

A contour line of the continuum Gaussian free field

Oded Schramm Scott Sheffield

Abstract

Consider an instance h of the Gaussian free field on a simply connected planar domain D with boundary conditions $-\lambda$ on one boundary arc and λ on the complementary arc, where λ is the special constant $\sqrt{\pi/8}$. We argue that even though h is defined only as a random distribution, and not as a function, it has a well-defined zero level line γ connecting the endpoints of these arcs, and the law of γ is SLE(4). We construct γ in two ways: as the limit of the chordal zero contour lines of the projections of h onto certain spaces of piecewise linear functions, and as the only path-valued function on the space of distributions with a natural Markov property.

We also show that, as a function of h , γ is “local” (it does not change when h is modified away from γ) and derive some general properties of local sets.

Contents

1	Introduction	3
1.1	Overview	3
1.2	Notation and definitions	5
1.3	GFF approximations and discrete level lines	9
1.4	Main results	11
2	Coupling SLE and the GFF	13
2.1	GFF on subdomains: projections and restrictions	13
2.2	Constructing the coupling	15
3	Local sets	21
3.1	Absolute continuity	21
3.2	Local sets for discrete fields	23
3.3	Local sets for the GFF	26
4	Fine grid preliminaries	30
4.1	Subspaces are asymptotically dense	30
4.2	Topological and measure theoretic preliminaries	32
4.3	Limits of discrete local sets are local	34
4.4	Statement of the height gap lemma	36
5	Proofs of main results	37
6	Remark on other contour lines	42

1 Introduction

1.1 Overview

The two dimensional Gaussian free field (GFF) is an object of central importance in mathematics and physics. It appears frequently as a model for random surfaces and height interfaces and as a tool for studying two-dimensional statistical physics models that are not obviously random surfaces (e.g., Ising and Potts models, $O(n)$ loop models). It appears in random matrix theory (e.g. [RV07]), as a random electrostatic potential in Coulomb gas theory, and as the starting point for many constructions in conformal field theory and string theory. It is related to the Schramm-Loewner evolution (SLE) in a range of ways [She05, Dub09, IK10, HBB10] (see Section 1.4) and it represents the logarithmic conformal metric distortion in critical Liouville quantum gravity (see details and references in [DS08]).

This paper is a sequel to another paper by the current authors [SS09], which contains a much more detailed overview of the topics mentioned above, and many more references. That paper studied the discrete Gaussian free field (DGFF), which is a random function on a graph that (when defined on increasingly fine lattices) has the GFF as a scaling limit. The authors showed in [SS09] that a certain level line of the DGFF has SLE(4) as a scaling limit.

More precisely, when the DGFF is defined on triangulated two-dimensional lattice graphs (and interpolated linearly on each triangle to give a continuous function on a domain in \mathbb{R}^2), with positive boundary conditions on one boundary arc and negative boundary conditions on the complementary arc, the chordal contour line connecting the two endpoints of these arcs converges in law (as the lattice size gets finer) to a variant of SLE(4). In particular, there is a special constant $\lambda > 0$ such that if boundary conditions are set to $\pm\lambda$ on the two boundary arcs, then the zero chordal contour line converges to SLE(4) itself as the lattice size tends to zero. The exact value of λ was not determined in [SS09], but we will determine it here. Specifically, we will show that if the DGFF is scaled in such a way that its fine-mesh limit is the continuum Gaussian free field (GFF), defined by the Dirichlet norm $\int |\nabla\phi|^2$, then $\lambda = \sqrt{\pi/8}$.

It was observed in [SS09] that one can project an instance h of the GFF on a simply connected planar domain onto a sequence of subspaces to obtain a sequence of successively closer continuous and piecewise linear approximations to h , each of which has the law of a discrete Gaussian free field. Al-

though h is defined only as a distribution, and not as a continuous function, one might hope to define the “contour lines” of h as the limits of the contour lines of these approximations. The work in [SS09] implies that (when the boundary conditions are $\pm\lambda$ on complementary arcs) the zero chordal contour lines of the approximations converge *in law* to SLE(4). Our goal is to strengthen these results and show that they converge *in probability* to a path-valued function γ of h whose law is SLE(4). We will also characterize γ directly by showing that it is the unique path-valued function of h with a certain Markov property. In Section 3, we also show that, as a function of h , γ is “local.” Very roughly speaking, this means that γ can be determined without observing the values of h at points away from γ — and thus, γ is not changed if h is modified away from γ . We will also give a discrete definition of local and show that both discrete and continuum local sets have nice properties.

The reader may consult [SS09] for background and a historical introduction to the contour line problem for discrete and continuum Gaussian free fields. We will forego such an introduction here and proceed directly to the notation and the main results.

Acknowledgments. We wish to thank Vincent Beffara, Richard Kenyon, Jané Kondev, Julien Dubédat, and David Wilson for inspiring and useful conversations. In particular, Kenyon and the first author worked out a coupling of the GFF and SLE(8) early on, which is quite similar to the coupling between the GFF and SLE(4) presented here. We also thank Jason Miller for helpful comments on a draft of this paper.

This paper was mostly finished well before the tragic 2008 death of the first author, Oded Schramm. The results were originally part of our first draft of [SS09] (begun in 2003) and were separated once it became apparent that [SS09] was becoming quite long for a journal paper. We presented them informally in slides and online talks (e.g. [She05]) and mostly finished the writing together. The completion of this project is a somewhat melancholy occasion for the second author, as it concludes a long and enjoyable collaboration with an inspiring mathematician and a wonderful friend. It is also another occasion to celebrate the memory of Oded Schramm.

1.2 Notation and definitions

Let D be a planar domain (i.e., a connected open subset of \mathbb{R}^2 , which we sometimes identify with the complex plane \mathbb{C}). Assume further that D is a subset of a simply connected domain that is not all of \mathbb{C} . (In particular, this implies that D can be mapped conformally to a subset of the unit disc.)

For each $x \in D$, the point in ∂D at which a Brownian motion started at x first exits D is a random variable whose law we denote by ν_x and call the **harmonic measure of ∂D viewed from x** . It is not hard to show that for any $y \in D$ the measures ν_x and ν_y are absolutely continuous with respect to one another, with a Radon Nikodym derivative bounded between two positive constants. In particular, this implies that if a function $f_\partial : \partial D \rightarrow \mathbb{R}$ lies in $L^1(\nu_x)$ then it lies in $L^1(\nu_y)$ for any $y \in D$. We refer to a function with this property as a **harmonically integrable** function of ∂D . If $f_\partial : \partial D \rightarrow \mathbb{R}$ is harmonically integrable, then the function

$$f(x) = \int f_\partial(y) d\nu_x(y),$$

defined on D , is called the **harmonic extension** of f_∂ to D .

Given a planar domain D and $f, g \in L^2(D)$, we denote by (f, g) the inner product $\int f(x)g(x)dx$ (where dx is Lebesgue measure on D). Let $H_s(D)$ denote the space of real-valued C^∞ , compactly supported functions on D . We also write $H(D)$ (a.k.a., $H_0^1(D)$ — a Sobolev space) for the Hilbert space completion of $H_s(D)$ under the **Dirichlet inner product**

$$(f, g)_\nabla := \int_D \nabla f \cdot \nabla g \, dx.$$

We write $\|f\| := (f, f)^{1/2}$ and $\|f\|_\nabla := (f, f)_\nabla^{1/2}$. If $f, g \in H_s(D)$, then integration by parts gives $(f, g)_\nabla = (f, -\Delta g)$.

Let ϕ be a conformal map from D to another domain D' . Then an elementary change of variables calculation shows that

$$\int_{D'} \nabla(f_1 \circ \phi^{-1}) \cdot \nabla(f_2 \circ \phi^{-1}) \, dx = \int_D (\nabla f_1 \cdot \nabla f_2) \, dx.$$

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

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 $\phi_k \rightarrow 0$

in $H_s(D)$ if and only if there is a compact set on which all of the ϕ_k are supported and the m th 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 $\mathbf{p} \mapsto (h, \mathbf{p})$. We will frequently abuse notation and use h — or more precisely the map denoted by $\mathbf{p} \rightarrow (h, \mathbf{p})$ — to represent a general distribution (which is a functional of \mathbf{p}), even though h may not correspond to an element in $L^2(D)$. We define partial derivatives and integrals of distributions in the usual way (via integration by parts), i.e., $(\frac{\partial}{\partial x} h, \mathbf{p}) := -(h, \frac{\partial}{\partial x} \mathbf{p})$; in particular, if h is a distribution then Δh is a distribution defined by $(\Delta h, \mathbf{p}) := (h, \Delta \mathbf{p})$. When h is a distribution and $g \in H_s(D)$, we also write

$$(h, g)_{\nabla} := (-\Delta h, g) = (h, -\Delta g).$$

When $x \in D$ is fixed, we let $\tilde{G}_x(y)$ be the harmonic extension to D of the function on ∂D given by $(2\pi)^{-1} \log |y - x|$. (It is not hard to see that this function is harmonically integrable.) Then **Green's function in the domain D** is defined by

$$G(x, y) = (2\pi)^{-1} \log |y - x| - \tilde{G}_x(y).$$

When $x \in D$ is fixed, Green's function may be viewed as a distributional solution of $\Delta G(x, \cdot) = -\delta_x(\cdot)$ with zero boundary conditions. Note, however, that under the above definition it is not generally true that $G(x, \cdot)$ extends continuously to a function on \bar{D} that vanishes on ∂D (e.g., if ∂D contains isolated points, then $G(x, \cdot)$ need not tend to zero at those points) although this will be the case if D is simply connected. For any $\mathbf{p} \in H_s(D)$, we write $-\Delta^{-1}\mathbf{p}$ for the function $\int_D G(\cdot, y) \mathbf{p}(y) dy$. This is a C^∞ (though not necessarily compactly supported) function in D whose Laplacian is $-\mathbf{p}$.

If $f_1 = -\Delta^{-1}\mathbf{p}_1$ and $f_2 = -\Delta^{-1}\mathbf{p}_2$, then integration by parts gives $(f_1, f_2)_{\nabla} = (\mathbf{p}_1, -\Delta^{-1}\mathbf{p}_2)$. By the definition of $-\Delta^{-1}\mathbf{p}_2$ above, the latter expression may be rewritten as

$$\int_{D \times D} \mathbf{p}_1(x) \mathbf{p}_2(y) G(x, y) dx dy, \tag{1.1}$$

where $G(x, y)$ is Green's function in D . Another way to say this is that, since $\Delta G(x, \cdot) = -\delta_x(\cdot)$, integration by parts gives $\int_D G(x, y) \mathbf{p}_2(y) dy =$

$-\Delta^{-1}\mathbf{p}_2(x)$, and we obtain (1.1) by multiplying each side by $\mathbf{p}_1(x)$ and integrating with respect to x .

We next observe that every $h \in H(D)$ is naturally a distribution, since we may define the map (h, \cdot) by $(h, \mathbf{p}) := (h, -\Delta^{-1}\mathbf{p})_{\nabla}$. (It is not hard to see that $-\Delta^{-1}\mathbf{p} \in H(D)$, since its Dirichlet energy is given explicitly by (1.1).)

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 $\sum |\alpha_j|^2$ is a.s. infinite). However, it does converge almost surely within the space of distributions — that is, the limit $(\sum_{i=1}^{\infty} \alpha_j f_j, \mathbf{p})$ almost surely exists for all $\mathbf{p} \in H_s(D)$, and the limiting value as a function of \mathbf{p} is almost surely a continuous functional on $H_s(D)$ [She03]. We view h as a sample from the measure space (Ω, \mathcal{F}) where $\Omega = \Omega_D$ is the set of distributions on D and \mathcal{F} is the smallest σ -algebra that makes (h, \mathbf{p}) measurable for each $\mathbf{p} \in H_s(D)$, and we denote by $\mu = \mu_D$ the probability measure which is the law of h . The following is a standard and straightforward result about Gaussian processes [She03]:

Proposition 1.1. *The GFF with zero boundary conditions on D is the only random distribution h on D with the property that for each $\mathbf{p} \in H_s(D)$ the random variable $(h, \mathbf{p}) = (h, -\Delta^{-1}\mathbf{p})_{\nabla}$ is a mean zero Gaussian with variance $(-\Delta^{-1}\mathbf{p}, -\Delta^{-1}\mathbf{p})_{\nabla} = (\mathbf{p}, -\Delta^{-1}\mathbf{p})$. In particular, the law of h is independent of the choice of basis $\{f_j\}$ in the definition above.*

Given a harmonically integrable function $h_{\partial} : \partial D \rightarrow \mathbb{R}$, the **GFF with boundary conditions** h_{∂} is the random distribution whose law is that of the GFF with zero boundary conditions plus the deterministic function \underline{h}_{∂} which is the harmonic interpolation of h_{∂} to D (viewing \underline{h}_{∂} as a distribution).

Suppose $g \in H_s(D)$ and write $\mathbf{p} := -\Delta g$. If h is a GFF with boundary conditions h_{∂} (and \underline{h}_{∂} is the harmonic extension of h_{∂} to D), then we may integrate by parts and write $(h, \mathbf{p}) = (\underline{h}_{\partial}, \mathbf{p}) + (h, g)_{\nabla}$. Note that $(\underline{h}_{\partial}, \mathbf{p})$ is deterministic while the latter term, $(h, g)_{\nabla}$ does not depend on h_{∂} . Then the random variables (h, \mathbf{p}) , for $\mathbf{p} \in \Delta H_s(D)$, have means $(\underline{h}_{\partial}, \mathbf{p})$ and covariances given by

$$\text{Cov}((h, \mathbf{p}_1), (h, \mathbf{p}_2)) = (\mathbf{p}_1, -\Delta^{-1}\mathbf{p}_2) = \int_{D \times D} \mathbf{p}_1(x) \mathbf{p}_2(y) G(x, y) dx dy,$$

where $G(x, y)$ is Green's function in D .

