the law of the AC surface $(\mathbb{H}, h)$ is invariant under the operation of independently sampling f_T using a Brownian motion and (7.5), transforming the AC surface via the coordinate change f_T^{-1} (going from right to left in Figure 7.4 [11] — see also Figure 7.3) in the manner of (8.4), and erasing the path η_T (to obtain an AC surface parameterized by $\mathbb{H}$ instead of $\mathbb{H} \setminus \eta_T$). We discuss the geometric intuition behind the alternative statement in Section 7.1.

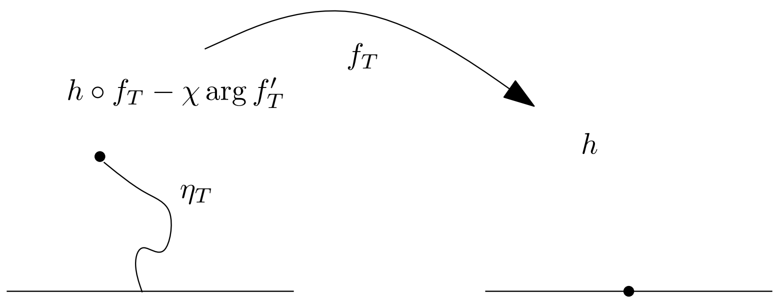

Figure 7.4: Forward coupling.

Note that, as a function, $\mathbf{H}_T$ is not defined on η_T itself. However, we will see in Section 7.3 that $\mathbf{H}_T$ is a.s. well defined as a distribution, independently of how we define it as a function on η_T itself. This will follow from the fact that, when $\kappa = 4$, this $\mathbf{H}_T$ is almost surely a bounded function off of η_T, and when $\kappa \neq 4$, the restriction of $\mathbf{H}_T$ to any compact subset of $\mathbb{H}$ is almost surely in L^p for each $p < \infty$. The fact that $\widetilde{h} \circ f_T$ is well defined as a distribution on $\mathbb{H}$ (not just as a distribution on $\mathbb{H} \setminus \eta_T$) follows from conformal invariance of the GFF, and the fact (mentioned above, proved in [SS10]) that a zero boundary GFF instance on a subdomain can be understood as a distribution on the larger domain.

Another standard approach for generating a segment η_T of an SLE curve is via the reverse Loewner flow, whose definition is recalled in the statement of the following theorem. (Note that if T is a fixed constant, then the law of the η_T generated by reverse Loewner evolution is the same as that generated by forward Loewner evolution; see Figures 7.4 and 7.5.)

Theorem 7.2. *Fix $\kappa > 0$ and let η_T be the segment of SLE$_\kappa$ generated by a reverse Loewner flow*

$$df_t(z) = \frac{-2}{f_t(z)} dt - \sqrt{\kappa} dB_t, \qquad f_0(z) = z \tag{7.6}$$

[11] All figures in this paper are sketches, not representative simulations.

up to a fixed time $T > 0$. Write

$$\begin{aligned} \mathbf{H}_0(z) &:= \frac{2}{\sqrt{\kappa}} \log|z|, \qquad\qquad Q := \frac{2}{\sqrt{\kappa}} + \frac{\sqrt{\kappa}}{2}, \\ \mathbf{H}_t(z) &:= \mathbf{H}_0(f_t(z)) + Q\log|f_t'(z)|, \end{aligned}$$

and let $\widetilde{h}$ be an instance of the free boundary GFF on $\mathbb{H}$, independent of B_t. Then the following two random distributions (modulo additive constants) on $\mathbb{H}$ agree in law:[12]

$$\begin{aligned} h &:= \mathbf{H}_0 + \widetilde{h}. \\ h \circ f_T + Q \log|f_T'| &= \mathbf{H}_T + \widetilde{h} \circ f_T. \end{aligned}$$

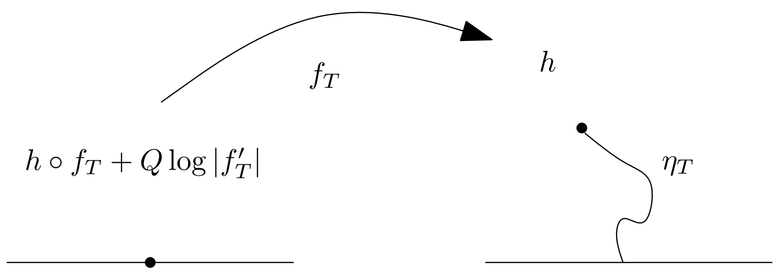

Figure 7.5: Reverse coupling.

Alternative statement of Theorem 7.2: A more elegant way to state the theorem is that the law of $(\mathbb{H}, h)$ is invariant under the operation of independently sampling f_T, cutting out K_T (equivalent to η_T when $\kappa \le 4$), and transforming via the coordinate change f_T^{-1} (going from right to left in Figure 7.5) in the manner of (8.3).

Both theorems give us an *alternate* way of sampling a distribution with the law of h — i.e., by first sampling the B_t process (which determines η_T), then sampling a (fixed or free boundary) GFF $\widetilde{h}$ and taking

$$h = \mathbf{H}_T + \widetilde{h} \circ f_T.$$

This two part sampling procedure produces a coupling of η_T with h. In the forward SLE setting of Theorem 7.1, it was shown in [Dub09c] that in any such coupling, η_T is almost surely equal to a particular path-valued *function* of h. (This was also done in [SS10] in the case $\kappa = 4$.) In other words, in such a coupling, h determines η_T almost surely. This is important for our geometric interpretations. Even though h is not defined pointwise as a function, we would like to geometrically interpret η as a level set of h (when $\kappa = 4$) or a flow line of $e^{ih/\chi}$ (when $\kappa < 4$), as we stated above

[12] Note that f_T maps $\mathbb{H}$ to $\mathbb{H} \setminus K_T$, so $(\mathbb{H}, h \circ f_T + Q\log|f_T'|)$ and $(\mathbb{H} \setminus K_T, h)$ describe equivalent quantum surfaces by (8.3). Indeed, $(\mathbb{H}, h \circ f_T + Q\log|f_T'|) = f_T^{-1}(\mathbb{H} \setminus K_T, h)$.

and will explain in more detail in Section 7.1. It is thus conceptually natural that such curves are uniquely determined by h (as they would be if h were a smooth function, see Section 7.1).

As mentioned earlier, this paper introduces and proves Theorem 7.2 while highlighting its similarity to Theorem 7.1. Indeed, it won't take us much more work to prove Theorems 7.1 and 7.2 together than it would take to prove one of the two theorems alone. It turns out that in both Figure 7.4 (which illustrates Theorem 7.1) and Figure 7.5 (which illustrates Theorem 7.2), the field illustrated on the left hand side of the figure (which agrees with h in law) actually determines η_T and the map f_T, at least when $\kappa < 4$. In the former context (Figure 7.4) this a major result due to Dubédat [Dub09c] (see also the exposition on this point in [MS12a]). It says that a certain "flow line" is a.s. uniquely determined by h. The statement in the latter context is a major result obtained in this paper, stated in Theorems 8.1 and 8.2. With some hard work, we will be able to show that the map f_T describes a conformal welding in which boundary arcs of equal quantum boundary length are "welded together". Once we have this, the fact that the boundary measure uniquely characterizes f_T will be obtained by applying a general "removability" result of Jones and Smirnov, as we will explain in Section 8.3.

7.3 Proofs of coupling theorems

This section will simultaneously prove Theorem 7.1 and Theorem 7.2. It is instructive to prove them together, and we will put the relevant calculations in tables, with those for the forward SLE coupling of Theorem 7.1 on the left side and those for the reverse SLE coupling of Theorem 7.2 on the right.

Now, using the language of stochastic differential equations and applying Itô/'s formula in the case $W_t = \sqrt{\kappa} B_t$, we compute the time derivatives of the four processes $f_t(z)$, $\log f_t(z)$, $f_t'(z)$, and $\log f_t'(z)$ in both forward and reverse SLE settings. Here $f_t'(z)$ denotes the spatial derivative $\frac{\partial}{\partial z} f_t$. (Similar calculations appear in [SS10] in the case $\kappa = 4$.)

FORWARD FLOW SLE	**REVERSE FLOW SLE**
$df_t(z) = \frac{2}{f_t(z)} dt - \sqrt{\kappa} dB_t$	$df_t(z) = \frac{-2}{f_t(z)} dt - \sqrt{\kappa} dB_t$
$d \log f_t(z) = \frac{(4-\kappa)}{2 f_t(z)^2} dt - \frac{\sqrt{\kappa}}{f_t(z)} dB_t$	$d \log f_t(z) = \frac{-(4+\kappa)}{2 f_t(z)^2} dt - \frac{\sqrt{\kappa}}{f_t(z)} dB_t$
$df_t'(z) = \frac{-2 f_t'(z)}{f_t(z)^2} dt$	$df_t'(z) = \frac{2 f_t'(z)}{f_t(z)^2} dt$
$d \log f_t'(z) = \frac{-2}{f_t(z)^2} dt$	$d \log f_t'(z) = \frac{2}{f_t(z)^2} dt$

We next define the martingales $\mathbf{H}_t$ in both settings and compute their stochastic derivatives. The purpose of the stochastic calculus below is to show that the quantities $(\mathbf{H}_t, \rho)$ are continuous local martingales (the fact that they are martingales will become apparent later) and to explicitly computing their quadratic variations, so that they can be understood as Brownian motions subject to an explicit time change. Ultimately, we will use the properties of these Brownian motions to establish couplings between SLE and the Gaussian free field.

Note that while the two columns have differed only in signs until now, the definitions of $\mathbf{H}_t$ below will diverge in that one involves the imaginary and one the real part of $\mathbf{H}_t^*$. We will write $\gamma := \sqrt{\min(\kappa, 16/\kappa)} \in [0, 2]$.

FORWARD FLOW SLE	**REVERSE FLOW SLE**
$\chi := \frac{2}{\sqrt{\kappa}} - \frac{\sqrt{\kappa}}{2}$	$Q := \frac{2}{\sqrt{\kappa}} + \frac{\sqrt{\kappa}}{2} = \frac{2}{\gamma} + \frac{\gamma}{2}$
$\mathbf{H}_t^*(z) := \frac{-2}{\sqrt{\kappa}} \log f_t(z) - \chi \log f_t'(z)$	$\mathbf{H}_t^*(z) := \frac{2}{\sqrt{\kappa}} \log f_t(z) + Q \log f_t'(z)$
$d\mathbf{H}_t^*(z) = \frac{2}{f_t(z)} dB_t$	$d\mathbf{H}_t^*(z) = \frac{-2}{f_t(z)} dB_t$
$\mathbf{H}_t(z) := \operatorname{Im} \mathbf{H}_t^*(z)$	$\mathbf{H}_t(z) := \operatorname{Re} \mathbf{H}_t^*(z)$
$d\mathbf{H}_t(z) = \operatorname{Im} \frac{2}{f_t(z)} dB_t$	$d\mathbf{H}_t(z) = \operatorname{Re} \frac{-2}{f_t(z)} dB_t$

Before continuing with the calculation, we make several remarks.

Remark 7.3. The form of $d\mathbf{H}_t(z)$ in the forward case is significant. At time $t = 0$, the function $-2\mathrm{Im}\,(f_t(z)^{-1})$ is simply $-2\mathrm{Im}\,(z^{-1})$. This is a positive harmonic function whose level sets are circles in $\mathbb{H}$ that are tangent to $\mathbb{R}$ at the origin. It is a multiple of the so-called Poisson kernel, and it is a derivative of the Green's function $G(y, z) = G^{\mathbb{H}_0}(y, z) = \log \left| \frac{y - \bar{z}}{y - z} \right|$ in the following sense:

$$[\frac{\partial}{\partial s} G(is, z)]_{s=0} = \frac{\partial}{\partial s} \left| \frac{z + is}{z - is} \right|_{s=0} = \operatorname{Re} \frac{2iz}{|z^2|} = 2\mathrm{Im}\,(z^{-1}).$$

Intuitively, the value of $-2\mathrm{Im}\,(f_t(z)^{-1})$ represents the harmonic measure of the tip of $\eta_t := \eta([0, t])$ as seen from the point z. Roughly speaking, as one makes observations of the GFF at points near the tip of η_t, the conditional expectation of h goes up or down by multiples of this function.

Remark 7.4. Also, in the forward case, $\mathbf{H}_0$ is the harmonic function on $\mathbb{H}$ with boundary conditions $-2\pi/\sqrt{\kappa}$ on the negative real axis and 0 on the positive real axis. We could have (for sake of symmetry) added a constant to $\mathbf{H}_0$ (and general $\mathbf{H}_t$) so that $\mathbf{H}_0$ is equal to $-\lambda$ on the negative real axis and λ on the positive real axis, where $\lambda := \pi/\sqrt{\kappa}$. Observe that when $\kappa = 4$, we have $\chi = 0$ and hence each $\mathbf{H}_t$ would be the harmonic function on $\mathbb{H}\backslash\eta_t$ with boundary conditions $-\lambda$ on the left side of the tip of η_t and λ on the right side. In this case, the $\lambda = \pi/2$ is the same (up to a $\sqrt{2\pi}$ factor stemming from a different choice of normalization for the GFF) as the value $\lambda = \sqrt{\pi/8}$ obtained in [SS10].

Remark 7.5. In the reverse case, the expression for $d\mathbf{H}_t$ has $\mathrm{Re}\,\frac{-2}{f_t(z)}$ in place of $\mathrm{Im}\,\frac{2}{f_t(z)}$. Intuitively, at time zero, when one observes what f_t looks like for small t, one learns something about the *difference* between h just to the left of 0 and h just to the right of 0. (It is this difference that determines the ratio of the ν_h densities to the left and to the right of zero, which is what determines how the zipping-up should behave in the short term.) The conditional expectation of h thus changes by a small multiple of $\mathrm{Re}\,\frac{2}{f_t(z)}$, which is negative on one side of the imaginary axis and positive on the other side. Unlike $\mathrm{Im}\,\frac{2}{f_t(z)}$, the function $\mathrm{Re}\,\frac{2}{f_t(z)}$ is non-zero on $\mathbf{R}$.

We use $\langle X_t, Y_t\rangle$ to denote cross variation between processes X_t and Y_t up to time t, so that $\langle X_t, X_t\rangle$ represents the quadratic variation of the process X_t up to time t. (The cross variation $\langle X_t, Y_t\rangle$ is also often written as $\langle X, Y\rangle_t$.) In both forward and reverse flow settings, $\mathbf{H}_t(z)$ is a continuous local martingale for each fixed z and is thus a Brownian motion under the quadratic variation parameterization, which we can give explicitly:

FORWARD FLOW SLE	**REVERSE FLOW SLE**
$C_t(z) := \log \mathrm{Im}\, f_t(z) - \mathrm{Re}\, \log f_t'(z)$	$C_t(z) := -\log \mathrm{Im}\, f_t(z) - \mathrm{Re}\, \log f_t'(z)$
$d\langle \mathbf{H}_t(z), \mathbf{H}_t(z)\rangle = -dC_t(z)$	$d\langle \mathbf{H}_t(z), \mathbf{H}_t(z)\rangle = -dC_t(z)$

If z is a point in a simply connected domain D, and ϕ conformally maps the unit disc to D, with $\phi(0) = z$, then we refer to the quantity $|\phi'(0)|$ as the *conformal radius* of D viewed from z. If, in the above definition of conformal radius, we replaced the unit disc with $\mathbb{H}$ and 0 with i, this would only change the definition by an additive constant. Thus, in the forward flow case, $C_t(z)$ is (up to an additive constant) the log of the conformal radius of $\mathbb{H}\backslash\eta([0,t])$ viewed from z. In both cases $\mathbf{H}_t(z)$ is a Brownian motion when parameterized by the time parameter $-C_t(z)$ (which is increasing as a function of t). The fact that $d\langle \mathbf{H}_t(z), \mathbf{H}_t(z)\rangle = -dC_t(z)$ may be computed directly via Itô/'s formula but it is also easy to see by taking $y \to z$ in the formulas for $\langle \mathbf{H}_t(y), \mathbf{H}_t(z) \rangle$ and $-dG_t(y,z)$ that we will give below.

We will now show that weighted averages of $\mathbf{H}_t$ over multiple points in $\mathbb{H}$ are also continuous local martingales (and hence Brownian motions when properly parameterized). The calculation will make use of the function $G(y,z)$, which we take to be the zero-boundary Green's function $G^{\mathbb{H}_0}(y,z)$ on $\mathbb{H}$ in the forward case and the free-boundary Green's function $G^{\mathbb{H}_F}(y,z)$ in the reverse case.

Now write $G_t(y,z) = G(f_t(y), f_t(z))$ in the reverse case. In the forward case, write $G_t(y,z) = G(f_t(y), f_t(z))$ when y and z are both in the infinite component of $\mathbb{H} \backslash \eta_t$ — otherwise, let $G_t(y,z)$ be the limiting value of $G_s(y,z)$ as s approaches the first time at which one of y or z ceases to be in this infinite component. The reader may check that for fixed y and z, this limit exists almost surely when $4 < \kappa < 8$: it is equal to zero when y and z are in different connected components of $\mathbb{H} \backslash \eta_t$, and when y and z lie in the same component, it is simply the Green's function of y and z on this bounded domain. Now we let ρ be a smooth compactly supported function on $\mathbb{H}$ (which we will assume has mean zero in the reverse case) and do some more calculations.

FORWARD FLOW SLE	**REVERSE FLOW SLE**
$G(y,z) := \log\lvert y-\bar{z}\rvert - \log\lvert y-z\rvert$	$G(y,z) := -\log\lvert y-z\rvert - \log\lvert y-\bar{z}\rvert$
$G_t(y,z) := G(f_t(y), f_t(z))$	$G_t(y,z) := G(f_t(y), f_t(z))$
$dG_t(y,z) = -\mathrm{Im}\,\frac{2}{f_t(y)}\mathrm{Im}\,\frac{2}{f_t(z)}dt$	$dG_t(y,z) = -\mathrm{Re}\,\frac{2}{f_t(y)}\mathrm{Re}\,\frac{2}{f_t(z)}dt$
$d\langle \mathbf{H}_t(y), \mathbf{H}_t(z)\rangle = -dG_t(y,z)$	$d\langle \mathbf{H}_t(y), \mathbf{H}_t(z)\rangle = -dG_t(y,z)$
$E_t(\rho) := \int_{\mathbb{H}} \rho(y) G_t(y,z) \rho(z) dy dz$	$E_t(\rho) := \int_{\mathbb{H}} \rho(y) G_t(y,z) \rho(z) dy dz$
$d\langle (\mathbf{H}_t,\rho), (\mathbf{H}_t,\rho)\rangle = -dE_t(\rho)$	$d\langle (\mathbf{H}_t,\rho), (\mathbf{H}_t,\rho)\rangle = -dE_t(\rho)$

Each of the equations above comes from a straightforward Itô/ calculation. To explain their derivation, we begin by expanding the dG_t computation in the forward case (the reverse case is similar):

$$dG_t(x,y) \quad = \quad -d\,\mathrm{Re}\,\log[f_t(x) - f_t(y)] + d\,\mathrm{Re}\,\log[f_t(x) - \overline{f_t(y)}]$$

$$
\begin{aligned}
&= -2\,\mathrm{Re}\,\frac{f_t(x)^{-1}-f_t(y)^{-1}}{f_t(x)-f_t(y)}\,dt + \\
&\qquad 2\,\mathrm{Re}\,\frac{f_t(x)^{-1}-\overline{f_t(y)^{-1}}}{f_t(x)-\overline{f_t(y)}}\,dt \\
&= 2\,\mathrm{Re}\left(f_t(x)^{-1}f_t(y)^{-1}\right)dt - 2\,\mathrm{Re}\left(f_t(x)^{-1}\left(\overline{f_t(y)}\right)^{-1}\right)dt \\
&= 2\,\mathrm{Re}\left(i\,f_t(x)^{-1}\,\mathrm{Im}\,[2f_t(y)^{-1}]\right)dt \\
&= -\mathrm{Im}\,\frac{2}{f_t(x)}\,\mathrm{Im}\,\frac{2}{f_t(y)}\,dt\,.
\end{aligned}
$$

The fact that $d\langle \mathbf{H}_t(y), \mathbf{H}_t(z)\rangle = -dG_t(y,z)$ is then immediate from our calculation of $d\mathbf{H}_t$.

The fact that $d\langle (\mathbf{H}_t,\rho),(\mathbf{H}_t,\rho)\rangle = -dE_t(\rho)$ is essentially a Fubini calculation but it requires some justification. First, we claim that the $(\mathbf{H}_t,\rho)$ are continuous martingales. We begin by considering $\mathbf{H}_t(z)$ for a fixed z in the support of ρ. We have shown above that the quantity $\mathbf{H}_t(z)$ is a Brownian motion under a certain parameterization. In the reverse case, the Loewner evolution gives that $|\frac{\partial}{\partial t}C_t(z)|$ is uniformly bounded above for z in the support of ρ and for all times t. (Note that $\mathrm{Im}\, f_t(z)$ is strictly increasing in t.) This immediately implies that $\mathbf{H}_t(z)$ is a martingale (not merely a local martingale) because for each z and t, $\mathbf{H}_t(z)$ represents the value of a Brownian motion stopped at a random time that is strictly less than a constant times t. The fact that the expectation of $\mathbf{H}_t(z)$ — given the filtration up to time $s<t$ — is $\mathbf{H}_s(z)$ is then immediate from the optional stopping theorem.

In the forward case, one obtains something similar by noting that the law of the conformal radius r of z in $\mathbb{H}\setminus\eta([0,t])$ has a power law decay as $r\to 0$ — i.e., the probability that $-C_t(z)>c$ decays exponentially in c, and is in fact bounded by an exponentially decaying function that is independent of z, for z in the support of ρ. (A precise description of the law of the conformal radius at time infinity appears as the main construction in [SSW09].) This implies that $\mathbf{H}_t(z)$ is a Brownian motion stopped at a time whose law decays exponentially (uniformly over z in the support of ρ) which is again enough to apply the optional stopping theorem and conclude that $\mathbf{H}_t(z)$ is martingale. In both cases, we obtain that for any t, the probability distribution function for $|\mathbf{H}_t(z)|$ decays exponentially fast, uniformly for z in the support of ρ. In both cases, we also see that (for any fixed t), $\mathbf{H}_t(z)$ is an L^1 function of z and the probability space, which allows us to use Fubini's theorem and conclude that the $(\mathbf{H}_t,\rho)$ are martingales.

Let L^p_{loc} denote the set of ψ for which the integral of $|\psi|^p$ over every compact subset of $\mathbb{H}$ is finite. The exponential decay above implies that $\mathbf{H}_t$ is almost surely in L^1_{loc}, since the expected integral of $|\mathbf{H}_t|$ over any compact set is finite. (Note that we can define $\mathbf{H}_t$ arbitrarily on the measure zero set $\eta([0,t])$ without affecting the definition of $\mathbf{H}_t$ as an element of $L^1_{\mathrm{loc}}(\mathbb{H})$.) In fact, since $\mathbb{E}|\mathbf{H}_t(z)|^p$ is bounded uniformly for z in a compact set, it follows that $\mathbf{H}_t$ is almost surely in $L^p_{\mathrm{loc}}(\mathbb{H})$ for any $p\in(1,\infty)$. The fact that $\mathbf{H}_t$ is almost surely in L^1_{loc} also implies that it can be understood as a random distribution on $\mathbb{H}$.

Moreover,

$$\sup_{s\in[0,t]} |\mathbf{H}_s(z)| \tag{7.7}$$

also has, by Doob's inequality, a law that decays exponentially, uniformly in z. Thus (7.7) also belongs a.s. to $L^p_{\mathrm{loc}}(\mathbb{H})$ for any $p \in (1,\infty)$. From this and the a.s. continuity of SLE it follows that $(\mathbf{H}_t, \rho)$ is a.s. continuous in t. (This continuity is obvious in the reverse case; in the forward case, it is also obvious if one replaces ρ by ρ_ε, which we define to be zero on an ε neighborhood of η and ρ elsewhere. The fact that (7.7) belongs to $L^p_{\mathrm{loc}}(\mathbb{H})$ implies that the $(\mathbf{H}_t, \rho_\varepsilon)$ converge to $(\mathbf{H}_t, \rho)$ uniformly, for almost all η, and a uniform limit of continuous functions is continuous.)

Now we can show $d\langle(\mathbf{H}_t,\rho),(\mathbf{H}_t,\rho)\rangle = -dE_t(\rho)$, as noted in [SS10], either via a stochastic Fubini's theorem (see e.g. [Pro90, §IV.4]) or by using the following simpler approach proposed in private communication by Jason Miller.

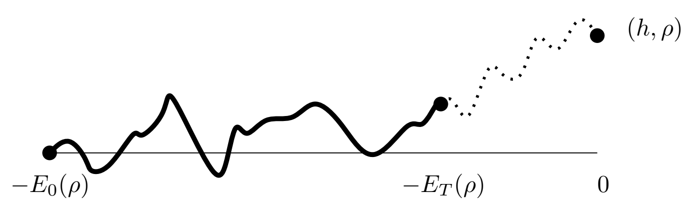

Figure 7.6: The pair $\big(-E_t(\rho), (\mathfrak{h}_t,\rho)\big)$ traces the graph of a Brownian motion (solid curve) as t ranges from 0 to T. Conditioned on this, the difference between (h,ρ) and $(\mathfrak{h}_T,\rho)$ is a centered Gaussian of variance $E_T(\rho)$. Choosing (h,ρ) to be $(\mathfrak{h}_T,\rho)$ plus a Gaussian of this variance is equivalent to continuing the Brownian motion parameterized by $-E_t(\rho)$ time (solid curve) all the way to time zero (dotted curve) and letting (h,ρ) be its value at time zero.

First note that $\langle(\mathbf{H}_t,\rho_1),(\mathbf{H}_t,\rho_2)\rangle$ is characterized by the fact that

$$(\mathbf{H}_t,\rho_1)(\mathbf{H}_t,\rho_2) - \langle(\mathbf{H}_t,\rho_1),(\mathbf{H}_t,\rho_2)\rangle$$

is a local martingale. Thus it suffices for us to show that

$$(\mathbf{H}_t,\rho_1)(\mathbf{H}_t,\rho_2) + \int \rho_1(x)\,\rho_2(y)\,G_t(x,y)\,dx\,dy \tag{7.8}$$

is a martingale. We know from the above calculations that

$$\mathbf{H}_t(x)\mathbf{H}_t(y) + G_t(x,y)$$

is a martingale for fixed x and y in $\mathbb{H}$. Since $G_t(x,y)$ is non-increasing and the $\mathbf{H}_t(z)$ have laws that decay exponentially, uniformly in z, we can use Fubini's theorem to conclude that (7.8) is a martingale. Thus we have that $(\mathbf{H}_t, \rho)$ is a Brownian motion when parameterized by time $-E_t(\rho)$. To complete the proofs of Theorems 7.1 and 7.2, recall that in the theorem statements $\widetilde{h}$ denotes an instance of the free boundary GFF on $\mathbb{H}$, and note that since each $(\mathbf{H}_T + \widetilde{h} \circ f_T, \rho)$ is a sum of a standard Brownian motion stopped at time $E_0(\rho) - E_T(\rho)$ and a conditionally independent Gaussian of variance $E_T(\rho)$, it has the same law as a Gaussian of variance $E_0(\rho)$ and mean $(\mathbf{H}_0, \rho)$. (See Figure 7.6.) For future reference, we note that in the reverse flow case one may integrate the expression for $d\mathbf{H}_t(z)$ above to find (using the stochastic Fubini's theorem) that

$$d(\mathbf{H}_t, \rho) = \big(-2\mathrm{Re}\,(f_t)^{-1}, \rho\big) dB_t. \tag{7.9}$$

Remark 7.6. The statement of Theorem 7.1 excluded the case $\kappa \geq 8$, since SLE_κ is space-filling in that case and $\mathbf{H}_t$ cannot be defined as a function almost everywhere. Nonetheless, we may still define $(\mathbf{H}_t, \rho)$ to be the solution to the stochastic differential equation $d(\mathbf{H}_t, \rho) = \big(-2\mathrm{Im}\,(f_t)^{-1}, \rho\big) dB_t$. In this case, the calculations above again yield that $d\langle(\mathbf{H}_t, \rho), (\mathbf{H}_t, \rho)\rangle = -dE_t(\rho)$, which as before implies that $(\mathbf{H}_T + \widetilde{h} \circ f_T, \rho)$ and $(\mathbf{H}_0 + \widetilde{h}, \rho)$ agree in law for each ρ, just as in the $\kappa < 8$ case, which yields a $\kappa \geq 8$ analog of Theorem 7.1. (Figure 7.6 still makes sense then $\kappa \geq 8$.)

It will be useful for later purposes to note that (at least in the reverse SLE case) the graph in Figure 7.6 actually uniquely determines (and is uniquely determined by) the process $W_t = \sqrt{\kappa} B_t$ almost surely, see Figure 7.7. This is a special case of a much more general theorem about stochastic processes (see Chapter IX, Theorem 2.1 of [RY99a] — it suffices that $(\mathfrak{h}_t, \rho)$ satisfies an SDE in W_t with a diffusive coefficient that remains strictly bounded away from zero and infinity, at least as long as we stop at any time strictly before $t = \infty$). This means that the evolution of η can be described by the Brownian motion in Figure 7.6, as well as by the Brownian motion B_t. Context will determine which description is more convenient to work with.

7.4 Flow lines starting from the boundary

7.4.1 GFF flow line overview

Given an instance h of the Gaussian free field (GFF), constants $\chi > 0$ and $\theta \in [0, 2\pi)$, and an initial point z, is there always a canonical way to define the flow lines of the complex vector field $e^{i(h/\chi+\theta)}$, i.e., solutions to the ODE

$$\eta'(t) = e^{i(h(\eta(t))/\chi+\theta)} \quad \text{for} \quad t > 0, \tag{7.10}$$

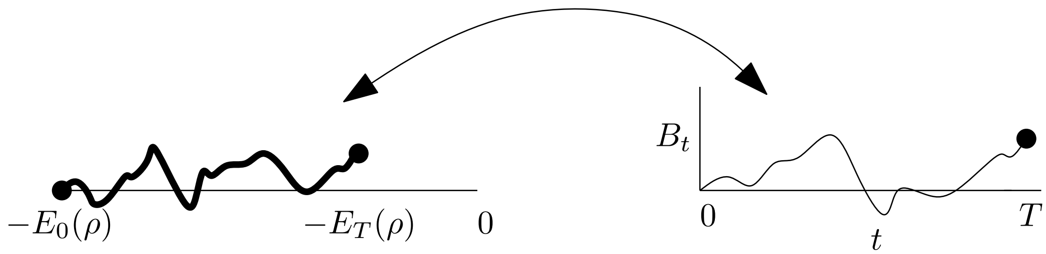

Figure 7.7: The graph traced by $\big(-E_t(\rho), (\mathfrak{h}_t, \rho)\big)$ as t ranges from 0 to T (left) and the graph traced by (t, B_t) (right), where $W_t = \sqrt{\kappa} B_t$. The left graph uniquely determines the right graph, and vice versa, almost surely. Each has the law of a standard Brownian motion (up to a stopping time).

beginning at z? The answer would obviously be yes if h were a smooth function (Figure 7.8), but it is less obvious for an instance of the GFF, which is a distribution (a.k.a. a generalized function), not a function (Figures 7.9–7.12).

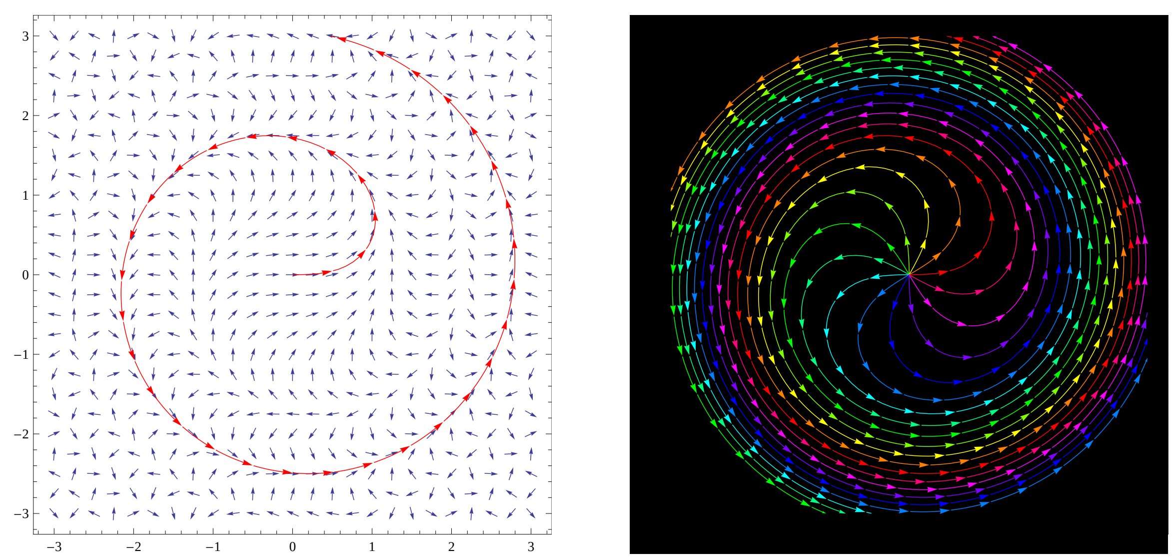

(a) The vector field $e^{ih(z)}$ where $h(z) = |z|^2$, together with a flow line started at zero.

(b) Flow lines of $e^{i(h(z)+\theta)}$ for 12 uniformly spaced θ values.

Figure 7.8

Several works in recent years have addressed special cases and variants of this question [She, Dub09d, MS10, SS13, HBB10b, IK13, She15] and have shown that in certain circumstances there is a sense in which the paths are well-defined (and uniquely determined) by h, and are variants of the Schramm-Loewner evolution (SLE). In this chapter, we will focus on the case that z is point on the boundary of the domain where h is defined and establish a more general set of results.

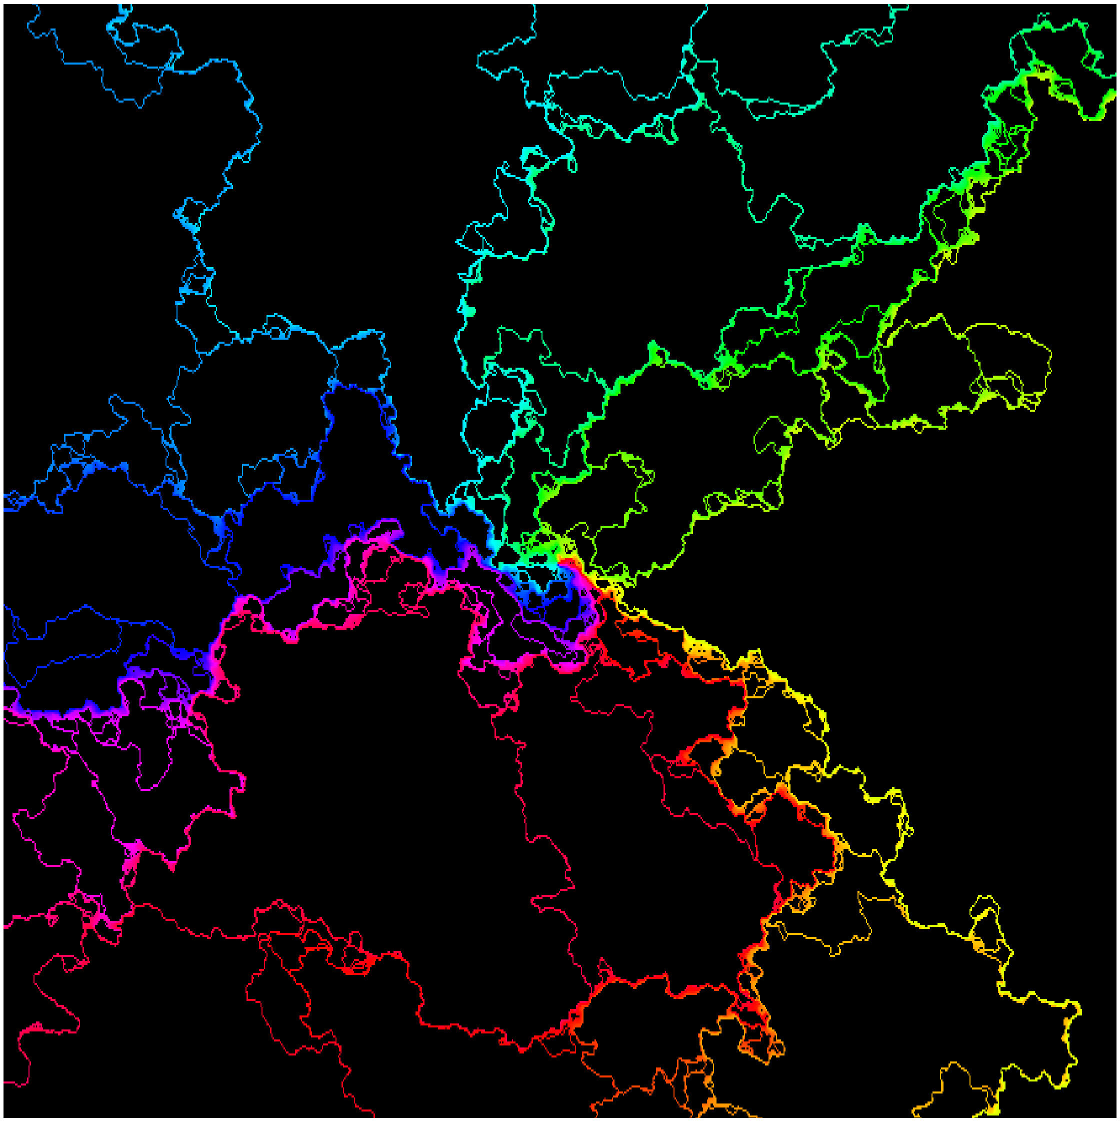

Figure 7.9: Numerically generated flow lines, started at a common point, of $e^{i(h/\chi+\theta)}$ where h is the projection of a GFF onto the space of functions piecewise linear on the triangles of a 300×300 grid; $\kappa = 4/3$ and $\chi = 2/\sqrt{\kappa} - \sqrt{\kappa}/2 = \sqrt{4/3}$. Different colors indicate different values of $\theta \in [0, 2\pi)$. We expect but do not prove that if one considers increasingly fine meshes (and the same instance of the GFF) the corresponding paths converge to limiting continuous paths (an analogous result was proven for $\kappa = 4$ [SS09a, SS13]).

We will fix $\chi > 0$ and interpret the paths corresponding to different θ values as "rays of a random geometry" angled in different directions and show that different paths started

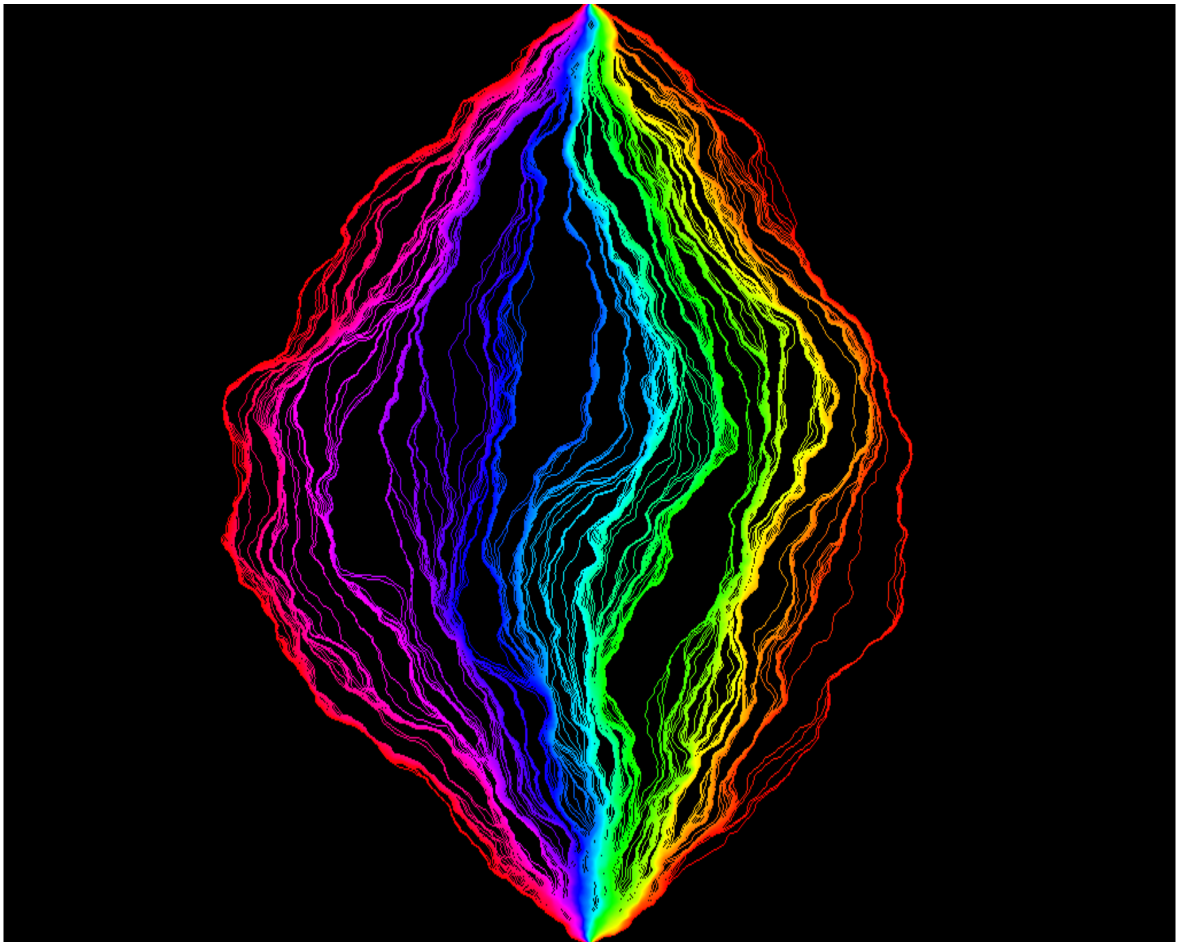

Figure 7.10: Numerically generated flow lines, started at $-i$ of $e^{i(h/\chi+\theta)}$ where h is the projection of a GFF on $[-1,1]^2$ onto the space of functions piecewise linear on the triangles of a 300×300 grid; $\kappa = 1/8$. Different colors indicate different values of $\theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$. The boundary data for h is chosen so that the central ("north-going") curve shown should approximate an $\mathrm{SLE}_{1/8}$ process.

at a common point never cross one another. Note that these are the rays of ordinary Euclidean geometry when h is a constant.

Theorem 7.7 and Theorem 7.8 establish the fact that the flow lines are well-defined and uniquely determined by h almost surely. Theorem 7.7 is the same as a theorem proved in [Dub09d, MS12a]. For convenience, we have restated it here. This theorem establishes the existence of a coupling between h and the path with certain properties. Theorem 7.8 then shows that in this coupling, the path is almost surely determined by the field. Theorem 7.8 is an extension of a result in [Dub09d]. Unlike the result in [Dub09d], our Theorem 7.8 applies to paths that interact with the domain boundaries in non-trivial ways, and this requires new tools.

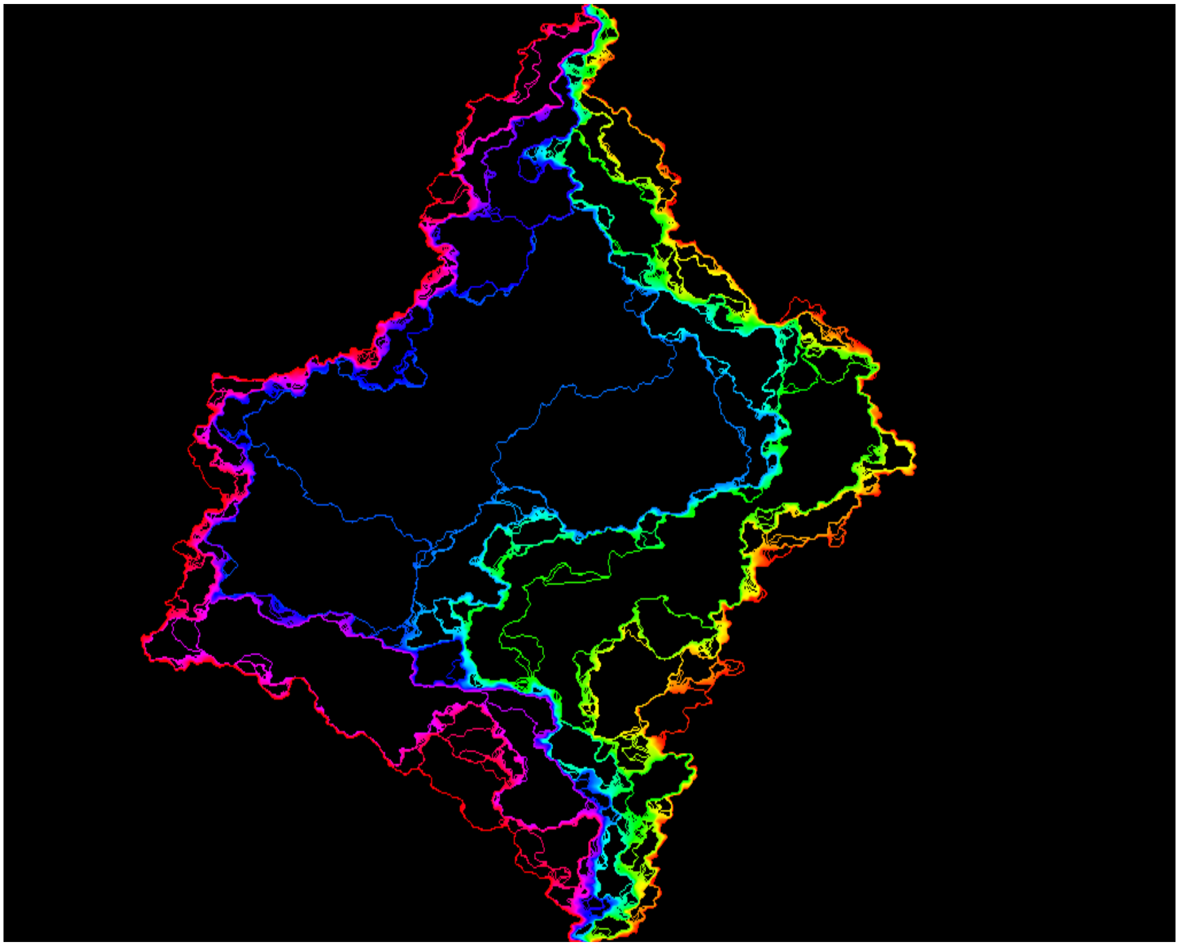

Figure 7.11: Numerically generated flow lines, started at $-i$ of $e^{i(h/\chi+\theta)}$ where h is the projection of a GFF on $[-1,1]^2$ onto the space of functions piecewise linear on the triangles of a 300×300 grid; $\kappa = 1$. Different colors indicate different values of $\theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$. The boundary data for h is chosen so that the central ("north-going") curve shown should approximate an SLE_1 process.

The boundary-intersecting case of Theorem 7.8 and other ideas will then be used to describe the way that distinct flow lines interact with one another when they intersect (see Figure 7.28). We show that the flow lines started at the same point, corresponding to different θ values, may bounce off one another (depending on the angle difference) but almost surely do not cross one another, that flow lines started at distinct points with the same angle can "merge" with each other, and that flow lines started at distinct points with distinct angles almost surely cross at most once. We give a complete description of the *conditional* law of h given a finite collection of (possibly intersecting) flow lines. (The conditional law of h given multiple flow line segments is discussed in [Dub09d], but the results there only apply to *non-intersecting* segments. Extending

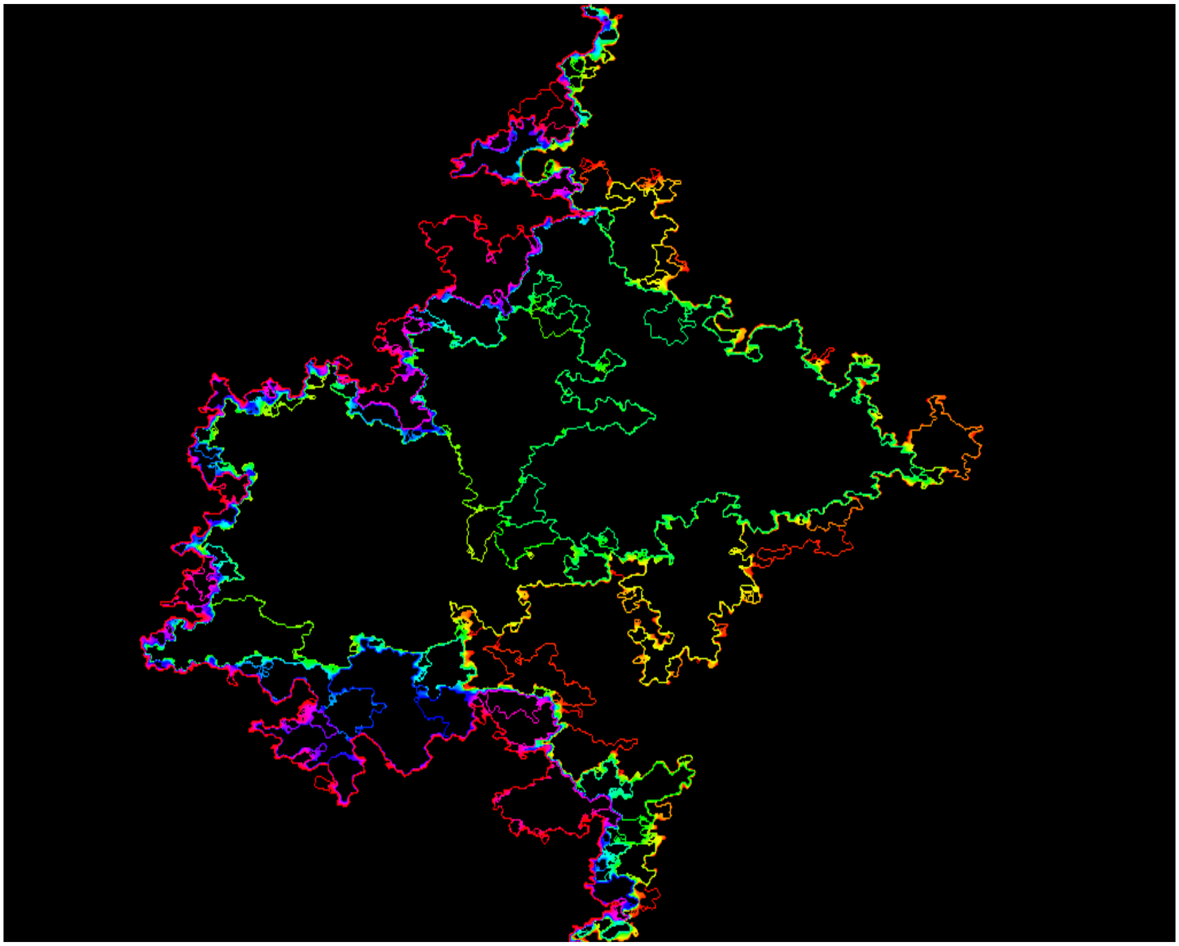

Figure 7.12: Numerically generated flow lines, started at $-i$ of $e^{i(h/\chi+\theta)}$ where h is the projection of a GFF on $[-1,1]^2$ onto the space of functions piecewise linear on the triangles of a 300×300 grid; $\kappa = 2$. Different colors indicate different values of $\theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$. The boundary data for h is chosen so that the central ("north-going") curve shown should approximate an SLE_2 process.

these results requires, among other things, ruling out pathological behavior of the conditional expectation of the field — given the paths — near points where the paths intersect.) These are some of the fundamental results one needs to begin to understand (continuum analogs of) Figures 7.9–7.12, 7.14, and 7.15.

As mentioned above, we also establish some new results in classical SLE theory. For example, the flow line technology enables us to show in Theorem 7.9 that the so-called $\mathrm{SLE}_\kappa(\underline{\rho})$ curves are a.s. continuous even when they hit the boundary. Rohde and Schramm proved that ordinary SLE_κ on a Jordan domain is continuous when $\kappa \neq 8$ [RS05a]; the continuity of SLE_8 was proved by Lawler, Schramm, and Werner in [LSW04a] (extensions to more general domains are proved in [GRS08]) but their techniques do not readily apply to boundary intersecting $\mathrm{SLE}_\kappa(\underline{\rho})$, and the lack of a proof for $\mathrm{SLE}_\kappa(\underline{\rho})$ has been a persistent gap in the literature. Another approach to proving Theorem 7.9 in the case of a single force point, based on extremal length arguments, has been proposed (though not yet published) by Kemppainen, Schramm, and Sheffield [KSS].

The random geometry point of view also gives us a new way of understanding other random objects with conformal symmetries. For example, we will use the flow-line geometry to construct so-called *counterflow lines*, which are forms of $\mathrm{SLE}_{16/\kappa}$ ($\kappa \in (0,4)$) that arise as the "light cones" of points accessible by certain angle-restricted SLE_κ trajectories. To use another metaphor, we say that a point y is "downstream" from another point x if it can be reached from x by an angle-varying flow line whose angles lie in some allowed range; the counterflow line is a curve that traces through all the points that are downstream from a given boundary point x, but it traces them in an "upstream" (or "counterflow") direction. This is the content of Theorem 7.10, which is stated somewhat informally. (More precise and general statements of Theorem 7.10 appear in [MS12a].) In contrast to what happens when h is smooth, the light cones thus constructed are not simply connected sets when $\kappa \in (2,4)$. It also turns out that one can reach all points in the light cone by considering paths that alternate between the two extreme angles. See Figures 7.20–7.25 for discrete simulations of light cones generated in this manner (the two extreme angles differ by π; see also Figure 7.26 for an explanation of the fact that a path with angle changes of size π does not just retrace itself).

It is also shown [MS12a] that, for any $\kappa \in (0,4)$, the closure of the union of all the flow lines starting at a given point z with angles in a countable, dense set (as depicted in Figures 7.9–7.12) almost surely has Lebesgue measure zero. (It is easy to see that the resulting object does not depend on the choice of countable, dense set.) Put somewhat fancifully, this states that when a person holds a gun at a point z in the imaginary geometry, there are certain other points (in fact, almost all points) that the gun cannot hit no matter how carefully it is aimed. (One might guess this to be the case from the amount of black space in Figures 7.9–7.12, 7.23.) Generally, random imaginary geometry yields many natural ways of coupling and understanding multiple SLEs on the same domain, as well as SLE variants on non-simply-connected domains.

The flow lines constructed here also turn out to be relevant to the study of Liouville quantum gravity. For example, we plan to show in a subsequent joint work with Duplantier that the rays in Figures 7.9–7.12 arise when gluing together independent Liouville quantum gravity surfaces via the conformal welding procedure presented in [She15]. The tools developed here are essential for that program.

7.4.2 Background and setting

Let $D \subseteq \mathbf{C}$ be a domain with harmonically non-trivial boundary (i.e., a Brownian motion started at a point $z \in D$ almost surely hits ∂D) and let $C_0^\infty(D)$ denote the space of compactly supported C^∞ functions on D. For $f, g \in C_0^\infty(D)$, let

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

denote the Dirichlet inner product of f and g where dx is the Lebesgue measure on D. Let $H(D)$ be the Hilbert space closure of $C_0^\infty(D)$ under $(\cdot,\cdot)_\nabla$. The continuum Gaussian free field h (with zero boundary conditions) is the so-called standard Gaussian on $H(D)$. It is given formally as a random linear combination

$$h = \sum_n \alpha_n \phi_n, \tag{7.11}$$

where (α_n) are i.i.d. $N(0,1)$ and (ϕ_n) is an orthonormal basis of $H(D)$.

The GFF is a two-dimensional-time analog of Brownian motion. Just as many random walk models have Brownian motion as a scaling limit, many random (real or integer valued) functions on two dimensional lattices have the GFF as a scaling limit [BAD96, NS97, Ken01a, RV07, Mil10a].

The GFF can be used to generate various kinds of random geometric structures, including both Liouville quantum gravity and the imaginary geometry discussed here [She15]. Roughly speaking, the former corresponds to replacing a Euclidean metric $dx^2 + dy^2$ with $e^{\gamma h}(dx^2 + dy^2)$ (where $\gamma \in (0,2)$ is a fixed constant and h is the Gaussian free field). The latter is closely related, and corresponds to considering $e^{ih/\chi}$, for a fixed constant $\chi > 0$. Informally, as discussed above, the "rays" of the imaginary geometry are flow lines of the complex vector field $e^{i(h/\chi+\theta)}$, i.e., solutions to the ODE (7.10), for given values of $\eta(0)$ and θ.

A brief overview of imaginary geometry (as defined for general functions h) appears in [She15], where the rays are interpreted as geodesics of a variant of the Levi-Civita connection associated with Liouville quantum gravity. One can interpret the e^{ih} direction as "north" and the $e^{i(h+\pi/2)}$ direction as "west", etc. Then h determines a way of assigning a set of compass directions to every point in the domain, and a ray is determined by an initial point and a direction. (We have not described a Riemannian geometry, since we have not introduced a notion of length or area.) When h is constant, the rays correspond to rays in ordinary Euclidean geometry. For more general continuous h, one can still show that when three rays form a triangle, the sum of the angles is always π [She15].

Throughout the rest of this article, when we say that η is a flow line of h it is to be interpreted that η is a flow line of the vector field $e^{ih/\chi}$; both h and χ will be clear from the context. In particular, the statement that η is a flow line of h with angle θ is equivalent to the statement that η is a flow line of $h + \theta\chi$.

We next remark that if h is a smooth function on D, η a flow line of $e^{ih/\chi}$, and $\psi\colon \widetilde{D} \to D$ a conformal transformation, then by the chain rule, $\psi^{-1} \circ \eta$ is a flow line of $h \circ \psi - \chi \arg \psi'$ (note that a reparameterization of a flow line remains a flow line), as in Figure 7.13. With this in mind, we define an **imaginary surface** to be an equivalence class of pairs (D, h) under the equivalence relation

$$(D, h) \to (\psi^{-1}(D), h \circ \psi - \chi \arg \psi') = (\widetilde{D}, \widetilde{h}). \tag{7.12}$$

Note that this makes sense even for h which are not necessarily smooth. We interpret ψ as a (conformal) *coordinate change* of the imaginary surface. In what follows, we will generally take D to be the upper half plane, but one can map the flow lines defined there to other domains using (7.12).

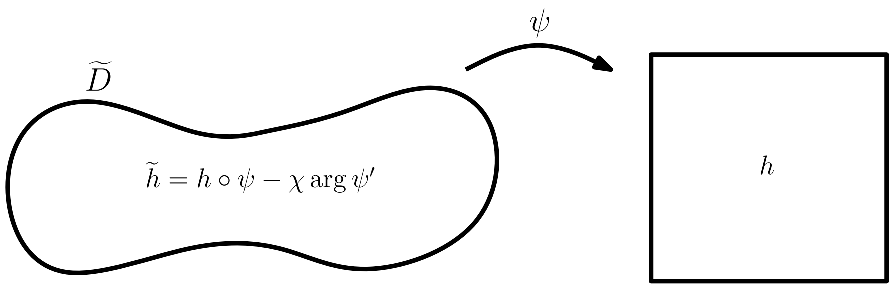

Figure 7.13: The set of flow lines in $\widetilde{D}$ will be the pullback via a conformal map ψ of the set of flow lines in D provided h is transformed to a new function $\widetilde{h}$ in the manner shown.

When h is an instance of the GFF on a planar domain, the ODE (7.10) is not well-defined, since h is a distribution-valued random variable and not a continuous function. One could try to approximate one of these rays by replacing the h in (7.10) by its projection onto a space of continuous functions — for example, the space of functions that are piecewise linear on the triangles of some very fine lattice. This approach (and a range of θ values) was used to generate the rays in Figures 7.9–7.12, 7.14, 7.15, 7.20-7.25, and 7.28. We expect that these rays will converge to limiting path-valued functions of h as the mesh size gets finer. This has not been proved, but an analogous result has been shown for level sets of h [SS09a, SS13].

As we discussed briefly in Section 7.4.1, it turns out that it is possible to make sense of these flow lines and level sets directly in the continuum, without the discretizations mentioned above. The construction is rather interesting. One begins by constructing explicit couplings of h with variants of the Schramm-Loewner evolution and showing that these couplings have certain properties. Namely, if one conditions on part of the curve, then the conditional law of h is that of a GFF in the complement of

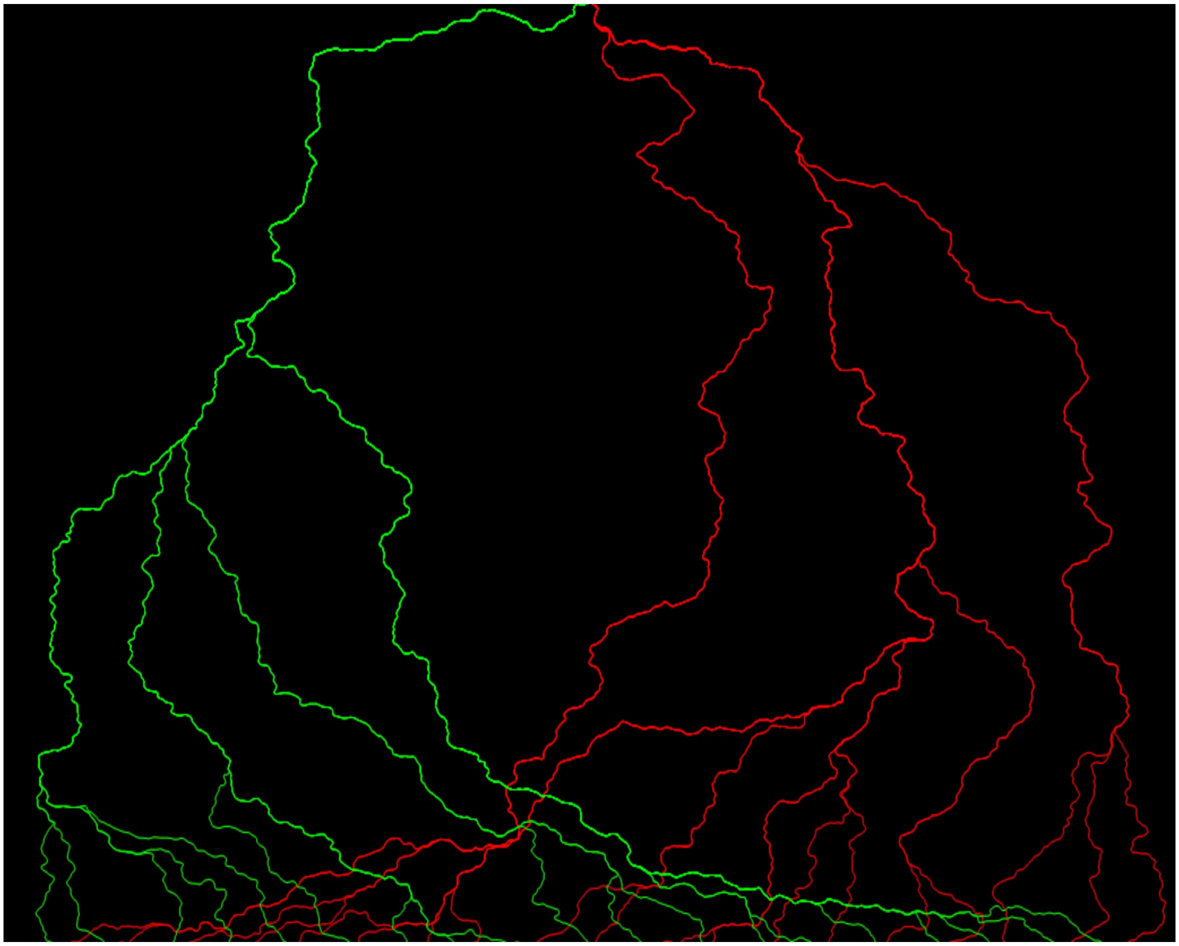

Figure 7.14: Numerically generated flow lines, started at evenly spaced points on $[-1-i, 1-i]$ of $e^{ih/\chi}$ where h is the projection of a GFF on $[-1,1]^2$ onto the space of functions piecewise linear on the triangles of a 300×300 grid; $\kappa = 1/2$. The angle of the green lines is $\frac{\pi}{4}$ and the angle of the red lines is $-\frac{\pi}{4}$. Flow lines of the same color appear to merge, but the red and green lines always cross at right angles. The boundary data of h was given by taking 0 boundary conditions on $\mathbf{H}$ and then applying the transformation rule (7.12) with a conformal map $\psi\colon \mathbf{H} \to [-1,1]^2$ where $\psi(0) = -i$ and $\psi(\infty) = i$.

the curve with certain boundary conditions. Examples of these couplings appear in [She, Dub09d, SS13, She15] as well as variants in [MS10, HBB10b, IK13]. This step is carried out in some generality in [Dub09d, She15]. A second step (implemented only for some particular boundary value choices in [Dub09d] and [SS13]) is to show that in such a coupling, the path is actually completely *determined* by h, and thus can be interpreted as a path-valued function of h.

Before we describe the rigorous construction of the flow lines of $e^{i(h/\chi+\theta)}$, let us offer

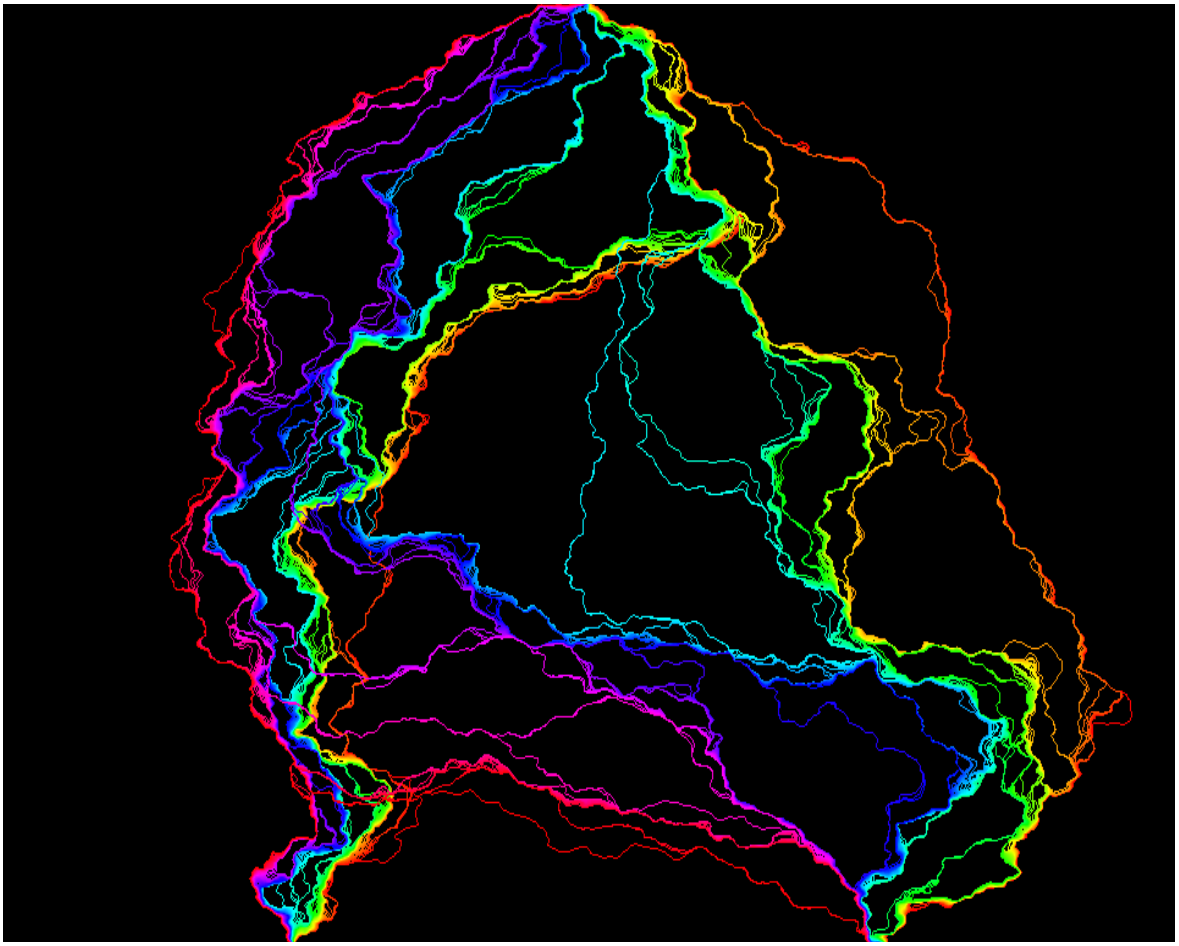

Figure 7.15: Numerically generated flow lines, started at $-1/2-i$ and $1/2-i$ of $e^{i(h/\chi+\theta)}$ with angles evenly spaced in $[-\frac{\pi}{4},\frac{\pi}{4}]$ where h is the projection of a GFF on $[-1,1]^2$ onto the space of functions piecewise linear on the triangles of a 300×300 grid; $\kappa=1/2$. Flow lines of different colors appear to cross at most once and flow lines of the same color appear to merge. The boundary data for h is the same as in Figure 7.14.

some geometric intuition. Suppose that h is a continuous function and consider a flow line of the complex vector field $e^{ih/\chi}$ in $\mathbf{H}$ beginning at 0. That is, $\eta\colon[0,\infty)\to\mathbf{H}$ is a solution to the ODE

$$\eta'(t)=e^{ih(\eta(t))/\chi}\quad\text{for}\quad t>0,\quad \eta(0)=0. \tag{7.13}$$

Note that $\|\eta'(t)\|=1$. Thus, the time derivative $\eta'(t)$ moves continuously around the unit circle $\mathbf{S}^1$ and $\big(h(\eta(t))-h(\eta(0))\big)/\chi$ describes the net amount of *winding* of η' around $\mathbf{S}^1$ between times 0 and t. Let g_t be the Loewner map of η. That is, for each t, g_t is the unique conformal transformation of the unbounded connected component of $\mathbf{H}\setminus\eta([0,t])$ to $\mathbf{H}$ that looks like the identity at infinity: $\lim_{z\to\infty}|g_t(z)-z|=0$.

Loewner's theorem says that g_t is a solution to the equation

$$\partial_t g_t(z) = \frac{2}{g_t(z) - W_t}, \quad g_0(z) = z, \tag{7.14}$$

where $W_t = g_t(\eta(t))$, provided η is parameterized appropriately. It will be convenient for us to consider the centered Loewner flow $f_t = g_t - W_t$ of η in place of g_t. The reason for this particular choice is that f_t maps the tip of $\eta|_{[0,t]}$ to 0. Note that

$$df_t(z) = \frac{2}{f_t(z)}dt - dW_t. \tag{7.15}$$

We may assume that η starts out in the vertical direction, so that the winding number is approximately $\pi/2$ as $t \downarrow 0$. We claim that the statement that $\eta|_{[0,t]}$ is a flow line of $e^{i(h/\chi+\pi/2)}$ is equivalent to the statement that for each x on $\eta((0,t))$, we have

$$\chi \arg f_t'(z) \to -h(x) - \chi\pi/2 \tag{7.16}$$

as z approaches from the left side of η and

$$\chi \arg f_t'(z) \to -h(x) + \chi\pi/2 \tag{7.17}$$

as z approaches from the right side of η. To see this, first note that both $s \mapsto f_t^{-1}(s)|_{(0,s_+)}$ and $s \mapsto f_t^{-1}(-s)|_{(s_-,0)}$ are parameterizations of $\eta|_{[0,t]}$ where s_-, s_+ are the two images of 0 under f_t. One then checks (7.16) (and (7.17) analogously) by using that $\eta(s) = f_t^{-1}(\phi(s))$ for $\phi\colon (0,\infty) \to (0,\infty)$ a smooth decreasing function and applying (7.13). If $\chi = 0$, then (7.16) and (7.17) hold if and only if h is identically zero along the path, which is to say that η is a zero-height contour line of h. Roughly speaking, the flow lines of $e^{i(h/\chi+\pi/2)}$ and level sets of h are characterized by (7.16) and (7.17), though it turns out that the "angle gap" must be modified by a constant factor in order to account for the roughness of the field. In a sense there is a constant "height gap" between the two sides of the path, analogous to what was shown for level lines of the GFF in [SS09a, SS13]. The law of the flow line of h starting at 0 is determined by the boundary conditions of h. It turns out that if the boundary conditions of h are those shown in Figure 7.16, then the flow line starting at 0 is an SLE_κ process (with $\underline{\rho} \equiv 0$). Namely, one has $-\lambda$ and λ along the left and right sides of the axis and along the path one has $-\lambda'$ plus the winding on the left and λ' plus the winding on the right, for the particular values of λ and λ' described in the caption. Each time the path makes a quarter turn to the left, heights go up by $\frac{\pi}{2}\chi$. Each time the path makes a quarter turn to the right, heights go down by $\frac{\pi}{2}\chi$.

7.4.3 Coupling of paths with the GFF

We now extend the GFF coupling results to more general setting. For convenience and concreteness, we take D to be the upper half-plane $\mathbf{H}$. Couplings for other simply

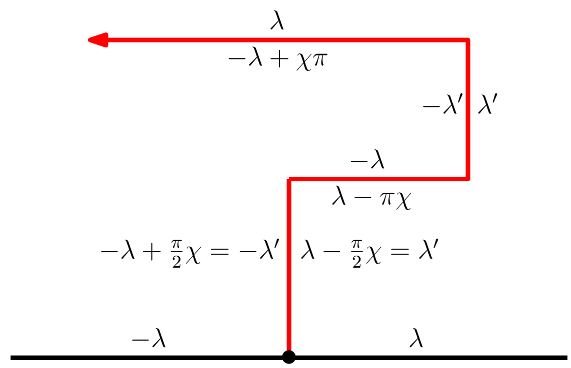

Figure 7.16: Fix $\kappa \in (0,4)$ and set $\lambda = \lambda(\kappa) = \frac{\pi}{\sqrt{\kappa}}$. Write $\lambda' = \lambda(16/\kappa) = \frac{\pi\sqrt{\kappa}}{4}$. Conditioned on a flow line, the heights of the field are given by (a constant plus) χ times the winding of the path minus λ' on the left side and χ times the winding plus λ' on the right side. For a fractal curve, these heights are not pointwise defined (though their harmonic extension is well-defined). The figure illustrates these heights for a piecewise linear curve. In Figure 7.17, we will describe a more compact notation for indicating the boundary heights in figures/.

connected domains are obtained using the change of variables described in Figure 7.13. Recall that SLE_κ is the random curve described by the centered Loewner flow (7.15) where $W_t = \sqrt{\kappa}B_t$ and B_t is a standard Brownian motion. More generally, an $\mathrm{SLE}_\kappa(\underline{\rho})$ process is a variant of SLE_κ in which one keeps track of multiple additional points, which we refer to as force points. Throughout the rest of the article, we will denote configurations of force points as follows. We suppose $\underline{x}^L = (x^{k,L} < \cdots < x^{1,L})$ where $x^{1,L} \leq 0$, and $\underline{x}^R = (x^{1,R} < \cdots < x^{\ell,R})$ where $x^{1,R} \geq 0$. The superscripts L, R stand for "left" and "right," respectively. If we do not wish to refer to the elements of $\underline{x}^L, \underline{x}^R$, we will denote such a configuration as $(\underline{x}^L; \underline{x}^R)$. Associated with each force point $x^{i,q}$, $q \in \{L, R\}$ is a weight $\rho^{i,q} \in \mathbf{R}$ and we will refer to the vector of weights as $\underline{\rho} = (\underline{\rho}^L; \underline{\rho}^R)$. An $\mathrm{SLE}_\kappa(\underline{\rho})$ process with force points $(\underline{x}^L; \underline{x}^R)$ corresponding to the weights $\underline{\rho}$ is the measure on continuously growing compact hulls K_t — compact subsets of $\overline{\mathbf{H}}$ so that $\mathbf{H} \setminus K_t$ is simply connected — such that the conformal maps $g_t \colon \mathbf{H} \setminus K_t \to \mathbf{H}$, normalized so that $\lim_{z\to\infty} |g_t(z) - z| = 0$, satisfy (7.15) with W_t replaced by the solution to the system of (integrated) SDEs

$$W_t = \sqrt{\kappa}B_t + \sum_{q\in\{L,R\}} \sum_i \int_0^t \frac{\rho^{i,q}}{W_s - V_s^{i,q}} ds, \tag{7.18}$$

$$V_t^{i,q} = \int_0^t \frac{2}{V_s^{i,q} - W_s} ds + x^{i,q}, \quad q \in \{L, R\}. \tag{7.19}$$

Additional discussion of both SLE_κ and $\mathrm{SLE}_\kappa(\underline{\rho})$ processes appears in [MS12a]. The general coupling statement below applies for all $\kappa > 0$. Theorem 7.7 below gives a

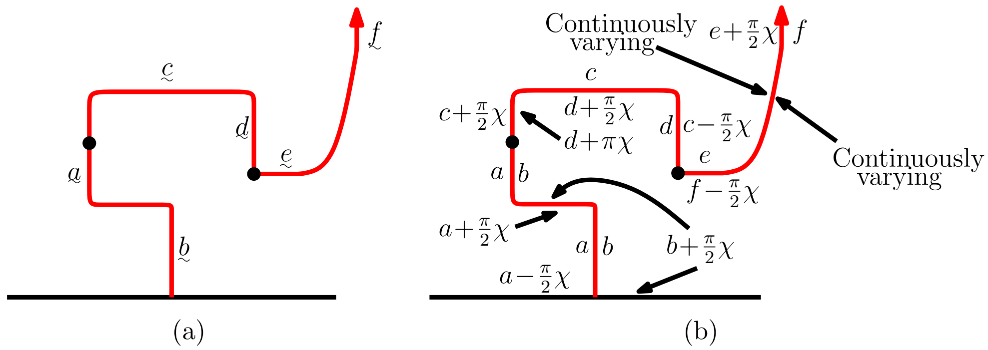

Figure 7.17: Throughout this article, we will need to consider Gaussian free fields whose boundary data changes with the winding of the boundary. In order to indicate this succinctly, we will often make use of the notation depicted on the left hand side. Specifically, we will delineate the boundary ∂D of a Jordan domain D with black dots. On each arc L of ∂D which lies between a pair of black dots, we will draw either a horizontal or vertical segment L_0 and label it with $\underset{\sim}{x}$ where $x \in \mathbf{R}$. This serves to indicate that the boundary data along L_0 is given by x as well as describe how the boundary data depends on the winding of L. Whenever L makes a quarter turn to the right, the height goes down by $\frac{\pi}{2}\chi$ and whenever L makes a quarter turn to the left, the height goes up by $\frac{\pi}{2}\chi$. More generally, if L makes a turn which is not necessarily at a right angle, the boundary data is given by χ times the winding of L relative to L_0. When we just write x next to a horizontal or vertical segment, we mean to indicate the boundary data at that segment and nowhere else. The right panel above has exactly the same meaning as the left panel, but in the former the boundary data is spelled out explicitly everywhere. Even when the curve has a fractal, non-smooth structure, the *harmonic extension* of the boundary values still makes sense, since one can transform the figure via the rule in Figure 7.13 to a half plane with piecewise constant boundary conditions. The notation above is simply a convenient way of describing the values of the constants. We will often include horizontal or vertical segments on curves in our figures/ (even if the whole curve is known to be fractal) so that we can label them this way.

general statement of the existence of the coupling. Essentially, the theorem states that if we sample a particular random curve on a domain D — and then sample a Gaussian free field on D minus that curve with certain boundary conditions — then the resulting field (interpreted as a distribution on all of D) has the law of a Gaussian free field on D with certain boundary conditions.

It is proved in [Dub09d] that Theorem 7.7 holds for any κ and $\underline{\rho}$ for which a solution to (7.18) exists (this can also be extended to a continuum of force points; this is done for a time-reversed version of SLE in [She15]). The special case of $\pm\lambda$ boundary conditions also appears in [She]. (See also [She15] for a more detailed version of the argument in

[She] with additional figures/ and explanation.)

The question of when (7.18) has a solution is not explicitly addressed in [Dub09d]. In [MS12a] the authors prove (adapting some results from [Dub09d, Theorem 6.4]) the existence of a unique solution to (7.18) up until the **continuation threshold** is hit — the first time t that $W_t = V_t^{j,q}$ where $\sum_{i=1}^{j} \rho^{i,q} \le -2$, for some $q \in \{L, R\}$.

All of our results will hold for $\mathrm{SLE}_\kappa(\underline{\rho})$ processes up until (and including) the continuation threshold. It turns out that the continuation threshold is infinite almost surely if and only if

$$\sum_{i=1}^{j} \rho^{i,L} > -2 \quad \text{for all} \quad 1 \le j \le k \quad \text{and} \quad \sum_{i=1}^{j} \rho^{i,R} > -2 \quad \text{for all} \quad 1 \le j \le \ell.$$

Theorem 7.7. *Fix $\kappa > 0$ and a vector of weights $(\underline{\rho}^L; \underline{\rho}^R)$. Let K_t be the hull at time t of the $\mathrm{SLE}_\kappa(\underline{\rho})$ process generated by the Loewner flow (7.15) where $(W, V^{i,q})$ solves (7.18), (7.19). Let $\mathfrak{h}_t^0$ be the function which is harmonic in $\mathbf{H}$ with boundary values*

$$\begin{aligned} -\lambda\left(1 + \sum_{i=0}^{j} \rho^{i,L}\right) \quad & \text{if} \quad s \in [V_t^{j+1,L}, V_t^{j,L}), \\ \lambda\left(1 + \sum_{i=0}^{j} \rho^{i,R}\right) \quad & \text{if} \quad s \in [V_t^{j,R}, V_t^{j+1,R}), \end{aligned}$$

where $\rho^{0,L} = \rho^{0,R} = 0$, $x^{0,L} = 0^-$, $x^{k+1,L} = -\infty$, $x^{0,R} = 0^+$, and $x^{\ell+1,R} = \infty$. (See Figure 7.18.) Let

$$\mathfrak{h}_t(z) = \mathfrak{h}_t^0(f_t(z)) - \chi \arg f_t'(z), \quad \chi = \frac{2}{\sqrt{\kappa}} - \frac{\sqrt{\kappa}}{2}.$$

Let $(\mathcal{F}_t)$ be the filtration generated by $(W, V^{i,q})$. There exists a coupling (K, h) where $\widetilde{h}$ is a zero boundary GFF on $\mathbf{H}$ and $h = \widetilde{h} + \mathfrak{h}_0$ such that the following is true. Suppose τ is any $\mathcal{F}_t$-stopping time which almost surely occurs before the continuation threshold is reached. Then K_τ is a local set for h and the conditional law of $h|_{\mathbf{H}\setminus K_\tau}$ given $\mathcal{F}_\tau$ is equal to the law of $\mathfrak{h}_\tau + \widetilde{h} \circ f_\tau$.

We will give a review of the theory of local sets [SS13] for the GFF in Section 3.3.

Notice that $\chi > 0$ when $\kappa \in (0, 4)$, $\chi < 0$ when $\kappa > 4$, and that $\chi(\kappa) = -\chi(\kappa')$ for $\kappa' = 16/\kappa$ (though throughout the rest of this article, whenever we write χ it will be assumed that $\kappa \in (0, 4)$). This means that in the coupling of Theorem 7.7, the conditional law of h given either an SLE_κ or an $\mathrm{SLE}_{\kappa'}$ curve transforms in the same way under a conformal map, up to a change of sign. Using this, we are able to construct

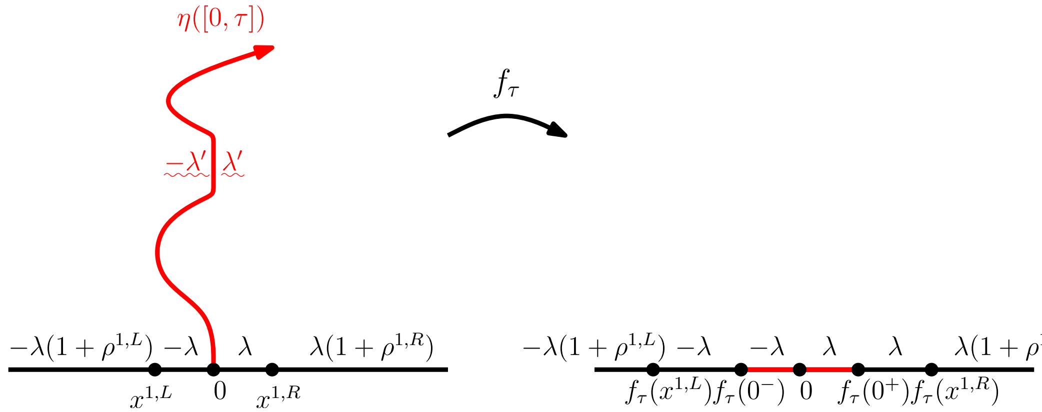

Figure 7.18: The function $\mathfrak{h}_\tau^0$ in Theorem 7.7 is the harmonic extension of the boundary values depicted in the right panel in the case that there are two boundary force points, one on each side of 0. The function $\mathfrak{h}_\tau = \mathfrak{h}_\tau^0 \circ f_\tau - \chi \arg f_\tau'$ in Theorem 7.7 is the harmonic extension of the boundary data specified in the left panel. (Recall the relationship between λ and λ' indicated in Figure 7.16.)

$\eta \sim \mathrm{SLE}_\kappa$, $\kappa \in (0,4)$, and $\eta' \sim \mathrm{SLE}_{\kappa'}$ curves within the same imaginary geometry (see Figure 7.19). We accomplish this by taking η to be coupled with h and η' to be coupled with $-h$, as in the statement of Theorem 7.7 (this is the reason we can always take $\chi > 0$).

Definition. When $\kappa \in (0,4)$, we will refer to an $\mathrm{SLE}_\kappa(\underline{\rho})$ curve (if it exists) coupled with a GFF h on $\mathbf{H}$ with boundary conditions as in Theorem 7.7 as a **flow line** of h. One can use the conformal coordinate change of Figure 7.13 to extend this definition to simply connected domains other than $\mathbf{H}$. To spell out this point explicitly, suppose that D is a simply connected domain homeomorphic to the disk, $x, y \in \partial D$ are distinct, and $\psi\colon D \to \mathbf{H}$ is a conformal transformation with $\psi(x) = 0$ and $\psi(y) = \infty$. Let us assume that we have fixed a branch of $\arg \psi'$ that is defined continuously on all of D. We assume further that $\underline{x}^L$ (resp. $\underline{x}^R$) consists of k (resp. ℓ) distinct marked prime ends in the clockwise (resp. counterclockwise) segment of ∂D (as defined by ψ) which are in clockwise (resp. counterclockwise) order. We take $x^{0,L} = x = x^{0,R} = x$ and $x^{k+1,L} = x^{\ell+1,R} = y$. We then suppose that h is a GFF on D with boundary conditions in the clockwise (resp. counterclockwise) segment of ∂D from $x^{j,L}$ to $x^{j+1,L}$ (resp. $x^{j,R}$ to $x^{j+1,R}$) given by $-\lambda\left(1+\sum_{i=0}^{j} \rho^{i,L}\right) - \chi \arg \psi'$ (resp. $\lambda\left(1+\sum_{i=0}^{j} \rho^{i,R}\right) - \chi \arg \psi'$). We refer to an $\mathrm{SLE}_\kappa(\underline{\rho})$ curve η (if it exists) from x to y on D, $\kappa \in (0,4)$, coupled with h as a **flow line** of h if the curve $\psi(\eta)$ in $\mathbf{H}$ is coupled as a flow line of the GFF $h \circ \psi^{-1} - \chi \arg(\psi^{-1})'$ on $\mathbf{H}$. (Recall (7.12) and Figure 7.13.)

Remark. Observe that in the discussion above, the choice of the branch of $\arg \psi'$ was important. Changing the branch chosen would in some sense correspond to adding a multiple of $2\pi\chi$ to either side of the $\mathrm{SLE}_\kappa(\underline{\rho})$ curve, and if one did this then (in order for the curve to remain a flow line) one would have to compensate by adding the same quantity to the boundary data. In some sense, changing the branch of $\arg\psi'$ is equivalent to adding a multiple of $2\pi\chi$ to the boundary data. If one wishes to be fully concrete, one can fix the branch of $\arg\psi'$ in an arbitrary way — say, so that $\arg\psi'(\psi^{-1}(i)) \in (-\pi,\pi]$ — and then assume that the boundary data is adjusted accordingly. In practice, when we discuss flow lines (in the half plane or elsewhere) we will usually specify boundary data using a figure and the notation explained in Figure 7.17 (or in Figure 7.18). This approach will avoid any "multiple of $2\pi\chi$" ambiguity and will make it completely clear exactly what the boundary data is along the curve. This remark also applies to the definition of counterflow line given below.

There are several examples of coordinate changes in the section on Dubédat's approach in [MS12a]. See also Figure 7.16 and Figure 7.17 for an illustration of how the boundary data for the GFF changes when applying (7.12).

The fact that $\mathrm{SLE}_\kappa(\underline{\rho})$ is generated by a continuous curve up until hitting the continuation threshold will be established for general ρ values in Theorem 7.9. It is not obvious from the coupling described in Theorem 7.7 that such paths are deterministic functions of h. That this is in fact the case is given in Theorem 7.8.

As mentioned earlier, we will sometimes use the phrase *flow line of angle θ* to denote the corresponding curve that one obtains when $\theta\chi$ is added to the boundary data (so that h is replaced by $h+\theta\chi$).

Definition. We will refer to an $\mathrm{SLE}_{\kappa'}(\underline{\rho})$ curve (if it exists), $\kappa' \in (4,\infty)$, coupled with a GFF $-h$ (note the sign change here; this accounts for the $\chi(\kappa)$ vs. $\chi(\kappa')$ issue discussed just above) as in Theorem 7.7 as a **counterflow line** of h. Again, one can use conformal maps to extend this definition to simply connected domains other than $\mathbf{H}$. Suppose that D is a non-trivial simply connected domain, $x,y \in \partial D$ are distinct, and $\psi\colon D \to \mathbf{H}$ is a conformal transformation with $\psi(x)=0$ and $\psi(y)=\infty$, and that a branch of $\arg\psi'$ has been fixed (as in the flow line definition above). We assume further that $\underline{x}^L$ (resp. $\underline{x}^R$) consists of k (resp. ℓ) distinct marked prime ends in the clockwise (resp. counterclockwise) segment of ∂D (as defined by ψ) which are in clockwise (resp. counterclockwise) order. We take $x^{0,L}=x=x^{0,R}=x$ and $x^{k+1,L}=x^{\ell+1,R}=y$. We then suppose that h is a GFF on D with boundary conditions in the clockwise (resp. counterclockwise) segment of ∂D from $x^{j,L}$ to $x^{j+1,L}$ (resp. $x^{j,R}$ to $x^{j+1,R}$) given by $\lambda'\left(1+\sum_{i=0}^j \rho^{i,L}\right) - \chi\arg\psi'$ (resp. $-\lambda'\left(1+\sum_{i=0}^j \rho^{i,R}\right) - \chi\arg\psi'$); here $\chi=\chi(\kappa)>0$. We refer to an $\mathrm{SLE}_{\kappa'}(\underline{\rho})$ curve η' (if it exists) from x to y on D, $\kappa' \in (4,\infty)$, coupled with h as a **counterflow line** of h if the curve $\psi(\eta')$ in $\mathbf{H}$ is coupled as a counterflow line of the GFF $h\circ\psi^{-1} - \chi\arg(\psi^{-1})'$ on $\mathbf{H}$; here $\chi=\chi(\kappa)>0$. (Recall (7.12) and Figure 7.13.)