Green's function in the upper half plane is $G(x, y) = (2\pi)^{-1} \log \left| \frac{\bar{x}-y}{x-y} \right|$ (where \bar{x} is the complex conjugate of x), since the right hand side is zero when $y \in \mathbb{R}$ and is the sum of $-(2\pi)^{-1} \log |x-y|$ and a function harmonic in y . (Note that $G(x, y) > 0$ for all $x, y \in \mathbb{H}$.) In physical Coulomb gas models, when \mathbf{p} represents the density function of a (signed) electrostatic charge distribution, the quantity $(\mathbf{p}, -\Delta^{-1}\mathbf{p}) = \int_{D \times D} \mathbf{p}(x) \mathbf{p}(y) G(x, y) dx dy$ is sometimes called the **electrostatic potential energy** or **energy of assembly** of \mathbf{p} (assuming the system is grounded at the boundary of D).

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 the distribution \tilde{h} on \tilde{D} for which $(\tilde{h}, \tilde{\mathbf{p}}) = (h, \mathbf{p})$ whenever $\mathbf{p} \in H_s(D)$ and $\tilde{\mathbf{p}} = |\phi'|^{-2} \mathbf{p} \circ \phi^{-1}$, where ϕ' is the complex derivative of ϕ (viewing the latter as an analytic function on a subset of \mathbb{C}).

Note that if \mathbf{p} is interpreted in the physical sense as an electrostatic charge density on D and ϕ is a change-of-coordinates map, then $\tilde{\mathbf{p}}$ is the corresponding charge density on \tilde{D} . The reader may also observe that if h is a function in $H(D)$, interpreted as a distribution as discussed above, then the function $h \circ \phi^{-1} \in H(\tilde{D})$, interpreted as a distribution, agrees with the definition given above.

Also, if $f \in H(D)$ and $\mathbf{p} = -\Delta f \in H_s(D)$, then we have the following:

$$(h, f)_{\nabla} = (h, \mathbf{p}) = (h \circ \phi^{-1}, |\phi'|^{-2} \mathbf{p} \circ \phi^{-1}) = (h \circ \phi^{-1}, f \circ \phi^{-1})_{\nabla},$$

where the last equality follows from the simple calculus fact that

$$-\Delta(f \circ \phi^{-1}) = |\phi'|^{-2} \mathbf{p} \circ \phi^{-1}.$$

This and Proposition 1.1 imply the conformal invariance of the GFF: i.e., if h has the law of a GFF on D , then $h \circ \phi^{-1}$ has the law of a GFF on \tilde{D} .

In addition to defining an instance h of the GFF with zero boundary conditions as a random distribution, it is also occasionally useful to define the random variables $(h, g)_{\nabla}$ for all $g \in H(D)$ (and not merely $g \in H_s(D)$). Suppose that an orthonormal basis $\{f_j\}$ for $H(D)$, made up of elements in $H_s(D)$, is fixed and we write $\alpha_j = (h, f_j)_{\nabla}$. Then the α_j are i.i.d. Gaussians and for each fixed $g \in H(D)$ the sum

$$(h, g)_{\nabla} := \sum_{j=1}^{\infty} \alpha_j (f_j, g)_{\nabla}$$

converges almost surely. The collection of Gaussian random variables $(h, g)_\nabla$, with $g \in H(D)$ (defined as limits of partial sums) is a Hilbert space under the covariance inner product. Replacing $\{f_j\}$ with another orthonormal basis for $H(D)$ only changes the random variable $(h, g)_\nabla$ on a set of measure zero [She03]. Thus, for each fixed $g \in H(D)$, it makes sense to think of the random variable $(h, g)_\nabla$ as a real-valued function of Ω (which is canonically defined up to redefinition on a set of measure zero). This Hilbert space of random variables is a closed subspace of $L^2(\Omega, \mathcal{F}, \mu)$. The covariance of $(h, g_1)_\nabla$ and $(h, g_2)_\nabla$ is equal to $(g_1, g_2)_\nabla$, so this Hilbert space is naturally isomorphic to $H(D)$ [She03, Jan97]. (In general, a collection of centered Gaussian random variables which is a Hilbert space under the covariance inner product is called a **Gaussian Hilbert space** [Jan97].)

1.3 GFF approximations and discrete level lines

In order to construct contour lines of an instance h of the Gaussian free field, it will be useful to approximate h by a continuous function which is linear on each triangle in a triangular grid. We now describe a special case of the approximation scheme used in [SS09]. (The results in [SS09] also apply to more general periodic triangulations of \mathbb{R}^2 .) Let TG be the triangular grid in the plane $\mathbb{C} \cong \mathbb{R}^2$, i.e., the graph whose vertex set is the integer span of 1 and $e^{\pi i/3} = (1 + \sqrt{3}i)/2$, with straight edges joining v and w whenever $|v - w| = 1$. A **TG-domain** $D \subset \mathbb{R}^2$ is a domain whose boundary is a simple closed curve comprised of edges and vertices in TG .

Fix some $\lambda > 0$. Let D be any TG -domain, and let $\partial_+ \subset \partial D$ be an arc whose endpoints are distinct midpoints of edges of TG . Let V denote the vertices of TG in \overline{D} . Set $\partial_- := \partial D \setminus \partial_+$. Let $h_\partial : V \cap \partial D \rightarrow \mathbb{R}$ take the value $-\lambda$ on $\partial_- \cap V$ and λ on $\partial_+ \cap V$. Let ϕ_D be any conformal map from D to the upper half-plane \mathbb{H} that maps ∂_+ onto the positive real ray.

Let h^0 be an instance of the GFF on \mathbb{H} with zero boundary conditions. By conformal invariance of the Dirichlet inner product, $h^0 \circ \phi_D$ has the law of a GFF on D . Let $H_{TG}(D)$ be the subspace of $H(D)$ comprised of continuous functions that are affine on each TG triangle in D , and let h_D^0 be the orthogonal projection of h^0 onto $H_{TG}(D)$ with respect to the inner product $(\cdot, \cdot)_\nabla$. That is, h_D^0 is the random element of $H_{TG}(D)$ for which $(h^0, \cdot)_\nabla$ and $(h_D^0, \cdot)_\nabla$ are equivalent as linear functionals on $H_{TG}(D)$. (The former is defined almost surely from the Gaussian Hilbert space perspective, as discussed above; the given definition of this projection depends on the choice of basis $\{f_j\}$, but

changing the basis affects the definition only on a set of measure zero.)

We may view h_D^0 as a function on the vertices of $TG \cap \overline{D}$, linearly interpolated on each triangle [SS09]. Up to multiplicative constant, the law of h_D^0 is that of the discrete Gaussian free field on the graph $TG \cap \overline{D}$ with zero boundary conditions on the vertices in ∂D [SS09]. Let h_D be h_D^0 plus the (linearly interpolated) discrete harmonic (w.r.t the usual discrete Laplacian) extension of h_∂ to $V \cap D$. Thus, h_D is precisely the discrete field for which [SS09] proves that the chordal interface of $h_D \circ \phi_D^{-1}$ converges to SLE(4) for an appropriate choice of λ .

Now $h(z) = h^0(z) + \lambda(1 - 2\pi^{-1} \arg(z))$ is an instance of the GFF on \mathbb{H} with boundary conditions $-\lambda$ on the negative reals and λ on the positive reals, as $\lambda(1 - 2\pi^{-1} \arg(z))$ is the harmonic function with this boundary data. The functions $h_D \circ \phi_D^{-1}$ may be viewed as approximations to h .

There is almost surely a zero contour line γ_D (with no fixed parametrization) of h_D on D that connects the endpoints of ∂_- and ∂_+ , and $\hat{\gamma}_D := \phi_D \circ \gamma_D$ is a random path in \mathbb{H} connecting 0 to ∞ . We are interested in the limit of the paths $\hat{\gamma}_D$ as D gets larger. The correct sense of “large” is measured by

$$r_D := \text{rad}_{\phi_D^{-1}(i)}(D),$$

where $\text{rad}_x(D)$ denotes the radius of D viewed from x , i.e., $\inf_{y \notin D} |x - y|$. Of course, if $\phi_D^{-1}(i)$ is at bounded distance from ∂D , then the image of the triangular grid under ϕ_D is not fine near i , and there is no hope for approximating SLE(4) by $\hat{\gamma}_D$.

We have chosen to use \mathbb{H} as our canonical domain (mapping all other paths into \mathbb{H}), because it is the most convenient domain in which to define chordal SLE. When describing distance between paths, it is often more natural to use the disc or another bounded domain. To get the best of both worlds, we will endow \mathbb{H} with the metric it inherits from its conformal map onto the disc \mathbb{U} . Namely, we let $d_*(\cdot, \cdot)$ be the metric on $\overline{\mathbb{H}} \cup \{\infty\}$ given by $d_*(z, w) = |\Psi(z) - \Psi(w)|$, where $\Psi(z) := (z - i)/(z + i)$ maps $\overline{\mathbb{H}} \cup \{\infty\}$ onto $\overline{\mathbb{U}}$. (Here $\overline{\mathbb{H}}$ denotes the Euclidean closure of \mathbb{H} .) If $z \in \overline{\mathbb{H}}$, then $d_*(z_n, z) \rightarrow 0$ is equivalent to $|z_n - z| \rightarrow 0$, and $d_*(z_n, \infty) \rightarrow 0$ is equivalent to $|z_n| \rightarrow \infty$.

Let Λ be the set of continuous functions $W : [0, \infty) \rightarrow \mathbb{R}$ which satisfy $W_0 = 0$. We endow Λ with the topology of uniform convergence on compact intervals—i.e., the topology generated by the sets of the form $\{W : |W_t - V_t| < \epsilon \text{ when } 0 \leq t \leq T\}$ for some $V_t \in \Lambda$ and $\epsilon > 0$ and $T > 0$. Let \mathcal{L} be the corresponding Borel σ -algebra.

For each $W \in \Lambda$, we may define the Loewner maps g_t from subsets of \mathbb{H} to \mathbb{H} via the ODE

$$\partial_t g_t(z) = \frac{2}{g_t(z) - W_t}, \quad g_0(z) = z. \quad (1.2)$$

When g_t^{-1} extends continuously to W_t , we write $\gamma(t) := g_t^{-1}(W_t)$. Generally, we write $\tau(z)$ to be the largest t for which the ODE describing $g_t(z)$ is defined, and $K_t = \{z \in \overline{\mathbb{H}} : \tau(z) \leq t\}$. So K_t is a closed subset of $\overline{\mathbb{H}}$, and g_t is a conformal map from $\mathbb{H} \setminus K_t$ to \mathbb{H} .

Let $\Lambda_C \subset \Lambda$ be the set of W for which $\gamma : [0, \infty) \rightarrow \overline{\mathbb{H}}$ is well defined for all $t \in [0, \infty)$ and is a continuous path. The metric on paths given by $d_{\text{STRONG}}(\gamma_1, \gamma_2) = \sup_{t \geq 0} d_*(\gamma_1(t), \gamma_2(t))$ induces a corresponding metric on Λ_C .

1.4 Main results

The first of our two main results constructs a “zero contour line of h ” as the limit of the contour lines of the approximations described above:

Theorem 1.2. *Suppose that $\lambda = \sqrt{\pi/8}$. As $r_D \rightarrow \infty$, the random paths $\hat{\gamma}_D = \phi_D \circ \gamma_D$ described above, viewed as Λ_C -valued random variables on (Ω, \mathcal{F}) , converge in probability (with respect to the metric d_{STRONG} on Λ_C) to an (Ω, \mathcal{F}) -measurable random path $\gamma \in \Lambda_C$ that is distributed like SLE(4). In other words, for every $\epsilon > 0$ there is some $R = R(\epsilon)$ such that if $r_D > R$, then*

$$\mathbf{P} \left[d_{\text{STRONG}}(\gamma(t), \hat{\gamma}_D(t)) > \epsilon \right] < \epsilon.$$

The theorem shows, in particular, that the random path γ is a.s. determined by h . It will be called the zero contour line of h . We will actually initially construct γ in Section 2 in a way that does not involve discrete approximations. This construction is rather interesting. In some sense, we reverse the procedure. Rather than constructing γ from h , we start with a chordal SLE(4) path γ and define \tilde{h} as the field that is the GFF with boundary values λ in the connected component of $\mathbb{H} \setminus \gamma$ having \mathbb{R}_+ on its boundary plus a conditionally independent (given γ) GFF with boundary values $-\lambda$ in the other connected component of $\mathbb{H} \setminus \gamma$. We then show that \tilde{h} has the same law as h , effectively giving a coupling of h and γ . Much later, in Section 5, we will show that in this coupling γ is a.s. determined by h ; that is, it is

equal a.s. to a certain \mathcal{F} -measurable Λ_C -valued function of h . We will then prove Theorem 1.3 (stated below), which characterizes the zero contour line of h in terms of conditional expectations.

In order to make sense of the theorem stated below, we will need the fact that the GFF on a subdomain of \mathbb{H} makes sense as a random distribution on \mathbb{H} —and not just a random distribution on the subdomain. We will see in Section 2.2 that the GFF on a simply connected subdomain of \mathbb{H} is indeed canonically defined as a random distribution on \mathbb{H} .

Consider a random variable (\tilde{h}, \tilde{W}) in $(\Omega \times \Lambda, \mathcal{F} \times \mathcal{L})$. Let $\tilde{g}_t : \mathbb{H} \setminus \tilde{K}_t \rightarrow \mathbb{H}$ denote the Loewner evolution corresponding to \tilde{W} . Write

$$\tilde{h}_t(z) := \lambda \left(1 - 2\pi^{-1} \arg(\tilde{g}_t(z) - W_t) \right). \quad (1.3)$$

(This is the harmonic function on $\mathbb{H} \setminus \tilde{K}_t$ with boundary values $-\lambda$ on the left side of the tip of the Loewner evolution and $+\lambda$ on the right side.)

Theorem 1.3. *Let $D = \mathbb{H}$ and suppose that for some $\lambda > 0$ a random variable (\tilde{h}, \tilde{W}) in $(\Omega \times \Lambda, \mathcal{F} \times \mathcal{L})$ satisfies the following conformal Markov property. For every fixed $T \in [0, \infty)$ we have that given the restriction of \tilde{W}_t to $[0, T]$, a regular conditional law for \tilde{h} on $H_s(\mathbb{H} \setminus \tilde{K}_T)$ is given by \tilde{h}_T , as in (1.3), plus a zero boundary GFF on $\mathbb{H} \setminus \tilde{K}_T$. Then the following hold*

1. $\lambda = \sqrt{\pi/8}$.
2. *The trace of the Loewner evolution is almost surely a path $\tilde{\gamma}$ with the law of an SLE(4) (i.e., \tilde{W} is 2 times a standard Brownian motion).*
3. *Conditioned on \tilde{h} , the function \tilde{W} is almost surely completely determined (i.e., there exists a Λ -valued function on Ω such that \tilde{W} is almost surely equal to the value of that function applied to \tilde{h}).*
4. *The pair $(\tilde{h}, \tilde{\gamma})$ has the same law as the pair (h, γ) from Theorem 1.2.*

Another derivation of the SLE(κ)-GFF coupling in Theorem 1.3 appears in [Dub09], which references our work in progress and also explores relationships between this coupling and continuum partition functions, Laplacian determinants, and the Polyakov-Alvarez conformal anomaly formula. The couplings in [She05, Dub09] are in fact more general than those given here; they are defined for general κ , and are characterized by conformal Markov properties similar to those in Theorem 1.3. The SLE(κ) curves in these couplings are local sets (in the sense of Section 3) and are interpreted in [She05]

as “flow lines of e^{ih} ” where h is a multiple of the GFF and e^{ih} is viewed as a complex unit vector field (which is well defined when h is smooth, and in a certain sense definable even when h is a multiple of the GFF). Indeed, these examples were the primary motivation for our definition of local sets.

2 Coupling SLE and the GFF

2.1 GFF on subdomains: projections and restrictions

Proposition 2.1. *Suppose that $H'(D)$ is a closed subspace of $H(D)$, that P is the orthogonal projection onto that subspace, and that $\{f_j\}$ is an orthonormal basis for $H(D)$. Let α_j be i.i.d. mean zero, unit variance normal random variables. Then the sum $\sum P(\alpha_j f_j)$ converges almost surely in the space of distributions on D . The law of the limit is the unique law of a random distribution h on D with the property that for each $\mathbf{p} \in H_s(D)$, the random variable (h, \mathbf{p}) is a centered Gaussian with variance $\|P(-\Delta^{-1}\mathbf{p})\|_{\nabla}^2$.*

Proof. For each fixed $\mathbf{p} \in H_s(D)$, the fact that the partial sums of

$$\begin{aligned} \left(\sum P(\alpha_j f_j), \mathbf{p} \right) &= \left(\sum P(\alpha_j f_j), -\Delta^{-1}\mathbf{p} \right)_{\nabla} \\ &= \left(\sum \alpha_j f_j, P(-\Delta^{-1}\mathbf{p}) \right)_{\nabla} = \sum \alpha_j \left(f_j, P(-\Delta^{-1}\mathbf{p}) \right)_{\nabla} \end{aligned}$$

converge to a random variable with variance $\|P(-\Delta^{-1}\mathbf{p})\|_{\nabla}^2$ is immediate from the fact that $P(-\Delta^{-1}\mathbf{p}) \in H(D)$.

To complete the proof, we recall the approach of [She03] (which uses earlier work in [Gro67]) for showing that $\sum \alpha_j f_j$ converges in the space of distributions. Consider a norm $\|\cdot\|_*$ of the form $\|f\|_*^2 = (Tf, f)_{\nabla}$ where T is a **Hilbert-Schmidt operator** mapping $H(D)$ to itself—i.e., for each basis $\{f_j\}$ of $H(D)$ we have $\sum \|Tf_j\|_{\nabla}^2 < \infty$. (The sum is independent of the choice of basis.) Then Gross’s classical **abstract Wiener space** result implies that the closure of $H(D)$ with respect to the norm $\|\cdot\|_*$ is a space in which $\sum \alpha_j f_j$ converges almost surely [Gro67]. In [She03] a Hilbert-Schmidt operator T was given such that this closure was contained in the space of distributions (and such that convergence in the norm implied weak convergence in the space of distributions). For our purposes, it suffices to observe that if T is a Hilbert-Schmidt operator on $H(D)$ then $T' = PT$ is a Hilbert-Schmidt operator and $\|f\|'_* = (T'f, f)_{\nabla}$ is a measurable norm on

$H'(D)$. The existence and uniqueness of the random distribution described in the proposition statement then follows by the standard abstract Wiener space construction in [Gro67] as described in [She03]. \square

For any deterministic open subset B of D , we denote by Supp_B the closure in $H(D)$ of the space of C^∞ functions compactly supported on B and we denote by Harm_B the orthogonal complement of Supp_B , so that

$$H(D) = \text{Supp}_B \oplus \text{Harm}_B.$$

Note that since $(f, g)_\nabla = (-\Delta f, g)$, a smooth function $f \in H(D)$ satisfies $f \in \text{Harm}_B$ if and only if $-\Delta f = 0$ (so that f is harmonic) on B . In general, if f is any distribution on D , we say that f is **harmonic** on B if $(-\Delta f, g) := (f, -\Delta g) = 0$ for all $g \in \text{Supp}_B$. By this definition Harm_B consists of precisely of those elements of $H(D)$ which are harmonic on B . If f is a distribution which is harmonic on B , then the restriction of f to B may be viewed as an actual harmonic function—i.e., there is a harmonic function \tilde{f} on B such that for all $\mathbf{p} \in H_s(B)$ we have $(f, \mathbf{p}) = (\tilde{f}, \mathbf{p})$. We may construct \tilde{f} explicitly as follows. For each $z \in B$, write $\tilde{f}(z) = (f, \mathbf{p})$ where \mathbf{p} is any positive radially symmetric bump function centered at z whose integral is one and whose support lies in B . By harmonicity, this value is the same for any \mathbf{p} with these properties, since the difference of two such \mathbf{p} is the Laplacian of a (radially symmetric about z) function in $H_s(B)$. The fact that this function \tilde{f} is harmonic is easy to verify from its definition.

Suppose that B is a subdomain of a simply connected planar domain D . It will often be useful for us to interpret the GFF with zero boundary conditions on B as a random distribution on D . This means that we have to make sense of (h, \mathbf{p}) when $\mathbf{p} \in H_s(D)$ but we do not necessarily have $\mathbf{p} \in H_s(B)$. The following is immediate from Proposition 2.1, taking $H'(D)$ to be Supp_B .

Proposition 2.2. *Let B be a subdomain of a planar domain D . There is a unique law for a random distribution h_B on D with the property that for each $\mathbf{p} \in H_s(D)$ the random variable (h_B, \mathbf{p}) is a mean zero Gaussian with variance*

$$\int_{B \times B} \mathbf{p}(x) \mathbf{p}(y) G_B(x, y) dx dy.$$

If we restrict the pairing (h, \cdot) of the above proposition to functions on $H_s(B)$, then by definition h is the GFF on B with zero boundary conditions.

The projection of the GFF onto Harm_B is a random distribution which is almost surely harmonic on B . Applying Proposition 2.1 to Harm_B gives the following straightforward analog of Proposition 2.2. (Here Δ_B denotes the Laplacian restricted to B , and $\Delta_B^{-1}\mathbf{p}$ has Laplacian \mathbf{p} on B and zero boundary conditions on ∂B .)

Proposition 2.3. *Let B be a subdomain of a planar domain D . There is a unique random distribution h_B^* on D with the property that for each $\mathbf{p} \in H_s(D)$ the random variable (h_B^*, \mathbf{p}) is a mean zero Gaussian with variance*

$$(\mathbf{p}, -\Delta^{-1}\mathbf{p}) - (\mathbf{p}, -\Delta_B^{-1}\mathbf{p})_B. \quad (2.1)$$

An instance of the GFF on D may be written as $h = h_B + h_B^*$ where h_B is the zero boundary GFF on B and h_B^* and h_B are independent.

Although the above defines h_B^* as a random distribution and not as a function, we may, as discussed above, consider the restriction of h_B^* to B as a harmonic function. This function intuitively represents the conditional expectation of h at points of B given the values of h off of B .

2.2 Constructing the coupling

The overall goal of the paper is to recognize SLE(4) as an interface determined by the GFF. In this subsection we show how to start with an SLE(4) and use it to explicitly construct an instance of the GFF such that the resulting coupling of SLE(4) and the GFF satisfies the hypothesis of Theorem 1.3.

For any Loewner evolution W_t on the half plane \mathbb{H} we may write $f_t(z) := g_t(z) - W_t$. Since $\arg z$ is a harmonic function on \mathbb{H} with boundary values 0 on $(0, \infty)$ and π on $(-\infty, 0)$, the value $\pi^{-1} \arg f_t(z)$ is the probability that a two dimensional Brownian motion starting at z first exits $\mathbb{H} \setminus \gamma[0, t]$ either in $(-\infty, 0)$ or on the left hand side of $\gamma[0, t]$. In other words, for fixed t , the function $\pi^{-1} \arg f_t(z)$ is the bounded harmonic function on $\mathbb{H} \setminus \gamma[0, t]$ with boundary values given by 1 on the left side of the tip $\gamma(t)$ and 0 on the right side.

Lemma 2.4. *Let B_t be a standard 1-dimensional Brownian motion, and let $W_t = 2B_t$ be the driving parameter of an SLE(4) evolution g_t . Set $f_t(x) := g_t(z) - W_t$ and $h_t := \lambda(1 - 2\pi^{-1} \arg(f_t))$. Then for each fixed $z \in \mathbb{H}$,*

$$dh_t(z) = \frac{4\lambda}{\pi} \text{Im}(f_t(z)^{-1}) dB_t, \quad (2.2)$$

(in the sense of Itô differentials) and $h_t(z)$ is a martingale.

The function $\frac{4\lambda}{\pi}\text{Im}(f_t(z)^{-1})$ is significant. At time $t = 0$, it is a negative harmonic function whose level sets are circles in \mathbb{H} that are tangent to \mathbb{R} at $W_0 = 0$. Intuitively, it represents the harmonic measure (times a negative constant) of the tip of K_t as seen from the point z . When W_t moves an infinitesimal amount to the left or right, h_t changes by an infinitesimal multiple of this function.

Proof. If $dW_t = \sqrt{\kappa} dB_t$, then the Itô derivatives of f_t and $\log f_t$ are as follows:

$$df_t = \frac{2}{f_t} dt - dW_t,$$

$$d \log f_t = 2 f_t^{-2} dt - f_t^{-1} dW_t - \frac{\kappa}{2} f_t^{-2} dt = \frac{(4 - \kappa)}{2 f_t^2} dt - f_t^{-1} dW_t.$$

It is a special feature of $\kappa = 4$ that $d \log f_t = -2 f_t^{-1} dB_t$; hence $h_t(z)$ is a local martingale and (since it is bounded) a martingale. \square

In what follows we use bracket notation $\langle X_t, Y_t \rangle := \langle X, Y \rangle_t$ to denote the cross-variation product of time-varying processes X and Y up to time t , i.e.,

$$\langle X, Y \rangle_t := \lim \sum_{i=1}^k (X_{s_i} - X_{s_{i-1}})(Y_{s_i} - Y_{s_{i-1}}),$$

where the limit is taken over increasingly dense finite sequences $s_0 = 0 < s_1 < \dots < s_k = t$. Sometimes, we find it convenient to write $\langle X_t, Y_t \rangle$ in place of $\langle X, Y \rangle_t$. In particular, we have almost surely $\langle B_t, B_t \rangle = t$ and $\langle B_t, t \rangle = \langle t, B_t \rangle = \langle t, t \rangle = 0$ for all $t \geq 0$.