Again, the fact that $\mathrm{SLE}_{\kappa'}(\underline{\rho})$ is generated by a continuous curve up until hitting the continuation threshold is established for general ρ values in Theorem 7.9.

As in the setting of flow lines, it is not obvious from the coupling described in Theorem 7.7 that such paths are deterministic functions of h. That this is in fact the case is given in Theorem 7.8. The reason for the terminology "counterflow line" is that, as briefly mentioned earlier, it will turn out that the set of the points hit by an $\mathrm{SLE}_{\kappa'}$ counterflow line can be interpreted as a "light cone" of points accessible by certain angle-restricted SLE_κ flow lines; the $\mathrm{SLE}_{\kappa'}$ passes through the points on each of these flow lines in the opposite ("counterflow") direction. We will provide some additional explanation near the statement of Theorem 7.10.

The correction $-\chi \arg f_t'$ which appears in the statement Theorem 7.7 has the interpretation of being the harmonic extension of χ times the **winding** of $\partial(\mathbf{H} \setminus \eta([0,\tau]))$. We will use the informal notation $\chi \cdot$ winding for this function throughout this article and employ a special notation to indicate this in figures/. See Figure 7.17 for further explanation of this point.

Similar couplings are constructed in [IK13] for the GFF with Neumann boundary data on part of the domain boundary, and [HBB10b] couples the GFF on an annulus with annulus SLE. Makarov and Smirnov extend the SLE_4 results of [She, SS13] to the setting of the massive GFF and a massive version of SLE in [MS10].

7.4.4 Main results

In the case that $\rho = 0$ and η is ordinary SLE, Dubédat showed in [Dub09d] that in the coupling of Theorem 7.7 the path is actually a.s. determined by the field. A $\kappa = 4$ analog of this statement was also shown in [SS13]. In this paper, we will extend these results to the more general setting of Theorem 7.7.

Theorem 7.8. *Suppose that h is a GFF on $\mathbf{H}$ and that $\eta \sim \mathrm{SLE}_\kappa(\underline{\rho})$. If (η, h) are coupled as in the statement of Theorem 7.7, then η is almost surely determined by h.*

The basic idea of our proof is as follows. First, we extend the argument of [Dub09d] for SLE_κ, $\kappa \in (0,4]$, to the case of $\eta \sim \mathrm{SLE}_\kappa(\underline{\rho})$ with $\underline{\rho} = (\rho^L; \rho^R)$ where ρ^L and ρ^R are real numbers satisfying $\rho^L \geq \frac{\kappa}{2} - 2$ and $\rho^R \geq 0$. This condition implies that η almost surely does not intersect $\partial \mathbf{H}$ after time 0 and allows us to apply the argument from [Dub09d] with relatively minor modifications. We then reduce the more general case that $\rho^L, \rho^R > -2$ to the former setting by studying the flow lines η_θ of $e^{i(h/\chi + \theta)}$ emanating from 0. In this case, these are also $\mathrm{SLE}_\kappa(\underline{\rho})$ curves with force points at 0^- and 0^+. We will prove that if $\theta_1 < 0 < \theta_2$, then η_{θ_1} almost surely lies to the right of η which in turn almost surely lies to the right of η_{θ_2}. We will next show that the conditional law of η given $\eta_{\theta_1}, \eta_{\theta_2}$ is an $\mathrm{SLE}_\kappa(\rho^L(\theta_1); \rho^R(\theta_2))$ process independently in each of the connected components of $\mathbf{H} \setminus (\eta_{\theta_1} \cup \eta_{\theta_2})$ which lie between η_{θ_1} and η_{θ_2}. By adjusting θ_1, θ_2, we can obtain any combination of $\rho^L(\theta_1), \rho^R(\theta_2) > -2$. We then extend this result to the setting of many force points by systematically studying the case with

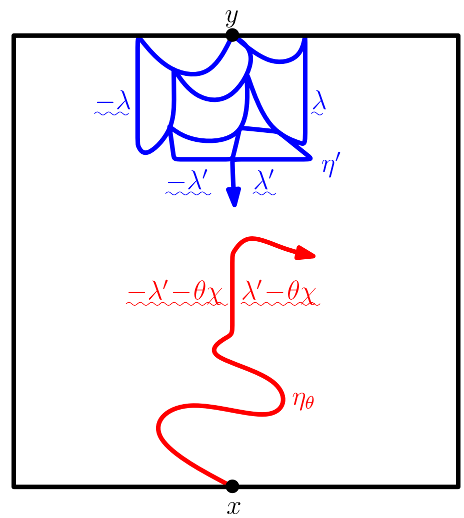

Figure 7.19: We can construct SLE_κ flow lines, $\kappa \in (0,4)$, and $\mathrm{SLE}_{\kappa'}$, $\kappa' = 16/\kappa$, counterflow lines within the same imaginary geometry. This is depicted above for a single counterflow line η' emanating from y and a flow line η_θ with angle θ starting from x. In this coupling, η_θ is coupled with $h + \theta\chi$ and η' is coupled with $-h$ as in Theorem 7.7. Also shown is the boundary data for h in $D \setminus (\eta'([0,\tau']) \cup \eta_\theta([0,\tau]))$ conditional on $\eta_\theta([0,\tau])$ and $\eta'([0,\tau'])$ where τ and τ' are stopping times for η_θ and η' respectively (we intentionally did not specify the boundary data of h on ∂D). Assume that η' is non-boundary filling. Then if $\theta = \frac{1}{\chi}(\lambda' - \lambda) = -\frac{\pi}{2}$ so that the boundary data on the right side of η_θ matches that on the right side of η', then η_θ will almost surely hit and then "merge" into the right boundary of η'. The analogous result holds if $\theta = \frac{1}{\chi}(\lambda - \lambda') = \frac{\pi}{2}$ so that the boundary data on the left side of η_θ matches that on the left side of η'. This fact is known as Duplantier duality (or SLE duality). More generally, if $\theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$ then η_θ is almost surely contained in η' but the union of the traces of η_θ as θ ranges over the entire interval $[-\frac{\pi}{2}, \frac{\pi}{2}]$ is almost surely a strict subset of the range of η'. We will show, however, that the range of η' can be constructed as a "light cone" of SLE_κ trajectories whose angle is allowed to vary in time but is restricted to $[-\frac{\pi}{2}, \frac{\pi}{2}]$.

two boundary force points which are both to the right of 0 and then employing the absolute continuity properties of the GFF combined with an induction argument. The idea for $\kappa > 4$ follows from a more elaborate variant of this general strategy.

By applying the same set of techniques used to prove Theorem 7.8, we also obtain the continuity of the $\mathrm{SLE}_\kappa(\underline{\rho})$ trace.

Theorem 7.9. *Suppose that $\kappa > 0$. If $\eta \sim \mathrm{SLE}_\kappa(\underline{\rho})$ on $\mathbf{H}$ from 0 to ∞ then η is almost surely a continuous path, up to and including the continuation threshold. On the event that the continuation threshold is not hit before η reaches ∞, we have a.s. that $\lim_{t\to\infty} |\eta(t)| = \infty$.*

The continuity of SLE_κ (with $\rho = 0$) was first proved by Rohde and Schramm in [RS05a]. By invoking the Girsanov theorem, one can deduce from [RS05a] that $\mathrm{SLE}_\kappa(\underline{\rho})$ processes are also continuous, but only up until just before the first time that a force point is absorbed. The main idea of the proof in [RS05a] is to control the moments of the derivatives of the reverse SLE_κ Loewner flow near the origin. These estimates involve martingales whose corresponding PDEs become complicated when working with $\mathrm{SLE}_\kappa(\underline{\rho})$ in place of usual SLE_κ. Our proof uses the Gaussian free field as a vehicle to construct couplings which allow us to circumvent these technicalities.

Another achievement of this paper will be to show how to jointly construct all of the flow lines emanating from a single boundary point. This turns out to give us a flow-line based construction of $\mathrm{SLE}_{16/\kappa}(\underline{\rho})$, $\kappa \in (0,4)$. That is, $\mathrm{SLE}_{16/\kappa}$ variants occur naturally within the same imaginary geometry as SLE_κ. Note that $16/\kappa$ assumes all possible values in $(4,\infty)$ as κ ranges over $(0,4)$. Imprecisely, we have that the set of all points reachable by proceeding from the origin in a possibly varying but always "northerly" direction (the so-called "light cone") along SLE_κ flow lines is a form of $\mathrm{SLE}_{16/\kappa}$ for $\kappa \in (0,4)$ generated in the reverse direction (see Figure 7.19).

Theorem 7.10 below is stated somewhat informally. More precise statements appear in [MS12a].

Theorem 7.10. *Suppose that h is a GFF on $\mathbf{H}$ with piecewise constant boundary data. Let η' be the counterflow line of h starting at ∞ targeted at 0. Assume that the continuation threshold for η' is almost surely not hit. Then the range of η' is almost surely equal to the set of points accessible by SLE_κ trajectories of h starting at 0 whose angles are restricted to be in $[-\frac{\pi}{2}, \frac{\pi}{2}]$ but may change in time. Let η_L be the flow line of h with angle $\frac{\pi}{2}$ starting at 0 and η_R the flow line of h with angle $-\frac{\pi}{2}$. It is almost surely the case that if η' is nowhere boundary filling (i.e., $\eta' \cap \mathbf{R}$ has empty interior), then η_L and η_R do not hit the continuation threshold before reaching ∞ and are the left and right boundaries of η'.*

A similar statement holds on the event that η' is boundary filling on one or more segments of $\mathbf{R}$. In this case, η_L and η_R hit their continuation thresholds before reaching ∞, but they can be extended to describe the entire left and right boundaries of η'. (See additional discussion in [MS12a].)

The light cone construction of $\mathrm{SLE}_{16/\kappa}$ processes described in the statement of Theorem 7.10 includes what is known as *Duplantier duality* or SLE *duality* — that the outer boundary of an $\mathrm{SLE}_{16/\kappa}$ process is equal in law to a kind of SLE_κ process. This was proven in certain cases by Zhan [Zha08a, Zha10] and Dubédat [Dub09a]. Theorem 7.10 provides a more general version of this duality. It shows that the law of the right boundary of any $\mathrm{SLE}_{16/\kappa}(\underline{\rho}')$ process η' from ∞ to 0 in $\mathbf{H}$ is given by the flow line of angle $-\frac{\pi}{2}$ in the same imaginary geometry. Analogously, the law of the left boundary of any $\mathrm{SLE}_{16/\kappa}(\underline{\rho}')$ process η' is given by the flow line of angle $\frac{\pi}{2}$ in the same imaginary geometry. We can also compute the conditional law of η' given either η_L or η_R. These

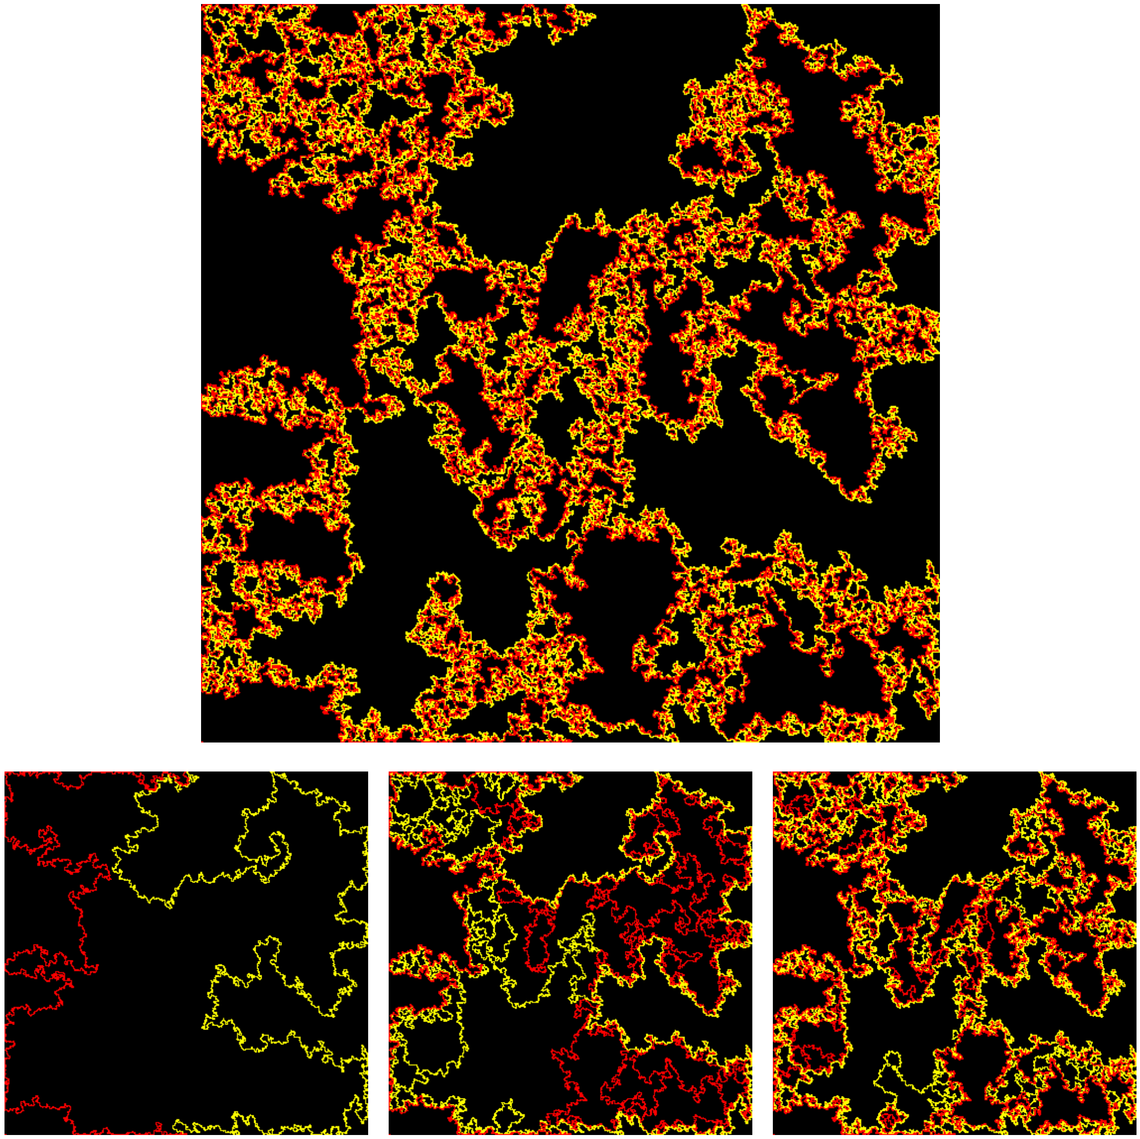

Figure 7.20: Simulation of the light cone construction of an SLE_6 curve η' in $[-1,1]^2$ from i to $-i$, generated using a projection h of a GFF on $[-1,1]^2$ onto the space of functions piecewise linear on the triangles of an 800×800 grid. The lower left panel shows left and right boundaries of η', which consist of points accessible by flowing in the vector field $e^{ih/\chi}$ for $\chi = 2/\sqrt{8/3} - \sqrt{8/3}/2$ at angle $\frac{\pi}{2}$ (red) and $-\frac{\pi}{2}$ (yellow), respectively, from $-i$. The lower middle panel shows points accessible by flowing at angle $\frac{\pi}{2}$ (red) or angle $-\frac{\pi}{2}$ (yellow) from the yellow and red points, respectively, of the left picture; the lower right shows another iteration of this. The top picture illustrates the light cone, the limit of this procedure. (All paths are red or yellow; any shade variation is a rendering artifact.)

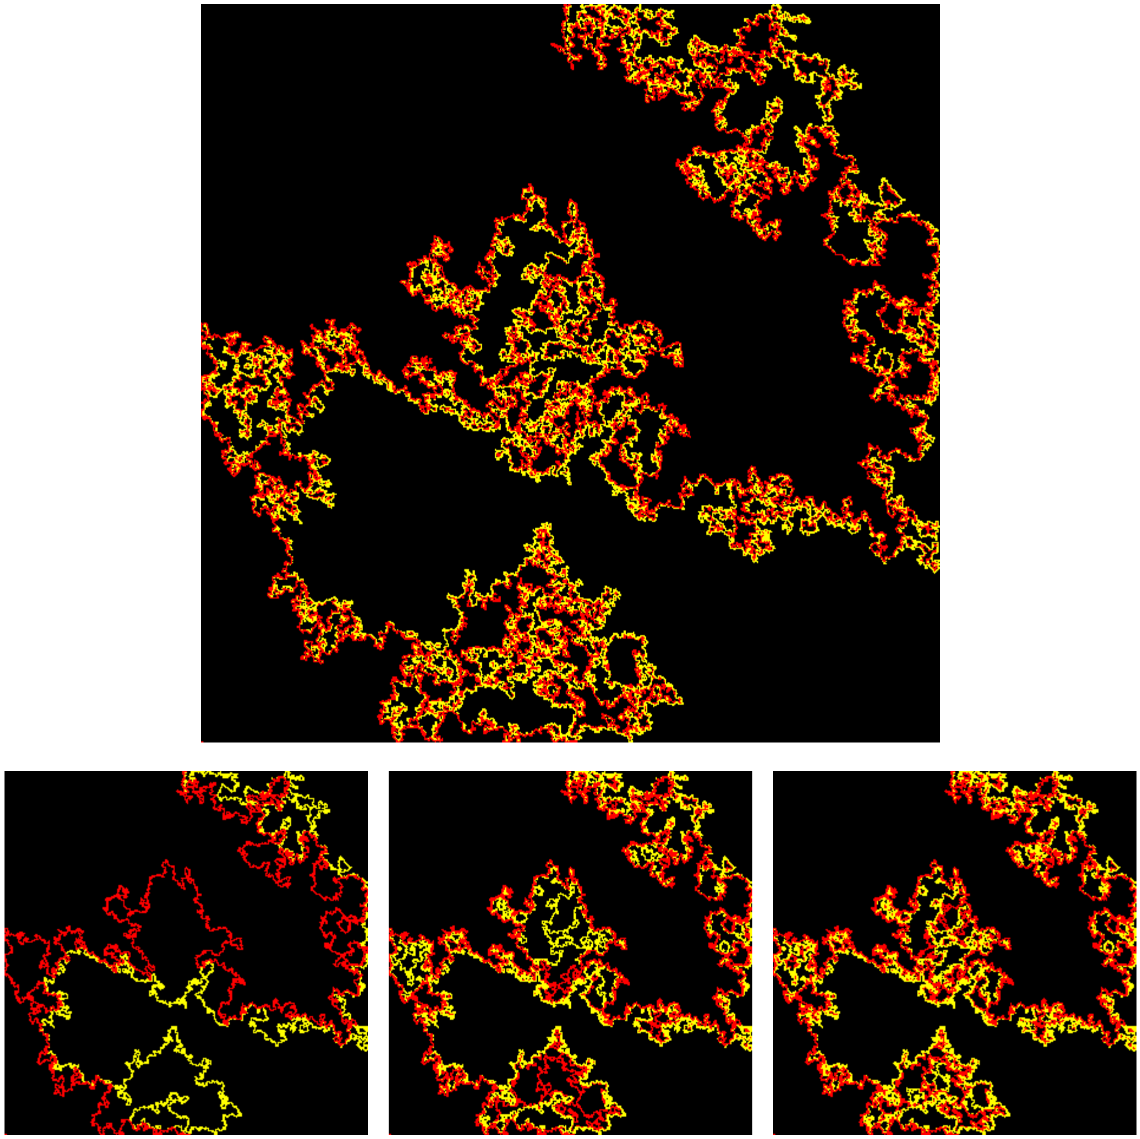

Figure 7.21: Numerical simulation of the light cone construction of an $\mathrm{SLE}_{16/3}$ process η' in $[-1,1]^2$ from i to $-i$ generated using a projection h of a GFF on $[-1,1]^2$ onto the space of functions piecewise linear on the triangles of an 800×800 grid. The lower left panel depicts the left and right boundaries of η', which correspond to the set of points accessible by flowing in the vector field $e^{ih/\chi}$ for $\chi = 2/\sqrt{3} - \sqrt{3}/2$ at angle $\frac{\pi}{2}$ (red) and $-\frac{\pi}{2}$ (yellow), respectively, from $-i$. The lower middle panel shows the set of points accessible by flowing at angle $\frac{\pi}{2}$ (red) or angle $-\frac{\pi}{2}$ (yellow) from the yellow and red points, respectively, of the left picture and the lower right panel depicts another iteration of this. The top picture illustrates the light cone, which is the limit of this procedure.

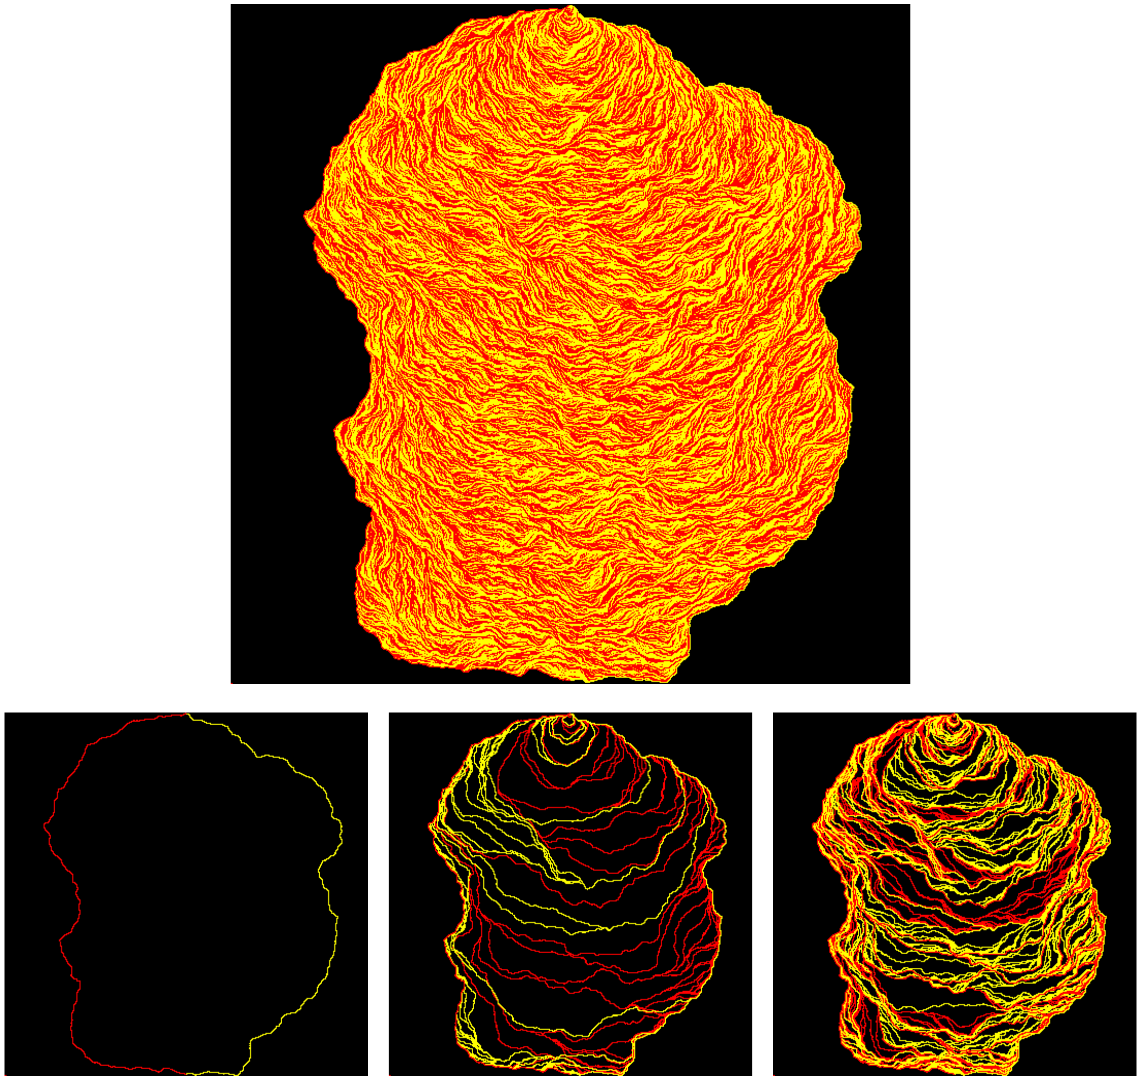

Figure 7.22: Numerical simulation of the light cone construction of an $\mathrm{SLE}_{64}(32;32)$ process η' in $[-1,1]^2$ from i to $-i$ generated using a projection h of a GFF on $[-1,1]^2$ onto the space of functions piecewise linear on the triangles of an 800×800 grid. (It turns out that $\mathrm{SLE}_{64}(\rho_1;\rho_2)$ processes are boundary filling only when $\rho_1, \rho_2 \le 28$.) The lower left panel depicts the left and right boundaries of η', which correspond to the set of points accessible by flowing in the vector field $e^{ih/\chi}$ for $\chi = 2/\sqrt{1/4} - \sqrt{1/4}/2$ at angle $\frac{\pi}{2}$ (red) and $-\frac{\pi}{2}$ (yellow), respectively, from $-i$. The lower middle panel shows the set of points accessible by flowing at angle $\frac{\pi}{2}$ (red) or angle $-\frac{\pi}{2}$ (yellow) from the yellow and red points, respectively, of the left picture and the lower right panel depicts another iteration of this. The top picture illustrates the light cone, which is the limit of this procedure.

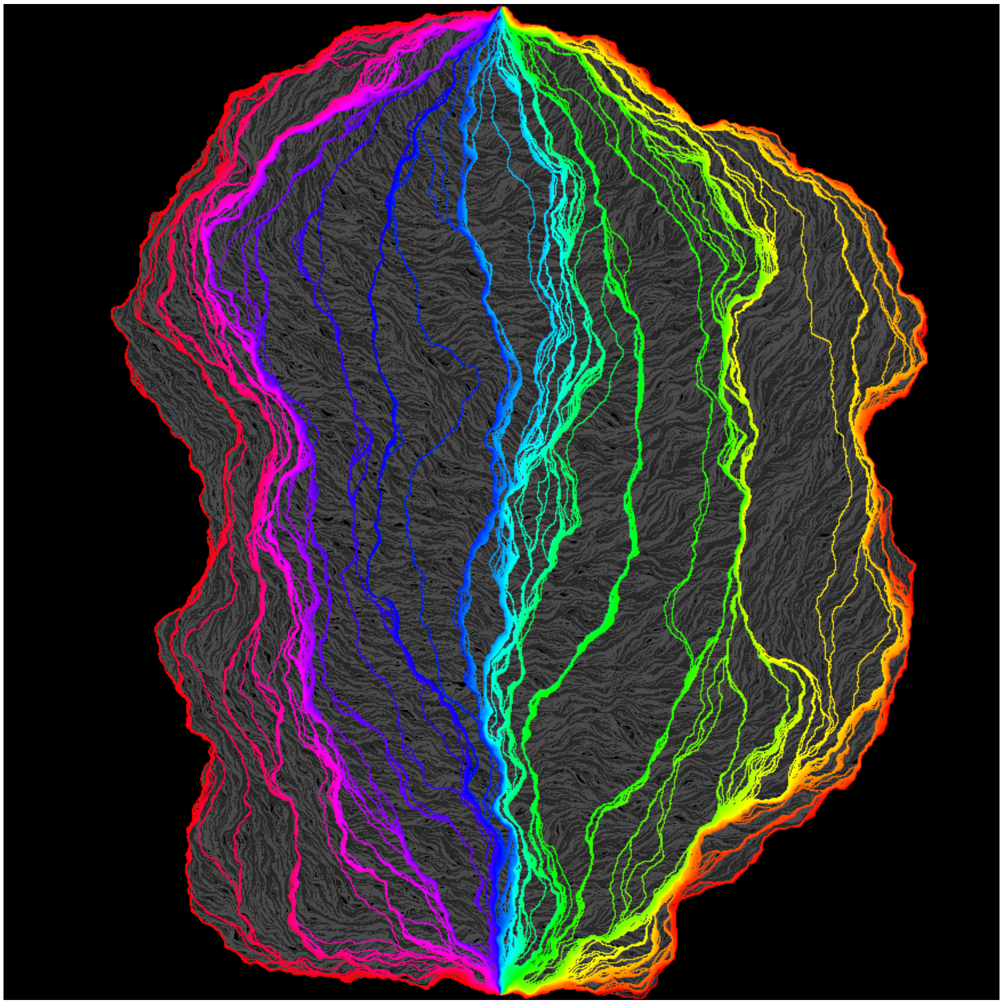

Figure 7.23: The simulation of the light cone from the top panel of Figure 7.22 where trajectories which flow at angle $\frac{\pi}{2}$ are dark gray and those which flow at angle $-\frac{\pi}{2}$ are depicted in a medium-dark gray. The fan from $-i$ — the set of all points accessible by fixed-angle trajectories with angles in $[-\frac{\pi}{2}, \frac{\pi}{2}]$ starting at $-i$ — is drawn on top of the light cone. The different colors indicate trajectories with different angles. The simulation shows that the fan does not fill the light cone.

results are described in more detail in [MS12a]. (One version of this statement also appears in [Dub09d, Section 8], where it is called "strong duality".) We will also describe the law of η' conditioned on the boundaries of the portions of η' traced before and after η' hits a given boundary point. This result will be of particular interest to us in a subsequent work, in which we will prove the time reversal symmetry of $\mathrm{SLE}_{16/\kappa}$ processes when $\kappa \in (2,4)$ (so that $16/\kappa \in (4,8)$).

The final result we wish to state concerns the interaction of imaginary rays with different

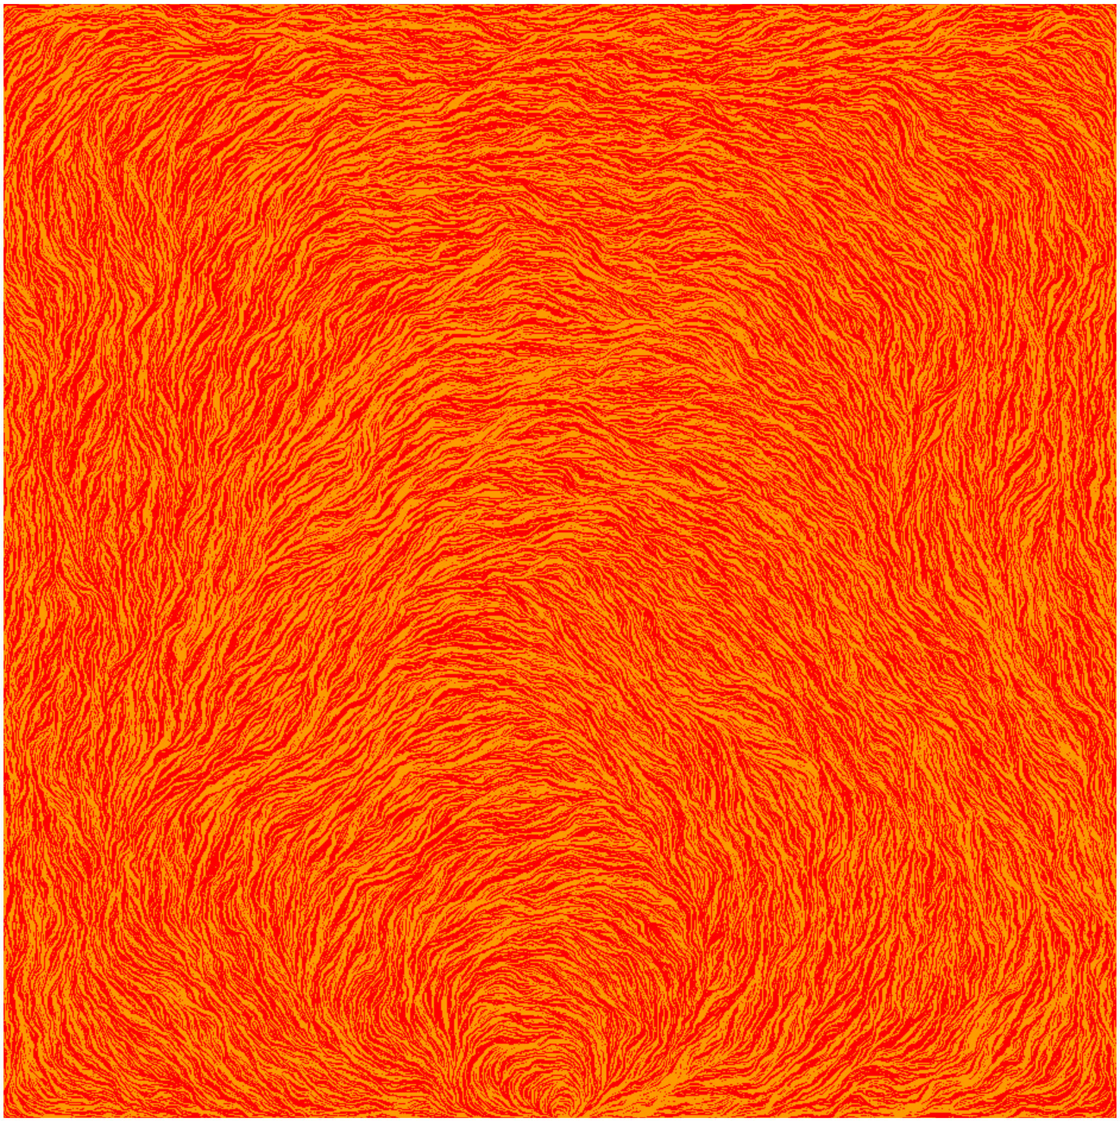

Figure 7.24: Numerical simulation of the light cone construction of an SLE_{128} process η' in $[-1,1]^2$ from i to $-i$ generated using a projection h of a GFF on $[-1,1]^2$ onto the space of functions piecewise linear on the triangles of an 800×800 grid. The red and yellow curves depict the left and right boundaries, respectively, of the time evolution of η' as it traverses $[-1,1]^2$.

angle and starting point. In contrast with the case that h is smooth, these rays may bounce off of each other and even merge, but they have the same monotonicity behavior in their starting point and angle as in the smooth case. This result leads to a theoretical understanding of the phenomena simulated in Figures 7.9–7.12, 7.14, and 7.15. The following statement is somewhat imprecise (as it does not describe all the constraints on boundary data that affect whether the distinct flow lines are certain to intersect before

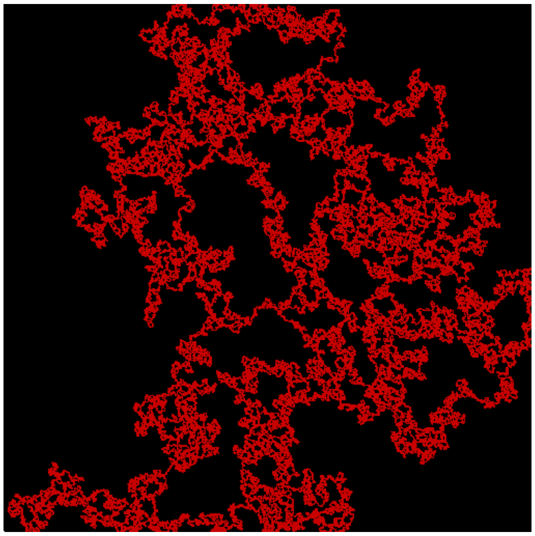

(a) An SLE_6 process η' from i to $-i$ generated using the light cone construction.

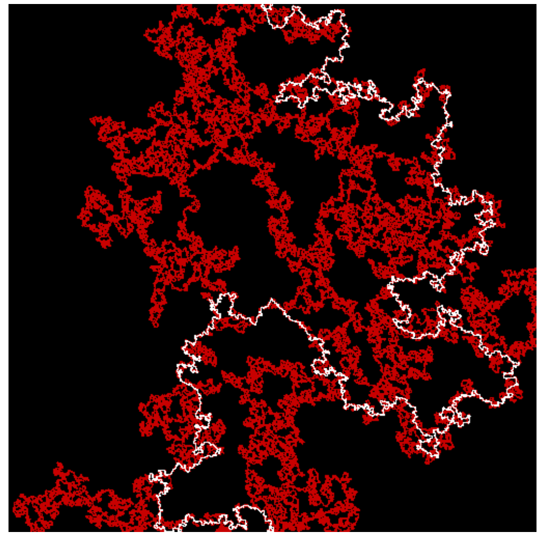

(b) The zero angle flow line η from $-i$ to i drawn on top of η'.

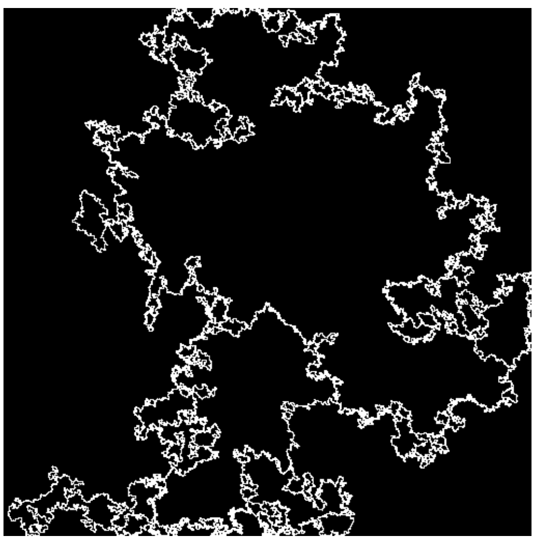

(c) The fan from $-i$ to i. The rays are $\mathrm{SLE}_{8/3}(\rho_1; \rho_2)$ processes.

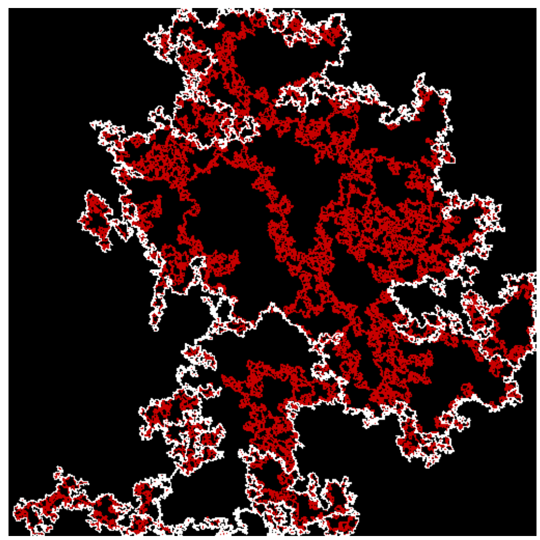

(d) The fan drawn on top of η'. It does not cover the range of η'.

Figure 7.25: Numerical simulation of the light cone construction of an SLE_6 process η' in $[-1,1]^2$ from i to $-i$ and its interaction with the zero angle flow line $\eta \sim \mathrm{SLE}_{8/3}(-1;-1)$ and the fan starting from $-i$, generated using a projection h of a GFF on $[-1,1]^2$ onto the space of functions piecewise linear on the triangles of an 800×800 grid. In the top right panel, the conditional law of the restrictions of η' given η to the left and right sides of $[-1,1]^2 \setminus \eta$ are independent $\mathrm{SLE}_6(-\frac{3}{2})$ processes.

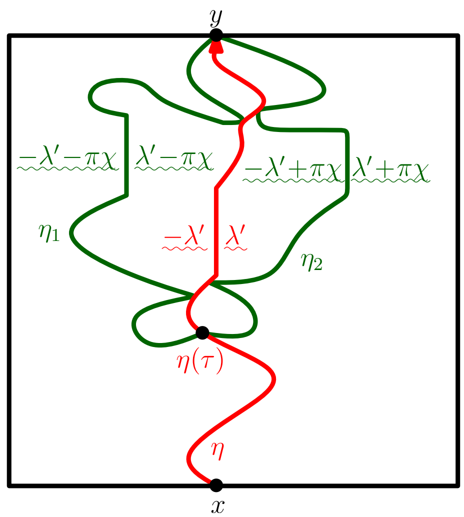

Figure 7.26: Let h be a GFF on a Jordan domain D, fix $x, y \in \partial D$ distinct, and let η be the flow line of h starting at x targeted at y. Let τ be any stopping time for η and let η_1 and η_2 be the flow lines of h conditional on η starting at $\eta(\tau)$ with angles π and $-\pi$, respectively, in the sense shown in the figure. If h were a smooth function, then we would have $\eta_1 = \eta_2$ and since π and $-\pi$ are the same modulo 2π, both paths would trace $\eta([0,\tau])$ in the reverse direction. For the GFF, we think of η_1 (resp. η_2) as starting infinitesimally to the left (resp. right) of $\eta(\tau)$; due to the roughness of the field, η_1 and η_2 do not merge into (and in fact cannot hit) $\eta([0,\tau])$. If $\kappa \in (2,4)$, then η_1 and η_2 can hit $\eta|_{(\tau,\infty)}$ and if $\kappa \in (0,2]$ then η_1 and η_2 do not hit $\eta|_{(\tau,\infty)}$. If $\kappa \in (8/3,4)$, then η_1 can hit η_2 and if $\kappa \in (0,8/3]$ then η_1 cannot hit η_2. This, in particular, explains why the yellow and red curves of Figures 7.20–7.22 do not trace each other.

getting trapped at other boundary points) but a more detailed discussion appears in [MS12a]; see also Figure 7.27 and Figure 7.28.

Theorem 7.11. *Suppose that h is a GFF on $\mathbf{H}$ with piecewise constant boundary data. For each $\theta \in \mathbf{R}$ and $x \in \partial\mathbf{H}$ we let η_θ^x be the flow line of h starting at x with angle θ. Fix $x_1, x_2 \in \partial\mathbf{H}$ with $x_1 \geq x_2$.*

(i) If $\theta_1 < \theta_2$ then $\eta_{\theta_1}^{x_1}$ almost surely stays to the right of $\eta_{\theta_2}^{x_2}$. If, in addition, $\theta_2 - \theta_1 < \pi\kappa/(4-\kappa)$, then $\eta_{\theta_1}^{x_1}$ and $\eta_{\theta_2}^{x_2}$ can bounce off of each other; otherwise the paths almost surely do not intersect (except possibly at their starting point).

(ii) If $\theta_1 = \theta_2$, then $\eta_{\theta_1}^{x_1}$ may intersect $\eta_{\theta_2}^{x_2}$ and, upon intersecting, the two curves merge and never separate.

(iii) Finally, if $\theta_2 + \pi > \theta_1 > \theta_2$, then $\eta_{\theta_1}^{x_1}$ may intersect $\eta_{\theta_2}^{x_2}$ and, upon intersecting, crosses and then never crosses back. If, in addition, $\theta_1 - \theta_2 < \pi\kappa/(4-\kappa)$, then

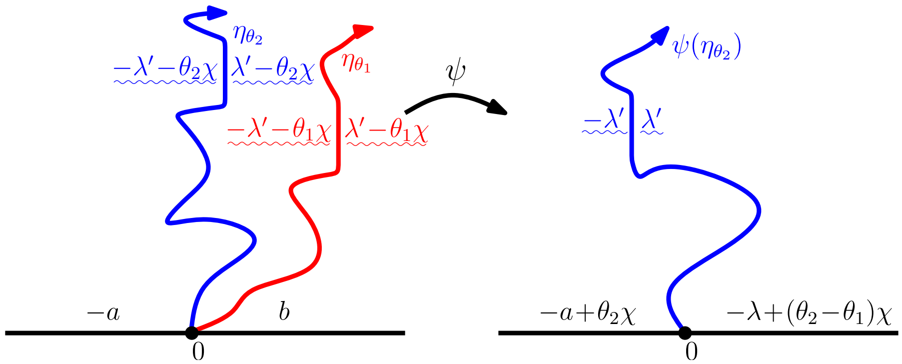

Figure 7.27: Suppose that h is a GFF on $\mathbf{H}$ with the boundary data on the left panel. For each $\theta \in \mathbf{R}$, let η_θ be the flow line of the GFF $h+\theta\chi$. This corresponds to setting the angle of η_θ to be θ. Just as if h were a smooth function, if $\theta_1 < \theta_2$ then η_{θ_1} lies to the right of η_{θ_2}. The conditional law of h given η_{θ_1} and η_{θ_2} is a GFF on $\mathbf{H}\setminus\bigcup_{i=1}^2 \eta_{\theta_i}$ whose boundary data is shown above. By applying a conformal mapping and using the transformation rule (7.12), we can compute the conditional law of η_{θ_2} given the realization of η_{θ_1} and vice-versa. That is, η_{θ_2} given η_{θ_1} is an $\mathrm{SLE}_\kappa((a-\theta_2\chi)/\lambda - 1; (\theta_2-\theta_1)\chi/\lambda - 2)$ process independently in each of the connected components of $\mathbf{H} \setminus \eta_{\theta_1}$ which lie to the left of η_{θ_1}. Moreover, η_{θ_1} given η_{θ_2} is an $\mathrm{SLE}_\kappa((\theta_2-\theta_1)\chi/\lambda - 2; (b+\theta_1\chi)/\lambda - 1)$ independently in each of the connected components of $\mathbf{H} \setminus \eta_{\theta_2}$ which lie to the right of η_{θ_2}. Versions of this result also hold for flow lines which start at different points as well as in the setting where the boundary data is piecewise constant (see Theorem 7.11).

$\eta_{\theta_1}^{x_1}$ and $\eta_{\theta_2}^{x_2}$ can bounce off of each other; otherwise the paths almost surely do not subsequently intersect.

The monotonicity component of Theorem 7.11 (i.e., the fact that $\eta_{\theta_1}^{x_1}$ almost surely stays to the right of $\eta_{\theta_2}^{x_2}$) is proved in [MS12a] first in settings where $\eta_{\theta_1}^{x}, \eta_{\theta_2}^{x}$ almost surely do not intersect $\partial\mathbf{H}$ after time 0 (and have the same starting point) in [MS12a]. It is then extended to the boundary intersecting regime and establish the merging and crossing statements. It is further shown that in [MS12a] how in the setting of Theorem 7.11 one can compute the conditional law of $\eta_{\theta_1}^{x_1}$ given $\eta_{\theta_2}^{x_2}$ and vice-versa (see Figure 7.27 for an important special case of this).

Note that the angle restriction $\theta_2 < \theta_1 < \theta_2 + \pi$ is also the one that allows the Euclidean lines to cross (i.e., would allow for η_{θ_2} to cross from the left side of η_{θ_1} to the right side if h were constant). Although we will not explore this issue here, we remark that it is also interesting to consider what would happen if we took $\theta_1 \geq \theta_2 + \pi$. It turns out that in this regime extra crossings can occur at points where both paths intersect $\mathbf{R}$, which is somewhat more complicated to describe.

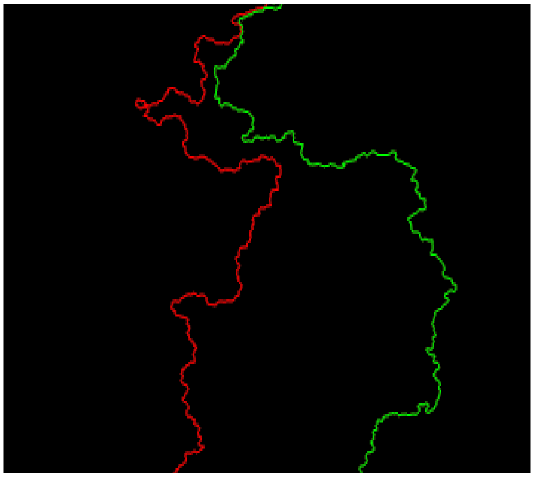

(a) If $\theta_1 < \theta_2$, then $\eta_{\theta_1}^{x_1}$ stays to the right of $\eta_{\theta_2}^{x_2}$.

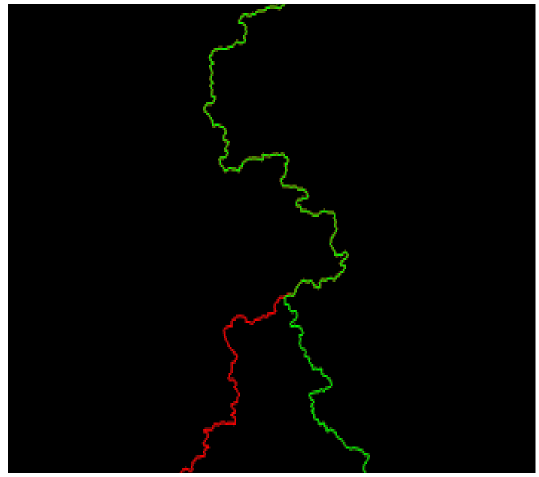

(b) If $\theta_1 = \theta_2$, then $\eta_{\theta_1}^{x_1}$ merges with $\eta_{\theta_2}^{x_2}$ upon intersecting.

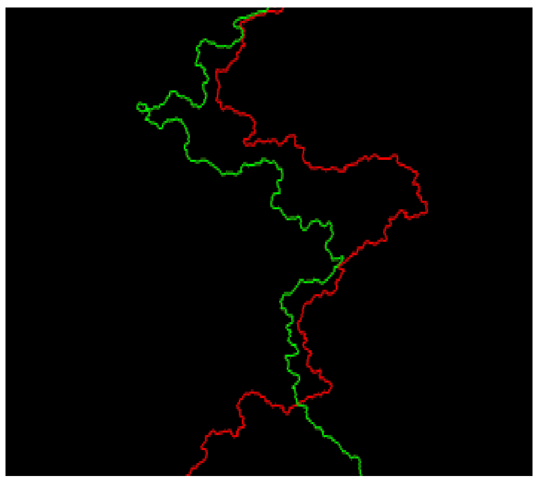

(c) If $\theta_2 < \theta_1 < \theta_2 + \pi$, then $\eta_{\theta_1}^{x_1}$ crosses $\eta_{\theta_2}^{x_2}$ upon interesting but does not cross back.

Figure 7.28: Numerical simulations which depict the three types of flow line interaction, as described in the statement of Theorem 7.11. In each of the simulations, we fixed $x_2 < x_1$ in $[-1-i, 1-i]$, $\theta_1, \theta_2 \in \mathbf{R}$, and took $\eta_{\theta_1}^{x_1}$ (resp. $\eta_{\theta_2}^{x_2}$) to be the flow line of a projection of a GFF on $[-1,1]^2$ to the space of functions piecewise linear on the triangles of a 300×300 grid starting at x_1 (resp. x_2) with angle θ_1 (resp. θ_2).

7.5 Interior flow lines

It is similarly possible to make sense of flow lines of $e^{ih(z)/\chi}$ starting from interior points of a planar domain.

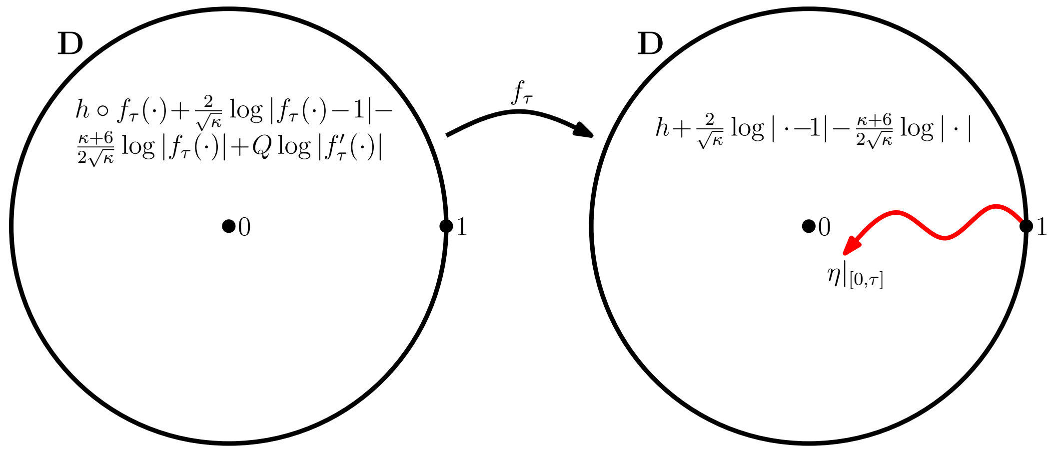

Figure 7.29: Illustration of the coupling of reverse radial SLE_κ in $\mathbf{D}$ starting from 1 and targeted at 0 with a free boundary GFF h on $\mathbf{D}$. Here, $Q = 2/\gamma + \gamma/2$ for $\gamma = \min(\sqrt{\kappa}, \sqrt{16/\kappa})$ and f_τ is the centered reverse radial SLE_κ Loewner flow evaluated at a stopping time τ. Theorem 7.12 implies that the distributions on the left and right above have the same law. (The reverse radial coupling of SLE_κ with the free boundary GFF can also be formulated using forward SLE_κ; see Figure 10.32 for an illustration.)

7.6 Counterflow lines and space-filling SLE

The tree and dual tree of flow lines have an interface that can be described as a space-filling curve.

7.7 Time reversal symmetries

Imaginary geometry can be used to prove several basic facts about SLE, including time reversal symmetry for several forms of SLE_κ with $\kappa < 4$ and $\mathrm{SLE}_{\kappa'}$ with $\kappa' > 4$.

7.8 The reverse radial SLE/GFF coupling

The purpose of this section is to establish the radial version of the reverse coupling of SLE_κ with the free boundary GFF. It is a generalization of the coupling with reverse chordal SLE_κ with the free boundary GFF established in [She10, Theorem 1.2]. Suppose that B_t is a standard Brownian motion, $W_t = \sqrt{\kappa} B_t$, and $U_t = e^{iW_t}$. Let (g_t) solve the reverse radial Loewner ODE (4.4) driven by U_t. The centered reverse SLE_κ is given by the centered conformal maps $f_t = U_t^{-1} g_t$. We note that

$$\begin{aligned} df_t(z) &= U_t^{-1} dg_t(z) - iU_t^{-1} g_t(z) dW_t - \frac{\kappa}{2} U_t^{-1} g_t(z) dt \\ &= -f_t(z) \left(\frac{1+f_t(z)}{1-f_t(z)} + \frac{\kappa}{2} \right) dt - i f_t(z) dW_t \\ &= -\left(\Phi(1, f_t(z)) + \frac{\kappa}{2} f_t(z) \right) dt - i f_t(z) dW_t \end{aligned} \tag{7.20}$$

(recall (4.2)).

Theorem 7.12. *Fix $\kappa > 0$. Suppose that h is a free boundary GFF on $\mathbf{D}$, let B be a standard Brownian motion which is independent of h, and let (f_t) be the centered reverse radial SLE_κ Loewner flow which is driven by $U_t = e^{iW_t}$ where $W = \sqrt{\kappa} B$ as in (7.20). For each $t \geq 0$ and $z \in \mathbf{D}$ we let*[13]

$$\mathfrak{h}_t(z) = \frac{2}{\sqrt{\kappa}} \log |f_t(z) - 1| - \frac{\kappa + 6}{2\sqrt{\kappa}} \log |f_t(z)| + Q \log |f_t'(z)| \tag{7.21}$$

where $Q = 2/\gamma + \gamma/2$ and $\gamma = \min(\sqrt{\kappa}, \sqrt{16/\kappa})$. Let τ be an almost surely finite stopping time for the filtration generated by W. Then

$$h + \mathfrak{h}_0 \stackrel{d}{=} h \circ f_\tau + \mathfrak{h}_\tau \tag{7.22}$$

where we view the left and right sides as distributions defined modulo additive constant.

Theorem 7.12 states that the law of $h + \mathfrak{h}_0$ is invariant under the operation of sampling an independent SLE_κ process η and then drawing it on top of $h + \mathfrak{h}_0$ up until some time t and then applying the change of coordinates formula for quantum surfaces using the forward radial Loewner flow for η at time t. An illustration of the setup for Theorem 7.12 is given in Figure 7.29.

We include the following self-contained proof of Theorem 7.12 for the convenience of the reader which follows the strategy of [She10]. The first step, carried out in Lemma 7.13, is to compute the Ito derivatives of some quantities which are related to the right side of (7.22). Next, we show in Lemma 7.14 that the random variable on the right hand side of (7.22) takes values in the space of distributions and, when integrated against a given smooth mean-zero test function, yields a process which is continuous in time. We then compute another Ito derivative in Lemma 7.15 and afterwards combine the different steps to complete the proof.

Let G denote the Neumann Green's function for Δ on $\mathbf{D}$ given in (3.10). Suppose that (g_t) is the reverse radial SLE_κ Loewner flow and (f_t) is the corresponding centered flow as in Theorem 7.12. Throughout, we let

$$G_t(y, z) = G(f_t(y), f_t(z)) = G(g_t(y), g_t(z)) \quad \text{for each} \quad t \geq 0.$$

[13]The function $\mathfrak{h}_t$ in the statement of Theorem 7.12 is not the same as the harmonic component in the definition of QLE. We are using this notation in this section to be consistent with the notation used in [She10].

We also let $\mathcal{P}$ (resp. $\overline{\mathcal{P}}$) denote 2π times the Poisson (resp. conjugate Poisson) kernel on $\mathbf{D}$. Explicitly,

$$\big(\mathcal{P} + i\overline{\mathcal{P}}\big)(z,w) = \frac{w+z}{w-z} = \Psi(w,z). \tag{7.23}$$

That is, $\mathcal{P}$ (resp. $\overline{\mathcal{P}}$) is given by the real (resp. imaginary) part of the expression in the right side above.

Lemma 7.13. *Suppose that we have the same setup as in Theorem 7.12. There exists a smooth function $\phi\colon \mathbf{D} \to \mathbf{R}$ such that the following is true. For each $y, z \in \mathbf{D}$ we have that*

$$dG_t(y,z) = \Big(\phi(f_t(y)) + \phi(f_t(z)) - \overline{\mathcal{P}}(1,f_t(y))\overline{\mathcal{P}}(1,f_t(z))\Big)dt \quad \textit{and} \tag{7.24}$$

$$d\mathfrak{h}_t(z) = \frac{1}{\sqrt{\kappa}}dt - \overline{\mathcal{P}}(1,f_t(z))dB_t. \tag{7.25}$$

When we apply Lemma 7.13 later in this section, we will consider $G_t(y,z)$ and $\mathfrak{h}_t(z)$ integrated against mean zero test functions. In particular, the terms involving ϕ for $dG_t(y,z)$ and the term $(1/\sqrt{\kappa})dt$ in $d\mathfrak{h}_t(z)$ will drop out.

Proof of Lemma 7.13. Both (7.24) and (7.25) follow from applications of Ito's formula. In particular,

$$d\log(g_t(y) - g_t(z)) = \frac{g_t(y)g_t(z) - U_t(g_t(y)+g_t(z)) - U_t^2}{(U_t - g_t(z))(U_t - g_t(y))}dt \quad \text{and} \tag{7.26}$$

$$d\log(1 - \overline{g_t(y)}g_t(z)) = \frac{2g_t(z)\overline{g_t(y)}}{(\overline{U}_t - \overline{g_t(y)})(U_t - g_t(z))}dt. \tag{7.27}$$

We note that (7.26) and (7.27) do not depend on the choice of driving function. A tedious calculation thus shows that $dG_t(y,z) + \overline{\mathcal{P}}(1,f_t(y))\overline{\mathcal{P}}(1,f_t(z))dt$ can be written as $\phi(f_t(y)) + \phi(f_t(z))$ where ϕ is a smooth function. This gives (7.24).

For (7.25), we fix $z \in \mathbf{D}$ and write $f_t = f_t(z)$. Then we can express $\mathfrak{h}_t(z)$ in terms of the real part of

$$\log(f_t - 1), \quad \log(f_t), \quad \text{and} \quad \log(f_t'). \tag{7.28}$$

The Ito derivative of f_t is given in (7.20). Differentiating this with respect to z yields

$$df_t' = -f_t'\left(\frac{1+f_t}{1-f_t} + \frac{2f_t}{(1-f_t)^2} + \frac{\kappa}{2}\right)dt - if_t'dW_t. \tag{7.29}$$

Applying (7.20) and (7.29), we see that the Ito derivatives of the terms in (7.28) are given by

$$d\log(f_t - 1) = \left(\frac{(1+\frac{\kappa}{2})f_t + f_t^2}{(1-f_t)^2}\right)dt + \frac{if_t}{1-f_t}dW_t,$$

$$\begin{aligned} d\log(f_t) &= -\left(\frac{1+f_t}{1-f_t}\right)dt - idW_t, \quad \text{and} \\ d\log(f_t') &= \left(1 - \frac{2}{(1-f_t)^2}\right)dt - idW_t. \end{aligned}$$

This implies that $d\mathfrak{h}_t(z)$ is given by the real part of

$$\frac{1}{\sqrt{\kappa}}dt + i\left(\frac{1+f_t}{1-f_t}\right)dB_t,$$

from which (7.25) follows. □

Lemma 7.14. *Suppose that we have the same setup as in Theorem 7.12. For each $t \geq 0$, the random variable $h \circ f_t + \mathfrak{h}_t$ takes values in the space of distributions defined modulo additive constant. Moreover, for any fixed $\rho \in C_0^\infty(\mathbf{D})$ with $\int_{\mathbf{D}} \rho(z)dz = 0$, both $(h \circ f_t + \mathfrak{h}_t, \rho)$ and $(\mathfrak{h}_t, \rho)$ are almost surely continuous and the latter is a square-integrable martingale.*

Proof. Fix $\rho \in C_0^\infty(\mathbf{D})$ with $\int_{\mathbf{D}} \rho(z)dz = 0$. We first note that it is clear that $h \circ f_t$ takes values in the space of distributions modulo additive constant. Moreover, $t \mapsto (h \circ f_t, \rho)$ is almost surely continuous from how it is defined. Indeed, by definition we have that

$$(h \circ f_t, \rho) = (h, \rho_t) \quad \text{where} \quad \rho_t = |(f_t^{-1})'|^2 \rho \circ f_t^{-1}.$$

Fix a value of $t > 0$ and $\delta \in (0, t)$ and note that there exists a compact set K such that the support of ρ_s is contained in K for all $s \in (t-\delta, t+\delta)$. Moreover, it is clear that $\rho_s \to \rho_t$ uniformly as $s \to t$ as well as all of its derivatives. This proves the continuity of $(h \circ f_t, \rho)$.

We are left to deal with $\mathfrak{h}_t$. It follows from (7.25) of Lemma 7.13 that

$$d\langle \mathfrak{h}_t(z) \rangle = \left(\overline{\mathcal{P}}(1, f_t(z))\right)^2 dt. \tag{7.30}$$

By the Schwarz lemma, we note that $|f_t(z)| \leq |z|$ for all $z \in \mathbf{D}$ and $t \geq 0$. Consequently, it follows from (7.23) that for each $r \in (0,1)$ there exists $C_r \in (0, \infty)$ such that

$$\sup_{z \in r\mathbf{D}} \langle \mathfrak{h}_t(z) - \mathfrak{h}_u(z) \rangle \leq C_r(t-u) \quad \text{for all} \quad 0 \leq u \leq t < \infty. \tag{7.31}$$

It therefore follows from the Burkholder-Davis-Gundy inequality that for each $p \geq 1$ and $r \in (0,1)$ there exists $C_p, C_{\kappa,r,p} \in (0,\infty)$ such that for all $0 \leq u \leq t$ we have

$$\begin{aligned} &\sup_{z \in r\mathbf{D}} \mathbf{E}\left[\sup_{u \leq s \leq t} |\mathfrak{h}_s(z) - \mathfrak{h}_u(z)|^p\right] \\ \leq &C_p\left(\sup_{z \in r\mathbf{D}} \mathbf{E}\left[\langle \mathfrak{h}_t(z) - \mathfrak{h}_u(z)\rangle^{p/2}\right] + \frac{1}{\kappa^{p/2}}(t-u)^p\right) \end{aligned}$$

$$\leq C_{\kappa,r,p}\bigg((t-u)^p + (t-u)^{p/2}\bigg). \tag{7.32}$$

It is easy to see from (7.32) with $u = 0$ and Fubini's theorem that for each $r \in (0,1)$ we have $\mathfrak{h}_t|_{r\mathbf{D}}$ is almost surely in $L^p(r\mathbf{D})$. By combining (7.32) with a large enough value of $p > 1$ and the Kolmogorov-Čentsov theorem, it is also easy to see that $t \mapsto (\mathfrak{h}_t, \rho)$ is almost surely continuous for any $\rho \in C_0^\infty(\mathbf{D})$ with $\int_\mathbf{D} \rho(z)dz = 0$. Lastly, it follows from (7.32) and (7.25) of Lemma 7.13 that $(\mathfrak{h}_t, \rho)$ is a square-integrable martingale. This completes the proof of both assertions of the lemma. □

For each $\rho \in C_0^\infty(\mathbf{D})$ with $\int_\mathbf{D} \rho(z)dz = 0$ and $t \geq 0$ we let

$$E_t(\rho) = \int_\mathbf{D}\int_\mathbf{D} \rho(y) G_t(y,z)\rho(z)dydz$$

be the conditional variance of $(h \circ f_t, \rho)$ given f_t.

Lemma 7.15. *For each $\rho \in C_0^\infty(\mathbf{D})$ with $\int_\mathbf{D} \rho(z)dz = 0$ we have that*

$$d\langle(\mathfrak{h}_t, \rho)\rangle = -dE_t(\rho).$$

Proof. Since $(\mathfrak{h}_t, \rho)$ is a continuous L^2 martingale, the process $\langle(\mathfrak{h}_t, \rho)\rangle$ is characterized by the property that

$$(\mathfrak{h}_t, \rho)^2 - \langle(\mathfrak{h}_t, \rho)\rangle$$

is a continuous local martingale in $t \geq 0$. Thus to complete the proof of the lemma, it suffices to show that

$$(\mathfrak{h}_t, \rho)^2 + E_t(\rho)$$

is a continuous local martingale. It follows from (7.24) and (7.25) of Lemma 7.13 that

$$\mathfrak{h}_t(y)\mathfrak{h}_t(z) + G_t(y,z)$$

evolves as the sum of a martingale in $t \geq 0$ plus a drift term which can be expressed as a sum of terms one of which depends only on y and the other only on z. These drift terms cancel upon integrating against $\rho(y)\rho(z)dydz$ which in turn implies the desired result. □

Proof of Theorem 7.12. Fix $\rho \in C_0^\infty(\mathbf{D})$ with $\int_\mathbf{D} \rho(z)dz = 0$. Let $\mathcal{F}_t$ be the filtration generated by f_t. Note that $\mathfrak{h}_t$ is $\mathcal{F}_t$-measurable and that, given $\mathcal{F}_t$, $(h \circ f_t, \rho)$ is a Gaussian random variable with mean zero and variance $E_t(\rho)$. Let $I_t(\rho) = (h \circ f_t + \mathfrak{h}_t, \rho)$. For $\theta \in \mathbf{R}$ we have that:

$$\begin{aligned}\mathbf{E}[\exp(i\theta I_t(\rho))] &= \mathbf{E}[\mathbf{E}[\exp(i\theta I_t(\rho))|\mathcal{F}_t]]\\ &= \mathbf{E}[\mathbf{E}[\exp(i\theta(h \circ f_t, \rho))|\mathcal{F}_t]\exp(i\theta(\mathfrak{h}_t, \rho))]\end{aligned}$$

$$=\mathbf{E}[\exp(i\theta(\mathfrak{h}_t,\rho) - \tfrac{\theta^2}{2}E_t(\rho))]$$
$$=\exp(i\theta(\mathfrak{h}_0,\rho) - \tfrac{\theta^2}{2}E_0(\rho)).$$

Therefore $I_t(\rho) \stackrel{d}{=} I_0(\rho)$ for each $\rho \in C_0^\infty(\mathbf{D})$ with $\int_{\mathbf{D}} \rho(z)dz = 0$. The result follows since this holds for all such test functions ρ and $\rho \mapsto I_0(\rho)$ has a Gaussian distribution. $\square$

Reverse radial $\mathrm{SLE}_\kappa(\rho)$ is a variant of reverse radial SLE_κ in which one keeps track of an extra marked point on $\partial\mathbf{D}$. It is defined in an analogous way to reverse radial SLE_κ except the driving function U_t is taken to be a solution to the SDE:

$$\begin{aligned} dU_t &= -\frac{\kappa}{2}U_t dt + i\sqrt{\kappa}U_t dB_t + \frac{\rho}{2}\Phi(V_t,U_t)dt \\ dV_t &= -\Phi(U_t,V_t)dt. \end{aligned} \tag{7.33}$$

Observe that when $\rho = 0$ this is the same as the driving SDE for ordinary reverse radial SLE_κ. This is analogous to the definition of forward radial $\mathrm{SLE}_\kappa(\rho)$ up to a change of signs (see, for example, [SW05, Section 2]). In analogy with Theorem 7.12, it is also possible to couple reverse radial $\mathrm{SLE}_\kappa(\rho)$ with the GFF (the chordal version of this is [She10, Theorem 4.5]).

Theorem 7.16. *Fix $\kappa > 0$. Suppose that h is a free boundary GFF on $\mathbf{D}$ and let (f_t) be the centered reverse radial $\mathrm{SLE}_\kappa(\rho)$ Loewner flow which is driven by the solution U as in (7.33) with $V_0 = v_0 \in \partial\mathbf{D}$ taken to be independent of h. For each $t \geq 0$ and $z \in \mathbf{D}$ we let*

$$\begin{aligned} \mathfrak{h}_t(z) =& \frac{2}{\sqrt{\kappa}}\log|f_t(z)-1| - \frac{\kappa+6-\rho}{2\sqrt{\kappa}}\log|f_t(z)| - \\ & \frac{\rho}{\sqrt{\kappa}}\log|f_t(z)-V_t| + Q\log|f_t'(z)| \end{aligned} \tag{7.34}$$

where $Q = 2/\gamma + \gamma/2$ and $\gamma = \min(\sqrt{\kappa}, \sqrt{16/\kappa})$. Let τ be an almost surely finite stopping time for the filtration generated by W which occurs before the first time t that $f_t(v_0) = 1$. Then

$$h + \mathfrak{h}_0 \stackrel{d}{=} h \circ f_\tau + \mathfrak{h}_\tau \tag{7.35}$$

where we view the left and right sides as distributions defined modulo additive constant.

Proof. This result is proved in the same manner as Theorem 7.12; the only difference is that the calculations needed to verify that the analogy of the assertion of (7.25) from Lemma 7.13 also holds in the setting of the present theorem. As in the proof of Lemma 7.13, we will not spell out all of the calculations but only indicate the high level steps. Fix $z \in \mathbf{D}$ and write $f_t = f_t(z)$. We also let

$$Z_t = U_t^{-1}V_t \quad \text{and} \quad A_t = \frac{\rho}{2}\Phi(Z_t,1).$$

We will now explain how to show that

$$d\mathfrak{h}_t(z) = -\mathrm{Re}\left(\frac{(A_t-1)(2-\rho)}{2\sqrt{\kappa}}\right)dt - \overline{\mathcal{P}}(1,f_t)dB_t. \tag{7.36}$$

Note that the diffusion term does not depend on ρ. Moreover, the drift term does not depend on z and so integrates to zero against any mean-zero test function.

First, we note that

$$df_t = -f_t\left(\frac{1+f_t}{1-f_t} + A_t + \frac{\kappa}{2}\right)dt - i\sqrt{\kappa}f_t dB_t. \tag{7.37}$$

Applying this for $z = v_0$ also gives dZ_t. Differentiating both sides with respect to z yields

$$df_t' = -f_t'\left(\frac{1+f_t}{1-f_t} + \frac{2f_t}{(1-f_t)^2} + A_t + \frac{\kappa}{2}\right)dt - i\sqrt{\kappa}f_t' dB_t. \tag{7.38}$$

Using (7.37) and (7.38), we thus see that

$$\begin{aligned}
d\log(f_t-1) &= \left(\frac{(1+\frac{\kappa}{2})f_t + f_t^2}{(1-f_t)^2} + \frac{A_t f_t}{1-f_t}\right)dt + \frac{if_t}{1-f_t}\sqrt{\kappa}dB_t,\\
d\log(f_t) &= -\left(\frac{1+f_t}{1-f_t} + A_t\right)dt - i\sqrt{\kappa}dB_t,\\
d\log(f_t') &= \left(1 - \frac{2}{(1-f_t)^2} - A_t\right)dt - i\sqrt{\kappa}dB_t, \quad \text{and}\\
d\log(f_t - Z_t) &= \left(\frac{Z_t+1}{Z_t-1}\cdot\frac{1}{1-f_t} - \frac{f_t}{1-f_t} - A_t\right)dt - i\sqrt{\kappa}dB_t.
\end{aligned}$$

Adding these expressions up gives (7.36). □

8 Conformal welding and the quantum zipper

8.1 Welding simple quantum wedges

One can "conformally weld" two quantum wedges to each other to obtain a new thicker quantum wedge. The first version of this story (which applies to two wedges of a particular thickness) was described in [She10]

As we have discussed already, Liouville quantum gravity and the Schramm-Loewner evolution (SLE) rank among the great mathematical physics discoveries of the last few decades. Liouville quantum gravity, introduced in the physics literature by Polyakov in 1981 in the context of string theory, is a canonical model of a random two dimensional Riemannian manifold [Pol81b, Pol81c]. The Schramm-Loewner evolution, introduced by Schramm in 1999, is a canonical model of a random path in the plane that doesn't cross itself [Sch00b]. Each of these models is the subject of a large and active literature spanning physics and mathematics.

Our goal here is to connect these two objects to each other in the simplest possible way. Roughly speaking, we will show that if one glues together two independent Liouville quantum gravity random surfaces along boundary segments (in a boundary-length-preserving way) — and then conformally maps the resulting surface to a planar domain — then the interface between the two surfaces is an SLE.

Peter Jones conjectured several years ago that SLE could be obtained in a similar way — specifically, by gluing (what in our language amounts to) one Liouville quantum gravity random surface and one deterministic Euclidean disc. Astala, Jones, Kupiainen, and Saksman showed that the construction Jones proposed produces a well-defined curve [AJKS09, AJKS10], but Binder and Smirnov recently announced a proof (involving multifractal exponents) that this curve is not a form of SLE, and hence the original Jones conjecture is false [Smi] (see Section 8.4). Our construction shows that a simple variant of the Jones conjecture is in fact true.

Beyond this, we discover some surprising symmetries. For example, it turns out that there is one particularly natural random simply connected surface (called a γ-*quantum wedge*) that has an infinite-length boundary isometric to $\mathbf{R}$ (almost surely) which contains a distinguished "origin." Although this surface is simply connected, it is almost surely highly non-smooth and it has a random fractal structure. We will explain precisely how it is defined in Section 8.5. The origin divides the boundary into two infinite-length boundary arcs. Suppose we glue (in a boundary-length preserving way) the right arc of one such surface to the left arc of an independent random surface with the same law, then conformally map the combined surface to the complex upper half plane $\mathbb{H}$ (sending the origin to the origin and ∞ to ∞ — see figure below), and then *erase* the boundary interface. The geometric structure of the combined surface can be pushed forward to give geometric structure (including an area measure) on $\mathbb{H}$. It is natural to wonder how well one can guess, from this geometric structure on $\mathbb{H}$, where the now-erased interface used to be.

We will show that the geometric structure yields no information at all. That is, the conditional law of the interface is that of an SLE in $\mathbb{H}$ *independently* of the underlying geometry (a fact formally stated as part of Theorem 8.6). Another way to put this is that conditioned on the combined surface, all of the information about the interface is contained in the *conformal structure* of the combined surface, which determines the embedding in $\mathbb{H}$ (up to rescaling $\mathbb{H}$ via multiplication by a positive constant, which does not affect the law of the path, since the law of SLE is scale-invariant).

This apparent coincidence is actually quite natural from one point of view. We recall that one reason (among many) for studying SLE is that it arises as the fine mesh "scaling limit" of random simple paths on lattices. Liouville quantum gravity is similarly believed (though not proved) to be the scaling limit of random discretized surfaces and random

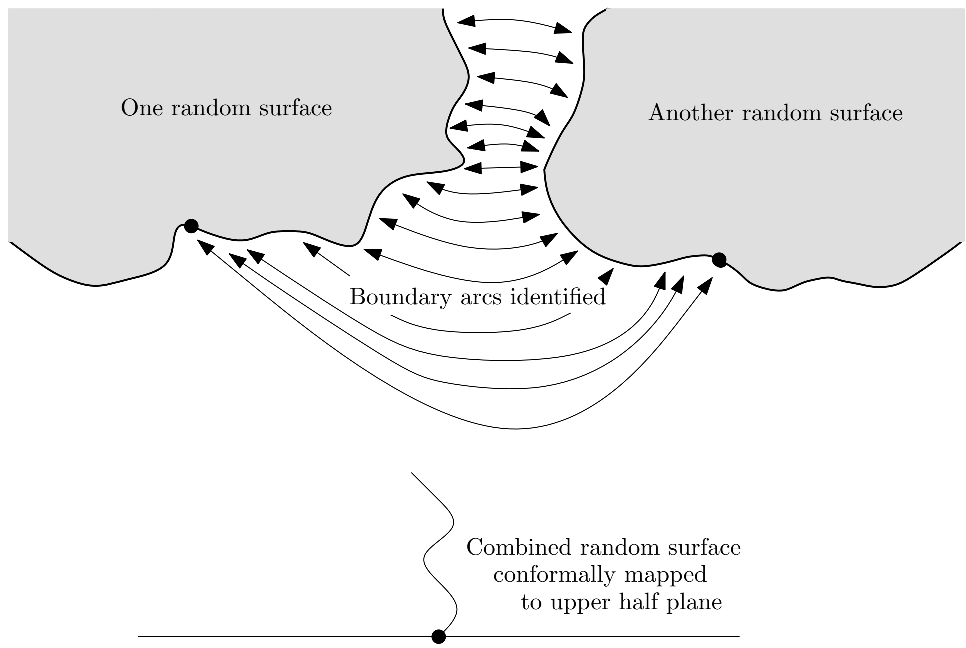

planar maps. The independence mentioned above turns out to be consistent with (indeed, at least heuristically, a *consequence of*) certain scaling limit conjectures (and a related conformal invariance Ansatz) that we will formulate precisely (in Section 8.6) for the first time here.

Polyakov initially proposed Liouville quantum gravity as a model for the intrinsic Riemannian manifold parameterizing the space-time trajectory of a string [Pol81b]. From this point of view, the welding/subdivision of such surfaces is analogous to the concatenation/subdivision of one-dimensional time intervals (which parameterize point-particle trajectories). It seems natural to try to understand complicated string trajectories by decomposing them into simpler pieces (and/or gluing pieces together), which should involve subdividing and/or welding the corresponding Liouville quantum gravity surfaces. The purpose of this section is to study these weldings and subdivisions mathematically. We will not further explore the physical implications here.

In a recent memoir [Pol08b], Polyakov writes that he first became convinced of the connection between the discrete models and Liouville quantum gravity in the 1980's after jointly deriving, with Knizhnik and Zamolodchikov, the so-called *KPZ formula* for certain Liouville quantum gravity scaling dimensions and comparing them with known combinatorial results for the discrete models [KPZ88b]. With Duplantier, the present author recently formulated and proved the KPZ formula in a mathematical way [DS11a] (see also [BS09b, RV08b]). We refer the reader there for references and history.

We will find it instructive to develop Liouville quantum gravity along with a closely related construction called the *AC geometry* or *imaginary geometry*. Both Liouville quantum gravity and the imaginary geometry are based on a simple object called the Gaussian free field.

8.2 Random geometries from the Gaussian free field

The two dimensional Gaussian free field (GFF) is a natural higher dimensional analog of Brownian motion that plays a prominent role in mathematics and physics. See the survey [She07] and the introductions of [SS09c, SS10] for a detailed account. On a planar domain D, one can define both a zero boundary GFF and a free boundary GFF (the latter being defined only modulo an additive constant, which we will sometimes fix arbitrarily). In both cases, an instance of the GFF is a random sum

$$h = \sum_i \alpha_i f_i,$$

where the α_i are i.i.d. mean-zero unit-variance normal random variables, and the f_i are an orthonormal basis for a Hilbert space of real-valued functions on D (or in the free boundary case, functions modulo additive constants) endowed with the Dirichlet inner product

$$(f_1, f_2)_\nabla := (2\pi)^{-1} \int_D \nabla f_1(z) \cdot \nabla f_2(z) dz.$$

The Hilbert space is the completion of either the space of smooth compactly supported functions $f : D \to \mathbf{R}$ (zero boundary) or the space of all smooth functions $f : D \to \mathbf{R}$ modulo additive constants with $(f, f)_\nabla < \infty$ (free boundary). In each case, h is understood not as a random function on D but as a random distribution or generalized function on D. (Mean values of h on certain sets are also defined, but the value of h at a particular point is not defined.) One can fix the additive constant for the free boundary GFF in various ways, e.g., by requiring the mean value of h on some set to be zero.

There are two natural ways to produce a "random geometry" from the Gaussian free field. The first construction is (critical) **Liouville quantum gravity**. Here, one replaces the usual Lebesgue measure dz on a smooth domain D with a random measure $\mu_h = e^{\gamma h(z)} dz$, where $\gamma \in [0, 2)$ is a fixed constant and h is an instance of (for now) the free boundary GFF on D (with an additive constant somehow fixed — there are various ways of fixing the additive constant; one way is to require the mean value of h on some fixed set to be 0). Since h is not defined as a function on D, one has to use a regularization procedure to be precise:

$$\mu = \mu_h := \lim_{\varepsilon \to 0} \varepsilon^{\gamma^2/2} e^{\gamma h_\varepsilon(z)} dz, \tag{8.1}$$

where dz is Lebesgue measure on D, $h_\varepsilon(z)$ is the mean value of h on the circle $\partial B_\varepsilon(z)$ and the limit represents weak convergence (on compact subsets) in the space of measures on D. (The limit exists almost surely, at least if ε is restricted to powers of two [DS11a].) We interpret μ_h as the area measure of a random surface conformally parameterized by D. When $x \in \partial D$, we let $h_\varepsilon(x)$ be the mean value of h on $D \cap \partial B_\varepsilon(x)$. On a linear segment of ∂D, we may define a boundary length measure by

$$\nu = \nu_h := \lim_{\varepsilon \to 0} \varepsilon^{\gamma^2/4} e^{\gamma h_\varepsilon(x)/2} dx, \tag{8.2}$$

where dx is Lebesgue measure on ∂D. (For details see [DS11a], which also relates the above random measures to the curvature-based action used to define Liouville quantum gravity in the physics literature.)

We could also parameterize the same surface with a different domain $\widetilde{D}$, and our regularization procedure implies a simple rule for changing coordinates. Suppose that ψ is a conformal map from a domain $\widetilde{D}$ to D and write $\widetilde{h}$ for the distribution on $\widetilde{D}$ given by $h \circ \psi + Q \log |\psi'|$ where

$$Q := \frac{2}{\gamma} + \frac{\gamma}{2},$$

as in Figure 9.8 [14]. Then μ_h is almost surely the image under ψ of the measure $\mu_{\widetilde{h}}$. That is, $\mu_{\widetilde{h}}(A) = \mu_h(\psi(A))$ for $A \subset \widetilde{D}$. Similarly, ν_h is almost surely the image under ψ of the measure $\nu_{\widetilde{h}}$ [DS11a]. In fact, [DS11a] formally defines a **quantum surface** to be an equivalence class of pairs (D, h) under the equivalence transformations (see Figure 9.8)

$$(D, h) \to \psi^{-1}(D, h) := (\psi^{-1}(D), h \circ \psi + Q \log |\psi'|) = (\widetilde{D}, \widetilde{h}), \tag{8.3}$$

noting that both area and boundary length are well defined for such surfaces. The invariance of ν_h under (8.3) actually yields a definition of the quantum boundary length measure ν_h when the boundary of D is not piecewise linear—i.e., in this case, one simply maps to the upper half plane (or any other domain with a piecewise linear boundary) and computes the length there.[15]

[14] We use the same distribution composition notation as [DS11a]: i.e., If ϕ is a conformal map from D to a domain $\tilde{D}$ and h is a distribution on D, then we define the pullback $h \circ \phi^{-1}$ of h to be a distribution on $\tilde{D}$ defined by $(h \circ \phi^{-1}, \tilde{\rho}) = (h, \rho)$ whenever $\rho \in H_s(D)$ and $\tilde{\rho} = |\phi'|^{-2} \rho \circ \phi^{-1}$. (Here ϕ' is the complex derivative of ϕ, and (h, ρ) is the value of the distribution h integrated against ρ.) Note that if h is a continuous function (viewed as a distribution via the map $\rho \to \int_D \rho(z) h(z) dz$), then the distribution $h \circ \phi^{-1}$ thus defined is the ordinary composition of h and ϕ^{-1} (viewed as a distribution).

[15] It remains an open question whether the *interior* of a quantum surface is canonically a metric space. A pair (D, h) is a metric space parameterized by D when, for distinct $x, y \in D$ and $\delta > 0$, one defines the distance $d_\delta(x, y)$ to be the smallest number of Euclidean balls in D of μ_h mass δ required to cover *some* continuous path from x to y in D. We conjecture but cannot prove that for some constant β the limiting metric

$$\lim_{\delta \to 0} \delta^\beta d_\delta$$

exists a.s. and is invariant under the transformations described by (8.3).

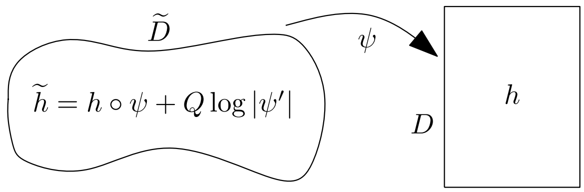

Figure 8.1: A quantum surface coordinate change.

The second construction involves "flow lines" of the unit vector field $e^{ih/\chi}$ where $\chi \neq 0$ is a fixed constant (see Figure 7.1), or alternatively flow lines of $e^{i(h/\chi + c)}$ for a constant $c \in [0, 2\pi)$. The author has proposed calling this collection of flow lines the **AC geometry**[16] of h, but a recent series of works uses the term **imaginary geometry** [MS12a, MS12b, MS12c, MS13a]. Makarov once proposed the term "magnetic gravity" in a lecture, suggesting that in some sense the AC geometry is to Liouville quantum gravity as electromagnetism is to electrostatics.

Although h is a distribution and not a function, one can make sense of flow lines using the couplings between the Schramm-Loewner evolution (SLE) and the GFF in [She05, SS10], which were further developed in [Dub09c] and more recently in [MS09, HBB10a, IK10]. The paths in these couplings are generalizations of the GFF contour lines of [SS10].

We define an **AC surface** to be an equivalence class of pairs under the following variant of (8.3):

$$(D, h) \to (\psi^{-1}(D), h \circ \psi - \chi \arg \psi') = (\widetilde{D}, \widetilde{h}), \tag{8.4}$$

as in Figure 8.2. The reader may observe that (at least when h is smooth) the flow lines of the LHS of (8.4) are the ψ images of the flow lines of the RHS. To check this, first consider the simplest case: if ψ^{-1} is a rotation (i.e., multiplication by a modulus-one complex number), then (8.4) ensures that the unit flow vectors $e^{ih/\chi}$ (as in Figure 7.1) are rotated by the same amount that D is rotated. The general claim follows from this, since every conformal map looks locally like the composition of a dilation and a rotation (see Section 7.1).

Recalling the conformal invariance of the GFF, if the h on the left side of (8.3) and (8.4) is a centered (expectation zero) Gaussian free field on D then the distribution on the right hand side is a centered (expectation zero) GFF on $\widetilde{D}$ *plus* a deterministic function. In other words, *changing the domain of definition is equivalent to recentering the GFF.* The deterministic function is harmonic if D is a planar domain, but it can also be defined (as a non-harmonic function) when D is a surface with curvature (see [DS11a]). In what follows, we will often find it convenient to define quantum and AC surfaces

[16] AC stands for "altimeter-compass." If the graph of h is viewed as a mountainous terrain, then a hiker holding an analog altimeter—with a needle indicating altitude modulo $2\pi\chi$—in one hand and a compass in the other can trace an AC ray by walking at constant speed (continuously changing direction as necessary) in such a way that the two needles always point in the same direction.

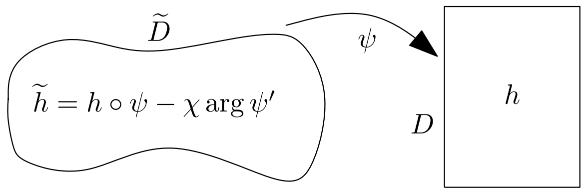

Figure 8.2: An AC surface coordinate change.

on the complex half plane $\mathbb{H}$ using a (free or zero boundary) GFF on $\mathbb{H}$, sometimes recentered by the addition of a deterministic function that we will call $\mathbf{H}_0$. We will state our main results in the introduction for fairly specific choices of $\mathbf{H}_0$.

8.3 Theorem statements: conformal weldings

We will now try to better understand Theorem 7.2 in the special case $\kappa < 4$. Note that a priori the h in Theorem 7.2 is defined only up to additive constant. We can either choose the constant arbitrarily (e.g., by requiring that the mean value of h on some set be zero) or avoid specifying the additive constant and consider the measures μ_h and ν_h to be defined only up to a global multiplicative constant. The choice does not affect the theorem statement below.

Theorem 8.1. *Suppose that $\kappa < 4$ and that h and η_T are coupled in the way described at the end of the previous section, i.e., h is generated by first sampling the B_t process up to time T in order to generate f_T via a reverse Loewner flow, and then choosing $\widetilde{h}$ independently and writing $h = \mathbf{H}_T + \widetilde{h} \circ f_T$, and $\eta_T\big((0,T]\big) = \mathbb{H} \setminus f_T\mathbb{H}$.*[17] *Given a point z along the path η_T, let $z_- < 0 < z_+$ denote the two points in $\mathbf{R}$ that f_T (continuously extended to $\mathbf{R}$) maps to z. Then almost surely*

$$\nu_h([z_-,0]) = \nu_h([0,z_+])$$

for all z on η_T.

Theorem 8.1 is a relatively difficult theorem, and it will be the last thing we prove.We next define $R = R_h : (-\infty, 0] \to [0,\infty)$ so that $\nu_h([x,0]) = \nu_h([0,R(x)])$ for all x (recall that ν is a.s. atom free [DS11a]). This R gives a homeomorphism from $[0_-,0]$ to $[0,0_+]$ that we call a **conformal welding** of these two intervals. We stress that the values 0_- and 0_+ depend on T, but the overall homeomorphism R between $(-\infty,0]$ and $[0,\infty)$ is determined by the boundary measure ν_h, whose law does not depend on T (although

[17]It is not known whether an analog of Theorem 8.1 can obtained in the case $\kappa = 4$. The standard procedure for constructing the boundary measure ν_h breaks down when $\kappa = 4, \gamma = 2$, but a scheme was introduced [DRSV12a, DRSV12b] to create a non-trivial boundary measure ν_h. The open problems listed in [She10] also also address a related question in the $\kappa > 4$ setting.

the coupling between h, $\tilde{h}$, and η_T described in the theorem statement clearly depends on T). Since η_T is simple, it clearly determines the restriction of R to $[0_-, 0]$. (See Figure 8.3.) It turns out that R also determines η_T:

Theorem 8.2. *For $\kappa < 4$, in the setting of Theorem 8.1, the homeomorphism R from $[0_-, 0]$ to $[0, 0_+]$ uniquely determines the curve η_T. In other words, it is almost surely the case that if $\widetilde{\eta}_{\tilde{T}}$ is any other simple curve in $\mathbb{H}$ such that the homeomorphism induced by its reverse Loewner flow is the same as R on $[0_-, 0]$, then $\widetilde{\eta}_{\tilde{T}} = \eta_T$. In particular, h determines η_T almost surely.*

Proof. The author learned from Smirnov that Theorem 8.2 follows almost immediately from Theorem 8.1 together with known results in the literature. If there were a distinct candidate $\widetilde{\eta}_T$ with a corresponding $\widetilde{f}_T$, then $\phi = \widetilde{f}_T \circ f_T^{-1}$ — extended from $\mathbb{H}$ to $\mathbf{R}$ by continuity, and to all of $\mathbf{C}$ by Schwarz reflection — would be a non-trivial homeomorphism of $\mathbf{C}$ (with $\lim_{z\to\infty} \phi(z) - z = 0$) which was conformal on $\mathbf{C} \setminus (\eta_T \cup \bar{\eta}_T)$, where $\bar{\eta}_T$ denotes the complex conjugate of η_T. Thus, to prove Theorem 8.2, it suffices to show that no such map exists. In complex analysis terminology, this is equivalent by definition to showing that the curve $\eta_T \cup \bar{\eta}_T$ is *removable.* Rohde and Schramm showed that the complement of $\eta([0, T])$ is a.s. a Hölder domain for $\kappa < 4$ (see Theorem 5.2 of [**?**]) and that η is a.s. a simple curve in this setting. In particular, $\eta_T \cup \bar{\eta}_T$ is almost surely the boundary of its complement, and this complement is a Hölder domain. (More about Hölder continuity appears in work of Beliaev and Smirnov [BS09a] and Kang [Kan07] and Lind [**?**].) Jones and Smirnov showed generally that boundaries of Hölder domains are removable (Corollary 2 of [JS00a]). The same observations are used in [AJKS09]. □

We remark that the above arguments also show that $\eta \cup \bar{\eta}$ is removable when η is the entire SLE path. In the coming sections, we will often interpret the left and right components of $\mathbb{H} \setminus \eta$ as distinct quantum surfaces, where the right boundary arc of one surface is welded (along η) to the left boundary arc of another surface in a quantum-boundary-length-preserving way. When the law of η is given by SLE_κ with $\kappa < 4$, removability implies that η is almost surely determined (up to a constant rescaling of $\mathbb{H}$) by the way that these boundary arcs are identified. In other words, aside from constant rescalings, there is no homeomorphism of $\mathbb{H}$, fixing 0 and ∞, whose restriction to $\mathbb{H} \setminus \eta$ is conformal.

8.4 Corollary: capacity stationary quantum zipper

This subsection contains some discussion and interpretation of some simple consequences of Theorems 8.1 and 8.2, in particular Corollary 8.3 below. We first observe that for $\kappa < 4$, Theorem 8.2 implies that R determines η_T almost surely for any given $T > 0$. In particular, this means that R determines an entire reverse Loewner evolution $f_t = f_t^h$

for all $t \geq 0$, and that this f_t^h is (in law) a *reverse* SLE_κ flow. Similarly, given a chordal curve η from 0 to ∞ in $\mathbb{H}$, we denote by f_t^η the *forward* Loewner flow corresponding to η. The following is now an immediate corollary of the domain Markov property for SLE and Theorems 7.2, 8.1 and 8.2. As usual, transformations $f(D,h)$ are defined using (8.3).

Corollary 8.3. *Fix $\kappa \in (0,4)$. Let $h = \mathbf{H}_0 + \widetilde{h}$ be as in Theorem 7.2 and let η be an SLE_κ on $\mathbb{H}$ chosen independently of h. Let D_1 be the left component of $\mathbb{H} \setminus \eta$ and h^{D_1} the restriction of h to D_1. Let D_2 be the right component of $\mathbb{H} \setminus \eta$ and h^{D_2} the restriction of h to D_2. For $t \geq 0$, write*

$$\begin{aligned} \mathcal{Z}_t^{\mathrm{CAP}}\big((D_1,h^{D_1}),(D_2,h^{D_2})\big) &= \big(f_t^h(D_1,h^{D_1}), f_t^h(D_2,h^{D_2})\big), \\ \mathcal{Z}_{-t}^{\mathrm{CAP}}\big((D_1,h^{D_1}),(D_2,h^{D_2})\big) &= \big(f_t^\eta(D_1,h^{D_1}), f_t^\eta(D_2,h^{D_2})\big). \end{aligned}$$

Note that both h and η are determined by the pair $\big((D_1,h^{D_1}),(D_2,h^{D_2})\big)$, and that f_t^h and f_t^η are also a.s. determined by this pair, so that the maps $\mathcal{Z}_t^{\mathrm{CAP}}$ and $\mathcal{Z}_{-t}^{\mathrm{CAP}}$ are well defined for almost all pairs $\big((D_1,h^{D_1}),(D_2,h^{D_2})\big)$ chosen in the manner described above. Then the law of $\big((D_1,h^{D_1}),(D_2,h^{D_2})\big)$ is invariant under $\mathcal{Z}_t^{\mathrm{CAP}}$ for all t. Also, for all s and t,

$$\mathcal{Z}_{s+t}^{\mathrm{CAP}} = \mathcal{Z}_s^{\mathrm{CAP}} \mathcal{Z}_t^{\mathrm{CAP}}$$

almost surely.

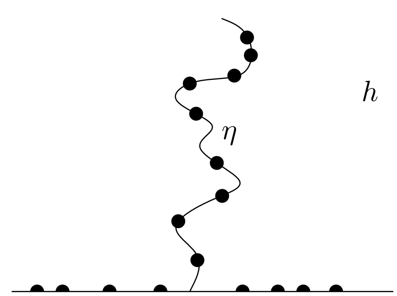

Figure 8.3: Sketch of η with marks spaced at intervals of the same ν_h length along ∂D_1 and ∂D_2. Here $(-\infty, 0]$ and $[0,\infty)$ are the two open strands of the "zipper" while η is the closed (zipped up) strand. Semicircular dots on $\mathbf{R}$ are "zipped together" by f_t^h. Circular dots on η are "pulled apart" by f_t^η. (Recall that under the reverse Loewner flow f_t^h, the center of a semicircle on the negative real axis will reach the origin at the same time as the center of the corresponding semicircle on the positive real axis.) The law of $\big((D_1,h^{D_1}),(D_2,h^{D_2})\big)$ is invariant under "zipping up" by t capacity units or "zipping down" by t capacity units.

Because the forward and reverse Loewner evolutions are parameterized according to half plane capacity, we refer to the group of transformations $\mathcal{Z}_t^{\mathrm{CAP}}$ as the **capacity quantum zipper**, see Figure 8.3. (The term "zipper" in the Loewner evolution context has been used before; see the "zipper algorithm" for numerically computing conformal mappings in [MR07] and the references therein.) When $t > 0$, applying $\mathcal{Z}_t^{\mathrm{CAP}}$ is called "zipping up" the pair of quantum surfaces by t capacity units and applying $\mathcal{Z}_{-t}^{\mathrm{CAP}}$ is called "zipping down" or "unzipping" by t capacity units.

To begin to put this construction in context, we recall that the general *conformal welding* problem is usually formulated in terms of identifying unit discs $\mathbf{D}_1$ and $\mathbf{D}_2$ along their boundaries via a given homeomorphism ϕ from $\partial\mathbf{D}_1$ to $\partial\mathbf{D}_2$ to create a sphere with a conformal structure. Precisely, one wants a simple loop η in the complex sphere, dividing the sphere into two pieces such that if conformal maps ψ_i from the $\mathbf{D}_i$ to the two pieces are extended continuously to their boundaries, then $\psi_1 \circ \psi_2^{-1}$ is ϕ. In general, not every homeomorphism ϕ between disc boundaries is a conformal welding in this way, and when it is, it does not always come from an η that is (modulo conformal automorphisms of the sphere) unique; in fact, arbitrarily small changes to ϕ can lead to large changes in η and some fairly exotic behavior (see e.g. [Bis07]).

The theorems of this paper can also be formulated in terms of a sphere obtained by gluing two discs along their boundaries: in particular, one can zip up the quantum surfaces of Corollary 8.3 "all the way", which could be viewed as welding two Liouville quantum surfaces (each of which is topologically homeomorphic to a disc) to obtain an SLE loop in the sphere, together with an instance of the free boundary GFF on the sphere.

Note that in the construction described above, the quantum surfaces are defined only modulo an additive constant for the GFF, and we construct the two surfaces together in a particular way. In Section 8.5 (Theorem 8.6), we will describe a related construction in which one takes two *independent* quantum surfaces (each with its additive constant well-defined) and welds them together to obtain SLE.

Peter Jones conjectured several years ago that an SLE loop could be obtained by (what in our language amounts to) welding a quantum surface to a deterministic Euclidean disc. (The author first learned of this conjecture during a private conversation with Jones in early 2007 [Jon].) Astala, Jones, Kupiainen, and Saksman recently showed that such a welding exists and determines a unique loop (up to conformal automorphism of the sphere) [AJKS09, AJKS10]. Binder and Smirnov recently announced (to the author, in private communication [Smi]) that they have obtained a proof that the original conjecture of Jones is false. By computing a multifractal spectrum, they showed that the loop constructed in [AJKS09, AJKS10] does not look locally like SLE. However, our construction, together with Theorem 8.6 below, shows that a natural variant of the Jones conjecture — involving two independent quantum surfaces instead of one quantum surface and one Euclidean disc — is in fact true.

We also remark that the "natural" d-dimensional measure on (or parameterization of) an SLE curve of Hausdorff dimension d was only constructed fairly recently [LS09, LZ10, **?**], and it was shown to be uniquely characterized by certain symmetries, in particular the requirement that it transforms like a d-dimensional measure under the maps f_t (i.e., if the map locally stretches space by a factor of r, then it locally increases the measure by a factor of r^d). Our construction here can be viewed as describing, for $\kappa < 4$, a natural "quantum" parameterization of SLE_κ, which is similarly characterized by transformation laws, in particular the requirement that adding C to h — which scales area by a factor of $e^{\gamma C}$ — scales length by a factor of $e^{\gamma C/2}$. These ideas are discussed further in [DS11b].

The relationship between Euclidean and quantum natural fractal measures and their evolution under capacity invariant quantum zipping is developed in [DS11b] in a way that makes use of the KPZ formula [KPZ88b, DS11a].

8.5 Quantum wedges and quantum length stationarity

This subsection contains ideas and definitions that are important for the proofs of Theorem 8.1 and 8.2, as well as the statement of another of this paper's main results, Theorem 8.6, which we will actually prove before Theorem 8.1 in [She10].

Theorem 8.6 includes a variant of Corollary 8.3 in which one parameterizes time by "amount of quantum length zipped up" instead of by capacity. The "stationary" picture will be described as a particular random quantum surface $\mathcal{S}$ with two marked boundary points and a chordal SLE η connecting the two marked points. The theorem will state that this η divides $\mathcal{S}$ into two quantum surfaces $\mathcal{S}_1$ and $\mathcal{S}_2$ that are *independent* of each other. (One can also reverse the procedure and first choose the $\mathcal{S}_i$ — these are the so-called γ-quantum wedges mentioned earlier — and then weld them together to produce $\mathcal{S}$ and the interface η.) As we have already mentioned, this independence appears at first glance to be a rather bizarre coincidence. However, as we will see in Section 8.6, this kind of result is to be expected if SLE-decorated Liouville quantum gravity is (as conjectured) the scaling limit of path-decorated random planar maps.

Before we state Theorem 8.6 formally, we will need to spend a few paragraphs constructing a particular kind of scale invariant random quantum surface that we will call an "α quantum wedge." The reader who has never encountered quantum wedges before before may wish to first read Section 1.4 of [DMS14], which contains a more recent and better illustrated discussion of the quantum wedge construction.

We begin this construction by making a few general remarks. Recall that given any quantum surface represented by $(\widetilde{D}, \widetilde{h})$ — with two distinguished boundary points — we can change coordinates via (8.3) and represent it as the pair $(\mathbb{H}, h)$ for some h, where $\mathbb{H}$ is the upper half plane, and the two marked points are taken to be 0 and ∞. We will represent the "quantum wedges" we construct in this way, and we will focus on constructions in which there is almost surely a finite amount of μ_h mass and ν_h mass in each bounded neighborhood of 0 and an infinite amount in each neighborhood of ∞. In this case, the corresponding quantum surface is **half-plane-like** in the sense that it has one distinguished boundary point "at infinity" and one distinguished "origin" — and each neighborhood of "infinity" includes infinite area and an infinite length portion of the surface boundary, while the complement of such a neighborhood contains only finite area and a finite-length portion of the surface boundary. We will let $\mathcal{S}_h$ denote the doubly marked quantum surface described by h in this way.

The h describing $\mathcal{S}_h$ is canonical except that we still have one free parameter corresponding to constant rescalings of $\mathbb{H}$ by (8.3). For each $a > 0$, such a rescaling is given by

$$(\mathbb{H}, h) \to (\mathbb{H}, h(a\cdot) + Q \log |a|). \tag{8.5}$$

We can fix this parameter by requiring that $\mu_h\big(B_1(0) \cap \mathbb{H}\big) = 1$. We will let μ_h be zero on the negative half plane so that we write this slightly more compactly as $\mu_h\big(B_1(0)\big) = 1$. (Alternatively, one could normalize so that $\nu_h\big([-1, 1]\big) = 1$.) We call the h for which this holds the **canonical description** of the doubly marked quantum surface.

Now to construct a "quantum wedge" it will suffice to give the law of the corresponding h. To this end, we first recall that one can decompose the Hilbert space for the free boundary GFF into an orthogonal sum of the space of functions which are radially symmetric about zero and the space of functions with zero mean about all circles centered at zero [DS11a]. Consequently, we can write $h(\cdot) = h_{|\cdot|}(0) + h^\dagger(\cdot)$, where $h^\dagger_\varepsilon(0) = 0$ for all ε, and $h_{|z|}(0)$ is (of course) a continuous and radially symmetric function of z. This is a decomposition of the GFF h into its projection onto two $(\cdot, \cdot)_\nabla$ orthogonal subspaces, so $h_{|\cdot|}(0)$ and $h^\dagger(\cdot)$ are independent of each other [She07]; the latter is a scale invariant random distribution and defined without an additive constant (since its mean is set to be zero on all circles centered at the origin). Now we define three types of quantum surfaces (the first two being defined only up an additive constant for h, which corresponds to a constant-factor rescaling of the surface itself). The third may seem unmotivated; however, the reader may note that it is similar in the spirit to the second, except that the third h is actually a well defined random distribution (as opposed to a random distribution modulo additive constant), so that $(\mathbb{H}, h)$ is a well-defined quantum surface.

1. **Definition — unscaled quantum wedge on $\mathbb{H}$:** the quantum surface $(\mathbb{H}, h)$ where h is an instance of the free boundary GFF (which is defined up to additive constant, so that the quantum surface is defined only up to rescaling). In this case, $h_{|\cdot|}$ agrees in law with $B_{-\log|\cdot|}$ when B_t, $t \in \mathbf{R}$ is $\sqrt{2}$ times a standard Brownian motion defined up to a global additive constant). We think of B_t as a Brownian motion with diffusive rate 2, which will be understood throughout the discussion below. We can write

$$h = h^\dagger(\cdot) + B_{-\log|\cdot|},$$

where $h^\dagger(\cdot)$ and $B_{-\log|\cdot|}$ are independent.

2. **Definition — α-log-singular free quantum surface on $\mathbb{H}$:** the quantum surface $(\mathbb{H}, h)$ where

$$h = h^\dagger(\cdot) + \alpha\big(-\log|\cdot|\big) + B_{-\log|\cdot|}, \tag{8.6}$$

with $h^\dagger$ and B as above (and h also defined only up to additive constant).

3. **Definition — α-quantum wedge:** for $\alpha < 0$, the quantum surface $(\mathbb{H}, h)$ where

$$h = h^\dagger(\cdot) + Q\big(-\log|\cdot|\big) + A_{-\log|\cdot|}, \tag{8.7}$$

and the process $A_t,\ \ t \in \mathbf{R}$ is defined in a particular way: namely, for $t \geq 0$, A_t is a Brownian motion with drift $\alpha - Q$, i.e., $A_t = B_t + (\alpha - Q)t$. Also, for $t \geq 0$, the negative-time process A_{-t} is chosen independently as a Brownian motion with drift $-(\alpha - Q)$ *conditioned* not to revisit zero. This involves conditioning on a probability zero event, so let us state this another way to be clear. Note that $\tilde{B}_t = B_t - (\alpha - Q)t$ has positive drift and hence a.s. $s_0 = \sup\{s : \tilde{B}_s = 0\} < \infty$. Then the law of A_{-t} (for $t \geq 0$) is the law of $\tilde{B}_{t+s_0}$, for $t \geq 0$.

To begin to motivate the definition above, note that applying the coordinate transformation (8.5) to the α-quantum wedge defined by (8.7), where the coordinate change map is a rescaling by a factor of a, amounts to replacing (8.7) with

$$h^\dagger(a\cdot) + Q\big(-\log|a\cdot|\big) + A_{-\log|a\cdot|} + Q\log|a| = h^\dagger(a\cdot) + Q\big(-\log|\cdot|\big) + A_{\log a - \log|\cdot|}.$$

Since the law of $h^\dagger$ is scaling invariant, we find that the coordinate change described amounts to a horizontal translation of A by $-\log a$. That is, the quantum surface obtained by sampling A and then sampling $h^\dagger$ independently agrees in law with the quantum surface obtained by sampling A, translating the graph of A horizontally by some (possibly random) amount, and then sampling $h^\dagger$ independently.

We think of A_t as a Brownian process that drifts steadily as a Brownian motion with drift $(\alpha - Q)$ from $-\infty$, reaches zero at some point, and then subsequently evolves as a regular Brownian motion with the same drift. Since translating the graph of A_t horizontally doesn't affect the law of the quantum surface obtained, we choose (for concreteness) the translation for which $\inf\{t : A_t = 0\} = 0$. (We remark that the process A_t can also be interpreted as the log of a Bessel process, reparameterized by quadratic variation, noting that the graph of such a reparameterization is *a priori* only defined up to a horizontal translation; this point of view is explained and used extensively in [DMS14].)

Now we make another simple claim: the α-quantum wedge is a doubly marked quantum surface whose law is invariant under the multiplication of its area by a constant. To explain what this means, let us observe that when $C \in \mathbf{R}$, we can "multiply the surface area by the constant e^C" by replacing h with $h + C/\gamma$, or equivalently, by replacing A

with $A + C/\gamma$. Let $t_0 = \inf\{t : \tilde{A}_t = 0\}$ and write $\tilde{A}_t = A_{t_0+t} + C/\gamma$. By the definition of t_0, we find that $\tilde{A}_t$ (like A_t) is a process that drifts up from $-\infty$, reaches zero for the first time when $t = 0$, and then subsequently evolves as a Brownian motion with drift. Indeed, it is not hard to see that $\tilde{A}_t$ has the same law as A_t. To deduce the claim, we then observe that the distribution of $h^\dagger$ is fixed; and since the radial parts $h_{|\cdot|}(0)$ of the GFF are continuous and independent of $\mu_{h^\dagger}$ and converge to a limit in law, we may conclude that $e^{\gamma h_{|\cdot|}(0)} d\mu_{h^\dagger}$ converges in law.

For future reference, we mention that one has a natural notion of "convergence" for quantum surfaces of this type: if $h^1, h^2, \ldots$ are the canonical descriptions of a sequence of doubly marked quantum surfaces and h is the canonical description of $\mathcal{S}_h$, then we say that the sequence $\mathcal{S}_{h^i}$ **converges** to $\mathcal{S}_h$ if the corresponding measures μ_{h^i} converge weakly to μ_h on all bounded subsets of $\mathbb{H}$.

One motivation for the definition of a quantum wedge is the following, which can be deduced from the description of quantum typical points given in Section 6 of [DS11a]. It says (in a certain special setting; for a stronger result) that if one zooms in near a "quantum-boundary-measure-typical" point, one finds that the quantum surface looks like a γ-quantum wedge near that point.

Proposition 8.4. *Fix $\gamma \in [0, 2)$ and let D be a bounded subdomain of $\mathbb{H}$ for which $\partial D \cap \mathbf{R}$ is a segment of positive length. Let $\widetilde{h}$ be an instance of the GFF with zero boundary conditions on $\partial D \setminus \mathbf{R}$ and free boundary conditions on $\partial D \cap \mathbf{R}$. Let $[a, b]$ be any sub-interval of $\partial D \cap \mathbf{R}$ and let $\mathbf{H}_0$ be a continuous function on D that extends continuously to the interval (a, b). Let dh be the law of $\mathbf{H}_0 + \widetilde{h}$, and let $\nu_h[a, b]dh$ denote*

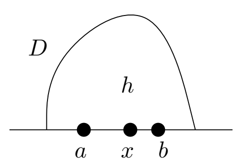

Figure 8.4: Point x sampled from ν_h (restricted to $[a, b]$).

the measure whose Radon-Nikodym derivative w.r.t. dh is $\nu_h[a, b]$. (Assume that this is a finite measure — i.e., the dh expectation of $\nu_h[a, b]$ is finite.) Now suppose we

1. *sample h from $\nu_h[a, b]dh$ (normalized to be a probability measure),*
2. *then sample x uniformly from ν_h restricted to $[a, b]$ (normalized to be a probability measure),*
3. *and then let h^* be h translated by $-x$ units horizontally (i.e., recentered so that x becomes the origin).*

Then as $C \to \infty$ the random quantum surfaces $\mathcal{S}_{h^*+C/\gamma}$ converge in law (w.r.t. the topology of convergence of doubly-marked quantum surfaces) to a γ-quantum wedge.

Proof. We first recall that in this setting the description of quantum typical points in Section 6 of [DS11a] implies a very explicit description of the joint law of the pair x and h sampled in Proposition 8.4. The marginal law of x is absolutely continuous with respect to Lebesgue measure, and conditioned on x the law of h is that of its original law *plus* a deterministic function that has the form $-\gamma \log |x - \cdot|$ plus a deterministic smooth function. In a small neighborhood of x, this deterministic smooth function is approximately constant, which means that h^* looks like (up to additive constant) the h used to define an α-log-singular free quantum surface in (8.6), with $\alpha = \gamma$. If we write $A'_t = B_t + (\alpha - Q)t$, then we find that h^* looks like the h used to define a γ-quantum wedge in (8.7), except with A replaced by A'.

Now replacing h^* by $h^* + C/\gamma$ corresponds to adding C/γ to the process B from (8.6), and hence also corresponds to adding C/γ to the process A', which translates the graph of A' vertically. Recall from above that translating the graph of A' horizontally corresponds to a coordinate change; so we can translate A' so that it hits zero for the first time at the origin. It is not hard to see that as $C \to \infty$, the law of A' thus translated converges to the law of A. Since the law of $h^\dagger$ is scale invariant and can be chosen independently, this implies the proposition statement. □

It is shown in [She10] that the conclusion of the proposition still holds if (when generating x and h) we condition on particular values for $\nu_h[a, x]$ and $\nu_h[x, b]$.

The following is an immediate consequence of Proposition 8.4. It tells us that the γ-quantum wedge is stationary with respect to shifting the origin by a given amount of quantum length. (When $\gamma = 0$, the proposition simply states that $\mathbb{H}$ itself is invariant under horizontal translations. Proposition 8.5 is the general quantum analog of this invariance.)

Proposition 8.5. *Fix a constant $L > 0$. Suppose that $(\mathbb{H}, h)$ is a γ-quantum wedge. Then choose $y > 0$ so that $\nu_h[0, y] = L$, and let h^* be h translated by $-y$ units horizontally (i.e., recentered so that y becomes the origin). Then $(\mathbb{H}, h^*)$ is a γ-quantum wedge.*

Proof. Suppose that x is the point chosen uniformly from the quantum boundary measure in Proposition 8.4, and x' is the point translated δL quantum length units to the right from x, so that $\nu_h[x, x'] = \delta L$. Note that such an x' exists with a probability that tends to 1 as $\delta \to 0$, and that the law of x' converges (in total variation sense) to the law of x as $\delta \to 0$. In the rescaled surfaces in Proposition 8.4, boundary lengths are scaled by $e^{C/2}$, so if we set $\delta = e^{-C/2}$, then the distance between x and x' is L after the rescaling. Since this δ tends to zero as $C \to \infty$ we conclude that the limiting surface law is (as desired) invariant under the operation that translates the origin by L units of quantum boundary length. □

Theorem 8.6. Wedge decomposition: *Fix $\gamma \in (0, 2)$, and let $\mathcal{S}$ be a $(\gamma - 2/\gamma)$-quantum wedge with canonical description h. Let η be a chordal SLE_κ in $\mathbb{H}$ from 0* *to ∞, with $\kappa = \gamma^2$, chosen independently of h. Let D_1 and D_2 be the left and right components of $\mathbb{H} \setminus \eta$, and let h^{D_1} and h^{D_2} be the restrictions of h to these domains. Then the quantum surfaces represented by (D_1, h^{D_1}) and (D_2, h^{D_2}) are* independent *γ-quantum wedges (marked at 0 and ∞), and their quantum boundary lengths along η agree.*

Zipper stationarity: *Moreover, suppose we define*

$$\mathcal{Z}_{-t}^{\mathrm{LEN}}\big((D_1, h^{D_1}), (D_2, h^{D_2})\big)$$

as follows. First find z on η for which the quantum boundary lengths along D_1 and D_2 (which are well defined by unzipping) along η between 0 and z are both equal to t. Let t' be the time that η hits z (when η is parameterized by capacity) and define

$$\mathcal{Z}_{-t}^{\mathrm{LEN}}\big((D_1, h^{D_1}), (D_2, h^{D_2})\big) = \text{ rescaling of } \big(f_{t'}^{\eta}(D_1, h^{D_1}), f_{t'}^{\eta}(D_2, h^{D_2})\big),$$

where the rescaling is done via (8.5) with the parameter a chosen so that $B_1(0)$ has area one in the transformed quantum measure. Then the following hold:

1. *The inverse $\mathcal{Z}_t^{\mathrm{LEN}}$ of the operation $\mathcal{Z}_{-t}^{\mathrm{LEN}}$ is a.s. uniquely defined (via conformal welding).*
2. *$\mathcal{Z}_{s+t}^{\mathrm{LEN}} = \mathcal{Z}_s^{\mathrm{LEN}} \mathcal{Z}_t^{\mathrm{LEN}}$ almost surely for $s, t \in \mathbf{R}$.*
3. *The law of the pair $\big((D_1, h^{D_1}), (D_2, h^{D_2})\big)$ is invariant under $\mathcal{Z}_t^{\mathrm{LEN}}$ for all $t \in \mathbf{R}$.*

It also follows from Theorem 8.2 and the subsequent discussion that the two independent γ-quantum wedges uniquely determine h and η almost surely. We refer to the group of transformations $\mathcal{Z}_t^{\mathrm{LEN}}$ as the **length quantum zipper**. When $t > 0$, applying $\mathcal{Z}_t^{\mathrm{LEN}}$ is called "zipping up" the pair of quantum surfaces by t quantum length units and applying $\mathcal{Z}_{-t}^{\mathrm{LEN}}$ is called "zipping down" or "unzipping" by t quantum length units. When we defined the operations $\mathcal{Z}_t^{\mathrm{CAP}}$, h was defined only up to additive constant, and the zipping maps f_t were independent of that constant. By contrast, $\mathcal{Z}_t^{\mathrm{LEN}}$ represents zipping by an actual quantity of quantum length and hence cannot be defined without the additive constant being fixed.

In these notes we present just a brief overview of the proof and the relationship to our other results. We will start with the scenario described in Figure 8.3, with h normalized to have mean zero on $\partial B_1(0)$, except that the measure dh on h is replaced by the probability measure whose Radon-Nikodym derivative w.r.t. dh is $\nu_h(-\delta, 0)$ for some fixed δ (see Figure 8.5).

Then we will sample x from ν_h restricted to $(-\delta, 0)$ (normalized to be a probability measure) and "zip up" until x hits the origin (to obtain a "quantum-length-typical" configuration). We then zoom in near the origin (multiplying the area by $\tilde{\varepsilon}^{-1}$ — and hence the boundary length by $\tilde{\varepsilon}^{-1/2}$ — say). We then use a variant of Proposition 8.4

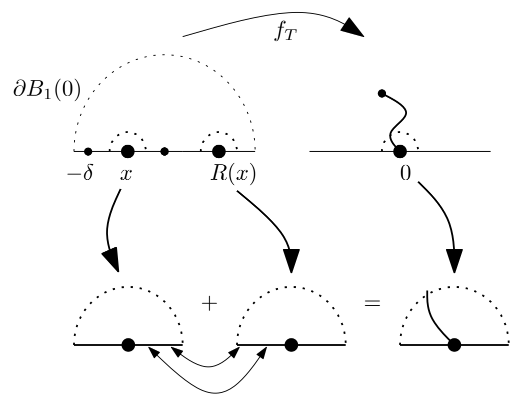

Figure 8.5: Choose h as in Theorem 7.2 (normalized by $h_1(0) = 0$) except with the law of h *weighted* by $\nu_h\big([-\delta, 0]\big)$ for some fixed $\delta \in (0, 1)$. Then conditioned on h, sample x from ν_h restricted to $[-\delta, 0]$ (normalized to be a probability measure). Take T so that f_T is the map zipping up $[x, 0]$ with $[0, R(x)]$. Consider the three random surfaces obtained by choosing a semi-disc of quantum area $\tilde{\varepsilon}$ centered at each of x and $R(x)$ (on the left side) and 0 (on the right side), and multiplying areas by $1/\tilde{\varepsilon}$ (zooming in) so that all three balls have quantum area 1. In the $\tilde{\varepsilon} \to 0$ limit, the left two quantum surfaces become independent γ-quantum wedges, and the right is the conformal welding of these two.

to show that (in the $\tilde{\varepsilon} \to 0$ limit) the lower two rescaled surfaces on the lower left of Figure 8.5 become independent γ-quantum wedges.

The fact that the curve on the right in Figure 8.5 is (in the $\tilde{\varepsilon} \to 0$ limit) an SLE_κ independent of the canonical description h on the right is shown in [She10] by directly calculating the law of the process that "zips up" $[x, 0]$ with $[0, R(x)]$. It could also be seen by showing that we can construct an equivalent pair of glued surfaces by beginning with Figure 8.3 (with h normalized to have mean zero on $\partial B_1(0)$) and then zipping *down* by a random amount (chosen uniformly from an interval) of quantum length, then zooming in by multiplying lengths by $1/\tilde{\varepsilon}$, and then taking the $\tilde{\varepsilon} \to 0$ limiting law. (In this case, the domain Markov property of the original SLE, and its independence from the original GFF, would imply that the conditional law of the still-zipped portion of the curve is an SLE_κ, independent of h.)

Similar arguments to those in [DS11a] will show that the procedure in Figure 8.5 produces a configuration related to the one in Figure 7.5 except that it is in some sense *weighted* by the amount of quantum mass near zero. It will turn out that this weighting effectively adds $-\gamma \log|\cdot|$ to the $\mathbf{H}_0$ of Theorem 7.2 and Corollary 8.3. This is why Theorem 8.6 involves a $(\gamma - 2/\gamma)$-quantum wedge, instead of a $(-2/\gamma)$-quantum wedge, as one might initially guess based on Theorem 7.2. Once we have all of this structure in place, the really crucial step will be showing that parameterizing time by the amount of "left boundary quantum length" zipped up yields the same stationary picture as parameterizing time the amount of "right boundary quantum length" zipped up. Given this, we will then use the ergodic theorem to show that over the long term, the amount of left bounday quantum length zipped up approximately agrees with the amount of right boundary length zipped up. Using scale invariance symmetries we will then deduce that this agreement almost surely holds exactly on all scales.

8.6 Reverse coupling: planar maps and scaling limits

In this section, we conjecture a connection between path-decorated planar maps and SLE-decorated Liouville quantum gravity (in particular, the quantum-length-invariant decorated quantum wedge of Theorem 8.6). We will explain the details in just one example based on the uniform spanning tree. (Variants based on Ising and $O(n)$ and FK models on random planar maps — or on random planar maps without additional decoration besides the chordal paths — are also possible. Many rigorous results for percolation and the Ising model have been obtained for deterministic graphs in [Smi01b, Smi05, Smi06, Smi07, CN08, CS09] (and in many other papers we will not survey here), and one could hope to extend these results to random graphs. One could also consider discrete random surfaces decorated by loops and in the continuum replace SLE decorations with CLE decorations [She09b].) As mentioned earlier, we will see that the more surprising elements of Theorem 8.6 are actually quite natural from the discrete random surface point of view.

Let G be a planar map with exactly n edges (except that each edge on the outer face is counted as *half* an edge) and let T be a subgraph consisting of a single boundary cycle, a chordal path from one boundary vertex a to another boundary vertex b that otherwise does not hit the boundary cycle, and a spanning forest rooted at this "figure 8" structure. (See Figure 8.6.) Here T is like the wired spanning tree (in which the entire boundary is considered to be one vertex), except that there is also one chord connecting a pair of boundary vertices. What happens if we consider the uniform measure on all pairs (G, T) of this type? This model is fairly well understood combinatorially (tree-rooted maps on the sphere are in bijective correspondence with certain walks in $\mathbf{Z}^2$ — see, e.g., [Mul67b] as well as [Ber07] and the references therein — and our model is a simple variant of this) and in particular, it follows from these bijections that the length of the boundary of the outer face of this map will be of order $\sqrt{n}$ with high probability when n is large. Now, can we understand the scaling limit of the random pair (G, T) as $n \to \infty$?

There are various ways to pose this problem. For example, one could consider G as a metric space and aim for convergence in law w.r.t. the Gromov-Hausdorff metric on metric spaces. The reader is probably aware that there is a sizable literature on the realization of a random metric space called the Brownian map as a Gromov-Hausdorff scaling limit of random planar maps of various types. However, since this paper is

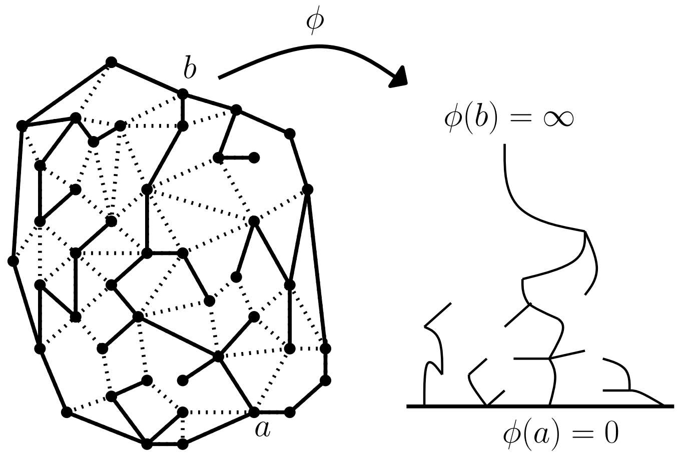

Figure 8.6: Planar map with a distinguished outer-boundary-plus-one-chord-rooted spanning tree (solid black edges), with chord joining marked boundary points a and b, plus image of tree under conformally uniformizing map ϕ to $\mathbb{H}$ (sketch).

concerned with the conformal structure of random geometr, we will try to phrase the the problem in a way that keeps track of that structure.

First, we would like to understand how to conformally map the planar map to the half plane, as in Figure 8.6. We may consider G as embedded in a two-dimensional manifold with boundary in various ways, one of which we sketch here: first add an interior vertex to each face of G and an edge joining it to each vertex of that face (as in Figure 8.7). Each interior edge of G is now part of a quadrilateral (containing one vertex for each interior face of G and one for each vertex of G) and we will endow that quadrilateral with the metric of a unit square $[0,1] \times [0,1]$. Similarly, the triangle containing an exterior edge of G is endowed with the metric of half a unit square (split on its diagonal, with the exterior edge as the hypotenuse). When two squares or half squares share an edge, the points along that edge are identified with one another in a length preserving way. We may view the collection of (whole and half) unit squares, glued together along boundaries, as a manifold (with isolated conical singularities at vertices whose number of incident squares is not four) with a uniquely defined conformal structure (note that it is trivial to define a Brownian motion on the manifold, since it a.s. never hits the singularities). We may choose a conformal map ϕ from this manifold to $\mathbb{H}$, sending a to 0 and b to ∞, as sketched in Figure 8.6.

This ϕ is determined only up to scaling, but we can fix the scaling in many ways. We will do so by considering a number $k < n$ and requiring that the area of $\phi^{-1}(B_1(0))$ be equal to k. Then ϕ determines a random measure on $\mathbb{H}$ (the image of the area measure on the manifold) in which the measure of $B_1(0)$ is deterministically equal to k; let $\mu_{n,k}$ denote this random measure divided by k, so that $\mu_{n,k}(B_1(0)) = 1$. We expect that if one lets n and k tend to ∞ in such a way that n/k tends to ∞, then the random

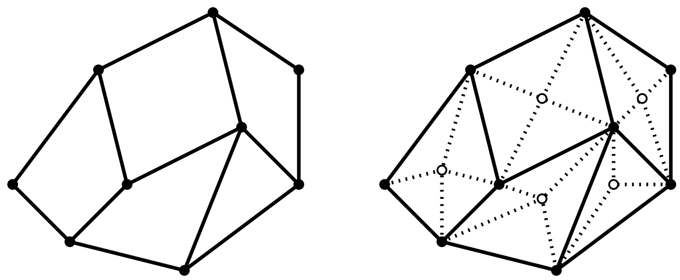

Figure 8.7: An arbitrary planar map can be used to construct a collection of stitched-together unit squares and half unit squares. The result is viewed as a two dimensional manifold with boundary.

measures $\mu_{n,k}$ will converge in law with respect to the metric of weak convergence on bounded subsets of $\mathbb{H}$ to the $\mu = \mu_h$ corresponding to the canonical description h of the $(\gamma - 2/\gamma)$-quantum wedge of Theorem 8.6. (By compactness, the laws of the $\mu_{n,k}$ restricted to the closure of $B_1(0)$ have at least a subsequential limit.) We similarly conjecture that $\nu_{n,k}$ — defined to be $1/\sqrt{k}$ times the image of the manifold's boundary measure — will converge in law to the corresponding ν_h. (We remark that one could alternatively formulate the conjecture by taking an infinite volume limit first — i.e., letting n go to infinity while keeping k constant to define a limiting measure $\mu_{\infty,k} := \lim_{n\to\infty} \mu_{n,k}$. This kind of infinite volume limit of random planar maps was constructed in [AS03]. One can subsequently take $k \to \infty$ and conjecture that the limit is μ_h. A similar conjecture in [DS11a] was formulated in terms of infinite volume limits.)

We are currently unable to prove these conjectures, but related questions about Brownian motion on random surfaces have been explored in [GR], where it was shown that certain infinite random triangulations and quadrangulations (*without* boundaries) are parabolic (as opposed to hyperbolic) Riemann surfaces [GR]. (This is equivalent to showing that a Brownian motion visits each face infinitely often almost surely; see analogous discrete results in [AS03].)

Now let us make some more observations. If we take k, n, and n/k to be large and condition on G, a, and b, then what is the conditional law of $\phi(T)$, as depicted in Figure 8.6? The conditional law of T itself is uniform among all valid 8-rooted spanning forest configurations. The physics literature frequently invokes a kind of "conformal invariance Ansatz" which suggests that this random path (and many other random sets in critical two dimensional statistical physics) should be a conformally invariant object.

In this case, we claim that the law of the chordal path should be approximately that of a chordal SLE_2 *even after we have conditioned on G, a, and b, which determine the measure $\mu_{n,k}$*. The reason for our claim is that a related SLE_2 convergence result is obtained in [LSW04c] in the case that G is a deterministic lattice graph, and this was generalized substantially in [YY08] where it was shown that if a graph can be embedded in the plane in such a way that simple random walk approximates Brownian motion, then the uniform spanning tree paths approximate a form of SLE_2. We do not know whether the hypotheses of [YY08] hold in our setting. Brownian motion is conformally invariant, but it is not clear whether simple random walk on our random G approximates Brownian motion on the corresponding quadrangulated manifold with high probability. However, it seems very natural to conjecture that the hypotheses hold. In any case, we stress the following: if our scaling limit conjecture holds, then the asymptotic independence of the chordal path from $\mu_{n,k}$ would be consistent with the independence of η and h in Theorem 8.6.

Next let D_1 and D_2 be the wired-spanning-tree decorated manifolds to the left and right of the chordal path. Note that once we condition on the length of the chordal path in (G, T) and the number of edges on each side of it, the laws of D_1 and D_2 are independent of one another. We might guess that the local behavior of D_1 and D_2 near a would be approximately independent of these global numbers. We expect a similar property to hold in the scaling limit, which would be consistent with the independence of the left and right quantum surfaces described in Theorem 8.6. (The idea of gluing together independent discrete surfaces in this manner has been explored in many works by Duplantier and others, beginning perhaps in [DK88]. The idea of gluing a whole series of discrete surfaces was used in [**?**] to heuristically derive certain "cascade relations" via the KPZ formula.)

Finally, if we condition on the point b and on D_1 and D_2, then the length of the path along which D_1 and D_2 are glued to each other is uniform among all possibilities (which range between 1 and the minimum M of the boundary lengths of the two D_i's minus 1). In other words, once D_1 and D_2 and b are all fixed, we can randomly decide how far to "zip up" or "unzip" these two surfaces (moving the vertex a accordingly). If r is the random number of steps we zip, then r and $r + m$ have approximately the same law (as long as m/M is small). We expect a similar property to hold in the scaling limit, which would be consistent with the quantum-length-zipper invariance described in Theorem 8.6.

8.6.1 Alternative underlying geometries and $\mathrm{SLE}_{\kappa,\rho}$

Both Theorems 7.1 and 7.2 can be generalized to other values of $\mathbf{H}_0$ using the so-called $\mathrm{SLE}_{\kappa,\rho}$ processes. (As discussed at the end of Section 8.2, changing $\mathbf{H}_0$ can be interpreted as changing the underlying geometry on which Liouville quantum gravity is defined.) We generalize the latter here (see [Dub09c] for the former). We take $G = G^{\mathbb{H}_F}$ in the following. Slightly abusing notation, we will consider situations where $\rho(y)dy$ represents a general signed measure (instead of requiring that ρ be a smooth test function). In this case $(F, \rho) = \int F(y)\rho(y)dy$ represents integration of F w.r.t. this measure.

Theorem 8.7. *Fix $\kappa > 0$ and a signed measure $\rho(y)dy$ on $\overline{\mathbb{H}}$ with finite positive and finite negative mass supported on some closed $\mathcal{C} \subset \mathbb{H}$. Write*

$$\hat{\mathbf{H}}_t(z) = \mathbf{H}_t(z) + \frac{1}{2\sqrt{\kappa}} \int G_t(y,z)\rho(y)dy, \tag{8.8}$$

where $\mathbf{H}_t(z)$ is as in Theorem 7.2, and let $\widetilde{h}$ be an instance of the free boundary GFF on $\mathbb{H}$, independent of B_t. Let η_T be the segment generated by the reverse Loewner flow

$$df_t(z) = \frac{-2}{f_t(z)}dt - dW_t, \tag{8.9}$$

where

$$dW_t = \Big(\int \mathrm{Re}\, \frac{-1}{f_t(y)}\rho(y)dy\Big)dt + \sqrt{\kappa}dB_t = \big(-\mathrm{Re}\,(f_t)^{-1}, \rho\big)dt + \sqrt{\kappa}dB_t, \tag{8.10}$$

up to any stopping time $T \geq 0$ at or before the smallest t for which $0 \in f_t(\mathcal{C})$ (here f_t is extended continuously from $\mathbb{H}$ to $\overline{\mathbb{H}}$). Then the following two random distributions (modulo additive constants) on $\mathbb{H}$ agree in law: $h = \hat{\mathbf{H}}_0 + \widetilde{h}$ and $\hat{\mathbf{H}}_T + \widetilde{h} \circ f_T$.

We will make several observations before we prove Theorem 8.7. First, if ρ is supported on a set of n points $y_1, \ldots, y_n$ in $\mathbb{H}$, with masses given by real numbers $\rho_1, \rho_2, \ldots, \rho_n$, then the process defined by (8.9) and (8.10) is (the reverse form of) what is commonly called an $\mathrm{SLE}_{\kappa,\rho}$ process in the literature: in this case, (8.10) takes the form

$$dW_t = \sum_{i=1}^{n} \mathrm{Re}\, \frac{-\rho_i}{f_t(y_i)}dt + \sqrt{\kappa}dB_t, \tag{8.11}$$

which is the same expression one finds in the usual definition of the forward-flow $\mathrm{SLE}_{\kappa,\rho}$ process. (Note that the notation here differs from [**?**], since here we use ρ to denote the measure, not the vector of mass values ρ_i.) In the special case that ρ is supported at a single point $x \in \mathbf{R}$, with mass ρ_1 we find that

$$df_t(x) = \frac{-2}{f_t(x)}dt - dW_t = \frac{-2+\rho_1}{f_t(x)}dt - \sqrt{\kappa}dB_t, \tag{8.12}$$

so that $f_t(x)/\sqrt{\kappa}$ is a Bessel process of dimension δ satisfying $(\delta-1)/2 = (\rho_1 - 2)/\kappa$, i.e.,

$$\delta = 1 + \frac{2(\rho_1 - 2)}{\kappa}. \tag{8.13}$$

For this and future discussion it will be useful to recall a few standard facts:

1. The Bessel process X_t of dimension δ by definition satisfies $dX_t = dB_t + \frac{\delta-1}{2}X_t^{-1}dt$. Hence $d\log X_t = \frac{1}{X_t}dB_t + \frac{\delta-1}{2X_t^2}dt - \frac{1}{2X_t^2}dt$. The process $\log X_t$, when parameterized by its quadratic variation, is a Brownian motion with a constant drift of magnitude $\frac{\delta-2}{2}$.

2. If X_t is a Bessel process of dimension δ started at $X_0 = x$ and run until the first time T that it reaches zero, then the time reversal X_{T-t} has the law of a Bessel process of dimension δ', started at zero and run until the last time that it hits x, where δ' is the dimension one gets by changing the sign of the drift in $\log X_t$. That is, $\frac{\delta-2}{2} = -\frac{\delta'-2}{2}$, so that
$$\delta = 4 - \delta'. \tag{8.14}$$

3. In the usual forward flow definition of $\mathrm{SLE}_{\kappa,\rho_1'}$ the function $f_t(x)$ is a Bessel process of dimension
$$\delta' = 1 + \frac{2(\rho_1' + 2)}{\kappa}. \tag{8.15}$$
The reason for the difference from (8.13) can be seen by considering the case $\rho_1 = \rho_1' = 0$. In the reverse process, the Loewner drift is pulling $f_t(x)$ toward the origin, while in the forward process the Loewner drift is pushing $f_t(x)$ away from the origin. In both cases ρ_1 (or ρ_1') indicates a quantity of additional force pushing $f_t(x)$ away from the origin.

4. Combining (8.13), (8.14), and (8.15) gives a relationship between ρ_1 and ρ_1'. Namely, $1 + \frac{2(\rho_1'+2)}{\kappa} = 4 - (1 + \frac{2(\rho_1-2)}{\kappa})$, so that
$$\rho_1' = \kappa - \rho_1. \tag{8.16}$$
This means that if we run a reverse $\mathrm{SLE}_{\kappa,\rho_1'}$ until the time T at which $f_t(x)$ hits zero, then f_T maps $\mathbb{H}$ to $\mathbb{H} \setminus \eta_T$ where η_T has the law of an initial segment of a forward $\mathrm{SLE}_{\kappa,\rho_1}$. In particular, if $\rho_1' = \kappa$, then η_T has the law of an ordinary SLE stopped at a time T (which corresponds to the last time that a Bessel process hits a certain value). This will be important later.

Recall from Section 8.2 that changing $\mathbf{H}_0$ to $\hat{\mathbf{H}}_0$ can be interpreted as changing the underlying geometry on which Liouville quantum gravity is defined. Moreover, ρ is proportional to $-\Delta(\hat{\mathbf{H}}_0 - \mathbf{H}_0)$, and $-\Delta\hat{\mathbf{H}}_0$ is proportional to the overall Gaussian curvature density (see the appendix).

We will give a formal proof of Theorem 8.7 below using Itô/ calculus, but first let us offer an informal explanation of why the result is true. The idea behind Theorem 8.7 is to interpret (8.8) as the expectation of h in a certain weighted measure and (8.10) as the description of the law of W_t in that measure. This is easiest to understand when we first switch coordinates using the correspondence shown in Figure 7.7. Suppose first that ρ is such that (3.6) is finite with $\rho_1 = \rho_2 = \rho$ and that the total integral of ρ is zero. If dh is the law of a (centered or not centered) GFF then the standard Gaussian complete-the-square argument shows that $e^{(h,\rho)}dh = e^{(h,-2\pi\Delta^{-1}\rho)_\nabla}dh$ (normalized to be a probability measure) is the law of the standard GFF plus $-2\pi\Delta^{-1}\rho$. When we weight the law of the Brownian motion in Figure 7.6 by $e^{\alpha(h,\rho)}$ for some constant α (note that (h,ρ) is the terminal value that the Brownian motion in that figure reaches at time zero) this is equivalent to adding a constant drift term to the Brownian motion (parameterized by $-E_t(\rho)$) in Figure 7.6.

We take $\alpha = \frac{1}{2\sqrt{\kappa}}$ and weight by

$$e^{(h, \frac{1}{2\sqrt{\kappa}}\rho)}, \tag{8.17}$$

which, as explained above, modifies the law in a way that amounts to adding the drift term of $\frac{1}{2\sqrt{\kappa}}\big(E_0(\rho) - E_t(\rho)\big)$ to the Brownian motion in Figure 7.6. Recalling the correspondence shown in Figure 7.7, the fact that the left figure is a Brownian motion with this constant drift (up to a stopping time) completely determines the law of W_t up to that stopping time. Indeed, recalling (7.9), we find that the law of W_t is necessarily the one described by (8.10).

We have now related the weighted measure to (8.10), but what does this have to do with (8.8)? Observe that

$$(\hat{\mathbf{H}}_t, \rho) = (\mathbf{H}_t, \rho) + \frac{1}{2\sqrt{\kappa}} E_t(\rho)$$

represents the conditional expectation (in the weighted measure) of (h, ρ) given $B_\cdot$ up to time t. In fact, by the standard complete-the-square argument, the function $\hat{\mathbf{H}}_t$ in Theorem 8.7 represents the conditional expectation (in the weighted measure) of h, given $B_\cdot$ up to time t, and is thus a martingale in t. The above construction (and a bit of thought) actually constitutes a proof of Theorem 8.7 when (3.6) is finite and the total integral of ρ is zero.

The argument above can be adapted to more general ρ. If the total integral of ρ is not zero, we may modify ρ by adding some mass very far from the origin, so that the total integral becomes zero but the drift in (8.10) does not change very much. If (3.6) is infinite, we may be able to modify it to make it finite: for example, if ρ is a point mass, then we may replace it with a uniform measure on a tiny ball centered at that point mass, and the harmonicity of $G_t(\cdot, z)$ in (8.8) and of $\mathrm{Re}\,(f_t)^{-1}$ in (8.10) show that (outside of this small ball) neither (8.8) nor (8.10) is affected by this replacement. We will present a more direct Itô/ calculation below, which also applies when (3.6) is infinite.

Proof of Theorem 8.7. We will follow the calculations of Theorem 7.2 and check where differences appear. First, we find the following:

REVERSE FLOW SLE
$df_t(z) = \frac{-2}{f_t(z)}dt + \left(\int \mathrm{Re}\,\frac{1}{f_t(y)}\rho(y)dy\right)dt - \sqrt{\kappa}dB_t$
$d\log f_t(z) = \frac{-(4+\kappa)}{2f_t(z)^2}dt + f_t(z)^{-1}\left(\int \mathrm{Re}\,\frac{1}{f_t(y)}\rho(y)dy\right)dt - \frac{\sqrt{\kappa}}{f_t(z)}dB_t$
$df_t'(z) = \frac{2f_t'(z)}{f_t(z)^2}dt$
$d\log f_t'(z) = \frac{2}{f_t(z)^2}dt$

Also, as before, we compute

$$dG_t(y,z) = -\mathrm{Re}\,\frac{2}{f_t(y)}\mathrm{Re}\,\frac{2}{f_t(z)}dt,$$

and recall that

$$\hat{\mathbf{H}}_t(z) := \frac{2}{\sqrt{\kappa}}\log|f_t(z)| + Q\log|f_t'(z)| + \frac{1}{2\sqrt{\kappa}}\int G_t(y,z)\rho(y)dy.$$

We then find that when computing $d\hat{\mathbf{H}}_t(z)$ the extra term in $d\frac{2}{\sqrt{\kappa}}\mathrm{Re}\,\log f_t(z)$ cancels the term $d\frac{1}{2\sqrt{\kappa}}\int G_t(y,z)\rho(y)dy$ so that $d\hat{\mathbf{H}}_t(z) = \mathrm{Re}\,\frac{-2}{f_t(z)}dB_t$, just as in the proof of Theorem 7.2. The remaining calculations are the same as in the proof of Theorem 7.2. $\square$

We next remark, in the context of Theorem 8.7, that if $(\mathbb{H}, \hat{\mathbf{H}}_T + \widetilde{h}\circ f_T)$ is a quantum surface, then

$$f_T(\mathbb{H}, \hat{\mathbf{H}}_T + \widetilde{h}\circ f_T) = f_T(\mathbb{H}, \mathbf{H}_T + \frac{1}{2\sqrt{\kappa}}\int G_T(y,\cdot)\rho(y)dy + \widetilde{h}\circ f_T), \tag{8.18}$$

and this can be written

$$(\mathbb{H}\setminus K_T, \mathbf{H}_0 + \frac{1}{2\sqrt{\kappa}}\int G(f_T(y),\cdot)\rho(y)dy + \widetilde{h}). \tag{8.19}$$

When $\kappa < 4$, this suggests the following interpretation of Theorem 8.7. We start with $\hat{\mathbf{H}}_0 + \widetilde{h}$, which is actually only defined up to additive constant, so it determines a quantum surface up to a multiplicative constant. We zip up this (modulo multiplicative constant) quantum surface until a stopping time T, and condition on the zipper map f_T. Then the conditional law of the new zipped-up quantum surface (which is also defined only up to a multiplicative constant) is the same as the original law except that ρ is replaced by the f_T image of ρ. We have not fully established that this interpretation is correct, because we have not yet shown that $\hat{\mathbf{H}}_0 + \widetilde{h}$ uniquely determines f_T. As in Theorem 7.2, we have only shown that sampling h from $\hat{\mathbf{H}}_0 + \widetilde{h}$ is equivalent to first zipping up according to a given law, then sampling the field from a putative conditional law in the zipped up picture, and then unzipping.

8.7 Welding more general quantum wedges

Additional welding constructions are described in [DMS14]. These allow one to weld together two wedges of weights W_1 and W_2 to produce a new wedge of weight $W_1 + W_2$. One can also weld the left and right sides of a single quantum wedge to each other, to produce a quantum cone.

9 Mating trees and the peanosphere

9.1 Overview

The idea of topologically mating Julia sets is given an overview in [Mil04, Mil06] and the references therein. In this section we will discuss how to glue to two continuum random trees produces a topological sphere decorated by a space-filling path. This section closely follows and adapts the presentation in [DMS14].

We describe the connections between the Schramm-Loewner evolution (SLE) [Sch00a] and the "random surfaces" associated with Liouville quantum gravity (LQG). We derive an extensive collection of results about weldings and decompositions of surfaces of various types (spheres, disks, wedges, chains of disks, trees of disks, etc.) along with fundamental structural results about LQG. These results comprise a sort of "calculus of random surfaces" that combines the conformal welding theory of [She10] with the imaginary geometry theory of [MS12a, MS12b, MS12c, MS13a].

However, we begin with a construction that can be described without any *a priori* reference to SLE or LQG. Section 9.2 explains how to "mate" a pair of continuum random trees (CRTs) to produce what we will call a *peanosphere*. A peanosphere is a topological sphere endowed with a measure and a space-filling path in a certain random way. We then pose the question of how to put an explicit conformal structure on the peanosphere — i.e., how to explicitly embed the peanosphere in the complex plane $\mathbf{C} \cup \{\infty\}$. This paper shows that an instance of the peanosphere can a.s. be canonically embedded in the plane in a particular way, and that the embedded peanosphere has the law of an LQG quantum surface parameterized by $\mathbf{C} \cup \{\infty\}$ and decorated by a space-filling form of SLE. We will ultimately introduce a correlation parameter to the peanosphere construction that corresponds to the parameters $\gamma \in (0, 2)$, $\kappa = \gamma^2 \in (0, 4)$, and $\kappa' = 16/\kappa \in (4, \infty)$ that appear in the LQG and SLE constructions. Once the parameter is fixed, our result produces a one-to-one correspondence between instances of the peanosphere and instances of SLE-decorated LQG (defined for almost all instances).

One reason that this result is intriguing is that it allows one to convert many questions about SLE and LQG into questions about the peanosphere that can be asked without reference to conformal structure. In principle, this could allow future researchers to derive properties of SLE, LQG, and related structures in papers that never mention either

Loewner evolution or the Gaussian free field. For example, the major open problem of endowing LQG with a metric space structure for general γ (see the discussion in [MS13b]) might turn out to be more easily addressed from the peanosphere perspective than from the conventional LQG perspective.

Another such open problem (which we state but will not solve here) is the following. Consider the graph whose vertices are the components of the complement of a chordal (non-space-filling) $\mathrm{SLE}_{\kappa'}$ trace with $\kappa' \in (4,8)$, where two such components are adjacent if their boundaries intersect; is this graph a.s. connected? We will explain in Section 9.8 how topological properties of $\mathrm{SLE}_{\kappa'}$ (along with some geometric properties like quantum path length) are directly encoded by the CRTs used to construct the peanosphere, or by certain stable Lévy processes derived from these CRTs. This encoding reduces problems like the one just stated to questions about stable Lévy processes that can be asked without reference to SLE.

We also stress that, as explained in [She09c], there is a simple procedure for constructing a conformal loop ensemble (CLE) from space-filling $\mathrm{SLE}_{\kappa'}$ when $\kappa' \in (4,8)$, and that the procedure itself is topological, making no use of conformal structure. Hence the peanosphere encodes the CLE loops in a straightforward way, and many properties of CLE (e.g., quantum dimensions of different types of special points, quantum lengths of loops and areas of enclosed regions, properties of the graph whose vertices are loops with two loops adjacent when they intersect, etc.) could also be addressed directly from the peanosphere construction.

Another reason for interest in the peanosphere is discussed in Section 9.3, which explains at a high level how, in light of related recent work, our result completes a solution to one form of the problem of showing that CLE-decorated LQG is a scaling limit of random cluster-decorated planar maps. Essentially, the work in [She11b] shows that certain trees associated to the latter scale to the pair of CRTs that this paper associates to the former, a fact that can be interpreted as scaling limit convergence within a topology native to the peanosphere.[18]

Update: As mentioned above this paper suggests that one natural way to study discrete statistical physics models on a planar maps (which we will discuss in Section 9.3) is try

[18]Our scaling limit theorem does not by itself resolve the scaling limit conjectures of [DS11a, She10] (which require convergence in a topology that encodes conformal structure), because it does not address how the discrete models are embedded in the plane. There are a number of natural ways to embed the discrete models "conformally" in the plane; Figures 9.4 and 9.5 illustrate one way to do this using circle packings. Another approach is to view the faces of the planar map as regular unit-side-length polygons, stitched together along their boundaries to form a "quilt," and then to conformally map the quilt onto the plane. Proving that the curves in these embeddings are close to SLE curves appears to require a deeper understanding of the embeddings themselves, including the establishment of some technical lemmas that seem out of reach at the moment. However, Nicolas Curien, and Stanislav Smirnov and Roman Boikii, have announced some progress in this direction in private communication. See [DMS14] for discussion about how one might try to "strengthen the topology" of the convergence theorem in this paper to establish the conjectures of [DS11a, She10] in a different way.

to identify a natural pair of trees induced by those models, and then use a combinatorial understanding of the trees to prove a relationship to SLE/LQG. Since this paper was posted to the arXiv, this approach has been successfully applied by a number of authors to several well known combinatorial objects that had not previously been studied in the conformal probability context, including bipolar orientations and Schnyder woods as well as certain generalizations of the FK cluster model; see [KMSW15, GKMW16] (combined with [GHMS15]) and work in progress by Li and Sun, along with additional work in progress by Bernardi, Li and Sun. The new results allow us to add these models to the pantheon of canonical SLE/LQG-related models, which already includes loop-erased random walk, percolation, the Ising model, the uniform spanning tree, and the Gaussian free field. In particular, these are the first elements of the pantheon corresponding to $\kappa < 2$ and $\kappa' > 8$. We expect that these ideas will in time lead to additional bridges between combinatorics and conformal probability — in part because the "peanosphere convergence" results for random planar maps (as discussed in Section 9.3) are often more accessible than the analogous results for deterministic lattices, and can be established without a precise understanding of the discrete conformal structure.

A third reason for interest in these results is that they turn out to have consequences for the quantum zipper, as described in [She10], for $\mathrm{SLE}_{\kappa'}$ with $\kappa' \in (4,8)$. We do not prove that $\mathrm{SLE}_{\kappa'}$ is removable (an issue raised as an open question in [She10]), but we obtain another way of showing that when $\kappa' \in (4,8)$ the embedding of an $\mathrm{SLE}_{\kappa'}$ in the plane is a.s. determined by the two "quantum disk trees" it cuts out, one on its left and one on its right, which resolves other open problems from [She10, Section 6, Question 8]. This implies that the "welding" or "conformal mating" of the pair of trees is uniquely determined by the trees themselves. (See Figure 9.6.)

A fourth reason is that we intend to heavily use the matings mentioned just above, along with several other results from this paper, in forthcoming work about the quantum Loewner evolution (QLE), as introduced in [MS13b], and its relationship to the Brownian map. Both QLE and several approximations to QLE have analogs that can be constructed directly on the peanosphere itself. It will also turn out that the time-reversals of some of the QLE processes described in [MS13b] (the ones constructed using $\mathrm{SLE}_{\kappa'}$ for $\kappa' \in (4,8)$) can be reformulated as ways to construct spheres by gluing trees of disks (more precisely, so-called *Lévy trees* of disks, of the sort described in Figure 9.6 and 9.7) to themselves. The Brownian map itself will also be obtained as a sort of "reshuffling" of the entire SLE_6 exploration tree associated to the peanosphere.

A fifth and more speculative reason to study the peanosphere is that developing objects like SLE, CLE and LQG in a framework that does not explicitly reference conformal structure may help us figure out how to construct interesting analogs of these objects in higher dimensions, where the tools of planar conformal geometry (such as the Riemann mapping theorem) no longer apply.

The remainder of this section is structured as follows. In Section 9.2, we describe how it is possible to construct a topological sphere by gluing a pair of trees to each other.

Section 9.3 explains how this article fits into a larger program to relate the scaling limits of random planar maps to LQG. Section 9.4 describes a variant of the construction of Section 9.2 to the setting of gluing trees of disks. In Section 9.5, we give an overview of the different types of quantum surfaces which will arise in this article. Section 9.7 gives our main results about welding quantum cones and wedges. Section 9.8 then describes our main results regarding different types of matings.

9.2 Constructing a topological sphere from a pair of trees

Let us begin with something we can do easily. We will construct a topological sphere as a quotient of a pair of continuum random trees. This sphere will turn out to be naturally endowed with a measure and a space-filling curve that represents (intuitively) the interface between the two identified trees.

Let X_t and Y_t be independent Brownian excursions, both indexed by $t \in [0,T]$. Thus $X_0 = X_T = 0$ and $X_t > 0$ for $t \in (0,T)$ (and similarly for Y_t). Once X_t and Y_t are chosen, choose C large enough so that the graphs of X_t and $C - Y_t$ do not intersect. (The precise value of C does not matter.) Let $R = [0,T] \times [0,C]$, viewed as a Euclidean metric space.

Let $\cong$ denote the smallest equivalence relation on R that makes two points equivalent if they lie on the same vertical line segment with endpoints on the graphs of X_t and $C - Y_t$, or they lie on the same horizontal line segment that never goes above the graph of X_t (or never goes below the graph of $C - Y_t$). Maximal segments of this type are shown in Figure 9.1.

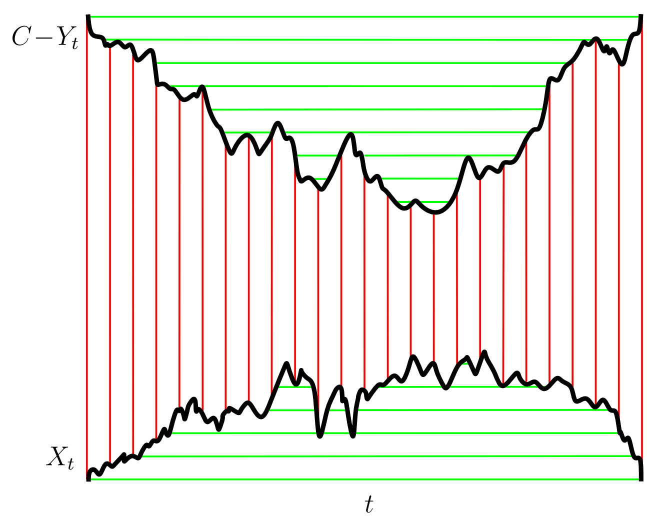

Figure 9.1: Points on the same vertical (or horizontal) line segment are equivalent.

Observe that the vertical line segment corresponding to time $t \in [0, T]$ shares its lower (resp. upper) endpoint with a horizontal line segment if and only if $X_s > X_t$ (resp. $Y_s > Y_t$) for all s in a neighborhood of the form $(t, t+\epsilon)$ or $(t-\epsilon, t)$ for $\epsilon > 0$. It is easy to see that (aside from the $t = 0$ and $t = T$ segments) there almost surely exists no vertical line segment that shares both an upper and a lower endpoint with a horizontal line segment. [19] It is also easy to see that each of X_t and Y_t a.s. has at most countably many local minima, and that almost surely the values obtained by X_t or Y_t at these minima are distinct. It follows that a.s. each of the maximal horizontal segments of Figure 9.1 intersects the graph of X_t or $C - Y_t$ in two or three places (the two endpoints plus at most one additional point).

Thus the equivalence classes a.s. all have one of the following types:

Type 0: The outer boundary rectangle ∂R.

Type 1: A single vertical segment that does not share an endpoint with a horizontal segment.

Type 2: A single maximal horizontal segment beneath the graph of X_t or above the graph of $C - Y_t$, together with the two vertical segments with which it has an endpoint in common.

Type 3: A single maximal horizontal segment beneath the graph of X_t or above the graph of $C - Y_t$, together with the two vertical segments with which it has an endpoint in common, and one additional vertical segment with an endpoint in the interior of the horizontal segment.

Proposition 9.1. *The relation $\cong$ is a.s.* topologically closed. *That is, if $x_j \to x$, and $y_j \to y$, and $x_j \cong y_j$ for all j, then $x \cong y$.*

Proof. By passing to a subsequence, one may assume that the type of the equivalence class of $x_j \cong y_j$ is the same for all j. The equivalence class can be described by a finite collection of numbers (the endpoint coordinates for the line segments), and compactness implies that we can find a subsequence along which this collection of coordinates converges to a limit. We then observe that these limiting numbers describe an equivalence class (or a subset of an equivalence class) comprised of zero or one horizontal lines that lie either below the graph of X_t or above the graph of $C - Y_t$, and one or more incident vertical lines that lie between these graphs. Since x and y necessarily belong to this limiting equivalence class, we have $x \cong y$. □

Let $\widetilde{R}$ be the topological quotient $R/\cong$. Then we have the following:

[19] It can be shown the set of record minima of a one dimensional Brownian motion agrees in law with the zero set of a one dimensional Brownian motion; the fact that no such minima occur simultaneously is related to the fact that planar Brownian motion a.s. does not hit any specific point. More general statements, related to so-called "cone times," appear in [Eva85] and are discussed in [DMS14].

Proposition 9.2. *The space $\widetilde{R}$ is a.s. topologically equivalent to the sphere $\mathbf{S}^2$. The map ϕ that sends t to the equivalence class containing (t, X_t) and $(t, C - Y_t)$ is a continuous surjective map from $[0, T]$ to $\widetilde{R}$.*

Proof. Let R' be obtained from R by identifying the outer boundary of R with a single point. Then R' is topologically a sphere, and $\cong$ induces a topologically closed relation on R'. The fact that $\widetilde{R}$ is topologically a sphere is immediate from the properties observed above and Proposition 9.3, stated just below. The continuity of ϕ is immediate from the fact that the quotient map from R to $\widetilde{R}$ is continuous. $\square$

The following was established by R.L. Moore in 1925 [Moo25] (the formulation below is lifted from [Mil04]):

Proposition 9.3. *Let $\cong$ be any topologically closed equivalence relationship on the sphere $\mathbf{S}^2$. Assume that each equivalence class is connected, but is not the entire sphere. Then the quotient space $\mathbf{S}^2/\cong$ is itself homeomorphic to $\mathbf{S}^2$ if and only if no equivalence class separates the sphere into two or more connected components.*

If R_1 is the set of points in R on or under the graph of X, then $\widetilde{R}_1 = R_1/\cong$ is a random metric space called a *continuum random tree*, constructed and studied by Aldous in the early 1990's [Ald91a, Ald91b, Ald93]. Similarly, if R_2 is the set of points on or above the graph of $C - Y$, then $\widetilde{R}_2/\cong$ is another continuum random tree. Thus, $\widetilde{R}$ can be understood as a topological quotient of this pair of metric trees. We can define a measure ν on $\widetilde{R}$ by letting $\nu(A)$ be the Lebesgue measure of $\phi^{-1}(A)$. Observe that for Lebesgue almost all times t, the value $\phi^{-1}\big(\phi(t)\big)$ consists of the single point t. Thus ν is supported on points that are hit by the space-filling path exactly once. The set of double points (i.e., points in $\widetilde{R}$ hit twice by the path) a.s. has ν measure zero but is uncountable. The set of triple points (i.e. points in $\widetilde{R}$ hit three times by the path) a.s. has ν measure zero and is countable.

We next claim that given $\widetilde{R}$, ν, and ϕ (with the latter given only modulo monotone reparameterization), it is a.s. possible to reconstruct X_t and Y_t. We will not need this fact, but we nonetheless sketch the proof here. Observe first that we recover the parameterization of ϕ by requiring $\nu(\phi([0, t])) = t$. Second, given any t_0, the union of equivalence classes in R corresponding to $\phi([0, t_0])$ is given by the union of $[0, t_0] \times [0, C]$ and

1. all horizontal line segments with exactly one endpoint in $[0, t_0] \times [0, C]$.
2. all vertical segments that share an endpoint with one of these horizontal segments.

One may check that the complement of this set is connected and open, which implies that $\widetilde{R} \setminus \phi([0, t_0])$ is connected and open, hence homeomorphic to a disk. Moreover, the boundary of $\phi([0, t_0])$ consists of two arcs, separated from each other by $\phi(0)$ and $\phi(t_0)$, each corresponding to a branch of one of the two trees; the lengths of these arcs, which are given by X_t and Y_t, can then be determined as the (1/2)-Minkowski measure of the set of times $t > t_0$ at which these arcs are visited; this is related to the well known fact that the set of record times of a Brownian motion (i.e., the times at which a new positive value is reached for the first time) agrees in law with the set of zero times of a Brownian motion, and both of these sets have a natural *local time* structure (much more about Brownian motion local times can be found in the reference text [RY99b]; see also [IM74, Section 2.2, 5) and 6)]).

Now observe the following:

Proposition 9.4. *The measure ν described above is a.s. non-atomic (i.e., assigns zero mass to each individual element of the sphere) and assigns positive measure to each non-empty open set of the sphere.*

Proof. If $x \in \widetilde{R}$ then $\phi^{-1}(x)$ consists of one, two or three points in $[0, T]$, hence has Lebesgue measure zero. If A is a non-empty open subset of $\widetilde{R}$, then since ϕ is continuous, we have that $\phi^{-1}(A)$ is a non-empty open subset of $[0, T]$, so the Lebesgue measure of $\phi^{-1}(A)$ is non-zero. $\square$

A non-atomic measure on $\mathbf{S}^2$ that assigns positive mass to each open set is sometimes called a *good measure*. By the so-called von Neumann-Oxtoby-Ulam theorem, derived in 1941 [OU41] (and an extension in [Fat80], see also [MP83]), any topological sphere endowed with a good measure is isomorphic to any other topological sphere endowed with a good measure.

Proposition 9.5. *Suppose that μ and ν are two good measures on $\mathbf{S}^2$, both with total mass T. Then there exists a homeomorphism of $\mathbf{S}^2$ w.r.t. which the pullback of μ is ν.*

Let $\mathcal{G}_T$ denote the canonical (up to measure-preserving homeomorphism) good-measure endowed topological sphere. Proposition 9.5 implies that $\widetilde{R}$ is (as a good-measure endowed sphere) equivalent to $\mathcal{G}_T$. What makes $\widetilde{R}$ more interesting is that it comes with additional structure: namely, the space-filling path ϕ, which we call an *exploration structure* on $\mathcal{G}_T$. The pair $(\mathcal{G}_T, \phi)$ is what we call a *peanosphere*[20].

[20] This term emerged during a private discussion with Richard Kenyon, who also proposed that the term *peanofold* be used for variants of the peanosphere that are homeomorphic to the disk, or to higher genus surfaces. The term peanofold could also apply to analogous constructions in dimensions higher than two. It is an open question whether the construction presented in this section can be generalized to produce a three dimensional manifold as a topological quotient of three trees. Moore's theorem does not apply as generally as one might expect in higher dimensions (see [HR12] for a recent survey of what is and is not known to be true), and of course the Riemann mapping theorem fails in higher dimensions. However, the von Neumann-Oxtoby-Ulam theorem remains valid for compact manifolds in any dimension [Fat80].

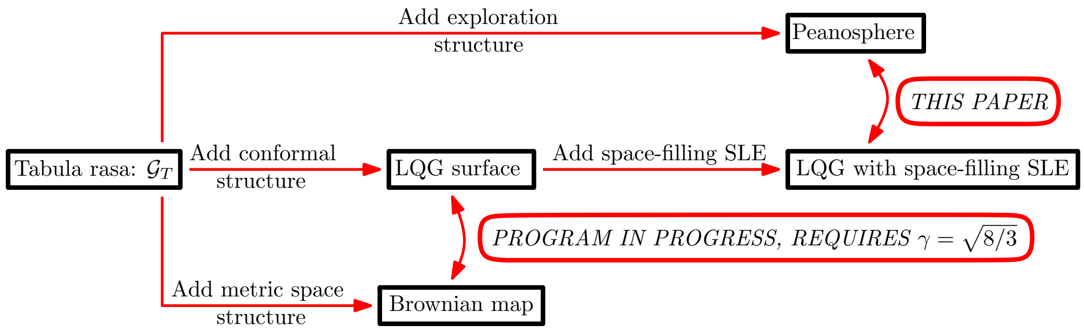

Figure 9.2: The two boxes on the right contain more information than the two in the middle column. As explained in Section 9.3, the objects on the right should be scaling limits of discrete models *decorated* by random statistical physics structure (such as an Ising spin or percolation cluster or spanning tree configuration) on top of the random planar map. The two objects in the center column (LQG surface and Brownian map) do not encode this extra decoration. The program to connect LQG surfaces (with $\gamma^2 = 8/3$) to the Brownian map has been partially begun and explained in [MS13b], has been advanced in [MS15c, MS15a, MS15b] and will be completed in [MS15d, MS15e].

Now let us make a couple of points about the role of the peanosphere in this paper. First, because it turns out to be simpler, we will work initially with a straightforward "infinite volume" variant of the peanosphere discussed above in which X_t and Y_t are Brownian motions indexed by all of $\mathbf{R}$.[21]

Second, throughout the remainder of this paper, we will actually work on a more general case in which (X_t, Y_t) is an affine transformation of a standard two dimensional Brownian motion such that

1. $X_t + Y_t$ is a standard Brownian motion,
2. $X_t - Y_t$ is an independent *constant multiple* of a standard Brownian motion.

The constant in question encodes how *correlated* the Brownian processes X_t and Y_t are with one another.[22] We remark that, as one may observe from Figure 9.1, replacing

[21] In this case, X_t and Y_t describe infinite-diameter CRTs, and the quotient is taken the same way as before. To clarify the construction, we note that we can draw it inside a finite rectangle as before if we rescale both time and space; if ϕ is a monotone continuous map from $\mathbf{R}$ onto $(0,1)$, we can replace X_t and Y_t with $\widetilde{X}_t = \phi(X_{\phi^{-1}(t)})$ and $\widetilde{Y}_t = \widetilde{X}_t = \phi(X_{\phi^{-1}(t)})$, and then use $\widetilde{X}_t$ and $\widetilde{Y}_t$ in the construction described by Figure 9.1 to generate $\widetilde{R}$. The outer boundary of R again forms a single equivalence class, which corresponds to a "point at ∞" in the sense that every neighborhood of that point has infinite measure. It is straightforward to extend Proposition 9.2 to this setting.

[22] The curious reader may try to repeat our proof of Proposition 9.2 in the setting where X_t and Y_t are correlated in this way. The key step is to show that, regardless of correlation constant, there is always a.s. some upper bound on the number of vertical line segments that can belong to a single $\cong$

X_t and Y_t with aX_t and bY_t, where a and b are positive constants, does not actually affect the construction. Thus, if one assumes that (X_t, Y_t) is an affine transformation of a standard two dimensional Brownian motion, the correlation constant is the only parameter that is really relevant. Theorem 9.13 gives a relationship between this constant and the parameter $\kappa' = 16/\gamma^2$, a relationship proved to hold at least in the range $\kappa' \in (4, 8]$.

Using this generality, Figure 9.2 gives a sense of what will be accomplished in this paper, and how the corresponding peanospheres fit into the context of other objects the reader may have studied.[23] As suggested in Figure 9.2, we will show in this paper that an exploration structure on $\mathcal{G}_T$ actually determines a conformal structure on $\mathcal{G}_T$ in a natural way. Another way to say this is that we give a canonical way of defining a Brownian diffusion on the peanosphere. For us, "endowing $\widetilde{R}$ with a conformal structure" will mean explicitly embedding $\widetilde{R}$ in the complex sphere $\mathbf{C} \cup \{\infty\}$ and describing the corresponding measure and space-filling path on $\mathbf{C}$. The measure turns out to be Liouville quantum gravity with parameter $\gamma \in (0, 2)$ and the space-filling path is a form of the Schramm-Loewner evolution (SLE) with parameter $16/\gamma^2$. The parameter γ depends on the level of correlation between X_t and Y_t, and is $\sqrt{2}$ in the case of independent CRTs discussed above.

Finally, let us stress that we will actually derive the upward direction of the double arrow in Figure 9.2 before the downward direction. That is, we proceed by studying SLE/LQG until we are able to show that an SLE/LQG pair encodes a pair of trees, each of which comes with a natural metric structure; then checking that the joint law of the pair of trees is the same as the law of the pair of CRTs discussed above; and only then showing *a posteriori* that the pair of CRTs almost surely determines the SLE/LQG structure in this construction.

equivalence class. Each of the countable set of equivalence classes that include local minima of X_t or Y_t involves exactly three vertical lines a.s., and other equivalence classes are alternating chains of vertical and horizontal line segments. One can use a "cone time" analysis following [Eva85] to derive an a.s. upper bound on the lengths of these chains. We will not give details of this argument here. Instead we note that this more general version of Proposition 9.2 will also follow as a consequence of our explicit embedding construction.

[23] We have *not* endowed a peanosphere with a non-trivial metric space structure, because we have not ruled out the possibility that the metric space quotient of $R/\cong$ is trivial (so that all points are distance zero from each other). Indeed, we believe that the metric space quotient actually *is* trivial but will not prove this here. The Brownian map is actually constructed in much the same way we constructed $\widetilde{R}$ above, except that while X_t is still taken to be a Brownian excursion, Y_t is defined using a Brownian motion indexed by the CRT that X_t describes. This guarantees that for every equivalence class containing a horizontal line segment below X_t, the points at the upper endpoints of the vertical lines in the equivalence class (points on the graph of $C - Y_t$) have the same y coordinates. In this case, one can start with the metric on the upper tree, take its quotient with respect to the induced equivalence on that tree, and note that the branches in the upper tree continue to be geodesics in the new quotient space (because the added equivalences introduce no shortcuts). Thus the quotient metric is a non-trivial metric on the corresponding topological quotient; this metric generates the spherical topology [LG07a] and the two trees obtained correspond to the geodesic tree of the metric (rooted at a distinguished point in the sphere) and its planar dual.

9.3 Liouville quantum gravity as a scaling limit

In this section, we say more about how the result we began to describe in Section 9.2 fits into the context of Liouville quantum gravity, conformal loop ensembles, and random planar maps. As illustrated in Figure 9.3, this paper can be understood, in some sense, as a culmination of a multi-paper project aimed at establishing a connection between

1. random planar maps "decorated" by FK distinguished edge subsets (and the loops forming cluster/dual-cluster interfaces), and
2. Liouville quantum gravity (LQG) random surfaces "decorated" by independently sampled conformal loop ensembles (CLE).

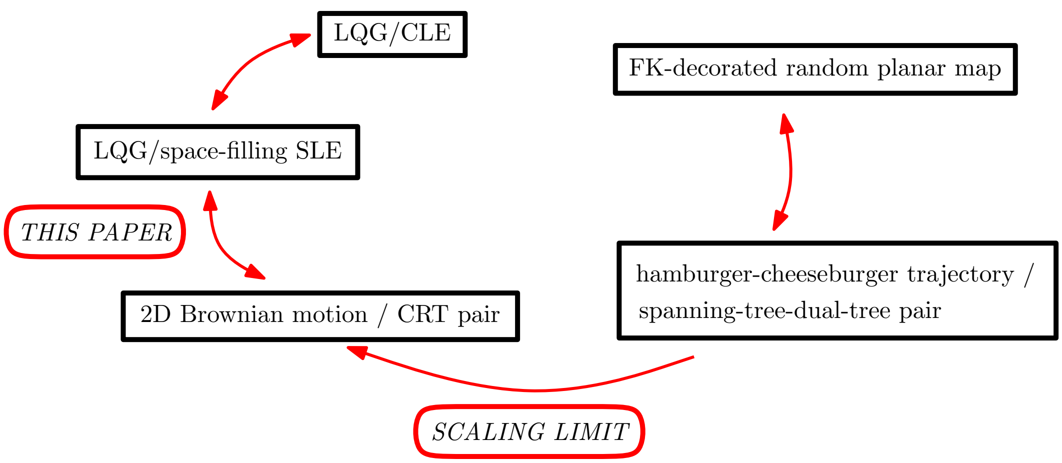

Figure 9.3: This paper focuses on the double arrow indicated above. The LQG/space-filling SLE box represents a space-filling $\mathrm{SLE}_{\kappa'}$ drawn on top of a γ-Liouville quantum gravity surface, where $\gamma \in (0,2)$, $\kappa' = 16/\kappa$, and $\kappa = \gamma^2$. The 2D Brownian motion/CRT pair box represents the coupled pair of CRTs described in Section 9.2.

This paper will focus narrowly on the double arrow of Figure 9.3 as the figure label indicates, and as was discussed in Section 9.2. We will not discuss the other arrows in Figure 9.3 outside of a few high level contextual remarks and references, which form the remainder of Section 9.3. The reader who prefers to focus on what is accomplished in this paper may skip these remarks and proceed directly to Section 9.4.

1. The double arrow on the right in Figure 9.3 is explained in detail in [She11b]. In that paper, one considers a rooted planar map M with n edges, together with a distinguished subset E of those edges associated to the self-dual FK random cluster model with parameter q. The interfaces between clusters and dual clusters are understood as loops. Given an (M,E) pair, there is a canonical space-filling

path that "explores" the loops, and also represents the interface between a special spanning tree T_1 and a special dual spanning tree T_2. The structure of these trees (and the way that they are "glued" together) is encoded by a walk in $\mathbf{Z}^2$, which is related in [She11b] to a two-product LIFO (last in first out) inventory management ("hamburger-cheeseburger") model. The trees T_1 and T_2 arising in a computer simulation of this construction (together with the "space-filling loop" in between them) appear in Figures 9.4 and 9.5. These resources/wedges/figures/ also illustrate a method of embedding the planar maps in the disk "conformally" using circle packing computed with [Ste10].

2. The "scaling limit" arrow in Figure 9.3 is also explained in [She11b], where it is shown that the above mentioned walk in $\mathbf{Z}^2$ scales to a two-dimensional Brownian motion, with a diffusion rate depending on q.

3. The upper left double arrow in Figure 9.3 is explained in more detail in [She09c] and [MS13a]. These works show that there is a space-filling variant of SLE_κ that "explores" the entire CLE_κ loop ensemble, and which may be understood as a continuum analog of the loop exploration path discussed in [She11b]. In this context, the analog of the pair (T_1, T_2) discussed above is a pair $(\mathcal{T}_1, \mathcal{T}_2)$ of space-filling trees, which can be interpreted as flow-line and dual-flow-line trees within an "imaginary geometry" based on a Gaussian free field.

More references to the physics literature and discussion of the overall project are also found in [She11b, She09c, MS13a]. As discussed in [She11b], the combination of the results summarized in Figure 9.3 implies that certain structures encoded by the discrete models, namely the trees (T_1, T_2), have scaling limits that agree in law with the analogous structures encoded by CLE-decorated LQG. Another way to state this fact is to say that CLE-decorated LQG is the scaling limit of FK-weighted random planar maps in a topology where two loop-decorated metric surfaces are considered close when their associated tree/dual-tree pairs are close as measure-endowed metric spaces.[24]

9.4 Gluing trees of disks

Although we do not give any details in this subsection, we mention that we will also present another variant of the construction from Section 9.2 in which the CRTs represented by Figure 9.1 are replaced by "trees of quantum disks" that correspond

[24] Readers familiar with SLE scaling limit results will recall that every statement of the form "these measures on discretized random paths converge to this measure on continuum random paths" is made w.r.t. to some topology on the set of measures. A standard approach is to begin with a topology on the set of paths themselves, and then consider the corresponding weak topology on the corresponding space of Borel measures. The choice of topology determines what properties of the paths we keep track of: whether the paths are treated as unparameterized paths, or natural-time-parameterized paths, or sets, or Loewner driving functions, etc.

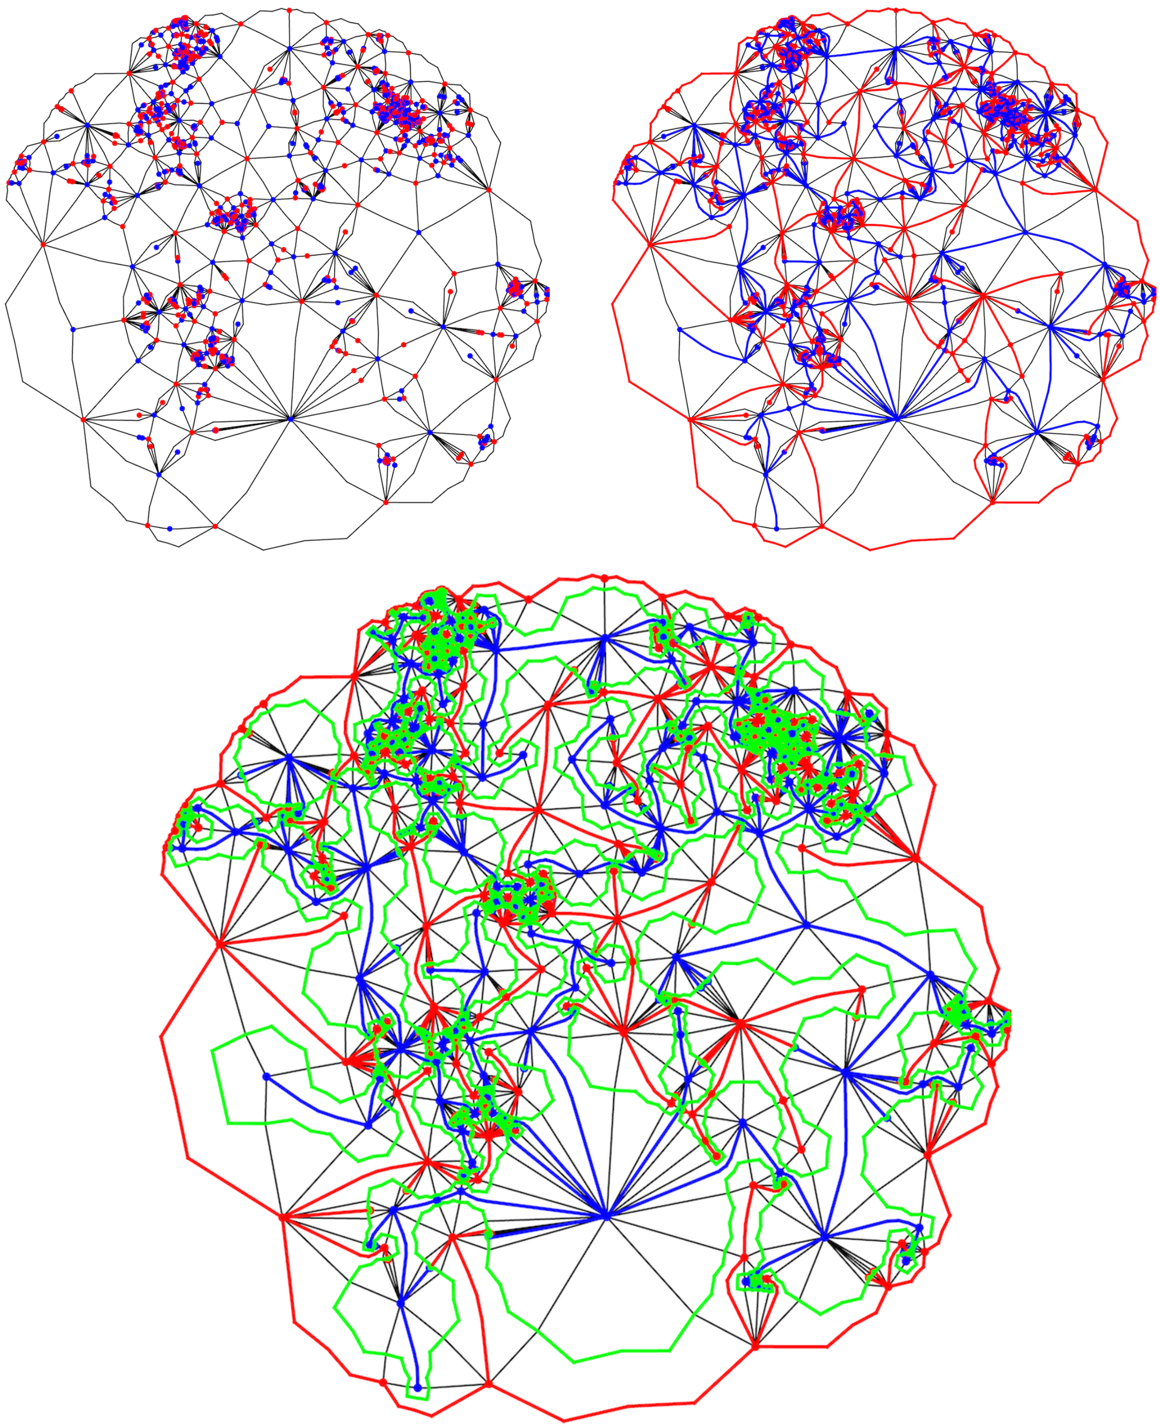

Figure 9.4: **Upper left:** random planar map sampled using the bijection from [She11b] with $p = 0$ and embedded into $\mathbf{C}$ using [Ste10]. **Upper right:** same map with distinguished tree/dual tree pair. **Bottom:** map with path which snakes between the trees and visits all of the edges.