Lemma 2.5. *Suppose that W_t is the driving parameter of an SLE(4) and*

$$h_t := \lambda (1 - 2\pi^{-1} \arg(f_t))$$

(as in Lemma 2.4). Let $G(x, y) = (2\pi)^{-1} \log \left| \frac{x-\bar{y}}{x-y} \right|$ be Green's function in the upper half plane and write $G_t(x, y) = G(f_t(x), f_t(y))$ when x and y are both in $\mathbb{H} \setminus K_t$. If $x, y \in \mathbb{H}$ are fixed and $\lambda = \sqrt{\pi/8}$, then the following holds for all t almost surely:

$$dG_t(x, y) = -d\langle h_t(x), h_t(y) \rangle. \tag{2.3}$$

Recall that for $\kappa = 4$ we have $x \notin \bigcup_{t>0} K_t$ a.s. for every $x \in \mathbb{H}$ [RS05]. Therefore, $G_t(x, y)$ is a.s. well defined for all t .

Proof. We first claim that

$$2\pi dG_t(x, y) = -4 \operatorname{Im} (f_t(x)^{-1}) \operatorname{Im} (f_t(y)^{-1}) dt. \quad (2.4)$$

We derive (2.4) explicitly using Itô calculus as follows:

$$\begin{aligned} 2\pi dG_t(x, y) &= -d \operatorname{Re} \log[g_t(x) - g_t(y)] + d \operatorname{Re} \log[g_t(x) - \overline{g_t(y)}] \\ &= -2 \operatorname{Re} \frac{f_t(x)^{-1} - f_t(y)^{-1}}{g_t(x) - g_t(y)} dt + \\ &\quad 2 \operatorname{Re} \frac{(g_t(x) - W_t)^{-1} - (\overline{g_t(y)} - W_t)^{-1}}{g_t(x) - \overline{g_t(y)}} dt \\ &= 2 \operatorname{Re} (f_t(x)^{-1} f_t(y)^{-1}) dt - 2 \operatorname{Re} (f_t(x)^{-1} (\overline{f_t(y)})^{-1}) dt \\ &= 4 \operatorname{Re} (i f_t(x)^{-1} \operatorname{Im} [f_t(y)^{-1}]) dt \\ &= -4 \operatorname{Im} (f_t(x)^{-1}) \operatorname{Im} (f_t(y)^{-1}) dt. \end{aligned}$$

Using (2.2), we have $d\langle h_t(x), h_t(y) \rangle = \left(\frac{4\lambda}{\pi}\right)^2 \operatorname{Im} (f_t(x)^{-1}) \operatorname{Im} (f_t(y)^{-1}) dt$. Setting $\lambda = \sqrt{\pi/8}$, we obtain (2.3). \square

Next fix some $\mathbf{p} \in H_s(\mathbb{H})$ and write $E_t(\mathbf{p}) := \int G_t(x, y) \mathbf{p}(x) \mathbf{p}(y) dx dy$ for the electrostatic potential energy of \mathbf{p} in $\mathbb{H} \setminus K_t$. For each t , the function h_t is not well defined on all of \mathbb{H} (since it is not defined on K_t), but it is defined (and harmonic and bounded between $-\lambda$ and λ) almost everywhere almost surely. In particular, when $\mathbf{p} \in H_s(\mathbb{H})$, the integral $\langle h_t, \mathbf{p} \rangle = \int_{\mathbb{H}} h_t(x) \mathbf{p}(x) dx$ is well defined, so we may view h_t as a distribution.

Lemma 2.6. *In the setting of Lemma 2.5, assume $\lambda = \sqrt{\pi/8}$. If $\mathbf{p} \in H_s(\mathbb{H})$, then (h_t, \mathbf{p}) is a martingale. Moreover,*

$$d\langle (h_t, \mathbf{p}), (h_t, \mathbf{p}) \rangle = -d \left(\int \mathbf{p}(x) \mathbf{p}(y) G_t(x, y) dx dy \right) = -dE_t(\mathbf{p}). \quad (2.5)$$

In other words, (h_t, \mathbf{p}) is a Brownian motion when parameterized by minus the electrostatic potential energy of \mathbf{p} in $\mathbb{H} \setminus K_t$. More generally, when $\mathbf{p}_1, \mathbf{p}_2 \in H_s(D)$, we have

$$d\langle (h_t, \mathbf{p}_1), (h_t, \mathbf{p}_2) \rangle = -d \left(\int \mathbf{p}_1(x) \mathbf{p}_2(y) G_t(x, y) dx dy \right). \quad (2.6)$$

Proof. The right equality in (2.5) is true by definition, so it is enough to prove (2.6). Since both sides of (2.6) are bilinear in \mathbf{p}_1 and \mathbf{p}_2 , we lose no generality in assuming that \mathbf{p}_1 and \mathbf{p}_2 are non-negative. (Note that any $\mathbf{p} \in H_s(D)$ can be written $\mathbf{p}_1 - \mathbf{p}_2$ for some non-negative $\mathbf{p}_1, \mathbf{p}_2 \in H_s(D)$; simply choose \mathbf{p}_1 to be any non-negative element of $H_s(D)$ which dominates \mathbf{p} and set $\mathbf{p}_2 = \mathbf{p}_1 - \mathbf{p}$.)

Jason Miller has pointed out in private communication a very simple way to obtain (2.6). Since (h_t, \mathbf{p}_i) are continuous martingales, the (non-decreasing) process on the LHS of (2.6) is characterized by the fact that

$$(h_t, \mathbf{p}_1)(h_t, \mathbf{p}_2) - \langle (h_t, \mathbf{p}_1), (h_t, \mathbf{p}_2) \rangle$$

is a martingale. Plugging in the RHS, we have now only to show that

$$(h_t, \mathbf{p}_1)(h_t, \mathbf{p}_2) + d \left(\int \mathbf{p}_1(x) \mathbf{p}_2(y) G_t(x, y) dx dy \right) \quad (2.7)$$

is a martingale. By (2.3)

$$h_t(x)h_t(y) + G_t(x, y) dx dy$$

is a martingale for fixed x and y in \mathbb{H} . These martingales are uniformly bounded above for x and y in the support of the \mathbf{p}_i (since $G_t(x, y)$ is non-increasing and $h_t(\cdot)$ is bounded). Since (2.7) is a weighted average of these martingales, Fubini's theorem implies that (2.7) is itself a martingale.

Our original argument invoked a stochastic Fubini theorem (we used the one in [Pro90, §IV.4]) and was longer and less self contained. \square

Now, define $h_\infty(z) = \lim_{t \rightarrow \infty} h_t(z)$, $G_\infty(x, y) = \lim_{t \rightarrow \infty} G_t(x, y)$, and $E_\infty(\mathbf{p}) = \lim_{t \rightarrow \infty} E_t(\mathbf{p})$. (The limit exists almost surely for fixed z since $h_t(z)$ is a bounded martingale; it can also be deduced for all z from the continuity of the SLE trace.) The reader may check that for fixed x , $h_\infty(x)$ is almost surely $\pm\lambda$ depending on whether x is to the left or right of γ .

Similarly, since $G_t(x, y)$ and $E_t(\mathbf{p})$ are decreasing functions of t , these limits also exist almost surely. The statement of the following lemma makes use of these definitions and implicitly assumes Proposition 2.2 (namely, the fact that the zero boundary GFF on an arbitrary subdomain of \mathbb{H} has a canonical definition as a random distribution on \mathbb{H}).

Lemma 2.7. *Assume the setting of Lemma 2.5 and $\lambda = \sqrt{\pi/8}$. Let \tilde{h} be equal to h_∞ (as defined above) plus a sum of independent zero-boundary*

GFF's, one in each component of $\mathbb{H} \setminus \gamma$. Then the law of \tilde{h} is that of a GFF in \mathbb{H} with boundary conditions $-\lambda$ and λ on the negative and positive real axes. In fact, the pair (\tilde{h}, W) constructed in this way (where W is the Loewner driving parameter of γ) satisfies the hypothesis of Theorem 1.3.

Proof. For each $\mathbf{p} \in H_s(D)$ the random variable (\tilde{h}, \mathbf{p}) is a sum of $(\underline{h}_\partial, \mathbf{p})$, a Brownian motion started at time zero and stopped at time $E_0(\mathbf{p}) - E_\infty(\mathbf{p})$, and a Gaussian of variance $E_\infty(\mathbf{p})$. Thus its law is Gaussian with mean $(\underline{h}_\partial, \mathbf{p})$ and variance $E_0(\mathbf{p})$. The fact that random variables (\tilde{h}, \mathbf{p}) have these laws for any $\mathbf{p} \in H_s(\mathbb{H})$ implies that \tilde{h} is a GFF with the given boundary conditions by Proposition 1.1. A similar argument applies if we replace h_0 with h_t for any stopping time t of W_t , and this implies that (\tilde{h}, γ) satisfies the hypothesis of Theorem 1.3. \square

We can now prove part of Theorem 1.3.

Lemma 2.8. *Suppose that (\tilde{h}, W) is a random variable whose law is a coupling of the GFF on \mathbb{H} (with $\pm\lambda$ boundary conditions as above) and a real-valued process $W = W_t$ defined for $t \geq 0$ that satisfies the hypothesis of Theorem 1.3. Then the marginal law of W is that of $\sqrt{4}$ times a Brownian motion (so that the Loewner evolution generated by W is SLE(4)) and $\lambda = \sqrt{\pi/8}$.*

Proof. Let $\mathbf{p} \in H_s(\mathbb{H})$ be non-negative but not identically zero. We first claim that the hypothesis of Lemma 2.8 implies the conclusion of Lemma 2.6, namely that $(\tilde{h}_t, \mathbf{p})$ is a Brownian motion when parameterized by $-E_t(\mathbf{p})$.

Write $u = u(t) = -E_t(\mathbf{p})$. It is easy to see that $u(t)$ is continuous and strictly increasing in t , at least up to the first time that K_t intersects the support of \mathbf{p} . Write $F(\cdot)$ for the inverse of $u(\cdot)$.

Fix some constant $T > 0$. We define a process B by writing $B(u) = (\tilde{h}_{F(u)}, \mathbf{p})$ whenever $F(u) < T$. After time $u_0 = \sup\{u : F(u) < T\}$, we let B evolve (independently of W) according to the law of a standard Brownian motion until time $E_0(\mathbf{p})$ (so that given B up until time u_0 , the conditional law of $B(u) - B(u_0)$ —for each $u \geq u_0$ —is a centered Gaussian of variance $u - u_0$). Because, given W , the conditional law of (h, \mathbf{p}) is a Gaussian with variance $E_0(\mathbf{p}) - u_0$, we may couple this process B with \tilde{h} in such a way $B(E_0(\mathbf{p})) = (h, \mathbf{p})$ almost surely.

Now, we claim that B is a standard Brownian motion on the interval $[0, E_0(\mathbf{p})]$. To see this, note that for each fixed $U > 0$, the conditional law of $B(E_0(\mathbf{p})) - B(U)$ (given $B(u)$ for $u \leq U$) is that of a Gaussian of variance

$E_0(\mathbf{p}) - U$ (independently of $B(u)$ for $u \leq U$). It is a general fact (easily seen by taking characteristic functions) that if X and Y are independent random variables, and X and $X + Y$ are Gaussian, then so is Y . Thus, $B(U)$ is a Gaussian of variance U that is independent of $B(V) - B(U)$ for each $V > U$, and $B(V) - B(U)$ is a Gaussian of variance $V - U$. Since B is clearly almost surely continuous, this implies that B is a Brownian motion on $[0, E_0(\mathbf{p})]$.

Now, for each fixed $z \in \mathbb{H}$, if we take \mathbf{p} to be a positive, symmetric bump function centered at z (with total integral one), then the harmonicity of \tilde{h}_t implies that $(\tilde{h}_t, \mathbf{p}) = \tilde{h}_t(z)$ provided that the support of \mathbf{p} does not intersect K_t . This implies that $\tilde{h}_t(z)$ is a continuous martingale up until the first time that K_t intersects the support of \mathbf{p} . Since we may take the support of \mathbf{p} to be arbitrarily small (and since \tilde{h}_t is bounded between $\pm\lambda$), this implies that $\tilde{h}_t(z)$ is a continuous martingale and is in fact a Brownian motion when parameterized by $E_0(\mathbf{p}) - E_t(\mathbf{p})$. The latter quantity satisfies

$$E_0(\mathbf{p}) - E_t(\mathbf{p}) = (\mathbf{p}, -\Delta^{-1}\mathbf{p} + \Delta_t^{-1}\mathbf{p}),$$

where Δ_t is the Laplacian restricted to $\mathbb{H} \setminus K_t$. Since $-\Delta^{-1}\mathbf{p} + \Delta_t^{-1}\mathbf{p}$ is harmonic on $\mathbb{H} \setminus K_t$, we have (provided K_t does not intersect the support of \mathbf{p}),

$$\begin{aligned} E_0(\mathbf{p}) - E_t(\mathbf{p}) &= (\Delta^{-1}\mathbf{p} - \Delta_t^{-1}\mathbf{p}, \delta_z) \\ &= (\mathbf{p}, \Delta^{-1}\delta_z - \Delta_t^{-1}\delta_z), \end{aligned}$$

which is the value of the harmonic function $\Delta^{-1}\delta_z - \Delta_t^{-1}\delta_z$ at the point z , which is easily seen to be the log of the modulus of the derivative at z of a conformal map from $\mathbb{H} \setminus K_t$ to \mathbb{H} that fixes z .

In order to prove the lemma, by Lemma 2.5 it is now enough to show that the fact that $\tilde{h}_t(z)$ is a Brownian motion under the time parameterization described above determines the law of W_t . At this point it is convenient to change to radial coordinates. Let Ψ be a conformal map from \mathbb{H} to the unit disc \mathbb{D} sending z to the origin and \widehat{K}_t the image of K_t under Ψ . Let $\widehat{g}_t : \mathbb{D} \setminus \widehat{K}_t \rightarrow \mathbb{D}$ be the conformal map normalized to fix 0 and have positive derivative at 0, and let \widehat{W}_t and \widehat{O}_t be the arguments of the images of W_t and ∞ , respectively, under the map $\widehat{g}_t \circ \Psi \circ g_t^{-1}$. Then $\tilde{h}_t(z)$ is an affine function of $\widehat{W}_t - \widehat{O}_t$, and the time parameterization described above is the standard radial Loewner evolution parameterization. By Loewner's equation, $\partial_t \widehat{O}_t$ is a function of $\widehat{W}_t - \widehat{O}_t$. Therefore, the process $\widehat{W}_t - \widehat{O}_t$, together with \widehat{W}_0 and

\widehat{O}_0 determines \widehat{O}_t and thus also \widehat{W}_t . This then determines W_t (see [SW05] for more details about changing between radial and chordal coordinates). \square

3 Local sets

3.1 Absolute continuity

We begin this section with two simple results about singularity and absolute continuity of the GFF.

Lemma 3.1. *Suppose that D is a simply connected planar domain and that \underline{h}_∂ is a deterministic non-identically-zero harmonic function on D and that h is an instance of the (zero boundary) GFF on D . Then h and $\underline{h}_\partial + h$ (both of which are random distributions on D) have mutually singular laws.*

Proof. We may assume without loss of generality that D is the disc of radius 1 centered at the origin (otherwise, we can conformally map D to this disc) and that $h_\partial(0) \neq 0$. For each $\epsilon > 0$, let \mathbf{p}_ϵ be a radially symmetric positive function in $H_s(D)$ which is supported in the annulus $\{z : 1 - \epsilon < |z| < 1\}$ and has integral one. If h is an instance of the GFF, then for each ϵ , the expected value of the Gaussian random variable $(h + h_\partial, \mathbf{p}_\epsilon)$ is $(h_\partial, \mathbf{p}_\epsilon) = h_\partial(0)$. It is easy to see from (1.1) that the variance of this random variable tends to zero as $\epsilon \rightarrow 0$. It follows from Borel-Cantelli that for any deterministic sequence of ϵ which tend to zero quickly enough, we will have $(h + h_\partial, \mathbf{p}_\epsilon) \rightarrow h_\partial(0)$ almost surely and $(h, \mathbf{p}_\epsilon) \rightarrow 0$ almost surely, and this implies the singularity. \square

Say two coupled variables X and Y are **almost independent** if their joint law is absolutely continuous with respect to the product of the marginal laws. We now prove the following:

Lemma 3.2. *Suppose that D is the unit disc and that S_1 and S_2 are connected closed subsets of D such that*

$$\text{dist}(S_1, S_2) := \inf\{d(x, y) : x \in S_1, y \in S_2\} = \epsilon > 0.$$

Then the projections of the GFF on D onto $\text{Harm}_{D \setminus S_1}$ and $\text{Harm}_{D \setminus S_2}$ are almost independent.

Informally, Lemma 3.2 says that the values of an instance h of the GFF on (an infinitesimal neighborhood of) S_1 are almost independent of the values of h on (an infinitesimal neighborhood of) S_2 .

Proof. Since the distance between S_1 and S_2 is positive, there exists a path γ in \bar{D} which is either simple or a simple closed loop, such that the distance δ from γ to $S_1 \cup S_2$ is positive and γ separates \bar{S}_1 from \bar{S}_2 in \bar{D} . Let D_1 and D_2 be the connected components of $D \setminus \gamma$ containing S_1 and S_2 .

Let $h_{\bar{D}}$ be an instance of the GFF in $\tilde{D} = \cup D_j$, and let $h_{\tilde{D}}^*$ be an independent instance of the projection of the GFF in D onto $\text{Harm}_{\bar{D}}$, as in Proposition 2.3. Then $h = h_{\bar{D}} + h_{\tilde{D}}^*$ is an instance of the GFF on D , by Proposition 2.3. As discussed in Section 2.1, $h_{\tilde{D}}^*$ restricted to \tilde{D} is a random harmonic (though not bounded) function on \tilde{D} , and $h_{\bar{D}}$ is a sum of independent zero boundary GFFs on D_1 and D_2 .

Next, we will construct continuous functions $h_1, h_2 \in H(D)$ such that each h_i is equal to $h_{\tilde{D}}^*$ on a $\delta/3$ neighborhood of S_i but vanishes on the component of \tilde{D} not including S_i . For $i \in \{1, 2\}$, the function $h_{\tilde{D}}$ is Lipschitz on S_i . To see this, observe that since $h_{\tilde{D}}$ is harmonic and zero on $S_i \cap \partial D$, its gradient on \tilde{D} extends continuously to all points on $\partial D \setminus \gamma$. Thus it has a maximum on each S_i (since each S_i is a positive distance from γ). We can then let h_i be the optimal Lipschitz extension to all of D of the function which is defined to be $h_{\tilde{D}}$ on the set of points in D of distance at most $\delta/3$ to S_i and 0 on $\partial\tilde{D} \cup D_{3-i}$. Since h_i is Lipschitz, it must, in particular, belong to $H(D)$.

Now, if P_i denotes the projection onto the space $\text{Harm}_{D \setminus (S_i)}$, then we have

$$(P_1(h), P_2(h)) = (P_1(h_1) + P_1(h_{\bar{D}}), P_2(h_2) + P_2(h_{\bar{D}})).$$

However, $P_i(h_{\bar{D}})$, for $i \in \{1, 2\}$ are, independent of one another.

The reader may easily check that if h' is any projection of the GFF onto a closed subspace of $H(D)$ and a is any fixed element of that subspace, then the law of $a + h'$ is absolutely continuous with respect to that of h' . In this case $G(h') := (h', a)_{\nabla} / \|a\|_{\nabla}$ is almost surely well defined (once we fix a basis of $H(D)$ comprised of members of $H_s(D)$; recall Section 1.2) and is a Gaussian with zero mean and unit variance. The Radon-Nikodym derivative is then given by

$$\exp(-\|G(h')^2/2\|) / \exp(-(G(h') + \|a\|_{\nabla})^2).$$

Absolute continuity similarly follows if a is random and independent of h' . Thus the law of

$$(P_1(h_1) + P_1(h_{\bar{D}}), P_2(h_2) + P_2(h_{\bar{D}})).$$

is absolutely continuous with respect to the law of the (independent, as discussed above) pair

$$(P_1(h_{\bar{D}}), P_2(h_{\bar{D}})),$$

which is absolutely continuous with respect to the independent product of the marginals of $(P_1(h), P_2(h))$ by the same argument applied to each component separately. \square

3.2 Local sets for discrete fields

In this subsection only we will use the symbol h to denote an instance of the discrete GFF instead of the continuum GFF. We will prove some basic results that will have analogs when h is an instance of the continuum GFF. If D is a TG -domain, then a random subset A of the set V of TG vertices in \bar{D} — coupled with an instance h of the discrete Gaussian free field on these vertices with some boundary conditions — is called **local** if conditioned on A and the restriction of h to A , the law of h is almost surely the discrete GFF whose boundary conditions are the given values of h on $A \cup \partial D$.

We will now observe some simple facts about discrete local sets; we will extend these facts in the next section to the continuum setting.

Recall that for any deterministic subset B of vertices, the space of functions supported on B is orthogonal (in the Dirichlet inner product) to the space of functions that are harmonic at every vertex in B . Denote by B^c the set of vertices in \bar{D} that do not lie in B .

Lemma 3.3. *Let h_B denote h restricted to a subset B of the vertices of D . Let A be a random subset of the vertices of D , coupled with an instance h of the discrete Gaussian free field on D with boundary conditions h_∂ . Then the following are equivalent:*

1. A is local.
2. For each fixed subset $B \subset V \cap \bar{D}$, the following holds: conditioned on h_{B^c} (for almost all choices of h_{B^c} in any version of the conditional probability), the event $A \cap B = \emptyset$ and the random variable h_B are independent.
3. For each fixed subset $B \subset V \cap \bar{D}$, the following holds: let S be the event that A intersects B , and let \tilde{A} be equal to A on the event S^c and \emptyset

otherwise. Then conditioned on h_{B^c} , the pair (S, \tilde{A}) is independent of h_B .

Proof. To show that (1) locality implies (3), it is enough to note that if we condition on h_{B^c} and A , then (if A does not intersect B), the conditional law of h_B is its conditional law given just h_{B^c} . We will show this by first sampling A , then the values of h on A , then h_{B^c} , then h_B . By the locality definition, the conditional law of h , given A and h_A , is that of a DGFF whose boundary conditions are the given values of h on $A \cup \partial D$. Therefore, after further conditioning on the event S and h_{B^c} , the conditional law of h will be that of the DGFF with the given heights on $A \cup \partial D \cup B^c$. In particular, once h_{B^c} and S^c are given, the conditional law of h does not depend on A . Since this conditional law of h is the same as the conditional law given only h_{B^c} (and no information about S or A), it follows that given h_{B^c} , the conditional law of h is independent of (S, \tilde{A}) .

Next, (3) clearly implies (2). We will now assume (2) and derive (1). To do this, we prove the statement “If $\mathbf{P}[A = C] \neq 0$, then conditioned on $A = C$ and the heights on $C \cup \partial D$, the law of h is the law of a DGFF given those heights” by induction on the size of C . The statement is clearly true if C is empty. For general C , we know from our assumption that conditioned on h_C and $A \subset C$, the law of h is the law of a DGFF with the given heights on C . By the inductive hypothesis, we know that if we condition on h_C and any particular choice for A which is a proper subset of C , we will have this same conditional law; considering separately the cases $A = C$ and A a proper subset of C , it follows that we also have this law if we condition on h_C and $A = C$, provided that $\mathbf{P}[A = C] \neq 0$. \square

The inductive technique used to derive (1) from (2) above is a fairly general one. We will make reference to it later in the paper as well.

We say that A is **algorithmically local** if there is a (possibly random) algorithm for generating A from the values of h on vertices of D such that almost surely every vertex whose height $h(v)$ is used by the algorithm is included in A . It is easy to see that algorithmically local sets are local:

Proposition 3.4. *If a random set A coupled with the discrete GFF on the vertices of D with boundary conditions h_∂ is algorithmically local, then it is also local.*

For a trivial example, A is algorithmically local whenever its law is independent of h —in particular, if A is deterministic. The set $\{v : h(v) < 0\}$

is not local; however, the set of v that are in or adjacent to a boundary-intersecting component of the subgraph of TG induced by $\{v : h(v) < 0\}$ is algorithmically local. Another algorithmically local set is the set of hexagons on either side of the discrete interface in Figure 4.1.

Given two distinct random sets A_1 and A_2 (each coupled with a discrete GFF h), we can construct a three way coupling (h, A_1, A_2) such that the marginal law of (h, A_i) (for $i \in \{1, 2\}$) is the given one, and conditioned on h , the sets A_1 and A_2 are independent of one another. This can be done by first sampling h and then sampling A_1 and A_2 independently from the regular conditional probabilities. The union of A_1 and A_2 is then a new random set coupled with h . We denote this new random set by $A_1 \check{\cup} A_2$ and refer to it as the **conditionally independent union** of A_1 and A_2 .

Lemma 3.5. *Suppose that \mathcal{A} and \mathcal{B} and \mathcal{C} are σ -algebras on which a probability measure is defined for which*

1. \mathcal{A} is independent of \mathcal{B} ,
2. \mathcal{A} is independent of \mathcal{C} , and
3. given \mathcal{A} , the σ -algebras \mathcal{B} and \mathcal{C} are independent of each other.

Then \mathcal{A} is independent of the σ -algebra generated by both \mathcal{B} and \mathcal{C} .

Proof. Let A, B, C be events in $\mathcal{A}, \mathcal{B}, \mathcal{C}$, respectively, then

$$\begin{aligned} \mathbf{P}[A \cap B \cap C] &= \mathbf{E}[\mathbf{P}[B \cap C \cap A \mid \mathcal{A}]] = \mathbf{E}[\mathbf{P}[B \cap C \mid \mathcal{A}] 1_A] \\ &= \mathbf{E}[\mathbf{P}[B \mid \mathcal{A}] \mathbf{P}[C \mid \mathcal{A}] 1_A] = \mathbf{E}[\mathbf{P}[B] \mathbf{P}[C] 1_A] = \mathbf{P}[B] \mathbf{P}[C] \mathbf{P}[A]. \quad \square \end{aligned}$$

Lemma 3.6. *If A_1 and A_2 are local sets coupled with h , then $A := A_1 \check{\cup} A_2$ is also local. In fact, we have the slightly stronger statement that given the pair (A_1, A_2) and the values of h on A , the conditional law of h off of A is that of a DGFF with the given boundary values.*

Proof. Let S_1 and S_2 be the events that A_1 and A_2 hit $B \subset \overline{D} \cap V$, respectively. Let \tilde{A}_i be equal to A_i on the event S_i^c and \emptyset otherwise. For almost all h_{B^c} we have (by Lemma 3.3) that conditioned on h_{B^c}

1. (S_1, \tilde{A}_1) is independent of h_B ,
2. (S_2, \tilde{A}_2) is independent of h_B , and

3. given h_B , the events (S_1, \tilde{A}_1) and (S_2, \tilde{A}_2) are independent of each other.