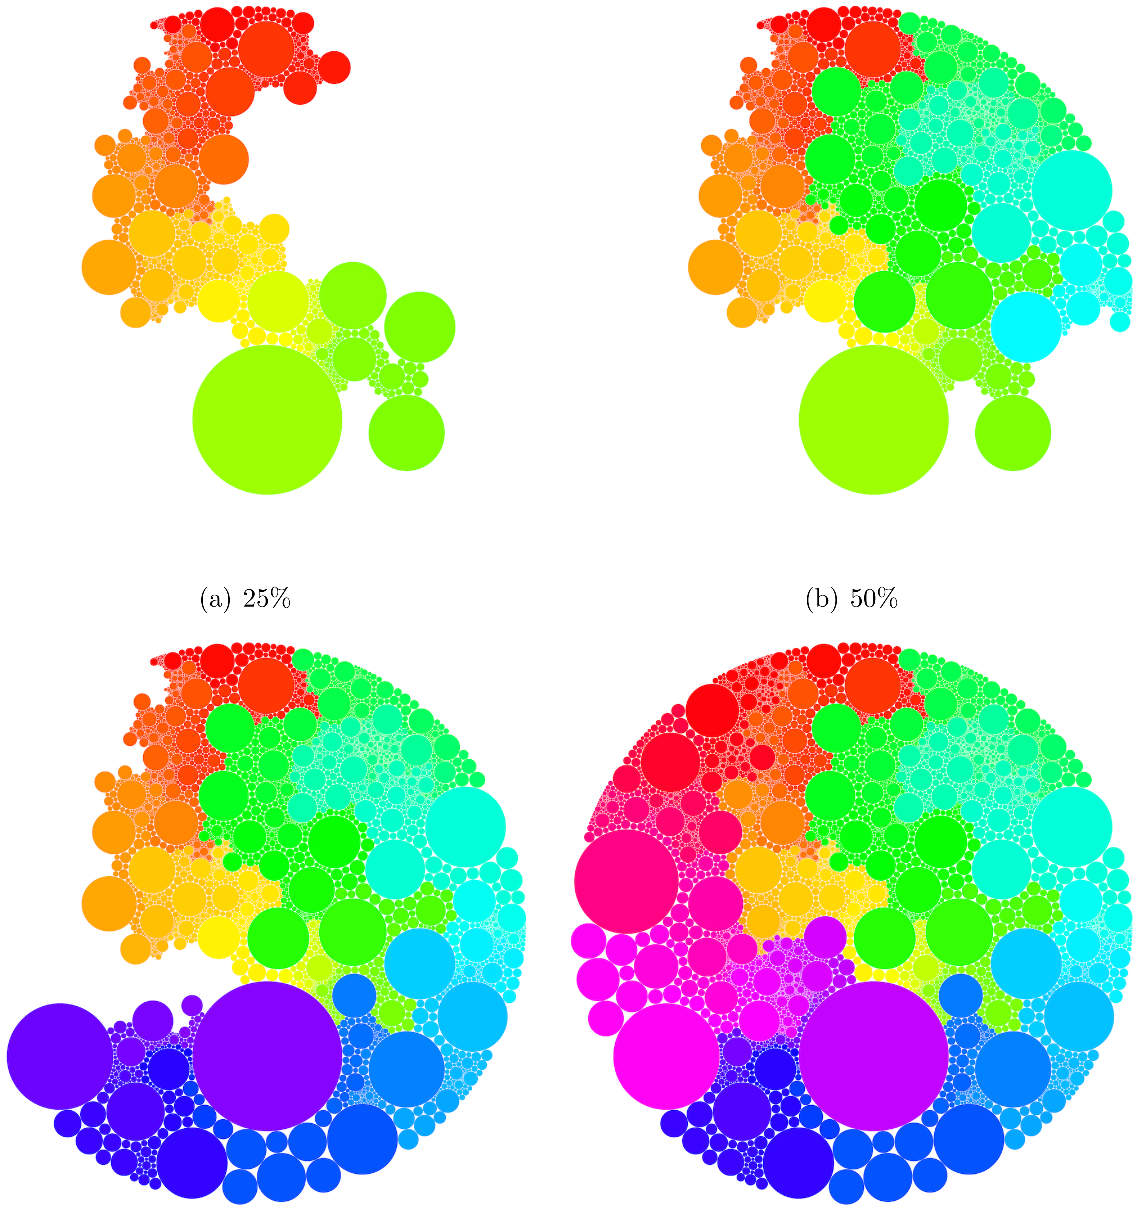

(a) 25%

(b) 50%

(c) 75%

(d) 100%

Figure 9.5: Shown is the circle packing associated with the map from Figure 9.4. The circles are colored according to the order in which they are visited by the space-filling path. The different panels show the circles visited by the path up to the first time that they cover 25%, 50%, 75%, and 100% of the total area.

to the so-called dual LQG surfaces and that can be obtained through a construction that assigns an independent disk to each jump of a stable Lévy excursion. The rough idea is illustrated in Figure 9.6 and Figure 9.7. (If we cut out the grey filled regions

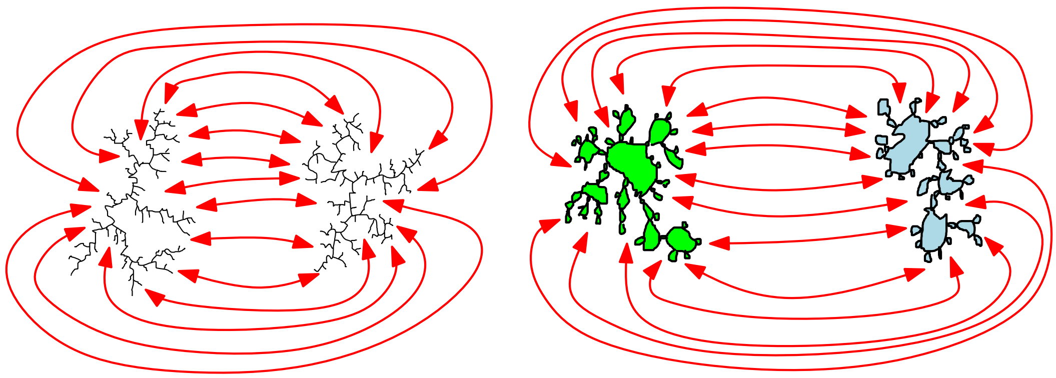

Figure 9.6: **Left:** If the two trees represented by the upper and lower graphs in Figure 9.1 are somehow embedded in the plane in a way that preserves the cyclic order around the tree (let us not worry about precisely how) then the vertical lines of Figure 9.1 correspond to the red lines shown, identifying points along one tree with points along the other. **Right:** Similar, but the trees are replaced with "trees of quantum disks", each somehow embedded in the plane (again, let us not worry about how). Both left and right identification procedures produce a topological sphere, which we embed canonically in $\mathbf{C} \cup \{\infty\}$. In both cases the cyclical ordering on the set of red lines describes a loop on the embedded sphere. In the left scenario, this loop is a space-filling form of $\mathrm{SLE}_{\kappa'}$, with $\kappa' > 4$. In the right scenario, it is a non-space-filling $\mathrm{SLE}_{\kappa'}$ with $\kappa' \in (4, 8)$.

from one the trees in Figure 9.7, we obtain a tree of loops; this tree of loops was also analyzed in the context of scaling limits of random planar maps in [CK13].) In some ways, this construction may seem more intuitive than the one from Section 9.2, since it produces a sphere as a quotient of objects that have non-empty interior to begin with. However, we will actually derive this construction as a consequence of the one described in Section 9.2, with $\kappa' \in (4, 8)$ taken to be the same value in the two pictures, as we explain in more detail in Section 9.8.

9.5 Quantum wedges, cones, disks, and spheres

A **quantum surface**, as in [DS11a, She10], is formally represented by a pair (D, h) where D is planar domain and h is an instance of (some form of) the Gaussian free field (GFF) on D. Intuitively, the quantum surface is the manifold conformally parameterized by D, with a metric tensor given by $e^{\gamma h(z)}$ times the Euclidean one, where $\gamma \in [0, 2)$ is a fixed constant. Since h is a distribution and not a function, this notion requires interpretation.

We can use a regularization procedure to define an area measure on D:

$$\mu = \mu_h := \lim_{\varepsilon \to 0} \varepsilon^{\gamma^2/2} e^{\gamma h_\varepsilon(z)} dz, \tag{9.1}$$

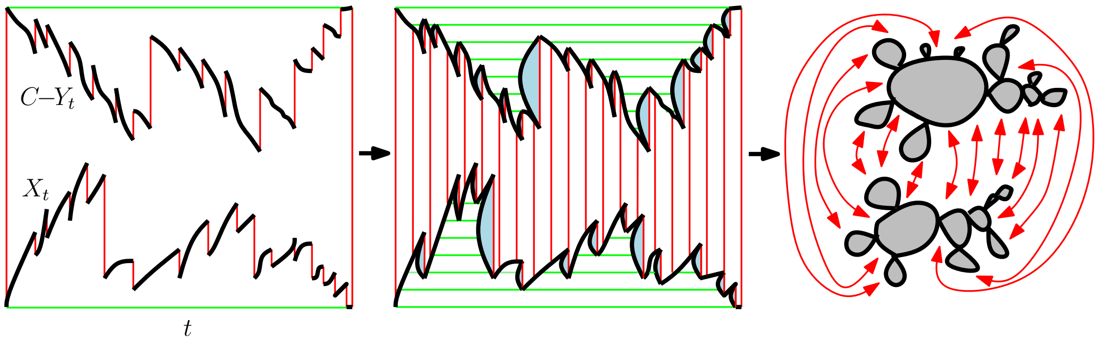

Figure 9.7: **Left:** X_t and Y_t are i.i.d. Lévy excursions, each with only negative jumps. Graphs of X_t and $C - Y_t$ are sketched; red segments indicate jumps. **Middle:** Add a black curve to the left of each jump, connecting its two endpoints; the precise form of the curve does not matter (as we care only about topology for now) but we insist that it intersect each horizontal line at most once and stay below the graph of X_t (or above the graph of $C - Y_t$) except at its endpoints. We also draw the vertical segments that connect one graph to another, as in Figure 9.1, declaring two points equivalent if they lie on the same such segment (or on the same jump segment). Shaded regions (one for each jump) are topological disks. **Right:** By collapsing green segments and red jump segments, one obtains two trees of disks with outer boundaries identified, as on the right side of Figure 9.6.

where dz is Lebesgue measure on D, $h_\varepsilon(z)$ is the mean value of h on the circle $\partial B(z, \epsilon)$ and the limit represents weak convergence in the space of measures on D. (The limit exists almost surely, at least if ε is restricted to powers of two [DS11a].) We interpret μ_h as the area measure of a random surface conformally parameterized by D. When $x \in \partial D$, we let $h_\varepsilon(x)$ be the mean value of h on $D \cap \partial B(x, \epsilon)$. Similarly, on a linear segment of ∂D, where h has free boundary conditions, we may define a boundary length measure by

$$\nu = \nu_h := \lim_{\varepsilon \to 0} \varepsilon^{\gamma^2/4} e^{\gamma h_\varepsilon(x)/2} dx, \tag{9.2}$$

where dx is Lebesgue measure on ∂D [DS11a].

One can parameterize the same quantum surface with a different domain $\widetilde{D}$, and our regularization procedure implies a simple rule for changing coordinates. Suppose that ψ is a conformal map from a domain $\widetilde{D}$ to D and write $\widetilde{h}$ for the distribution on $\widetilde{D}$ given by $h \circ \psi + Q \log |\psi'|$ where

$$Q := \frac{2}{\gamma} + \frac{\gamma}{2},$$

as in Figure 9.8. Then μ_h is almost surely the image under ψ of the measure $\mu_{\widetilde{h}}$. That is, $\mu_{\widetilde{h}}(A) = \mu_h(\psi(A))$ for $A \subseteq \widetilde{D}$. Similarly, ν_h is almost surely the image under ψ of the measure $\nu_{\widetilde{h}}$ in the case that ν_h and $\nu_{\widetilde{h}}$ are both defined and ψ extends to be a homeomorphism on the closure of the domain. In fact, [DS11a] formally defines a **quantum surface** to be an equivalence class of pairs (D, h) under the equivalence transformations (see Figure 9.8)

$$(D, h) \to \psi^{-1}(D, h) := (\psi^{-1}(D), h \circ \psi + Q \log |\psi'|) = (\widetilde{D}, \widetilde{h}), \tag{9.3}$$

noting that both area and boundary length are well-defined for such surfaces. The invariance of ν_h under (9.3) actually yields a definition of the quantum boundary length measure ν_h when the boundary of D is not piecewise linear—i.e., in this case, one simply maps to the upper half plane (or any other domain with a piecewise linear boundary) and computes the length there.

One can also define a **quantum surface with k marked points** to be an equivalence class of elements of the form $(D, h, x_1, x_2, \ldots, x_k)$, with each $x_i \in \overline{D}$, under maps of the form (9.3), where (9.3) is understood to map the $x_i \in \overline{D}$ to $\psi^{-1}(x_i)$.

We emphasize that the definition of equivalence for quantum surfaces does not require that D be simply connected or even connected.

Before we state our main results, we will first need to introduce certain types of random infinite LQG surfaces called *quantum wedges* and *quantum cones*, as defined in [She10]. Much of the paper will be concerned with studying the properties of these objects. The statement illustrated in Figure 9.3 will ultimately emerge as a consequence of this analysis. These objects are informally introduced in the present section; careful definitions are given in Section 5.

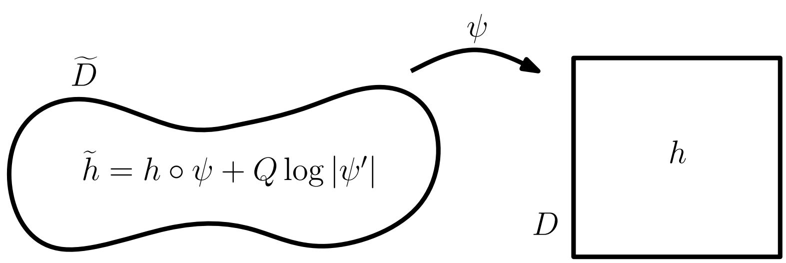

Figure 9.8: A quantum surface coordinate change.

A quantum wedge is a kind of quantum surface, with two marked points, which we define below and discuss in more detail in Section 5. The two marked points divide the boundary of the quantum wedge into an infinite-length left boundary and an infinite-length right boundary, each isometric to $[0, \infty)$. It was shown in [She10] that if one glues (according to boundary length) the right boundary of one particular type of quantum wedge to the left boundary of a quantum wedge with the same law — and then conformally maps the combined surface to the complex upper half plane $\mathbf{H}$, then the image of the interface becomes an SLE curve. We generalize these results to different types of quantum wedges and to sequences of more than two independent quantum wedges. The interfaces one obtains from such weldings turn out to be closely related to the rays in the so-called "imaginary geometry" associated to a GFF [She10, MS12a, MS12b, MS12c, MS13a].

Before we define quantum wedges, let us give a closely related construction. We present below a list of three equivalent ways to parameterize a quantum surface (D, h) in which the h is only defined modulo a global additive constant (so that the surface is only defined modulo constant rescaling). We refer to the quantum surface described in the list below (defined modulo an additive constant for h) as an **unscaled quantum wedge**.

1. Let D be the infinite wedge $\mathcal{W}_\theta = \{z \in \mathbf{C} : \arg(z) \in [0, \theta]\}$ for some $\theta > 0$. Then take h to be an instance of the free boundary GFF on $\mathcal{W}_\theta$.

2. Parameterize instead by the upper half plane $\mathbf{H}$. If we conformally map $\mathbf{H}$ to $\mathcal{W}_\theta$ via the map $\psi_\theta(\widetilde{z}) = \widetilde{z}^{\theta/\pi}$, then the coordinate change rule (9.3) for quantum surfaces implies that we can represent this as $\widetilde{D} = \mathbf{H}$ and let $\widetilde{h}$ be an instance of the free boundary GFF on $\mathbf{H}$ *plus* the deterministic function $Q \log |\psi'_\theta|$, which up to additive constant is $Q \log |\widetilde{z}^{\theta/\pi - 1}| = Q(\frac{\theta}{\pi} - 1) \log |\widetilde{z}| = -\alpha \log |\widetilde{z}|$ for $\alpha := Q(1 - \theta/\pi)$.

3. Parameterize instead by the infinite strip $\mathcal{S} = \mathbf{R} \times [0, \pi]$. This can be mapped to $\mathcal{W}_\theta$ via the map $\phi_\theta(\widetilde{z}) = e^{(\theta/\pi)\widetilde{z}}$, and then the coordinate change rule (9.3) for quantum surfaces implies that we can represent the quantum surface by $\mathcal{S}$ using an instance of the free boundary GFF on $\mathcal{S}$ plus the deterministic function $Q \log |\phi'_\theta| = Q(\theta/\pi)\mathrm{Re}\, \widetilde{z} = (Q - \alpha)\mathrm{Re}\, \widetilde{z}$ (modulo additive constant).

The strip approach is convenient because it turns out that an instance of the free boundary GFF on $\mathcal{S}$ can be represented as $B_{2\mathrm{Re}\, z}$ (where B is standard Brownian motion defined up to additive constant — this function is constant on vertical line segments of the strip) plus a "lateral noise" given by the $(\cdot, \cdot)_\nabla$-orthogonal projection of the free boundary GFF on $\mathcal{S}$ onto the subspace of functions on $\mathcal{S}$ that have mean zero on each vertical segment. One obtains an unscaled quantum wedge by replacing B_{2t} with $B_{2t} + at$ where $a = Q - \alpha \geq 0$ (again defined modulo additive constant). We remark that the parameter space $\mathcal{W}_\theta$ is actually less convenient to work with than either the strip of the half plane, and it will not be used outside of this subsection. (Moreover, in the future, the parameter θ will be used for another purpose: namely, to describe an angle from the theory of imaginary geometry.)

A quantum wedge is the same as an unscaled quantum wedge except that B_t is replaced by a closely related process whose graph is defined modulo horizontal translation, instead of modulo vertical translation. One simple way to describe this process is to say that it is the log of a Bessel process (started at 0, with dimension $\delta = 2 + \frac{2}{\gamma}a$, with $\delta \geq 2$) parameterized to have quadratic variation $2du$ (see Figure 9.9). This process is

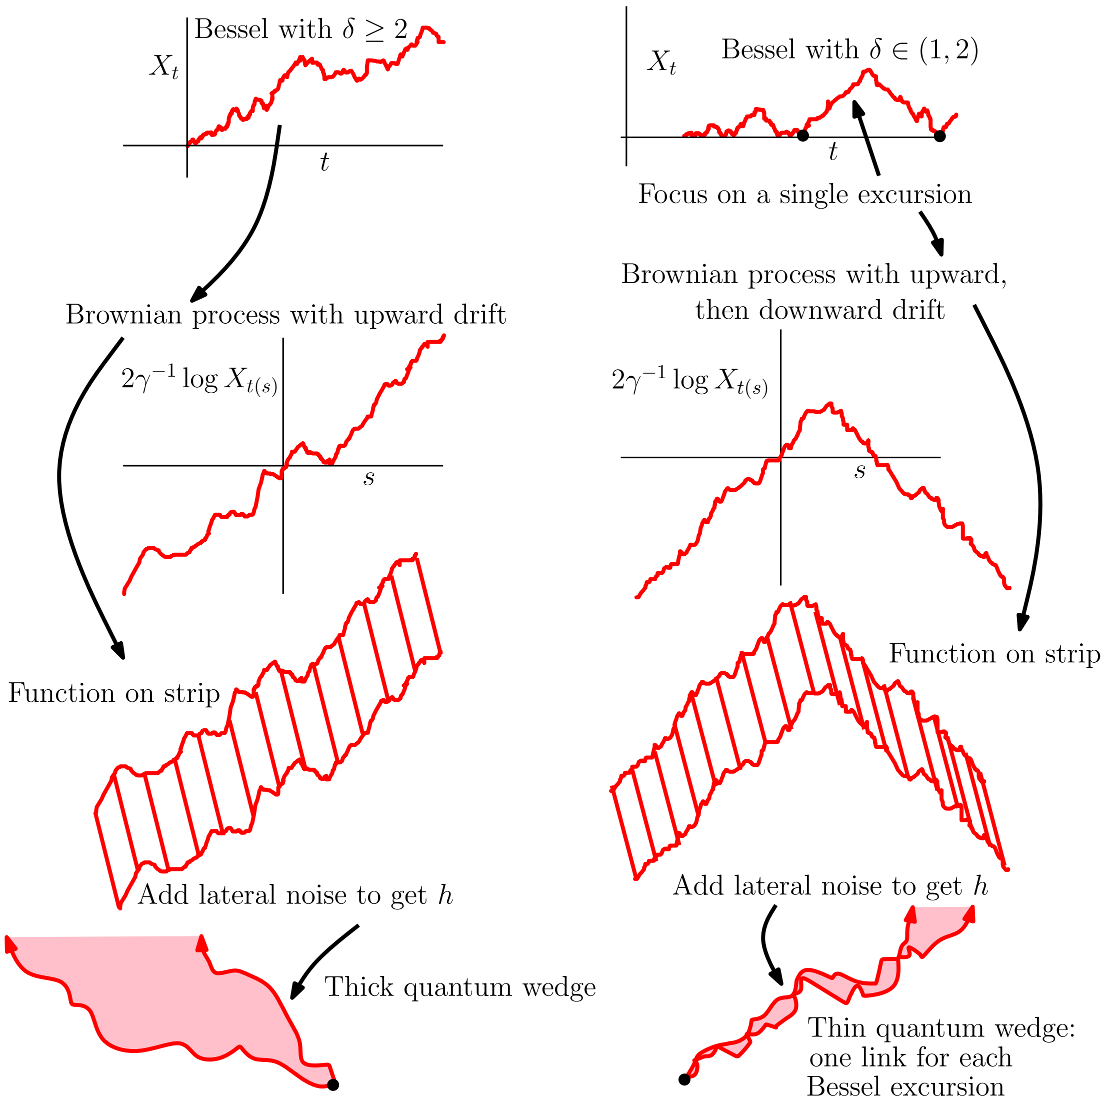

Figure 9.9: If X_t is a Bessel process with $\delta \geq 2$, then $2\gamma^{-1} \log X_t$ can be reparameterized to have quadratic variation $2du$ so that it evolves as $B_{2u} + au$ where B is a standard Brownian motion and a is a constant. The reparameterized graph is determined up to horizontal translation. When $\delta \in (1,2)$, the restriction of $2\gamma^{-1} \log X_t$ to a single Bessel excursion can be similarly reparameterized to obtain a Brownian process with a single maximum and a positive/negative drift on the left/right of that maximum (conditioned not to exceed the maximum); we obtain a map from the infinite strip $\mathcal{S} = \mathbf{R} \times [0, \pi]$ to $\mathbf{R}$ by applying this function to the first coordinate. After adding "lateral noise" obtained by $(\cdot,\cdot)_\nabla$-orthogonally projecting the free boundary GFF on $\mathcal{S}$ onto the subspace of functions with mean zero on all vertical segments, one obtains a field h. If $\delta \geq 2$ then $(\mathcal{S}, h)$ is a **thick quantum wedge**. If $\delta \in (1,2)$, one generates a quantum disk this way for each Bessel excursion, and the resulting chain is a **thin quantum wedge**.

essentially a Brownian motion with drift but the choice of how to translate the graph horizontally is somewhat arbitrary (and was chosen in Figure 9.9 to make the process reach 0 for the first time at time 0). In fact, if we use this approach, we can also consider Bessel processes of dimension $\delta \in (1,2)$, and define a quantum surface for each component of the Bessel process, as illustrated in Figure 9.9.

The quantum surfaces obtained for $\delta \geq 2$, or equivalently, $\alpha \leq Q$, are called **thick quantum wedges**. A thick quantum wedge is a random quantum surface with two marked points (0 and ∞) whose law is invariant under constant rescalings. Each quantum wedge—when parameterized by $\mathcal{S}$ as discussed above—has an infinite amount of quantum mass, almost surely, but only a finite amount corresponding to any particular bounded subset of $\mathcal{S}$. (In particular, the law of a quantum wedge is not symmetric under reversing the two marked points, since every neighborhood of its second point has infinite mass, and this is not true of the first point.) A **thin quantum wedge**, obtained for $\delta \in (1,2)$, or equivalently, $\alpha \in (Q, Q+\frac{\gamma}{2})$, is an infinite Poissonian "chain" (concatenation) of finite volume quantum surfaces, each with two marked boundary points; there is one quantum surface for each excursion of a Bessel process from 0, as Figure 9.9 illustrates. The thick-quantum wedge condition $\alpha \leq Q$ corresponds in Liouville quantum theory to the so-called *Seiberg bound* [Sei90].[25]

We will assign a "weight" parameter to each quantum wedge and show that these weights are additive under gluing operations. In particular, a wedge of weight W can be produced by welding together n independent wedges of weight W/n. Taking the $n \to \infty$ limit, we will obtain a way to construct a wedge by gluing together countably many quantum surfaces that can in some sense be interpreted as wedges of weight zero. Explicitly, the **weight** of an α-quantum wedge is the number defined from α as follows:

$$W := \gamma\left(\gamma + \frac{2}{\gamma} - \alpha\right) = \gamma\left(\frac{\gamma}{2} + Q - \alpha\right). \tag{9.4}$$

We will usually describe an α-quantum wedge in terms of either α or its weight W depending on the context. Table 1 summarizes the relationships between several ways of parameterizing the space of wedges (all equivalent up to affine transformation, once γ is fixed). Two of these will be introduced at later points in this paper. We have already introduced α, W, δ, and a. The value Δ is a "quantum scaling exponent." It turns out that at a point conditioned to intersect a random fractal of quantum scaling exponent Δ (boundary or interior) in the KPZ framework, the surface typically has a logarithmic singularity or magnitude $\alpha = \gamma - \gamma\Delta$, see [DS11a, Equation (63)] and Section **??**. The value θ is an imaginary geometry angle.

Taking $\alpha \leq Q$ corresponds to taking $W \geq \frac{\gamma^2}{2}$, which thus corresponds to a thick wedge. If we extend the formula relating W to α beyond this range, we find that taking

[25] For an introduction, see, e.g., [GM93, DFGZJ95, Nak04]. The thin wedges correspond to values of α above this bound; however, as we shall see shortly, each of the concatenated finite volume quantum surfaces which form the thin wedge locally looks like an $\widetilde{\alpha}$-wedge with the reflected value $\widetilde{\alpha} = 2Q - \alpha$ near each of its endpoints, i.e., a thick wedge with $\widetilde{\alpha} \in (\frac{2}{\gamma}, Q)$.

	α	W	θ	δ	a	Δ
α	—	$\frac{\gamma}{2}+Q-\frac{1}{\gamma}W$	$\frac{\gamma}{2}+Q-\chi\frac{\theta}{\pi}$	$Q+\frac{\gamma}{2}(2-\delta)$	$Q-a$	$\gamma(1-\Delta)$
W	$\gamma(\frac{\gamma}{2}+Q-\alpha)$	—	$\chi\gamma\frac{\theta}{\pi}$	$\frac{\gamma^2}{2}(\delta-1)$	$\gamma a+\frac{\gamma^2}{2}$	$2+\gamma^2\Delta$
θ	$\frac{\pi}{\chi}(\frac{\gamma}{2}+Q-\alpha)$	$\frac{\pi}{\chi\gamma}W$	—	$\frac{\pi\gamma}{2\chi}(\delta-1)$	$\frac{\pi}{\chi}(a+\frac{\gamma}{2})$	$\frac{\pi}{\chi}(\frac{2}{\gamma}+\gamma\Delta)$
δ	$2+\frac{2}{\gamma}(Q-\alpha)$	$1+\frac{2}{\gamma^2}W$	$1+\chi\frac{2}{\gamma}\frac{\theta}{\pi}$	—	$2+\frac{2}{\gamma}a$	$2(\Delta+\frac{1}{\gamma}Q)$
a	$Q-\alpha$	$\frac{1}{\gamma}W-\frac{\gamma}{2}$	$\chi\frac{\theta}{\pi}-\frac{\gamma}{2}$	$\frac{\gamma}{2}(\delta-2)$	—	$\chi+\gamma\Delta$
Δ	$1-\frac{\alpha}{\gamma}$	$\frac{1}{\gamma^2}(W-2)$	$\frac{\chi}{\gamma}\frac{\theta}{\pi}-\frac{2}{\gamma^2}$	$\frac{1}{2}\delta-\frac{1}{\gamma}Q$	$\frac{1}{\gamma}(a-\chi)$	—

Table 1: We can parameterize the space of **quantum wedges** using a number of different variables that are affine transformations of each other once γ is fixed. Shown are six such possibilities which are important for this article: multiple (α) of $-\log|\cdot|$ used when parameterizing by $\mathbf{H}$, weight (W), angle gap in imaginary geometry (θ), dimension of Bessel process (δ), drift for the corresponding Brownian motion (a) when parameterizing by the strip $\mathcal{S} = \mathbf{R} \times [0,\pi]$, and boundary quantum exponent at a quantum typical point in a fractal conditioned to intersect $\partial\mathbb{H}$ (Δ). (**Important:** The θ here is *not the same* as the one used to define $\mathcal{W}_\theta$ above.) Each element of the table gives the variable corresponding to the row as a function of the variable which corresponds to the column. Here, $Q = 2/\gamma + \gamma/2$ and $\chi = 2/\gamma - \gamma/2$ is a constant from imaginary geometry [MS12a, MS12b, MS12c, MS13a]. Recall that Theorem 9.9 gives the additivity of weights under the welding operation; by combining this with the table above, we can see how the other ways of parameterizing a wedge transform under the welding operation. In particular, the only other parameterization given above which is additive is θ.

$\alpha \in (Q, Q+\frac{\gamma}{2})$ formally corresponds to taking $W \in (0, \frac{\gamma^2}{2})$, which thus corresponds to a thin wedge.

For the reader who has read [She10], we briefly mention a few of the special wedge types discussed there:

1. Wedge obtained by zooming in at a boundary-measure-typical point: $\alpha = \gamma$, $W = 2$.
2. Wedge obtained by zooming in near the origin of a capacity invariant SLE_κ quantum zipper: $\alpha = -2/\gamma$, $W = 4 + \gamma^2$.
3. Wedge obtained by zooming in near the origin of a quantum-length invariant

	α	W	θ	δ	a	Δ
α	—	$Q-\frac{1}{2\gamma}W$	$Q-\chi\frac{\theta}{2\pi}$	$Q-\frac{\gamma}{4}(\delta-2)$	$Q-a$	$\gamma(1-\Delta)$
W	$2\gamma(Q-\alpha)$	—	$\gamma\chi\frac{\theta}{\pi}$	$\frac{\gamma^2}{2}(\delta-2)$	$2a\gamma$	$\gamma^2(2\Delta-1)+4$
θ	$\frac{2\pi}{\chi}(Q-\alpha)$	$\frac{\pi}{\gamma\chi}W$	—	$\gamma\frac{\pi}{2\chi}(\delta-2)$	$\frac{2\pi}{\chi}a$	$2\pi(1+\frac{\gamma}{\chi}\Delta)$
δ	$2+\frac{4}{\gamma}(Q-\alpha)$	$2+\frac{2}{\gamma^2}W$	$2+\chi\frac{2}{\gamma}\frac{\theta}{\pi}$	—	$2+\frac{4}{\gamma}a$	$4\Delta+\frac{8}{\gamma^2}$
a	$Q-\alpha$	$\frac{1}{2\gamma}W$	$\chi\frac{\theta}{2\pi}$	$\frac{\gamma}{4}(\delta-2)$	—	$\chi+\gamma\Delta$
Δ	$1-\frac{\alpha}{\gamma}$	$\frac{1}{2}+\frac{1}{2\gamma^2}(W-4)$	$\frac{\chi}{\gamma}(\frac{\theta}{2\pi}-1)$	$\frac{1}{4}\delta-\frac{2}{\gamma^2}$	$\frac{1}{\gamma}(a-\chi)$	—

Table 2: As in the case of quantum wedges (Table 1), we can parameterize the space of **quantum cones** using a number of different variables. Shown are six such possibilities which are important for this article: multiple of $-\log|\cdot|$ (α) when parameterizing by $\mathbf{C}$, weight (W), "space of angles" in imaginary geometry (θ), dimension of Bessel process (δ), and drift for the corresponding Brownian motion (a) when parameterizing by the cylinder $\mathcal{Q} = \mathbf{R}\times[0,2\pi]$ (with $\mathbf{R}\times\{0\}$ and $\mathbf{R}\times\{2\pi\}$ identified), and quantum exponent at a quantum typical point in a fractal in $\mathbb{H}$ (Δ). Each element of the table gives the variable corresponding to the row as a function of the variable which corresponds to the column. Here, $Q = 2/\gamma+\gamma/2$ and $\chi = 2/\gamma-\gamma/2$ is a constant from imaginary geometry [MS12a, MS12b, MS12c, MS13a].

SLE$_\kappa$ quantum zipper: $\alpha = -2/\gamma+\gamma$, $W=4$.

We will indicate a quantum wedge $\mathcal{W}$ parameterized by D and described by a distribution h with the notation (D,h). If we also wish to emphasize the distinguished origin (say z_1) and ∞ (say z_2) we will denote the wedge by (D,h,z_1,z_2).

Roughly, an α-**quantum cone** is a quantum surface obtained by taking $e^{\gamma h(z)}dz$ where $\gamma\in(0,2)$ is a fixed constant and h is an instance of the GFF on the infinite cone with opening angle θ, i.e. the surface that arises by starting with $\mathcal{W}_\theta$ and then identifying its left and right sides according to Lebesgue measure. Analogous to the case of a quantum wedge, by performing a change of coordinates via the map $\psi_\theta(\widetilde{z}) = \widetilde{z}^{\theta/2\pi}$ and applying (9.3), we can represent an unscaled quantum cone in terms of the sum of an instance $\widetilde{h}$ of the whole-plane GFF plus the deterministic function $-\alpha\log|\widetilde{z}|$ where $\alpha = Q(1-\frac{\theta}{2\pi})$. A quantum cone can be defined formally and precisely using the analog of Figure 9.9 in which the strip is replaced by the cylinder, see Section 5. As in the case of quantum wedges, there are a number of different ways of parameterizing the space of quantum cones; see Table 2.

We define the **weight** W of an α-quantum cone to be the number

$$W := 2\gamma(Q - \alpha). \tag{9.5}$$

Theorem 9.12 below will show that one can "glue" the two sides of a weight W quantum wedge to obtain a weight W quantum cone.

As in the case of wedges, we will indicate a quantum cone $\mathcal{C}$ parameterized by a domain D and described by a distribution h with the notation (D, h). If we also wish to emphasize the distinguished origin (say z_1) and ∞ (say z_2) we will denote the cone by (D, h, z_1, z_2).

We conclude this subsection by pointing out that the Bessel process construction as described in Figure 9.9 gives us a way to define a family of natural *infinite* measures on the space of *finite*-volume quantum disks (or spheres). Recall that for any $\delta < 2$ (which corresponds to a slope $a < 0$) one can define the **Bessel excursion measure** $\nu_\delta^{\mathrm{BES}}$ (i.e., the Itô excursion measure associated with the excursions that a Bessel process makes from 0), which is an infinite measure on the space of continuous processes X_t indexed by $[0, T]$ (for some random T) satisfying $X_0 = X_T = 0$ and $X_t > 0$ for $t \in (0, T)$. The construction of this measure is recalled in Remark **??**.

When $\delta \in (0, 2)$, one can consider a Bessel process X_t and let ℓ_t denote the local time the processes has spent at 0 between times 0 and t. If we parameterize time by the right-continuous inverse of ℓ_t, then the excursions appear as a Poisson point process (p.p.p.) on $\mathbf{R}_+ \times \mathcal{E}$ where $\mathcal{E}$ is the space of continuous functions $\phi\colon \mathbf{R}_+ \to \mathbf{R}_+$ such that $\phi(0) = 0$ and $\phi(t) = 0$ for all $t \geq T$, some T, and it is possible to recover the Bessel process from the p.p.p. from concatenating the countable collection of excursions. When $\delta \leq 0$, the corresponding p.p.p. is still well-defined, but it a.s. assigns, to any finite time interval, a countable collection of excursions whose lengths sum to ∞ (so that it is not possible to define the reflecting Bessel process in the same way, essentially because there are "too many small excursions").

Given any Bessel excursion measure $\nu_\delta^{\mathrm{BES}}$, with $\delta < 2$, one can construct a doubly marked quantum surface using the procedure described in Figure 9.9. Observe that this surface looks like a thick quantum wedge with a value given by $-a$ (or δ value given by $4 - \delta$) near each of the two endpoints.

This suggests a way of parameterizing the family of measures on disks (those induced by the measures $\nu_\delta^{\mathrm{BES}}$ and the procedure from Figure 9.9). We define a **quantum disk of weight** W to be a sample from the infinite measure $\mathcal{M}_W$ on quantum disks induced by $\nu_\delta^{\mathrm{BES}}$ with δ chosen so that surface looks like a quantum wedge of weight W near each of its two endpoints. We define parameters α and θ similarly. A similar procedure allows us to define a **quantum sphere of weight** W to be a sample from the infinite measure $\mathcal{N}_W$ on quantum spheres induced by $\nu_\delta^{\mathrm{BES}}$ with δ chosen so that the surface looks like a weight W quantum cone near each of its two endpoints. We similarly define the parameters α and θ for quantum disks/spheres to be those of the corresponding wedges/cones. (We will describe these constructions in more detail in Section **??**.)

Note that by definition, a thin quantum wedge of weight W (recall that the wedge being thin means $W \in (0, \frac{\gamma^2}{2}), \alpha \in (Q, Q + \frac{\gamma}{2})$) is obtained as a Poissonian concatenation of disks of weight $\widetilde{W}$, where the a corresponding to W is -1 times the a corresponding to $\widetilde{W}$. That is, $\widetilde{W} = \gamma^2 - W$, and $\widetilde{\alpha} = 2Q - \alpha$ (Table 1). Thus, as mentioned in Footnote 25, these marked disks locally look like thick wedges near their marked points and obey the Seiberg bound, $\widetilde{\alpha} \leq Q$.

The measure $\mathcal{M}_W$ has special significance when $W = 2$ so that $\alpha = \gamma$ and $\theta = 2\pi/(2-\frac{\gamma^2}{2})$. The measure $\mathcal{N}_W$ has special significance when $W = 4 - \gamma^2$, so that the cone value for α is γ and $\theta = 2\pi$. In these cases, the two marked points look like "typical" points chosen from the boundary or bulk quantum measures. We sometimes use the terms **quantum disk** or **quantum sphere** (without specifying weight) and the symbols $\mathcal{M}$ and $\mathcal{M}^c$ to denote samples from the infinite measures with these special weights.

One may obtain a **unit boundary length quantum disk** by sampling from $\mathcal{M}$ conditioned to have total boundary length 1. (This is now a sample from a finite measure, which can be normalized to be a probability measure.) Similarly, a **unit area quantum disk** and **unit area quantum sphere** can respectively be defined from $\mathcal{M}$ and $\mathcal{M}^c$ in a similar way. We define a quantum disk with an arbitrary fixed boundary length (or a quantum disk or quantum sphere with a fixed area) similarly. (Alternatively, a quantum disk with arbitrary boundary length can be defined by taking the unit boundary length quantum disk and then scaling it to have the given boundary length; the same is also true in the case of the disk or sphere with fixed area.)

9.6 Conformal structure and removability

As mentioned earlier, an LQG surface can be obtained by endowing a topological surface with both a good measure and a conformal structure in a random way. (And we can imagine that these two structures are added in either order.) Given two topological disks with boundary (each endowed with a good area measure in the interior, and a good length measure on the boundary) it is a simple matter to produce a new good-measure-endowed topological surface by taking a quotient that involves gluing (all or part of) the boundaries to each other in a boundary length preserving way.

The problem of *conformally welding* two surfaces is the problem of obtaining a conformal structure on the combined surface, given the conformal structure on the individual surfaces. (See, e.g., [Bis07] for further discussion and references.) This is closely related to the problem of defining a Brownian motion (up to monotone reparameterization) on the combined surface, given the definition of a Brownian motion on each of the individual pieces. We will now briefly (and somewhat informally) describe this connection. It is well known that a conformal structure and a choice of initial point determine a Brownian diffusion process β_t (up to monotone reparameterization). Similarly, if one is given the diffusion process (for all starting points), one can recover the conformal structure in a neighborhood of a point as follows: consider any Jordan curve surrounding that point, with three marked points on the curve dividing it into three segments E_1, E_2, E_3; then for each z in the region surrounded by the curve, consider the triple (p_1, p_2, p_3) where p_i for $i = 1, 2, 3$ is the probability that the Brownian motion starting from z first hits the curve along E_i. Note that $p_1 + p_2 = p_3 = 1$. One can then "conformally map" this region to a Euclidean triangle by sending each such z to the point in the triangle such that a standard two-dimensional Brownian motion from that point has probability p_i of first hitting the triangle boundary along the ith edge. The image of β_t in the triangle will then be two-dimensional Brownian motion (up to a time change).

This recipe can be carried out for any continuous diffusion process β_t to produce a map to the triangle; however it is not true that for all diffusion processes the image process in the triangle will be a standard Brownian motion. For example, it is necessary that the stationary measure for β_t (when one allows reflection at the boundary) is a good measure, and that β_t is reversible with respect to this measure. The diffusions of *Brownian type* — i.e., diffusions that correspond to time-changed Brownian motions w.r.t. *some* conformal structure — are a rather special subset of the set of all of the continuous diffusion processes one might produce on a topological sphere. "Conformally welding" the two conformal surfaces is equivalent to producing a continuous diffusion process of Brownian type that agrees with the Brownian motions on the individual surfaces at times when it is away from the boundary interface. To make sense of this idea, we will draw from the theory of *removability sets*, as explained below.

A compact subset K of a domain $D \subseteq \mathbf{C}$ is called (conformally) **removable** if every homeomorphism from D to D that is conformal on $D \setminus K$ is also conformal on all of D. In probabilistic language, the fact that K is removable means that once we are given the Brownian motion behavior off K, there is — among all the possible ways of extending this process to a continuous diffusion on D, involving various types of local time pushes along K, etc. — only one of Brownian type. A Jordan domain $D \subseteq \mathbf{C}$ is said to be a **Hölder domain** if any conformal transformation from $\mathbf{D}$ to D is Hölder continuous all of the way up to $\partial \mathbf{D}$. It was shown by Jones and Smirnov [JS00a] that if $K \subseteq D$ is the boundary of a Hölder domain, then K is removable; it is also noted there that if a compact set K is removable as a subset of D, then it is removable in any domain containing K, including all of $\mathbf{C}$. Thus, at least for compact sets K, one can speak of removability without specifying a particular domain D.

It was further shown by Rohde and Schramm [RS05b] that the following is true. Suppose that $g_t \colon \mathbf{H} \setminus \eta([0, t]) \to \mathbf{H}$ is the forward Loewner flow associated with an SLE_κ curve with $\kappa < 4$. Then for each $t \geq 0$, the map g_t^{-1} is a.s. Hölder continuous. (A more general result, proved in a different way, which implies this is also given in [MS13b, Section 8.1].) If η is an SLE_κ in $\mathbf{H}$ from 0 to 1, then the union of η and its reflection across the real axis is the boundary of a bounded Hölder domain, hence is removable by [JS00a]. It follows immediately that an SLE_κ path η from 0 to ∞ is removable in $\mathbf{H}$, and that $\eta([t_1, t_2])$ is removable for any $t_1 < t_2$. In fact, we can also say the following:

Proposition 9.6. *Suppose that η is an SLE_κ curve in $\mathbf{H}$ from 0 to 1 with $\kappa \in (0,4)$. Then it is a.s. the case that for a dense set of times $t_1 < t_2$ the segment $\eta([t_1,t_2])$ is removable in the domain $\mathbf{H} \setminus \eta\big([0,t_1] \cup [t_2,\infty]\big)$.*

Proof. As explained above, we already have the removability for the overall curve from 0 to 1. If we observe the path starting from 0 (up to any stopping time) and then observe the path from 1 (up to a reverse stopping time) then the conditional law of the remaining path is that of an SLE_κ in the remaining domain [Zha08b, MS12b], which implies the result. $\square$

Proposition 9.6 also implies the removability of random paths η that look locally like SLE_κ, such as the $\mathrm{SLE}_\kappa(\rho)$ processes with $\rho > -2$ and flow lines of the GFF. To see why, observe that for any point z on such a path, one can find times t_1 and t_2 such that $z \in \eta([t_1,t_2])$ and $\eta([t_1,t_2])$ is removable in $\mathbf{H} \setminus \eta\big([0,t_1] \cup [t_2,\infty]\big)$. It thus follows that any homeomorphism ϕ which is conformal on the complement of η will also be conformal in a neighborhood containing $\eta((t_1,t_2))$. Since this holds for any z, we find that any such map must be conformal in a neighborhood every point on η, hence conformal everywhere.

To our knowledge, it is not known in general that the union of two (non-disjoint) removable sets is removable. However, this is not a problem when one of the sets is of SLE type:

Proposition 9.7. *Suppose that $\widetilde{\eta}$ is a random segment of an SLE_κ curve η with $\kappa \in (0,4)$ and suppose that K is removable. Then $\widetilde{\eta} \cup K$ is a.s. removable.*

Proof. Let ϕ be any homeomorphism which is conformal off $\widetilde{\eta} \cup K$. Fix $z \in \widetilde{\eta} \setminus K$. By Proposition 9.6, we can find an interval (t_1,t_2) of time such that $z \in \eta((t_1,t_2))$, $\eta([t_1,t_2])$ is removable in $\mathbf{H} \setminus \eta\big([0,t_1] \cup [t_2,\infty]\big)$, and $\eta([t_1,t_2])$ is at a positive distance from K. From this one can easily see that $\eta([t_1,t_2])$ is removable in $\mathbf{H} \setminus \big(\eta\big([0,t_1] \cup [t_2,\infty]\big) \cup K\big)$ since K has positive distance from $\eta([t_1,t_2])$. It then follows that ϕ must be conformal in a neighborhood of z. Since $z \in \widetilde{\eta} \setminus K$ was arbitrary, we have that ϕ is conformal on the complement of K, hence conformal everywhere by the removability of K. $\square$

A similar argument implies the following:

Proposition 9.8. *A finite union of $\mathrm{SLE}_\kappa(\rho)$ curves with $\kappa \in (0,4)$ and $\rho > -2$ or GFF flow lines in $\mathbb{H}$ or $\mathbf{C}$ is a.s. removable. The same holds for an infinite collection of compact SLE_κ segments (coupled in some way) with the property that a.s. the set of accumulation points of these intervals (i.e., the set of points any neighborhood of which intersects infinitely many segments) is discrete.*

This implies in particular that flow lines of the GFF (whole plane SLE_κ processes started at interior points) are removable.

9.7 Welding and cutting quantum wedges and cones

The current subsection and Section 9.8 present several theorems involving the "gluing together" of infinite volume objects, including quantum wedges, infinite-volume continuum random trees (c.f. Section 9.2) and infinite-volume trees of disks (c.f. Section 9.4). We include the following chart (whose entries will all be explained later) to help the reader keep track of some of these statements. In this chart, the parameters given for the quantum cones and wedges indicate *weight*.

Theorem	**Objects to be glued**	**New object/new interface**
9.9, 9.11	W_1-wedge and W_2-wedge	$(W_1{+}W_2)$-wedge/$\mathrm{SLE}_\kappa(W_1-2;W_2-2)$
9.12	W-wedge, self	W-cone/whole-plane $\mathrm{SLE}_\kappa(W-2)$
9.13, 9.14	coupled pair of CRTs	$2\gamma\chi$-cone/space-filling $\mathrm{SLE}_{\kappa'}$
9.17	forested W_1-, W_2-wedges	$(W_1{+}\gamma\chi{+}W_2)$-wedge/$\mathrm{SLE}_{\kappa'}(\rho_1;\rho_2)$
9.18	bi-forested W-wedge, self	$(W{+}\gamma\chi)$-cone/whole-plane $\mathrm{SLE}_{\kappa'}(\rho)$

Note that for cones, the weight $2\gamma\chi = 4-\gamma^2$ that appears in the third row corresponds to $\alpha=\gamma$ and $\theta = 2\pi$. The value $\gamma\chi$ that appears in the last two rows corresponds to $\theta=\pi$. The so-called "forested quantum wedges" are later introduced and explored through Theorems 9.16–9.18 and Corollary 9.20. These theorems describe the structure of the pair of trees of disks produced by cutting a quantum surface with a form of $\mathrm{SLE}_{\kappa'}$ with $\kappa' \in (4,8)$, which are then related to (an infinite time version of) the pair of Lévy processes described in Figure 9.7. The values of the ρ_i and ρ in the last two rows of the table appear in the theorem statements.

Our first main result describes the welding of wedges with general positive weights. (See Figure 9.10 and Figure 9.11 for illustrations.)

Theorem 9.9 (Welding and cutting for quantum wedges)**.** *Fix $\gamma \in (0,2)$ and choose a quantum wedge $\mathcal{W}$ of weight $W > 0$, represented by $(\mathbf{H}, h, 0, \infty)$. Suppose $W = W_1 + W_2 \geq \frac{\gamma^2}{2}$ for $W_1, W_2 > 0$ and then independently choose an $\mathrm{SLE}_\kappa(\rho_1;\rho_2)$, for $\rho_i = W_i - 2$ and $\kappa = \gamma^2 \in (0,4)$, from 0 to ∞ with force points located at 0^- and 0^+. Let D_1 and D_2 denote left and right components of $\mathbf{H}\setminus\eta$. Then the quantum surfaces $\mathcal{W}_1 = (D_1, h, 0, \infty)$ (with h restricted to D_1) and $\mathcal{W}_2 = (D_2, h, 0, \infty)$ (with h restricted to D_2) are independent. Moreover, each $\mathcal{W}_i$ has the law of a quantum wedge with weight W_i.*

In the case that $W \in (0, \frac{\gamma^2}{2})$, the same statement holds except we take η to be a concatenation of independent $\mathrm{SLE}_\kappa(\rho_1;\rho_2)$ processes (one from the opening point to the

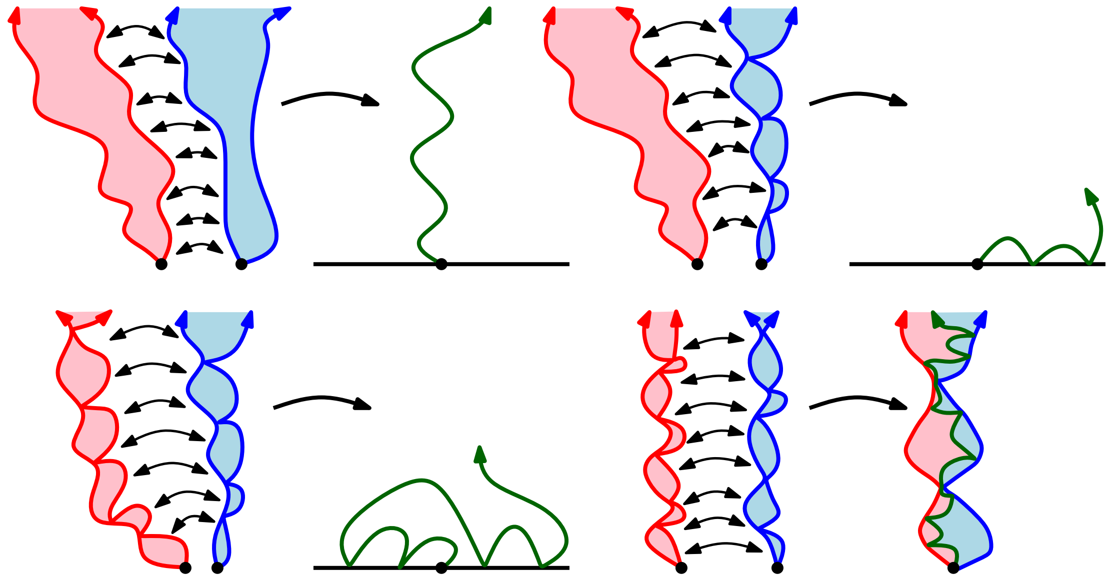

Figure 9.10: Suppose that $\mathcal{W}_1, \mathcal{W}_2$ are independent quantum wedges with respective weights $W_1, W_2 > 0$. Recall that if $W_i \geq \frac{\gamma^2}{2}$ then $\mathcal{W}_i$ is homeomorphic to $\mathbf{H}$ and if $W_i \in (0, \frac{\gamma^2}{2})$ then $\mathcal{W}_i$ consists of an ordered, countable sequence of beads each of which is topologically a disk. Illustrated are the four possible scenarios considered in Theorem 9.9. Let $W = W_1 + W_2$. **Top left:** $W_1, W_2 \geq \frac{\gamma^2}{2}$ so that both wedges are homeomorphic to $\mathbf{H}$. Their conformal welding is a wedge $\mathcal{W}$ of weight W and the interface between them is an $\mathrm{SLE}_\kappa(W_1 - 2; W_2 - 2)$ process independent of $\mathcal{W}$. **Top right:** The same is true if $W_1 \geq \frac{\gamma^2}{2}$, $W_2 \in (0, \frac{\gamma^2}{2})$. Since $W_2 \in (0, \frac{\gamma^2}{2})$, the interface intersects the right boundary of $\mathcal{W}$. **Bottom left:** The same is true if $W_1, W_2 \in (0, \frac{\gamma^2}{2})$ and $W \geq \frac{\gamma^2}{2}$. In this case, the interface intersects both the left and right boundaries of $\mathcal{W}$. **Bottom right:** If $W \in (0, \frac{\gamma^2}{2})$, the welding of $\mathcal{W}_1$ and $\mathcal{W}_2$ is still a wedge $\mathcal{W}$ of weight W. In this case, $\mathcal{W}$ is not homeomorphic to $\mathbf{H}$. Nevertheless, the interface between $\mathcal{W}_1$ and $\mathcal{W}_2$ is independently an $\mathrm{SLE}_\kappa(W_1 - 2; W_2 - 2)$ in each of the beads of $\mathcal{W}$.

closing point of each of the beads of $\mathcal{W}$ with the force points located immediately to the left and right of the opening point) and we take D_1 (resp. D_2) to be the chain of surfaces which are to the left (resp. right) of η.

Remark 9.10. We remark that this "linearity of wedge weights under gluing" (which can be extended to multiple wedges, see Figure 9.11) is pre-figured by certain results on non-intersection exponents that have appeared in the physics and math literatures. For example, suppose that on a random infinite planar triangulation of the half plane, one starts n simple random walks at far away locations and conditions on having them all reach the same single boundary point without intersecting each other. One expects that the (infinite volume, fine mesh) scaling limit should consist of $(2n+1)$ independent

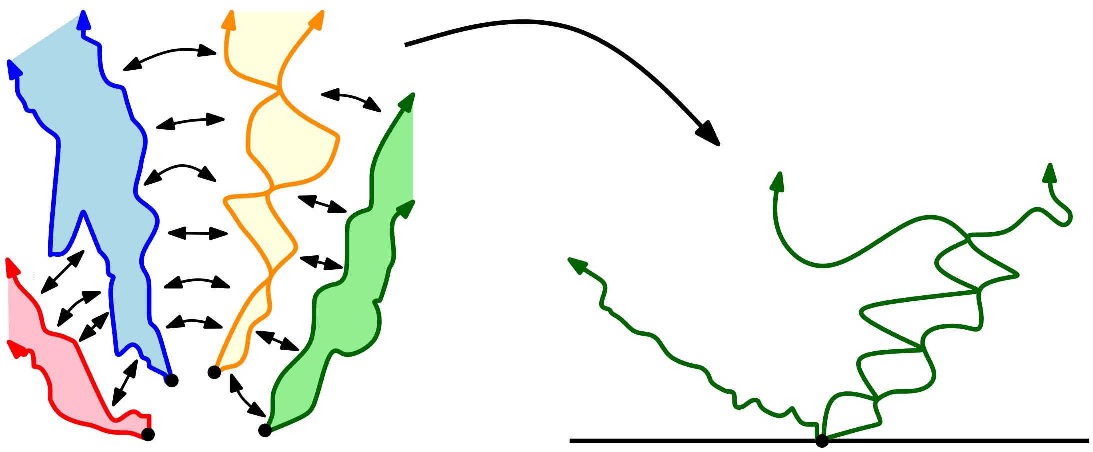

Figure 9.11: Four quantum wedges with respective weights $W_1, \ldots, W_4$ conformally welded along their boundaries and conformally mapped to $\mathbf{H}$. The resulting surface is a wedge of weight $W_1 + \cdots W_4$. In the illustration, $W_1, W_2, W_4 \geq \frac{\gamma^2}{2}$ and $W_3 \in (0, \frac{\gamma^2}{2})$. The images of the interfaces are coupled $\mathrm{SLE}_\kappa(\rho_1; \rho_2)$ processes, corresponding to rays in a so-called imaginary geometry [MS12a, MS12b, MS12c, MS13a]. Since $W_3 \in (0, \frac{\gamma^2}{2})$, the middle interface intersects the rightmost interface.

quantum wedges, with n of them corresponding to a region between the left and right boundaries of a single path, and $n+1$ corresponding to a region between two paths, or between a path and the boundary. The $n+1$ of the latter type should all have weight 2 (essentially because this is the weight of a wedge obtained by zooming in at a "typical boundary measure" point, as mentioned above — see Section **??**). The n wedges of the former type should all have some weight W_0 (which we do not specify for now). Thus the *total weight* of the wedge obtained by gluing these individual wedges together should be $W = nW_0 + 2(n+1) = (W_0 + 2)n + 2$. In particular, this implies that $W - 2$ should be a linear function of n. Equivalently, if we take the formula $W = 2 + \gamma^2 \Delta$ (from Table 1) this means that Δ (the so-called "boundary quantum scaling exponent" of the non-intersecting path event, see Section **??**) is a linear function of n. The fact that this *should* be the case was predicted and advocated by the first co-author using properties of discrete quantum gravity models and discrete analogs of the wedge weldings discussed above [Dup98, Dup00, Dup04, Dup06]. The KPZ relation expresses that Euclidean exponents are given by a quadratic function of their quantum analogs. In particular, this suggests that an inverse quadratic function of the analogous Euclidean exponents should be a linear function of n; this latter fact was obtained (and termed the "cascade relation") in the rigorous work by Lawler and Werner on Brownian intersection exponents [LW99, LW00]

We also have the following:

Theorem 9.11. *In the construction of Theorem 9.9, both $\mathcal{W}$ and the interface η are uniquely determined by the $\mathcal{W}_i$ and may be obtained by a* conformal welding *of the right side of $\mathcal{W}_1$ with the left side of $\mathcal{W}_2$, where each is parameterized by γ-LQG boundary length.*

We will not give a separate proof of Theorem 9.11 since it follows from the same argument used to prove [She10, Theorems 1.3, 1.4] and the removability results described in Section 9.6. (The proof of [She10, Theorem 1.4] is given in [She10, Section 1.4].)

As mentioned briefly above, there is a natural generalization of Theorem 9.9 in which one cuts a quantum wedge $\mathcal{W}$ with n SLE processes (as opposed to cutting with a single path) coupled together as flow lines with varying angles of a GFF [MS12a, MS12b, MS12c, MS13a] which is independent of $\mathcal{W}$. In this case, one obtains $n+1$ independent wedges and the sum of their weights is equal to the weight of $\mathcal{W}$. This result is illustrated in Figure 9.11 and is stated precisely in Proposition **??**. If one lets the number of paths tend to ∞ (with the spacing between the angles going to zero), then the collection of paths converges to the so-called **fan $\mathbf{F}$**. (See Figures 1.2–1.5 of [MS12a] for simulations of $\mathbf{F}$.) This is a random closed set which almost surely has zero Lebesgue measure [MS12a, Proposition 7.33]. As shown in [DMS14], we are able to deduce an exact Poissonian structure of the countable collection of surface "beads" parameterized by the complement of $\mathbf{F}$. In particular, we can interpret $\mathbf{F}$ as describing the interface that arises when one glues together a Poissonian collection of wedges with weight 0.

Our next main result implies that a quantum cone can be constructed by identifying the left and right sides of a quantum wedge according to γ-LQG boundary length.

Theorem 9.12 (Welding and cutting for quantum cones)**.** *Fix $\gamma \in (0,2)$, let $\kappa = \gamma^2 \in (0,4)$, and suppose that $\mathcal{C} = (\mathbf{C}, h, 0, \infty)$ is a quantum cone of weight $W > 0$. Let $\rho = W - 2$ and suppose that η is a whole-plane $\mathrm{SLE}_\kappa(\rho)$ process independent of $\mathcal{C}$ starting from 0. Then the quantum surface $\mathcal{W}$ described by $(\mathbf{C} \setminus \eta, h, 0, \infty)$ is a quantum wedge of weight W. Moreover, the pair consisting of $\mathcal{C}$ and η is almost surely determined by $\mathcal{W}$ and can be obtained by conformally welding the left boundary of $\mathcal{W}$ with its right boundary according to γ-LQG boundary length.*

See Figure 9.12 for an illustration of Theorem 9.12. This result is stated in the case of simple $\mathrm{SLE}_\kappa(\rho)$ processes ($W \geq \frac{\gamma^2}{2}$ so that $\rho \geq \frac{\kappa}{2} - 2$) in Proposition **??** and in the case of self-intersecting $\mathrm{SLE}_\kappa(\rho)$ processes ($W \in (0, \frac{\gamma^2}{2})$ so that $\rho \in (-2, \frac{\kappa}{2} - 2)$) in [DMS14]. Moreover, in [DMS14] it is explained that slicing a quantum cone with a collection of whole-plane $\mathrm{SLE}_\kappa(\rho)$ processes coupled together as flow lines of a whole-plane GFF starting from the origin [MS13a] yields a collection of independent wedges.

If we parameterize our cone in terms of α rather than W, then it follows from (9.5) that the value of ρ from Theorem 9.12 is given by

$$\rho = 2 + \gamma^2 - 2\alpha\gamma. \tag{9.6}$$

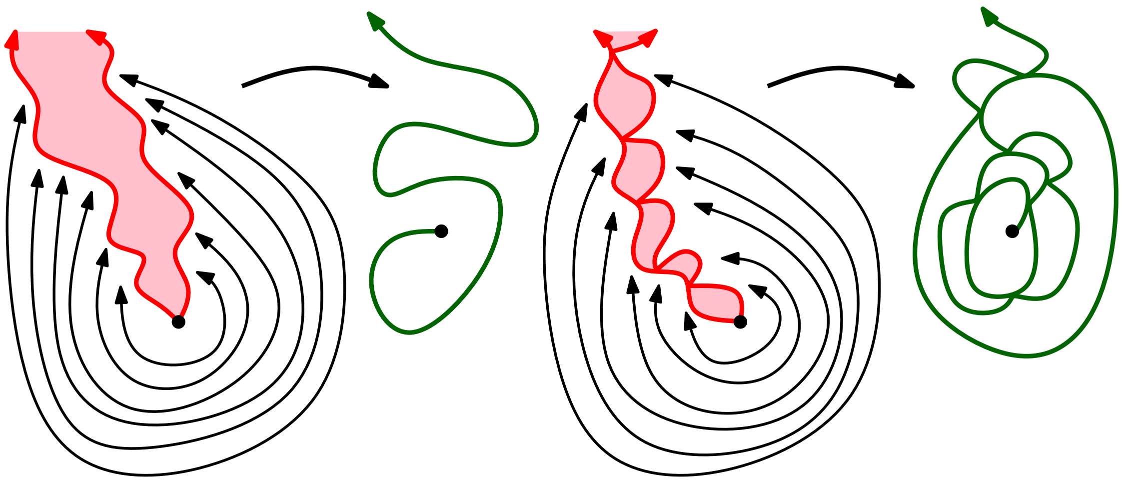

Figure 9.12: Illustration of Theorem 9.12. **Left:** A quantum wedge of weight $W \geq \frac{\gamma^2}{2}$ (hence homeomorphic to $\mathbf{H}$). If we weld together its left and right sides according to γ-LQG boundary length, then the resulting surface is a quantum cone of the same weight W (right side) decorated with an independent whole-plane $\mathrm{SLE}_\kappa(W-2)$ process. **Right:** The same statement holds in the case that $W \in (0, \frac{\gamma^2}{2})$ so that the wedge is not homeomorphic to $\mathbf{H}$ but rather is given by a Poissonian sequence of disks; in this case, $W - 2 \in (-2, \frac{\kappa}{2} - 2)$ so that the $\mathrm{SLE}_\kappa(W-2)$ process is self-intersecting.

Theorem 9.12 combined with (9.4) and (9.5) tells us that zipping up the left and right sides of an α'-quantum wedge yields an α-quantum cone with

$$\alpha = \frac{\alpha'}{2} + \frac{1}{\gamma}. \tag{9.7}$$

See Table 3 for the conversion between several parameterizations for quantum wedges and cones when performing the welding or cutting operation from Theorem 9.12.

9.8 Matings of trees and trees of loops

9.8.1 Main result on gluing infinite volume CRTs

In this section, we describe in much greater detail the manner in which the Brownian motion pair (X_t, Y_t), as discussed in Section 9.2, encodes various objects within LQG and SLE. Here and throughout much of the rest of the paper, we will use the symbols (L_t, R_t) in place of (X_t, Y_t) in order to highlight the fact that (as we will explain) in many circumstances, L_t and R_t can be interpreted as left and right boundary lengths of the quantum surface parameterized by $\{s : s \leq t\}$, or of the quantum surface parameterized by $\{s : s \geq t\}$. We include many resources/wedges/figures/ (Figures 9.13–9.22) which

(a) Conversion from wedge to cone parameterizations when zipping up a quantum wedge to obtain a quantum cone.

	α	W	θ	δ	a
α'	$\frac{1}{2}\alpha + \frac{1}{\gamma}$	$Q - \frac{1}{2\gamma}W$	$Q - \chi\frac{\theta}{2\pi}$	$Q + \frac{\gamma}{4}(1-\delta)$	$Q - \frac{\gamma}{4} - \frac{1}{2}a$
θ'	$\frac{\pi}{\chi}(\frac{\gamma}{2} + Q - \alpha)$	$\frac{\pi}{\chi\gamma}W$	θ	$\frac{\pi\gamma}{2\chi}(\delta - 1)$	$\frac{\pi}{\chi}(a + \frac{\gamma}{2})$
δ'	$3 + \frac{2}{\gamma}(Q - \alpha)$	$2 + \frac{2}{\gamma^2}W$	$2 + \chi\frac{2}{\gamma}\frac{\theta}{\pi}$	$\delta + 1$	$3 + \frac{2}{\gamma}a$
a'	$Q - \frac{1}{\gamma} - \frac{1}{2}\alpha$	$\frac{1}{2\gamma}W$	$\chi\frac{\theta}{2\pi}$	$\frac{\gamma}{4}(\delta - 1)$	$\frac{\gamma}{4} + \frac{1}{2}a$
ρ'	$\gamma^2 - \alpha\gamma$	$W - 2$	$\chi\gamma\frac{\theta}{\pi} - 2$	$\frac{\gamma^2}{2}(\delta - 1) - 2$	$\gamma a + \frac{\gamma^2}{2} - 2$

(b) Conversion from cone to wedge parameterizations when cutting a quantum cone to obtain a quantum wedge.

	α'	θ'	δ'	a'	ρ'
α	$2\alpha' - \frac{2}{\gamma}$	$Q + \frac{\gamma}{2} - \chi\frac{\theta'}{\pi}$	$Q + \frac{\gamma}{2}(3 - \delta')$	$Q + \frac{\gamma}{2} - 2a'$	$\gamma - \frac{1}{\gamma}\rho'$
W	$2\gamma(Q - \alpha')$	$\gamma\chi\frac{\theta'}{\pi}$	$\frac{\gamma^2}{2}(\delta' - 2)$	$2\gamma a'$	$2 + \rho'$
θ	$\frac{2\pi}{\chi}(Q - \alpha')$	θ'	$\frac{\pi\gamma}{2\chi}(\delta' - 2)$	$\frac{2\pi}{\chi}a'$	$\frac{\pi}{\chi\gamma}(\rho' + 2)$
δ	$1 + \frac{4}{\gamma}(Q - \alpha')$	$1 + \chi\frac{2}{\gamma}\frac{\theta'}{\pi}$	$\delta' - 1$	$1 + \frac{4}{\gamma}a'$	$1 + \frac{2}{\gamma^2}(\rho' + 2)$
a	$\frac{2}{\gamma} + Q - 2\alpha'$	$\chi\frac{\theta'}{\pi} - \frac{\gamma}{2}$	$\frac{\gamma}{2}(\delta' - 3)$	$2a' - \frac{\gamma}{2}$	$\frac{1}{\gamma}(\rho' + 2) - \frac{\gamma}{2}$

Table 3: If we zip up the left and right sides of a quantum wedge according to γ-LQG boundary length, then by Theorem 9.12 we get a quantum cone decorated with an independent whole-plane $\mathrm{SLE}_\kappa(\rho)$ process, and, conversely, if we cut a quantum cone with an independent whole-plane $\mathrm{SLE}_\kappa(\rho)$ then we get a quantum wedge. **a)** Each of the entries in the first row gives a way of parameterizing the wedge and each entry in the first column gives a way of parameterizing the resulting cone. The variable ρ' refers to the ρ-value for the whole-plane $\mathrm{SLE}_\kappa(\rho)$ process which is the image of $\mathbf{R}$ under the zipping up map. Each entry of the table gives the type of cone that one gets by zipping up a wedge where the cone is described using the parameterization corresponding to the row of the entry and the wedge is described using the parameterization of the column of the entry. Note that the row corresponding to θ' is identical to the row corresponding to θ in Table 1. This follows because the "space of angles" in the imaginary geometry for a given wedge does not change under the zipping up operation. **b)** Each entry of the table gives the type of wedge one gets by cutting a cone with an appropriate SLE process where the cone is described using the parameterization from the column and the wedge is parameterized using the variable from the corresponding row.

illustrate our most fundamental infinite volume peanosphere construction, as stated in Theorem 9.13 and Theorem 9.14 below.

Recall that space-filling $\mathrm{SLE}_{\kappa'}$ for $\kappa' \in (4,8)$ is a variant of $\mathrm{SLE}_{\kappa'}$ which fills up the components that it separates from its target point. This process was first constructed and analyzed in [MS13a]. Space-filling $\mathrm{SLE}_{\kappa'}$ for $\kappa' \geq 8$ is the same as ordinary $\mathrm{SLE}_{\kappa'}$. A precise definition of space-filling $\mathrm{SLE}_{\kappa'}$ is given in [MS13a, Section 1.2.3]; see in particular [MS13a, Figure 1.16] and the surrounding text. Many of the important properties of space-filling $\mathrm{SLE}_{\kappa'}$ are also established in [MS13a] (existence, continuity, reversibility, and its relationship to $\mathrm{CLE}_{\kappa'}$). It is also explained in [MS13a] how space-filling $\mathrm{SLE}_{\kappa'}$ can be interpreted as tracing (clockwise or counterclockwise) the outside of a (space-filling) "flow line tree" corresponding to a given instance of the Gaussian free field. In particular, if one considers the whole-plane GFF, then the corresponding path is what we will call a **whole-plane space-filling** $\mathrm{SLE}_{\kappa'}$ **from ∞ to ∞**.[26] Given this structure, one can also recover the non-space-filling $\mathrm{SLE}_{\kappa'}$ curves, which are called **counterflow lines** in [MS12a, MS12b, MS12c, MS13a] and studied in detail there. Theorem 9.13 describes the evolution of the left and right boundaries of a whole-plane space-filling $\mathrm{SLE}_{\kappa'}$ from ∞ to ∞ drawn on top of a quantum cone (see Figure 9.13).

Theorem 9.13. *Let $\mathcal{C} = (\mathbf{C}, h, 0, \infty)$ be a γ-quantum cone (which corresponds to $W = 4 - \gamma^2$ and $\theta = 2\pi$) together with a space-filling* $\mathrm{SLE}_{\kappa'}$ *process η' from ∞ to ∞ sampled independently of $\mathcal{C}$ and then reparameterized according to γ-LQG area. That is, for $s,t \in \mathbf{R}$ with $s < t$ we have that $\mu_h(\eta'([s,t])) = t - s$. Let L_t (resp. R_t) denote the change in the length of the left (resp. right) boundary of η' relative to time 0. Then (L,R) evolves as a two-dimensional Brownian motion. In the case that $\kappa' \in (4,8]$, we have (up to a non-random linear reparameterization of time) that*

$$\mathrm{Var}(L_t) = t, \quad \mathrm{Var}(R_t) = t, \quad \textit{and} \quad \mathrm{Cov}(L_t, R_t) = -\cos\left(\frac{4\pi}{\kappa'}\right) t \geq 0 \quad \textit{for} \quad t \geq 0.$$

Moreover, the joint law of (h, η') is invariant under shifting by t units of (γ-LQG area) time and then recentering. That is, for each $t \in \mathbf{R}$ we have that

$$(h, \eta') \stackrel{d}{=} (h(\cdot - \eta'(t)), \eta'(\cdot - t) - \eta'(t)). \tag{9.8}$$

[26] The chordal version of space-filling $\mathrm{SLE}_{\kappa'}$ (i.e., in $\mathbf{H}$ from 0 to ∞) is discussed in detail in [MS13a]. The whole-plane version from ∞ to ∞ is not explicitly constructed and shown to be continuous in [MS13a]. However, it can be easily be constructed and shown to be continuous using chordal space-filling $\mathrm{SLE}_{\kappa'}$. To accomplish this, one first starts flow and dual flow lines of a whole-plane GFF starting from 0. If $\kappa' \geq 8$, these flow lines will partition space into two regions which are each homeomorphic to $\mathbf{H}$. In this case, a whole-plane space-filling $\mathrm{SLE}_{\kappa'}$ from ∞ to ∞ can be constructed by splicing together two chordal space-filling $\mathrm{SLE}_{\kappa'}$'s, one for each of the two regions. The first path is taken to run from ∞ to 0 and the second from 0 to ∞. If $\kappa' \in (4,8)$, then the flow and dual flow lines started from 0 will partition space into a countable collection of pockets. In this case, a whole-plane space-filling $\mathrm{SLE}_{\kappa'}$ from ∞ to ∞ can be constructed by splicing together a countable collection of chordal space-filling $\mathrm{SLE}_{\kappa'}$'s, one for each of the pockets.

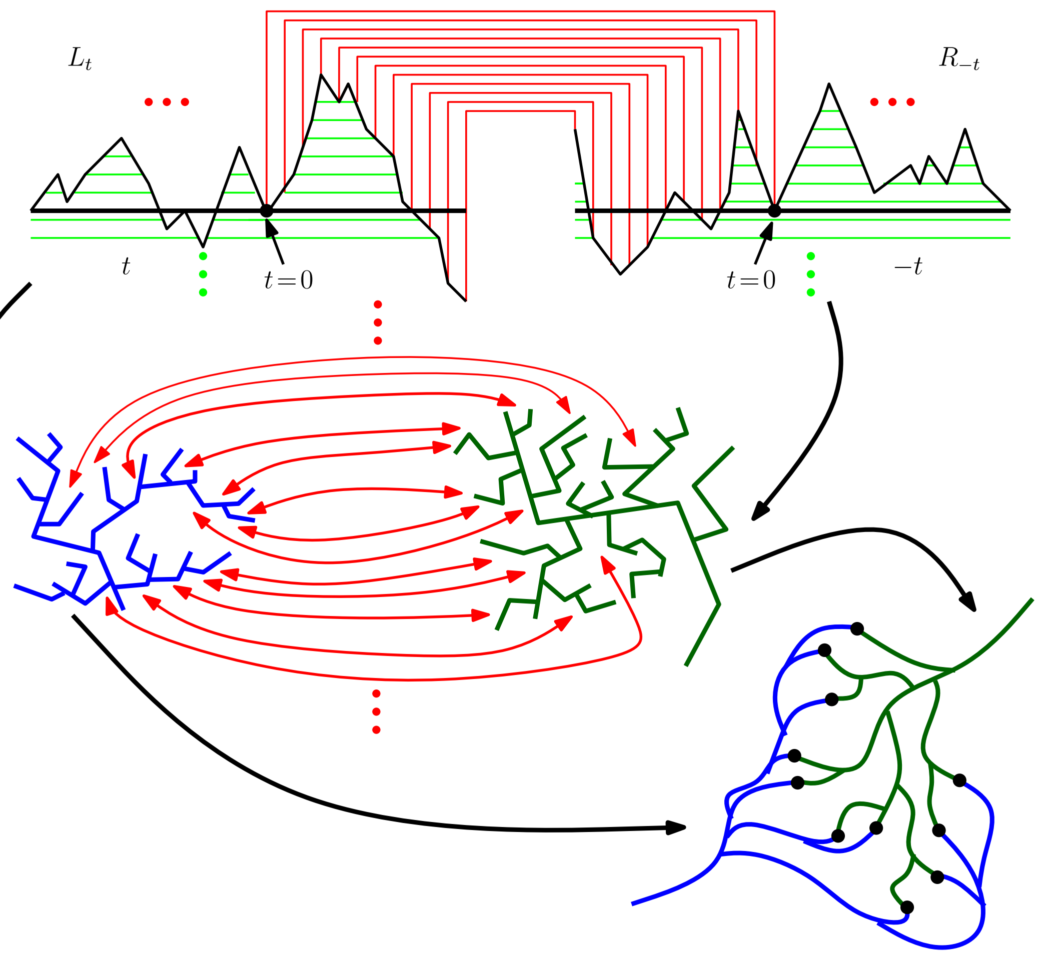

Figure 9.13: Gluing together two infinite volume space-filling trees, encoded by correlated Brownian motions L_t and R_t, produces a γ-quantum cone (i.e., $W = 4 - \gamma^2$ and $\theta = 2\pi$). This cone is decorated by a space-filling $\mathrm{SLE}_{\kappa'}$ from ∞ to ∞, which in turn encodes the east-going and west-going rays of an imaginary geometry.

Finally, the quantum surfaces parameterized by $\eta'([0, \infty])$ and $\eta'([-\infty, 0])$ are independent quantum wedges, each with parameter $\theta = \pi$, and $W = 2 - \frac{\gamma^2}{2}$.

As we explain in [DMS14], the main inputs into the proof of Theorem 9.13 are Theorem 9.9 and Theorem 9.12. (See also Figure 9.14.) Indeed, these results imply that drawing a certain pair of whole-plane $\mathrm{SLE}_\kappa(2 - \kappa)$ processes coupled together as flow lines of a whole-plane GFF [MS13a] on top of an independent γ-quantum cone ($W = 4 - \gamma^2$) yields a pair of independent quantum wedges of weight $2 - \frac{\gamma^2}{2}$. These flow lines give the left and right boundaries of η' stopped upon hitting 0. This, combined with the invariance statement (9.8) implies that (L, R) has independent increments and it does not require much additional work to extract from this that (L, R) must be *some*

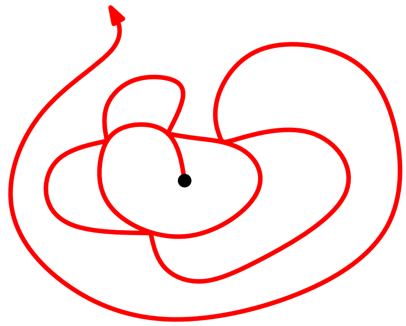

(a)

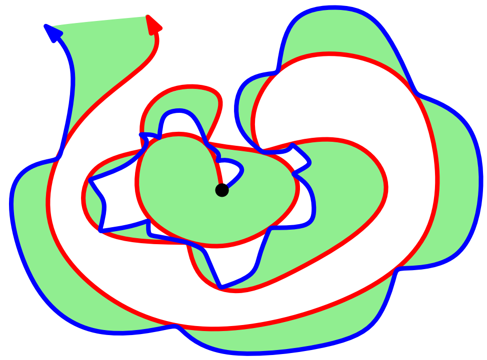

(b)

Figure 9.14: Illustration of the key step in the proof of Theorem 9.13 from Theorem 9.9 and Theorem 9.12. **a)** A γ-quantum cone sliced by an independent whole-plane $\mathrm{SLE}_\kappa(\rho)$ process η_1 with $\kappa = \gamma^2$ and $\rho = 2 - \kappa$. (Whether or not η_1 is self-intersecting depends on whether $\rho \in (-2, \frac{\kappa}{2} - 2)$ or $\rho \geq \frac{\kappa}{2} - 2$.) By Theorem 9.12, the sequence of quantum surfaces corresponding to $\mathbf{C} \setminus \eta_1$ ordered according to when their boundary is first drawn by η_1 is a wedge of weight $4 - \gamma^2$. **b)** Conditional on η_1, we draw in each of the components of $\mathbf{C} \setminus \eta_1$ an independent $\mathrm{SLE}_\kappa(-\frac{\kappa}{2}; -\frac{\kappa}{2})$ process; call their concatenation η_2. By the results of [MS13a], we can view (η_1, η_2) as flow lines of a common whole-plane GFF with an angle gap of π and (η_1, η_2) give the outer boundary of a space-filling $\mathrm{SLE}_{\kappa'}$, $\kappa' = 16/\kappa$, process η' stopped upon hitting 0. Theorem 9.9 implies that the pair (η_1, η_2) divides the plane into independent wedges of weight $2 - \frac{\gamma^2}{2}$. (These correspond to the regions of $\mathbf{C}$ visited by η' before and after it visits 0 and are respectively colored green and white.) This is the key observation that leads to the statement that the γ-LQG length of the left and right boundaries of η' (when parameterized by γ-LQG area) has independent increments and, ultimately, Theorem 9.13, which states that they evolve as a certain two-dimensional Brownian motion.

two-dimensional Brownian motion. In the case that $\kappa' \in (4, 8)$, we then compute the almost sure Hausdorff dimension of the set of local cut times for η' (with the γ-LQG area parameterization). These local cut times turn out to correspond to so-called "cone times" for (L, R), so we are able to determine the covariance matrix for $\kappa' \in (4, 8)$ (and for the limiting case $\kappa = 8$) by matching the dimension that we find with the dimension given in the main result of [Eva85]. In [DMS14], additional explanation is provided as to how this result relates to the scaling limits of discrete random planar map models [She11b] described earlier. We will not identify the correlation constant for $\kappa' > 8$ here, though it is natural to guess that the formula given in Theorem 9.13 continues to hold when $\kappa' > 8$.

Our next result is that the pair (L, R) from Theorem 9.13 almost surely determines both the space-filling $\mathrm{SLE}_{\kappa'}$ exploration path and the entire LQG surface.

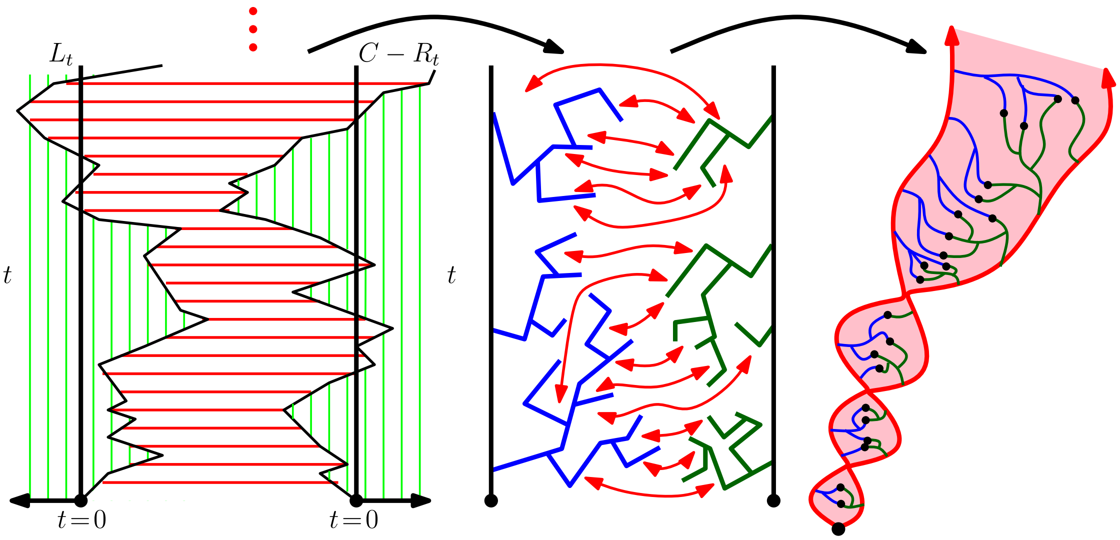

Figure 9.15: If we restrict the time in Figure 9.13 to $t \geq 0$, then the same Brownian motion encodes a $\theta = \pi$ wedge, which corresponds to $W = 2 - \frac{\gamma^2}{2}$. (Note that the vertical and horizontal axes have been swapped from what they were in Figure 9.13.) The left and right boundaries of the wedge correspond to the record minima of L_r and R_t (see also Figure 9.16, which in turn correspond to the vertical green segments that reach all the way to the bottom $t = 0$ line.

Theorem 9.14. *In the setting of Theorem 9.13, the pair (L, R) almost surely determines both η' and h (up to a rigid rotation of the complex plane about the origin).*

9.8.2 Non-space-filling counterflow lines

The peanosphere constructions are particularly interesting in the case that L_t and R_t are positively correlated, which corresponds to $\gamma \in (\sqrt{2}, 2)$. Recall that this is the range for which the corresponding $\mathrm{CLE}_{\kappa'}$ exist with $\kappa' \in (4, 8)$, where $\kappa' = 16/\gamma^2$. This is also the range in which $\mathrm{SLE}_{\kappa'}$ is itself non-space-filling and hence differs from space-filling $\mathrm{SLE}_{\kappa'}$.

Given $0 < s < t$, we say that s is an **ancestor** of t, and we write $s \prec t$, if for all $r \in (s, t]$ we have $L_r > L_s$ and $R_r > R_s$. The following facts are obvious from this definition:

1. $s \prec t$ implies $s < t$.
2. $s \prec t$ and $t \prec u$ implies $s \prec u$.
3. $s \prec t$ implies $s \prec u$ for all $u \in (s, t)$.

If s is an ancestor of any $t > s$, then s is called a **cone time** of the Brownian process. Figure 9.16 illustrates one such cone time. As the figure illustrates, the set of points that have a cone time s as an ancestor is an open set (s, s') for some s', and between s and s' the Brownian path (L_t, R_t) traces out an excursion into the quadrant $[L_s, \infty) \times [R_s, \infty)$ that begins at the corner and ends on one of the two sides. A point $t \geq 0$ is called **ancestor free** if there is no $s \in (0, t)$ that is an ancestor of t. The properties above imply that if s is ancestor free then a given $t > s$ is ancestor free if and only if t has no ancestor in (s, t). This implies that the set of ancestor free times is a regenerative process, and by scale invariance, we may conclude that it has the law of the range of a stable subordinator (which agrees in law with the zero set of a certain Bessel process, and which can be parameterized by a local time). Write $t(s)$ for the infimum of times at which this local time exceeds s (noting that $t(s)$ is necessarily an ancestor free time itself). Then we have the following:

Proposition 9.15. *The processes $L_{t(s)}$ and $R_{t(s)}$ parameterized by time s are independent totally asymmetric $\frac{\kappa'}{4}$-stable processes.*

The statement itself is a straightforward observation (it is clear from Figure 9.16 that each jump of the stable subordinator corresponds to a positive jump in precisely one of the processes $L_{t(s)}$ and $R_{t(s)}$, and that the measure on jumps has a power law distribution) and the parameter $\frac{\kappa'}{4}$ is identified in [DMS14]. As Figure 9.16 and 9.17 explain, the set of ancestor free times corresponds to the set of points in the space-filling curve from ∞ to 0 that lie on the outer boundary of the origin-containing component of the not-yet-explored region. The restriction of the path to these times is the ordinary $\mathrm{SLE}_{\kappa'}$ counterflow line from ∞ to 0. The following list summarizes a few special subsets of the wedge parameterized by $[0, \infty)$, which are obtained by restricting to special times in $[0, \infty)$:

1. **Counterflow line from ∞ to 0:** parameterized by the set of ancestor free times t.

2. **Left (resp. right) boundary of entire wedge:** parameterized by times at which L_t (resp. R_t) attains a record minimum.

3. **Cut points of entire wedge:** times at which L_t and R_t simultaneously achieve record minima. These times also correspond to points that are ancestors of 0 w.r.t. the time reversed process (L_{-t}, R_{-t}) parameterized by $t \leq 0$. See Figure 9.17.

The processes described in Proposition 9.15 encode two so-called **Lévy trees** of disks, after the manner outlined in Figure 9.7. (Lévy trees are studied in detail in [DLG02].) The procedure for obtaining one of these trees is explained in the top row of Figure 9.18. (The later rows contain related constructions that will relevant for Theorem 9.19.) Note that in a Lévy tree there are in fact a countably infinite number of small loops along the branch connecting any two given loops (i.e., it is almost surely the case that no two loops are adjacent). Each loop comes with a well-defined boundary length, which is the magnitude of the corresponding jump in the stable Lévy process. The outer boundary of the tree of loops also comes with a natural time parameterization, which is the time of the corresponding Lévy process. Intuitively, if for some tiny ϵ one keeps track of the number of loops of size between ϵ and 2ϵ that are encountered as one traces the boundary of the tree, then that number (times an appropriate power of ϵ) is a good approximation for this natural time. A **forested line** is the object obtained by beginning with a stable Lévy process as on the top row of Figure 9.18, and then filling each of the circles illustrated on the top row of Figure 9.18 with an independent quantum disk of the given boundary length, glued onto the circle in a boundary length preserving way.

Suppose that $\mathcal{W}$ is quantum wedge of weight $W > 0$. A **forested quantum wedge of weight** W is the beaded random surface which arises by gluing the line of a forested line (i.e., an independent forest of quantum disks as constructed above where the value of α for the stable process is taken to be $\frac{\kappa'}{4} = \frac{4}{\gamma^2} \in (1,2)$) to either the left or the right side of $\mathcal{W}$. A **doubly forested quantum wedge of weight** W is obtained by gluing an independent forested line to each of the two sides of $\mathcal{W}$. Illustrations of doubly forested wedges appear, e.g., in Figures 9.21 and 9.22.