(The last item follows from the definition of conditionally independent union.) Then Lemma 3.5 implies that conditioned on h_{B^c} , the random variable h_B must be independent of $(S_1, S_2, \tilde{A}_1, \tilde{A}_2)$. In particular, this shows that h_B is independent of the union of the events S_1 and S_2 , which implies that the conditionally independent union of A_1 and A_2 is local by Lemma 3.3. It also shows the final claim in Lemma 3.6, namely that conditioned on $A_1 \cup A_2 = C$ and the values of h on C and on the pair (A_1, A_2) , the law of h is the law of a DGFF given those values. \square

The following is an immediate consequence of the definition of a local set.

Proposition 3.7. *If A is a local set, then conditioned on A (in any regular conditional probability) the expected value of $h(v)$ is, as a function of v , almost surely harmonic in the complement of A .*

Note that conditioned on A , the restriction of h to A is not deterministic; thus we would not expect the expectation of h conditioned on A to be the same as the expectation conditioned on A and the values of h in A (though something like this will turn out to hold for our continuum level sets).

Remark 3.8. Although we will not use this fact here, we remark that Lemmas 3.3 and 3.6 are true in much greater generality. Suppose that h is any random function from a finite set V to a measure space X , and that for each $f : V \rightarrow X$ and each subset B of V we are given a probability measure $\Phi(f, B)$ on functions from V to X , and for each B , the measure $\Phi(f, B)$ is a regular conditional probability for h given its values on $V \setminus B$. (In the case of the DGFF on a graph G with vertex set V , this $\Phi(f, B)$ is simply the DGFF with boundary conditions given by the values of f on $V \setminus B$ and the original boundary vertices.) We can then define a random set A coupled with h to be **local** if $\Phi(h, A)$ is a regular version of the conditional probability of h given A and the values of h on A . The proofs of Lemmas 3.6 and 3.3 apply in this generality without modification.

3.3 Local sets for the GFF

Let Γ be the space of all closed (with respect to the d_* metric) nonempty subsets of $\overline{\mathbb{H}} \cup \{\infty\}$. We will always view Γ as a metric space, endowed

with the **Hausdorff** metric induced by d_* , i.e., the distance between sets $S_1, S_2 \in \Gamma$ is

$$d_{\text{HAUS}}(S_1, S_2) := \max \left\{ \sup_{x \in S_1} d_*(x, S_2), \sup_{y \in S_2} d_*(y, S_1) \right\}.$$

Note that Γ is naturally equipped with the Borel σ -algebra on Γ induced by this metric. It is well known (and the reader may easily verify) that Γ is a compact metric space. Note that the elements of Γ are themselves compact in the d_* metric.

Given $A \subset \Gamma$, let A_δ denote the closed set containing all points in \mathbb{H} whose d_* distance from A is at most δ . Let \mathcal{A}_δ be the smallest σ algebra in which A and the restriction of h (as a distribution) to the interior of A_δ are measurable. Let $\mathcal{A} = \bigcap_{\delta \in \mathbb{Q}, \delta > 0} \mathcal{A}_\delta$. Intuitively, this is the smallest σ -field in which A and the values of h in an infinitesimal neighborhood of A are measurable.

Lemma 3.9. *Let D be a simply connected planar domain, suppose that (h, A) is a random variable which is a coupling of an instance h of the GFF with a random element A of Γ . Then the following are equivalent:*

1. *For each deterministic open $B \subset D$, we have that given the projection of h onto Harm_B , the event $A \cap B = \emptyset$ is independent of the projection of h onto Supp_B . In other words, the conditional probability that $A \cap B = \emptyset$ given h is a measurable function of the projection of h onto Harm_B .*
2. *For each deterministic open $B \subset D$, we have that given the projection of h onto Harm_B , the pair (S, \tilde{A}) (defined as in Lemma 3.3) is independent of the projection of h onto Supp_B .*
3. *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 via Proposition 2.2) and h_1 is an \mathcal{A} -measurable random distribution (i.e., as a distribution-valued function on the space of distribution-set pairs (h, A) , h_1 is \mathcal{A} -measurable) which is a.s. harmonic on $D \setminus A$.*
4. *A sample with the law of (h, A) can be produced as follows. First choose the pair (h_1, A) 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 $h = h_1 + h_2$.*

Following the discrete definitions, we say a random closed set A coupled with an instance h of the GFF, is **local** if one of the equivalent items in Lemma 3.9 holds. For any coupling of A and h , we use the notation \mathbf{C}_A to describe the conditional expectation of the distribution h given \mathcal{A} . When A is local, \mathbf{C}_A is the h_1 described in item (3) above.

Proof. Trivially, (2) implies (1). Next, suppose A satisfies (1). We may assume that D is bounded (applying a conformal map if necessary to make this the case). Fix δ and let \widehat{A}_δ denote the intersection of D with the union of all closed squares of the grid $\delta\mathbb{Z}^2$ that intersect A_δ . Then we claim that \widehat{A}_δ satisfies (1) as well for each deterministic choice of δ . This can be seen by replacing B with $B' := D \setminus (D \setminus B)_\delta$, and noting that A intersects B' if and only if A_δ intersects B' . Since $B' \subset B$, conditioning on $\text{Harm}_{B'}$ is equivalent to conditioning on Harm_B and then conditioning on a function of Supp_B (and $\text{Supp}_{B'}$ is also a function of Supp_B), which proves the claim.

There are only finitely many possible choices for \widehat{A}_δ , so the fact that \widehat{A}_δ satisfies (3) follows by the inductive argument used in the proof of Lemma 3.3.

To be precise, we prove the statement “If $\mathbf{P}[\widehat{A}_\delta = C] \neq 0$, then conditioned on $\widehat{A}_\delta = C$ and the projection h_1 of h onto the space of functions harmonic off of C , the law of h is the law of a zero-boundary DGFF on $D \setminus C$ plus h_1 ” by induction on the size of C . The statement is clearly true if C is empty. For general C , we know from our assumption that conditioned on h_C and $A \subset C$, the law of h is the law of a DGFF with the given heights on C . By the inductive hypothesis, we know that if we condition on h_C and any particular choice for A which is a proper subset of C , we will have this same conditional law; it follows that we also have this law if we condition on h_C and $A = C$, provided that $\mathbf{P}[A = C] \neq 0$.

Since \mathcal{A} is the intersection of the \widehat{A}_δ (defined analogously to \mathcal{A}_δ), for $\delta > 0$, the reverse martingale convergence theorem implies the almost sure convergence $\mathbf{C}_{\widehat{A}_\delta} \rightarrow \mathbf{C}_A$ as $\delta \rightarrow 0$ in the weak sense, i.e., for each fixed \mathbf{p} , we have a.s.

$$(\mathbf{C}_{\widehat{A}_\delta}, \mathbf{p}) \rightarrow (\mathbf{C}_A, \mathbf{p}).$$

This and the fact that (3) holds for every \widehat{A}_δ implies that it must hold for A as well. Since this holds for every fixed \mathbf{p} , we may extend to all $\mathbf{p} \in H_s(D)$ and obtain (3) by Proposition 1.1 and Proposition 2.2.

Now (4) is immediate from (3) when we set $h_1 = \mathbf{C}_A$. To obtain (2) from

(4), it suffices to show that given the projection of h onto Harm_B and the pair (S, \tilde{A}) , the conditional law of the projection of h onto Supp_B is the same as its a priori law (or its law conditioned on only the projection of h onto Harm_B), namely the law of the zero boundary GFF on B . To see this, we may first sample A and h_1 and then—conditioned on $A \cap B = \emptyset$ —sample the projection of $h - h_1$ onto Supp_B . Since the law of $h - h_1$ is the GFF on $D \setminus A$ by assumption, this projection is the GFF on B , as desired. \square

Lemma 3.10. *Lemma 3.6 applies in the continuum setting as well. That is, if A_1 and A_2 are local sets coupled with the GFF h on D , then their conditionally independent union $A = A_1 \dot{\cup} A_2$ is also local. The analog of the slightly stronger statement in Lemma 3.6 also holds: given \mathcal{A} and the pair (A_1, A_2) , the conditional law of h is given by \mathbf{C}_A plus an instance of the GFF on $D \setminus A$.*

Proof. The proof is essentially identical to the discrete case. We use characterization (2) for locality as given in Lemma 3.9 and observe that Lemma 3.5 implies the analogous result holds for the quadruple $(S_1, \tilde{A}_1, S_2, \tilde{A}_2)$ —namely, that for each deterministic open $B \subset D$, we have that *given* the projection of h onto Harm_B and the quadruple $(S_1, \tilde{A}_1, S_2, \tilde{A}_2)$, the conditional law of the projection of h onto Supp_B (assuming $A \cap B = \emptyset$) is the law of the GFF on B . The proof that this analog of (2) implies the corresponding analog of (3) in the statement of Lemma 3.10 is also essentially as the discrete case. \square

Lemma 3.11. *Let A_1 and A_2 be connected local sets. Then $\mathbf{C}_{A_1 \dot{\cup} A_2} - \mathbf{C}_{A_2}$ is almost surely a harmonic function in $D \setminus (A_1 \dot{\cup} A_2)$ that tends to zero on all sequences of points in $D \setminus (A_1 \dot{\cup} A_2)$ that tend to a limit in $A_2 \setminus A_1$ (unless A_2 is a single point).*

Proof. By Lemma 3.10, the union $A_1 \dot{\cup} A_2$ is itself a local set, so $\mathbf{C}_{A_1 \dot{\cup} A_2}$ is well defined. Now, conditioned on \mathcal{A}_1 the law of the field in $D \setminus A_1$ is given by a GFF in $D \setminus A_1$ plus \mathbf{C}_{A_1} . We next claim that $\overline{A_2 \setminus A_1}$ is a local subset of $D \setminus A_1$, with respect to this GFF on $D \setminus A_1$. To see this, note that characterization (1) for locality from Lemma 3.9 follows from the latter statement in Lemma 3.10.

By replacing D with $D \setminus A_1$ and subtracting \mathbf{C}_{A_1} , we may thus reduce to the case that A_1 is deterministically empty and $\mathbf{C}_{A_1} = 0$. What remains to show is that if A is any local set on D then \mathbf{C}_A (when viewed as a harmonic

function on $D \setminus A$) tends to zero almost surely along all sequences of points in $D \setminus A$ that approach a point x that lies on a connected components of $\partial D \setminus A$ that consists of more than a single point.

If we fix a neighborhood B_1 of x and another neighborhood B_2 whose distance from B_1 is positive, then the fact that the statement holds on the event $A \subset B_2$ is immediately from Lemma 3.1 and Lemma 3.2. Since this holds for arbitrary B_1 and B_2 , the result follows. \square

To conclude, we note the following is immediate from the definition of local and Theorem 1.3.

Lemma 3.12. *In the coupling between h and γ of Theorem 1.3, the set $\gamma([0, T])$ is local for every $T \geq 0$. The same is true if T is a non-deterministic stopping time of the process W_t .*

4 Fine grid preliminaries

4.1 Subspaces are asymptotically dense

As $r_D \rightarrow \infty$, the subspaces $\{g \circ \phi_D^{-1} : g \in H_{TG}(D)\}$, where ϕ_D and r_D are as defined in Section 1.3, become asymptotically dense in $H(\mathbb{H})$ in the following sense. When D is a TG -domain and $g \in H_s(D)$, let $P_D(g)$ denote the orthogonal projection of g (with respect to the inner product $(\cdot, \cdot)_\nabla$) onto the space of continuous functions which are affine on each triangle of TG . For $f \in H(\mathbb{H})$ set $f_D := f \circ \phi_D$.

Lemma 4.1. *Let $D \subset \mathbb{C}$ denote a TG -domain, and assume the notation above. For each $f \in H(\mathbb{H})$, the values $\|P_D(f_D) \circ \phi_D^{-1} - f\|_\nabla = \|P_D(f_D) - f_D\|_\nabla$ tend to zero as $r_D \rightarrow \infty$. In fact, if $f \in H_s(\mathbb{H})$, then $\|P_D(f_D) - f_D\|_\nabla = O(\frac{1}{r_D})$, where the implied constant may depend on f .*

Proof. Since $H_s(\mathbb{H})$ is dense in $H(\mathbb{H})$, the former statement follows from the latter. Suppose that $f \in H_s(\mathbb{H})$. Then it is supported on a compact subset K of \mathbb{H} .

When z ranges over values in K and ϕ ranges over *all* conformal functions that map a subdomain of \mathbb{C} onto \mathbb{H} , standard distortion theorems for conformal functions (e.g., Proposition 1.2 and Corollary 1.4 of [Pom92]) imply the following (where the implied constants may depend on K):

1. The ratio of $|(\phi^{-1})'(i)|$ and $r_D := \text{rad}_{\phi^{-1}(i)}(\phi^{-1}\mathbb{H})$ is bounded between two positive constants.
2. $|(\phi^{-1})'(z)| = O(|(\phi^{-1})'(i)|) = O(r_D)$.
3. $\text{diam}(\phi^{-1}(K)) = O(|(\phi^{-1})'(i)|) = O(r_D)$.
4. $|\phi'(\phi^{-1}(z))| = O(\frac{1}{r_D})$.
5. $|\phi''(\phi^{-1}(z))| = O(\frac{1}{(\phi^{-1})'(z)^2}) = O(r_D^{-2})$.