Theorem 9.16. *Consider a quantum wedge of weight $W = 2 - \frac{\gamma^2}{2}$ (which corresponds to $\theta = \pi$) and the counterflow line from ∞ to 0 as depicted in Figure 9.17 and Figure 9.19. Then this counterflow line divides the wedge into two independent forested lines, whose boundaries are identified with one another according to the natural time parameterization of the outer boundary of the corresponding Lévy trees. Moreover, given the two forested lines, it is a.s. possible to uniquely recover the quantum wedge and the counterflow line.*

Theorem 9.16 has many variants.

The following theorem is one fairly easy generalization of Theorem 9.16, which can be deduced as a consequence of Theorem 9.16 and the additivity of wedges developed in Theorems 9.9 and 9.11. It states essentially that one can zip together a right-forested wedge of weight W_1 and a left-forested wedge of weight W_2 (zipping along the forested sides) in order to obtain a wedge of weight $W_1 + W_2 + 2 - \frac{\gamma^2}{2}$ decorated by an appropriate $\mathrm{SLE}_{\kappa'}(\rho_1;\rho_2)$ process. We include this result in order to illustrate that general non-boundary-filling $\mathrm{SLE}_{\kappa'}(\rho_1;\rho_2)$ processes can be obtained by zipping together forested wedges.[27]

[27] It is also natural to give a description of the structure of the surfaces cut out by an $\mathrm{SLE}_{\kappa'}(\rho_1;\rho_2)$ process, $\kappa' \in (4,8)$, drawn on top of a wedge $\mathcal{W}$ of weight W as in (9.9) when one or both of ρ_1, ρ_2 are in $(-2, \frac{\kappa'}{2} - 4)$ so that η' almost surely fills part of the domain boundary. In this case, the quantum surfaces which are completely surrounded by η' and are on its left (resp. right) side still have a tree structure. It seems, however, that these trees are not independent of each other. Describing the law of this pair of trees falls outside of the scope of this article.

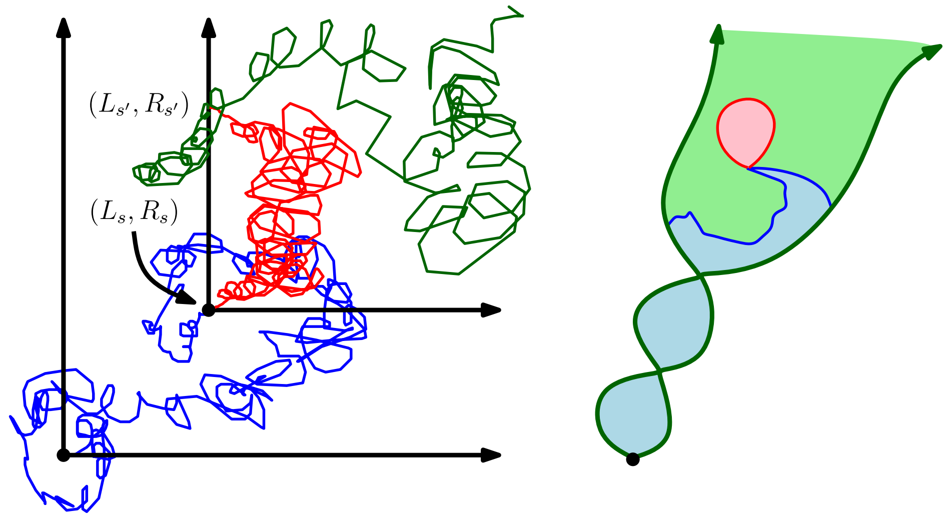

Figure 9.16: When L_t and R_t are positively correlated, we have $\gamma \in (\sqrt{2}, 2)$. In this case, there a.s. exist cone times like the time s illustrated on the left. As noted on the right, an interval of this type is "cut off" by the space-filling path from ∞ to 0 before it is filled up. (The path from ∞ to 0 fills first the green region, then the red region, then the blue region.)

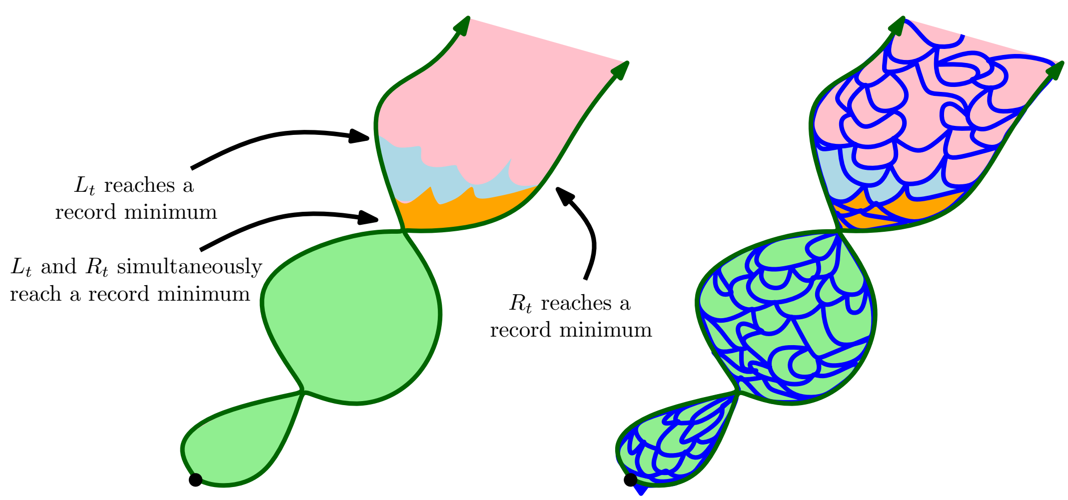

Figure 9.17: Shown on the left is a $2 - \frac{\gamma^2}{2}$ wedge. When the space-filling $\mathrm{SLE}_{\kappa'}$ process η' from 0 to ∞ hits a pinch point in the wedge, then L_t and R_t simultaneously hit a record minimum (green). When η' hits the left (resp. right) side of the wedge then L_t (resp. R_t) hits a record minimum. Shown on the right is the set of ancestor free times which corresponds to the counterflow line from ∞ to 0.

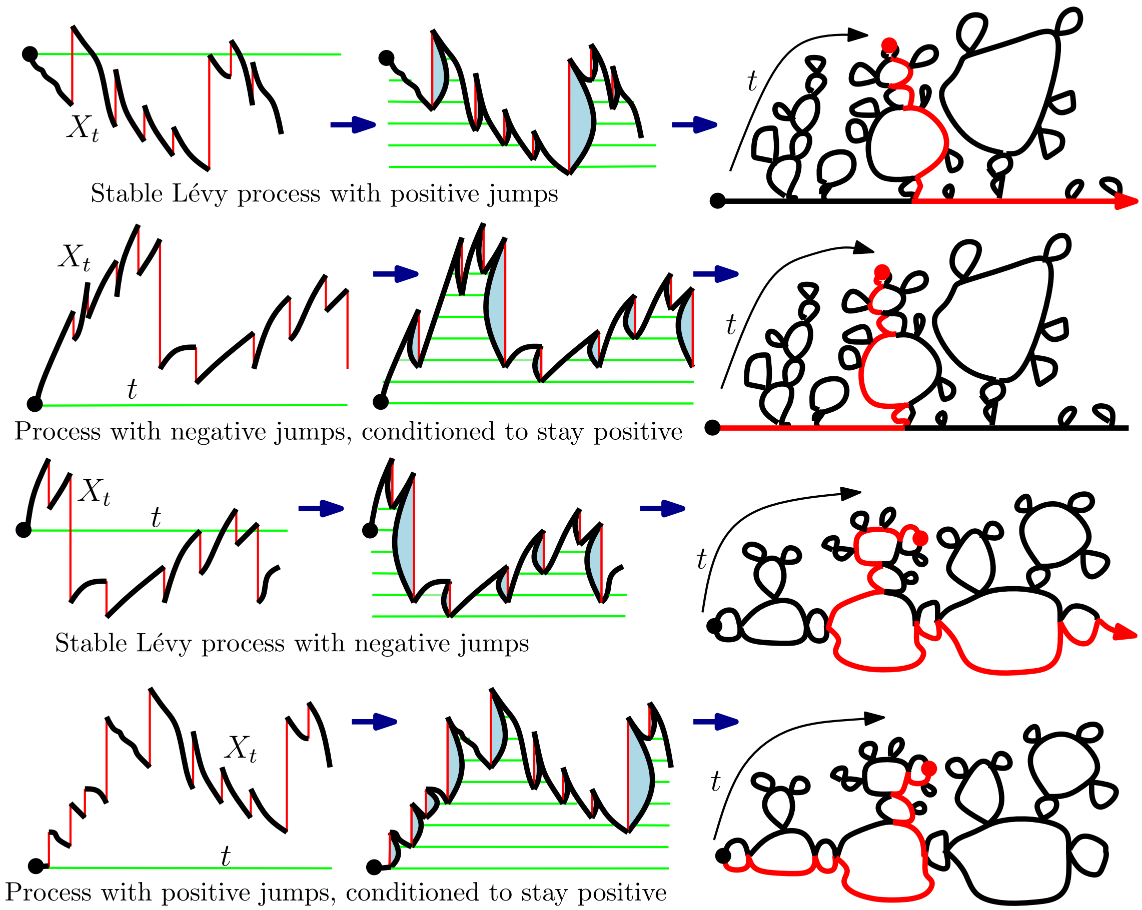

Figure 9.18: **First row:** A stable Lévy process X_t with positive jumps encodes a forested line via the procedure explained in Figure 9.7 (top right). Unmatched green segments (record minima of X_t) map to points on the line. As t increases, the red dot on right traces the forest boundary clockwise. X_t encodes the net change in length of the red path (which traces the right side of the branch of disks containing the point hit at time t, and continues right to ∞) since time 0. A jump occurs when a disk is hit for the first time, with jump size given by the boundary length of the disk. **Second row:** The same forested line corresponds to a stable Lévy process with negative jumps conditioned to stay positive. Unmatched green rays (last hitting times of X_t) map to points on the line. The value of X_t encodes the left boundary length. A jump occurs when a disk is hit for the last time. **Third row:** A stable Lévy process with negative jumps encodes a forested wedge. Unmatched green segments map to the underside of the wedge. A jump occurs when a disk is hit for the last time. **Fourth row:** The same forested wedge corresponds to a positive-jump stable Lévy process conditioned to stay positive.

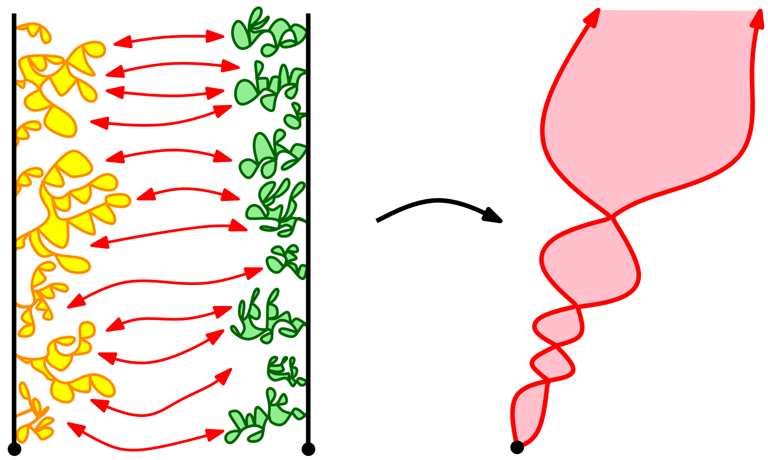

Figure 9.19: Two forested lines can be welded together according to quantum length to produce a single $\theta = \pi$ quantum wedge. Recall that $\theta = \pi$ corresponds to $W = 2 - \frac{\gamma^2}{2}$.

Theorem 9.17. *Fix $\gamma \in (0,2)$ and let $\kappa' = 16/\gamma^2 \in (4,\infty)$. Fix $\rho_1, \rho_2 \geq \frac{\kappa'}{2} - 4$. Let*

$$W_i = \gamma^2 - 2 + \frac{\gamma^2}{4}\rho_i \geq 0 \quad \text{for} \quad i = 1,2 \quad \text{and} \quad W = W_1 + W_2 + 2 - \tfrac{\gamma^2}{2}. \tag{9.9}$$

Let $\mathcal{W} = (\mathbf{H}, h, 0, \infty)$ be a quantum wedge of weight W and let η' be an $\mathrm{SLE}_{\kappa'}(\rho_1;\rho_2)$ process in $\mathbf{H}$ from 0 to ∞ with force points located at $0^-, 0^+$ which is independent of $\mathcal{W}$. (In the case that $W \in (0, \frac{\gamma^2}{2})$ so that $\mathcal{W}$ is not homeomorphic to $\mathbf{H}$, we take η' be to be given by a concatenation of independent $\mathrm{SLE}_{\kappa'}(\rho_1;\rho_2)$ processes: one for each of the bubbles of $\mathcal{W}$ starting at the opening point of the bubble and targeted at the terminal point.) Then the quantum surface $\mathcal{W}_1$ (resp. $\mathcal{W}_2$) which consists of those components of $\mathbf{H} \setminus \eta'$ which are to the left (resp. right) of η' is a quantum wedge of weight W_1 (resp. W_2) and the quantum surface $\mathcal{W}_3$ which is between the left and right boundaries of η' is a quantum wedge of weight $2 - \frac{\gamma^2}{2}$. (Note that $\mathcal{W}_1$ is beaded if $\rho_1 \in (\frac{\kappa'}{2} - 4, \frac{\kappa'}{2} - 2)$ and likewise $\mathcal{W}_3$ is beaded if $\rho_2 \in (\frac{\kappa'}{2} - 4, \frac{\kappa'}{2} - 2)$.) Moreover, $\mathcal{W}_1, \mathcal{W}_2, \mathcal{W}_3$ are independent.

Suppose further that $\gamma \in (\sqrt{2}, 2)$ so that $\kappa' \in (4,8)$. Then the beaded surface $\widetilde{\mathcal{W}}_1$ (resp. $\widetilde{\mathcal{W}}_2$) which consists of those components of $\mathbf{H} \setminus \eta'$ whose boundary is drawn by the left (resp. right) side of η' is a forested wedge of weight W_1 (resp. W_2). Moreover, $\widetilde{\mathcal{W}}_1$ and $\widetilde{\mathcal{W}}_2$ are independent and, together, almost surely determine both $\mathcal{W}$ and η'.

The following is another easy consequence of Theorem 9.16, together with Theorem 9.18. It states that we can zip up *both* sides of a doubly forested wedge of general weight, as illustrated in Figure 9.21, in order to obtain a quantum cone decorated by an appropriate whole-plane $\mathrm{SLE}_{\kappa'}(\rho)$ process.

Theorem 9.18. *Fix $\gamma \in (\sqrt{2}, 2)$ and suppose that $\mathcal{C} = (\mathbf{C}, h, 0, \infty)$ is a quantum cone of weight $W \geq 2 - \frac{\gamma^2}{2}$ (so that $\theta \geq \pi$). Let $\rho = \frac{4W}{\gamma^2} - 2$ and suppose that η' is a*

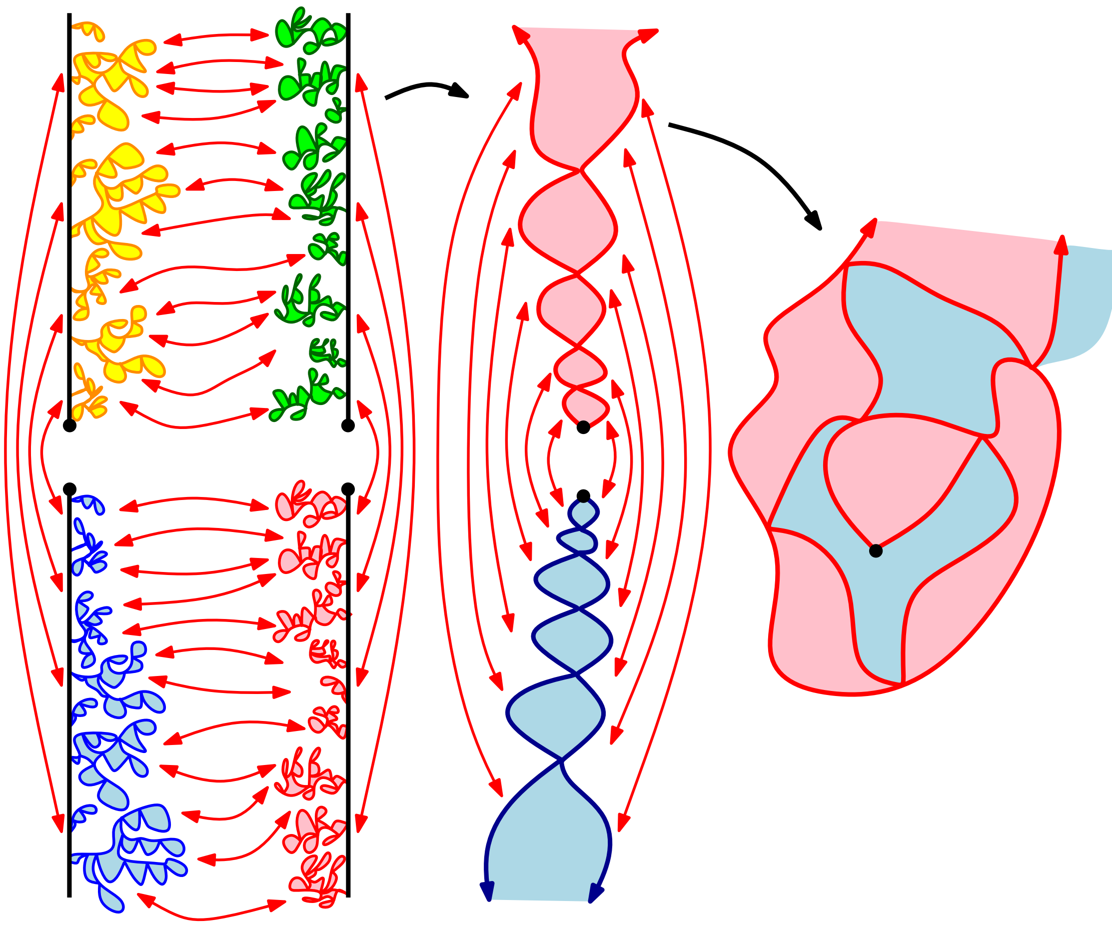

Figure 9.20: Repeating the procedure of Figure 9.19 for negative time, we can glue together four i.i.d. forested wedges to obtain the entire $\theta = 2\pi$ quantum cone. The interfaces are given by the flow line and dual flow line from 0 to ∞, and by the two counterflow lines coming from ∞ to 0 that have these paths as their boundaries.

whole-plane $\mathrm{SLE}_{\kappa'}(\rho)$ *process starting from* 0 *independent of* $\mathcal{C}$. *Then the beaded surface* $\mathcal{W}_1$ *(resp.* $\mathcal{W}_2$*) which consists of those components of* $\mathbf{C} \setminus \eta'$ *which are surrounded by* η' *on its left (resp. right) side when viewed as a path in the universal cover of* $\mathbf{C} \setminus \{0\}$ *has the structure of a forested line and the (possibly beaded) surface* $\mathcal{W}_3$ *which consists of the remaining components of* $\mathbf{C} \setminus \eta'$ *is a quantum wedge of weight*

$$W - \left(2 - \frac{\gamma^2}{2}\right) = \gamma^2 - 2 + \frac{\gamma^2}{4}\rho.$$

Moreover, $\mathcal{W}_1$, $\mathcal{W}_2$, *and* $\mathcal{W}_3$ *are independent and together almost surely determine both* $\mathcal{C}$ *and* η'.

Observe that if we express the ρ in the statement of Theorem 9.18 in terms of κ' and α

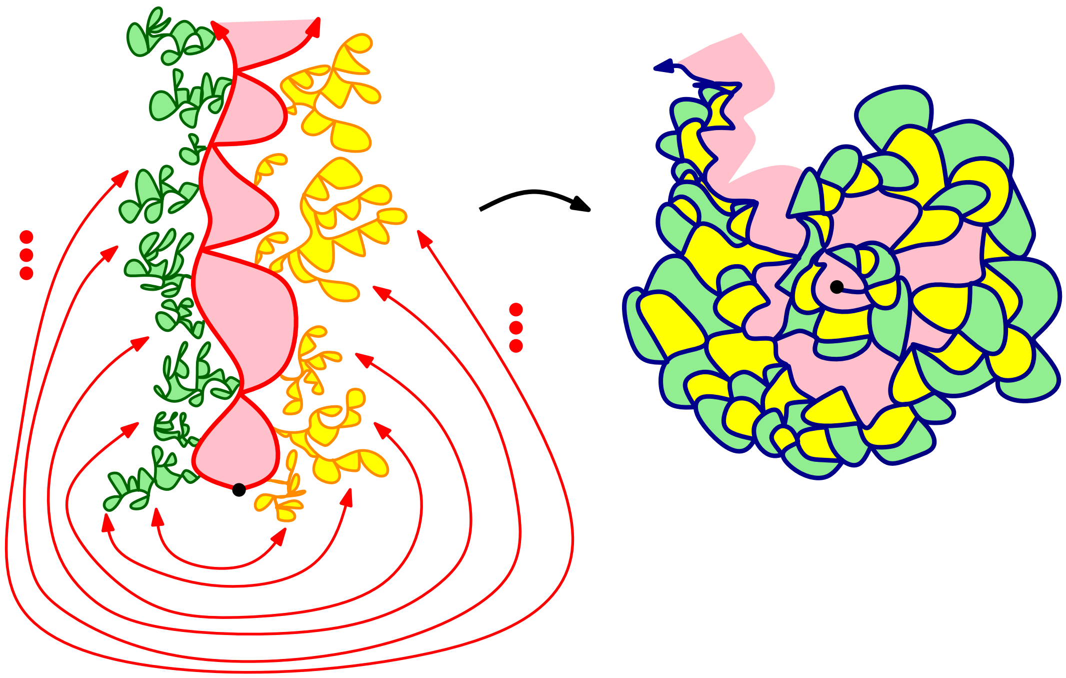

Figure 9.21: Alternatively, one can subdivide only one of the two $\theta = \pi$ wedges within a counterflow line. In this case, one has a doubly forested $\theta = \pi$ quantum wedge that zips up to become a $\theta = 2\pi$ quantum cone.

then we get

$$\rho = 2 + \kappa' - 2\alpha\sqrt{\kappa'}. \tag{9.10}$$

Note that (9.10) depends on κ' and α in the same way that (9.6) (with $\gamma^2 = \kappa$) from Theorem 9.12 depends on κ and α.

Theorem 9.19 is one of the more interesting and important consequences of Theorem 9.16. As illustrated in Figure 9.22, it implies a "quantum zipper" invariance principle for $\mathrm{SLE}_{\kappa'}$ analogous to the principle established for SLE_κ in [She10].

Combining Theorem 9.16 with Theorem 9.17 leads to another notion of time parameterization for an $\mathrm{SLE}_{\kappa'}$ process η', namely the time parameterization associated with the pair of independent stable Levy processes which encode the boundary lengths of the bubbles cut off by η'. We will refer to this time-parameterization as **quantum natural time**. It is the quantum version of the so-called "natural" parameterization for SLE [LS11, LZ13, LR12], see also [DS11b]. (The quantum analog of the natural parameterization for $\kappa \in (0,4)$ is γ-LQG length and for $\kappa' \geq 8$ is γ-LQG area.) We will write $\mathfrak{q}_u$ for the function which converts from quantum natural time u to capacity time. That is, if η' is parameterized by capacity time then $\eta'(\mathfrak{q}_u)$ is parameterized by quantum natural time.

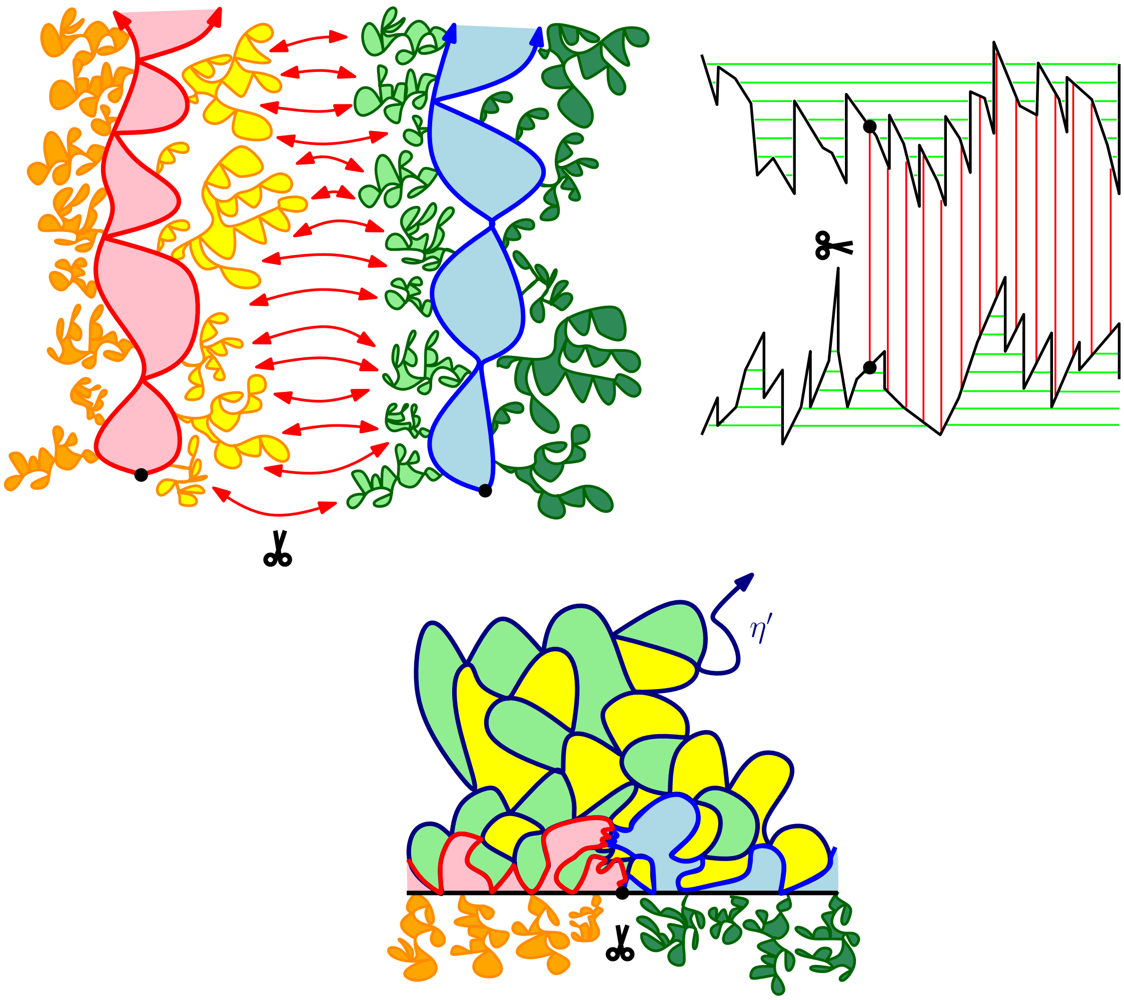

Figure 9.22: The quantum length invariant quantum zipper from [She10] involved two weight 2 quantum wedges which were zipped together on one side, and could be further zipped or unzipped; in the embedding in the upper half plane, the interface is an SLE_κ curve for $\kappa < 4$. In the analog described here for $\kappa' \in (4, 8)$ the two weight 2 quantum wedges are replaced by two doubly forested weight $(\gamma^2 - 2)$ wedges. The interface between them is a counterflow line (i.e., $\mathrm{SLE}_{\kappa'}$ curve), and once again the figure is invariant w.r.t. zipping or unzipping by a fixed amount of quantum natural time. Note that the entire collection of loops can be constructed from a pair of stable Lévy processes with negative jumps indexed by all time $\mathbf{R}$. The jumps before zero encode the loops in the two forests that have not yet been zipped up. The jumps after zero encode the loops in the upper half plane. Those jumps at times $t > 0$ at which L_t (resp. R_t) achieves a record minimum correspond to the disks that share boundary segments with the left (resp. right) real axis.

Theorem 9.19. *Fix $\gamma \in (\sqrt{2}, 2)$ and let $\mathcal{W} = (\mathbf{H}, h, 0, \infty)$ be a quantum wedge of weight $\frac{3\gamma^2}{2} - 2$ and let η' be an independent $\mathrm{SLE}_{\kappa'}$ process, $\kappa' = 16/\gamma^2 \in (4, 8)$, in $\mathbf{H}$ from 0 to ∞. Then η' divides $\mathcal{W}$ into two independent forested wedges both with weight $\gamma^2 - 2$. Moreover, the joint law of (h, η') is invariant under the operation of cutting along η' until a given quantum natural time and then conformally mapping back and applying (9.3). That is, if (f_t) denotes the centered chordal Loewner flow associated with η' (with the capacity time parameterization) and $\mathfrak{q}_u$ is as above, then we have that*

$$(h, \eta') \stackrel{d}{=} (h \circ f_{\mathfrak{q}_u}^{-1} + Q \log |(f_{\mathfrak{q}_u}^{-1})'|, f_{\mathfrak{q}_u}(\eta')) \quad \textit{for each} \quad u \geq 0.$$

Indeed, the entire image shown in Figure 9.22 is invariant with respect to the operation of zipping/unzipping according to quantum natural time.

If we start with the top row of Figure 9.18, and "zoom in" near a typical non-zero time, then we obtain a stable Lévy processes indexed by all of $\mathbf{R}$. We can then consider a pair of processes of this type. This is equivalent to drawing an SLE until a typical quantum time, zooming in, and "unzipping" up to the distinguished time, which produces a figure like Figure 9.22. Each of the forested lines hanging off the bottom of the image in Figure 9.22 is an independent forested line of the sort described by the top row of Figure 9.18. If we focus on the quantum wedge parameterized by the upper half plane in Figure 9.22 and then cut this wedge along the illustrated counterflow line, then we obtain a pair of forested quantum wedges, each of which corresponds, by construction, to the quantum wedge illustrated in the third line of Figure 9.18. In particular, this analysis tells us how the γ-LQG length of $\mathbf{R}_-$ and $\mathbf{R}_+$ changes when we zip and unzip the image shown in Figure 9.22.

Corollary 9.20. *In the context of Figure 9.22, the net change in the γ-LQG length of $\mathbf{R}_-$ and $\mathbf{R}_+$ as one "zips up" is given by an independent pair of totally asymmetric $\frac{\kappa'}{4}$-stable Lévy processes with positive jumps, each like the one illustrated in the first row of Figure 9.18. In the context of Figure 9.22, the net change in left and right boundary lengths as one "unzips" is given by an independent pair of totally asymmetric $\frac{\kappa'}{4}$-stable Lévy processes with negative jumps, each like the one illustrated in the third row of Figure 9.18.*

Recall that the first row of Figure 9.18 describes a forested line. Moreover, Corollary 9.20 implies that the third row of Figure 9.18 encodes a forested wedge of weight $\gamma^2 - 2$ corresponding to (by Theorem 9.17) those bubbles which are either to the right but not surrounded by an $\mathrm{SLE}_{\kappa'}$ process (a $\gamma^2 - 2$ wedge) or those bubbles which are on the right and completely surrounded by an $\mathrm{SLE}_{\kappa'}$ process (a forested line). One can glue the former forested line to the latter forested wedge to obtain a doubly forested wedge of weight $\gamma^2 - 2$. The law of this doubly forested wedge is invariant under the operation of "sliding the tip" a fixed amount of quantum time units along the boundary of the forest. This simply corresponds to the fact that the $\frac{\kappa'}{4}$-stable Lévy process indexed by all of $\mathbf{R}$ (which encodes this doubly forested wedge) has a stationary law. The second and fourth rows of Figure 9.18 describe different ways to encode the same process. Although we will not need this point here, we remark (and leave it to the reader to check) that as one moves from right to left along the processes in either the second or fourth row, one encounters or completes disks (which correspond to jumps) at a Poisson rate, and the processes shown can be understood as stable Lévy processes conditioned to stay positive.

We next remark that it is not hard to derive analogs of Theorem 9.19 and Corollary 9.20 that involve radial and whole-plane $\mathrm{SLE}_{\kappa'}$ processes. For these variants, one considers a single doubly forested wedge of some weight, zipped up like the doubly forested wedge in Figure 9.21, so that the interface becomes the $\mathrm{SLE}_{\kappa'}(\rho)$ curve described in Theorem 9.18. We then keep track of what happens as we begin to cut with scissors along the counterflow line, starting from the origin (always conformally mapping the infinite unexplored region conformally to $\mathbf{C} \setminus \mathbf{D}$). (This is the usual setup used to construct a whole-plane $\mathrm{SLE}_{\kappa'}(\rho)$ curve.) The case in which the wedge (along which the two forested lines are added) has weight $\gamma^2 - 2$ turns out to be particularly interesting, because in this case, the disks in the wedge turn out to play (in some sense) the same role as the disks in the forests rooted on that wedge.

9.8.3 Discrete intuition

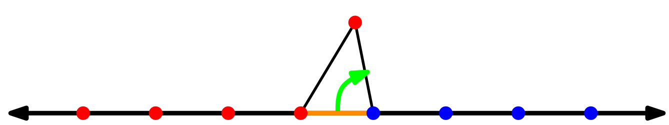

Figure 9.23: The horizontal line represents the (infinite) boundary of a domain Markov half-planar map with a marked boundary edge (orange). The boundary vertices to the left (resp. right) of this marked edge are colored red (resp. blue). The rest of the vertices in the map are colored red or blue i.i.d. with probability $\frac{1}{2}$. Shown is the first step of the percolation exploration starting from the marked edge with red (resp. blue) on the left (resp. right) side of the path. Triangles which have an edge on the boundary either have their third vertex on the boundary of the map or in the interior of the map. The triangle shown is of the latter type. By the domain Markov property, the conditional law of the map in the unbounded component given the revealed triangle is the same as the law of the original map.

Consider the **half-planar** random planar triangulation as defined in [AS03, Ang03,

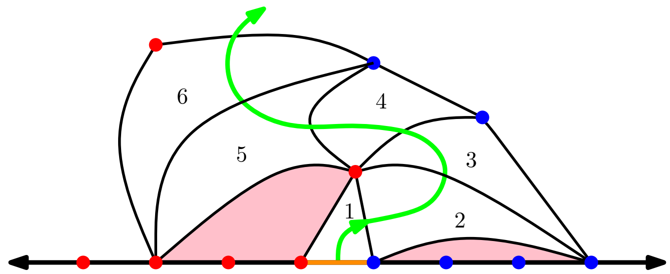

Figure 9.24: (Continuation of Figure 9.23.) Shown are five more steps of the percolation exploration (for a total of six steps). The triangles are numbered according to the order in which they are visited by the exploration path. When the exploration visits the second and fourth triangles it disconnects regions of the map from ∞; these are colored light red in the illustration. By the domain Markov property, the law of the sequence of these regions is i.i.d. Moreover, the domain Markov property implies that the incremental net changes to the left (right) boundary length are i.i.d.

AC13, AR13]. This is a simply-connected infinite planar map with an infinite boundary and a distinguished "origin" edge; it has the *domain Markov property*, which means that if we condition on any finite collection of triangles discovered by exploring from the boundary, then the conditional law of the infinite connected component of the unexplored region has the same law as the original map. (See [AR13, Definition 1.1] for a more careful definition.) Suppose that we pick a distinguished edge on the boundary of the map and then color the vertices which are to the left (resp. right) of this edge red (resp. blue) as illustrated in Figure 9.23. We then color the remaining vertices in the map i.i.d. red or blue with equal probability $\frac{1}{2}$. Consider the percolation exploration which starts at the marked boundary edge with red vertices on its left side and blue vertices on its right side. (See Figures 9.23–9.24.) By the domain Markov property, it follows that:

(i) The components which are cut off by the exploration form an i.i.d. sequence.

(ii) The change in the lengths of the left (resp. right) boundary of the unbounded component of the map evolve as random walks with independent increments.

Theorem 9.16 and Theorem 9.17 give the continuum analog of (i) and Corollary 9.20 gives the continuum analog of (ii).

The stable Lévy processes described in Proposition 9.15 and Corollary 9.20 for $\kappa' = 6$ are consistent with a heuristic argument made by Angel just before the statement of [Ang03, Lemma 3.1] for the scaling limit of the boundary length process associated with a percolation exploration on the uniform infinite planar triangulation. It is also consistent with [Cur13, Question 5]. The law of the boundary length process associated with the growth of a metric ball on $\sqrt{8/3}$-LQG should be related to the boundary length process associated with a percolation exploration process on $\sqrt{8/3}$-LQG via a certain time-change. In particular, if A_t denotes the boundary length process associated with the latter then the former should be given by $A_{t(s)}$ where $t(s) = \int_0^s A_u du$ (this says that the rate of growth should be proportional to boundary length, as in first passage percolation with i.i.d. exp(1) edge weights). If A is a totally asymmetric $\frac{3}{2}$-stable process conditioned to be non-negative, then $A_{t(s)}$ has the law of the time-reversal of a continuous-state branching process with branching mechanism $\psi(u) = u^{3/2}$ [Lam67b]. This is consistent with a result due to Krikun [Kri05, Theorem 4] for the evolution of the length of the boundary of a metric ball as its diameter increases in the setting of the uniform infinite planar quadrangulation. This is also consistent with a calculation carried out in the continuum for the Brownian plane [CL12] recently announced by Curien and Le Gall in [Le 14, CL14].

The results described in Section 9.8 will be important ingredients for a work in progress by the second two co-authors concerning the so-called quantum Loewner evolution (QLE) [MS13b]. In particular, they allow us to define a "quantum natural time" version of QLE. Since QLE$(8/3, 0)$ is constructed by applying a certain transformation to SLE_6 (so-called "tip-rerandomization"), Theorem 9.19 gives us that the Poissonian structure of the complementary components of a QLE$(8/3, 0)$ exploration and an SLE_6 exploration of $\sqrt{8/3}$-LQG are the same. Also, Theorem **??** gives us that the evolution of the $\sqrt{8/3}$-LQG length of the outer boundary of a QLE$(8/3, 0)$ is the same as the corresponding boundary length process for a metric ball in the Brownian plane. Finally, Theorem 9.13 gives many distributional identities between the Poissonian structure of the complementary components of a QLE$(8/3, 0)$ and a metric ball in the Brownian plane. As a simple example, it implies that the conditional law of the area of such a component given its boundary length is the same as in the setting of the Brownian plane.

10 Quantum Loewner evolution

10.1 QLE Overview

The mathematical physics literature contains several simple "growth models" that can be understood as random increasing sequences of clusters on a fixed underlying graph G, which is often taken to be a lattice such as $\mathbf{Z}^2$. These models are used to describe crystal formation, electrodeposition, lichen growth, lightning formation, coral reef formation, mineral deposition, cancer growth, forest fire progression, Hele-Shaw flow, water seepage, snowflake formation, oil dissipation, and many other natural processes. Among the most famous and widely studied of these models are the Eden model (1961), first passage percolation (1965), diffusion limited aggregation (1981), the dielectric breakdown model (1984), and internal diffusion limited aggregation (1986)

[Ede61, HW65, WJS81, WS83, NPW84, MD86], each of which was originally introduced with a different physical motivation in mind.

Here we mainly treat the dielectric breakdown model (DBM), which is a family of growth processes, indexed by a parameter η, in which new edges are added to a growing cluster according to the η-th power of harmonic measure, as we explain in more detail in Section 10.2[28]. DBM includes some of the other models mentioned above as special cases: when $\eta = 0$, DBM is equivalent to the Eden model, and when $\eta = 1$, DBM is equivalent to diffusion limited aggregation (DLA), as noted in [NPW84]. Moreover, first passage percolation (FPP) is a growing family of metric balls in a metric space obtained by assigning i.i.d. positive weights to the edges of G — and when the law of the weights is exponential, FPP is equivalent (up to a time change) to the Eden model (see Section 10.2.1).

We would like to consider what happens when G is taken to be a *random* graph embedded in the plane. Specifically, instead of using $\mathbf{Z}^2$ or another deterministic lattice (which in some sense approximates the Euclidean structure of space) we will define the DBM on random graphs that in some sense approximate the random measures that arise in Liouville quantum gravity.

Liouville quantum gravity (LQG) was proposed in the physics literature by Polyakov in 1981, in the context of string theory, as a canonical model of a random two-dimensional Riemannian manifold [Pol81b, Pol81c], although it is too rough to be defined as a manifold in the usual sense. By Riemann uniformization, any two-dimensional simply connected Riemannian manifold $\mathcal{M}$ can be conformally mapped to a planar domain D. If μ is the pullback to D of the area measure on $\mathcal{M}$, then the pair consisting of D and μ completely characterizes the manifold $\mathcal{M}$. One way to define an LQG surface is as the

[28]In [NPW84] growth is based on harmonic measure viewed from a specified boundary *set* within a regular lattice like $\mathbf{Z}^2$. For convenience, one may identify the points in the boundary set and treat them as a single vertex v. A cluster grows from a fixed interior vertex, and at each growth step, one considers the function ϕ that is equal to 1 at v and 0 on the vertices of the growing cluster — and is discrete harmonic elsewhere. The harmonic measure (viewed from v) of an edge $e = (v_1, v_2)$, with v_1 in the cluster and v_2 not in the cluster, is defined to be proportional to $\phi(v_2) - \phi(v_1) = \phi(v_2)$. We claim this is in turn proportional to the probability that a random walk started at v first reaches the cluster via e (which is the definition of harmonic measure we use for general graphs in this paper). We sketch the proof of this standard observation here in this footnote. On $\mathbf{Z}^2$, $\phi(v_2)$ is the probability that a random walk from v_2 reaches v before the cluster boundary, i.e., $\phi(v_2) = \sum_P 4^{-|P|}$ where P ranges over paths from v_2 to v that do not hit the cluster or v (until the end), and $|P|$ denotes path length. Also, for each P, the probability that a walk from v traces P in the reverse direction and then immediately follows e to hit the cluster is given by $4^{-|P|}/\deg(v)$. Summing over P proves that $\phi(v_2)$ is proportional to the probability that a walk starting from v exits at e without hitting v a second time; this is in turn proportional to the overall probability that a walk from v exits at e, which proves the claim. **Variants:** One common variant is to consider the first time a walk from v hits a cluster-adjacent vertex (instead of the first time it crosses a cluster-adjacent edge); this induces a fharmonic measure on cluster-adjacent vertices and one may add new vertices via the η-th power of this measure. The difference is analogous to the difference between site percolation and bond percolation. Also, it is often natural to consider harmonic measure viewed from ∞ instead of from a fixed vertex v.

pair D and μ with $\mu = e^{\gamma h(z)}dz$, where dz is Lebesgue measure on D, h is an instance of some form of the Gaussian free field (GFF) on D, and $\gamma \in [0,2)$ is a fixed parameter. Since h is a distribution, not a function, a regularization procedure is needed to make this precise [DS11a]. It turns out that one can define the mean value of h on a circle of radius ϵ, call this $h_\epsilon(z)$, and then write $\mu = \lim_{\epsilon \to 0} \epsilon^{\gamma^2/2} e^{\gamma h_\epsilon(z)} dz$ [DS11a] (and a slightly different construction works when $\gamma = 2$ [DRSV12a, DRSV12b]).

Figure 10.1 illustrates one way to tile D with squares all of which have size of order δ (for some fixed $\delta > 0$) in the random measure μ. Given such a tiling, one can consider a growth model on the graph whose vertices are the squares of this grid. Another more isotropic approach to obtaining a graph from μ is to sample a Poisson point process with intensity given by some large multiple of μ, and then consider the Voronoi tessellation corresponding to that point process. A third approach, which we explain in more detail below, is to consider one of the *random planar maps* believed to converge to LQG in the scaling limit.

We are interested in all three approaches, but ultimately, the main purpose of this paper is to produce a *candidate* for the scaling limit of an η-DBM process on a γ-LQG surface (in the fine mesh, or $\delta \to 0$ scaling limit). We expect that there is a universal scaling limit that does not depend on which approach we take (at least if the discrete setup is sufficiently isotropic; see the discussion in Section 10.2.2). Our goal is to show that (at least for some choices of γ and η) there exists a process, which we call quantum Loewner evolution QLE(γ^2, η), that has the dynamic properties that we would expect a scaling limit to have.

For certain values of the parameters γ^2 and η, which are illustrated in Figure 10.3, we will be able to explicitly describe a stationary law of the growth process in terms of quantum gravity. We will see that this growth process is similar to SLE except that the point-valued "driving function" that one feeds into the Loewner differential equation to obtain SLE$_\kappa$ (namely $\sqrt{\kappa}$ times Brownian motion on a circle) is replaced by a *measure*-valued driving function ν_t whose stationary law corresponds to a certain boundary measure that appears in Liouville quantum gravity. The time evolution of this measure is not nearly as easy to describe as the time evolution of Brownian motion, and making sense of this evolution is one of the main goals of this paper.

Let us explain this point a bit further. We will fix γ and an instance h of a free boundary GFF (plus a deterministic multiple of the log function) on the unit disk $\mathbf{D}$. We will interpret the pair $(\mathbf{D}, h)$ as a γ-LQG quantum surface and seek to define an increasing collection (K_t) of closed sets, indexed by $t \in [0,T]$ for some T, starting with $K_0 = \partial \mathbf{D}$ and growing inward within $\mathbf{D}$ toward the origin. We assume that each K_t is a **hull**, i.e., a subset of $\overline{\mathbf{D}}$ whose complement is a simply connected open set containing the origin. (Note that if a growth model grows outward toward infinity, one can always apply a conformal inversion so that the growth target becomes the origin.) We will require that for each t the set K_t is a so-called *local set* of the GFF instance h. This is a natural technical condition (see the more detailed treatment in [SS13]) that essentially states

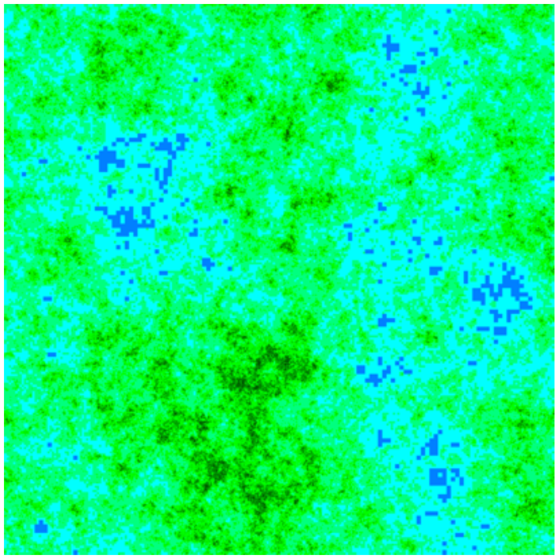

(a) $\gamma = 1/2$

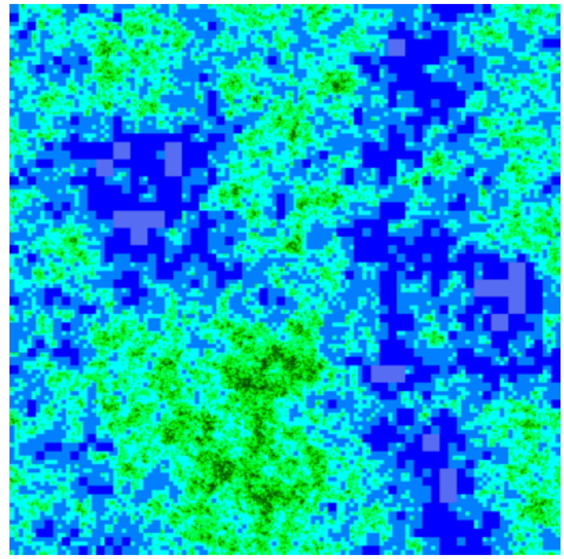

(b) $\gamma = 1$

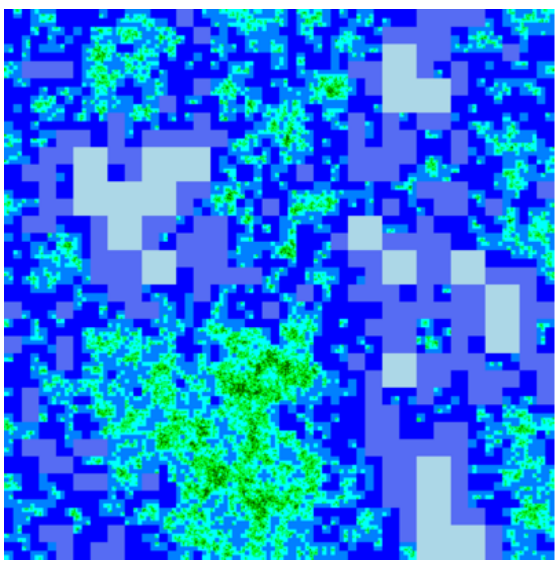

(c) $\gamma = 3/2$

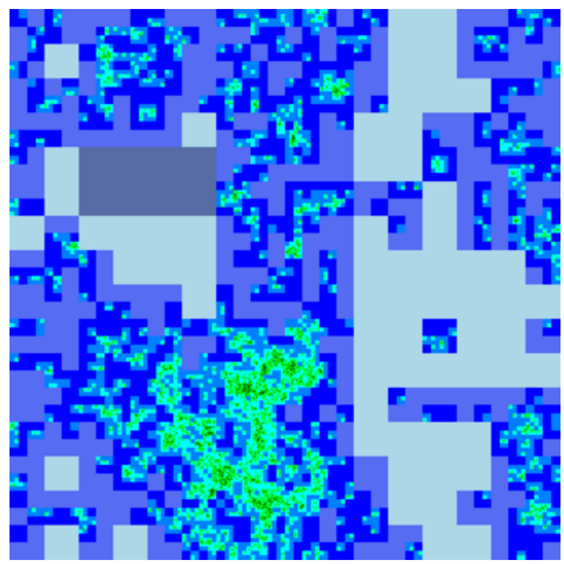

(d) $\gamma = 2$

(e) Number of subdivisions performed ranging from 0 (left) to 12 (right).

Figure 10.1: To construct the figures above, first an approximate γ-LQG measure μ was chosen by taking a GFF h on a fine $(4096 \times 4096 = 2^{12} \times 2^{12})$ lattice and constructing the measure $e^{\gamma h(z)}dz$ where dz is counting measure on the lattice (normalized so μ has total mass 1). Then a small constant δ is fixed (here $\delta = 2^{-16}$) and one divides the large square into four squares of equal Euclidean size, divides each of these into four squares of equal Euclidean size, etc., except that if any square's μ-area is less than δ, that square is not further divided. Each square remaining in the end has μ-area less than δ, but the μ-area of its dyadic parent is greater than δ. Squares are colored by Euclidean size.

that altering h on an open set $S \subseteq \mathbf{D}$ does not affect the way that K_t grows *before* the first time that K_t reaches S. In order to describe these growing sets K_t, we will construct a solution to a type of differential equation imposed on a triple of processes, each of which is indexed by a time parameter $t \in [0, T]$, for some fixed $T > 0$:

1. A measure ν on $[0, T] \times \partial\mathbf{D}$ whose first coordinate marginal is Lebesgue measure. We write ν_t for the conditional probability measure (defined for almost all t) obtained by restricting ν to $\{t\} \times \partial\mathbf{D}$. Let $\mathcal{N}_T$ be the space of measures ν of this type.

2. A family (g_t) of conformal maps $g_t \colon \mathbf{D} \setminus K_t \to \mathbf{D}$, where for each t the set K_t is a closed subset of $\overline{\mathbf{D}}$ whose complement is a simply connected set containing the origin. We require further that the sets K_t are increasing, i.e. $K_s \subseteq K_t$ whenever $s \leq t$, and that for all $t \in [0, T]$ we have $g_t(0) = 0$ and $g_t'(0) = e^t$. That is, the sets (K_t) are parameterized by the negative log conformal radius of $\mathbf{D} \setminus K_t$ viewed from the origin.[29] Let $\mathcal{G}_T$ be the space of families of maps (g_t) of this type.

3. A family $(\mathfrak{h}_t)$ of harmonic functions on $\mathbf{D}$ with the property that $\mathfrak{h}_t(0) = 0$ for all $t \in [0, T]$ and the map $[0, T] \times \mathbf{D} \to \mathbf{R}$ given by $(t, z) \to \mathfrak{h}_t(z)$ is jointly continuous in t and z. Let $\mathcal{H}_T$ be the space of harmonic function families of this type.

The differential equation on the triple $(\nu_t, g_t, \mathfrak{h}_t)$ is a triangle of maps between the sets $\mathcal{N}_T$, $\mathcal{G}_T$, and $\mathcal{H}_T$ that describes how the processes in the triple $(\nu_t, g_t, \mathfrak{h}_t)$ are required to be related to each other, as illustrated in Figure 10.2 and further explained below. The triple involves a constant α that for now is unspecified. The constant η will actually emerge *a posteriori* as a scaling symmetry of the map from $\mathcal{H}_T$ to $\mathcal{N}_T$ that applies almost surely to the triples we construct. We will see that when an LQG coordinate change is applied to $(\nu_t, \mathbf{D})$ (a change that preserves quantum boundary length but changes harmonic measure viewed from zero) ν_t is locally rescaled by the derivative of the map to the $2 + \eta$ power; Figure 10.4 explains heuristically why the scaling limit of η-DBM on a γ-LQG should have such a symmetry. The definition of η and its relationship to α will be explained in more detail in Section 10.4 and Section 10.10.

1. $\mathcal{N}_T \to \mathcal{G}_T$: the process (g_t) is obtained as the radial Loewner flow driven by (ν_t), as further explained in Section 10.3. It turns out (see Theorem 10.1) that Loewner evolution describes a one-to-one map from $\mathcal{N}_T$ to $\mathcal{G}_T$.

2. $\mathcal{G}_T \to \mathcal{H}_T$: for each t, the function $\mathfrak{h}_t$ is obtained from h by first letting $P^t_{\mathrm{harm}}(h)$ be the harmonic extension of the values of h from $\partial(\mathbf{D} \setminus K_t)$ to $\mathbf{D} \setminus K_t$, and then letting $\mathfrak{h}_t$ be the harmonic function on $\mathbf{D}$ defined by $\mathfrak{h}_t = P^t_{\mathrm{harm}}(h) \circ (g_t^{-1}) + Q \log |(g_t^{-1})'|$.

[29] (-1) times the log of the conformal radius of $\mathbf{D} \setminus K_t$, viewed from the origin, is also called the *capacity* of the K_t (though we caution that the term "capacity" has several other meanings in other contexts).

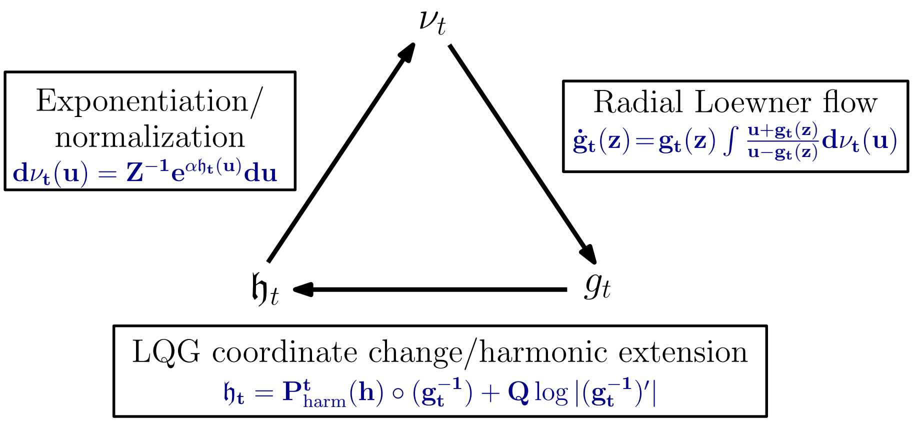

Figure 10.2: Visual sketch of the differential equation for the QLE dynamics. The map that takes the process ν_t to the process g_t is the most straightforward to describe. It is simply Loewner evolution, and works for any ν_t we would want to consider, see Theorem 10.1. This is the "differential" part of the equation, since ν_t determines the time derivative of g_t. The map from g_t to $\mathfrak{h}_t$ is also fairly straightforward, assuming h has been fixed in advance. Here $P^t_{\text{harm}}(h)$ is the harmonic extension of the values of h from $\partial(\mathbf{D} \setminus K_t)$ to $\mathbf{D} \setminus K_t$. (This notion is defined precisely in the case that K_t is a local set in Section **??**.) We usually choose an additive constant for $\mathfrak{h}_t$ so that $\mathfrak{h}_t(0) = 0$. Since the $\mathfrak{h}_t$ of interest tend to blow up to $\pm\infty$ as one approaches $\partial\mathbf{D}$, a limiting procedure is required to make sense of the map from $\mathfrak{h}_t$ to ν_t. One approach is to define a continuous approximation $\mathfrak{h}_t^n$ to $\mathfrak{h}_t$ using the first n terms in the power series expansion of the analytic function with real part $\mathfrak{h}_t$. One can then let ν_t be the weak $n \to \infty$ limit of the measures $e^{\alpha\mathfrak{h}_t^n(u)}du$ on $\partial\mathbf{D}$, normalized to be probability measures. Such an approach makes sense provided that the process $\mathfrak{h}_t$ almost surely spends almost all time on functions for which this limit exists.

(Here $Q = 2/\gamma + \gamma/2$ and the addition of $Q \log |(g_t^{-1})'|$ comes from the LQG coordinate change rule described in Section 10.2.5 below.) We note that $\mathfrak{h}_t$ is defined modulo a global additive constant. We fix this additive constant by setting $\mathfrak{h}_t(0) = 0$. Once h is fixed, this step essentially describes a map from $\mathcal{G}_T$ to $\mathcal{H}_T$. We say "essentially" because the harmonic extension $P^t_{\text{harm}}(h)$ is not necessarily well-defined for *all* h and g_t pairs, but it *is* almost surely defined under the above-mentioned assumption that K_t is local; see [SS13].

3. $\mathcal{H}_T \to \mathcal{N}_T$: ν_t is obtained by exponentiating $\alpha\mathfrak{h}_t$ on $\partial\mathbf{D}$, for a given value α (which depends on η and γ). Since the $\mathfrak{h}_t$ we will be interested in are almost surely harmonic functions that blow up to $\pm\infty$ as one approaches $\partial\mathbf{D}$, we will have to use a limiting procedure: $d\nu_t = \lim_{n\to\infty} \mathcal{Z}_{n,t}^{-1} e^{\alpha\mathfrak{h}_t^n(u)}du$ where du is Lebesgue measure on $\partial\mathbf{D}$ and $\mathfrak{h}_t^n$ is (the real part of) the sum of the first n terms in the power series expansion of the analytic function (with real part) $\mathfrak{h}_t$, and $\mathcal{Z}_{n,t} = \int_{\partial\mathbf{D}} e^{\alpha\mathfrak{h}_t^n(u)}du$.

We would like to say that this step provides a map from $\mathcal{H}_T$ to $\mathcal{N}_T$, but in fact the map is only defined on the subset of $\mathcal{H}_T$ for which these limits exist for almost all time.[30]

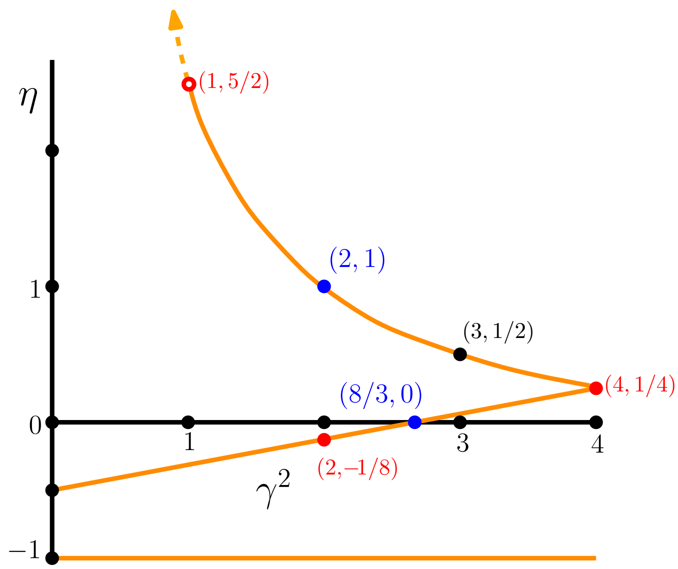

Figure 10.3: The solid orange curves illustrate the (γ^2, η) pairs for which we are able to construct and understand a QLE(γ^2, η) process most explicitly. The curves correspond to $\eta \in \{-1\,,\ \frac{3\gamma^2}{16} - \frac{1}{2}\,,\ \frac{3}{\gamma^2} - \frac{1}{2}\}$, where $\gamma^2 \in (0, 4]$. When (γ^2, η) is on the middle curve, our construction involves radial SLE$_\kappa$ with $\kappa = 16/\gamma^2$. When (γ^2, η) is on the upper curve, it involves radial SLE$_\kappa$ with $\kappa = \gamma^2$. The solid red dots are phase transitions of the SLE$_\kappa$ curves used to construct QLE: $(2, -1/8)$ corresponds to $\kappa = 8$ and $(4, 1/4)$ corresponds to $\kappa = 4$. The point $(1, 5/2)$ corresponds to $\kappa = 1$ and is a phase transition beyond which the QLE construction of this paper becomes trivial — i.e., when $\kappa \leq 1$, the construction (carried out naively) produces a simple radial SLE curve independent of h (and the measures ν_t are point masses for all t). The blue dots are points we are especially interested in. The point $(2, 1)$ is related to DLA on spanning-tree-decorated random planar maps. The point $(8/3, 0)$ is related to the Eden model on undecorated random planar maps, and to the problem of endowing pure LQG with a distance function.

---

[30] Alternatively, one could define $\nu \in \mathcal{N}_T$ as the weak $n \to \infty$ limit of the measures $\mathcal{Z}_{n,t}^{-1} e^{\alpha \mathfrak{h}_t^n(u)} dt du$ on $[0, T] \times \partial \mathbf{D}$. This limit could conceivably exist even in settings for which the ν_t did not exist for almost all t. To avoid making any assumptions at all about limit existence, one could alternatively define a one-to-(possibly)-many map from each $\mathfrak{h}_t$ process in $\mathcal{H}_T$ to the set of *all* $\nu \in \mathcal{N}_T$ obtained as weak $n \to \infty$ limit points of the sequence $\mathcal{Z}_{n,t}^{-1} e^{\alpha \mathfrak{h}_t^n(u)} dt du$ of measures on $[0, T] \times \partial \mathbf{D}$. With that approach, one might require only that ν be *one* of these limit points. Although we do not address this point in this paper, we believe that it might be possible, using these alternatives, to extend the solution existence results of this paper to additional values of η and γ.

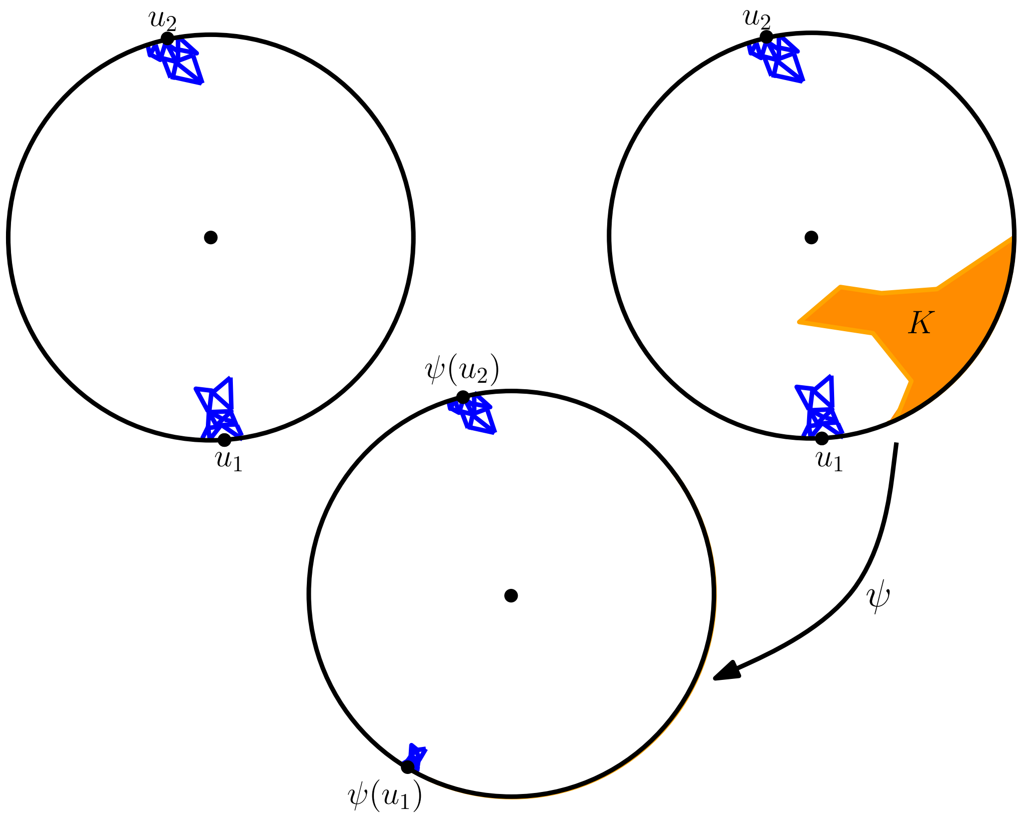

Figure 10.4: **Upper left:** Suppose a discrete random triangulation is conformally mapped to the disk, and the Eden model growing from the boundary inward takes about N time units to absorb the cluster of triangles shown near u_1, and also about N units to absorb the cluster near u_2. (Other not-shown triangles scattered around the boundary are also added during that time.) **Upper right:** Now suppose that we modify the initial setup by designating the hull K to be part of the boundary. Intuitively, if the regions near u_1 and u_2 are small, this modification should not affect the *relative* rate at which growth happens near u_1 and u_2. That is, there should be some N' such that both clusters take about N' steps to be absorbed. **Bottom:** A conformal map $\psi\colon \mathbf{D} \setminus K \to \mathbf{D}$ with $\psi(0) = 0$ scales the region near u_i by about $|\psi'(u_i)|$. The capacity corresponding to the shown blue cluster near $\psi(u_i)$ is approximately $|\psi'(u_i)|^2$ times that of the original blue cluster near u_i. This suggests that if (ν_t) is the driving measure in the bottom figure and $(\widetilde{\nu}_t)$ is the original driving measure in the upper left, and I_i is a small interval about u_i, then $\nu_0(\psi(I_i))$ should be roughly proportional to $|\psi'(u_i)|^2 \widetilde{\nu}_0(I_i)$. In the η-DBM model, one replaces $|\psi'(u_i)|^2$ by $|\psi'(u_i)|^{2+\eta}$ because the rate at which particles reach u_i should also change by a factor roughly proportional to $|\psi'(u_i)|^\eta$.

We remark that if we had $h = 0$, then the triangle in Figure 10.2 would say that $\mathfrak{h}_t = Q \log |(g_t^{-1})'|$ and that ν_t is given (up to multiplicative constant) by $|(g_t^{-1}(u))'|^{\alpha Q} du$. This is precisely the deterministic evolution associated with the DBM that is discussed, for example, in [RZ05] (except that the exponent αQ given here is replaced by a single parameter $-\alpha$). This deterministic evolution has some smooth trivial solutions (for example the constant circular growth given by letting ν be the uniform measure on $[0,T] \times \partial\mathbf{D}$, and taking $g_t(z) = e^t z$ and $\mathfrak{h}_t(z) = 0$). For these solutions, we would not need to use limits to construct ν_t from $\mathfrak{h}_t$, since the measures $e^{\alpha \mathfrak{h}_t(u)}du$ would be well defined. However, if one starts with a generic harmonic function for h that extends smoothly to $\partial\mathbf{D}$ (instead of simply $h = 0$) then the evolution can develop singularities in finite time, and once one encounters the singularities it is unclear how to continue the evolution; this issue and various regularization/approximation schemes to prevent singularity-formation are discussed in [CM01, RZ05]. Even in the $h = 0$ case, Figure 10.2 suggests an interesting alternative to the regularization approaches of [CM01, RZ05]. It suggests an *exact* (non-approximate) notion of what it means to be a solution to the dynamics that makes sense even when singularities are present; the approximation is only involved in making sense of the map from $\mathcal{H}_T$ to $\mathcal{N}_T$. Since the real aim in the $h = 0$ case is to define a natural probability measure on the space of fractal solutions to the dynamics (which should describe the scaling limit of DBM, at least in suitably isotropic formulations), one might hope that these solutions would have some nice properties (perhaps a sort of almost sure fractal self similarity, or long range approximate independence of $\mathfrak{h}_t$ boundary values) that would allow the map from $\mathcal{H}_T$ to $\mathcal{N}_T$ to be almost surely well defined.

In this paper, we will take h to be the GFF (plus a deterministic multiple of $\log|\cdot|$) and we will construct solutions to the dynamics of Figure 10.2 for α and Q values that correspond (in a way we explain later) to the (γ^2, η) values that lie on the upper two curves in Figure 10.3. We will also argue that $\eta = -1$ corresponds to $\alpha = 0$, which yields a trivial solution corresponding to the bottom curve in Figure 10.3. We remark that although this solution is "trivial" in the continuum, the analogous statement about discrete graphs (namely that if a random planar map model, conformally mapped to the disk in some appropriate way, scales to LQG on the disk, then the (-1)-DBM on the random planar map has the dilating circle process as a scaling limit) is still an open problem.

We will produce non-trivial continuum constructions for (the solid portions of) the upper two curves in Figure 10.3 by taking subsequential limits of certain discrete-time approximate processes defined using a radial version of the quantum gravity zipper defined in [She10]. These approximate processes can themselves be interpreted as non-lattice-based variants of η-DBM on a γ-LQG surface that are designed to be isotropic and to have some extra conformal invariance symmetries (here one grows small portions of SLE curves instead of adding small particles of fixed Euclidean shape). The similarity between our approximations and DBM seems to support the idea that (at least for some (γ^2, η) pairs) the processes we construct are the "correct" continuum analogs of η-DBM on a γ-LQG surface. The portion of the upper curve corresponding to $\gamma^2 \leq 1$ is degenerate in that the approximation procedure used to construct the process ν_t, as described in Section 10.14, would yield a point mass for almost all t (although we will discuss this case in detail in this paper).

To each of these processes, we associate a discrete-time approximation of the triple $(\nu_t, g_t, \mathfrak{h}_t)$, in which the time parameter takes values $0, \delta, 2\delta, \cdots$ for a constant δ. The most important property that these discrete-time processes have (which distinguishes them from, e.g., the Hastings-Levitov approximations described in [HL98]) is that the stationary law of the ν_t and the $\mathfrak{h}_t$ turn out to be *exactly* the same for each discrete-time approximation (even as the time step size varies). This rather surprising property is what allows us to understand the stationary law of the $\delta \to 0$ limit (something that has never been possible, in the Euclidean $\gamma = 0$ case, for DBM approximation schemes like Hastings-Levitov). We find that the limiting stationary law is exactly the same as the common stationary law of the approximations, and this allows us to prove that the limit satisfies the dynamics of Figure 10.2, and to prove explicit results about this limit, which we state formally in Section 10.4.

The procedure we use to generate the continuum process has discrete analogs, which give interesting relationships between percolation and the Eden model, and also between loop-erased random walk and DLA.

Before we state our results more precisely, we present in Section 10.2 an overview of several of the models and mathematical objects that will be treated in this work. We also present, in Figures 10.5 through 10.13, computer simulations of the Eden model and DLA on γ-LQG square tilings such as those represented in Figure 10.1. In each of these figures we have $\delta = 2^{-16}$ (as explained in the caption to Figure 10.1) which results in many squares of a larger Euclidean size (and hence a more pixelated appearance) for the larger γ values. Figures 10.14, 10.15, and 10.17 show instances with larger γ but smaller δ values. Generally, the DLA simulations for larger γ values appear to have characteristics in common with the $\gamma = 0$ case, but there is more variability to the shapes when γ is larger. The large-γ, small-δ DLA simulations such as Figure 10.17 sometimes look a bit like Chinese dragons, with a fairly long and windy backbone punctuated by shorter heavily decorated limbs.

Figures 10.18 and 10.19 show what happens when different instances of the Eden model or DLA are performed on top of the same instance of a LQG square decomposition. These figures address an interesting question: how much of the shape variability comes from the randomness of the underlying graph, and how much from the additional randomness associated with the growth process? We believe but cannot prove that in the Eden model case shown in Figure 10.18, the shape of the cluster is indeed determined, to first order (as δ tends to zero), by the GFF instance used to define the LQG measure. The deterministic (given h) shape should be the metric ball in a canonical continuum metric space determined by the GFF.

On the continuum level, the authors are in the process of carrying out a program for using QLE$(8/3, 0)$ to endow a $\gamma = \sqrt{8/3}$ Liouville quantum gravity surface with metric space structure, and to show that the resulting metric space is equivalent in law to a particular random metric space called the Brownian map. But this is not something we will fully explain in these notes.

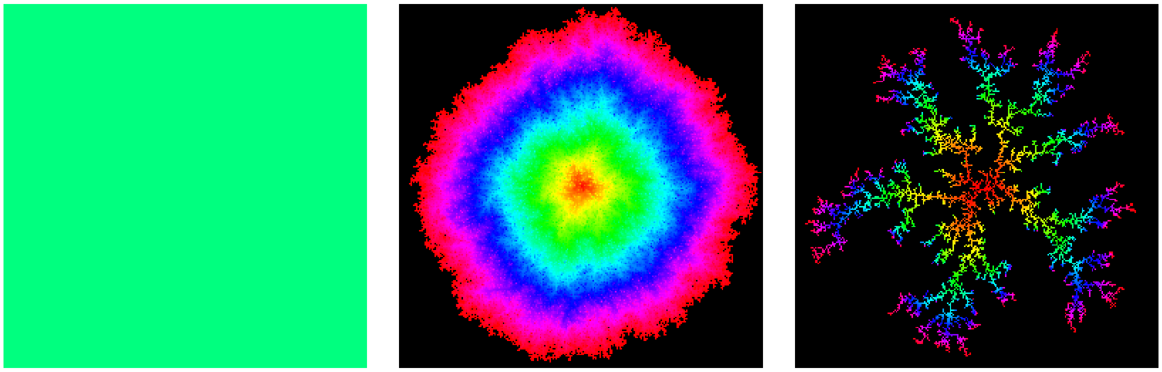

(a) Squares (b) Eden model (c) DLA

Figure 10.5: $\gamma = 0$

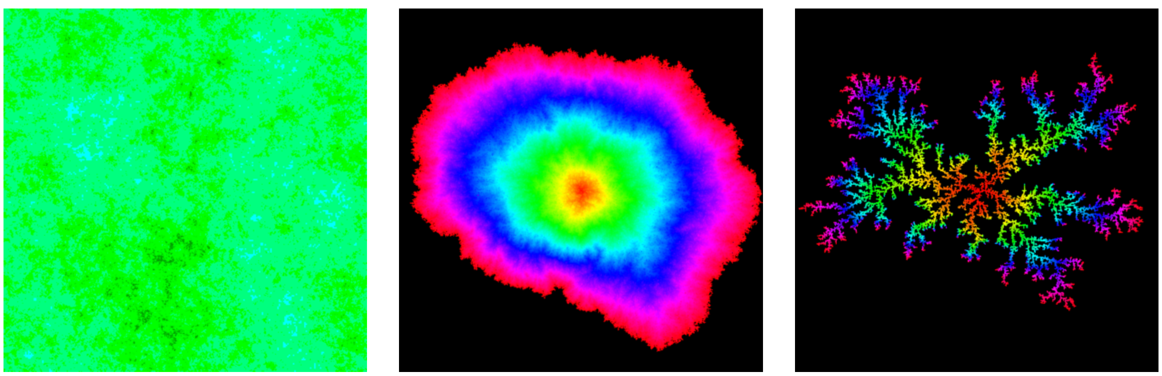

(a) Squares (b) Eden model (c) DLA

Figure 10.6: $\gamma = 1/4$

10.2 Background on several relevant models

10.2.1 First passage percolation and Eden model

The Eden growth model was introduced by Eden in 1961 [Ede61]. One constructs a randomly growing sequence of clusters C_n within a fixed graph $G = (V, E)$ as follows: C_0 consists of a single deterministic initial vertex v_0, and for each $n \in \mathbf{N}$, the cluster C_n is obtained by adding one additional vertex to C_{n-1}. To obtain this vertex, one samples uniformly from the set of edges that have exactly one endpoint in C_{n-1}, and adds the endpoint of this edge that does not lie in C_{n-1}.

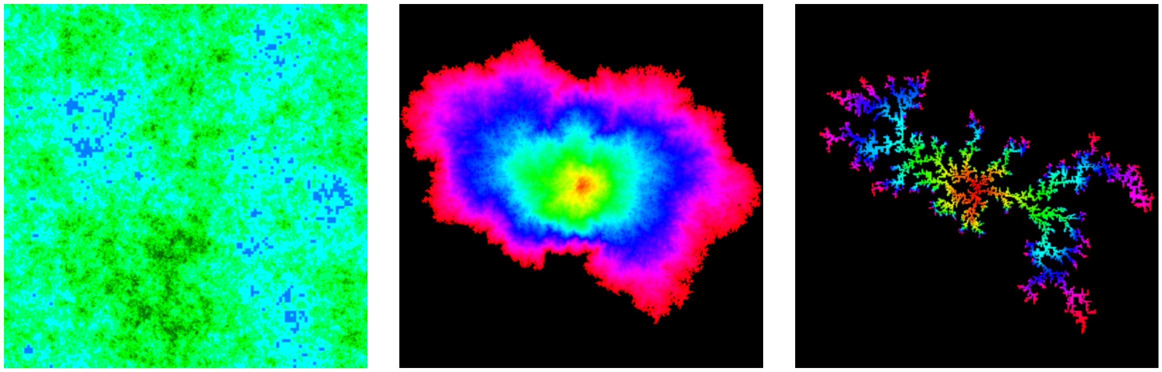

(a) Squares (b) Eden model (c) DLA

Figure 10.7: $\gamma = 1/2$

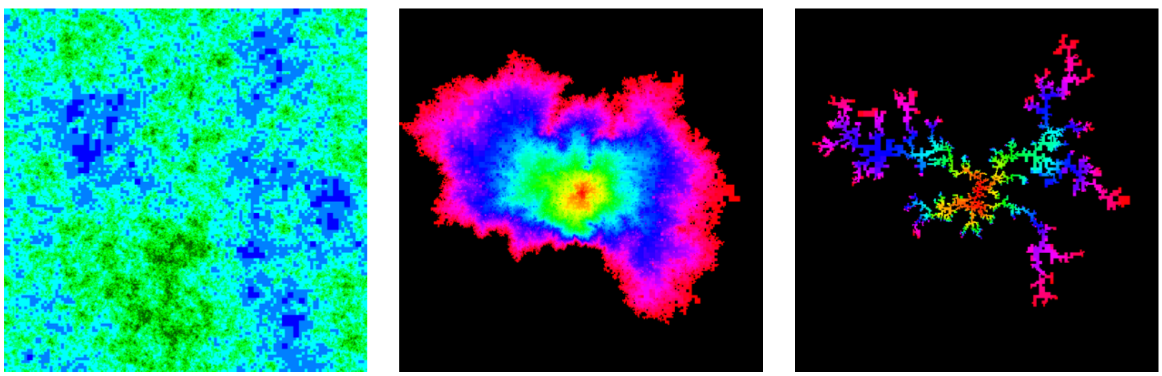

(a) Squares (b) Eden model (c) DLA

Figure 10.8: $\gamma = 3/4$

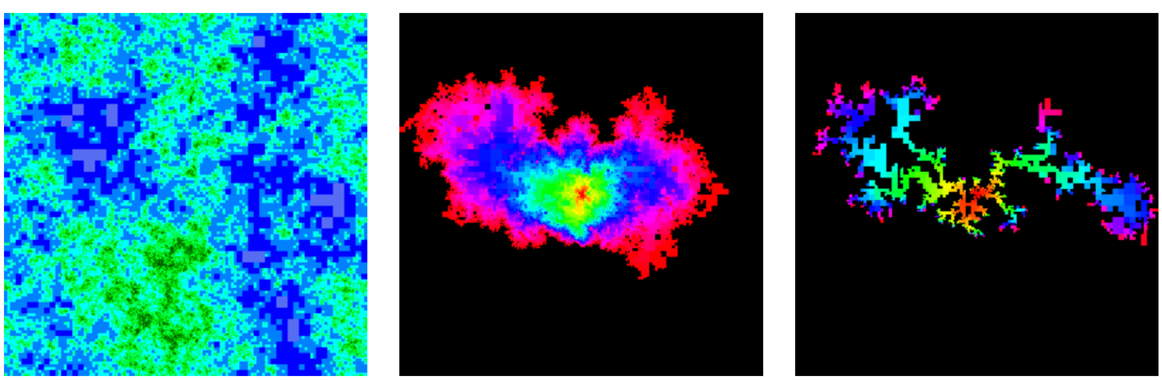

(a) Squares (b) Eden model (c) DLA

Figure 10.9: $\gamma = 1$

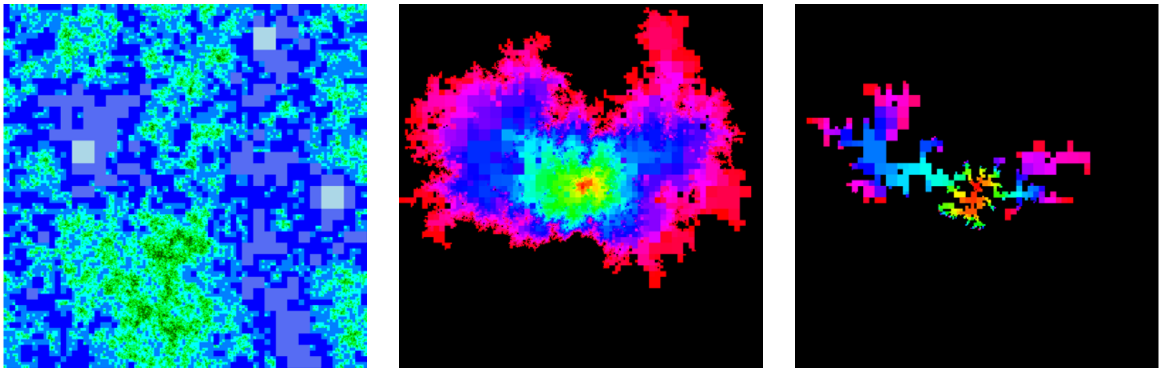

(a) Squares (b) Eden model (c) DLA

Figure 10.10: $\gamma = 5/4$

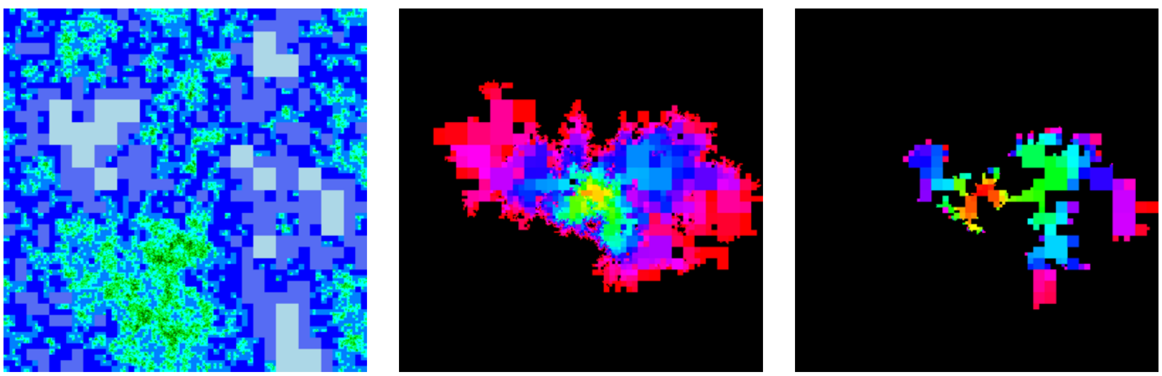

(a) Squares (b) Eden model (c) DLA

Figure 10.11: $\gamma = 3/2$

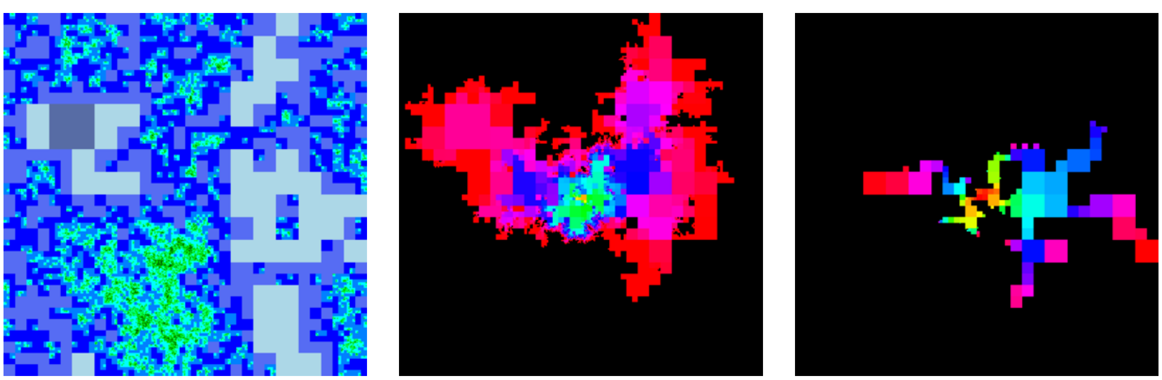

(a) Squares (b) Eden model (c) DLA

Figure 10.12: $\gamma = 7/4$

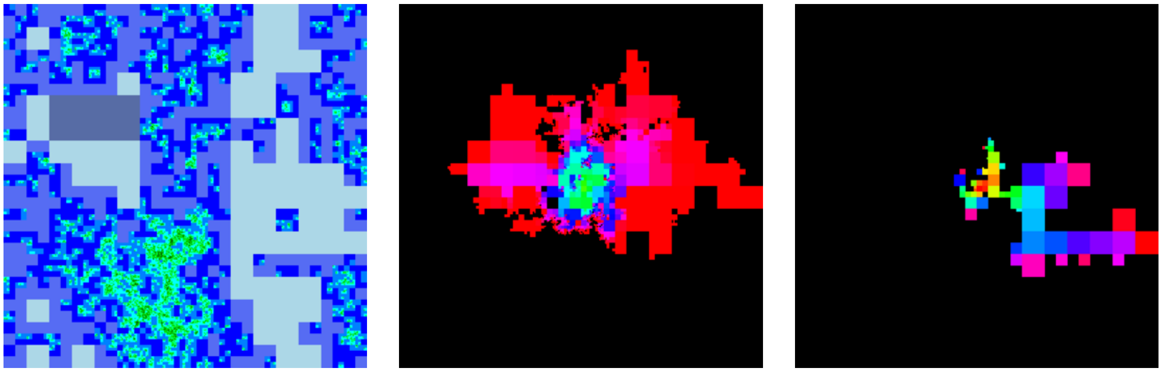

(a) Squares (b) Eden model (c) DLA

Figure 10.13: $\gamma = 2$

First passage percolation (FPP) in turn was introduced by Hammersley and Welsh in 1965 [HW65]. We can construct a random metric on the vertices of the graph G obtained by weighting all edges of G with i.i.d. positive weights; the distance between any two vertices is defined to be the infimum, over all paths between them, of the weight sum along that path. We can then let C_t be the set of all vertices whose distance from an initial vertex v_0 is at most t. If we think of the weight of an edge as representing the amount of time it takes a fluid to "percolate across" the edge, and we imagine that a fluid source is hooked up to a vertex v_0 at time 0, then C_t represents the set of vertices that "get wet" by time t. It is instructive to think of C_t as a growing sequence of balls in a random metric space obtained from the ordinary graph metric on G via independent local perturbations.

For a discrete time parameterization of FPP, we can instead let C_n be the set containing v_0 and the n vertices that are closest to v_0 in this metric space. When the law of the weight for each edge is that of an exponential random variable, it is not hard to see (using the "memoryless" property of exponential random variables) that the sequence of clusters C_n obtained this way agrees in law with the sequence obtained in the Eden growth model.

An overall shape theorem was given by Cox and Durrett in [CD81] in 1981, which dealt with general first passage percolation on the square lattice and showed that under mild conditions on the weight distribution (which are satisfied in the case of exponential weights described above) the set $t^{-1}C_t$ converges to a deterministic shape (though not necessarily exactly a disk) in the limit. Vahidi-Asl and Wierman proved an analogous result for first passage percolation on the Voronoi tessellation (and the related "Delaunay triangulation") later in the early 1990's [VAW90, VAW92] and showed that in this case the limiting shape is actually a ball.

With a very quick glance at Figure 10.5, one might guess that the limiting shape of the Eden model (whose existence is guaranteed by the Cox and Durrett theorem

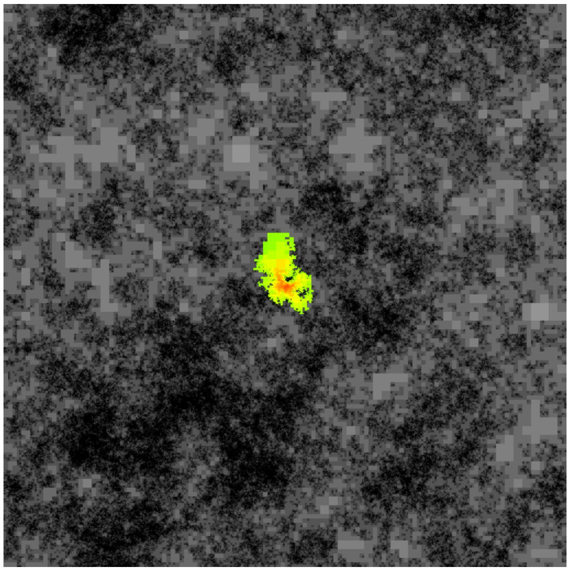

(a) 25%

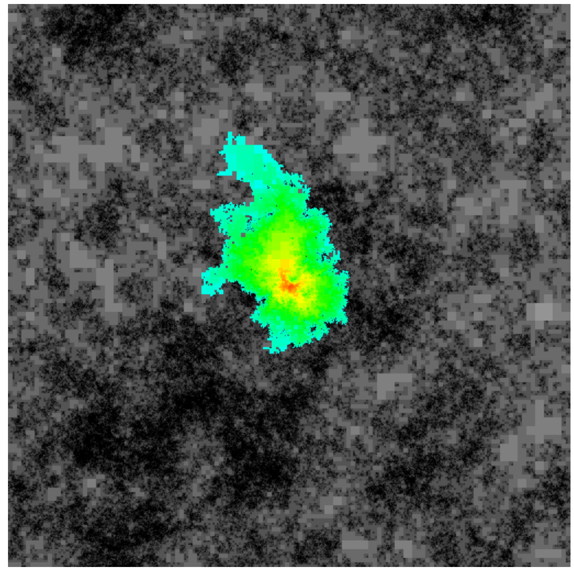

(b) 50%

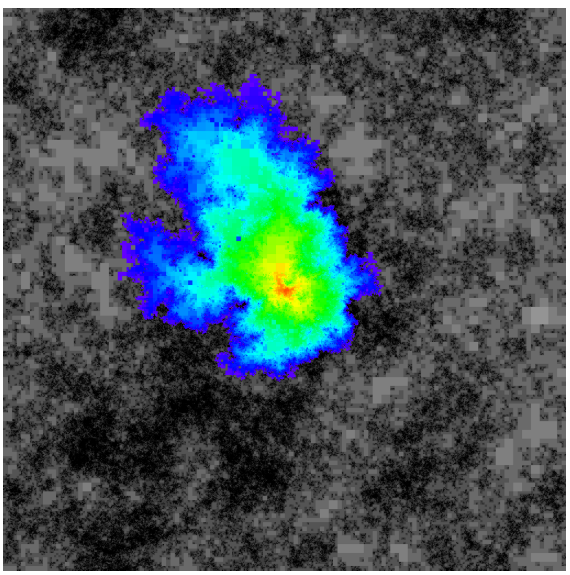

(c) 75%

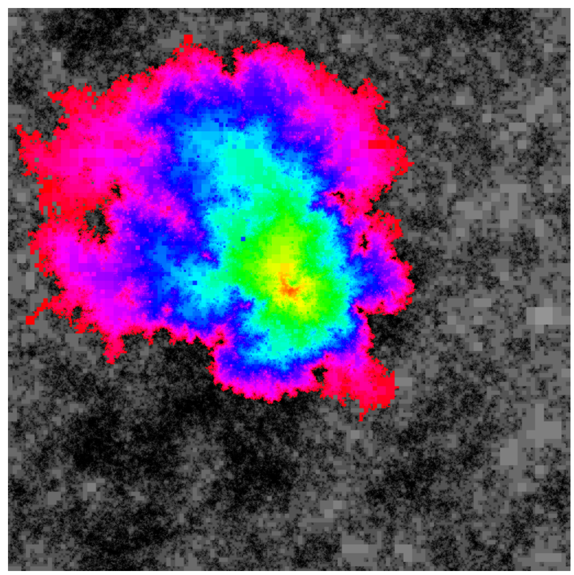

(d) 100%

(e) Time-parameterization.