Now $\|P_D(f_D) - f_D\|_{\nabla} = \inf\{\|g - f_D\|_{\nabla} : g \in H_{TG}(D)\}$. We will bound the latter by considering the case that g is the function $g_D \in H_{TG}(D)$ that agrees with f_D on TG — and then applying the above bounds with $\phi = \phi_D$.

Since the triangles of TG contained in D have side length one, the value $|\nabla g_D - \nabla f_D|$ on a triangle is bounded by a constant times the maximal norm of the second derivative matrix of $f_D = f \circ \phi_D$ in that triangle (where the latter is viewed as a function from \mathbb{R}^2 to \mathbb{R}). If f and ϕ_D were both functions from \mathbb{R} to \mathbb{R} , then the chain rule would give

$$(f \circ \phi_D)''(z) = [f'(\phi_D(z))\phi_D'(z)]' = f''(\phi_D(z))\phi_D'(z)^2 + f'(\phi_D(z))\phi_D''(z).$$

In our case, when we view f as a function from \mathbb{R}^2 to \mathbb{R} and ϕ_D' as a function from \mathbb{R}^2 to \mathbb{R}^2 , the chain rule yields the formulas: but now f' at a point is understood to be a linear map from \mathbb{R}^2 to \mathbb{R} , and ϕ_D' at a point is understood to be a linear map from \mathbb{R}^2 to \mathbb{R}^2 , etc.

Since all components of f' and f'' are bounded on K , the distortion bounds above give

$$|(f \circ \phi_D)''(z)| = O(|\phi_D'(z)^2 + \phi_D''(z)|) = O(r_D^{-2})$$

and hence

$$\|\nabla g_D - \nabla f_D\|_{\infty}^2 = O(r_D^{-4}).$$

The area of the support of $f \circ \phi_D$ is $O([\text{diam } \phi_D^{-1}(K)]^2) = O(r_D^2)$. Thus $\|g_D - f \circ \phi_D\|_{\nabla}^2 = O(r_D^2/r_D^4) = O(r_D^{-2})$. \square

4.2 Topological and measure theoretic preliminaries

In this section we assemble several simple topological facts that will play a role in the proof of Theorem 1.2. Up to this point, we have treated the space $\Omega = \Omega_D$ of distributions on a planar domain D as a measure space, using \mathcal{F} to represent the smallest σ -algebra that makes (\cdot, \mathbf{p}) measurable for each fixed $\mathbf{p} \in H_s(D)$. We have not yet explicitly introduced a metric on Ω . (When we discussed convergence of distributions, we implicitly used the weak topology—i.e., the topology in which $h_i \rightarrow h$ if and only if $(h_i, \mathbf{p}) \rightarrow (h, \mathbf{p})$ for all $\mathbf{p} \in H_s(D)$.) Although it does not play a role in our main theorem statements, the following lemma will be useful in the proofs. Recall that a topological space is called σ -**compact** if it is the union of countably many compact sets:

Lemma 4.2. *Let D be a simply connected domain. There exists a metric \widehat{d} on a subspace $\widehat{\Omega} \subset \Omega$ with $\widehat{\Omega} \in \mathcal{F}$ such that*

1. *An instance of the GFF on D lies in $\widehat{\Omega}$ almost surely.*
2. *The topology induced by \widehat{d} on $\widehat{\Omega}$ is σ -compact.*
3. *The Borel σ -algebra on $\widehat{\Omega}$ induced by \widehat{d} is the set of subsets of $\widehat{\Omega}$ that lie in \mathcal{F} .*

Proof. It is enough to prove Lemma 4.2 for a single bounded simply connected domain D (say the unit disc), since pulling back the metric \widehat{D} via a conformal map preserves the properties claimed in the Lemma. When f_i is an eigenvalue of the Laplacian with negative eigenvalue λ , then we may define $(-\Delta)^a f_i = (-\lambda)^a f_i$, and we may extend this definition linearly to the linear span of the f_i . Denote by $(-\Delta)^a L^2(D)$ the Hilbert space closure of the linear span of the eigenfunctions of the Laplacian on D (that vanish on ∂D) under the inner product $(f, g)_a := ((-\Delta)^{-a} f, (-\Delta)^{-a} g)$. In other words, $(-\Delta)^a L^2(D)$ consists of those f for which $(-\Delta)^{-a} f \in L^2(D)$. It follows immediately from Weyl's formula for bounded domains that $(-\Delta)^a L^2(D) \subset (-\Delta)^b L^2(D)$ when $a < b$, that each of these spaces is naturally a subset of Ω , and that when $a > 0$, an instance of the GFF almost surely lies in $(-\Delta)^a L^2(D)$. (See [She03] for details.) We can thus take $\widehat{\Omega} = (-\Delta)^a L^2(D)$ for some $a > 0$ and let \widehat{d} be the Hilbert space metric corresponding to $(-\Delta)^b L^2(D)$ for some $b > a$.

To see that the topology induced by \widehat{d} on $\widehat{\Omega}$ is σ -compact follows from the fact that with its usual Hilbert space metric, $(-\Delta)^a L^2(D)$ is separable (in particular can be covered with countably many translates of the unit ball), and that such a unit ball is compact w.r.t. \widehat{d} .

We next argue that the Borel σ -algebra $\widehat{\mathcal{F}}$ on $\widehat{\Omega}$ induced by \widehat{d} is the set of subsets of $\widehat{\Omega}$ that lie in \mathcal{F} . Recall that \mathcal{F} is the smallest σ -algebra that makes (h, \mathbf{p}) measurable for each $\mathbf{p} \in H_s(D)$. Clearly a unit ball of \widehat{d} is in this σ -algebra (since $H_s(D)$ is dense in such a unit ball), which shows that $\widehat{\mathcal{F}} \subset \mathcal{F}$. For the other direction, it suffices to observe that each generating subset of \mathcal{F} of the form $\{(\cdot, \mathbf{p}) \leq c\}$, with $c \in \mathbb{R}$ and $\mathbf{p} \in H_s(D)$ has an intersection with $\widehat{\Omega}$ that belongs to $\widehat{\mathcal{F}}$. \square

We now cite the following basic fact (see [Vai, Thm. 72,73] or [Dud02, Ch. 7]):

Lemma 4.3. *Every σ -compact metric space is separable. If μ is a Borel probability measure on a σ -compact metric space then μ is regular, i.e., for each Borel measurable set S , we have*

$$\mu(S) = \inf \mu(S') = \sup \mu(S''),$$

where S' ranges over open supersets of S and S'' ranges over compact subsets of S .

A family of probability measures μ on a separable topological space X is said to be *tight* if for every $\epsilon > 0$ there is a compact $X' \subset X$ such that $\mu(X') > 1 - \epsilon$ for every μ in the family. Prokhorov's theorem states that (assuming X is separable) every tight family of probability measures on X is weakly pre-compact. If X is a separable metric space, then the converse holds, i.e., every weakly pre-compact family of probability measures is also tight.

Lemma 4.4. *If Θ_1 and Θ_2 are two weakly pre-compact families of probability measures on complete separable metric spaces Z_1 and Z_2 , then the space of couplings between measures in the two families is weakly pre-compact.*

Proof. By the converse to Prokhorov's theorem, Θ_1 and Θ_2 are both tight. This implies that the space of couplings between elements of Θ_1 and Θ_2 is also tight, which in turn implies pre-compactness (by Prokhorov's theorem). \square

The following is another simple topological observation that will be useful later on:

Lemma 4.5. *Suppose that Z_1 and Z_2 are complete separable metric spaces, μ is a Borel probability measure on Z_1 and ψ_1, ψ_2, \dots is a sequence of measurable functions from Z_1 to Z_2 . Suppose further that when z is a random variable distributed according to μ , the law of $(z, \psi_i(z))$ converges weakly to that of $(z, \psi(z))$ as $i \rightarrow \infty$, where $\psi : Z_1 \rightarrow Z_2$ is Borel measurable. Then the functions ψ_i , viewed as random variables on the probability space Z_1 , converge to ψ in probability.*

Proof. By tightness of the set of measures in the sequence (recall Lemma 4.4), for each $\epsilon > 0$, we can find a compact $K \subset Z_2$ such that $\mu(\psi^{-1}(K)) > 1 - \epsilon$. Let B_1, \dots, B_k be a finite partition of K into disjoint measurable sets of diameter at most ϵ and write $C_j = \psi^{-1}B_j$ for each j . Then Lemma 4.3 implies that there exist open subsets C'_1, \dots, C'_k of Z_1 such that $C'_j \supset C_j$ for each j and $\sum_j \mu(C'_j \setminus C_j) \leq \epsilon$. For each j , let B'_j be the set of points of distance at most ϵ from B_j . Let $\tilde{\mu}_i$ denote the law of $(z, \psi_i(z))$ and $\tilde{\mu}$ the law of $(z, \psi(z))$. Set $A' := \bigcup_{j=1}^k C'_j \times B'_j$ and $A := \bigcup_{j=1}^k C_j \times B_j$. Then a standard consequence of weak convergence (Portmanteau's theorem [Dud02, Theorem 11.1.1]) implies $\liminf_{i \rightarrow \infty} \tilde{\mu}_i(A') \geq \tilde{\mu}(A)$. But

$$\tilde{\mu}(A') \geq \tilde{\mu}(A) = \sum_{j=1}^k \mu(C_j) > 1 - \epsilon.$$

Hence $\liminf_{i \rightarrow \infty} \tilde{\mu}_i(A') > 1 - \epsilon$. Since $\sum_{j=1}^k \mu(C'_j \setminus C_j) \leq \epsilon$, we have $\liminf_{i \rightarrow \infty} \tilde{\mu}_i(A) \geq 1 - 2\epsilon$. But when $(z, \psi_i(z)) \in A$, the distance between $\psi_i(z)$ and $\psi(z)$ is at most 2ϵ . Hence,

$$\liminf_{i \rightarrow \infty} \mu\{x \in Z_1 : d(\psi_i(x), \psi(x)) \leq 2\epsilon\} \geq 1 - 2\epsilon.$$

Since this holds for any $\epsilon > 0$, the result follows. \square

4.3 Limits of discrete local sets are local

Lemma 4.6. *Let D_n be a sequence of TG-domains with maps $\phi_n : D_n \rightarrow \mathbb{H}$ such that $r_{D_n} \rightarrow \infty$ as $n \rightarrow \infty$. Let an instance h of the GFF on \mathbb{H} be coupled with the discrete GFF on each D_n , as in Section 1.3. Let A_n be a sequence of discrete local subsets of $D_n \cap TG$. Then there is a subsequence along which the law of $(h, \phi_n A_n)$ converges weakly (in the space of measures on $\widehat{\Omega} \times \Gamma$) to a limiting coupling (h, A) with respect to the sum of the metric*

\widehat{d} on the first component and d_{HAUS} on the second component. In any such limit, A is local.

Proof. Lemmas 4.2, 4.3, and 4.4 imply the existence of the subsequential limit (h, A) , so it remains only to show that in any such limit A is local. We will prove that characterization (2) for locality as given in Lemma 3.9 holds. For this, it suffices to show that for every deterministic open $B \subset \mathbb{H}$ and function $\phi \in -\Delta H_s(B)$ (supported in a compact subset of B) the law of (h, ϕ) is independent of the pair (S, \tilde{A}) (as defined in Lemma 3.9) together with the projection of h onto Harm_B . Here we are using the fact that for Gaussian fields the marginals characterize the field; see Lemma 2.1. It is clearly enough to consider the case that B has compact closure in \mathbb{H} .

Fix $g \in H_s(B)$ and set $\phi = -\Delta g$. Let \mathbb{S}_n denote the space $\{f \circ \phi_{D_n}^{-1} : f \in H_{TG}(D_n)\}$. By Lemma 4.1, we can approximate g by elements g_n in \mathbb{S}_n in such a way that $\|g_n - g\|_{\nabla} \rightarrow 0$ as $n \rightarrow \infty$. Let B' be the set of points in B of distance at least ϵ from ∂B , where ϵ is small enough so that g is compactly supported in B' . In fact, the construction given in the proof of Lemma 4.1 ensures that each g_n will be supported in B' for all n sufficiently large.

Now, for each fixed n , let h_n^1 denote the projection of h onto the space of functions in \mathbb{S}_n that vanish outside of B' . Let h_n^2 and h_n^3 be such that $h_n^1 + h_n^2$ is the projection of h onto \mathbb{S}_n and $h_n^1 + h_n^2 + h_n^3 = h$. Clearly, h_n^1 , h_n^2 , and h_n^3 are mutually independent, since they are projections of h onto orthogonal spaces.

Following characterization 3 of Lemma 3.3, let S_n be the event that A_n includes a vertex of a triangle whose image under ϕ_n intersects B' , and let \tilde{A}_n be equal to A_n on the event S_n^c and \emptyset otherwise. By Lemma 3.3, conditioned on h_n^2 , the pair (S_n, \tilde{A}_n) is independent of h_n^1 . In fact, (since h_n^3 is a priori independent of the triple (A_n, h_n^1, h_n^2)), the pair (S_n, \tilde{A}_n) is independent of $h_n^1 + h_n^3$.

When n is large enough, the space Harm_B is orthogonal to \mathbb{S}_n . Thus, for each sufficiently large n , $(h, g_n)_{\nabla}$ is independent of the projection h_{B^c} of h onto Harm_B and the pair (S_n, \tilde{A}_n) .