Figure 10.14: An Eden model instance on a $\sqrt{8/3}$-LQG generated with an $8192 \times 8192 = 2^{13} \times 2^{13}$ DGFF, where $\delta = 2^{-24}$. Shown in greyscale is the original square decomposition (squares of larger Euclidean size are colored lighter). Using the scale shown above, the colors indicate the radius of the ball as it grows relative to the radius at which it first hits boundary of the square. This simulation is a discrete analog of QLE$(8/3, 0)$. See also Figure 10.15 and Figure 10.16.

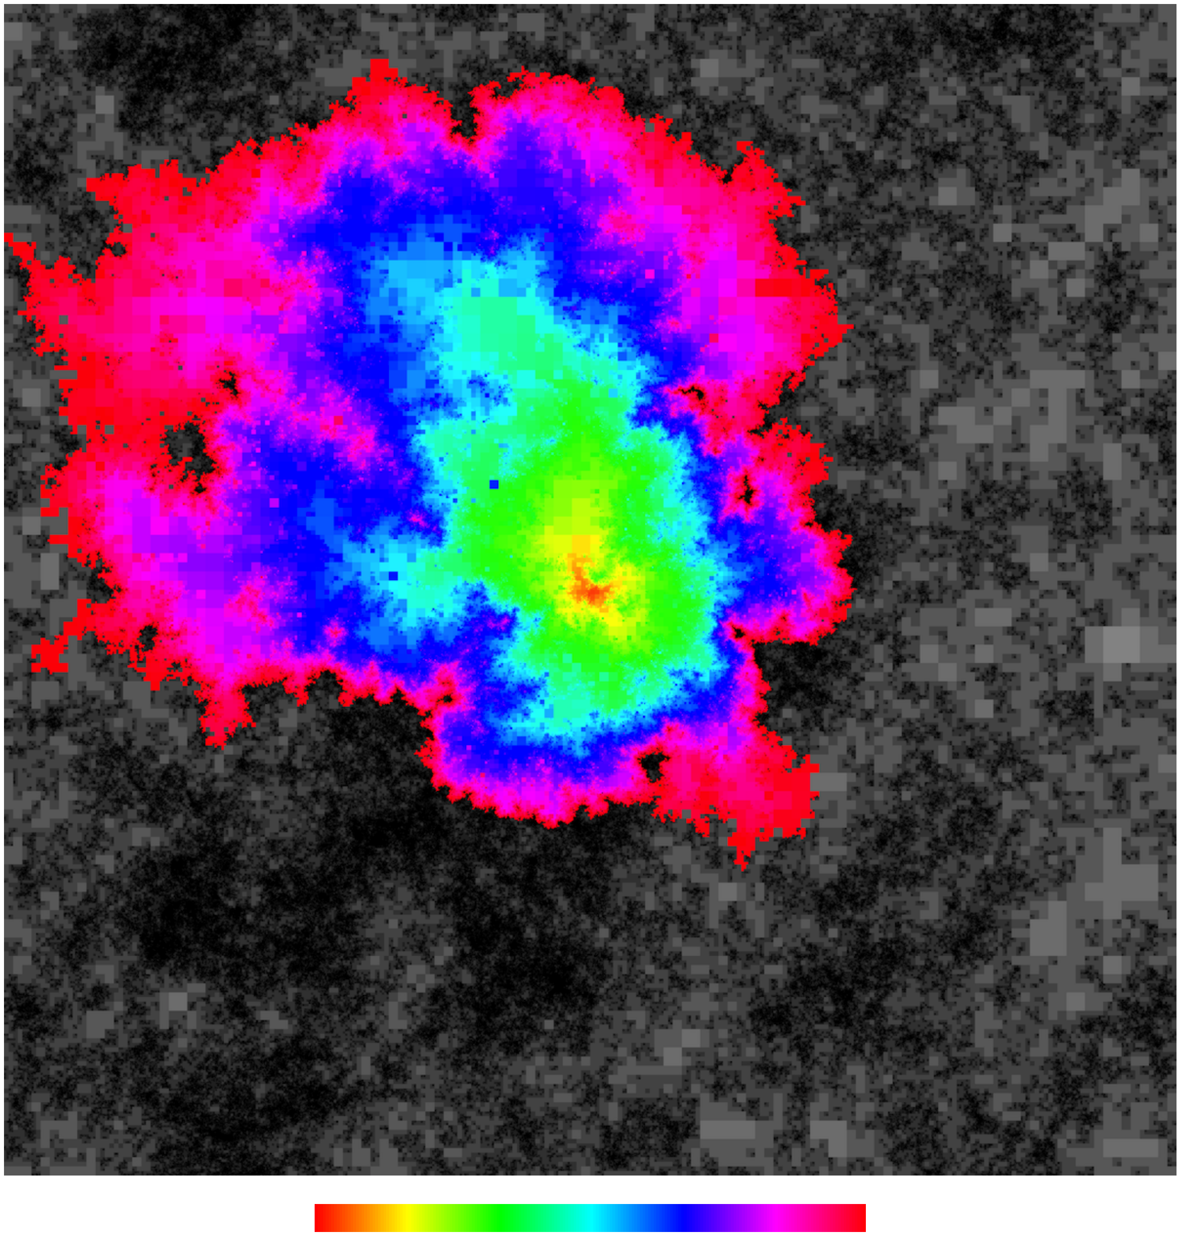

(b) Time-parameterization.

Figure 10.15: Enlargement of final box in Figure 10.14.

mentioned above) is circular; but early and subsequent computer experiments suggest that though the deterministic limit shape is "roundish" it is not exactly circular (e.g., [FSS85, BH91]).

The fluctuations of $t^{-1}C_t$ away from the boundary of the deterministic limit are of smaller order; with an appropriate rescaling, they are believed to have a scaling limit closely related to the KPZ equation introduced by Kardar, Parisi, and Zhang in 1986

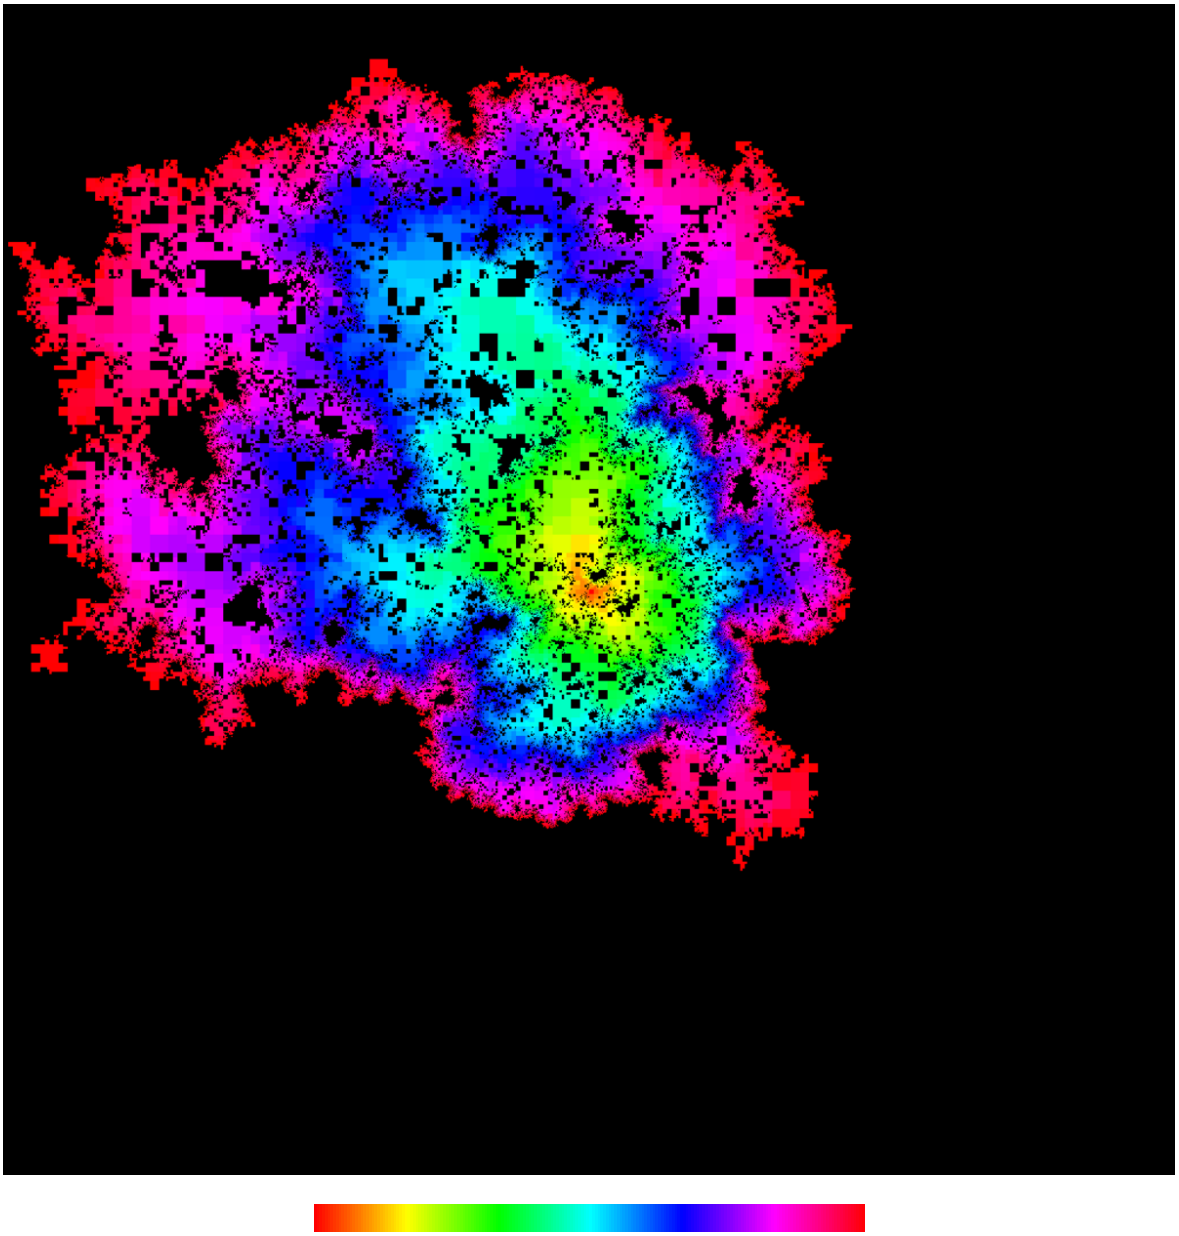

(b) Time-parameterization.

Figure 10.16: Eden model as in Figure 10.15 except that one only adds squares *on the outside* (i.e., reachable by paths from infinity that don't pass through the cluster). The cluster appears to tend to hit regions with big squares but circumvent regions with tiny squares. The number of colored squares is $213061 \approx 2^{17.7}$, and each has μ mass less than a $\delta = 2^{-24}$ fraction of the total, with one caveat: our simulation did not further subdivide the tiny $2^{-13} \times 2^{-13}$ squares, so these can have mass greater than a 2^{-24} fraction of the total. There are $3008224 \approx 2^{21.5}$ squares (colored and non-colored) in this figure, and most of the μ mass lies in the tiny ones.

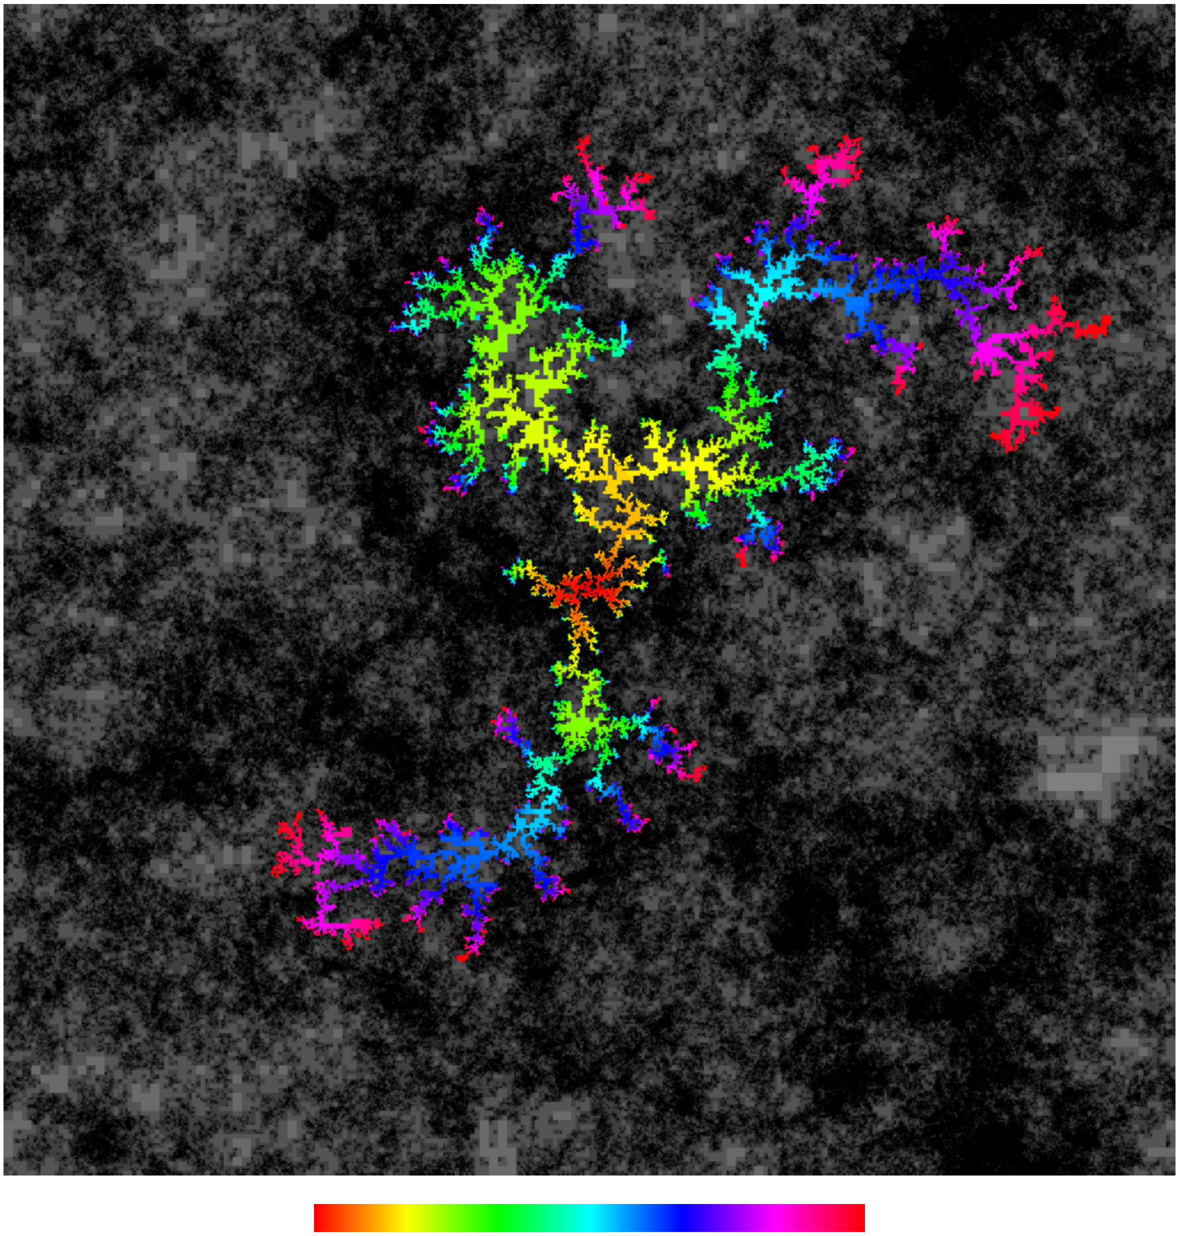

(b) Time-parameterization.

Figure 10.17: DLA on a $\sqrt{2}$-LQG generated with a $8192 \times 8192 = 2^{13} \times 2^{13}$ DGFF, with $\delta = 2^{-26}$. Time is parameterized by the ratio of the number of particles in the cluster over the number required for it to reach the concentric circle inside of the square and is indicated using the color scale shown above. This simulation is a discrete analog of QLE$(2, 1)$.

[KPZ86]. Indeed, understanding growth models of this form was the original motivation for the KPZ equation [KPZ86], see Section 10.2.8.

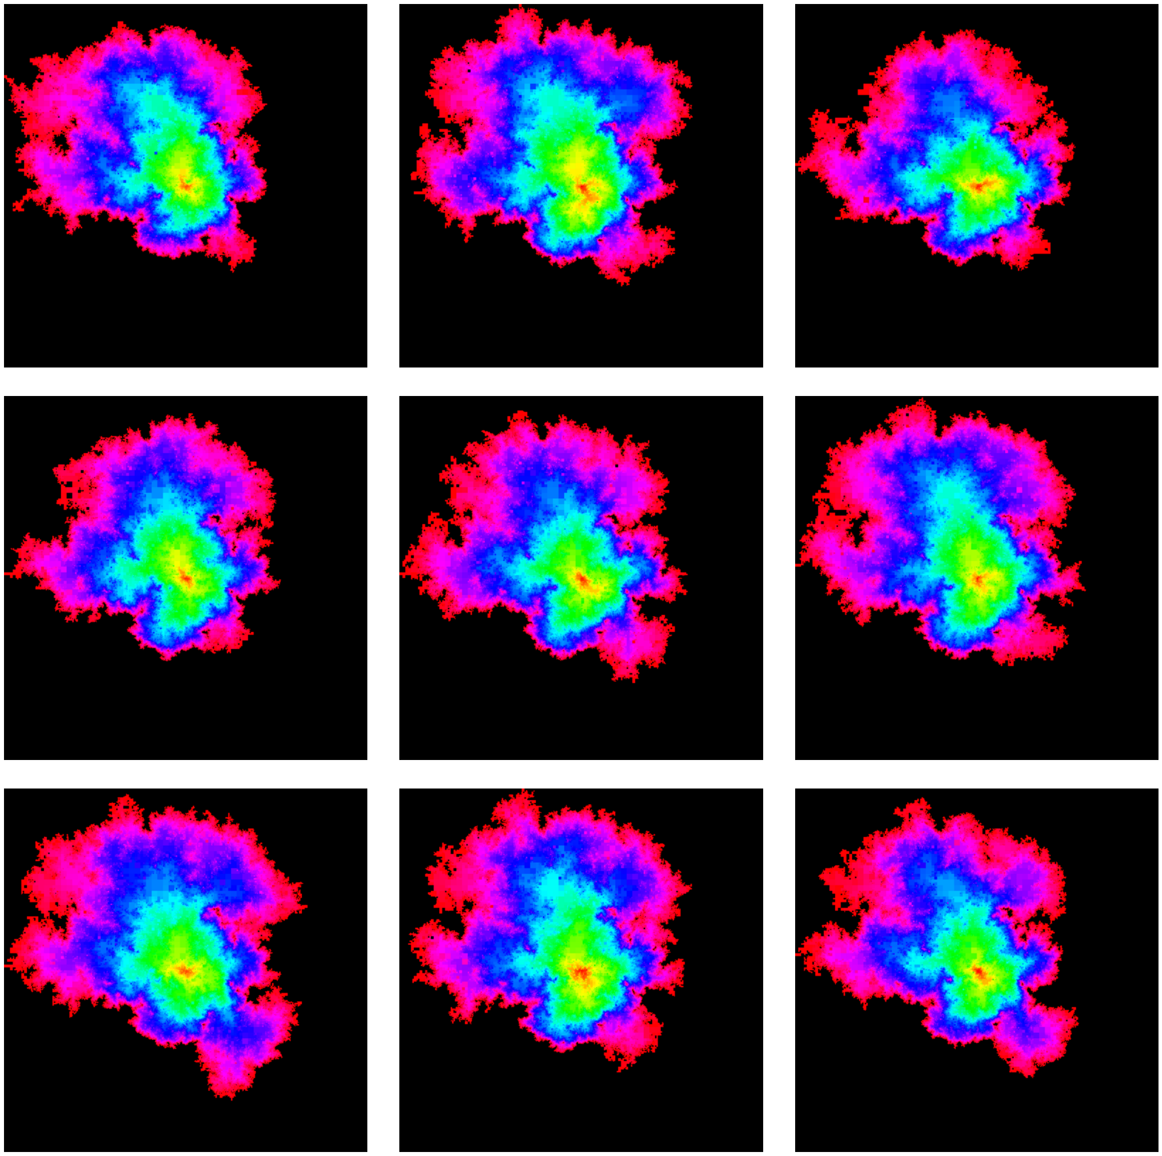

Figure 10.18: Different instances of the Eden model drawn on the square tiling shown in Figure 10.15. We expect that given an instance h of the GFF (which determines the square decomposition for all δ), it is almost surely the case that the shapes converge in probability to a limiting shape (depending only on h) as $\delta \to 0$. The KPZ dynamics are conjecturally related to the fluctuations from the limit shape when $\gamma = 0$ and δ tends to zero. We do not have an analog of this conjecture for general γ.

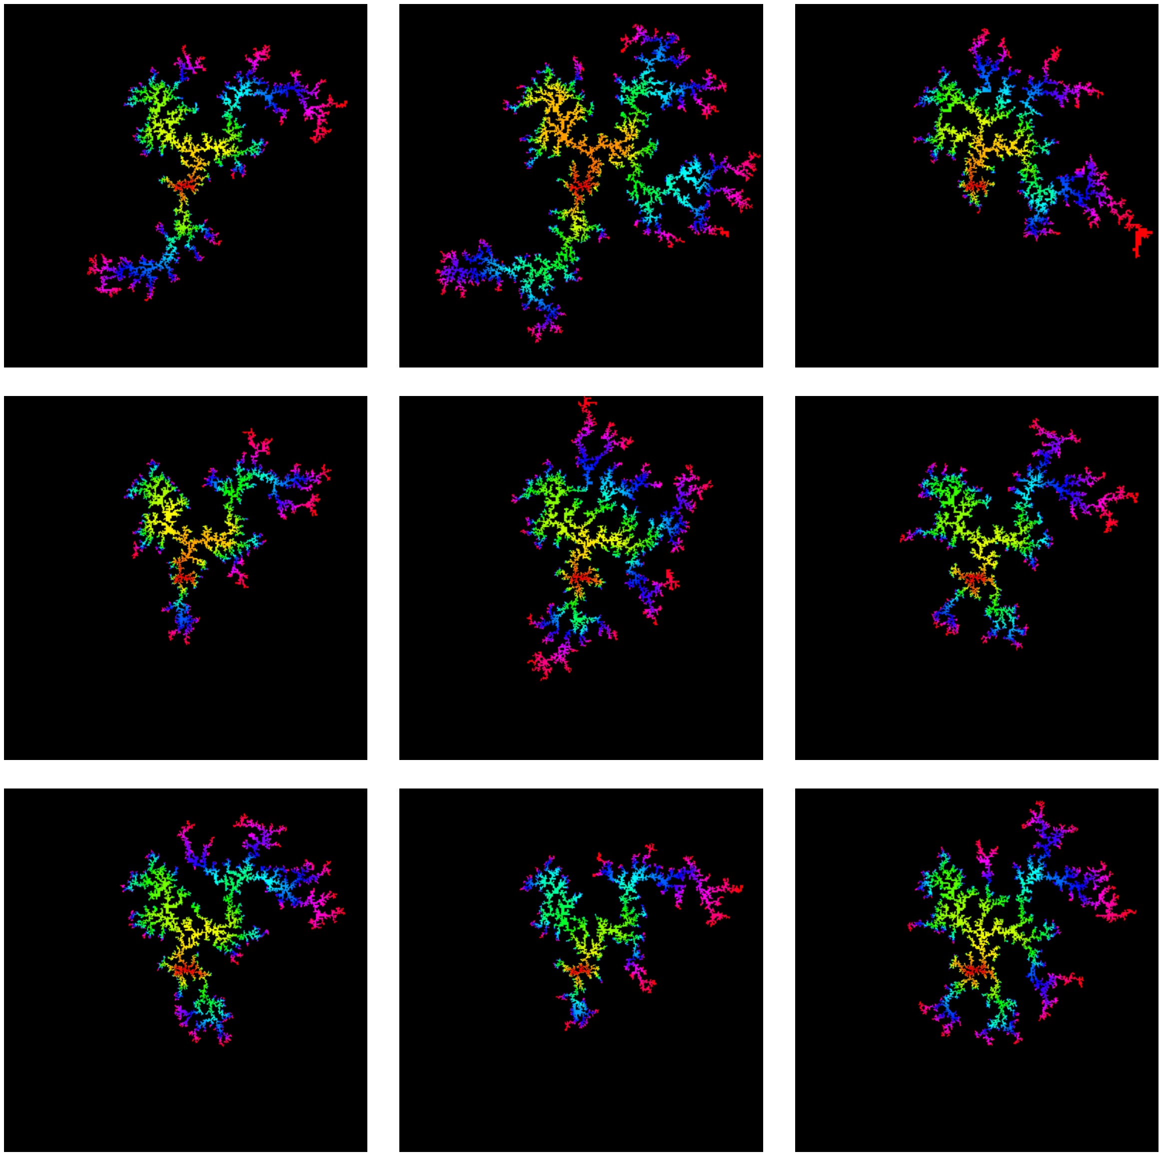

Figure 10.19: Different instances of DLA on a common $\gamma = \sqrt{2}$ LQG tiling. Same 8192×8192 DGFF as in Figure 10.17 with the same value of δ. There are some macroscopic differences between the instances, but we do not know whether these differences will remain macroscopic in the limit as $\delta \to 0$. Similarly, in our continuum formulation, we do not know whether QLE$(2,1)$ is determined by the quantum surfaces on which it is drawn.

10.2.2 Diffusion limited aggregation (DLA)

10.2.3 Dielectric breakdown model and the Hastings-Levitov model

As mentioned above, when FPP weights are exponential, the growth process selects new edges from counting measure on cluster-adjacent edges, i.e., according to the Eden model. DLA is the same but with counting measure replaced by harmonic measure viewed from a special point (or from infinity).

Niemeyer, Pietronero, and Wiesmann introduced the dielectric breakdown model (DBM) in 1984 [NPW84]. Like SLE and LQG, it is a family of models indexed by a single real parameter, which in [NPW84] is called η. As noted in [NPW84], η-DBM can be understood as a hybrid between DLA and the Eden model. If μ is counting measure on the harmonically exposed edges, and ν is harmonic measure, then the DBM model involves choosing a new edge from the measure μ weighted by $(\partial\nu/\partial\mu)^\eta$ (multiplied by a constant to produce a probability measure). Equivalently, we can consider ν weighted by $(\partial\mu/\partial\nu)^{1-\eta}$, also multiplied by a normalizing constant to produce a probability measure. Observe that 0-DBM is then the Eden growth model, while 1-DBM is external DLA.

The DBM models are believed to be related to the so-called α-Hastings-Levitov model when $\alpha = \eta + 1$ [HL98]. (The α used in Hastings-Levitov is not the same as the α used in this paper describe QLE dynamics.) The Hastings-Levitov model is constructed in the continuum using Loewner chains (rather than on a lattice). It was introduced by Hastings and Levitov in 1998 as a plausible and simpler alternative to DLA and DBM, with the expectation that it would agree with these other models in the scaling limit but that it might be simpler to analyze [HL98]. In the Hastings-Levitov model one always samples the location of a new particle from harmonic measure, but the size of the new particle varies as the α power of the derivative of the normalizing conformal map at the location where the point is added. This model itself is now the subject of a sizable literature that we will not attempt to properly survey here. See for example works of Carleson and Makarov [CM01] (obtaining growth bounds analogous to Kesten's bound for DLA), Rohde and Zinsmeister [RZ05] (analyzing scaling limit dimension and other properties for various $\alpha \in [0,2]$, discussing the possibility of an $\alpha = 1$ phase transition from smooth to turbulent growth), Norris and Turner [NT12] (proof of convergence in the $\alpha = 0$ case to a growing disk and a connection to the Brownian web), and the reference text [GV06]. In our terminology, the scaling limit of the α-Hastings-Levitov model should correspond to QLE$(0, \alpha - 1)$, and the $\alpha \in [0,2]$ family studied in [RZ05] should correspond to the points in Figure 10.3 along the vertical axis with $\eta \in [-1,1]$.

10.2.4 Gaussian free field

The Gaussian free field (GFF) is a Gaussian random distribution on a planar domain D, which can be interpreted as a standard Gaussian in the Hilbert space described by the so-called Dirichlet inner product. It has free and fixed boundary analogs, as well as discrete variants defined on a grid; see the GFF survey [She07]. We defer a more detailed discussion of the GFF until Section **??**.

10.2.5 Liouville quantum gravity

Liouville quantum gravity, introduced in the physics literature by Polyakov in 1981 in the context of string theory, is a canonical model of a random two-dimensional Riemannian manifold [Pol81b, Pol81c]. One version of this construction involves replacing the usual Lebesgue measure dz on a smooth domain D with a random measure $\mu_h = e^{\gamma h(z)}dz$, where $\gamma \in [0, 2]$ is a fixed constant and h is an instance of (for now) the free boundary GFF on D (with an additive constant somehow fixed). Since h is not defined as a function on D, one has to use a regularization procedure to be precise. Namely, one defines $h_\epsilon(z)$ to be the mean value of h on the circle $\partial B(z, \epsilon)$, and takes the measure μ to be the weak limit of the measures

$$\epsilon^{\gamma^2/2} e^{\gamma h_\epsilon(z)} dz$$

as ϵ tends to zero [DS11a]. On a linear segment of ∂D, a boundary measure ν_h on ∂D can be similarly defined as

$$\lim_{\epsilon \to 0} \epsilon^{\gamma^2/4} e^{(\gamma/2)h_\epsilon(u)} du,$$

where in this case h_ϵ is the mean of h on the semicircle $D \cap \partial B(u, \epsilon)$ [DS11a]. (A slightly different procedure is needed to construct the measure in the critical case $\gamma = 2$ [DRSV12a, DRSV12b].)

We could also parameterize the same surface with a different domain $\widetilde{D}$. Suppose $\psi \colon \widetilde{D} \to D$ is a conformal map. Write $\widetilde{h}$ for the distribution on $\widetilde{D}$ given by $h \circ \psi + Q \log |\psi'|$ where $Q := \frac{2}{\gamma} + \frac{\gamma}{2}$. Then it is shown in [DS11a] that μ_h is almost surely the image under ψ of the measure $\mu_{\widetilde{h}}$. That is, $\mu_{\widetilde{h}}(A) = \mu_h(\psi(A))$ for $A \subseteq \widetilde{D}$.[31] A similar argument to the one in [DS11a] mentioned above shows that the boundary length ν_h is almost surely the image under ψ of the measure $\nu_{\widetilde{h}}$. (This also allows us to make sense of ν_h on domains with non-linear boundary.)

We define a **quantum surface** to be an equivalence class of pairs (D, h) under the equivalence transformations

$$(D, h) \to (\psi^{-1}(D), h \circ \psi + Q \log |\psi'|) = (\widetilde{D}, \widetilde{h}). \tag{10.1}$$

The measures μ_h and ν_h are almost surely highly singular objects with fractal structure, and thus we cannot understand LQG random surfaces as smooth manifolds. Nonetheless, quantum surfaces come equipped with well-defined notions of area, boundary length, and conformal structure. (Recall that the conformal structure of a Riemannian manifold refers to the Riemannian metric modulo rescaling by a real-valued function. Note that, under such a rescaling, the angle at which two curves intersect as measured by the

[31] The reader can also verify this fact directly; the first term in Q is related to the ordinary change of measure formula, since the term $\frac{2}{\gamma} \log |\psi'|$ in $\widetilde{h}$ corresponds to a factor of $|\psi'|^2$ in the $\mu_{\widetilde{h}}$ definition. The term $\frac{\gamma}{2} \log |\psi'|$ compensates for the rescaling of the ϵ that appears in the definition of $\mu_{\widetilde{h}}$.

metric do not change. Since the approximations of a quantum surface involve rescaling the Euclidean metric by a real-valued function, they have the same conformal structure as the Euclidean metric.)

In addition to [DS11a, DRSV12a, DRSV12b], we also direct the reader to the surveys respectively by Garban and Rhodes and Vargas [Gar12, RV14] for more on Liouville quantum gravity.

10.2.6 Random planar maps

The number of planar maps with a fixed number of vertices and edges is finite, and there is an extensive literature on the enumeration of planar maps, beginning with the works of Mullin and Tutte in the 1960's [Tut62, Mul67b, Tut68b]. On the physics side, various types of random planar maps were studied in great detail throughout the 1980's and 1990's, in part because of their interpretation as "discretized random surfaces." (See [DS11a] for a more extensive bibliography on planar maps and Liouville quantum gravity.) The *metric space* theory of random quadrangulations begins with an influential bijection discovered by Schaeffer [Sch97], and earlier by Cori and Vauquelin [CV81]. Closely related bijections of Bouttier, Di Franceso, and Guitter [BDFG04] deal with planar maps with face size restrictions, including triangulations. Subsequent works by Angel [Ang03] and by Angel and Schramm [AS03] have explained the *uniform infinite planar triangulation* (UITP) as a subsequential limit of planar triangulations.

Although microscopic combinatorial details differ, there is one really key idea that underlies much of the combinatorial work in this subject: namely, that instead of considering a planar map alone, one can consider a planar map together with a spanning tree. Given the spanning tree, one often has a notion of a dual spanning tree, and a path that somehow goes between the spanning tree and the dual spanning tree. It is natural to fix a root vertex for the dual tree and an adjacent root vertex for the tree. Then as one traverses the path, one can keep track of a pair of parameters in $\mathbf{Z}_+^2$: one's distance from a root vertex within the tree, and one's distance from the dual root within the dual tree. Mullin in 1967 used essentially this construction to give a way of enumerating the pairs (M,T) where M is a rooted planar map on the sphere with n edges and T is distinguished spanning tree [Mul67b]. These pairs correspond precisely to walks of length $2n$ in $\mathbf{Z}_+^2$ that start and end at the origin. (The bijection between tree-decorated maps and walks in $\mathbf{Z}_+^2$ was more explicitly explained by Bernardi in [Ber07]; see also the presentation and discussion in [She11c], as well as the brief overview in Section 10.8.1.) As n tends to infinity and one rescales appropriately, one gets a Brownian excursion on $\mathbf{R}^2$ starting and ending at 0.

The Mullin bijection gives a way of choosing a uniformly random (M,T) pair, and if we ignore T, then it gives us a way to choose a random M where the probability of a given M is proportional to the number of spanning trees that M admits. If instead we had a way to choose randomly from a *subset* $\mathcal{S}$ of the set of pairs (M,T), with the property

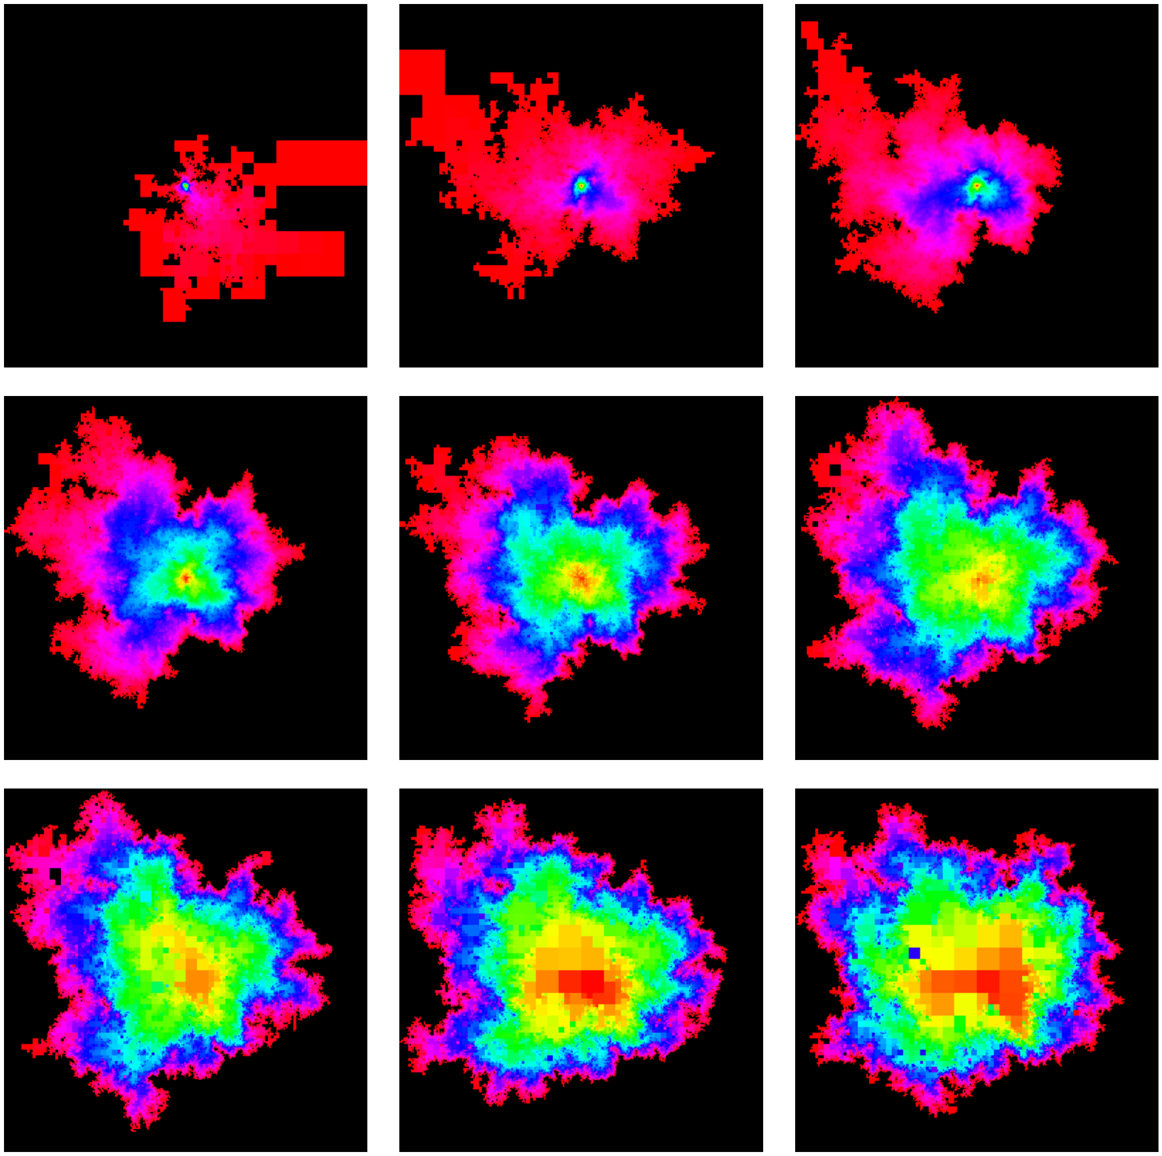

Figure 10.20: $\gamma = \sqrt{8/3}$ Eden model on graph obtained when h is the GFF plus $j \log|\cdot|$, where $j \in \{-4, -3, -2, \ldots, 2, 3, 4\}$ (read left to right, top to bottom). Upper left figure has smaller squares in center, bigger squares on outside. Bottom right has bigger boxes in center, smaller boxes outside.

that each M belonged to at most *one* pair $(M, T) \in \mathcal{S}$, then this would give us a way to sample *uniformly* from some collection of maps M. The Cori-Vauquelin-Schaeffer construction [CV81, Sch97] suggests a way to do this: in this construction, M is required to be a quadrangulation, and a "tree and dual tree" pair on M are produced from M in a deterministic way. (The precise construction is simple but a bit more complicated

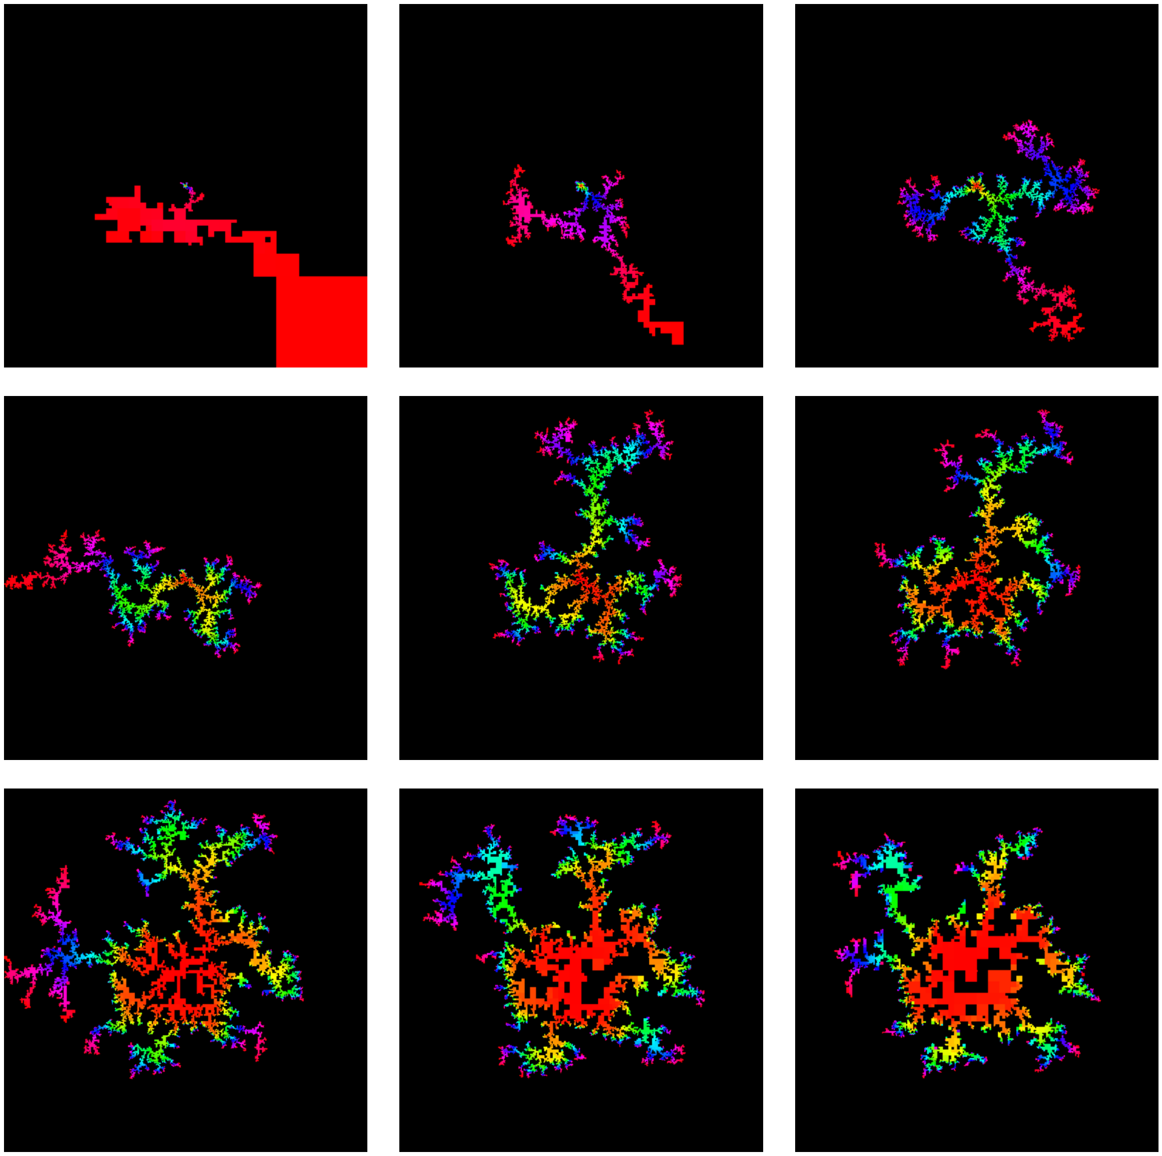

Figure 10.21: $\gamma = \sqrt{2}$ DLA drawn on graph obtained when h is the GFF plus $j \log|\cdot|$, where $j \in \{-4, -3, -2, \ldots, 2, 3, 4\}$ (read left to right, top to bottom). Upper left figure has smaller squares in center, bigger squares on outside. Bottom right has bigger boxes in center, smaller boxes outside.

than the Mullin bijection. One of the trees is a breadth first search tree of M consisting of geodesics, and the other is something like a dual tree defined on the same vertices, but with some edges that cross the quadrilaterals diagonally and some edges that overlap the tree edges.) As one traces the boundary of the dual tree, the distance from the root in the dual tree changes by ± 1 at each step, while the distance in the geodesic tree changes by either 0 or ± 1. Chassaing and Schaeffer showed that this distance function should scale to a two-dimensional continuum random path called the Brownian snake, in which the first coordinate is a Brownian motion (and the second coordinate comes from a Brownian motion indexed by the continuum random tree defined by the first Brownian motion) [CS04].

Another variant due to the second author appears in [She11c], where the trees are taken to be the exploration trees associated with a random planar map together with a random FK cluster on top of it. In fact, the construction in [She11c], described in terms of a "hamburgers and cheeseburgers" inventory management process, is a generalization of the work of Mullin [Mul67b]. We stress that the walks on $\mathbf{Z}_+^2$ that one finds in both [Mul67b] and [She11c] have as scaling limits forms of two-dimensional Brownian motion (in [She11c] the diffusion rate of the Brownian motion varies depending on the FK parameter), unlike the walks on $\mathbf{Z}_+^2$ given in [Sch97, CS04] (which scale to the Brownian snake described above).

10.2.7 The Brownian map

The Brownian map is a random metric space equipped with an area measure. It can be constructed from the Brownian snake, and is believed to be in some sense equivalent to a form of Liouville quantum gravity when $\gamma = \sqrt{8/3}$. The idea of the Brownian map construction has its roots in the combinatorial works of Schaeffer and of Chassaing and Schaeffer [Sch97, CS04], as discussed just above in Section 10.2.6, where it was shown that certain types of random planar maps could be described by a random tree together with a random labeling that determines a dual tree, and that this construction is closely related to the Brownian snake.

The Brownian map was introduced in works by Marckert and Mokkadem and by Le Gall and Paulin [MM06b, LGP08]. For a few years, the term "Brownian map" was often used to refer to any one of the subsequential Gromov-Hausdorff scaling limits of random planar maps. Because it was not known whether the limit was unique, the phrase "a Brownian map" was sometimes used in place of "the Brownian map". Works by Le Gall and by Miermont established the uniqueness of this limit and showed that it is equivalent to a natural metric space constructed directly from the Brownian snake [LG10, Mie13b, LG13b]. Infinite planar quadrangulations on the half plane or the plane and the associated infinite volume Brownian maps are discussed in [CM12, CLG12].

10.2.8 KPZ: Kardar-Parisi-Zhang

As mentioned briefly in Section 10.2.1, Kardar, Parisi, and Zhang introduced a formal stochastic partial differential equation in 1986 in order to describe the fluctuations from the deterministic limit shape that one finds in the Eden model on a grid (as in Figure 10.5) or in related models such as first passage percolation [KPZ86]. As described in [KPZ86], the equation is a type of ill-posed stochastic partial differential equation, but one can interpret the log of the stochastic heat equation with multiplicative noise as in some sense solving this equation (this is called the Hopf-Cole solution). The Eden model fluctuations are believed to scale to a "fixed point" of the dynamics defined this way; see the discussion by Corwin and Quastel in [CQ11], as well as the survey article [Cor12b]. Other recent discussions of this point include e.g. [CMB96, AOF11].

One interesting question for us is what the analog of the KPZ growth equation should be for the random graphs described in this paper. Figure 10.15 shows different instances of the Eden model drawn on the square tiling shown in Figure 10.15. Although they appear to be roughly the same shape, there are clearly random fluctuations and at present we do not have a way to predict the behavior or even the magnitude of these fluctuations (though we would guess that the magnitude decays like some power of δ).

10.2.9 KPZ: Knizhnik-Polyakov-Zamolodchikov

A natural question is whether discrete models for random surfaces (built combinatorially by randomly gluing together small squares or triangles) have Liouville quantum gravity as a scaling limit. Polyakov became convinced in the affirmative in the 1980's after jointly deriving, with Knizhnik and Zamolodchikov, the so-called *KPZ formula* for certain Liouville quantum gravity scaling dimensions and comparing them with known combinatorial results for the discrete models [KPZ88a, Pol08a]. Several precise conjectures along these lines appear in [DS11a, She10] and the KPZ formula was recently formulated and proved mathematically in [DS11a]; see also [DS09].

To describe what the KPZ formula says, suppose that a constant $\gamma \in [0, 2]$, a fractal planar set X, and an instance h of the GFF are all fixed. The set X can be either deterministic or random, as long as it is chosen independently from h. Then for any δ one can generate a square decomposition of the type shown in Figure 10.1 and ask whether the expected number of squares intersecting X scales like a power of δ. One form of the KPZ statement proved in [DS11a] is that if the expected number of squares (using the decomposition for $\gamma = 0$) intersecting X scales like a power of δ when $\gamma = 0$ (the Euclidean case) then, for any fixed $\gamma \in (0, 2)$, the expected number of squares (using the decomposition for the given value of γ) intersecting X also scales like a power of δ, and the Euclidean and quantum exponents satisfy a particular quadratic relationship (depending on γ). Formulations of this statement in terms of Hausdorff dimension (and a quantum-surface analog of Hausdorff dimension) in one and higher dimensions appear respectively in [BS09b, RV11]; see also [DRSV12a, DRSV12b] for the case $\gamma = 2$.

One important thing to recognize for this paper is that the KPZ formula only applies when X and h are chosen independently of one another. This independence assumption is natural in many contexts—for example, one sometimes expects the scaling limit of a random planar map decorated with a path (associated to some statistical physics model) to be an LQG surface decorated with an SLE-curve that is in fact independent of the field h describing the LQG surface [She10, DS11c]. However, we do not expect the Euclidean and quantum dimensions of the QLE traces constructed in this paper to be related by the KPZ formula, because these random sets are not independent of the GFF instance h. (See [Aru13] for an example of a set which is *not* independent of the field and does *not* satisfy the KPZ relation.)

10.2.10 Schramm Loewner evolution

SLE_κ ($\kappa > 0$) is a one-parameter family of conformally invariant random curves, introduced by Oded Schramm in [Sch00c] as a candidate for (and later proved to be) the scaling limit of loop erased random walk [LSW04d] and the interfaces in critical percolation [Smi01c, CN06b]. Schramm's curves have been shown so far also to arise as the scaling limit of the macroscopic interfaces in several other models from statistical physics: [Smi10b, CS12, SS05c, SS09d, Mil10c]. More detailed introductions to SLE can be found in many excellent survey articles of the subject, e.g., [Wer04b, Law05b].

Given a simply connected planar domain D with boundary points a and b and a parameter $\kappa \in [0, \infty)$, the *chordal* **Schramm-Loewner evolution** SLE_κ is a random non-self-crossing path in $\overline{D}$ from a to b. In this work, we will be particularly concerned with the so-called *radial* SLE_κ, which is a random non-self-crossing path from a fixed point on ∂D to a fixed interior point in D. Like chordal SLE, it is completely determined by certain conformal symmetries [Sch00c].

The construction of SLE is very interesting. When $D = \mathbf{D}$ is the unit disk, the radial SLE curve can be parameterized by a function $U\colon [0, \infty) \to \partial \mathbf{D}$. However, instead of constructing the curve directly, one constructs for each t the conformal map $g_t\colon \mathbf{D}_t \to \mathbf{D}$, where $\mathbf{D}_t$ is the complementary component[32] of the curve drawn up to time t which contains 0, with $g_t(0) = 0$ and $g_t'(0) > 0$. For $u \in \partial \mathbf{D}$ and $z \in \mathbf{D}$, let

$$\Psi(u, z) = \frac{u + z}{u - z} \quad \text{and} \quad \Phi(u, z) = z\Psi(u, z). \tag{10.2}$$

For each fixed z, the value $g_t(z)$ is defined as the solution to the ODE

$$\dot{g}_t(z) = \Phi(U_t, g_t(z)), \tag{10.3}$$

where $U_t = e^{i\sqrt{\kappa}B_t}$ and B_t is a standard Brownian motion. More introductory material about SLE appears in Section 4.10.

SLE is relevant to this paper primarily because of its relevance to Liouville quantum gravity and the so-called *quantum gravity zipper* described by the second author in [She10]. Roughly speaking, the constructions there allow one to form one LQG surface

[32] Here "complementary component of" means "component of the complement of".

by "cutting" another LQG surface along an SLE path. In fact, one can do this in such a way that the new (cut) surface has the same law as the original (uncut) surface. This will turn out to be extremely convenient as we construct and study the quantum Loewner evolution.

10.3 Measure-driven Loewner evolution

We consider an analog of Loewner evolution, also called the *Loewner-Kufarev evolution*, in which the point-valued driving function is replaced by a measure-valued driving function:

$$\dot{g}_t(z) = \int_{\partial \mathbf{D}} \Phi(u, g_t(z)) d\nu_t(u), \tag{10.4}$$

(recall (10.2)) where, for each time t, the measure ν_t is a probability measure on $\partial\mathbf{D}$. For each time t, the map g_t is the unique conformal map from $\mathbf{D} \setminus K_t$ to $\mathbf{D}$ with $g_t(0) = 0$ and $g_t'(0) > 0$, for some hull K_t. Time is parameterized so that $g_t'(0) = e^t$ (this is the reason that ν_t is normalized to be a probability measure). That is, the log conformal radius of $\mathbf{D} \setminus K_t$, viewed from the origin, is given by $-t$. Given any measure ν on $[0, T] \times \partial\mathbf{D}$ whose first coordinate is given by Lebesgue measure, we can define ν_t to be the conditional measure obtained on $\partial\mathbf{D}$ by restricting the first coordinate to t.

Unlike the space of point-valued driving functions indexed by $[0, T]$, the space of measure-valued driving functions indexed by $[0, T]$ has a natural topology with respect to which it is compact: namely the topology of weak convergence of measures on $[0, T] \times \partial\mathbf{D}$.

We now recall a standard result, which can be found, for example, in [JVST12]. (A slightly more restrictive statement is found in [Law05b].) Essentially it says that $\mathcal{N}_T$ together with the notion of weak convergence, corresponds to the space of capacity-parameterized growing hull processes in $\mathbf{D}$ (indexed by $t \in [0, T]$), with the notion of Carathéodory convergence for all t. (Recall that a sequence of hulls $K^1, K^2, \ldots$ converges to a hull K in the Carathéodory sense if the conformal normalizing maps from $\mathbf{D} \setminus K^j$ to $\mathbf{D}$ converge uniformly on compact subsets of $\mathbf{D} \setminus K$ to the conformal normalizing map from $\mathbf{D} \setminus K$ to $\mathbf{D}$.)

Theorem 10.1. *Consider the following:*

(i) A measure $\nu \in \mathcal{N}_T$.

(ii) An increasing family (K_t) of hulls in $\mathbf{D}$, indexed by $t \in [0, T]$, such that $\mathbf{D} \setminus K_t$ is simply connected and includes the origin and has conformal radius e^{-t}, viewed from the origin. (In other words, for each t, there is a unique conformal map $g_t \colon \mathbf{D} \setminus K_t \to \mathbf{D}$ with $g_t(0) = 0$ and $g_t'(0) = e^t$.)

There is a one-to-one correspondence between objects of type (i) *and* (ii)*. In this correspondence, the maps g_t are obtained from ν via* (10.4)*, where ν_t is taken to be the conditional law of the second coordinate of ν given that the first coordinate is equal to t. Moreover, a sequence of measures $\nu^1, \nu^2, \ldots$ in $\mathcal{N}_T$ converges weakly to a limit ν if and only if for each t the functions $g_t^1, g_t^2, \ldots$ corresponding to ν_i converge uniformly to the function g_t corresponding to ν on any compact set in the interior of $\mathbf{D} \setminus K_t$.*

For completeness, we will provide a proof of Theorem 10.1 in Section 10.14. The reader may observe that the notion of Carathéodory convergence for all $t \in [0, T]$ is equivalent to the notion of Carathéodory convergence for all t in a fixed countable dense subset of $[0, T]$. This can be used to give a direct proof of compactness of the set of hull families described Theorem 10.1, using the topology of Carathéodory convergence for all t.

10.4 Main results

10.4.1 Subsequential limits and compactness

The main purpose of this paper is to construct a candidate for what should be the scaling limit of η-DLA on a γ-LQG surface (at least in sufficiently isotropic formulations) for the (γ^2, η) pairs which lie on the top two solid curves from Figure 10.3.

Before presenting these results, let us explain one path that we will *not* pursue in this paper. One natural approach would be to take a subsequential limit of η-DLA on δ-approximations of γ-LQG (perhaps using an inherently isotropic setting, such as the one involving Voronoi cells of a Poisson point process associated with the LQG measure) and to simply *define* the limit to be a QLE(γ^2, η). Using Theorem 10.1 and the weak compactness of $\mathcal{N}_T$, it should not be hard to construct a triple $(\nu_t, g_t, \mathfrak{h}_t)$ coupled with a free field instance h, as in the context of Figure 10.2, with the property that

1. The sets K_t corresponding to g_t are local.
2. The maps from ν_t to g_t, and from g_t to $\mathfrak{h}_t$ are as described in Figure 10.2.

The natural next step would then be to show that $\mathfrak{h}_t$ determines ν_t in the manner of Figure 10.2. We consider this to be an interesting problem, and one that might potentially be solvable by understanding (using the discrete approximations) how ν_t restricted to a boundary arc would change when one added a constant to $\mathfrak{h}_t$ on that boundary arc.

However, we stress that even if this problem were solved, it would not immediately give us an explicit description of the stationary law of ν_t. The main contribution of this article is to construct a solution to the dynamics of Figure 10.2 for the (γ^2, η) pairs illustrated in Figure 10.3 and to explicitly describe the stationary law of the corresponding ν_t. The construction is explicit enough to enable us to describe basic properties of the QLE growth.

10.4.2 Theorem statements

Before presenting our main results, we need to formalize the scaling symmetry illustrated in Figure 10.4, which in the continuum should be a statement (which holds for any fixed t) about how the boundary measure ν_t changes when $\mathfrak{h}_t$ is locally transformed via an LQG coordinate change. It is a bit delicate to formulate this, since this should be an almost sure statement (i.e., it should hold almost surely for the $\mathfrak{h}_t$ that one observes in a random solution, but not necessarily for *all* possible $\mathfrak{h}_t$ choices) and one would not necessarily expect a coordinate change such as the one described in Figure 10.4 to preserve the probability measure on $\mathfrak{h}_t$, or even that the law of the image would be absolutely continuous with respect to the law of the original. However, we believe that it *would* be reasonable to expect the law of the restriction of $\mathfrak{h}_t$ to the intervals I_i in Figure 10.4 to change in an absolutely continuous way. (This is certainly the case when $\mathfrak{h}_t$ is a free boundary Gaussian free field plus a smooth deterministic function; see the many similar statements in [SS13].) In this case, one can couple two instances of the field in such a way that one looks like a quantum coordinate change of the other (via a map such as the one described in Figure 10.4) with positive probability. Given a coupling of this type one can formalize the η-DBM scaling symmetry, as we do in the following definition:

Definition 10.2. We say that a triple $(\nu_t, g_t, \mathfrak{h}_t)$ that forms a solution to the dynamics described in Figure 10.2 satisfies **η-DBM scaling** if the following is true. Suppose that we are given any two instances $(\nu_t, g_t, \mathfrak{h}_t)$ and $(\widetilde{\nu}_t, \widetilde{g}_t, \widetilde{\mathfrak{h}}_t)$ coupled in such a way that for a fixed value of $t_0 \geq 0$ and a fixed conformal map ψ from a subset of $\mathbf{D}$ to $\mathbf{D}$, there is a positive probability of the event $\mathcal{A}$ that $\widetilde{\mathfrak{h}}_{t_0}(u) = \mathfrak{h}_{t_0} \circ \psi(u) + Q \log |\psi'(u)|$ for all $u \in I$ where I is an arc of $\partial \mathbf{D}$ with $\psi(I) \subseteq \partial \mathbf{D}$. More precisely, this means that

$$\lim_{\substack{u \to I \\ u \in \mathbf{D}}} \left(\widetilde{\mathfrak{h}}_{t_0}(u) - \mathfrak{h}_{t_0} \circ \psi(u) - Q \log |\psi'(u)| \right) = 0$$

and it says that $\mathfrak{h}_{t_0}$ and $\widetilde{\mathfrak{h}}_{t_0}$ are related by an LQG quantum coordinate change (as in (10.1)). Then we have almost surely on $\mathcal{A}$ that

$$A \mapsto \nu_{t_0}(\psi(A)) \quad \text{and} \quad A \mapsto \int_A |\psi'(u)|^{2+\eta} d\widetilde{\nu}_{t_0}(u) \tag{10.5}$$

agree as measures on I, up to a global multiplicative constant.

Our first result is the existence of stationary solutions to the dynamics described in Figure 10.2 that satisfy η-DBM scaling for appropriate η values. (The existence of the trivial solution corresponding to $\alpha = 0$, ν_t given by uniform Lebesgue measure for all t, and $\eta = -1$, i.e. to the bottom line in Figure 10.3, is obvious and hence omitted from the theorem statement, since in this case the measures ν_t do not depend on h and (10.5) is a straightforward change of coordinates.)

Theorem 10.3. *For each $\gamma \in (0,2]$ and $\eta \in \{\frac{3\gamma^2}{16} - \frac{1}{2}, \frac{3}{\gamma^2} - \frac{1}{2}\}$ (so that (γ^2, η) lies on one of the upper two curves in Figure 10.3), there is a $(\nu_t, g_t, \mathfrak{h}_t)$ triple that forms a solution to the dynamics described in Figure 10.2 and that satisfies η-DBM scaling. Moreover,*

(i) The triple can be constructed using an explicit limiting procedure that involves radial SLE_κ, where $\kappa = \gamma^2$ when $\eta = \frac{3}{\gamma^2} - \frac{1}{2}$ and $\kappa = 16/\gamma^2$ when $\eta = \frac{3\gamma^2}{16} - \frac{1}{2}$. In this solution, the α appearing in Figure 10.3 is equal to $-\frac{1}{\sqrt{\kappa}}$ and

$$h = \widetilde{h} - \frac{\kappa + 6}{2\sqrt{\kappa}} \log|\cdot| + \frac{2}{\sqrt{\kappa}} \log|\cdot - u|$$

where $\widetilde{h}$ is a free boundary GFF on $\mathbf{D}$ where u is a uniformly chosen random point on $\partial \mathbf{D}$ independent of $\widetilde{h}$.

(ii) The pair $(\nu_t, \mathfrak{h}_t)$ is stationary with respect to capacity (i.e., minus log conformal radius) time.

We note that the case $\kappa = \gamma^2$ so that $\eta = \frac{3}{\gamma^2} - \frac{1}{2}$ corresponds to the upper curve in Figure 10.3 and the case $\kappa = 16/\gamma^2$ so that $\eta = \frac{3\gamma^2}{16} - \frac{1}{2}$ corresponds to the middle curve in Figure 10.3.

We also note that for the h in the statement of Theorem 10.3 we have that there is an infinite amount of quantum mass in any neighborhood of the origin.

The significance of the value of $\alpha = -\frac{1}{\sqrt{\kappa}}$ in the statement of Theorem 10.3 is that the law of the pair (h, u) described in the statement is invariant under the operation of resampling u from the measure (formally described by) $e^{\alpha h}$ on $\partial \mathbf{D}$. (See Proposition 5.1 below for more.)

The solutions described in Theorem 10.3 will be constructed as subsequential limits of certain approximations involving SLE. Although we cannot prove that the limits are unique, we can prove that *every* subsequential limit of these approximations has the properties described in Theorem 10.3 (and in particular has the same stationary distribution, described in terms in of the GFF). We will write $\mathrm{QLE}(\gamma^2, \eta)$ to refer to one of these solutions. That these solutions satisfy η-DBM scaling will turn out to follow easily from the fact that $\mathfrak{h}_t$, for each $t \geq 0$, is given by the harmonic extension of the boundary values of a form of the GFF and ν_t is simply a type of LQG boundary measure corresponding to that GFF instance; these points will be explained in Section 10.10 and Theorem 10.3.

Our next result is the Hölder continuity of the complementary component of a $\mathrm{QLE}(\gamma^2, \eta)$ which contains the origin.

Theorem 10.4. *Fix $\gamma \in (0,2)$, let $Q = 2/\gamma + \gamma/2$, and let*

$$\overline{\Delta} = \frac{Q-2}{Q+2\sqrt{2}}. \tag{10.6}$$

Fix $\Delta \in (0, \overline{\Delta})$. Suppose that $(\nu_t, g_t, \mathfrak{h}_t)$ is one of the $\mathrm{QLE}(\gamma^2, \eta)$ processes described in Theorem 10.3. For each $t \geq 0$, let $\mathbf{D}_t = \mathbf{D} \setminus K_t$. Then $\mathbf{D}_t$ is almost surely a Hölder domain with exponent Δ. That is, for each $t \geq 0$, $g_t^{-1} \colon \mathbf{D} \to \mathbf{D}_t$ is almost surely Hölder continuous with exponent Δ.

In fact, the proof of Theorem 10.4 will only use the fact the stationary law of $\mathfrak{h}_t$ is given by the harmonic extension of the boundary values of a form of the GFF; if we could somehow construct other solutions to the QLE dynamics with this property, then this theorem would apply to those solutions as well.

Theorem 10.4 is a special case of a more general result which holds for any random closed set A which is coupled with h in a certain manner. This is stated as Theorem 7.9 in Section **??**. Another special case of this result is the fact that the complementary components of SLE_κ for $\kappa \neq 4$ are Hölder domains. This fact was first proved by Rohde and Schramm in [RS05b, Theorem 5.2] in a very different way.

Suppose that $K \subseteq \mathbf{D}$ is a closed set. Then K is said to be **conformally removable** if the only maps $\varphi \colon \mathbf{D} \to \mathbf{C}$ which are homeomorphisms of $\mathbf{D}$ and conformal on $\mathbf{D} \setminus K$ are the maps which are conformal transformations of $\mathbf{D}$. The removability of the curves coupled with the GFF which arise in this theory is important because it is closely related to the question of whether the curve is almost surely determined by the GFF [She10]. One important consequence of Theorem 10.4 and [JS00b, Corollary 2] is the removability of component boundaries of a $\mathrm{QLE}(\gamma^2, \eta)$ when (γ^2, η) lies on one of the upper two curves of Figure 10.3 and $\gamma \in (0,2)$.

Corollary 10.5. *Suppose that we have the same setup as in Theorem 10.4. For each $t \geq 0$ we almost surely have that $\partial \mathbf{D}_t$ is conformally removable.*

The particular law of h described in the statements above (a free boundary GFF with certain logarithmic singularity at the origin and another logarithmic singularity at a prescribed boundary point) may seem fairly specific. Both singularities are necessary for our particular method of constructing a solution to the QLE dynamics (which uses ordinary radial SLE and the quantum gravity zipper). However, we stress that once one obtains a solution for this particular law for h, one gets for free a solution corresponding to *any* random h whose law is absolutely continuous with respect to that law, since one can always weight the law of the collection $\big(h, (\nu_t, g_t, \mathfrak{h}_t)\big)$ by a Radon-Nikodym derivative depending only on h without affecting any almost sure statements.

In particular, it turns out that adding the logarithmic singularity (which is not too large) centered at the uniformly chosen boundary point changes the overall law of h in an absolutely continuous way (in fact the Radon-Nikodym derivative has an explicit interpretation in terms of the total mass of a certain LQG boundary measure; see the discussion in [DS11a]). Also, adding any finite Dirichlet energy function to h changes the law in an absolutely continuous way. In particular, one could add to h a finite Dirichlet energy function that agrees with a multiple of $\log|\cdot|$ outside a neighborhood U of the origin; a corresponding QLE would then be well defined up until the process first reaches U. Since this can be done for any arbitrarily small U, one can obtain in this way a (not-necessarily-stationary) solution to the QLE dynamics that involves replacing the multiple of the logarithm in the definition of h with another multiple of the logarithm (or removing this term altogether).

Figures 10.20 and 10.21 illustrate the changes that occur in the simulations when different multiples of $\log|\cdot|$ are added to h. As explained in [She10, Section 1.6] (in the case of a wedge, which is in contrast to the case of a cone that we consider here), adding a $\log|\cdot|$ singularity to the GFF has the interpretation of first starting off with a cone and then conformally mapping to $\mathbf{C}$ with the conic singularity sent to the origin. Adding a negative multiple of $\log|\cdot|$ corresponds to an opening angle smaller than 2π and a positive multiple corresponds to an opening angle larger than 2π. This is why the simulations of the Eden model (resp. DLA) in Figure 10.20 (resp. Figure 10.21) appear more and more round (resp. have more arms) as one goes from left to right and then from top to bottom.

10.4.3 Sketch of constructions and proofs

We are going to provide a short sketch here of the arguments used to prove Theorem 10.3 and Theorem 10.4. The starting point is a radial version of the quantum zipper [She10], which is established in Section **??**. This result states that the following is true. Suppose that $\widetilde{h}$ is an instance of the free boundary GFF on $\mathbf{D}$, $u \in \partial\mathbf{D}$ is distributed according to Lebesgue measure on $\partial\mathbf{D}$ independently of $\widetilde{h}$, and

$$h = \widetilde{h} - \frac{\kappa+6}{2\sqrt{\kappa}}\log|\cdot| + \frac{2}{\sqrt{\kappa}}\log|\cdot - u|.$$

Then the law of (h,u) is invariant under the following operation. Suppose that η is a radial SLE_κ process in $\mathbf{D}$ from u to 0 sampled conditionally independently of h given u and let (g_t) be the corresponding Loewner flow (parameterized by capacity). Then the law of (h,u) is equal to the law of the pair consisting of the field

$$h \circ g_t^{-1} + Q\log|(g_t^{-1})'| \tag{10.7}$$

where $Q = 2/\gamma + \gamma/2$, $\gamma = \min(\sqrt{\kappa}, 4/\sqrt{\kappa})$, and the marked point on $\partial\mathbf{D}$ given by the image of η at time t under g_t, i.e. $g_t(\eta(t))$.

Let ν be the Liouville quantum gravity boundary measure on $\partial\mathbf{D}$ which is formally given by $\exp(-\frac{1}{\sqrt{\kappa}}h)$. The pair (h,u) is *also* invariant under the operation of resampling u according to ν.

Fix $\delta > 0$. Starting with the pair (h, u), we then:

1. Draw a radial SLE_κ path in $\mathbf{D}$ from u to 0 which is sampled conditionally independently of h given u for δ units of capacity time and then map back using the change of coordinates in (10.7). This gives us a new field, say $\widehat{h}$, and a marked boundary point, say $\widehat{v}$ (which corresponds to the image of the tip of the path at time δ under the Loewner flow).
2. Replace $\widehat{v}$ with $\widehat{u}$ which is sampled using the Liouville quantum gravity boundary measure $\exp(-\frac{1}{\sqrt{\kappa}}\widehat{h})$.
3. Repeat using the pair $(\widehat{h}, \widehat{u})$ as the starting point.

This defines a growth process that we call the δ-approximation to QLE. QLE itself is defined as a subsequential limit of δ-approximations along a sequence (δ_k) of positive numbers decreasing to 0. As we will explain in Section 10.14, it is not difficult to see that such a subsequential limit gives rise to a triple which satisfies the bottom and right arrows in Figure 10.2. The bulk of Section 10.14 is focused on proving that the subsequential limit satisfies the left arrow of Figure 10.2.

Theorem 10.4 is proved by using the stationarity of the construction and the change of coordinates formula in order to bound the derivative behavior of the inverse of the family conformal maps (g_t) associated with the QLE.

10.5 Interpretation and conjecture when η is large

In the physics literature, there has been some discussion and debate about what happens to the η-DBM model (in the Euclidean setting, i.e., $\gamma = 0$) when η is large. Generally, it is understood that when η is large, there could be a strong enough preference for growth to occur at the "tip" that the scaling limit of η-DBM could in principle be a simple path. There has also been some discussion on the matter of whether one actually obtains a one-dimensional path when η is above some critical value. Some support for this idea with a critical value of about $\eta = 4$ appears in [Has01] and [MJ02]. (The latter contains a figure depicting a simulation of the $\eta = 3$ DBM.) However, a later study estimates the dimension of η-DBM in more detail and does not find evidence for a phase transition at $\eta = 4$, and concludes that the dimension of η-DBM is about 1.08 when $\eta = 4$ [MJB08]. Another reasonable guess might be that the scaling limit of η-DBM is indeed a simple path when η is large enough, but that the simple path may be an SLE with a small value of κ (and not necessarily a straight line).

As a reasonable toy model for this scenario, and a model that is also interesting in its own right, one may consider a variant of η-DBM in which, at each step, one *conditions* on having the next edge added begin *exactly* at the tip of where the last edge was added (so that a simple path is produced in the end). That is, instead of choosing a new edge from the set of all cluster adjacent edges (with probability proportional to harmonic measure to the η power) one chooses a new edge from the set of edges beginning at the current tip of the path (again with probability proportional to harmonic measure to the η power). This random non-self-intersecting walk is sometimes called the *Laplacian random walk* (LRW) with parameter η.[33] Lawler has proposed (citing early calculations by Hastings) that the η-LRW should have SLE as a scaling limit (on an ordinary grid) with

$$\eta = \frac{6-\kappa}{2\kappa}, \tag{10.8}$$

[Law06, LEP86, Has02] at least when $\eta \geq 1/4$, which corresponds to $\kappa \in (0,4]$. Interestingly, if we set $\kappa = \gamma^2$, then (10.8) states that $\eta = 3/\gamma^2 - 1/2$, which corresponds to the upper curve in Figure 10.3. Simulations have shown that η-LRW for large η looks fairly similar to a straight line [BRH10], as one would expect.

At this point, there are two natural guesses that come to mind:

1. Maybe the conjecture about η-LRW on a grid scaling to SLE_κ holds in more generality, so that the scaling limit of η-LRW on a γ-LQG is also given by SLE_κ with κ as in (10.8). (Note that it is often natural to guess that processes that converge to SLE on fixed lattices also converge to SLE when drawn on γ-LQG type random graphs, assuming the latter are embedded in the plane in a conformal way [DS11a, She10].)

2. Maybe, for each fixed γ, it is the case that when η is large enough, the scaling limit of η-DBM on a γ-LQG is the same as the scaling of η-LRW on a γ-LQG. (If in the η-DBM model, the growth tends to take place near the tip, maybe the behavior does not change so much when one requires the growth to take place *exactly* at the tip.)

The authors do not have a good deal of evidence supporting these guesses. However it is interesting to observe that if these guesses are correct, then for sufficiently large η, the η-DBM on a γ-LQG has a scaling limit given by SLE_κ for the κ obtained from (10.8), and this scaling limit does not actually depend on the value of γ. If this is the case, then (at least for η sufficiently large) the dotted line in Figure 10.3 would represent (γ^2, η) pairs for which the scaling limit of η-DBM on a γ-LQG is described by the ordinary radial quantum gravity zipper. We remark that η-DBM scaling as formulated above might be satisfied in a fairly empty way for this process (when ν_t is a point mass for all t), but the property may have some content if one considers an initial configuration in which there are two or more distinct tips (i.e., ν_t contains atoms but is not entirely concentrated at a single point).

[33]The term "Laplacian-b random walk", with parameter $b = \eta$, is also used.

We will not further speculate on the large η case or further discuss scenarios in which ν_t might contain atoms. Indeed, throughout the remainder of the paper, we will mostly limit our discussion to the solid portions of the upper two curves in Figure 10.3, and the ν_t we construct in these settings will be almost surely non-atomic for all almost all t.

10.6 Reshuffled Markov chains

At the heart of our discrete constructions lies a very simple observation about Markov chains. Consider a measure space S which is a disjoint union of spaces $S_1, S_2, \ldots, S_N$. In the examples of this section, S will be a finite set. Suppose we have a measure μ_r defined on each S_r for $1 \le r \le N$. Let $X = (X_k)$ be any Markov chain on S with the property that for any r, j, and $\mathcal{S} \subseteq S_r$, we have

$$\mathbf{P}[X_j \in \mathcal{S} \,|\, X_j \in S_r] = \mu_r(\mathcal{S}).$$

This property in particular implies that the conditional law of X_0, given that it belongs to S_r, is given by μ_r.

Then there is a *reshuffled Markov chain* $Y = (Y_k)$ defined as follows. First, Y_0 has the same law as X_0. Then, to take a step in the reshuffled Markov chain from a point x, one

1. Chooses a point $y \in S$ according to the transition rule for the Markov chain X (from the point x), and then
2. Chooses a new point z from μ_r, where r is the value for which $y \in S_r$.

The step from x to z is a step in the reshuffled Markov chain (and subsequent steps are taken in the same manner). Intuitively, one can think of the reshuffled Markov chain as a Markov chain in which one imposes a certain degree of forgetfulness: if we are given the value Y_i, then in order to sample Y_{i+1} we can imagine that we first take a transition step from Y_i and then we "forget" everything we know about the new location except which of the sets S_r it belongs to — since we cannot remember where within S_r we are supposed to be, we resample a location randomly from the corresponding μ_r.

Now suppose that A is a union of some of the S_r and is a sink of the reshuffled Markov chain (i.e., once the Markov chain enters A, it almost surely does not leave). Then we have the following:

Proposition 10.6. *In the context described above, for each fixed $j \ge 0$, the law of X_j is equivalent to the law of Y_j. Moreover, the law of $\min\{j : X_j \in A\}$ agrees with the law of $\min\{j : Y_j \in A\}$.*

Proof. The first statement holds for $j = 0$ and follows for all $j > 0$ by induction. The second statement follows from the first, since for any j the probability that A has been reached by step j is the same for both Markov chains. □

Our aim in the next two subsections is to show two things:

1. The Eden model on a random triangulation can be understood as a reshuffled percolation interface exploration on that triangulation.

2. DLA on a random planar map can be understood as a reshuffled loop-erased random walk on that map.

In order to establish these results, and to apply Proposition 10.6, we will need to decide in each setting what information we keep track of (i.e., what information is contained in the state Y_i) and what information we forget (i.e., what information we lose when we remember which S_r the state Y_i belongs to and forget everything else — informally, the information we forget should be thought of as a subset of the information we keep track of). In both settings, the information we keep track of will be

1. the structure of the "unexplored" region of a random planar map,

2. the location of a "target" within that region,

3. the location of a "tip" on the boundary of the unexplored region.

In other words, an element of the state space will consist of a planar map (representing an unexplored region), a target location, and a tip location, and the Y_i will be elements of this state space. Also in both settings, the information that we "forget" (before immediately resampling) is the location of the tip. Note that in both settings, both the original Markov chain and the reshuffled Markov chain are defined on the same state space — i.e., the information one "keeps track of" is the same for both of them. In the constructions we introduce below, the location of the tip will be used to decide which triangle/edge to reveal in the growth process — it is part of the information needed to describe a step in the original Markov chain. In the reshuffled Markov chain, after revealing this triangle/edge, we will "forget" the location of the tip and immediately resample it from its conditional law before taking another step according to the original Markov chain.

Thus in both the percolation/Eden model and LERW/DLA settings we will replace a path that grows continuously with a growth process that grows from multiple locations. In both cases, a natural sink (to which one could apply Theorem 10.6) is the "terminal" state obtained when the exploration process reaches its target. The total number of steps in the exploration path agrees in law with the total number of steps in the reshuffled variant. We will also find additional symmetries (and a "slot machine" decomposition) in the percolation/Eden model setting.

10.7 The Eden model and percolation interface

10.7.1 Finite volume Eden/percolation relationship

The following definitions and basic facts are lifted from the overview of planar triangulations given by Angel and Schramm in [AS03] (which cites many of these results from other sources, including [Ang03]). Throughout this section we consider only so-called "type II triangulations," i.e., triangulations whose graphs have no loops but may have multiple edges. For integers $n, m \geq 0$, [AS03] defines $\phi_{n,m}$ to be the number of triangulations of a disk (rooted at a boundary edge) with $m+2$ boundary edges and n internal vertices, giving in [AS03, Theorem 2.1] the explicit formula[34]:

$$\phi_{n,m} = \frac{2^{n+1}(2m+1)!(2m+3n)!}{(m!)^2 n!(2m+2n+2)!}. \tag{10.9}$$

By convention $\phi_{0,0} = 1$ because when the external face is a 2-gon, one possible way to "fill in" the inside is simply to glue the external edges together, with no additional vertices, edges, or triangles inside (and this is in fact the only possibility). As $n \to \infty$,

$$\phi_{n,m} \sim C_m \alpha^n n^{-5/2}, \tag{10.10}$$

where $\alpha = 27/2$ and

$$C_m = \frac{\sqrt{3}(2m+1)!}{2\sqrt{\pi}(m!)^2}(9/4)^m \sim C 9^m m^{1/2}. \tag{10.11}$$

(Both (10.10) and (10.11) are stated just after [AS03, Theorem 2.1].)

Figure 10.22 shows a triangulation T of the sphere with two distinguished edges e_1 and e_2, and the caption describes a mechanism for choosing a random path in the dual graph of the triangulation, consisting of distinct triangles $t_1, t_2, \ldots, t_k$, that goes from e_1 to e_2. It will be useful to imagine that we begin with a single 2-gon and then grow the path dynamically, exploring new territory as we go. At any given step, we keep track of the total number edges on the boundary of the already-explored region and the number of vertices remaining to be seen in the component of the unexplored region that contains the target edge. The caption of Figure 10.23 explains one step of the exploration process. This exploration procedure is closely related to the peeling process described in [Ang03], which is one mechanism for sampling a triangulation of the sphere by "exploring" new triangles one at a time. The exploration process induces a Markov chain on the set of pairs (m, n) with $m \geq 0$ and $n \geq 0$. In this chain, the n coordinate is almost surely non-increasing, and the m coordinate can only increase by 1 when the n coordinate decreases by 1.

[34] In [AS03] a superscript 2 is added to $\phi_{n,m}$ to emphasize that the statement is for type II triangulations. We omit this superscript since we only work with triangulations of this type.

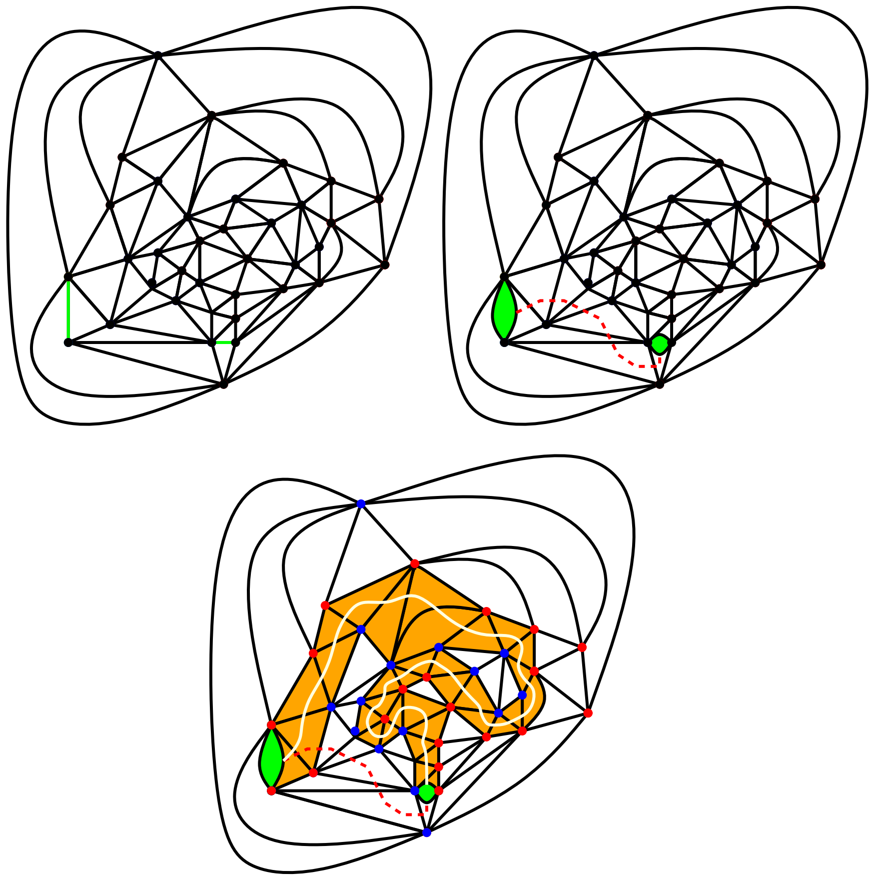

Figure 10.22: **Upper left:** a triangulation of the sphere together with two distinguished edges colored green. **Upper right:** It is conceptually useful to "fatten" each green edge into a 2-gon. We fix a distinguished non-self-intersecting dual-lattice path p (dotted red line) from one 2-gon to the other. **Bottom:** Vertices are colored red or blue with i.i.d. fair coins. There is then a unique dual-lattice path from one 2-gon to the other (triangles in the path colored orange) such that each edge it crosses either has opposite-colored endpoints and *does not* cross p, or has same-colored endpoints and *does* cross p. The law of the orange path does not depend on the choice of p, since shifting p across a vertex has the same effect as flipping the color of that vertex. (Readers familiar with this terminology will recognize the orange path as a percolation interface of an antisymmetric coloring of the double cover of the complement of the 2-gons. Here "antisymmetric" means the two liftings of a vertex have opposite colors.) When the triangulations are embedded in the sphere in a conformal way, the conjectural scaling limit of the path is a whole plane SLE_6 between the two endpoints.

Now consider the version of the Eden model in which new triangles are only added to the unexplored region containing the target edge, as illustrated Figure 10.24. In both Figure 10.22 and Figure 10.24, each time an exploration step separates the unexplored

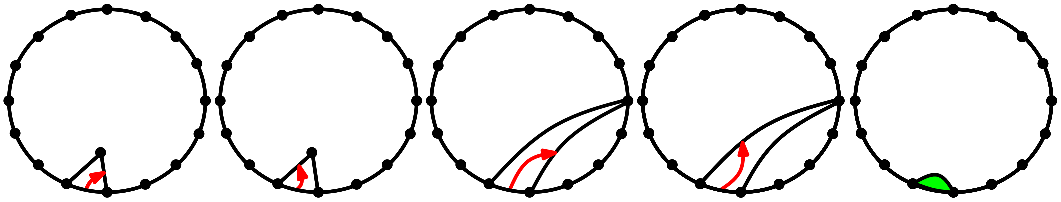

Figure 10.23: Begin with a polygon with $m+2$ edges (for some $m \geq 0$) and a fixed seed edge on the boundary (from which the exploration will take place). Suppose we wish to construct a triangulation of the polygon with $n \geq 0$ additional vertices in the interior. Observe by an easy induction argument that n and m together determine the number of triangles in this triangulation: $m + 2n$. They also determine the number of edges (including boundary edges): $m + 2 + 3n$. The total number of possible triangulations is $\phi_{m,n}$, and for each triangulation there are $(m + 2 + 3n)$ choices for the location of the green edge. The exploration ends if the face incident to the seed edge is the green 2-gon, as in the right figure, which has probability $(m + 2 + 3n)^{-1}$. Conditioned on this not occurring, the probability that we see a triangle with a new vertex (as in the left two figures) is given by $\phi_{m+1,n-1}/\phi_{m,n}$, and given this, the two directions are equally likely (and depend on the coin toss determining the vertex color). In the third and fourth pictures, the exploration step involves deciding both the location of the new vertex (how many steps it is away from the seed edge, counting clockwise) *and* how many of the remaining interior vertices will appear on the right side. We can work out the number of triangulations consistent with each choice: it is given by the product $\phi_{m_1,n_1}\phi_{m_2,n_2}$ where (m_i, n_i) are the new (m, n) values associated to the two unexplored regions. (The choices are constrained by $m_1 + m_2 = m - 1$ and $n_1 + n_2 = n$.) The probability of such a choice is therefore given by this value divided by $\phi_{m,n}$. Once that choice is made, we have to decide whether the step corresponds to the third or fourth figure shown — i.e., whether the green edge is somewhere in the left unexplored region or the right unexplored region. The probability of it being in the first region is the number of edges in that region divided by the total number of edges (excluding the seed edge, since we are already conditioning on the seed not being the target): $(m_1 + 2 + 3n_1)/(m + 1 + 3n)$. In each of the first four figures, we end up with a new unexplored polygon-bounded region known to contain the target green edge, and a new (m, n) pair. We may thus begin a new exploration step starting with this pair and continue until the target is reached.

region into two pieces (each containing at least one triangle) we refer to the one that does not contain the target as a *bubble*. The exploration process described in Figure 10.22 created two bubbles (the two small white components), and the exploration process described in Figure 10.24 created one (colored blue). We can interpret the bubble as a triangulation of a polygon, rooted at a boundary edge (the edge it shares with the triangle that was observed when the bubble was created).

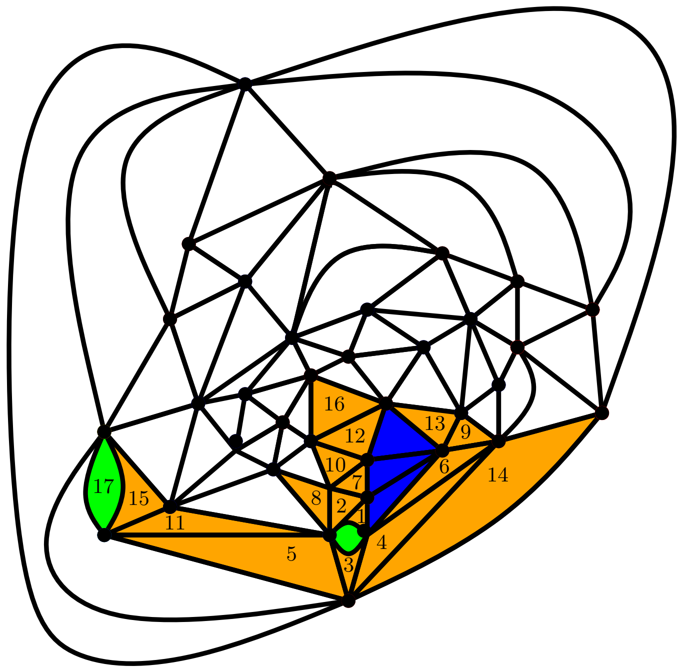

Figure 10.24: Same as Figure 10.22 except that one explores using the Eden model instead of percolation. At each step, one chooses a uniformly random edge on the boundary of the unexplored region containing the target and explores the face incident to that edge. The faces are numbered according to the order in which they were explored. When the unexplored region is divided into two pieces, each with one or more triangles, the piece without the target is called a *bubble* and is never subsequently explored by this process. In this figure there is only one bubble, which is colored blue.

The specific growth pattern in Figure 10.24 is very different from the one depicted in Figure 10.22. However, the analysis used in Figure 10.23 applies equally well to both scenarios. The only difference between the two is that in Figure 10.24 one re-randomizes the seed edge (choosing it uniformly from all possible values) after each step.

In either of these models, we can define C_k to be the boundary of the target-containing unexplored region after k steps. If (M_k, N_k) is the corresponding Markov chain (where M_k is the value of m and N_k is the value of n after k steps of the exploration), then the length of C_k is $M_k + 2$ for each k. Let D_k denote the union of the edges and vertices in C_k, the edges and vertices in C_{k-1} and the triangle and bubble (if applicable) added at step k, as in Figure 10.25. We refer to each D_k as a *necklace* since it typically contains a cycle of edges together with a cluster of one or more triangles hanging off of it. The analysis used in Figure 10.23 (and discussed above) immediately implies the following (parts of which could also be obtained from Proposition 10.6):

Proposition 10.7. *Consider a random rooted triangulation of the sphere with a fixed number $n > 2$ of vertices together with two distinguished edges chosen uniformly from* *the set of possible edges. (Using the Euler characteristic and the fact that edges and faces are in* 2 *to* 3 *correspondence, it is clear that this triangulation contains* $2(n-2)$ *triangles and* $3(n-2)$ *edges.) If we start at one edge and explore using the Eden model as in Figure 10.24, or if we explore using the percolation interface of Figure 10.22, we will find that the following are the same:*

(i) The law of the Markov chain (M_k, N_k) *(which terminates when the target 2-gon is reached).*

(ii) The law of the total number of triangles observed before the target is reached.

(iii) The law of the sequence D_k *of necklaces.*

Indeed, one way to construct an instance of the Eden model process is to start with an instance of the percolation interface exploration process and then randomly rotate the necklaces in the manner illustrated in Figure 10.25.

In the setting of Proposition 10.7, the two randomly chosen edges respectively play the role of the initial and terminal (i.e., target) edge in the percolation exploration. As the percolation and Eden exploration grow, they both separate the triangulation into multiple components and the target edge determines the component into which the growth continues.

10.7.2 Infinite volume Eden/percolation relationship

In [Ang03] Angel gives a very explicit construction of the *uniform infinite planar triangulation* (UIPT), which is further investigated by Angel and Schramm in [AS03]. The authors in [AS03] define τ_n to be the uniform distribution on rooted type II triangulations of the sphere with n vertices and show that the measures τ_n converge as $n \to \infty$ (in an appropriate topology) to an infinite volume limit called the *uniform infinite planar triangulation* (UIPT). (Related work on infinite planar quadrangulations appears in [CMM13].) The convenient property that this process possesses is that the number n of remaining vertices is always infinite, and hence, in the analog of the Markov chain described in Figure 10.23, it is only necessary to keep track of the single number m, instead of the pair (m, n). A very explicit description of this Markov chain and the law of the corresponding necklaces appears in [Ang03, AS03]. As in the finite volume case, the sequence of necklaces has the same law in the UIPT Eden model as in the UIPT percolation interface exploration. One can first choose the necklaces associated to a UIPT percolation interface model and then randomly rotate them (by "pulling the slot machine lever") to obtain an instance of the UIPT Eden model, as in Figure 10.25.

Later in this paper, we will interpret the version of QLE$(8/3, 0)$ that we construct as a continuum analog of the Eden model on the UIPT. The construction will begin with

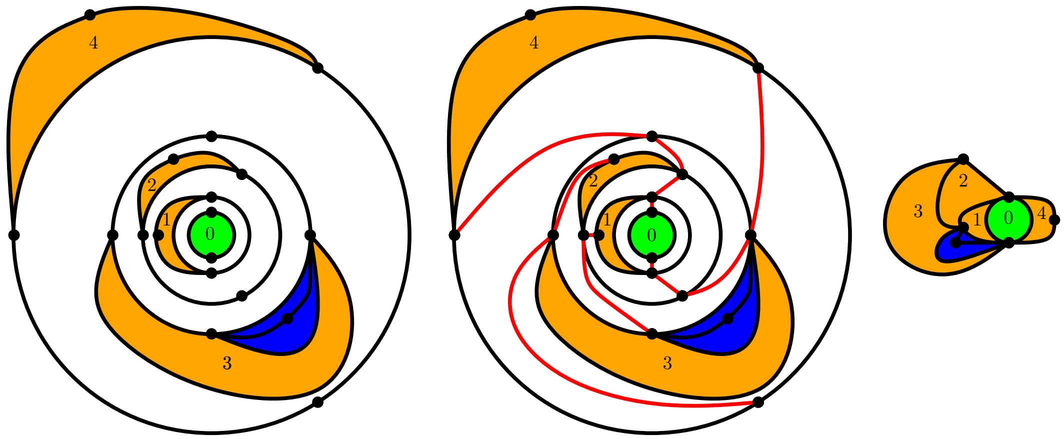

Figure 10.25: **Left:** the first four necklaces (separated by white space) generated by an Eden model exploration. **Middle:** one possible way of identifying the vertices on the outside of each necklace with those on the inside of the next necklace outward. **Right:** The map with exploration associated to this identification. If a necklaces has n vertices on its outer boundary, then there are n ways to glue this outer boundary to the inner boundary of the next necklace outward. It is natural to choose one of these ways uniformly at random, independently for each consecutive pair of necklaces. Intuitively, we imagine that before gluing them together, we randomly spin the necklaces like the reels of a slot machine, as in Figure 10.26. A fanciful interpretation of Proposition 10.7 is that if we take a percolation interface exploration as in Figure 10.22 (which describes a sequence of necklaces) and we pull the slot machine lever, then we end up with an Eden model exploration of the type shown in Figure 10.24. In later sections, this paper will discuss a continuum analog of "pulling the slot machine lever" that involves SLE and LQG.

the continuum analog of the percolation exploration process on the UIPT, which is a radial SLE_6 exploration on a certain type of LQG surface. We will then "rerandomize the tip" at discrete time intervals, and we will then find a limit of these processes when the interval size tends to zero.

Finally, we remark that Gill and Rohde have recently established parabolicity of the Riemann surfaces obtained by gluing triangles together [GR13], which implies that the UIPT as a triangulation can be conformally mapped onto the entire complex plane, as one would expect.

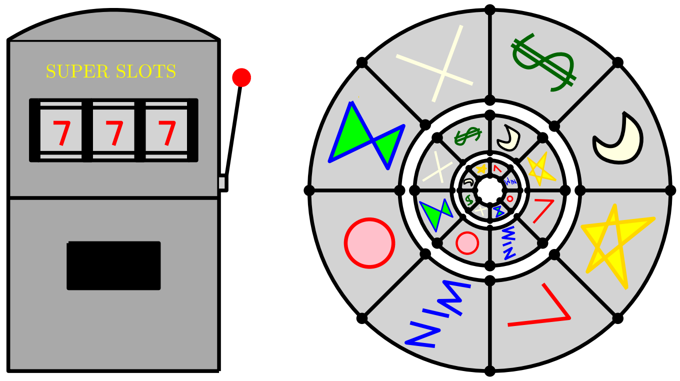

Figure 10.26: **Left:** sketch of an actual slot machine. **Right:** sketch of slot machine reels conformally mapped from cylinder to plane. When the lever is pulled, each of the reels rotates a random amount.

10.8 DLA and the loop-erased random walk

10.8.1 Finite volume DLA/LERW relationship

The uniform spanning-tree-decorated random planar map is one of the simplest and most elegant of the planar map models, due to the relationship with simple random walks described by Mullin in 1967 [Mul67b] (and explained in more detail by Bernardi in [Ber07]) which we briefly explain in Figure 10.27 and Figure 10.28 (which are lifted from a more detailed exposition in [She11c]). As the caption to Figure 10.27 explains, one first observes a correspondence between planar maps and quadrangulations: there is a natural quadrangulation such that each edge of the original map corresponds to a quadrilateral (whose vertices correspond to the two endpoints and the two dual endpoints of that edge). As the caption to Figure 10.28 explains, one may then draw diagonals in these quadrilaterals corresponding to edges of the tree or the dual tree.

If an adjacent vertex and dual vertex are fixed and designated as the root and dual root (big dots in Figure 10.28) then one can form a cyclic path starting at that edge that passes through each green edge once, always with blue on the left and red on the right. To the kth green edge that the path encounters (after one spanning root and dual root) we assign a pair of integers (x_k, y_k), where x_k is the distance of the edge's left vertex to the root within the tree, and y_k is the distance from its right vertex to the dual root within the dual tree. If n is the number of edges in the original map, then the sequence $(x_0, y_0), (x_1, y_1), (x_2, y_2), \ldots, (x_{2n}, y_{2n})$ is a walk in $\mathbf{Z}_+^2$ beginning and ending at the origin, and it is not hard to see that there is a one-to-one correspondence between walks of this type and rooted spanning-tree-decorated maps with n edges, such as the one illustrated in Figures 10.27 and 10.28. The walks of this type with $2m$ left-right steps and $2n - 2m$ up-down steps correspond to the planar maps with m edges in the tree (hence $m + 1$ vertices total in the original planar map) and $n - m$ edges in the dual tree (hence $n - m + 1$ faces total in the original planar map). Once m and n are fixed (it is natural to take $m \approx n/2$), it is easy to sample the spanning-tree-decorated rooted planar map by sampling the corresponding random walk.

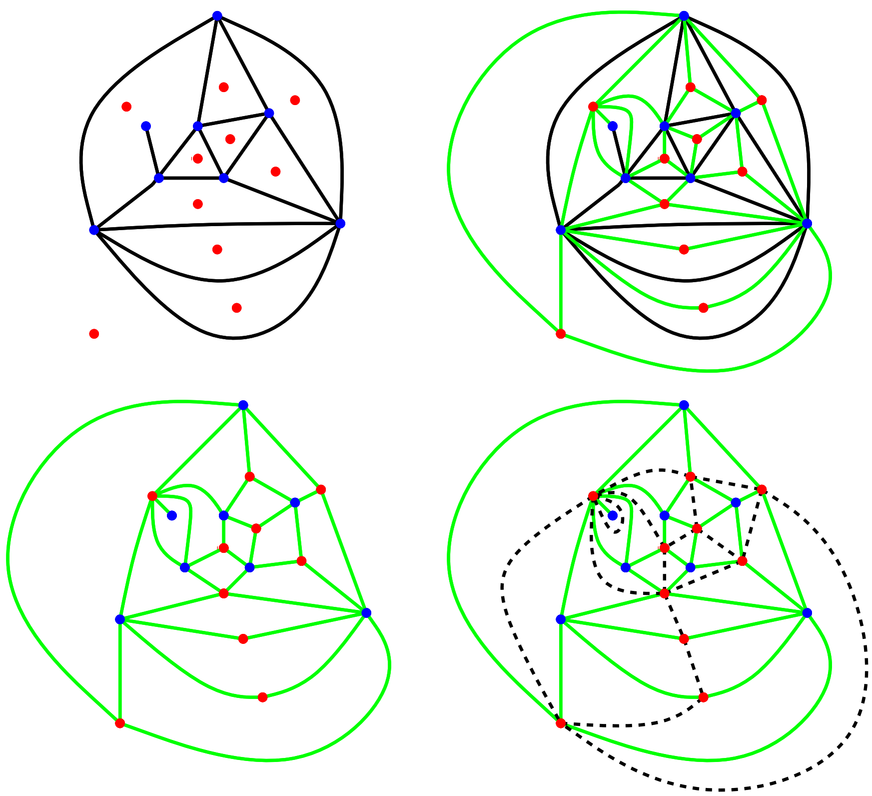

Figure 10.27: **Upper left:** a planar map M with vertices in blue and "dual vertices" (one for each face) shown in red. **Upper right:** the quadrangulation $Q = Q(M)$ formed by adding a green edge joining each red vertex to each of the boundary vertices of the corresponding face. **Lower left:** quadrangulation Q shown without M. **Lower right:** the dual map M' corresponding to the same quadrangulation, obtained by replacing the blue-to-blue edge in each quadrilateral with the opposite (red-to-red) diagonal.

As shown in Figure 10.29, if we endow the map with two distinguished vertices, a "seed" and a "target" then there is a path from the seed to the target and a deterministic procedure for "unzipping" the edges of the path one at a time, to produce (at each step) a new planar map with a distinguished grey polygon that has a marked tip vertex ("zipper handle") on its boundary. This procedure is also reversible — i.e., if we see one of the later decorated maps in Figure 10.29, then we have enough information to recover the earlier figures.

It is possible to consider the same procedure but keep track of less information: one can imagine a version of Figure 10.29 in which all of the edges colored black or green (except those on the boundary of the grey polygon) were colored red, like the first two

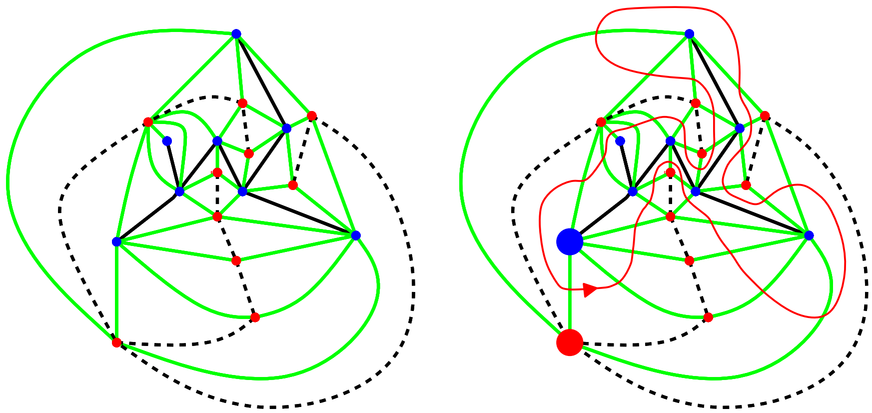

Figure 10.28: **Left:** in each quadrilateral we either draw an edge (connecting blue to blue) or the corresponding dual edge (connecting red to red). In this example, the edges drawn form a spanning tree of the original (blue-vertex) graph, and hence the dual edges drawn form a spanning tree of the dual (red-vertex) graph. **Right:** designate a "root" (large blue dot) and an adjacent "dual root" (large red dot). The red path starts at the midpoint of the green edge between the root and the dual root and crosses each of the green edges once, keeping the blue endpoint to the left and red endpoint to the right, until it returns to the starting position. Each endpoint corresponds to a pair of vertices

maps shown in Figure 10.30. To put ourselves in the context of Proposition 10.6, we can let X_i be the decorated planar map (the planar map endowed with a distinguished grey face with a marked blue tip on its boundary, and a distinguished green target vertex) obtained after unzipping i steps. By Wilson's algorithm [Wil96], if one is given the first k steps of the path from the seed to the target, then the conditional law of the remaining edges is the law of the loop erasure of a simple random walk started at the target and conditioned to hit the grey polygon for the first time at the blue tip vertex (whereupon the walk is terminated). In particular, this tells us how to perform the Markov transition step from X_i to X_{i+1}. Namely, one chooses an edge incident to the tip with probability proportional to its harmonic measure (viewed from the target), colors that edge green, and "unzips" it by sliding up the blue tip, as in the first transition step shown in Figure 10.30.

Using the notation of Proposition 10.6, we can let the sets S_r be the equivalence classes of decorated maps, where two maps are considered equivalent if they agree except that their blue tips are at different locations (on the boundary of the same grey polygon). Conditioned on $X_i \in S_r$, it is not hard to see that the conditional law of the tip location is given by harmonic measure viewed from the target. This is because, once we condition on S_r, one can treat the grey polygon as a single vertex, and note that all spanning trees of the collapsed graph are equally likely; hence, one can therefore use Wilson's algorithm to sample the path from the target to the grey polygon, and the law of the

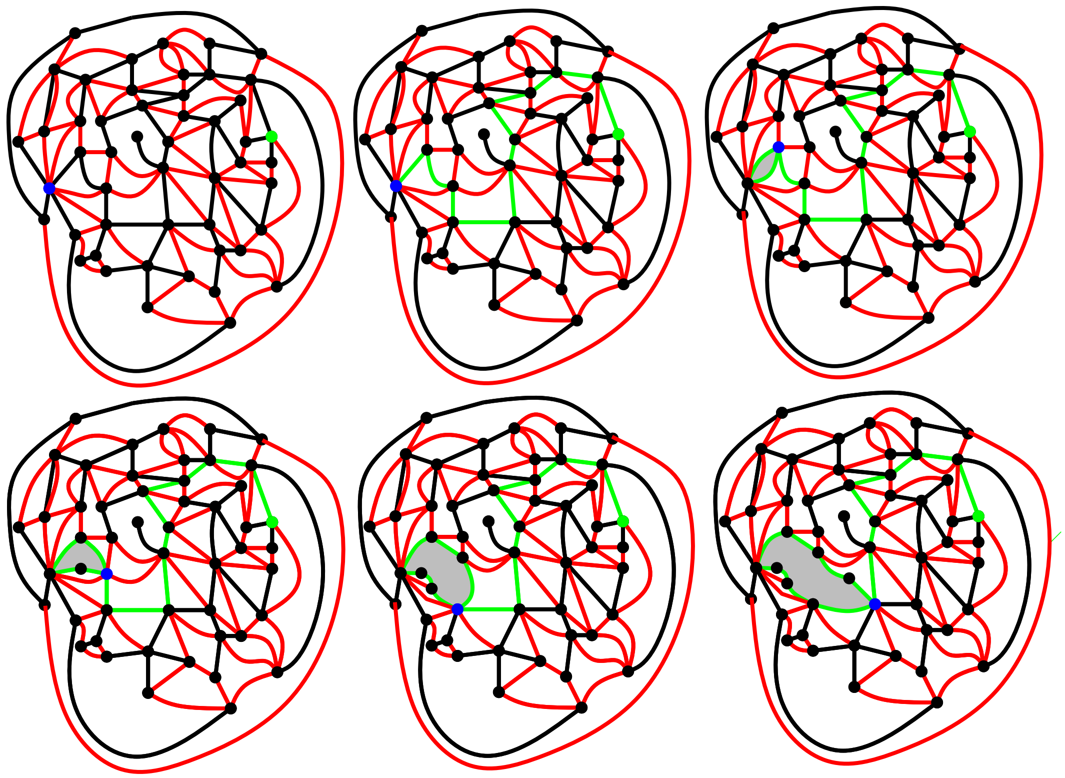

Figure 10.29: **Upper left:** A planar map with distinguished spanning tree (tree edges black, other edges red) along with distinguished "seed" and "target" vertices (colored green). Assume the tree-decorated map is chosen uniformly from the set of tree-decorated maps with a given number of vertices and edges, and that the seed and target are then uniformly chosen vertices. **Upper middle:** tree path from seed to target colored green. **Upper right:** think of the blue dot as a zipper handle, and the green path as the closed zipper; we slide the blue dot up one step and "unzip" the first edge by splitting it in two to form 2-gon (with inside colored grey). **Lower left to lower right:** second, third, fourth edges along path are similarly unzipped, to produce 4-gon, 6-gon, 8-gon. Given the initial tree-decorated map and seed/target vertices, the unzipping procedure is deterministic.

location at which it exits is indeed given by harmonic measure. Thus, within each S_r we can define the measure μ_r on decorated maps such that sampling from μ_r amounts to re-sampling the seed vertex from this harmonic measure.

One can now define the reshuffled Markov chain $Y_0, Y_1, \dots$ using precisely the procedure described in Proposition 10.6. This chain has the same transition law as the unreshuffled chain except that after each step we resample the blue tip from the harmonic measure viewed from the target, as explained in Figure 10.30. As explained in Figure 10.30, this reshuffling procedure converts loop-erased random walk (LERW) to diffusion limited

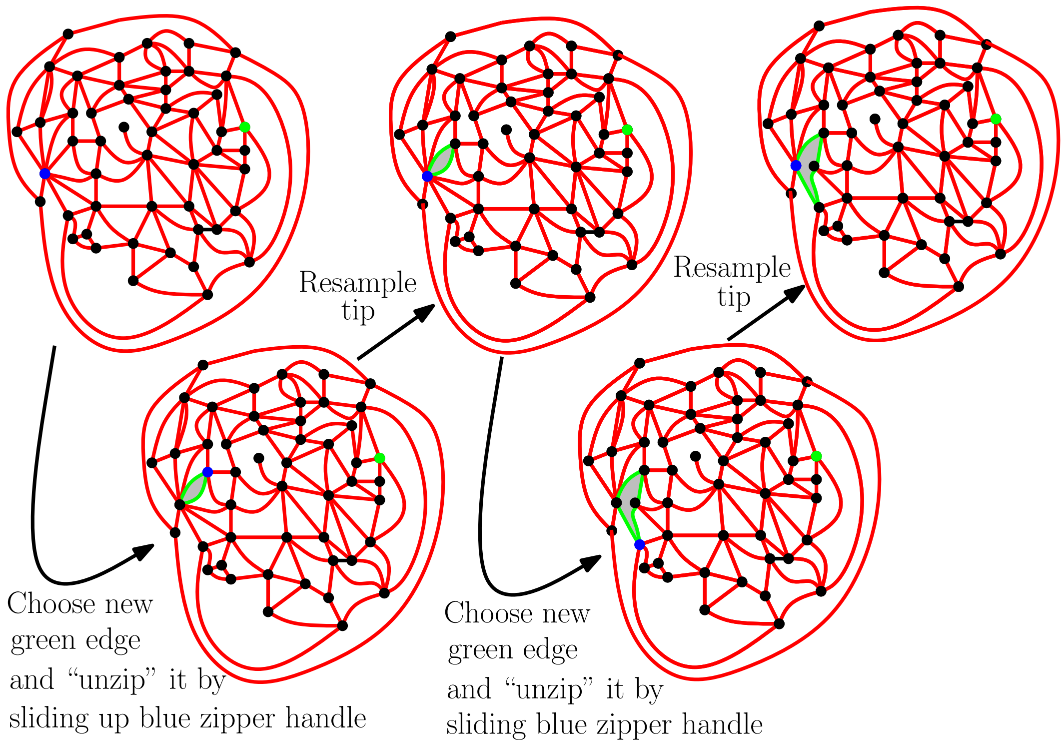

Figure 10.30: We could have drawn the images in Figure 10.29 with a different coloring — showing all edges red except for those around the grey polygon. With such a coloring, we could imagine that we do not know the tree in advance: we only discover the path from seed to tip one edge at a time. Conditioned on the first k edges, Wilson's algorithm implies that the probability that a given tip-adjacent edge e is the next edge in the path is proportional to the probability that a random walk from the target first reaches the grey polygon via e. After selecting an edge, we color it green and unzip it by sliding up the blue zipper handle, tracing a path whose overall law is that of a LERW from tip to seed. The process shown in the figure is a "reshuffled" version of the one just described. After an edge is drawn, we "resample" the blue vertex according to harmonic measure viewed from the target and then choose a sample green edge from that vertex. We can equivalently combine the resample-tip and pick-new-edge steps by performing a random walk from the target and picking the last edge traversed before the grey polygon is hit. The order in which edges are "unzipped" in this reshuffled form of LERW is the same as the order in which edges are discovered in edge-growth DLA. In this setting, DLA is nothing more than reshuffled LERW.

aggregation (DLA). The following is now immediate from Proposition 10.6.

Proposition 10.8. *Consider a random rooted planar map M with n edges, $m+1$ vertices, and $n-m+1$ faces, with two of the vertices designated "seed" and "target" chosen uniformly among all such decorated maps except that the probability of a given* *decorated map is proportional to the number of spanning trees the map has. Conditional on M, one may generate a loop-erased random walk L from the seed to target. Given M, one may also generate an edge-based DLA growth process, which yields a random tree D containing the seed and the target.*

(i) The number of edges in D agrees in law with the number of edges in L.

(ii) The law of the map obtained by unzipping the first k steps of L (to produce the grey polygon with distinguished tip, as in Figure 10.29) is the same as the law of the map obtained by unzipping the first k steps of D, as in Figure 10.29.

(iii) The law of the map obtained by unzipping all of the edges of L agrees in law with the map obtained by unzipping all of the edges of D.

10.8.2 Infinite volume DLA/LERW relationship

In this section, we will observe that the constructions of the previous section can be extended to the so-called *uniform infinite planar tree-decorated map* (UIPTM). We present this infinite volume construction partly because of its intrinsic interest, and partly because we believe that the form of QLE$(2,1)$ that we construct in this paper is the scaling limit of DLA on the UIPTM.

We define the UIPTM to be the infinite volume limit of the models of random rooted planar maps described in the previous section as $n \to \infty$ and $m = n/2$. More discussion of this model appears in [She11c] and in the work of Gill and Rohde in [GR13]. The latter showed that the Riemannian surface defined by gluing together the triangles in the UIPTM is parabolic (like the analogous surface defined using the UIPT). Gurel-Gurevich and Nachmias also recently proved a very general recurrence statement for random planar maps, which implies that if we forget the spanning tree on the UIPTM and simply run a random walk on the vertices of the underlying graph, then this walk is almost surely recurrent [GGN13]. (Their work also implies that random walk on the UIPT is almost surely recurrent, and extends the earlier recurrence results obtained by Benjamini and Schramm in [BS01].)

Given the walk (x_k, y_k) described in the previous section, we may write $I_k = (x_k, y_k) - (x_{k-1}, y_{k-1})$. The I_k are random variables taking values in

$$\{(-1,0),(1,0),(0,-1),(0,1)\}.$$

There is a one-to-one correspondence between steps of type $(1,0)$ and vertices v of the planar map (discounting the root vertex) since the red path in Figure 10.28 first encounters a green edge incident to a vertex v at step k if and only if $I_k = (1,0)$.

If k is such that I_k corresponds to a chosen seed vertex v, then we may recenter time so that this vertex corresponds to the first increment. That is, we define a new centered increment process: $\widetilde{I}_j = I_{j-k}$. It is not hard to see that in the limit as $n \to \infty$ and $m = n/2$, the $\widetilde{I}_j$ converge to a process indexed by $\mathbf{Z}$ in which $I_1 = (1,0)$ almost surely but the other I_i are i.i.d. uniformly chosen elements from $\{(-1,0),(1,0),(0,-1),(0,1)\}$. The use of doubly infinite sequences of this form to describe random surfaces is discussed in more detail in [She11c]. In this description, the I_i are the increments of a walk on $\mathbf{Z}^2$ (parameterized by $\mathbf{Z}$) and the x (resp. y) coordinate of this walk determines the structure of the infinite tree (resp. infinite dual tree). In particular, it is easy to see from this construction that the infinite tree (which can be described by a simple random walk on $\mathbf{Z}$) a.s. has a single end, so that there is a unique infinite simple path in the tree in the UIPTM that extends from the seed vertex to infinity.

Proposition 10.9. *If one samples a UIPTM (which is an infinite rooted planar map M endowed with a spanning tree T and a root vertex v) and then samples a tree T' on M according to Wilson's algorithm, then the law of (M, T', v) is again that of a UIPTM.*

Proof. It is shown in [BLPS01, Theorem 5.6] that for any recurrent graph the tree generated by Wilson's algorithm (with any choice of vertex order) agrees in law with the so-called wired spanning forest, and also with the so-called free spanning forest. In particular, this implies that Wilson's algorithm determines a unique random tree on M (independent of the vertex order) and we just have to show that the law of this tree agrees with the conditional law of T given M.

Let M_n be the random tree-decorated rooted planar map obtained with n edges and $m = n/2$ vertices, T_n the corresponding spanning tree, and v_n the corresponding seed vertex. The proposition will follow from the fact that (M_n, T_n, v) converges in law to (M, T, v), that M is almost surely recurrent, and that for any n one can first sample (M_n, v) and then use Wilson's algorithm to sample T_n.

To explain this in more detail, note that it suffices to show that for any $N > 0$ the law of (M, T) restricted to the ball of radius N about v agrees with the law of (M, T') restricted to the ball of radius N about v. Now, the recurrence of M implies that for any $\delta > 0$ we can choose N' large enough so that if we run Wilson's algorithm starting at all points within $B(v, N)$, to obtain the shortest tree path from each of these points to v, we find that the probability that any of these paths reaches distance N' from v is at most δ.

Take n large enough so that (M_n, T_n, v_n) and (M, T, v) can be coupled in such a way that their restrictions to the N' ball about their origin vertices agree with probability at least $1 - \delta$. It then follows that we can couple (M_n, T_n, v_n) with (M, T', v) so that they agree within a radius N ball of their origin vertices with probability at least $1 - 2\delta$. Since δ can be made arbitrarily small (by taking N' large enough) this completes the argument. □

Based on Proposition 10.9, we find that the law of the branch of the tree from the origin to ∞ can be obtained as a limit of the law of the loop erased random walk from w to v, as the distance from w to v tends to ∞. In particular, the limit of the harmonic measure of the possible next edges to be added to this LERW (as measured from w as this distance from v to w tends to infinity) exists, and one can grow the branch from v to ∞ one step at a time by sampling from the tip according to this measure, using the procedure indicated in Figure 10.29.

Noting that (M, T, v) is the limit of the (M_n, T_n, v_n), we also find that the conditional law of the location of the tip (given the grey polygon and the map but not tip location) is given by harmonic measure, and hence we can obtain the infinite volume analog of Proposition 10.8 using the same argument used in the proof of Proposition 10.8.

10.9 "Capacity" time parameterization

We discuss here a natural stochastic way of reparameterizing an η-DBM growth according to capacity (as opposed to according to the number of edges added, which is the natural choice of parameterization at the discrete level). Figure 10.4 will be closely related to this discussion.

In each of the models in this section, for any edge e on the boundary of the cluster, we can let $b(e)$ denote the harmonic measure at edge e as viewed from the target. When considering possible scaling limits of the discrete models in this section, we should keep in mind that heuristically the "capacity" added to the cluster by putting in the new edge should be roughly proportional to $b(e)^2$. (In the continuum, drawing a slit of length ϵ from $\partial\mathbf{D}$ towards the origin changes the conformal radius of the remaining domain viewed from the origin by order ϵ^2.) Thus, the amount of "capacity" time corresponding to a given step in a discrete model is random. One might therefore try to reparameterize time in the discrete models in such a way that one might expect to obtain a scaling limit parameterized by capacity (i.e., negative the log conformal radius).

One way to do this with the η-DBM model is as follows: suppose that $b(e)$ represents the harmonic measure at an edge e viewed from the target. Then at each step, we:

1. Choose an edge e with probability proportional $b(e)^{2+\eta}$.
2. Given e, we toss a coin that is heads with probability proportional to $b(e)^{-2}$.
3. Add e to the cluster only if the coin comes up heads.

Note that an edge e is added to the cluster with probability proportional to $b(e)^{\eta}$. Therefore (up to a random time change) this construction is equivalent to the usual η-DBM model. However, in contrast to the usual parameterization for the η-DBM model, the expected amount of capacity added to the cluster at each time step is of the same order.

Another approach is to say that after an edge is selected, instead of flipping a coin that is heads with probability proportional to $b(e)^{-2}$, we simply only add a "$b(e)^{-2}$ sized portion of the edge" (i.e., we don't consider an edge to have been "added" until it has been hit multiple times, and the sum of all of these fractional contributions exceeds some large constant).

We note that the probability $b(e)^{2+\eta}$ in the sampling procedure described just above is related to the factor $|\psi'|^{2+\eta}$ which appears in Figure 10.4 (and in other places in this article). Indeed, if X represents the η-DBM cluster at a given time and ψ is the unique conformal map which takes the complement of X to $\mathbf{C} \setminus \mathbf{D}$ and both fixes and has positive derivative at ∞, then $|\psi'|$ evaluated near e approximates the harmonic measure of e as seen from ∞.

We mention these alternatives, because the approximations to the continuum construction of QLE we present in this paper will involve random increments of constant capacity (i.e., constant change to the log conformal radius), and the scaling limit will be parameterized by capacity. One could modify the continuum construction (adding increments of constant quantum length instead of constant capacity) but this will not be our first approach.

10.10 Continuum interpretation of QLE

The QLE dynamics described in Figure 10.2 involves two parameters: γ and α. Here γ describes the type of LQG surface on which the growth process takes place and α determines the multiple of $\mathfrak{h}_t$ used in the exponentiation that generates ν_t. As discussed in Section 10.4, once one has a solution to the dynamics for a given α and γ pair, one can seek to verify that the solution satisfies η-DBM scaling, as defined in Definition 10.2, for some value of η.

It is natural to wonder whether, for each γ value, there is a one-to-one correspondence between α and η values (at least over some range of the parameters). This is not a question we will settle in this paper, as we will only construct (and determine α and η for) certain families of QLE processes, and these correspond to points on the curves in Figure 10.3.

However, in Section 10.11 we will propose a relationship between α, β, γ, and η where β (introduced in Section 10.11 below) is an additional parameter that appears in the regularization used to make sense of $e^{\alpha\mathfrak{h}_t}$, and in some sense encodes how fast $e^{\alpha\mathfrak{h}_t}$ blows up near $\partial\mathbf{D}$.

In full generality, this calculation should be taken as a heuristic (since we do not know that β is defined for general solutions to the QLE dynamics) but it can be made rigorous under some assumptions — for example, if one assumes that the stationary law of $\mathfrak{h}_t$ is given by a free boundary Gaussian free field (restricted to $\partial\mathbf{D}$ and harmonically extended to $\mathbf{D}$). This latter assumption will turn out to imply that $\beta = \alpha^2$ and hence (for each fixed γ) it determines a relationship between α and η. This assumption turns out to hold for the solutions we construct from the quantum zipper (corresponding to the upper two curves in Figure 10.3) and this gives us a way to recover η from α for these solutions, as we discuss in Section 10.12.

In Section 10.13 we will argue that when $\eta = 0$ the $\beta = \alpha^2$ assumption leads to a prediction of the dimension for the γ-LQG surfaces when these surfaces are understood as metric spaces. (We stress that endowing a γ-LQG surface with a metric space structure has never been done rigorously, but we *believe* that such a metric should exist and that a ball in this metric, whose radius increases in time, should be described by a QLE$(\gamma^2, 0)$ process.) The dimension prediction we obtain agrees with a prediction made in the physics literature by Watabiki in [Wat93]. (The fact that our formula agrees with Watabiki's derivation was pointed out to us by Duplantier.) As mentioned above, the $\beta = \alpha^2$ assumption would hold if the stationary law of $\mathfrak{h}_t$ for a QLE$(\gamma^2, 0)$ process were given by a free boundary Gaussian free field (harmonically extended from $\partial\mathbf{D}$ to $\mathbf{D}$). However, we do not currently have a compelling heuristic to suggest why a stationary law for QLE$(\gamma^2, 0)$ should have this form.

10.11 Scaling exponents: a relationship between α, β, γ, η

The caption to Figure 10.2 describes a particular way to make sense of the map from $\mathfrak{h}_t$ to ν_t. Precisely, we let ν_t be the $n \to \infty$ limit of the measures $e^{\alpha \mathfrak{h}_t^n(u)} du$ on $\partial\mathbf{D}$, normalized to be probability measures; recall that the $\mathfrak{h}_t^n$ are obtained by throwing out all but the first n terms in the power series expansion of the analytic function with real part $\mathfrak{h}_t$. (This can be understood as a projection of the GFF onto a finite dimensional subspace.)

Instead of using the power series approximations or other projections of the GFF onto finite dimensional subspaces, another natural approach would be to use approximations $\mathfrak{h}_t^\epsilon$ to $\mathfrak{h}_t$ defined by "something like" convolving $\mathfrak{h}_t$ with a bump function supported (or mostly supported) on an interval with length of order ϵ. For example, we could write $\mathfrak{h}_t^\epsilon(u) = \mathfrak{h}_t\big((1-\epsilon)u\big)$ for each $u \in \partial\mathbf{D}$. (This is equivalent to convolving with a bump function related to the Poisson kernel.) Or we could let $\mathfrak{h}_t^\epsilon(u)$ be the mean of $\mathfrak{h}_t$ on $\partial B(u,\epsilon) \cap \mathbf{D}$ for each $u \in \partial\mathbf{D}$. To describe another approach (which involves more of the unexplored field than just the harmonic projection), let us simplify notation for now by writing h for the sum of $\mathfrak{h}_t$ and an independent zero boundary GFF on $\mathbf{D}$ and let $\mathfrak{h}_t^\epsilon = h^\epsilon$ be the mean value of h on $\partial B(u,\epsilon) \cap \mathbf{D}$. (The latter definition of $\mathfrak{h}_t^\epsilon$ is essentially what is used in [DS11a] to define boundary measures when h is an instance of the free boundary GFF.)

For now, let us assume that the following are true:

1. The boundary values of h are such that it is possible to make sense of the average $h^\epsilon(z)$ of h on $\partial B(z,\epsilon)$ for each $z \in \overline{\mathbf{D}}$ and $\epsilon > 0$.
2. $h^\epsilon(u)$ blows up to $\pm\infty$ almost surely for each $u \in \partial\mathbf{D}$ as $\epsilon \to 0$ (as is the case when h is given by the form of the free boundary GFF considered in Theorem 10.3).
3. There exists a constant β such that the following limit exists and is almost surely a non-zero finite measure:

$$\nu_h = \lim_{\epsilon \to 0} \epsilon^\beta e^{\alpha h^\epsilon(u)} du. \tag{10.12}$$

(This limit turns out not to depend on the zero-boundary GFF used in the definition of h [DS11a].)

In a sense, β encodes the growth rate of $e^{\alpha \mathfrak{h}_t}$ near $\partial\mathbf{D}$. Note that when describing the dynamics of Figure 10.2, we avoided having to specify a regularizing factor such as ϵ^β (or an analogous factor depending on n) because we normalized to make each approximation a probability measure.

In the case that h is the free boundary GFF and $\alpha \in (-1,1)$ so that ν_h is given by the 2α-LQG boundary measure, β is given by $(2\alpha)^2/4 = \alpha^2$ [DS11a]. [35]

For $\gamma > 0$ given, let $Q_\gamma = 2/\gamma + \gamma/2$. Recall that Q_γ is the factor the appears in front of the log-derivative in the γ-LQG coordinate change described in (10.1). We are going to derive the following relationship between α, γ, η, and β:

$$\alpha Q_\gamma = \beta - \eta - 1. \tag{10.13}$$

Once three of the variables α, β, γ, and η are fixed we can use (10.13) to determine the fourth. Moreover, once γ is fixed, (10.13) gives an affine relationship between α, β, and η.

Let $\psi \colon \mathbf{D} \to \widetilde{D}$ be a conformal change of coordinates. Let

$$\widetilde{Q} := \frac{1}{\alpha} + \frac{\beta}{\alpha} = \frac{1+\beta}{\alpha}$$

and let $\widetilde{h}$ be the distribution on $\widetilde{D}$ given by

$$\widetilde{h} = h \circ \psi^{-1} + \widetilde{Q} \log |(\psi^{-1})'|. \tag{10.14}$$

Let $\nu_{\widetilde{h}}$ be the boundary measure as in (10.12) defined in terms of $\widetilde{h}$. Then it is not hard to see (at least if ψ is linear) from (10.12) that $\nu_{\widetilde{h}}$ is almost surely the image

[35] In [DS11a, Section 6], the existence of the limit (10.12) is proved when h is given by the free boundary GFF on a domain with piecewise linear boundary while here we are taking our domain to be $\mathbf{D}$. It is easy to see, however, that the argument of [DS11a] also goes through in the case that the domain is $\mathbf{D}$.

under ψ of ν_h. That is, $\nu_h(A) = \nu_{\widetilde{h}}(\psi(A))$ for $A \subseteq \partial\mathbf{D}$. To see this, observe that $e^{\alpha\tilde{Q}\log|(\psi^{-1})'|} = |(\psi^{-1})'|^{1+\beta}$, which is $|(\psi^{-1})'|$ (the ordinary coordinate change term) times $|(\psi^{-1})'|^{\beta}$.

When $\alpha = \gamma/2$ and $\beta = \alpha^2$, the definition (10.14) is the same as the usual change of coordinates formula for the LQG boundary measure [DS11a].

Let ν_γ be the measure on $\partial\widetilde{D}$ which is constructed by replacing $\widetilde{Q}$ in the definition (10.14) of $\widetilde{h}$ with Q_γ. Replacing $\widetilde{Q}$ with Q_γ makes it so that the change of coordinates by ψ preserves the γ-LQG boundary measure defined from h (as opposed to the boundary measure with scaling exponent β as defined in (10.12)). Then the Radon-Nikodym derivative between ν_γ and $\nu_{\widetilde{h}}$ is (formally) given by a constant times

$$\exp(\alpha(Q_\gamma - \widetilde{Q})\log|(\psi^{-1})'|) = |(\psi^{-1})'|^{\alpha(Q_\gamma - \widetilde{Q})}.$$

The application of the conformal transformation ψ scales the harmonic measure of a small region near $\partial\mathbf{D}$ by the factor $|\psi'|$. Recalling the discussion in the caption of Figure 10.4, we want ν_γ to be given by scaling $\nu_{\widetilde{h}}$ by the factor $|\psi'|^{2+\eta}$. We therefore want

$$-\alpha(Q_\gamma - \widetilde{Q}) = 2 + \eta.$$

Plugging in the definition for $\widetilde{Q}$, we have $-\alpha Q_\gamma + 1 + \beta = 2 + \eta$. Rearranging gives (10.13).

10.12 Free boundary GFF and quantum zipper α

Fix $\gamma \in (0, 2]$. Using the quantum zipper machinery, we will find in later sections that it is natural to consider a setting in which $\beta = \alpha^2$ and we have one additional constraint, namely, $\alpha \in \{-\gamma/4, -1/\gamma\}$. These two facts and (10.13) together imply the relationship between η and γ described by the upper two curves in Figure 10.3[36]. Very roughly speaking, the reason is that for these values the 2α-LQG boundary measure is supported on "thick points" u near which the field behaves like $-2\alpha \log|u - \cdot|$ where $2\alpha \in \{-\gamma/2, -2/\gamma\}$ (see [DS11a, Proposition 3.4] for the bulk version of this statement

[36] There is also another heuristic way to determine what α must be when η and γ are given (in the case that $\mathfrak{h}_0$ is a harmonically projected GFF, so that $\beta = \alpha^2$), which would give an alternate derivation of (10.13). This heuristic was shown to us by Bertrand Duplantier. Consider the discrete η-DBM interpretation in which one samples a boundary face (or edge) of the planar map from harmonic measure to the $\eta + 2$ power, and then adds a unit of capacity near the chosen face. Recall that the measure that assigns a unit mass to each face is (conjecturally) supposed to have approximately the form $e^{\gamma h(z)}dz$ for a type of free boundary GFF h. Now, what does the field look like near a "typical" face chosen from harmonic measure to the $\eta + 2$ power? According to the KPZ formalism as applied to "negative dimensional" sets (see the discussion in [DS11a] on non-intersecting Brownian paths), if the face is centered at a point u, then the field near u should look approximately like an ordinary free boundary GFF plus $2\alpha \log|u - \cdot|$, where α and η are related in precisely the manner described here. We hope to explain this point in more detail in a future joint work with Duplantier.

as well as Proposition 5.1 below for the version which will be relevant for this article), and these values have the form $-2/\sqrt{\kappa}$ for $\kappa \in \{16/\gamma^2, \gamma^2\}$, which correspond to the singularities that appear in the capacity invariant quantum zipper.

In these settings we will also find a stationary law of $\mathfrak{h}_t$ given by the harmonic extension of the boundary values of a form of the free boundary GFF on $\mathbf{D}$, and as mentioned earlier, in this setting one has $\beta = \alpha^2$.

If we plug in $\alpha = -1/\gamma$ and $\beta = \alpha^2$ into (10.13) then we obtain:

$$-\frac{1}{\gamma}\left(\frac{2}{\gamma} + \frac{\gamma}{2}\right) = \frac{1}{\gamma^2} - \eta - 1,$$

or equivalently

$$\eta = \frac{3}{\gamma^2} - \frac{1}{2}. \tag{10.15}$$

This describes the upper curve in Figure 10.3

If we plug $\alpha = -\gamma/4$ and $\beta = \alpha^2$ into (10.13) then we obtain

$$-\frac{\gamma}{4}\left(\frac{2}{\gamma} + \frac{\gamma}{2}\right) = \frac{\gamma^2}{16} - \eta - 1,$$

or equivalently

$$\eta = \frac{3\gamma^2}{16} - \frac{1}{2}. \tag{10.16}$$

This describes the middle curve in Figure 10.3. Note that the lower curve in Figure 10.3 corresponds to $\alpha = \beta = 0$ and $\eta = -1$, which is trivially a solution to (10.13) for any γ.

10.13 Free boundary scaling $\beta = \alpha^2$ and $\eta = 0$

In Theorem 10.3 we prove the existence of stationary QLE(γ^2, η) processes for (γ^2, η) pairs which are on one of the upper two curves in Figure 10.3 with $\beta = \alpha^2$. It is natural to wonder whether this is just a coincidence, or whether there are other (γ^2, η) pairs for which there exist QLE solutions with $\beta = \alpha^2$. (This would be the case, for example, if the $\mathfrak{h}_t$ turned out to have stationary laws described by the harmonic extension of the boundary values from $\partial\mathbf{D}$ to $\mathbf{D}$ of a form of the free boundary GFF.) We observe that if we simply plug in $\beta = \alpha^2$, then (10.13) becomes

$$\alpha Q_\gamma = \alpha^2 - \eta - 1,$$

or equivalently

$$\eta = \alpha^2 - \alpha Q_\gamma - 1.$$

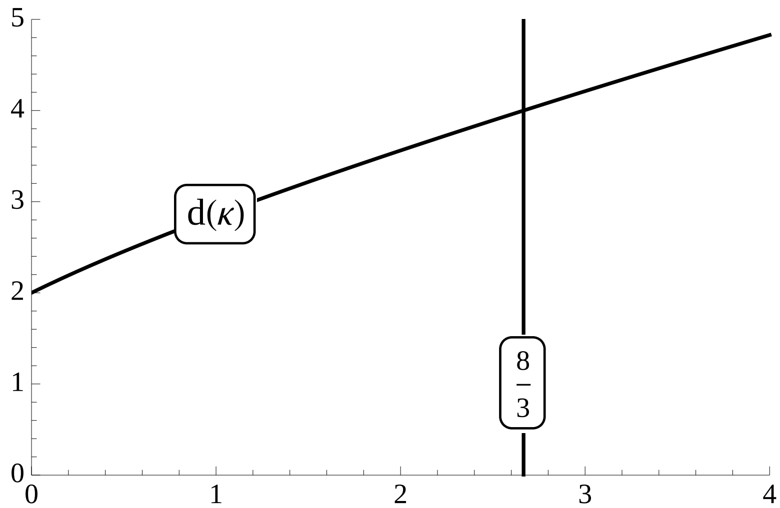

Figure 10.31: The value d as a function of $\kappa = \gamma^2$, as defined by (10.17). Although the graph is not a straight line, it appears "almost straight" and it takes the value 2 for $\kappa = 0$ and 4 for $\kappa = 8/3$.

One can also solve this for α to obtain

$$\alpha = \frac{Q_\gamma \pm \sqrt{Q_\gamma^2 + 4 + 4\eta}}{2}.$$

We now introduce a parameter $d = -\gamma/\alpha$, which can interpreted as a sort of "dimension", at least in the $\eta = 0$ case.[37] Let $A = e^{\gamma C}$. This represents the factor by which the γ-LQG area in a small neighborhood of a boundary point $u \in \partial\mathbf{D}$ changes when we add a function to h that is equal to a constant C in that neighborhood. Then $e^{\alpha C}$ represents the factor by which the QLE driving measure changes, which suggests that the time it takes to traverse the neighborhood should scale like $T = e^{-\alpha C}$. Now d is the value such that $A = T^{-\gamma/\alpha} = T^d$.

Computing this, we have

$$d = -\frac{2\gamma}{Q_\gamma \pm \sqrt{Q_\gamma^2 + 4 + 4\eta}} = \frac{2\gamma\left(Q_\gamma \pm \sqrt{Q_\gamma^2 + 4 + 4\eta}\right)}{4 + 4\eta}$$

[37] One way to define the dimension of a metric space is as the value d such that the number of radius δ balls required to cover the space scales as δ^{-d}. (Hausdorff dimension is a variant of this idea.) If the metric space comes endowed with a measure (and is homogeneous, in some sense) then one might guess that each of these balls would have area of order δ^d. In fact, if there is a natural notion of "rescaling" the metric space so that its diameter changes by a factor of δ (and the measure is also defined for the rescaled version), then one can define d to be such that the area scales as δ^d. In the QLE setting with $\eta = 0$, if we consider a small neighborhood U of a point $u \in \partial\mathbf{D}$, and we rescale the quantum surface restricted to U (by modifying h on U) then we expect the "length of time it takes a QLE to traverse U" to scale by approximately the same factor as the diameter of U (assuming a metric space structure on the quantum surface is defined).

$$= \frac{1}{1+\eta}\left(1 + \frac{\gamma^2}{4} \pm \sqrt{\frac{\gamma^4}{16} + \frac{3\gamma^2}{2} + 1 + \eta\gamma^2}\right).$$

Setting $\kappa = \gamma^2 \in (0,4)$ this is equivalently equal to

$$\frac{1}{1+\eta}\left(1 + \frac{\kappa}{4} \pm \sqrt{\frac{\kappa^2}{16} + \frac{3\kappa}{2} + 1 + \eta\kappa}\right).$$

In the case $\eta = 0$, the positive root can be written as

$$d = 1 + \frac{\kappa}{4} + \frac{1}{4}\sqrt{(4+\kappa)^2 + 16\kappa}. \tag{10.17}$$

The graph of d as a function of $\kappa = \gamma^2$ is illustrated in Figure 10.31. The plot matches a physics literature prediction made by Watabiki in 1993 for the fractal dimension of γ-LQG quantum gravity when understood as a metric space [Wat93, Equation (5.13)].[38] However, we stress again that our calculation was made under the assumption that $\beta = \alpha^2$, and that we do not currently have even a heuristic argument for why there should exist QLE processes satisfying this relationship for $\eta = 0$ and a given $\gamma \in (0,2]$ (though of course the reader may consult the explanation given in [Wat93]). The exception is the case $\gamma = \sqrt{8/3}$, since $(8/3, 0)$ is one of the (γ^2, η) pairs for which we construct solutions to the QLE dynamics. In this case, our arguments do support the notion the Hausdorff dimension of Liouville quantum gravity should be 4 for $\gamma = \sqrt{8/3}$, though we will not prove this statement in this paper. This is consistent with the dimension of the Brownian map [CS04, LG07b]

10.14 Existence of QLE

The purpose of this section is to prove Theorem 10.3. Throughout, we suppose that the pair (γ^2, η) is on one of the upper two lines from Figure 10.3. That is, we suppose that (γ^2, η) satisfy either

$$\eta = \frac{3\gamma^2}{16} - \frac{1}{2} \quad \text{or} \quad \eta = \frac{3}{\gamma^2} - \frac{1}{2}.$$

We are going to construct a triple $(\nu_t, g_t, \mathfrak{h}_t)$ which satisfies the dynamics described in Figure 10.2 where

$$\alpha_\kappa = -\frac{1}{\sqrt{\kappa}} \tag{10.18}$$

[38]The quantities α_1 and α_{-1} which appear in [Wat93, Equation (5.13)] are defined in [Wat93, Equation (4.15)]. These, in turn, are defined in terms of the central charge c. The central charge c corresponding to an SLE_κ is $(8-3\kappa)(\kappa-6)/(2\kappa)$; see the introduction of [LSW03b].

for $\kappa > 1$. We will first give a careful definition of the spaces in which our random variables take values in Section 10.14.1. We will then prove Theorem 10.1 in Section 10.14.2. We next introduce approximations $(\varsigma_t^\delta, g_t^\delta, \mathfrak{h}_t^\delta)$ to QLE(γ^2, η) in Section 10.14.3. Throughout, we reserve using the symbol ν to denote a measure which is constructed using exponentiation. This is why the Loewner driving measure for the approximation is referred to as ς_t^δ. We will then show that each of the elements of $(\varsigma_t^\delta, g_t^\delta, \mathfrak{h}_t^\delta)$ is tight on compact time intervals with respect to a suitable topology in Section 10.14.4. Finally, we will show that the subsequentially limiting triple $(\nu_t, g_t, \mathfrak{h}_t)$ satisfies the dynamics from Figure 10.2 in Section 10.15. This will complete the proof of Theorem 10.3.

10.14.1 Spaces, topologies, and σ-algebras

We are going to recall the spaces $\mathcal{N}_T$, $\mathcal{G}_T$, and $\mathcal{H}_T$ (and their infinite time versions) from the introduction as well as introduce a certain subspace of the space of distributions. We will then equip each of these spaces with a metric and the corresponding Borel σ-algebra. We emphasize that each of the spaces that we consider is separable. This will be important later since we will make use of the Skorohod representation theorem for weak convergence.

Measures. We let $\mathcal{N}_T$ be the space of measures ς on $[0,T] \times \partial \mathbf{D}$ whose marginal on $[0,T]$ is given by Lebesgue measure. We equip $\mathcal{N}_T$ with the topology given by weak convergence. That is, we say that a sequence (ς^n) in $\mathcal{N}_T$ converges to $\varsigma \in \mathcal{N}_T$ if for every continuous function ϕ on $[0,T] \times \partial \mathbf{D}$ we have that $\int_{[0,T]\times\partial\mathbf{D}} \phi(s,u) d\varsigma^n(s,u) \to \int_{[0,T]\times\partial\mathbf{D}} \phi(s,u) d\varsigma(s,u)$. Equivalently, we can equip $\mathcal{N}_T$ with the Levy-Prokhorov metric $d_{\mathcal{N},T}$. We let $\mathcal{N}$ be the space of measures ς on $[0,\infty) \times \partial\mathbf{D}$ whose marginal on $[0,\infty)$ is given by Lebesgue measure. Note that there is a natural projection $P_T \colon \mathcal{N} \to \mathcal{N}_T$ given by restriction. We equip $\mathcal{N}$ with the following topology. We say that a sequence (ς^n) in $\mathcal{N}$ converges to ς if $(P_T(\varsigma^n))$ converges to $P_T(\varsigma)$ as a sequence in $\mathcal{N}_T$ for each $T \geq 0$. Equivalently, we can equip $\mathcal{N}$ with the metric $d_{\mathcal{N}}$ given by $\sum_{n=1}^\infty 2^{-n} \min(d_{\mathcal{N},n}(\cdot,\cdot), 1)$. Then $(\mathcal{N}, d_{\mathcal{N}})$ is a separable metric space and we equip $\mathcal{N}$ with the Borel σ-algebra.

Families of conformal maps. We let $\mathcal{G}_T$ be the space of families of conformal maps (g_t) where, for each $0 \leq t \leq T$, $g_t \colon \mathbf{D} \setminus K_t \to \mathbf{D}$ is the unique conformal transformation with $g_t(0) = 0$ and $g_t'(0) > 0$. We assume further that $g_t'(0) = e^t$ so that time is parameterized by log conformal radius. We define $\mathcal{G}$ analogously except time is defined on the interval $[0,\infty)$. We say that a sequence of families (g_t^n) in $\mathcal{G}$ converges to (g_t) if $(g_t^n)^{-1} \to g_t^{-1}$ locally uniformly in space and time. In other words, for each compact set $K \subseteq \mathbf{D}$ and $T \geq 0$ we have that $(g_t^n)^{-1} \to g_t^{-1}$ uniformly on $[0,T] \times K$. We can construct a metric which is compatible with this notion of convergence by taking $d_{\mathcal{G},n}$ to be the uniform distance on functions defined on $B(0, 1-1/n) \times [0,n]$ and then taking $d_{\mathcal{G}}$ to be $\sum_{n=1}^\infty 2^{-n} \min(d_{\mathcal{G},n}(\cdot,\cdot),1)$. Then $(\mathcal{G}, d_{\mathcal{G}})$ is a separable metric space and we equip $\mathcal{G}$ with the corresponding Borel σ-algebra.

Families of harmonic functions. We let $\mathcal{H}_T$ be the space of families of harmonic functions $(\mathfrak{h}_t)$ where, for each $t \in [0,T]$, $\mathfrak{h}_t \colon \mathbf{D} \to \mathbf{R}$ is harmonic, $\mathfrak{h}_t(0) = 0$, and $(t,z) \mapsto \mathfrak{h}_t(z)$ is continuous. We define $\mathcal{H}$ similarly with $T = \infty$. We equip $\mathcal{H}$ with the topology of local uniform convergence. That is, if $(\mathfrak{h}_t^n)$ is a sequence in $\mathcal{H}$ then we say that $(\mathfrak{h}_t^n)$ converges to $(\mathfrak{h}_t)$ if for each compact set $K \subseteq \mathbf{D}$ and $T \geq 0$ we have that $\mathfrak{h}_t^n \to \mathfrak{h}_t$ uniformly on $[0,T] \times K$. We can construct a metric $d_{\mathcal{H}}$ which is compatible with this notion of convergence in a manner which is analogous to $d_{\mathcal{G}}$ and we equip $\mathcal{H}$ with the corresponding Borel σ-algebra.

Distributions. Suppose that (f_n) are the eigenvectors of Δ with Dirichlet boundary conditions on $\mathbf{D}$ with negative eigenvalues (λ_n). By the spectral theorem, (f_n) properly normalized gives an orthonormal basis of $L^2(\mathbf{D})$. Thus for $f \in C_0^\infty(\mathbf{D})$ we can write $f = \sum_n \alpha_n f_n$ and, for $a \in \mathbf{R}$, we define $(-\Delta)^a f = \sum_n \alpha_n (-\lambda_n)^a f_n$. We let $(-\Delta)^a L^2(\mathbf{D})$ denote the Hilbert space closure of $C_0^\infty(\mathbf{D})$ with respect to the inner product $(f,g)_a = ((-\Delta)^{-a} f, (-\Delta)^{-a} g)$ where $(\cdot,\cdot)$ is the $L^2(\mathbf{D})$ inner product; see [She07, Section 2.3] for additional discussion of this space. We equip $(-\Delta)^a L^2(\mathbf{D})$ with the Borel σ-algebra associated with the norm generated by $(\cdot,\cdot)_a$.

The GFF with zero boundary conditions takes values in $(-\Delta)^a L^2(\mathbf{D})$ for each $a > 0$ [She07] (see also [SS13, Section 4.2]). By Proposition 3.2, we can write the GFF on $\mathbf{D}$ with either mixed or free boundary conditions as the sum of a harmonic function and an independent zero-boundary GFF on $\mathbf{D}$. It therefore follows that for each $\epsilon > 0$, each of these fields restricted to $(1-\epsilon)\mathbf{D}$ take values in $(-\Delta)^a L^2((1-\epsilon)\mathbf{D})$. (In the case of free boundary conditions, we can either consider the space modulo additive constant or fix the additive constant in a consistent manner by taking, for example, the mean of the field on $\mathbf{D}$ to be zero.) We let $\mathcal{D}_a^\epsilon$ be the subspace of distributions on $\mathbf{D}$ which are elements of $(-\Delta)^a L^2((1-\epsilon)\mathbf{D})$ and let $d_{a,\epsilon}$ be the metric on $\mathcal{D}_a^\epsilon$ induced by the $(\cdot,\cdot)_a$ inner product. Let $\mathcal{D}_a = \cap_{\epsilon>0} \mathcal{D}_a^\epsilon$ and equip $\mathcal{D}_a$ with the metric given by $d_a(\cdot,\cdot) = \sum_n 2^{-n} \min(d_{a,n^{-1}}(\cdot,\cdot),1)$. Since each $\mathcal{D}_a^\epsilon$ is separable, so is $\mathcal{D}_a$ and we equip it with the Borel σ-algebra.

10.14.2 Proof of Theorem 10.1

Recall that Theorem 10.1 has three assertions. For the convenience of the reader, we restate them here and then give the precise location of where each is established below.

(i) For any $\varsigma \in \mathcal{N}$ there exists a unique solution to the radial Loewner equation (in integrated form) driven by ς. This is proved in Proposition 10.10.

(ii) If we have any increasing family of compact hulls (K_t) in $\overline{\mathbf{D}}$ parameterized by log conformal radius as seen from 0 then there exists a unique measure $\varsigma \in \mathcal{N}$ such that the complement of the domain in $\overline{\mathbf{D}}$ of the solution to the radial Loewner equation driven by ς at time t is given by K_t. This is proved in Proposition 10.13.

(iii) The convergence of a sequence (ς^n) in $\mathcal{N}$ to a limiting measure $\varsigma \in \mathcal{N}$ is equivalent to the Caratheodory convergence of the families of compact hulls in $\overline{\mathbf{D}}$ parameterized by log conformal radius associated with the corresponding radial Loewner chains. That the convergence of such measures implies the Caratheodory convergence of the hulls is proved as part of Proposition 10.10. The reverse implication is proved in Proposition 10.15.

We establish the first assertion of Theorem 10.1 in the following proposition.

Proposition 10.10. *Suppose that $\varsigma \in \mathcal{N}$. Then there exists a unique solution (g_t) to the radial Loewner evolution driven by ς. That is, (g_t) solves*

$$g_t(z) = \int_{[0,t]\times\partial\mathbf{D}} \Phi(u, g_s(u))d\varsigma(s,u), \quad g_0(z) = z. \tag{10.19}$$

Moreover, suppose that (ς^n) is a sequence in $\mathcal{N}$ converging to $\varsigma \in \mathcal{N}$. For each $n \in \mathbf{N}$, let (g_t^n) solve the radial Loewner equation driven by ς^n and likewise let (g_t) solve the radial Loewner equation driven by ς. Then $(g_t^n) \to (g_t)$ as $n \to \infty$ in $\mathcal{G}$.

Before we prove Proposition 10.10, we first collect the following two lemmas.

Lemma 10.11. *Suppose that $\varsigma \in \mathcal{N}$. Then there exists a sequence (ς_t^n) where, for each $n \in \mathbf{N}$ and $t \geq 0$, ς_t^n is a probability measure on $\partial\mathbf{D}$ such that the following are true.*

(i) For each $n \in \mathbf{N}$, $t \mapsto \varsigma_t^n$ is continuous with respect to the weak topology on measures on $\partial\mathbf{D}$.

(ii) We have that $d\varsigma_t^n dt \to \varsigma$ as $n \to \infty$ in $\mathcal{N}$.

Proof. We define ς_t^n by averaging the first coordinate of ς as follows: for $\phi\colon \partial\mathbf{D} \to \mathbf{R}$ continuous, we take

$$\int_{\partial\mathbf{D}} \phi(u)d\varsigma_t^n(u) = n\int_{[t,t+n^{-1}]\times\partial\mathbf{D}} \phi(u)d\varsigma(s,u).$$

Then it is easy to see that the sequence (ς_t^n) has the desired properties. □

Lemma 10.12. *If (ς^n) is a sequence in $\mathcal{N}$ and, for each $n \in \mathbf{N}$, $t \mapsto g_t^n$ solves the radial Loewner equation driven by ς^n, then the following is true. There exists a family of conformal transformations (g_t) which are continuous in both space and time and each of which maps $\mathbf{D}$ into itself and a subsequence $(g_t^{n_k})$ of (g_t^n) such that $(g_t^{n_k}) \to (g_t)$ in $\mathcal{G}$.*

Proof. Let $\psi_t^n = (g_t^n)^{-1}$. Then the chain rule implies that for each $z \in \mathbf{D}$ and $t \geq 0$ we have that

$$\psi_t^n(z) = -z \int_{[0,t]\times\partial\mathbf{D}} (\psi_s^n)'(z)\Psi(u,z)d\varsigma^n(s,u) + z; \tag{10.20}$$

see [Law05b, Remark 4.15]. The desired result follows because it is clear from the form of (10.20) that the family (ψ_t^n) is equicontinuous when restricted to a compact subset of $\mathbf{D}$ and compact interval of time in $[0,\infty)$. □

Proof of Proposition 10.10. We are first going to prove uniqueness of solutions to (10.19). Suppose that we have two solutions (g_t) and $(\widetilde{g}_t)$ to (10.19). Fix $T \geq 0$. Then the domain of g_T (resp. $\widetilde{g}_T$) contains $B(0, \frac{1}{4}e^{-T})$ by the Koebe one-quarter theorem since time is parameterized by log conformal radius. To show that $g_T = \widetilde{g}_T$, it suffices to show that $g_T(z) = \widetilde{g}_T(z)$ for all $z \in B(0, \frac{1}{16}e^{-T})$ because two conformal transformations with connected domain and whose values agree on an open set agree everywhere. For $0 \leq s \leq t \leq T$, we let $g_{s,t} = g_t \circ g_s^{-1}$ and $\widetilde{g}_{s,t} = \widetilde{g}_t \circ \widetilde{g}_s^{-1}$. From the form of the radial Loewner equation it follows that the maps $g_{s,t}$ are Lipschitz in $0 \leq s \leq t \leq T$ and $z \in B(0, \frac{1}{16}e^{-T})$ where the Lispchitz constant only depends on T. By estimating $g_{s,r}$ (resp. $\widetilde{g}_{s,r}$) by z in the integral below, it thus follows that there exists a constant $C > 0$ depending only on T such that

$$\begin{aligned}
&|g_{s,t}(z) - \widetilde{g}_{s,t}(z)| \\
\leq &\int_{[s,t]\times\partial\mathbf{D}} |\Phi(u, g_{s,r}(z)) - \Phi(u, \widetilde{g}_{s,r}(z))|\, d\varsigma(r,u) \\
\leq &C(t-s)^2.
\end{aligned}$$

Fix $\delta > 0$ and let $t_\ell = \delta\ell$ for $\ell \in \mathbf{N}_0$. Then

$$\begin{aligned}
g_T(z) &= g_{t_1,T} \circ g_{t_1}(z) \\
&= g_{t_1,T} \circ \widetilde{g}_{t_1}(z) + \big(g_{t_1,T}(g_{t_1}(z)) - g_{t_1,T}(\widetilde{g}_{t_1}(z))\big).
\end{aligned}$$

By the previous estimate and the Lipschitz property, the second term is of order $O(\delta^2)$ as $\delta \to 0$ where the implicit constant depends only on T. Iterating this procedure implies that $g_T(z) - \widetilde{g}_T(z) = O(\delta)$ as $\delta \to 0$ where the implied constant depends only on T. This implies uniqueness.

We are next going to show that if (ς^n) is a sequence in $\mathcal{N}$ converging to ς and, for each n, (g_t^n) is the solution to the radial Loewner equation with driving function (ς^n), then $(g_t^n) \to (g_t)$ in $\mathcal{G}$ where (g_t) is the radial Loewner equation driven by ς. By possibly passing to a subsequence, Lemma 10.12 implies that there exists a family of conformal maps (g_t) such that $(f_t = g_t^{-1})$ is a locally uniform subsequential limit of (f_t^n) in both space and time. To finish the proof, we just need to show that (g_t) satisfies the radial Loewner equation driven by ς. For each $t \geq 0$ and $z \in \mathbf{D}$ with positive distance from the complement of the domain of g_t, we can write:

$$g_t^n(z) = \int_{[0,t]\times\partial\mathbf{D}} \Phi(u, g_s^n(z))d\varsigma^n(s,u) + z$$

$$= O(t \times \sup_{s\in[0,t]} |g_s^n(z) - g_s(z)|) + \int_{[0,t]\times\partial\mathbf{D}} \Phi(u, g_s(z)) d\varsigma^n(s,u) + z.$$

Taking a limit as $n \to \infty$ of both sides proves the assertion.

It is left to prove existence. In the case that the radial Loewner evolution is driven by a family of measures $t \mapsto \varsigma_t$ on $\partial\mathbf{D}$ which is piecewise continuous with respect to the weak topology, the existence of a solution to the radial Loewner equation (g_t) driven by (ς_t) follows from standard existence results for ordinary differential equations (see, for example, [Law05b, Theorem 4.14]). The result in the general case follows by combining the previous assertion with Lemma 10.11. In particular, if $\varsigma \in \mathcal{N}$, then we let (ς_t^n) be a sequence as in Lemma 10.11. For each n, let (g_t^n) be the radial Loewner evolution driven by $t \mapsto \varsigma_t^n$. Then the previous assertion implies that (g_t^n) converges in $\mathcal{G}$ to the unique solution (g_t) driven by ς. □

To finish the proof of Theorem 10.1, we need to show that we can associate a growing family of hulls (K_t) in $\overline{\mathbf{D}}$ parameterized by log conformal radius with an element of $\mathcal{N}$ using the radial Loewner evolution and that the convergence of hulls with respect to the Caratheodory topology is equivalent to the convergence of measures in $\mathcal{N}$, also using radial Loewner evolution. This is accomplished in the following two propositions.

Proposition 10.13. *Suppose that (K_t) is a family of hulls in $\overline{\mathbf{D}}$ parameterized by* log *conformal radius as seen from* 0*. That is, the conformal radius of $D_t = \mathbf{D} \setminus K_t$ as seen from* 0 *is equal to e^{-t} for each $t \geq 0$. There exists a unique measure $\varsigma \in \mathcal{N}$ such that if (g_t) is the solution of the radial Loewner evolution driven by ς then, for each $t \geq 0$, K_t is the complement in $\overline{\mathbf{D}}$ of the domain of g_t.*

The main ingredient in the proof of Proposition 10.13 is the following lemma.

Lemma 10.14. *Suppose that $K \subseteq \overline{\mathbf{D}}$ is a compact hull and let $T = -\log \mathrm{C}R(0; \mathbf{D} \setminus K)$. Then there exists a measure $\varsigma \in \mathcal{N}_T$ such that if (g_t) is the radial Loewner evolution driven by ς then $\mathbf{D} \setminus K$ is the domain of g_T.*

Proof. Fix $\epsilon > 0$ and let $\gamma^\epsilon \colon [0, T_\epsilon] \to \overline{\mathbf{D}}$ be a simple curve starting from a point in $\partial\mathbf{D}$ such that the Hausdorff distance between K and $\gamma^\epsilon([0, T_\epsilon])$ is at most ϵ. Then (the radial version of) [Law05b, Proposition 4.4] implies that there exists a continuous function $U^\epsilon \colon [0, T_\epsilon] \to \partial\mathbf{D}$ such that if (g_t^ϵ) is the radial Loewner evolution driven by U^ϵ then, for each $t \in [0, T_\epsilon]$, $\gamma^\epsilon([0,t])$ is the complement in $\mathbf{D}$ of the domain of g_t^ϵ. Let $\varsigma_t^\epsilon = \delta_{U^\epsilon(t)}$. By possibly passing to a subsequence (ϵ_k) of positive numbers which decrease to 0 as $k \to \infty$, we have that $d\varsigma_t^\epsilon dt$ converges in $\mathcal{N}_T$ to $\varsigma \in \mathcal{N}_T$. Proposition 10.10 implies that the radial Loewner evolution (g_t) driven by ς has the property that the domain of g_T is $\mathbf{D} \setminus K$. □

Proof of Proposition 10.13. The uniqueness component of the proposition is obvious, so we will just give the proof of existence. Fix $\delta > 0$ and, for each $\ell \in \mathbf{N}_0$, let $K^{\delta,\ell} = g_{\delta\ell}(K_{\delta(\ell+1)})$. Let $\varsigma^{\delta,\ell}$ be a measure on $[\delta, \delta(\ell+1)] \times \partial\mathbf{D}$ as in Lemma 10.14 with respect to $K^{\delta,\ell}$, let $\varsigma^{\delta} = \sum_{\ell=0}^{\infty} \delta_{[\delta\ell,\delta(\ell+1))}(t)\varsigma^{\delta,\ell}$, and let (g_t^{δ}) be the radial Loewner evolution driven by ς^{δ}. Then the complement of the domain of $g_{\delta\ell}^{\delta}$ is equal to $K_{\delta\ell}$ for each $\ell \in \mathbf{N}_0$. The result follows by taking a limit along a sequence (δ_k) of positive numbers which decrease to 0 as $k \to \infty$ such that ς^{δ_k} converges in $\mathcal{N}$ to $\varsigma \in \mathcal{N}$. $\square$

Proposition 10.15. *Let (ς^n) be a sequence in $\mathcal{N}$. Suppose that, for each $n \in \mathbf{N}$ and $t \geq 0$, K_t^n is the complement in $\overline{\mathbf{D}}$ of the domain of g_t^n where $t \mapsto g_t^n$ is the radial Loewner evolution driven by ς^n. Then ς^n converges to an element ς of $\mathcal{N}$ if and only if (K_t^n) converges with respect to the Caratheodory topology to the growing sequence of compact hulls (K_t) in $\overline{\mathbf{D}}$ associated with the radial Loewner evolution driven by ς.*

Proof. That the convergence of $\varsigma^n \to \varsigma$ in $\mathcal{N}$ implies the Caratheodory convergence of the corresponding families of compact hulls is proved in Proposition 10.10. Therefore, we just have to prove the reverse implication. That is, we suppose that for each n, (K_t^n) is a family of compact hulls in $\mathbf{D}$ parameterized by log conformal radius as seen from 0 which converge in the Caratheodory sense to (K_t). For each $n \in \mathbf{N}$, let ς^n be the measure which drives the radial Loewner evolution associated with (K_t^n) and let ς be the measure which drives the radial Loewner evolution associated with (K_t). Let $\widetilde{\varsigma}$ be a subsequential limit in $\mathcal{N}$ of (ς^n). The Caratheodory convergence of (K_t^n) to (K_t) implies that $\widetilde{\varsigma}$ drives a radial Loewner evolution whose corresponding family of compact hulls is the same as (K_t), therefore $\varsigma = \widetilde{\varsigma}$. This implies that the limit of every convergent subsequence of (ς^n) is given by ς, hence $\varsigma^n \to \varsigma$ as $n \to \infty$ as desired. $\square$

Proof of Theorem 10.1. Combine Proposition 10.10, Proposition 10.13, and Proposition 10.15 as explained in the beginning of this section. $\square$

10.14.3 Approximations

We are now going to describe an approximation procedure for generating QLE(γ^2, η). Fix $\kappa > 1$. Let (h, u) have the law as described in Proposition 5.1 (where the role of γ in the application of the proposition is played by $2\alpha_\kappa$) plus $-\frac{\kappa+6}{2\sqrt{\kappa}} \log|\cdot|$. That is, u is sampled uniformly on $\partial\mathbf{D}$ from Lebesgue measure and, given u, the conditional law of h is that of a free boundary GFF on $\mathbf{D}$ plus $\frac{2}{\sqrt{\kappa}} \log|\cdot - u| - \frac{\kappa+6}{2\sqrt{\kappa}} \log|\cdot|$. Let $\nu_{h,\kappa}(\partial\mathbf{D})$ be the $-\frac{2}{\sqrt{\kappa}}$-quantum boundary length measure associated with h. Fix $\delta > 0$. We are now going to describe the dynamics of the triple $(\varsigma_t^{\delta}, g_t^{\delta}, \mathfrak{h}_t^{\delta})$ which will be an approximation to QLE(γ^2, η). The random variables $\varsigma_t^{\delta} dt$, (g_t^{δ}), and $(\mathfrak{h}_t^{\delta})$ will take values in $\mathcal{N}$, $\mathcal{G}$, and $\mathcal{H}$ respectively. The basic operation is illustrated in Figure 10.32. Consider the Markov chain in which we, starting with a field h on $\mathbf{D}$:

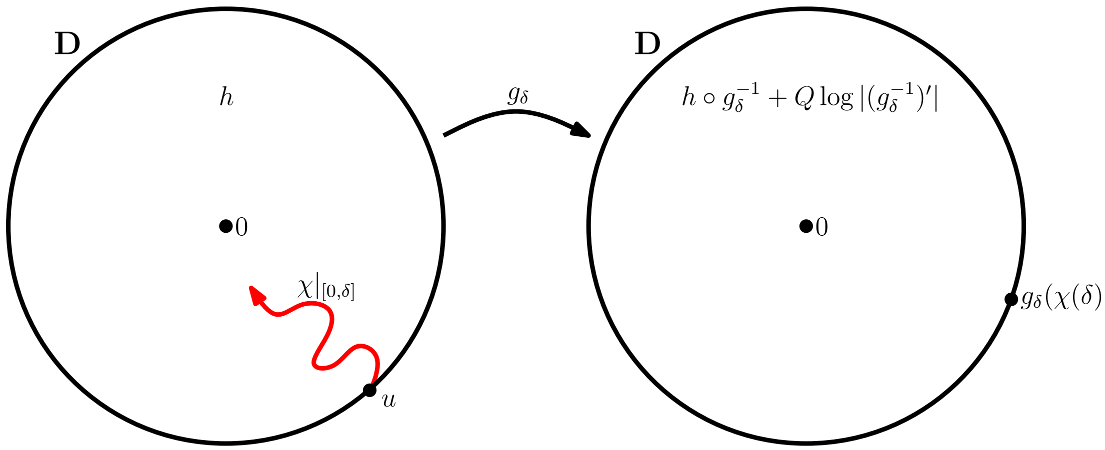

Figure 10.32: Fix $\kappa > 1$ and suppose that (h, u) has the law as described in Proposition 5.1 (where the role of γ in the application of the proposition is played by $2\alpha_\kappa$) plus $-\frac{\kappa+6}{2\sqrt{\kappa}} \log|\cdot|$ and let $\nu_{h,\kappa}$ be the $-\frac{2}{\sqrt{\kappa}}$-quantum boundary length measure associated with h. Then the conditional law of h given u is that of a free boundary GFF on $\mathbf{D}$ plus $-\frac{\kappa+6}{2\sqrt{\kappa}} \log|\cdot| + \frac{2}{\sqrt{\kappa}} \log|\cdot - u|$. By Theorem 7.12, the law of the pair (h, u) is invariant under the operation of sampling a radial SLE_κ process in $\mathbf{D}$ starting from u and targeted at 0 (which given u is conditionally independent of h) up to some fixed (log conformal radius) time δ, mapping back using the (forward) radial Loewner map g_δ as illustrated above, and applying the change of coordinates formula for quantum surfaces. Here, h is viewed as a modulo additive constant distribution. This is the basic operation which is used to construct QLE.

1. Pick $u \in \partial\mathbf{D}$ according to $\nu_{h,\kappa}$. By Proposition 5.1, the conditional law of h given u is equal to that of the sum of a free boundary GFF on $\mathbf{D}$ plus $-\frac{\kappa+6}{2\sqrt{\kappa}} \log|\cdot| + \frac{2}{\sqrt{\kappa}} \log|\cdot - u|$.
2. Sample a radial SLE_κ in $\mathbf{D}$ starting from u and targeted at 0 taken to be conditionally independent of h given u. Let (g_t) be the corresponding family of conformal maps which we assume to be parameterized by log conformal radius.
3. Replace h with $h \circ (g_\delta^{-1}) + Q \log|(g_\delta^{-1})'|$ where $Q = 2/\gamma + \gamma/2$ with $\gamma = \min(\sqrt{\kappa}, \sqrt{16/\kappa})$.

By Proposition 5.1 and Theorem 7.12, we know that this Markov chain preserves the law of h. (Recall from Remark 4.1 that g_δ^{-1} is equal in law to the reverse SLE_κ radial Loewner map run for δ units of time.) We use this construction to define the processes $(\varsigma_t^\delta, g_t^\delta, \mathfrak{h}_t^\delta)$ as follows.

- We sample $U^{\delta,0}$ from $\nu_{h,\kappa} = \exp(\alpha_\kappa h)$ and let $W^{\delta,0} = \exp(i\sqrt{\kappa}B^{\delta,0})$ where $B^{\delta,0}$ is a standard Brownian motion independent of h, and take $g^\delta|_{[0,\delta)}$ to be the radial Loewner evolution driven by $U^{\delta,0}W^{\delta,0}$.

- For each $t \in [0,\delta]$, we let $\mathfrak{h}^\delta|_{[0,\delta)}$ be given by[39] $P_{\rm harm}(h \circ (g_t^\delta)^{-1} + Q \log |((g_t^\delta)^{-1})'| + \frac{\kappa+6}{2\sqrt{\kappa}} \log|\cdot|)$ normalized so that $\mathfrak{h}_t^\delta(0) = 0$ for $t \in [0,\delta]$.
- Given that (g_t^δ) and $(\mathfrak{h}_t^\delta)$ have been defined for $t \in [0, \delta k)$, some $k \in \mathbf{N}$, we sample $U^{\delta,k}$ from $\exp(\alpha_\kappa \mathfrak{h}_{\delta k}^\delta)$ and let $W^{\delta,k} = \exp(i\sqrt{\kappa} B^{\delta,k})$ where $B^{\delta,k}$ is an independent standard Brownian motion defined in the time interval $[\delta k, \delta(k+1))$ (so that $B_{\delta k}^{\delta,k} = 0$).
- We then take $\widetilde{g}^\delta|_{[\delta k,\delta(k+1))}$ to be the radial Loewner evolution driven by $U^{\delta,k} W^{\delta,k}$ and $g^\delta|_{[\delta k,\delta(k+1))} = \widetilde{g}^\delta|_{[\delta k,\delta(k+1))} \circ g_{\delta k}^\delta$.
- Finally, we take $\mathfrak{h}^\delta|_{[\delta,\delta(k+1))}$ to be given by $P_{\rm harm}(h \circ (g_t^\delta)^{-1} + Q \log |((g_t^\delta)^{-1})'| + \frac{\kappa+6}{2\sqrt{\kappa}} \log|\cdot|)$ normalized so that $\mathfrak{h}_t^\delta(0) = 0$ for $t \in [\delta k, \delta(k+1))$.

Since $h \circ (g_t^\delta)^{-1} + Q \log |((g_t^\delta)^{-1})'| + \frac{\kappa+6}{2\sqrt{\kappa}} \log|\cdot|$ for each $t \geq 0$ has the law of a free boundary GFF on $\mathbf{D}$ plus $\frac{2}{\sqrt{\kappa}} \log|\cdot - u|$ where $u \in \partial\mathbf{D}$ is uniformly chosen from Lebesgue measure on $\partial\mathbf{D}$, the orthogonal projections used to define $\mathfrak{h}_t^\delta$ are almost surely defined for Lebesgue almost all $t \geq 0$ simultaneously; recall Proposition 3.2. We can extend the definition of $\mathfrak{h}_t^\delta$ so that it makes sense almost surely for all $t \geq 0$ simultaneously as follows. By induction, it is easy to see that the complement K_t^δ in $\overline{\mathbf{D}}$ of the domain of g_t^δ is a local set for (the GFF part of) h for each $t \geq 0$. Lemma 10.18 below (which is in some sense a strengthening of Proposition 3.6 stated above) implies that $\mathfrak{h}_t^\delta$ is almost surely continuous as a function $[0,\infty) \times \mathbf{D} \to \mathbf{R}$. Indeed, to see this we note that we have that (recall the definition of $\mathcal{C}_K$ from just before the statement of Proposition 3.6)

$$\mathfrak{h}_t^\delta = \mathcal{C}_{K_t} \circ (g_t^\delta)^{-1} + Q \log |(g_t^\delta)^{-1})'|$$

almost surely for each $t \geq 0$. Thus the claim follows since Lemma 10.18 implies that $\mathcal{C}_{K_t}$ is almost surely continuous in (t,z).

Let

$$\varsigma_t^\delta = \sum_{\ell=0}^\infty \mathbf{1}_{[\delta\ell,\delta(\ell+1))}(t) \delta_{U^{\delta,\ell} W^{\delta,\ell}}. \tag{10.21}$$

That is, ς_t^δ for $t \in [\delta\ell, \delta(\ell+1))$ and $\ell \in \mathbf{N}$ is given by the Dirac mass located at $U^{\delta,\ell} W^{\delta,\ell} \in \partial\mathbf{D}$. Then (g_t^δ) is the radial Loewner evolution driven by ς_t^δ. That is, (g_t^δ) solves

$$\dot{g}_t^\delta(z) = \int_{\partial\mathbf{D}} \Phi(u, g_t^\delta(z)) d\varsigma_t^\delta(u), \qquad g_0^\delta(z) = z. \tag{10.22}$$

(Recall (4.2).) We emphasize that by Theorem 7.12 we have

$$h \circ (g_t^\delta)^{-1} + Q \log |((g_t^\delta)^{-1})'| \stackrel{d}{=} h \quad \text{for all} \quad t \geq 0$$

[39]We add the term $\frac{\kappa+6}{2\sqrt{\kappa}} \log|\cdot|$ back into the GFF whenever applying $P_{\rm harm}$ because $P_{\rm harm}$ as defined in Remark 3.4 is defined only for the GFF.

as modulo additive constant distributions. In particular, the law of $(\mathfrak{h}_t^\delta)$ is stationary in t.

For our later arguments, it will be more convenient to consider the measure $d\varsigma_t^\delta dt$ on $[0,\infty) \times \partial\mathbf{D}$ in place of ς_t^δ, which for each $t \geq 0$ is a measure on $\partial\mathbf{D}$. Note that this is a random variable which takes values in $\mathcal{N}$. We also note that (g_t^δ) takes values in $\mathcal{G}$ and $(\mathfrak{h}_t^\delta)$ takes values in $\mathcal{H}$.

Definition 10.16. We call the triple $(\varsigma_t^\delta, g_t^\delta, \mathfrak{h}_t^\delta)$ constructed above the **δ-approximation** to $\mathrm{QLE}(\gamma^2,\eta)$.

We note that the operations that one uses to build the δ-approximation to QLE correspond to similar operations used to build the Hastings-Levitov growth model [HL98]. In particular, the SLE curves in QLE correspond to the slits in the Hastings-Levitov construction and the Liouville quantum gravity boundary measure in QLE plays the role of both the harmonic measure and the scaling factor (derivative of conformal map to a power) in Hastings-Levitov.

Note that the dynamics $(\varsigma_t^\delta, g_t^\delta, \mathfrak{h}_t^\delta)$ satisfy two of the arrows from Figure 10.2. Namely, g_t^δ is obtained from ς_t^δ by solving the radial Loewner equation and $\mathfrak{h}_t^\delta$ is obtained from g_t^δ using $P_{\mathrm{harm}}(h \circ (g_t^\delta)^{-1} + Q \log|((g_t^\delta)^{-1})'| + \frac{\kappa+6}{2\sqrt{\kappa}} \log|\cdot|)$ (and then normalized to vanish at the origin). However, ς_t^δ is *not* obtained from $\mathfrak{h}_t^\delta$ via exponentiation. (Rather, ς_t^δ is given by a Dirac mass at a point in $\partial\mathbf{D}$ which *is sampled from* the measure given by exponeniating $\mathfrak{h}_t^\delta$.) In Section 10.14.4, we will show that each of the elements of $(\varsigma_t^\delta, g_t^\delta, \mathfrak{h}_t^\delta)$ is tight (on compact time intervals) with respect to a suitable topology as $\delta \to 0$. In Section 10.15, we will show that both of the aforementioned arrows for the $\mathrm{QLE}(\gamma^2,\eta)$ dynamics still hold for the subsequentially limiting objects $(\varsigma_t, g_t, \mathfrak{h}_t)$. We will complete the proof by showing that ς_t is equal to the measure ν_t which is given by exponentiating $\mathfrak{h}_t$, hence the triple $(\nu_t, g_t, \mathfrak{h}_t)$ satisfies all three arrows of the $\mathrm{QLE}(\gamma^2,\eta)$ dynamics.

10.14.4 Tightness

The purpose of this section is to establish Proposition 10.19, which gives the existence of subsequential limits of the triple $(\varsigma_t^\delta, g_t^\delta, \mathfrak{h}_t^\delta)$ viewed as a random variable taking values in $\mathcal{N} \times \mathcal{G} \times \mathcal{H}$ as $\delta \to 0$. We begin with the following two lemmas which are general results about local sets for the GFF. Recall that $\mathcal{D}_a$ is defined at the end of Section 10.14.1.

Lemma 10.17. *Suppose that (h_n, K_n) is a sequence such that, for each n, h_n is a GFF on $\mathbf{D}$ (with Dirichlet, free, or mixed boundary conditions and the same boundary conditions for each n) and $K_n \subseteq \overline{\mathbf{D}}$ is a local set for h_n. Fix $a > 0$. Assume that (h_n, K_n) are coupled together so that $h_n \to h$ (resp. $K_n \to K$) almost surely as $n \to \infty$ in $\mathcal{D}_a$ (resp. the Hausdorff topology) where h is a GFF on $\mathbf{D}$ and $K \subseteq \overline{\mathbf{D}}$ is closed. Then K is local for h. If $\mathcal{C}_{K_n}$ (resp. $\mathcal{C}_K$) denotes the conditional expectation of h_n (resp.*

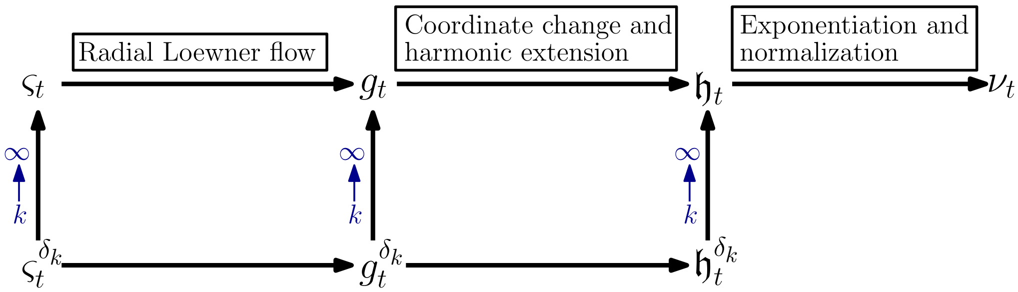

Figure 10.33: In Proposition 10.19 of Section 10.14.4, we prove the tightness of each of the elements of the triple $(\varsigma_t^\delta, g_t^\delta, \mathfrak{h}_t^\delta)$ (on compact time intervals) in the δ-approximation to $\mathrm{QLE}(\gamma^2,\eta)$ as $\delta \to 0$ with respect to the topologies introduced in Section 10.14.1. The subsequentially limiting objects $(\varsigma_t, g_t, \mathfrak{h}_t)$ are related to each other in the same way as $(\varsigma_t^\delta, g_t^\delta, \mathfrak{h}_t^\delta)$ and as indicated above. Namely, g_t is generated from ς_t by solving the radial Loewner equation and $\mathfrak{h}_t$ is related to g_t in that it is given by $P_{\mathrm{harm}}(h \circ g_t^{-1} + Q \log|(g_t^{-1})'| + \frac{\kappa+6}{2\sqrt{\kappa}} \log|\cdot|)$ (normalized to vanish at the origin). The measure ν_t is obtained from $\mathfrak{h}_t$ by exponentiation and is constructed in Proposition 10.20. Upon proving tightness, the existence of $\mathrm{QLE}(\gamma^2,\eta)$ is established by showing that $\nu_t = \varsigma_t$. This is completed in Section 10.15.

h) given K_n and $h|_{K_n}$ (resp. K and $h|_K$) and $\mathcal{C}_{K_n} \to F$ locally uniformly almost surely for some function $F\colon \mathbf{D} \setminus K \to \mathbf{R}$, then $F = \mathcal{C}_K$.

Proof. The proof that K is a local set for h is similar to that of [SS13, Lemma 4.6]. In particular, we will make use of the second characterization of local sets from Lemma 3.5. Fix a deterministic open set $B \subseteq \mathbf{D}$. For each $n \in \mathbf{N}$, we let S_n be the event that $K_n \cap B \neq \emptyset$ and let $\widetilde{K}_n = K_n$ on S_n^c and $\widetilde{K}_n = \emptyset$ otherwise. We also let S be the event that $K \cap B \neq \emptyset$ and let $\widetilde{K} = K$ on S^c and $\widetilde{K} = \emptyset$ otherwise. For each $n \in \mathbf{N}$, let $h_n^1 = P_{\mathrm{harm}}(h_n; B)$ and $h_n^2 = P_{\mathrm{supp}}(h_n; B)$ and define h^1, h^2 analogously for h. Since h^1 is independent of h^2 (recall Proposition 3.2), it suffices to show that h^2 is independent of the triple $(S, \widetilde{K}, h^1)$. Since K_n is local for h_n, the second characterization of local sets from Lemma 3.5 implies that h_n^2 is independent of the triple $(S_n, \widetilde{K}_n, h_n^1)$ for each $n \in \mathbf{N}$. The result therefore follows because this implies that the independence holds in the $n \to \infty$ limit.

Suppose that $\mathcal{C}_{K_n} \to F$ locally uniformly almost surely for some $F\colon \mathbf{D} \setminus K \to \mathbf{R}$. Then F is almost surely harmonic since each $\mathcal{C}_{K_n}$ is harmonic. Since K_n is local for h_n we can write $h_n = \widetilde{h}_n + \mathcal{C}_{K_n}$ where $\widetilde{h}_n$ is a zero-boundary GFF on $\mathbf{D} \setminus K_n$. Fix $\epsilon > 0$. Since $h_n \to h$ in $\mathcal{D}_a$ it follows that $h_n \to h$ in $(-\Delta)^a L^2((1-\epsilon)\mathbf{D} \setminus K)$ as $n \to \infty$. The local uniform convergence of $\mathcal{C}_{K_n}$ to F in $\mathbf{D} \setminus K$ as $n \to \infty$ implies that $\mathcal{C}_{K_n} \to F$ in $(-\Delta)^a L^2(V)$ for all $V \subseteq \overline{\mathbf{D}} \setminus K$ with $\mathrm{dist}(V, K \cup \partial \mathbf{D}) > 0$ as $n \to \infty$. Combining, we have that $\widetilde{h}_n$ converges to some $\widetilde{h}$ in $(-\Delta)^a L^2(V)$ as $n \to \infty$ for such V. Since this holds for all such V, we have that $h = \widetilde{h} + F$ and $\widetilde{h}$ is a zero-boundary GFF in $\mathbf{D} \setminus K$.

Since K is local for h, $\widehat{h} = h - \mathcal{C}_K = \widetilde{h} + F - \mathcal{C}_K$ is a zero-boundary GFF on $\mathbf{D} \setminus K$. Rearranging, we have that $\widetilde{h} - \widehat{h} = \mathcal{C}_K - F$. If ϕ is harmonic in $\mathbf{D} \setminus K$, then we have that $(\mathcal{C}_K - F, \phi)_\nabla = (\widetilde{h} - \widehat{h}, \phi)_\nabla = 0$. Since this holds for all such ϕ, we have that $\mathcal{C}_K - F = 0$, desired. $\square$

Proposition 3.6 gives that if (K_t) is an increasing family of local sets for a GFF h on $\mathbf{D}$ parameterized by log conformal radius as seen from a given point $z \in \mathbf{D}$, then $\mathcal{C}_{K_t}(z)$ evolves as a Brownian motion as t varies but z is fixed. This in particular implies that $\mathcal{C}_{K_t}(z)$ is continuous in t. We are now going to extend this to show that $\mathcal{C}_{K_t}(z)$ is continuous in both t and z.

Lemma 10.18. *Suppose that h is a GFF on $\mathbf{D}$ (with Dirichlet, free, or mixed boundary conditions) and that (K_t) is an increasing family of local sets for h parameterized so that $-\log \mathrm{C}R(0; \mathbf{D} \setminus K_t) = t$ for all $t \in [0,T]$ where $T > 0$ is fixed. Then the function $[0,T] \times B(0, \frac{1}{16}e^{-T}) \to \mathbf{R}$ given by $(t,z) \mapsto \mathcal{C}_{K_t}(z)$ has a modification which is Hölder continuous with any exponent strictly smaller than $1/2$. The Hölder norm of the modification depends only on T and the boundary data for h. (In the case that h has free boundary conditions, we fix the additive constant for h so that $\mathcal{C}_{K_0}(0) = 0$.)*

The reason that Lemma 10.18 is stated for $z \in B(0, \frac{1}{16}e^{-T})$ is that the Koebe one-quarter theorem implies that $B(0, \frac{1}{4}e^{-T}) \subseteq \mathbf{D} \setminus K_t$ for all $t \in [0,T]$. In particular, $B(0, \frac{1}{16}e^{-T})$ has positive distance from K_t for all $t \in [0,T]$. By applying Lemma 10.18 iteratively, we see that $\mathcal{C}_{K_t}(z)$ is in fact continuous for all z,t pairs such that z is contained in the component of $\mathbf{D} \setminus K_t$ containing the origin.

Proof of Lemma 10.18. We are going to prove the result using the Kolmogorov-Čentsov theorem. Fix $0 \leq s \leq t \leq T$ and $z, w \in B(0, \frac{1}{16}e^{-T})$. Since K_t is local for h, we can write $h = h_t + \mathcal{C}_{K_t}$ where h_t is a zero-boundary GFF on $\mathbf{D} \setminus K_t$. Re-arranging, we have that $\mathcal{C}_{K_t} = h - h_t$. Let h_ϵ (resp. $h_{\epsilon,t}$) be the circle average process for h (resp. h_t). Taking $\epsilon = \frac{1}{16}e^{-T}$ so that $z \in B(0, \frac{1}{16}e^{-T})$ implies $B(z, \frac{1}{16}e^{-T}) \subseteq B(0, \frac{1}{8}e^{-T})$ in what follows, we have that

$$\mathcal{C}_{K_t}(z) - \mathcal{C}_{K_t}(w) = \big(h_\epsilon(z) - h_\epsilon(w)\big) - \big(h_{\epsilon,t}(z) - h_{\epsilon,t}(w)\big).$$

The same argument as in the proof of [DS11a, Proposition 3.1] applied to both h_ϵ and $h_{\epsilon,t}$ implies that for each $p \geq 2$ there exists a constant $C > 0$ such that

$$\mathbf{E}\left[\big(\mathcal{C}_{K_t}(z) - \mathcal{C}_{K_t}(w)\big)^p\right] \leq C|z-w|^{p/2}. \tag{10.23}$$

Proposition 3.6 also implies that for each $p \geq 2$ there exists a constant $C > 0$ such that

$$\mathbf{E}\left[\big(\mathcal{C}_{K_t}(z) - \mathcal{C}_{K_s}(z)\big)^p\right] \leq C|t-s|^{p/2}. \tag{10.24}$$

Combining (10.23), (10.24) with the inequality $(a+b)^p \le 2^p(a^p+b^p)$ implies that for each $p \ge 2$ there exists a constant $C > 0$ such that

$$\mathbf{E}\left[\left(\mathcal{C}_{K_t}(z) - \mathcal{C}_{K_s}(w)\right)^p\right] \le C(|z-w|^{p/2} + |t-s|^{p/2}). \tag{10.25}$$

The desired result thus follows from the Kolmogorov-Čentsov theorem [RY99c, KS91b]. □

Proposition 10.19. *There exists a sequence (δ_k) in $(0,\infty)$ decreasing to 0 such that the following is true. There exists a coupling of the laws of h_k, $(\varsigma_t^{\delta_k})$, $(g_t^{\delta_k})$, and $(\mathfrak{h}_t^{\delta_k})$ as $k \in \mathbf{N}$ varies — where h_k denotes the GFF used to generate $(\varsigma_t^{\delta_k})$, $(\mathfrak{h}_t^{\delta_k})$, and $(g_t^{\delta_k})$ — and limiting processes $h \in \mathcal{D}_a$ (some $a > 0$), $\varsigma \in \mathcal{N}$, $(g_t) \in \mathcal{G}$, and $(\mathfrak{h}_t) \in \mathcal{H}$ such that h_k, $\varsigma_t^{\delta_k} dt$, $(g_t^{\delta_k})$, and $(\mathfrak{h}_t^{\delta_k})$ almost surely converge to h, ς, (g_t), and $(\mathfrak{h}_t)$ respectively, in $\mathcal{D}_a$, $\mathcal{N}$, $\mathcal{G}$, and $\mathcal{H}$. Moreover, (g_t) is the radial Loewner evolution generated by ς and $\mathfrak{h}_t$ for each $t \ge 0$ is almost surely given by $P_{\mathrm{harm}}(h \circ g_t^{-1} + Q \log|(g_t^{-1})'| + \frac{\kappa+6}{2\sqrt{\kappa}} \log|\cdot|)$ (normalized to vanish at the origin). Finally,*

$$h \circ g_t^{-1} + Q \log|(g_t^{-1})'| \stackrel{d}{=} h \quad \textit{for each} \quad t \ge 0 \tag{10.26}$$

as modulo additive constant distributions.

Proof. As explained in Section 10.14.1, the law of the free boundary GFF has separable support; see also [SS13, Lemma 4.2 and Lemma 4.3]. It is also explained in Section 10.14.1 that the same holds for the laws of $\varsigma_t^\delta dt$, (g_t^δ), and $(\mathfrak{h}_t^\delta)$ viewed as random variables taking values in $\mathcal{N}$, $\mathcal{G}$, and $\mathcal{H}$, respectively. The tightness of the law of h is obvious as is the tightness of the law of $\varsigma_t^\delta dt$. The tightness of the law of (g_t^δ) follows from Lemma 10.12 and the tightness of the law of $(\mathfrak{h}_t^\delta)$ follows from Lemma 10.18. This implies the existence of a sequence (δ_k) of positive real numbers along which the law $\mathcal{L}_\delta$ of $(h, \zeta_t^\delta dt, g_t^\delta, \mathfrak{h}_t^\delta)$ has a weak limit. The Skorohod representation theorem implies that we find a coupling $(h_k, \varsigma_t^{\delta_k} dt, g_t^{\delta_k}, \mathfrak{h}_t^{\delta_k})$ of the laws $\mathcal{L}_{\delta_k}$ such that $h_k \to h$, $\varsigma_t^{\delta_k} dt \to \varsigma$, $(g_t^{\delta_k}) \to (g_t)$, and $(\mathfrak{h}_t^{\delta_k}) \to (\mathfrak{h}_t)$ almost surely as $k \to \infty$ in the senses described in the statement of the proposition.

It is left to show that $(h, \varsigma, g_t, \mathfrak{h}_t)$ are related in the way described in the proposition statement. Theorem 10.1 implies that (g_t) is obtained from ς by solving the radial Loewner equation. Therefore we just need to show that

(i) $\mathfrak{h}_t$ can be obtained from g_t via coordinate change and applying P_{harm} and then

(ii) establish (10.26).

We will start with (i). For each $t \ge 0$, let K_t be the hull given by the complement in $\overline{\mathbf{D}}$ of the domain of g_t. The first step is to show that K_t is local for h. Let $(K_t^{\delta_k})$ denote the corresponding family of hulls associated with $(g_t^{\delta_k})$. By possibly passing to a subsequence of (δ_k) and using that the Hausdorff topology is compact hence separable, we can recouple so that, in addition to the above, we have that $K_t^{\delta_k} \to \widetilde{K}_t$ almost surely in the Hausdorff topology for all $t \in \mathbf{Q}_+$. Lemma 10.17 then implies that $\widetilde{K}_t$ is local for h for all $t \in \mathbf{Q}_+$. Combining this with the first characterization of local sets given in Lemma 3.5 implies that K_t is local for h for all $t \geq 0$. Lemma 10.17 implies that $\mathcal{C}_{K_t^{\delta_k}} \to \mathcal{C}_{K_t}$ locally uniformly almost surely for all $t \in \mathbf{Q}_+$. Lemma 10.18 implies that the family $(\mathcal{C}_{K_t^{\delta_k}})$ is equicontinuous in time and space. Lemma 10.17 implies that any locally uniform subsequential limit of $\mathcal{C}_{K_t}^{\delta_k}$ must in fact be equal to $\mathcal{C}_{K_t}$. Combining, we conclude that $\mathcal{C}_{K_t^{\delta_k}} \to \mathcal{C}_{K_t}$ locally uniformly in both space and time. Note that $\mathfrak{h}_t^{\delta_k}$ is given by $\mathcal{C}_{K_t^{\delta_k}} \circ (g_t^{\delta_k})^{-1} + Q \log |((g_t^{\delta_k})^{-1})'| + \frac{\kappa+6}{2\sqrt{\kappa}} \log |\cdot|$ (normalized to vanish at the origin). Thus since $\mathcal{C}_{K_t^{\delta_k}} \to \mathcal{C}_{K_t}$ locally uniformly and $(g_t^{\delta_k})^{-1} \to g_t^{-1}$ locally uniformly, we therefore have that $\mathfrak{h}_t$ is given by $\mathcal{C}_{K_t} \circ g_t^{-1} + Q \log |(g_t^{-1})'| + \frac{\kappa+6}{2\sqrt{\kappa}} \log |\cdot|$ (normalized to vanish at the origin).

The construction of the δ-approximation implies that

$$h_k \circ (g_t^{\delta_k})^{-1} + Q \log |((g_t^{\delta_k})^{-1})'| \stackrel{d}{=} h_k \quad \text{for each} \quad k \in \mathbf{N} \quad \text{and} \quad t \geq 0$$

as modulo additive constant distributions, hence the same holds in the limit as $k \to \infty$ due to the nature of the convergence. This gives (ii). $\square$

10.15 Subsequential limits solve the QLE dynamics

Throughout this section, we suppose that (δ_k) is a sequence in $(0, \infty)$ decreasing to 0 as in the statement of Proposition 10.19 and $(h_k, \varsigma_t^{\delta_k}, g_t^{\delta_k}, \mathfrak{h}_t^{\delta_k})$ are coupled together on a common probability space such that $h_k \to h$ in $\mathcal{D}_a$ for $a > 0$, $\varsigma_t^{\delta_k} dt \to \varsigma$ in $\mathcal{N}$, $(g_t^{\delta_k}) \to (g_t)$ in $\mathcal{G}$, and $(\mathfrak{h}_t^{\delta_k}) \to (\mathfrak{h}_t)$ in $\mathcal{H}$ as in the statement of Proposition 10.19. The purpose of this section is to construct a family of probability measures (ν_t) on $\partial \mathbf{D}$ from $(\mathfrak{h}_t)$ and then show that the triple $(\nu_t, g_t, \mathfrak{h}_t)$ satisfies the QLE dynamics illustrated in Figure 10.2. The measures ν_t will only be defined for almost all $t \geq 0$, so we will in fact think of (ν_t) as being given by a single measure $\nu \in \mathcal{N}$. (That is, the measure ν_t is given by the conditional measure associated with ν when a value of t has been fixed. This gives us a definition of ν_t for almost every $t \geq 0$. We will not address the continuity of ν_t in t, so it is not *a priori* possible to extend the definition of ν_t to all t values simultaneously.)

We will accomplish the above in two steps. We will first construct a measure $\nu \in \mathcal{N}$ which, for a given time $t \geq 0$, should be thought of as the $-\frac{2}{\sqrt{\kappa}}$-quantum boundary length measure (Proposition 10.20) generated from the boundary values of $\mathfrak{h}_t$ (normalized to be a probability). That is, $\nu \in \mathcal{N}$ is formally given by $\mathcal{Z}_t^{-1} \exp(\alpha_\kappa \mathfrak{h}_t(u)) du dt$ where $\mathcal{Z}_t$ is a normalization constant. This step is carried out in Section 10.15.1. The second step (Proposition 10.21) is to show that $\varsigma = \nu$. This is carried out in Section 10.15.2. As we explained earlier, this will complete the proof because it gives that $(\nu_t, g_t, \mathfrak{h}_t)$ satisfies all three arrows of the QLE(γ^2, η) dynamics described in Figure 10.2.

10.15.1 Construction of the QLE driving measure

We begin by defining the approximations we will use to construct ν. We first approximate $\mathfrak{h}_t$ by orthogonally projecting it to the subspace of $H(\mathbf{D})$ (recall the definition of $H(\mathbf{D})$ from Section 3.2.9) spanned by $\{f_1, \ldots, f_n\}$ where (f_n) is an orthonormal basis of the subspace of functions of $H(\mathbf{D})$ which are harmonic in $\mathbf{D}$. In what follows in this section, the precise choice of basis is not important (i.e., the resulting measure ν does not depend on the choice of basis). However, for our later arguments, it will be convenient to make a particular choice so that it is obvious that our approximations are continuous in t. Thus for each $n \in \mathbf{N}$ which is even (resp. odd) we take $f_n(z) = \beta_n^{-1} \mathrm{Re}(z^{n/2})$ (resp. $f_n = \beta_n^{-1} \mathrm{Im}(z^{(n+1)/2})$) where $\beta_n = \|\mathrm{Re}(z^{n/2})\|_\nabla$ (resp. $\beta_n = \|\mathrm{Im}(z^{(n+1)/2})\|_\nabla$) so that $\|f_n\|_\nabla = 1$. Indeed, an elementary calculation implies that (f_n) is orthonormal in $H(\mathbf{D})$ and one can see that (f_n) spans the subspace of harmonic functions in $H(\mathbf{D})$ by recalling that every harmonic function in $\mathbf{D}$ is the real part of an analytic function on $\mathbf{D}$. Note that (f_n) is part of an orthonormal basis of all of $H(\mathbf{D})$; we will use this in conjunction with (5.3) in what follows. For each $n \in \mathbf{N}$ and $t \geq 0$, we let $\mathfrak{h}_t^n$ be the orthogonal projection of $\mathfrak{h}_t$ onto the subspace of $H(\mathbf{D})$ spanned by $\{f_1, \ldots, f_n\}$, i.e. the real parts of polynomials in z of degree at most $n/2$. We let

$$d\nu_t^n(u) = \frac{1}{\mathcal{Z}_{n,t}} \exp(\alpha_\kappa \mathfrak{h}_t^n(u)) du \quad \text{for} \quad u \in \partial\mathbf{D} \quad \text{and} \quad t \geq 0 \tag{10.27}$$

where $\mathcal{Z}_{n,t}$ is a normalizing constant so that ν_t^n has unit mass. Note that $\mathfrak{h}_t^n$ varies continuously in t with respect to the uniform topology on continuous functions defined on $\overline{\mathbf{D}}$. One way to see this is to note that since $\mathfrak{h}_t$ is harmonic in $\mathbf{D}$ for each fixed t, it is equal to the real part of an analytic function F_t on $\mathbf{D}$. Then $\mathfrak{h}_t^n$ is given by the real part of the terms up to degree $n/2$ in the power series expansion for F_t. The claimed continuity follows because these coefficients for F_t are a continuous function of $\mathfrak{h}_t$ restricted to $\frac{1}{2}\mathbf{D}$ with respect to the uniform topology on continuous functions on $\frac{1}{2}\overline{\mathbf{D}} \to \mathbf{R}$. We also let

$$d\nu^n(t,u) = d\nu_t^n(u) dt \quad \text{for} \quad u \in \partial\mathbf{D} \quad \text{and} \quad t \geq 0. \tag{10.28}$$

Then $\nu^n \in \mathcal{N}$ for all $n \in \mathbf{N}$.

Proposition 10.20. *There exists a sequence (n_j) in $\mathbf{N}$ with $n_j \to \infty$ as $j \to \infty$ and a measure $\nu \in \mathcal{N}$ such that $\nu^{n_j} \to \nu$ in $\mathcal{N}$ almost surely. That is, we almost surely have for each $T \geq 0$ and continuous function $\phi \colon [0,T] \times \partial\mathbf{D} \to \mathbf{R}$ that*

$$\lim_{j\to\infty} \int_{[0,T]\times\partial\mathbf{D}} \phi(s,u) d\nu^{n_j}(s,u) = \int_{[0,T]\times\partial\mathbf{D}} \phi(s,u) d\nu(s,u).$$

In the proof that follows and throughout the rest of this section, for measures ς_1, ς_2, we will use the notation $d(\varsigma_1 - \varsigma_2)$ to denote integration against the signed measure $\varsigma_1 - \varsigma_2$.

Proof of Proposition 10.20. Fix $n, n' \in \mathbf{N}$, $T \geq 0$, and a continuous function $\phi \colon [0,T] \times \partial \mathbf{D} \to \mathbf{R}$. By Fubini's theorem, we have that

$$\begin{aligned} &\mathbf{E}\left[\left| \int_{[0,T]\times\partial\mathbf{D}} \phi(s,u) d\big(\nu^n(s,u) - \nu^{n'}(s,u)\big) \right| \right] \\ &\leq \int_0^T \mathbf{E}\left[\left| \int_{\partial\mathbf{D}} \phi(s,u) d\big(\nu_s^n(u) - \nu_s^{n'}(u)\big) \right| \right] ds. \end{aligned} \tag{10.29}$$

Note that the integral inside of the expectation converges to zero as $n, n' \to \infty$ for any fixed $s \geq 0$ because ν_s^n converges weakly almost surely as $n \to \infty$ to the $-\frac{2}{\sqrt{\kappa}}$-quantum boundary measure on $\partial\mathbf{D}$ associated with $\mathfrak{h}_s$ normalized to be a probability measure (recall (5.3)) and the quantity inside of the expectation is bounded by $2\|\phi\|_{L^\infty}$. Therefore it follows from the dominated convergence theorem that the expression in (10.29) converges to zero as $n, n' \to \infty$. Applying Markov's inequality and the Borel-Cantelli lemma gives the almost sure convergence of $\int_{[0,T]\times\partial\mathbf{D}} \phi(s,u) d\nu^n(s,u)$ provided we take a limit along a sequence (n_j) in $\mathbf{N}$ which tends to ∞ sufficiently quickly. By possibly passing to a further (diagonal) subsequence, this, in turn, gives us the almost sure convergence of $\int_{[0,T]\times\partial\mathbf{D}} \phi(s,u) d\nu^{n_j}(s,u)$ for any countable collection of continuous functions $\phi \colon [0,T] \times \partial\mathbf{D} \to \mathbf{R}$. This proves the result because we can pick a countable dense subset of continuous functions $\phi \colon [0,T] \times \partial\mathbf{D} \to \mathbf{R}$ with respect to the uniform topology on $[0,T] \times \partial\mathbf{D}$ and then use the continuity of the aforementioned integral with respect to the uniform topology on continuous functions. Passing to a final (diagonal) subsequence gives the convergence for all $T \geq 0$ simultaneously. $\square$

10.15.2 Loewner evolution driven by the QLE driving measure solves the QLE dynamics

Throughout, we let ν be the (random) element of $\mathcal{N}$ constructed in Proposition 10.20 and we let ς be the (random) element of $\mathcal{N}$ which drives (g_t). As explained in Proposition 10.19, we know that we can obtain $\mathfrak{h}_t$ from g_t by $P_{\text{harm}}(h \circ g_t^{-1} + Q \log |(g_t^{-1})'| + \frac{\kappa+6}{2\sqrt{\kappa}} \log|\cdot|)$ (normalized to vanish at the origin) and that we can obtain ν by exponentiating $\mathfrak{h}_t$. Therefore the proof of Theorem 10.3 will be complete upon establishing the following.

Proposition 10.21. *We almost surely have that $\varsigma = \nu$.*

A schematic illustration of the main steps in the proof of Proposition 10.21 is given in Figure 10.34. The strategy is to relate ς and ν using three approximating measures: $\varsigma_t^{\delta_k} dt$, $\varsigma_t^{\delta_k, n_j} dt$, and $\nu_t^{\delta_k, n_j} dt$. We introduced $\varsigma_t^{\delta_k}$ in (10.21) and we introduced $\nu_t^{n_j}$ in (10.27).

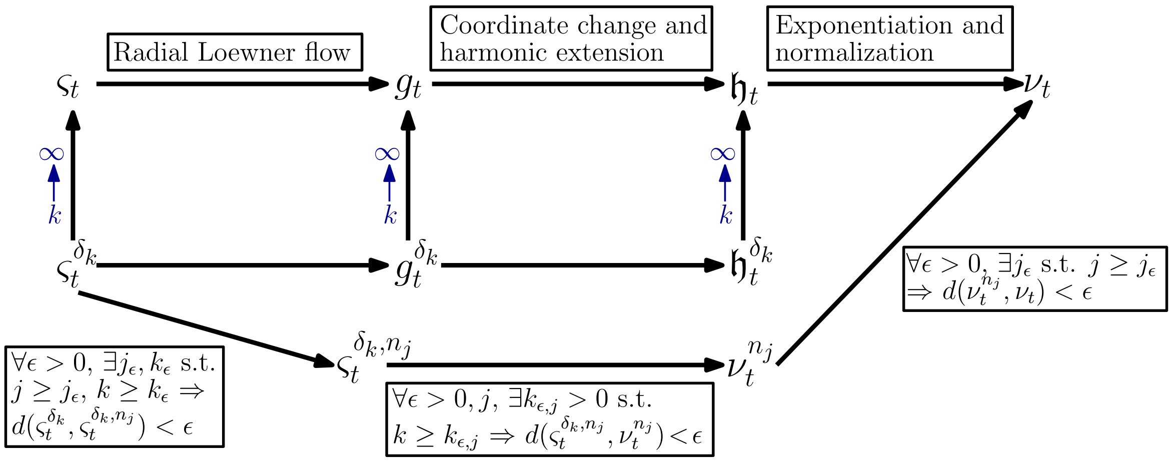

Figure 10.34: (Continuation of Figure 10.33.) Shown is the approximation scheme used to show that $\varsigma_t = \nu_t$ (Proposition 10.21) to complete the proof of the existence of QLE(γ^2, η) for (γ^2, η) on one of the upper two curves from Figure 10.3. The statements in each of the three boxes along the bottom of the figure from left to right are proved in Lemma 10.22, Lemma 10.23, and Proposition 10.20, respectively. The symbol d represents a notion of "closeness" which is related to the topology of $\mathcal{N}$. To show that $\varsigma_t = \nu_t$, we first pick j very large so that $d(\nu_t^{n_j}, \nu_t) < \epsilon$ and $d(\varsigma_t^{\delta_k}, \varsigma_t^{\delta_k,n_j}) < \epsilon$. We then pick k very large so that $d(\varsigma_t^{\delta_k}, \varsigma_t) < \epsilon$ and $d(\varsigma_t^{\delta_k,n_j}, \nu_t^{n_j}) < \epsilon$.

We know that $\varsigma_t^{\delta_k} dt \to \varsigma$ as $k \to \infty$ and $\nu_t^{n_j} dt \to \nu$ as $j \to \infty$ in $\mathcal{N}$. In the rest of this section, we will introduce $\varsigma_t^{\delta_k,n_j} dt$ and then show that $\varsigma_t^{\delta_k,n_j} dt$ is close to both $\varsigma_t^{\delta_k} dt$ and $\nu_t^{n_j} dt$ for large j and k.

We now give the definition of $\varsigma_t^{\delta,n}$. Fix $n \in \mathbf{N}$ and let

$$\nu_t^{\delta,n} = \mathcal{Z}_{n,t,\delta}^{-1} \exp(\alpha_\kappa \mathfrak{h}_t^{\delta,n}(u)) du \tag{10.30}$$

where $\mathfrak{h}_t^{\delta,n}$ is the orthogonal projection of $\mathfrak{h}_t^\delta$ onto the subspace spanned by $\{f_1, \ldots, f_n\}$ as defined above and $\mathcal{Z}_{n,t,\delta}$ is a normalization constant so that $\nu_t^{\delta,n}$ is a probability measure. For each $\ell \in \mathbf{N}_0$, let $U^{\delta,\ell,n}$ be a point picked from $\nu_t^{\delta,n}$ with $t = \delta\ell$. Fix $\zeta > 0$. For each $t \geq 0$, it follows from (5.3) that

$$\nu_t^{\delta,n} \to \nu_t^\delta \quad \text{as} \quad n \to \infty$$

weakly almost surely. By the stationarity of $\mathfrak{h}_t^\delta$, the rate of convergence is independent of t and δ. Recall the definition of the sequence $(U^{\delta,\ell})$ from the construction of the δ-approximation of QLE given in Section 10.14.3. It thus follows that there exists non-random $n_0 \in \mathbf{N}$ depending only on ζ such that for all $n \geq n_0$ we can couple the sequences $(U^{\delta,\ell})$ and $(U^{\delta,\ell,n})$ together so that

$$\mathbf{P}[E_\ell^{\delta,n}] \leq \zeta \quad \text{where} \quad E_\ell^{\delta,n} = \{|U^{\delta,\ell} - U^{\delta,\ell,n}| \geq \zeta\}. \tag{10.31}$$

We assume throughout that $U^{\delta,\ell,n}$ and $U^{\delta,n}$ are coupled as such. Let

$$\varsigma_t^{\delta,n} = \sum_{\ell=0}^{\infty} \mathbf{1}_{[\delta\ell,\delta(\ell+1))}(t)\delta_{U^{\delta,\ell,n}}.$$

That is, $\varsigma_t^{\delta,n}$ for $t \in [\delta\ell, \delta(\ell+1))$ with $\ell \in \mathbf{N}_0$ is given by the Dirac mass located at $U^{\delta,\ell,n}$. Note that $\varsigma_t^{\delta,n}$ is defined analogously to ς_t^{δ} except the $U^{\delta,\ell,n}$ are picked from $\nu_t^{\delta,n}$ in place of $\nu_t^{\delta} = \mathcal{Z}_{t,\delta}^{-1} \exp(\alpha_\kappa \mathfrak{h}_t^{\delta}(u))du$ and the Brownian motions have been omitted.

The proof of Proposition 10.21 has two steps.

The first step (Lemma 10.22) is to show that for each $\epsilon > 0$ there exists $j_\epsilon, k_\epsilon \in \mathbf{N}$ such that $\varsigma_t^{\delta_k} dt$ and $\varsigma_t^{\delta_k,n_j} dt$ are ϵ-close for all $j \geq j_\epsilon$ and $k \geq k_\epsilon$ (the result is stated for more general values of δ and n because it is not necessary in the proof to work along the sequences (δ_k) and (n_j)). We note that the choice of k determines the speed at which the location of the Dirac mass is resampled while the choice of j determines the expected fraction of the (U^{δ_k,ℓ,n_j}) which are close to the $(U^{\delta_k,\ell})$ (recall (10.31)).

The second step (Lemma 10.23) is to show that for each $\epsilon > 0$ and $j \in \mathbf{N}$ there exists $k_{\epsilon,j} > 0$ such that $\varsigma_t^{\delta_k,n_j} dt$ and $\nu_t^{n_j} dt$ are ϵ-close for all $k \geq k_{\epsilon,j}$ (the result is stated for more general values of n because in the proof it is not necessary to work along the sequence (n_j)). The proof is by a law of large numbers argument. By construction, we know that $t \mapsto \nu_t^{\delta_k,n_j}$ is continuous for a *fixed* value of j and the choice of j controls our estimate its modulus of continuity. When $\delta_k > 0$ is sufficiently small (depending on j), we can think of organizing the points (U^{δ_k,ℓ,n_j}) into groups, each of which is close to being i.i.d. (This follows because the rate at which the points (U^{δ_k,ℓ,n_j}) are being sampled is much faster than the rate at which $t \mapsto \nu_t^{\delta_k,n_j}$ is changing.) This is what leads to the law of large numbers effect.

Once these estimates have been established, we will pick j very large so that both $\nu_t^{n_j} dt$ is close to ν and $\varsigma_t^{\delta_k,n_j} dt$ is close to $\varsigma_t^{\delta_k} dt$. We will then choose k to be very large so that $\varsigma_t^{\delta_k} dt$ is close to ς and $\varsigma_t^{\delta_k,n_j} dt$ is close to $\nu_t^{n_j} dt$.

Lemma 10.22. *Fix $T > 0$ and suppose that $\phi\colon [0,T] \times \partial\mathbf{D} \to \mathbf{R}$ is continuous. For every $\epsilon > 0$ there exists $n_\epsilon \in \mathbf{N}$ and $\delta_\epsilon > 0$ such that $n \geq n_\epsilon$ and $\delta \in (0, \delta_\epsilon)$ implies that*

$$\mathbf{E}\left| \int_{[0,T]\times\partial\mathbf{D}} \phi(s,u) d\big(\varsigma_s^{\delta}(u) - \varsigma_s^{\delta,n}(u)\big) ds \right| < \epsilon.$$

Proof. We are first going to explain how to bound the difference when $s = \ell\delta$ for some $\ell \in \mathbf{N}_0$. Fix $\epsilon > 0$. By the continuity of ϕ, it follows from (10.31) that there exists $n_{1,\epsilon}$ such that $n \geq n_{1,\epsilon}$ implies that

$$\mathbf{E}\left| \int_{\partial\mathbf{D}} \mathbf{1}_{(E_\ell^{\delta,n})^c} \phi(\delta\ell, u) d\big(\varsigma_{\ell\delta}^{\delta}(u) - \varsigma_{\ell\delta}^{\delta,n}(u)\big) \right| < \frac{\epsilon}{4} \tag{10.32}$$

provided we choose $\zeta > 0$ small enough. Since the integrand is bounded, it also follows from (10.31) that there exists $n_{2,\epsilon}$ such that $n \geq n_{2,\epsilon}$ implies that

$$\mathbf{E}\left|\int_{\partial\mathbf{D}} \mathbf{1}_{E_\ell^{\delta,n}} \phi(\delta\ell, u) d\big(\varsigma_{\ell\delta}^{\delta}(u) - \varsigma_{\ell\delta}^{\delta,n}(u)\big)\right| < \frac{\epsilon}{4}. \tag{10.33}$$

Combining (10.32) and (10.33) gives that $n \geq n_\epsilon = \max(n_{1,\epsilon}, n_{2,\epsilon})$ implies that

$$\mathbf{E}\left|\int_{\partial\mathbf{D}} \phi(\delta\ell, u) d\big(\varsigma_{\ell\delta}^{\delta}(u) - \varsigma_{\ell\delta}^{\delta,n}(u)\big)\right| < \frac{\epsilon}{2}.$$

Using the continuity of Brownian motion, it follows that there exists $\delta_\epsilon > 0$ such that for all $n \geq n_\epsilon$ and $\delta \in (0, \delta_\epsilon)$ we have that

$$\sup_{s \in [\delta\ell, \delta(\ell+1))} \mathbf{E}\left|\int_{\partial\mathbf{D}} \phi(s, u) d\big(\varsigma_s^{\delta}(u) - \varsigma_s^{\delta,n}(u)\big)\right| < \epsilon.$$

This implies the desired result. □

Lemma 10.23. *Fix $T > 0$ and suppose that $\phi\colon [0,T] \times \partial\mathbf{D} \to \mathbf{R}$ is continuous. For each $n \in \mathbf{N}$ there exists $k_{\epsilon,n} \in \mathbf{N}$ such that $k \geq k_{\epsilon,n}$ implies that*

$$\mathbf{E}\left|\int_{[0,T]\times\partial\mathbf{D}} \phi(s, u) d\big(\varsigma_s^{\delta_k,n}(u) - \nu_s^n(u)\big) ds\right| < \epsilon. \tag{10.34}$$

It is important that the limit in the statement of Lemma 10.23 is along the sequence (δ_k) because then we have that $\mathfrak{h}_t^{\delta_k} \to \mathfrak{h}_t$ as $k \to \infty$ and ν_t^n is defined in terms of $\mathfrak{h}_t$.

Proof of Lemma 10.23. Let (recall (10.30))

$$\overline{\varsigma}_t^{\delta_k,n}(u) du = \sum_{\ell=0}^{\infty} \mathbf{1}_{[\ell\delta_k,(\ell+1)\delta_k)}(t) \nu_{\ell\delta_k}^{\delta_k,n}(u) du.$$

Note that the increments of

$$\int_0^{\delta_k \ell} \int_{\partial\mathbf{D}} \phi(s, u) d(\varsigma_s^{\delta_k,n}(u) - \overline{\varsigma}_s^{\delta_k,n}(u)) ds$$

as ℓ varies are uncorrelated given $\mathfrak{h}_t^{\delta_k,n}$. Consequently, we have that

$$\mathbf{E}\left[\left(\int_{[0,T]\times\partial\mathbf{D}} \phi(s, u) d(\varsigma_s^{\delta_k,n}(u) - \overline{\varsigma}_s^{\delta_k,n}(u)) ds\right)^2\right] = O(\delta_k)$$

where the implicit constant in $O(\delta_k)$ depends on T. It thus suffices to prove (10.34) with $\overline{\varsigma}_t^{\delta_k,n}$ in place of $\varsigma_t^{\delta_k,n}$. By the continuity of $\mathfrak{h}_t$ and the local uniform convergence of $\mathfrak{h}_t^{\delta_k}$ to $\mathfrak{h}_t$ as $k \to \infty$, it is easy to see that

$$\lim_{k\to\infty} \mathbf{E}\left|\int_{[0,T]\times\partial\mathbf{D}} \phi(s, u) d(\overline{\varsigma}_s^{\delta_k,n}(u) - \nu_s^{\delta_k,n}(u)) ds\right| = 0.$$

Combining gives (10.34). □

Proof of Proposition 10.21. Fix $T > 0$ and $\phi\colon [0,T] \times \partial\mathbf{D} \to \mathbf{R}$ continuous. It suffices to show that

$$\int_{[0,T]\times\partial\mathbf{D}} \phi(s,u)d(\nu(s,u) - \varsigma(s,u)) = 0 \tag{10.35}$$

almost surely. Fix $\epsilon > 0$. Then Proposition 10.19 implies that there exists $k_\epsilon \in \mathbf{N}$ such that $k \geq k_\epsilon$ implies that

$$\mathbf{E}\left|\int_{[0,T]\times\partial\mathbf{D}} \phi(s,u)\big(d\varsigma_s^{\delta_k}(u)ds - d\varsigma(s,u)\big)\right| < \frac{\epsilon}{4}. \tag{10.36}$$

Lemma 10.22 implies that there exists $j_\epsilon \in \mathbf{N}$ such that, by possibly increasing the value of k_ϵ, we have that $j \geq j_\epsilon$ and $k \geq k_\epsilon$ implies that

$$\mathbf{E}\left|\int_{[0,T]\times\partial\mathbf{D}} \phi(s,u)d\big(\varsigma_s^{\delta_k}(u) - \varsigma_s^{\delta_k,n_j}(u)\big)ds\right| < \frac{\epsilon}{4}. \tag{10.37}$$

Proposition 10.20 implies that, by possibly increasing the value of j_ϵ, we have that $j \geq j_\epsilon$ implies that

$$\mathbf{E}\left|\int_{[0,T]\times\partial\mathbf{D}} \phi(s,u)d\big(\nu^{n_j}(s,u) - \nu(s,u)\big)\right| < \frac{\epsilon}{4}. \tag{10.38}$$

Let $j = j_\epsilon$. Lemma 10.23 implies that there exists $k_{\epsilon,j} \in \mathbf{N}$ such that $k \geq k_{\epsilon,j}$ implies that

$$\mathbf{E}\left|\int_{[0,T]\times\partial\mathbf{D}} \phi(s,u)d\big(\varsigma_s^{\delta_k,n_j}(u) - \nu_s^{n_j}(u)\big)ds\right| < \frac{\epsilon}{4}. \tag{10.39}$$

Using the triangle inequality, (10.36)—(10.39), and that $\epsilon > 0$ was arbitrary implies (10.35), as desired. □

We can now complete the proof of Theorem 10.3.

Proof of Theorem 10.3. Proposition 10.21 gives us that the limiting triple $(\nu_t, g_t, \mathfrak{h}_t)$ satisfies all three arrows of the QLE dynamics as described in Figure 10.2. That the limiting triple $(\nu_t, g_t, \mathfrak{h}_t)$ satisfies η-DBM scaling as defined in Definition 10.2 follows from the argument explained in Section 10.12. Combining gives the desired result. □

References

[AC13] O. Angel and N. Curien. Percolations on random maps I: half-plane models. *ArXiv e-prints*, January 2013.

[AG13a] Amine Asselah and Alexandre Gaudillière. From logarithmic to subdiffusive polynomial fluctuations for internal DLA and related growth models. *Ann. Probab.*, 41(3A):1115–1159, 2013. 1009.2838.

[AG13b] Amine Asselah and Alexandre Gaudillière. Sublogarithmic fluctuations for internal DLA. *Ann. Probab.*, 41(3A):1160–1179, 2013. 1011.4592.

[AJKS09] Kari Astala, Peter Jones, Antti Kupiainen, and Eero Saksman. Random Conformal Weldings. *To appear in Acta Math. ArXiv e-prints*, September 2009.

[AJKS10] Kari Astala, Peter Jones, Antti Kupiainen, and Eero Saksman. Random curves by conformal welding. *C. R. Math. Acad. Sci. Paris*, 348(5-6):257–262, 2010.

[Ald91a] David Aldous. The continuum random tree. I. *Ann. Probab.*, 19(1):1–28, 1991.

[Ald91b] David Aldous. The continuum random tree. II. An overview. In *Stochastic analysis (Durham, 1990)*, volume 167 of *London Math. Soc. Lecture Note Ser.*, pages 23–70. Cambridge Univ. Press, Cambridge, 1991.

[Ald93] David Aldous. The continuum random tree. III. *Ann. Probab.*, 21(1):248–289, 1993.

[Ang03] O. Angel. Growth and percolation on the uniform infinite planar triangulation. *Geom. Funct. Anal.*, 13(5):935–974, 2003. math/0208123.

[AOF11] Sidiney G Alves, Tiago J Oliveira, and Silvio C Ferreira. Universal fluctuations in radial growth models belonging to the KPZ universality class. *EPL (Europhysics Letters)*, 96(4):48003, 2011. 1109.4901.

[AR13] O. Angel and G. Ray. Classification of Half Planar Maps. *ArXiv e-prints*, March 2013.

[Aru13] J. Aru. KPZ relation does not hold for the level lines and the SLE_κ flow lines of the Gaussian free field. *ArXiv e-prints*, December 2013.

[AS03] Omer Angel and Oded Schramm. Uniform infinite planar triangulations. *Comm. Math. Phys.*, 241(2-3):191–213, 2003. math/0207153.

[BAD96] G. Ben Arous and J.-D. Deuschel. The construction of the $(d+1)$-dimensional Gaussian droplet. *Comm. Math. Phys.*, 179(2):467–488, 1996.

[BDFG04] J. Bouttier, P. Di Francesco, and E. Guitter. Planar maps as labeled mobiles. *Electron. J. Combin.*, 11(1):Research Paper 69, 27, 2004. math/0405099.

[Ber] Nathanaël Berestycki. Introduction to the gaussian free field and liouville quantum gravity.

[Ber96] Jean Bertoin. *Lévy processes*, volume 121 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1996.

[Ber07] Olivier Bernardi. Bijective counting of tree-rooted maps and shuffles of parenthesis systems. *Electron. J. Combin.*, 14(1):Research Paper 9, 36 pp. (electronic), 2007. math/0601684.

[BH91] MT Batchelor and BI Henry. Limits to Eden growth in two and three dimensions. *Physics Letters A*, 157(4):229–236, 1991.

[Bis07] Christopher J. Bishop. Conformal welding and Koebe's theorem. *Ann. of Math. (2)*, 166(3):613–656, 2007.

[BLPS01] Itai Benjamini, Russell Lyons, Yuval Peres, and Oded Schramm. Uniform spanning forests. *Ann. Probab.*, 29(1):1–65, 2001.

[BN11] N Berestycki and JR Norris. Lectures on schramm–loewner evolution. 2011.

[BNMR01] Ole E. Barndorff-Nielsen, Thomas Mikosch, and Sidney I. Resnick, editors. *Lévy processes*. Birkhäuser Boston Inc., Boston, MA, 2001. Theory and applications.

[BRH10] J. . H. Bakke, P. Ray, and A. Hansen. Morphology of Laplacian random walks. *EPL (Europhysics Letters)*, 92(3):36004, 2010.

[BS01] Itai Benjamini and Oded Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6:no. 23, 13 pp. (electronic), 2001. math/0011019.

[BS09a] Dmitry Beliaev and Stanislav Smirnov. Harmonic measure and SLE. *Comm. Math. Phys.*, 290(2):577–595, 2009.

[BS09b] Itai Benjamini and Oded Schramm. KPZ in one dimensional random geometry of multiplicative cascades. *Comm. Math. Phys.*, 289(2):653–662, 2009. 0806.1347.

[Car06] J. Cardy. Lectures on Stochastic Loewner Evolution and other growth processes in two dimensions, March 2006.

[CD81] J. Theodore Cox and Richard Durrett. Some limit theorems for percolation processes with necessary and sufficient conditions. *Ann. Probab.*, 9(4):583–603, 1981.

[CK13] N. Curien and I. Kortchemski. Random stable looptrees. *ArXiv e-prints*, April 2013.

[CL12] N. Curien and J.-F. Le Gall. The Brownian plane. *ArXiv e-prints*, April 2012.

[CL14] N. Curien and J.-F. Le Gall. The hull process of the Brownian plane. *ArXiv e-prints*, September 2014.

[CLG12] Nicolas Curien and Jean-François Le Gall. The Brownian plane. *Journal of Theoretical Probability*, pages 1–43, 2012. 1204.5921.

[CLUB09] Ma. Emilia Caballero, Amaury Lambert, and Gerónimo Uribe Bravo. Proof(s) of the Lamperti representation of continuous-state branching processes. *Probab. Surv.*, 6:62–89, 2009.

[CM01] L. Carleson and N. Makarov. Aggregation in the plane and Loewner's equation. *Comm. Math. Phys.*, 216(3):583–607, 2001.

[CM12] N. Curien and G. Miermont. Uniform infinite planar quadrangulations with a boundary. *ArXiv e-prints*, February 2012. 1202.5452.

[CMB96] Marek Cieplak, Amos Maritan, and Jayanth R. Banavar. Invasion percolation and Eden growth: Geometry and universality. *Phys. Rev. Lett.*, 76:3754–3757, May 1996.

[CMM13] N. Curien, L. Ménard, and G. Miermont. A view from infinity of the uniform infinite planar quadrangulation. *ALEA Lat. Am. J. Probab. Math. Stat.*, 10(1):45–88, 2013. 1201.1052.

[CN06a] Federico Camia and Charles M. Newman. Two-dimensional critical percolation: the full scaling limit. *Comm. Math. Phys.*, 268(1):1–38, 2006.

[CN06b] Federico Camia and Charles M. Newman. Two-dimensional critical percolation: the full scaling limit. *Comm. Math. Phys.*, 268(1):1–38, 2006. math/0605035.

[CN08] Federico Camia and Charles M. Newman. SLE_6 and CLE_6 from critical percolation. In *Probability, geometry and integrable systems*, volume 55 of *Math. Sci. Res. Inst. Publ.*, pages 103–130. Cambridge Univ. Press, Cambridge, 2008.

[Cor12a] Ivan Corwin. The Kardar-Parisi-Zhang equation and universality class. *Random Matrices Theory Appl.*, 1(1):1130001, 76, 2012. 1106.1596.

[Cor12b] Ivan Corwin. The Kardar-Parisi-Zhang equation and universality class. *Random Matrices Theory Appl.*, 1(1):1130001, 76, 2012. 1106.1596.

[CQ11] I. Corwin and J. Quastel. Renormalization fixed point of the KPZ universality class. *ArXiv e-prints*, March 2011.

[CS] D. Chelkak and S. Smirnov. Universality in the 2D Ising model and conformal invariance of fermionic observables. *Invent. Math.*

[CS04] Philippe Chassaing and Gilles Schaeffer. Random planar lattices and integrated superBrownian excursion. *Probab. Theory Related Fields*, 128(2):161–212, 2004. math/0205226.

[CS09] Dmitry Chelkak and Stanislav Smirnov. Universality in the 2D Ising model and conformal invariance of fermionic observables. *ArXiv e-prints*, October 2009.

[CS12] Dmitry Chelkak and Stanislav Smirnov. Universality in the 2D Ising model and conformal invariance of fermionic observables. *Invent. Math.*, 189(3):515–580, 2012. 0910.2045.

[Cur13] N. Curien. A glimpse of the conformal structure of random planar maps. *ArXiv e-prints*, August 2013.

[CV81] Robert Cori and Bernard Vauquelin. Planar maps are well labeled trees. *Canad. J. Math.*, 33(5):1023–1042, 1981.

[DF91] P. Diaconis and W. Fulton. A growth model, a game, an algebra, Lagrange inversion, and characteristic classes. *Rend. Sem. Mat. Univ. Politec. Torino*, 49(1):95–119 (1993), 1991. Commutative algebra and algebraic geometry, II (Italian) (Turin, 1990).

[DFGZJ95] P. Di Francesco, P. Ginsparg, and J. Zinn-Justin. 2D gravity and random matrices. *Phys. Rep.*, 254:1–133, 1995.

[DFMS97] Philippe Di Francesco, Pierre Mathieu, and David Sénéchal. *Conformal field theory*. Graduate Texts in Contemporary Physics. Springer-Verlag, New York, 1997.

[DK88] Bertrand Duplantier and Ivan Kostov. Conformal spectra of polymers on a random surface. *Phys. Rev. Lett.*, 61(13):1433–1437, 1988.

[DLG02] Thomas Duquesne and Jean-François Le Gall. Random trees, Lévy processes and spatial branching processes. *Astérisque*, 281:vi+147, 2002. math/0509558.

[DLG05] Thomas Duquesne and Jean-François Le Gall. Probabilistic and fractal aspects of Lévy trees. *Probab. Theory Related Fields*, 131(4):553–603, 2005.

[DLG06] Thomas Duquesne and Jean-François Le Gall. The Hausdorff measure of stable trees. *ALEA Lat. Am. J. Probab. Math. Stat.*, 1:393–415, 2006.

[DLG09] Thomas Duquesne and Jean-François Le Gall. On the re-rooting invariance property of Lévy trees. *Electron. Commun. Probab.*, 14:317–326, 2009.

[DMS14] B. Duplantier, J. Miller, and S. Sheffield. Liouville quantum gravity as a mating of trees. *ArXiv e-prints*, September 2014.

[DRSV12a] B. Duplantier, R. Rhodes, S. Sheffield, and V. Vargas. Critical Gaussian Multiplicative Chaos: Convergence of the Derivative Martingale. *ArXiv e-prints*, June 2012.

[DRSV12b] B. Duplantier, R. Rhodes, S. Sheffield, and V. Vargas. Renormalization of Critical Gaussian Multiplicative Chaos and KPZ formula. *ArXiv e-prints*, December 2012.

[DRSV14a] Bertrand Duplantier, Rémi Rhodes, Scott Sheffield, and Vincent Vargas. Critical gaussian multiplicative chaos: Convergence of the derivative martingale. *The Annals of Probability*, 42(5):1769–1808, 09 2014.

[DRSV14b] Bertrand Duplantier, Rémi Rhodes, Scott Sheffield, and Vincent Vargas. Log-correlated gaussian fields: an overview. *arXiv preprint arXiv:1407.5605*, 2014.

[DRSV14c] Bertrand Duplantier, Rémi Rhodes, Scott Sheffield, and Vincent Vargas. Renormalization of Critical Gaussian Multiplicative Chaos and KPZ Relation. *Comm. Math. Phys.*, 330(1):283–330, 2014.

[DS09] Bertrand Duplantier and Scott Sheffield. Duality and the Knizhnik-Polyakov-Zamolodchikov relation in Liouville quantum gravity. *Phys. Rev. Lett.*, 102(15):150603, 4, 2009.

[DS11a] Bertrand Duplantier and Scott Sheffield. Liouville quantum gravity and KPZ. *Invent. Math.*, 185(2):333–393, 2011. 0808.1560.

[DS11b] Bertrand Duplantier and Scott Sheffield. Schramm-Loewner evolution and Liouville quantum gravity. *Phys. Rev. Lett.*, 107:131305, 2011. 1012.4800.

[DS11c] Bertrand Duplantier and Scott Sheffield. Schramm-Loewner evolution and Liouville quantum gravity. *Phys. Rev. Lett.*, 107:131305, Sep 2011. 1012.4800.

[Dub09a] Julien Dubédat. Duality of Schramm-Loewner evolutions. *Ann. Sci. Éc. Norm. Supér. (4)*, 42(5):697–724, 2009.

[Dub09b] Julien Dubédat. Sle and the free field: partition functions and couplings. *Journal of the American Mathematical Society*, 22(4):995–1054, 2009.

[Dub09c] Julien Dubédat. SLE and the free field: partition functions and couplings. *J. Amer. Math. Soc.*, 22(4):995–1054, 2009.

[Dub09d] Julien Dubédat. SLE and the free field: partition functions and couplings. *J. Amer. Math. Soc.*, 22(4):995–1054, 2009.

[Dup98] Bertrand Duplantier. Random walks and quantum gravity in two dimensions. *Phys. Rev. Lett.*, 81(25):5489–5492, 1998.

[Dup00] Bertrand Duplantier. Conformally invariant fractals and potential theory. *Phys. Rev. Lett.*, 84(7):1363–1367, 2000.

[Dup04] Bertrand Duplantier. Conformal fractal geometry & boundary quantum gravity. In M. L. Lapidus and M. van Frankenhuysen, editors, *Fractal geometry and applications: a jubilee of Benoît Mandelbrot, Part 2*, volume 72 of *Proc. Sympos. Pure Math.*, pages 365–482. Amer. Math. Soc., Providence, RI, 2004.

[Dup06] Bertrand Duplantier. Conformal random geometry. In A. Bovier, F. Dunlop, F. den Hollander, A. van Enter, and J. Dalibard, editors, *Mathematical statistical physics (Les Houches Summer School, Session LXXXIII, 2005)*, pages 101–217. Elsevier B. V., Amsterdam, 2006.

[Dur10] Rick Durrett. *Probability: theory and examples*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, fourth edition, 2010.

[Ede61] Murray Eden. A two-dimensional growth process. In *Proc. 4th Berkeley Sympos. Math. Statist. and Prob., Vol. IV*, pages 223–239. Univ. California Press, Berkeley, Calif., 1961.

[Eva85] Steven N. Evans. On the Hausdorff dimension of Brownian cone points. *Math. Proc. Cambridge Philos. Soc.*, 98(2):343–353, 1985.

[Fat80] A. Fathi. Structure of the group of homeomorphisms preserving a good measure on a compact manifold. *Ann. Sci. École Norm. Sup. (4)*, 13(1):45–93, 1980.

[FSS85] P Freche, D Stauffer, and HE Stanley. Surface structure and anisotropy of Eden clusters. *Journal of Physics A: Mathematical and General*, 18(18):L1163, 1985.

[Gar12] C. Garban. Quantum gravity and the KPZ formula. *ArXiv e-prints*, June 2012.

[GGN13] Ori Gurel-Gurevich and Asaf Nachmias. Recurrence of planar graph limits. *Ann. of Math. (2)*, 177(2):761–781, 2013. 1206.0707.

[GHMS15] E. Gwynne, N. Holden, J. Miller, and X. Sun. Brownian motion correlation in the peanosphere for $\kappa \geq 8$. *ArXiv e-prints*, October 2015.

[GKMW16] Ewain Gwynne, Adrien Kassel, Jason Miller, and David B. Wilson. Topologically-weighted spanning trees on planar maps and SLE-decorated Liouville quantum gravity for $\kappa \geq 8$. 2016. In preparation.

[GM93] P. Ginsparg and G. Moore. Lectures on 2D gravity and 2D string theory (TASI 1992). In J. Harvey and J. Polchinski, editors, *Recent direction in particle theory, Proceedings of the 1992 TASI*. World Scientific, Singapore, 1993.

[GR] James Gill and Steffen Rohde. On the Riemann surface type of limits of planar maps. In preparation.

[GR13] James T. Gill and Steffen Rohde. On the Riemann surface type of random planar maps. *Rev. Mat. Iberoam.*, 29(3):1071–1090, 2013. 1101.1320.

[GRS08] C. Garban, S. Rohde, and O. Schramm. Continuity of the SLE trace in simply connected domains. *ArXiv e-prints*, October 2008.

[GV06] Björn Gustafsson and Alexander Vasilćev. *Conformal and potential analysis in Hele-Shaw cells*. Advances in Mathematical Fluid Mechanics. Birkhäuser Verlag, Basel, 2006.

[Hal00] Thomas C Halsey. Diffusion-limited aggregation: a model for pattern formation. *Physics Today*, 53(11):36–41, 2000.

[Has01] M. B. Hastings. Fractal to nonfractal phase transition in the dielectric breakdown model. *Phys. Rev. Lett.*, 87:175502, Oct 2001.

[Has02] Matthew B Hastings. Exact multifractal spectra for arbitrary Laplacian random walks. *Physical review letters*, 88(5):055506, 2002. cond-mat/0109304.

[HBB10a] C. Hagendorf, D. Bernard, and M. Bauer. The Gaussian Free Field and SLE_4 on Doubly Connected Domains. *Journal of Statistical Physics*, 140:1–26, July 2010.

[HBB10b] Christian Hagendorf, Denis Bernard, and Michel Bauer. The Gaussian free field and SLE_4 on doubly connected domains. *J. Stat. Phys.*, 140(1):1–26, 2010.

[HK71] Raphael Høegh-Krohn. A general class of quantum fields without cut-offs in two space-time dimensions. *Communications in Mathematical Physics*, 21(3):244–255, 1971.

[HL98] Matthew B Hastings and Leonid S Levitov. Laplacian growth as one-dimensional turbulence. *Physica D: Nonlinear Phenomena*, 116(1):244–252, 1998. cond-mat/9607021.

[HR12] Denise M Halverson and Dušan Repovš. Survey on the Generalized RL Moore Problem. *arXiv e-prints*, January 2012.

[HW65] J. M. Hammersley and D. J. A. Welsh. First-passage percolation, subadditive processes, stochastic networks, and generalized renewal theory. In *Proc. Internat. Res. Semin., Statist. Lab., Univ. California, Berkeley, Calif*, pages 61–110. Springer-Verlag, New York, 1965.

[IK10] K. Izyurov and K. Kytölä. Hadamard's formula and couplings of SLEs with free field. *ArXiv e-prints*, June 2010.

[IK13] Konstantin Izyurov and Kalle Kytölä. Hadamard's formula and couplings of SLEs with free field. *Probab. Theory Related Fields*, 155(1-2):35–69, 2013.

[IM74] Kiyosi Itô and Henry P. McKean, Jr. *Diffusion processes and their sample paths*. Springer-Verlag, Berlin-New York, 1974. Second printing, corrected, Die Grundlehren der mathematischen Wissenschaften, Band 125.

[Jan97] Svante Janson. *Gaussian Hilbert spaces*, volume 129 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1997.

[Jiř58] Miloslav Jiřina. Stochastic branching processes with continuous state space. *Czechoslovak Math. J.*, 8 (83):292–313, 1958.

[JLS10] D. Jerison, L. Levine, and S. Sheffield. Internal DLA in Higher Dimensions. *Elec. Journ. Prob., to appear*, 2010. 1012.3453.

[JLS11] D. Jerison, L. Levine, and S. Sheffield. Internal DLA and the Gaussian free field. *Duke Math. Journ. to appear.*, 2011. 1101.0596.

[JLS12a] David Jerison, Lionel Levine, and Scott Sheffield. Logarithmic fluctuations for internal DLA. *J. Amer. Math. Soc.*, 25(1):271–301, 2012. 1010.2483.

[JLS12b] David Jerison, Lionel Levine, and Scott Sheffield. Logarithmic fluctuations for internal DLA. *J. Amer. Math. Soc.*, 25(1):271–301, 2012. 1010.2483.

[JLS13] David Jerison, Lionel Levine, and Scott Sheffield. Internal dla in higher dimensions. *Electron. J. Probab*, 18(98):1–14, 2013.

[JLS+14] David Jerison, Lionel Levine, Scott Sheffield, et al. Internal dla and the gaussian free field. *Duke Mathematical Journal*, 163(2):267–308, 2014.

[Jon] Peter Jones. Private communciation.

[JS98] Benjamin Jacquard and Gilles Schaeffer. A bijective census of nonseparable planar maps. *J. Combin. Theory Ser. A*, 83(1):1–20, 1998.

[JS00a] Peter W. Jones and Stanislav K. Smirnov. Removability theorems for Sobolev functions and quasiconformal maps. *Ark. Mat.*, 38(2):263–279, 2000.

[JS00b] Peter W. Jones and Stanislav K. Smirnov. Removability theorems for Sobolev functions and quasiconformal maps. *Ark. Mat.*, 38(2):263–279, 2000.

[JVST12] Fredrik Johansson Viklund, Alan Sola, and Amanda Turner. Scaling limits of anisotropic Hastings-Levitov clusters. volume 48, pages 235–257, 2012. 0908.0086.

[Kah85] J-P Kahane. Le chaos multiplicatif. *Comptes rendus de l'Académie des sciences. Série 1, Mathématique*, 301(6):329–332, 1985.

[Kan07] Nam-Gyu Kang. Boundary behavior of SLE. *J. Amer. Math. Soc.*, 20(1):185–210 (electronic), 2007.

[Ken00a] Richard Kenyon. The asymptotic determinant of the discrete laplacian. *Acta Mathematica*, 185(2):239–286, 2000.

[Ken00b] Richard Kenyon. Conformal invariance of domino tiling. *Ann. Probab.*, 28(2):759–795, 2000.

[Ken00c] Richard Kenyon. Conformal invariance of domino tiling. *Ann. Probab.*, 28(2):759–795, 2000.

[Ken01a] R. Kenyon. Dominos and the Gaussian free field. *Annals of Probability*, 29:1128–1137, 2001.

[Ken01b] Richard Kenyon. Dominos and the Gaussian free field. *Ann. Probab.*, 29(3):1128–1137, 2001.

[Ken08] Richard Kenyon. Height fluctuations in the honeycomb dimer model. *Comm. Math. Phys.*, 281(3):675–709, 2008.

[Kes87] Harry Kesten. Hitting probabilities of random walks on $\mathbf{Z}^d$. *Stochastic Process. Appl.*, 25(2):165–184, 1987.

[KMSW15] R. Kenyon, J. Miller, S. Sheffield, and D. B. Wilson. Bipolar orientations on planar maps and SLE_12. *ArXiv e-prints*, November 2015.

[KN04] Wouter Kager and Bernard Nienhuis. A guide to stochastic löwner evolution and its applications. *Journal of statistical physics*, 115(5-6):1149–1229, 2004.

[KPZ86] Mehran Kardar, Giorgio Parisi, and Yi-Cheng Zhang. Dynamic scaling of growing interfaces. *Phys. Rev. Lett.*, 56:889–892, Mar 1986.

[KPZ88a] V. G. Knizhnik, A. M. Polyakov, and A. B. Zamolodchikov. Fractal structure of 2D-quantum gravity. *Modern Phys. Lett. A*, 3(8):819–826, 1988.

[KPZ88b] Vadim G. Knizhnik, Alexander M. Polyakov, and Alexander B. Zamolodchikov. Fractal structure of 2D-quantum gravity. *Modern Phys. Lett. A*, 3(8):819–826, 1988.

[Kri05] M. Krikun. Local structure of random quadrangulations. *ArXiv e-prints*, December 2005.

[KS91a] Ioannis Karatzas and Steven E. Shreve. *Brownian motion and stochastic calculus*, volume 113 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1991.

[KS91b] Ioannis Karatzas and Steven E. Shreve. *Brownian motion and stochastic calculus*, volume 113 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1991.

[KSS] A. Kemppainen, S. Sheffield, and S. Schramm. TBD.

[Kyp06] Andreas E. Kyprianou. *Introductory lectures on fluctuations of Lévy processes with applications*. Universitext. Springer-Verlag, Berlin, 2006.

[Lam67a] John Lamperti. Continuous state branching processes. *Bull. Amer. Math. Soc.*, 73:382–386, 1967.

[Lam67b] John Lamperti. Continuous state branching processes. *Bull. Amer. Math. Soc.*, 73:382–386, 1967.

[Lam67c] John Lamperti. The limit of a sequence of branching processes. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete*, 7:271–288, 1967.

[Law05a] Gregory F. Lawler. *Conformally invariant processes in the plane*, volume 114 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2005.

[Law05b] Gregory F. Lawler. *Conformally invariant processes in the plane*, volume 114 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2005.

[Law06] Gregory F. Lawler. The Laplacian-b random walk and the Schramm-Loewner evolution. *Illinois J. Math.*, 50(1-4):701–746 (electronic), 2006.

[Law09a] G. Lawler. Conformal invariance and $2D$ statistical physics. *Bull. AMS*, 46:25–54, 2009.

[Law09b] Gregory F Lawler. Partition functions, loop measure, and versions of sle. *Journal of Statistical Physics*, 134(5-6):813–837, 2009.

[LBG92] Gregory F. Lawler, Maury Bramson, and David Griffeath. Internal diffusion limited aggregation. *Ann. Probab.*, 20(4):2117–2140, 1992.

[Le 14] J.-F. Le Gall. Random geometry on the sphere. *ArXiv e-prints*, March 2014.

[LEP86] JW Lyklema, C Evertsz, and L Pietronero. The Laplacian random walk. *EPL (Europhysics Letters)*, 2(2):77, 1986.

[LG99] Jean-François Le Gall. *Spatial branching processes, random snakes and partial differential equations*. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, 1999.

[LG07a] Jean-François Le Gall. The topological structure of scaling limits of large planar maps. *Invent. Math.*, 169(3):621–670, 2007.

[LG07b] Jean-François Le Gall. The topological structure of scaling limits of large planar maps. *Invent. Math.*, 169(3):621–670, 2007. math/0607567.

[LG10] Jean-François Le Gall. Geodesics in large planar maps and in the Brownian map. *Acta Math.*, 205(2):287–360, 2010. 0804.3012.

[LG13a] Jean-François Le Gall. Uniqueness and universality of the Brownian map. *Ann. Probab.*, 41(4):2880–2960, 2013. 1105.4842.

[LG13b] Jean-François Le Gall. Uniqueness and universality of the Brownian map. *Ann. Probab.*, 41(4):2880–2960, 2013. 1105.4842.

[LG14] J.-F. Le Gall. The Brownian map: a universal limit for random planar maps. In *XVIIth International Congress on Mathematical Physics*, pages 420–428. World Sci. Publ., Hackensack, NJ, 2014.

[LGLJ98] Jean-Francois Le Gall and Yves Le Jan. Branching processes in Lévy processes: the exploration process. *Ann. Probab.*, 26(1):213–252, 1998.

[LGM+12] Jean-François Le Gall, Grégory Miermont, et al. Scaling limits of random trees and planar maps. *Probability and statistical physics in two and more dimensions*, 15:155–211, 2012.

[LGP08] Jean-François Le Gall and Frédéric Paulin. Scaling limits of bipartite planar maps are homeomorphic to the 2-sphere. *Geom. Funct. Anal.*, 18(3):893–918, 2008. math/0612315.

[Lov71] C Lovelace. Pomeron form factors and dual regge cuts. *Physics Letters B*, 34(6):500–506, 1971.

[LR12] G. F. Lawler and M. A. Rezaei. Minkowski content and natural parameterization for the Schramm-Loewner evolution. *ArXiv e-prints*, November 2012.

[LS09] Gregory F. Lawler and Scott Sheffield. The natural parametrization for the Schramm-Loewner evolution. *ArXiv e-prints*, June 2009.

[LS11] Gregory F. Lawler and Scott Sheffield. A natural parametrization for the Schramm-Loewner evolution. *Ann. Probab.*, 39(5):1896–1937, 2011. 0906.3804.

[LSSW14] Asad Lodhia, Scott Sheffield, Xin Sun, and Samuel S Watson. Fractional gaussian fields: a survey. *arXiv preprint arXiv:1407.5598*, 2014.

[LSW03a] Gregory Lawler, Oded Schramm, and Wendelin Werner. Conformal restriction: the chordal case. *J. Amer. Math. Soc.*, 16(4):917–955 (electronic), 2003. math/0209343.

[LSW03b] Gregory Lawler, Oded Schramm, and Wendelin Werner. Conformal restriction: the chordal case. *J. Amer. Math. Soc.*, 16(4):917–955 (electronic), 2003. math/0209343.

[LSW04a] Gregory F. Lawler, Oded Schramm, and Wendelin Werner. Conformal invariance of planar loop-erased random walks and uniform spanning trees. *Ann. Probab.*, 32(1B):939–995, 2004.

[LSW04b] Gregory F. Lawler, Oded Schramm, and Wendelin Werner. Conformal invariance of planar loop-erased random walks and uniform spanning trees. *Ann. Probab.*, 32(1B):939–995, 2004. math/0112234.

[LSW04c] Gregory F. Lawler, Oded Schramm, and Wendelin Werner. Conformal invariance of planar loop-erased random walks and uniform spanning trees. *Ann. Probab.*, 32(1B):939–995, 2004.

[LSW04d] Gregory F. Lawler, Oded Schramm, and Wendelin Werner. Conformal invariance of planar loop-erased random walks and uniform spanning trees. *Ann. Probab.*, 32(1B):939–995, 2004. math/0112234.

[LTW09] S.-Y. Lee, R. Teodorescu, and P. Wiegmann. Shocks and finite-time singularities in Hele-Shaw flow. *Phys. D*, 238(14):1113–1128, 2009.

[LW99] G. F. Lawler and W. Werner. Intersection exponents for planar Brownian motion. *Ann. Probab.*, 27(4):1601–1642, 1999.

[LW00] Gregory F. Lawler and Wendelin Werner. Universality for conformally invariant intersection exponents. *J. Eur. Math. Soc. (JEMS)*, 2(4):291–328, 2000.

[LYTZC12] Deng Li, Wang Yan-Ting, and Ou-Yang Zhong-Can. Diffusion-limited aggregation with polygon particles. *Communications in Theoretical Physics*, 58(6):895, 2012. 1209.5958.

[LZ10] G. F. Lawler and W. Zhou. SLE curves and natural parametrization. *ArXiv e-prints*, June 2010.

[LZ13] Gregory F. Lawler and Wang Zhou. *SLE* curves and natural parametrization. *Ann. Probab.*, 41(3A):1556–1584, 2013. 1006.4936.

[MD86] Paul Meakin and JM Deutch. The formation of surfaces by diffusion limited annihilation. *The Journal of chemical physics*, 85:2320, 1986.

[Mea86] Paul Meakin. Universality, nonuniversality, and the effects of anisotropy on diffusion-limited aggregation. *Physical Review A*, 33(5):3371, 1986.

[Men12] Anton Menshutin. Scaling in the diffusion limited aggregation model. *Phys. Rev. Lett.*, 108:015501, Jan 2012.

[Mer94] Russell Merris. Laplacian matrices of graphs: a survey. *Linear algebra and its applications*, 197:143–176, 1994.

[Mie13a] Grégory Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.*, 210(2):319–401, 2013.

[Mie13b] Grégory Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.*, 210(2):319–401, 2013. 1104.1606.

[Mil04] John Milnor. Pasting together Julia sets: a worked out example of mating. *Experiment. Math.*, 13(1):55–92, 2004.

[Mil06] John Milnor. *Dynamics in one complex variable*, volume 160 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, third edition, 2006.

[Mil10a] J. Miller. Fluctuations for the Ginzburg-Landau model on a bounded domain. 2010.

[Mil10b] J. Miller. Universality for SLE(4). 2010.

[Mil10c] J. Miller. Universality for SLE(4). *ArXiv e-prints*, October 2010.

[MJ02] Joachim Mathiesen and Mogens H. Jensen. Tip splittings and phase transitions in the dielectric breakdown model: Mapping to the diffusion-limited aggregation model. *Phys. Rev. Lett.*, 88:235505, May 2002.

[MJB08] Joachim Mathiesen, Mogens H. Jensen, and Jan Øystein Haavig Bakke. Dimensions, maximal growth sites, and optimization in the dielectric breakdown model. *Phys. Rev. E*, 77:066203, Jun 2008.

[MM06a] Jean-François Marckert and Abdelkader Mokkadem. Limit of normalized quadrangulations: the Brownian map. *Ann. Probab.*, 34(6):2144–2202, 2006.

[MM06b] Jean-François Marckert and Abdelkader Mokkadem. Limit of normalized quadrangulations: the Brownian map. *Ann. Probab.*, 34(6):2144–2202, 2006. math/0403398.

[Moo25] R. L. Moore. Concerning upper semi-continuous collections of continua. *Trans. Amer. Math. Soc.*, 27(4):416–428, 1925.

[MP83] Sidney A. Morris and Vincent C. Peck. A note on homeomorphic measures on topological groups. *Proc. Edinburgh Math. Soc. (2)*, 26(2):169–171, 1983.

[MP10] Peter Mörters and Yuval Peres. *Brownian motion*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 2010. With an appendix by Oded Schramm and Wendelin Werner.

[MPST06] Joachim Mathiesen, Itamar Procaccia, Harry L Swinney, and Matthew Thrasher. The universality class of diffusion-limited aggregation and viscous fingering. *EPL (Europhysics Letters)*, 76(2):257, 2006. cond-mat/0512274.

[MR07] Donald E. Marshall and Steffen Rohde. Convergence of a variant of the zipper algorithm for conformal mapping. *SIAM J. Numer. Anal.*, 45(6):2577–2609 (electronic), 2007.

[MS09] Nicolai Makarov and Stanislav Smirnov. Off-critical lattice models and massive SLEs. *ArXiv e-prints*, September 2009.

[MS10] Nikolai Makarov and Stanislav Smirnov. Off-critical lattice models and massive SLEs. pages 362–371, 2010.

[MS11] A Yu Menshutin and LN Shchur. Morphological diagram of diffusion driven aggregate growth in plane: Competition of anisotropy and adhesion. *Computer Physics Communications*, 182(9):1819–1823, 2011. 1008.3449.

[MS12a] J. Miller and S. Sheffield. Imaginary Geometry I: Interacting SLEs. *ArXiv e-prints*, January 2012.

[MS12b] J. Miller and S. Sheffield. Imaginary geometry II: reversibility of $\mathrm{SLE}_\kappa(\rho_1;\rho_2)$ for $\kappa \in (0,4)$. *ArXiv e-prints*, January 2012.

[MS12c] J. Miller and S. Sheffield. Imaginary geometry III: reversibility of SLE_κ for $\kappa \in (4, 8)$. *ArXiv e-prints*, January 2012.

[MS13a] J. Miller and S. Sheffield. Imaginary geometry IV: interior rays, whole-plane reversibility, and space-filling trees. *ArXiv e-prints*, February 2013.

[MS13b] J. Miller and S. Sheffield. Quantum Loewner Evolution. *ArXiv e-prints*, December 2013.

[MS15a] J. Miller and S. Sheffield. An axiomatic characterization of the Brownian map. *ArXiv e-prints*, June 2015.

[MS15b] J. Miller and S. Sheffield. Liouville quantum gravity and the Brownian map I: The QLE(8/3,0) metric. *ArXiv e-prints*, July 2015.

[MS15c] J. Miller and S. Sheffield. Liouville quantum gravity spheres as matings of finite-diameter trees. *ArXiv e-prints*, June 2015.

[MS15d] Jason Miller and Scott Sheffield. Liouville quantum gravity and the Brownian map II: geodesics and continuity of the embedding. 2015. In preparation.

[MS15e] Jason Miller and Scott Sheffield. Liouville quantum gravity and the Brownian map III: the conformal structure is determined. 2015. In preparation.

[Mul67a] R. C. Mullin. On the enumeration of tree-rooted maps. *Canad. J. Math.*, 19:174–183, 1967.

[Mul67b] R. C. Mullin. On the enumeration of tree-rooted maps. *Canad. J. Math.*, 19:174–183, 1967.

[Nak04] Yu Nakayama. Liouville field theory: a decade after the revolution. *Internat. J. Modern Phys. A*, 19(17-18):2771–2930, 2004.

[NPW84] L. Niemeyer, L. Pietronero, and H. J. Wiesmann. Fractal dimension of dielectric breakdown. *Phys. Rev. Lett.*, 52(12):1033–1036, 1984.

[NS97] Ali Naddaf and Thomas Spencer. On homogenization and scaling limit of some gradient perturbations of a massless free field. *Comm. Math. Phys.*, 183(1):55–84, 1997.

[NT12] James Norris and Amanda Turner. Hastings-Levitov aggregation in the small-particle limit. *Comm. Math. Phys.*, 316(3):809–841, 2012. 1106.3546.

[OU41] J. C. Oxtoby and S. M. Ulam. Measure-preserving homeomorphisms and metrical transitivity. *Ann. of Math. (2)*, 42:874–920, 1941.

[Pol81a] A. M. Polyakov. Quantum geometry of bosonic strings. *Phys. Lett. B*, 103(3):207–210, 1981.

[Pol81b] A. M. Polyakov. Quantum geometry of bosonic strings. *Phys. Lett. B*, 103(3):207–210, 1981.

[Pol81c] A. M. Polyakov. Quantum geometry of fermionic strings. *Phys. Lett. B*, 103(3):211–213, 1981.

[Pol08a] A. M. Polyakov. From Quarks to Strings. *ArXiv e-prints*, November 2008.

[Pol08b] Alexander M. Polyakov. From Quarks to Strings. 2008.

[Pro90] Philip Protter. *Stochastic integration and differential equations*, volume 21 of *Applications of Mathematics (New York)*. Springer-Verlag, Berlin, 1990. A new approach.

[PW98] James Gary Propp and David Bruce Wilson. How to get a perfectly random sample from a generic markov chain and generate a random spanning tree of a directed graph. *Journal of Algorithms*, 27(2):170–217, 1998.

[RS05a] Steffen Rohde and Oded Schramm. Basic properties of SLE. *Ann. of Math. (2)*, 161(2):883–924, 2005.

[RS05b] Steffen Rohde and Oded Schramm. Basic properties of SLE. *Ann. of Math. (2)*, 161(2):883–924, 2005. math/0106036.

[RV07] Brian Rider and Bálint Virág. The noise in the circular law and the Gaussian free field. *Int. Math. Res. Not. IMRN*, (2):Art. ID rnm006, 33, 2007.

[RV08a] R. Rhodes and V. Vargas. KPZ formula for log-infinitely divisible multifractal random measures. *ArXiv e-prints:0807.1036*, 2008. To appear in ESAIM P&S.

[RV08b] R. Rhodes and V. Vargas. KPZ formula for log-infinitely divisible multifractal random measures. *ArXiv e-prints:0807.1036*, 2008. To appear in ESAIM P&S.

[RV11] Rémi Rhodes and Vincent Vargas. KPZ formula for log-infinitely divisible multifractal random measures. *ESAIM Probab. Stat.*, 15:358–371, 2011. 0807.1036.

[RV14] Rémi Rhodes and Vincent Vargas. Gaussian multiplicative chaos and applications: a review. *Probab. Surv.*, 11:315–392, 2014. 1305.6221.

[RY99a] Daniel Revuz and Marc Yor. *Continuous martingales and Brownian motion*, volume 293 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, third edition, 1999.

[RY99b] Daniel Revuz and Marc Yor. *Continuous martingales and Brownian motion*, volume 293 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, third edition, 1999.

[RY99c] Daniel Revuz and Marc Yor. *Continuous martingales and Brownian motion*, volume 293 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, third edition, 1999.

[RZ05] Steffen Rohde and Michel Zinsmeister. Some remarks on Laplacian growth. *Topology Appl.*, 152(1-2):26–43, 2005.

[San00] Leonard M Sander. Diffusion-limited aggregation: a kinetic critical phenomenon? *Contemporary Physics*, 41(4):203–218, 2000.

[Sar90] Peter Sarnak. Determinants of laplacians; heights and finiteness. *Analysis, et cetera, Academic Press, Boston, MA*, pages 601–622, 1990.

[Sat99] Ken-iti Sato. *Lévy processes and infinitely divisible distributions*, volume 68 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1999. Translated from the 1990 Japanese original, Revised by the author.

[Sch97] Gilles Schaeffer. Bijective census and random generation of Eulerian planar maps with prescribed vertex degrees. *Electron. J. Combin.*, 4(1):Research Paper 20, 14 pp. (electronic), 1997.

[Sch99] Gilles Schaeffer. Random sampling of large planar maps and convex polyhedra. In *Annual ACM Symposium on Theory of Computing (Atlanta, GA, 1999)*, pages 760–769 (electronic). ACM, New York, 1999.

[Sch00a] Oded Schramm. Scaling limits of loop-erased random walks and uniform spanning trees. *Israel J. Math.*, 118:221–288, 2000. math/9904022.

[Sch00b] Oded Schramm. Scaling limits of loop-erased random walks and uniform spanning trees. *Israel J. Math.*, 118:221–288, 2000.

[Sch00c] Oded Schramm. Scaling limits of loop-erased random walks and uniform spanning trees. *Israel J. Math.*, 118:221–288, 2000. math/9904022.

[Sch07] Oded Schramm. Conformally invariant scaling limits: an overview and a collection of problems. In *International Congress of Mathematicians. Vol. I*, pages 513–543. Eur. Math. Soc., Zürich, 2007. math/0602151.

[Sei90] Nathan Seiberg. Notes on quantum Liouville theory and quantum gravity. *Progr. Theoret. Phys. Suppl.*, (102):319–349 (1991), 1990. Common trends in mathematics and quantum field theories (Kyoto, 1990).

[She] S. Sheffield. Local sets of the gaussian free field: slides and audio. www.fields.utoronto.ca/0506/percolationsle/sheffield1, www.fields.utoronto.ca/audio/0506/percolationsle/sheffield2, www.fields.utoronto.ca/audio/0506/percolationsle/sheffield3.

[She05] Scott Sheffield. Local sets of the Gaussian free field: slides and audio. www.fields.utoronto.ca/audio/05-06/percolation_SLE/sheffield1, www.fields.utoronto.ca/audio/05-06/percolation_SLE/sheffield2, www.fields.utoronto.ca/audio/05-06/percolation_SLE/sheffield3, 2005.

[She07] Scott Sheffield. Gaussian free fields for mathematicians. *Probab. Theory Related Fields*, 139(3-4):521–541, 2007. math/0312099.

[She09a] Scott Sheffield. Exploration trees and conformal loop ensembles. *Duke Math. J.*, 147(1):79–129, 2009. math/0609167.

[She09b] Scott Sheffield. Exploration trees and conformal loop ensembles. *Duke Math. J.*, 147(1):79–129, 2009. math/0609167.

[She09c] Scott Sheffield. Exploration trees and conformal loop ensembles. *Duke Math. J.*, 147(1):79–129, 2009. math/0609167.

[She10] S. Sheffield. Conformal weldings of random surfaces: SLE and the quantum gravity zipper. *ArXiv e-prints*, December 2010.

[She11a] S. Sheffield. Quantum gravity and inventory accumulation. *ArXiv e-prints*, August 2011.

[She11b] S. Sheffield. Quantum gravity and inventory accumulation. *ArXiv e-prints*, August 2011. To appear in Annals of Probability.

[She11c] S. Sheffield. Quantum gravity and inventory accumulation. *ArXiv e-prints*, August 2011.

[She15] S. Sheffield. Conformal weldings of random surfaces: SLE and the quantum gravity zipper. *ArXiv e-prints*, December 2015. To appear in Annals of Probability.

[Shi85] Michio Shimura. Excursions in a cone for two-dimensional Brownian motion. *J. Math. Kyoto Univ.*, 25(3):433–443, 1985.

[Smi] Stanislav Smirnov. Private communciation.

[Smi01a] Stanislav Smirnov. Critical percolation in the plane: conformal invariance, Cardy's formula, scaling limits. *C. R. Acad. Sci. Paris Sér. I Math.*, 333(3):239–244, 2001.

[Smi01b] Stanislav Smirnov. Critical percolation in the plane: conformal invariance, Cardy's formula, scaling limits. *C. R. Acad. Sci. Paris Sér. I Math.*, 333(3):239–244, 2001.

[Smi01c] Stanislav Smirnov. Critical percolation in the plane: conformal invariance, Cardy's formula, scaling limits. *C. R. Acad. Sci. Paris Sér. I Math.*, 333(3):239–244, 2001. 0909.4499.

[Smi05] Stanislav Smirnov. Critical percolation and conformal invariance. In *XIVth International Congress on Mathematical Physics*, pages 99–112. World Sci. Publ., Hackensack, NJ, 2005.

[Smi06] Stanislav Smirnov. Towards conformal invariance of 2D lattice models. In *International Congress of Mathematicians. Vol. II*, pages 1421–1451. Eur. Math. Soc., Zürich, 2006.

[Smi07] Stanislav Smirnov. Conformal invariance in random cluster models. I. Holomorphic fermions in the Ising model. *Ann. of Math.*, 2007.

[Smi10a] Stanislav Smirnov. Conformal invariance in random cluster models. I. Holomorphic fermions in the Ising model. *Ann. of Math. (2)*, 172(2):1435–1467, 2010.

[Smi10b] Stanislav Smirnov. Conformal invariance in random cluster models. I. Holomorphic fermions in the Ising model. *Ann. of Math. (2)*, 172(2):1435–1467, 2010. 0708.0039.

[SS05a] Oded Schramm and Scott Sheffield. Harmonic explorer and its convergence to SLE_4. *Ann. Probab.*, 33(6):2127–2148, 2005.

[SS05b] Oded Schramm and Scott Sheffield. Harmonic explorer and its convergence to SLE_4. *Ann. Probab.*, 33(6):2127–2148, 2005.

[SS05c] Oded Schramm and Scott Sheffield. Harmonic explorer and its convergence to SLE_4. *Ann. Probab.*, 33(6):2127–2148, 2005. math/0310210.

[SS09a] Oded Schramm and Scott Sheffield. Contour lines of the two-dimensional discrete Gaussian free field. *Acta Math.*, 202(1):21–137, 2009.

[SS09b] Oded Schramm and Scott Sheffield. Contour lines of the two-dimensional discrete Gaussian free field. *Acta Math.*, 202(1):21–137, 2009.

[SS09c] Oded Schramm and Scott Sheffield. Contour lines of the two-dimensional discrete Gaussian free field. *Acta Math.*, 202(1):21–137, 2009.

[SS09d] Oded Schramm and Scott Sheffield. Contour lines of the two-dimensional discrete Gaussian free field. *Acta Math.*, 202(1):21–137, 2009. math/0605337.

[SS10] Oded Schramm and Scott Sheffield. A contour line of the continuum Gaussian free field. *ArXiv e-prints*, August 2010.

[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. 1008.2447.

[SSW09] Oded Schramm, Scott Sheffield, and David B. Wilson. Conformal radii for conformal loop ensembles. *Comm. Math. Phys.*, 288(1):43–53, 2009.

[Ste10] K Stephenson. CirclePack. `http://www.math.utk.edu/~kens/CirclePack/`, 1992–2010.

[SW05] Oded Schramm and David B. Wilson. SLE coordinate changes. *New York J. Math.*, 11:659–669 (electronic), 2005. math/0505368.

[SW12] Scott Sheffield and Wendelin Werner. Conformal loop ensembles: the Markovian characterization and the loop-soup construction. *Ann. of Math. (2)*, 176(3):1827–1917, 2012. 1006.2374.

[Tut62] W. T. Tutte. A census of planar triangulations. *Canad. J. Math.*, 14:21–38, 1962.

[Tut68a] W. T. Tutte. On the enumeration of planar maps. *Bull. Amer. Math. Soc.*, 74:64–74, 1968.

[Tut68b] W. T. Tutte. On the enumeration of planar maps. *Bull. Amer. Math. Soc.*, 74:64–74, 1968.

[VAW90] Mohammad Q. Vahidi-Asl and John C. Wierman. First-passage percolation on the Voronoĭ tessellation and Delaunay triangulation. In *Random graphs '87 (Poznań, 1987)*, pages 341–359. Wiley, Chichester, 1990.

[VAW92] Mohammad Q. Vahidi-Asl and John C. Wierman. A shape result for first-passage percolation on the Voronoĭ tessellation and Delaunay triangulation. In *Random graphs, Vol. 2 (Poznań, 1989)*, Wiley-Intersci. Publ., pages 247–262. Wiley, New York, 1992.

[Wat93] Yoshiyuki Watabiki. Analytic study of fractal structure of quantized surface in two-dimensional quantum gravity. *Progr. Theoret. Phys. Suppl.*, (114):1–17, 1993. Quantum gravity (Kyoto, 1992).

[Wer03] Wendelin Werner. Random planar curves and schramm-loewner evolutions. *arXiv preprint math/0303354*, 2003.

[Wer04a] Wendelin Werner. Random planar curves and Schramm-Loewner evolutions. In *Lectures on probability theory and statistics*, volume 1840 of *Lecture Notes in Math.*, pages 107–195. Springer, Berlin, 2004. math/0303354.

[Wer04b] Wendelin Werner. Random planar curves and Schramm-Loewner evolutions. In *Lectures on probability theory and statistics*, volume 1840 of *Lecture Notes in Math.*, pages 107–195. Springer, Berlin, 2004. math/0303354.

[Wil96] David Bruce Wilson. Generating random spanning trees fmore quickly than the cover time. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 296–303. ACM, 1996.

[WJS81] TA Witten Jr and Leonard M Sander. Diffusion-limited aggregation, a kinetic critical phenomenon. *Physical review letters*, 47(19):1400, 1981.

[WS83] T. A. Witten and L. M. Sander. Diffusion-limited aggregation. *Phys. Rev. B (3)*, 27(9):5686–5697, 1983.

[YY08] A. Yadin and A. Yehudayoff. Loop-Erased Random Walk and Poisson Kernel on Planar Graphs. *ArXiv e-prints*, September 2008.

[Zha08a] Dapeng Zhan. Duality of chordal SLE. *Invent. Math.*, 174(2):309–353, 2008.

[Zha08b] Dapeng Zhan. Reversibility of chordal SLE. *Ann. Probab.*, 36(4):1472–1494, 2008. 0808.3649.

[Zha10] Dapeng Zhan. Duality of chordal SLE, II. *Ann. Inst. Henri Poincaré Probab. Stat.*, 46(3):740–759, 2010.

Department of Mathematics
Massachusetts Institute of Technology
Cambridge, MA, USA