Now, since $\|g_n - g\|_{\nabla} \rightarrow 0$ as $n \rightarrow \infty$, the random variables $(g_n - g, h)_{\nabla}$ tend to zero in law as $n \rightarrow \infty$. Since weak limits of independent random variables are independent, we conclude that in any weak limit $(h, S_{\text{lim}}, \tilde{A}_{\text{lim}}, A)$ of the quadruple $(h, S_n, \phi_n \tilde{A}_n, \phi_n A_n)$ (again, subsequential limits exist by Lemmas 4.2, 4.3, and 4.4), the value $(h, g)_{\nabla}$ is independent of h_{B^c} and $(S_{\text{lim}}, \tilde{A}_{\text{lim}})$. The event S_{lim} contains the event S (since any Hausdorff limit of sets that




Figure 4.1: Gaussian free field on faces of hexagonal lattice—faces shaded by height—with boundary conditions equal to $-\lambda$ on the left boundary arc and λ on the right boundary arc, where $\lambda > 0$ is a constant. Thick line indicates chordal interface between positive and negative height hexagons. In the figure, λ is taken to be the special constant for which, as the mesh size is taken to zero, the law of the interface converges to that of SLE(4).

intersect B' must intersect B), and thus the pair $(S_{\text{lim}}, \tilde{A}_{\text{lim}})$ determines the pair (S, A) . This implies that $(h, g)_{\nabla}$ is independent of h_{B^c} and (S, \tilde{A}) . Since this is true for all B and g supported on B , we conclude that A is local. \square

4.4 Statement of the height gap lemma

We now state the special case of the **height gap lemma** (as proved in [SS09]) that is relevant to the current work. (The lemma in [SS09] applies to more general boundary conditions.)

As usual, we let D be a TG domain with boundary conditions $-\lambda$ on one arc ∂_- and λ on a complementary arc ∂_+ , and let x_{∂} and y_{∂} denote respectively the clockwise and counterclockwise endpoints of ∂_+ .

Let γ^T be the path in the dual lattice of TG from x_∂ to y_∂ that has adjacent to its right hand side vertices in ∂_+ or vertices where $h > 0$ and has adjacent to its left hand side vertices in ∂_- or vertices where $h < 0$, stopped at some stopping time T for the discrete exploration process. This is the path that traces the boundary between hexagons with positive sign and hexagons with negative sign in the dual lattice, as described in [SS09] and illustrated in Figure 4.1. Let v_0 be some vertex of TG in D .

Let V_- denote the vertices on the left side of γ^T together with the vertices in ∂_- and V_+ the vertices on the right side or in ∂_+ . Let F_T denote the function that is $+\lambda$ on $V_+(\gamma^T)$, $-\lambda$ on $V_-(\gamma^T)$ and discrete-harmonic at all other vertices in \overline{D} . Let h_T be the discrete harmonic interpolation of the values of h on $V_-(\gamma^T) \cup V_+(\gamma^T)$ and on all TG -vertices in ∂D .

Lemma 4.7. *For some fixed value of $\lambda > 0$, we have*

$$h_T(v_0) - F_T(v_0) \rightarrow 0$$

in probability as T , D and v_0 are taken so that $\text{dist}(v_0, \partial D) \rightarrow \infty$. Similarly, if v_0 is a random vertex (with law independent of h) supported on the set of points of distance at least r from ∂D , then as $r \rightarrow \infty$

$$\mathbf{E}[h_T(v_0) - F_T(v_0) \mid \gamma^T]$$

(viewed as a random variable depending on γ^T — the expectation is respect to both v_0 and h_T) tends to zero in probability.

5 Proofs of main results

Proof of Theorem 1.3. In Section 2.2 (Lemma 2.7), we explicitly produced a coupling of W (the Loewner driving parameter of an SLE(4)) and h (the GFF on \mathbb{H} with $\pm\lambda$ boundary conditions) with the conformal Markov property described in Theorem 1.3. Lemma 2.8 implies that any (\tilde{h}, \tilde{W}) that satisfies the hypotheses of Theorem 1.3 must have this same law—and that the value of λ is indeed $\sqrt{\pi/8}$.

All that remains to prove in Theorem 1.3 is items 3 and 4. To prove 3 we must show that W is equivalent (up to redefinition on a set of measure zero) to an \mathcal{F} -measurable function from Ω to Λ . In other words, we must show that given h , the conditional law of W (in any regular version of the conditional probability) is almost surely supported on a single element of Λ .

Let h be an instance of the GFF (with boundary conditions $-\lambda$ on $(-\infty, 0)$ and λ on $(0, \infty)$). Write $\Phi(z) = -z^{-1}$. Then Φ is a conformal automorphism of \mathbb{H} sending 0 to ∞ and ∞ to zero, and $-h \circ \Phi$ has the same law as h (where $-h \circ \Phi$ is the pullback of h as defined in Section 1.2). Let W and V be random elements of Λ coupled with h in such a way that

1. The pair (h, W) satisfies the hypotheses of Theorem 1.3.
2. The pair $(-h \circ \Phi, V)$ satisfies the hypotheses of Theorem 1.3.
3. Given h , V and W are independent of one another.

Let γ^1 be the path with Loewner evolution given by W and let γ^2 be the image of the Loewner evolution generated by V under Φ . Then the law of γ^1 is that of an SLE(4) from 0 to ∞ and that the law of γ^2 is that of an SLE(4) from ∞ to 0 .

For any fixed time T , conditioned on $\gamma^2([0, T])$, the law of h is that of a GFF on $\mathbb{H} \setminus \gamma^2([0, T])$ with boundary conditions of $-\lambda$ on the left side of $\gamma^2([0, T])$ and $(-\infty, 0)$ and λ on the right side of $\gamma^2([0, T])$ and $(0, \infty)$. In particular (recall Lemma 3.12) $\gamma^2([0, T])$ is local, and the same holds if T is a stopping time of $\gamma^2([0, T])$.

If we fix a stopping time T_1 for γ^1 , then $\gamma^1([0, T_1])$ is also local. On the event that $\gamma^1([0, T_1])$ and $\gamma^2([0, T])$ do not intersect each other, Lemma 3.11 yields that the conditional law of h given both sets is that of a GFF on the complement of these sets, with the expected $\pm\lambda$ boundary conditions. Since the same holds for any T_1 , Lemma 2.8 and Lemma 3.11 imply that conditioned on $\gamma^2([0, T])$, the law of γ^1 — up until the first time it hits $\gamma^2([0, T])$, is that of an SLE(4) in $\mathbb{H} \setminus \gamma^2([0, T])$, started at 0 and targeted at $\gamma^2(T)$.

It follows that almost surely γ^1 hits $\gamma^2([0, T])$ for the first time at $\gamma^2(T)$. Since this applies to any choice of T , we conclude that γ^1 hits a dense countable set of points along γ^2 , and (by symmetry) γ^2 hits a dense countable set of points along γ^1 , and hence the two paths (both of which are almost surely continuous simple paths, by the corresponding almost sure properties of SLE(4)) are equal almost surely. This implies that conditioned on V and h , the law of W is almost surely supported on a single element of Λ . Since V and W are conditionally independent given h , it follows that conditioned on h , the law of W is almost surely supported on a single element of Λ . The proof of item 4 will be established during the proof of Theorem 1.2 below. \square

In what follows, we use the notation introduced in Section 1.3 and the statement of Theorem 1.2. We will first give a proof of Theorem 1.2 that assumes the main result of [SS09], namely that the discrete interfaces converge in law to SLE(4) with respect to the metric d_{STRONG} . Afterwards, we will show how this convergence in law can be derived from Lemma 4.7 (the height gap lemma) and Theorem 1.2. That is, we give an alternate way of deriving the main result of [SS09] (in the case of $\pm\lambda$ boundary conditions) from Theorem 1.2, so that the only result from [SS09] that we really need for this paper is Lemma 4.7. (The proof of Lemma 4.7 admittedly takes about 2/3 of the body of [SS09], excluding the introduction, preliminaries, etc.)

Proof of Theorem 1.2. Let $\phi_n : D_n \rightarrow \mathbb{H}$ be a sequence of conformal homeomorphisms from TG -domains D_n to \mathbb{H} such that $\lim_{n \rightarrow \infty} r_{D_n} = \infty$. Let $\hat{\gamma}^n = \hat{\gamma}_{D_n}$ denote the image in \mathbb{H} of the interface of the coupled discrete GFF in D_n . For each fixed t , by Lemma 4.4, there is a subsequence of the D_n along which the pair $(\hat{\gamma}^n([0, t]), h)$ converges in law (with respect to the sum of the Hausdorff metric on the first component and the \hat{d} metric on the second component) to the law of a random pair $(\gamma([0, t]), h)$, where the marginal law of γ (by the main result of [SS09]) must be SLE(4).

By Theorem 1.3 and Lemma 4.5, it will be enough to show that any such limiting pair (γ, h) satisfies the hypotheses of Theorem 1.3. For each fixed t , by Lemma 4.6, $\gamma([0, t])$ is a local set in this limiting coupling. We next claim that $\mathbf{C}_{\gamma([0, t])}$ is almost surely given by

$$h_t := \lambda \left(1 - 2\pi^{-1} \arg(g_t - W_t) \right),$$

where g_t is the Loewner evolution, driven by W_t , that corresponds to γ . Note that since h_t is a bounded function that is defined almost everywhere in \mathbb{H} , it may be also viewed as a distribution on \mathbb{H} in the obvious way: $(h_t, \mathbf{p}) = \int h_t(z) \mathbf{p}(z) dz$.

Once this claim is proved, Theorem 1.2 is immediate from Theorem 1.3, since the claim implies that the limiting law of (h, W) satisfies the hypotheses of Theorem 1.3 and thus W is almost surely the Λ -valued function of h described in Theorem 1.3—and the fact that this convergence holds for any subsequence of the D_n implies that it must hold for the entire sequence.

Let A_n denote the set of vertices incident to the left or right of the preimage of the path $\hat{\gamma}^n([0, t])$ in D_n (so that each A_n is a discrete algorithmically local set, representing the set of vertices whose values are observed up to the

first point in the exploration algorithm that the capacity of the image of the level line in \mathbb{H} reaches t).

Let \mathbf{C}_t^n denote the conditional expectation of h given the values of h_{D_n} on vertices in A_n , viewed as a distribution—more precisely, $(\mathbf{C}_t^n, \mathbf{p}) = (\widehat{h}_{D_n} \circ \phi_n^{-1}, \mathbf{p})$, where \widehat{h}_{D_n} is the (piecewise affine interpolation of) the discrete harmonic interpolation to D_n of the values of h_{D_n} on the vertices of A_n and on the boundary vertices. Let W_t^n be the Loewner driving parameter for $\widehat{\gamma}^n([0, t])$. Fix $t \geq 0$ and consider now the triple:

$$(W_t^n, \mathbf{C}_t^n, \widehat{\gamma}^n([0, t])).$$

By Lemma 4.4, this converges along a subsequence in law (with respect to the d_{STRONG} metric on first coordinate plus the \widehat{d} metric on the second coordinate plus the Hausdorff d_* metric on the third coordinate) to a limit (W_t, \mathbf{C}_t, K_t) . We may define the Loewner evolution g_t^n in terms of W_t^n and analogously define

$$h_t^n := \lambda(1 - 2\pi^{-1} \arg(g_t^n - W_t^n)).$$

For each $\mathbf{p} \in H_s(\mathbb{H})$, we claim that the random quantity

$$(h_t^n, \mathbf{p}) - (\mathbf{C}_t^n, \mathbf{p})$$

is a continuous function of the triplet above, which implies that the difference between this quantity and $(h_t, \mathbf{p}) - (\mathbf{C}_t, \mathbf{p})$ converges in probability to zero. The continuity of the latter term holds simply since \mathbf{p} is smooth and compactly supported (and thus the $\Delta^a \mathbf{p}$ lies in L^2 for all a), while the former piece is continuous with respect to the d_{STRONG} metric on γ . Following the proof of Lemma 4.6, it is not hard to see that $\mathbf{C}_t = \mathbf{C}_{K_t}$, since on the discrete level, once one conditions on h_t^n and \mathbf{C}_t^n the conditional law of the field minus \mathbf{C}_t^n is that of a zero boundary DGFF on the set of unobserved vertices.

A harmonic function is determined by its values in an open set, since a harmonic function is the real part of a (possibly multi-valued) analytic function, so it is now enough to show that for each such \mathbf{p} , the difference between the conditional expectation of (h, \mathbf{p}) — given A_n and the value of h_{D_n} on A_n — and the value (h_t^n, \mathbf{p}) converges in probability to zero as $n \rightarrow \infty$ (along a subsequence for which a weak limit exists).

The expected value of (h, \mathbf{p}) given the values on A_n is given by $(\mathbf{C}_t^n, \mathbf{p})$, which is in turn a weighted average of the values of \widehat{h}_{D_n} on vertices of D_n which lie on triangles that intersect the image under ϕ_n of the support of

\mathfrak{p} . It follows from Lemma 4.7 that if \mathfrak{p} is such that the distance between these vertices and A_n necessarily tends to ∞ as $n \rightarrow \infty$ (which is the case if \mathfrak{p} is compactly supported in the complement of the set of points that can be reached by a Loewner evolution up to time t), then any subsequential weak limit of the law of the triplet above is the same as it would be if \widehat{h}_{D_n} were replaced by the discrete harmonic function which is $-\lambda$ on the left-side vertices of A_n and λ on the right side. By standard estimates relating discrete and continuous harmonic measure (it is enough here to recall that discrete random walk scales to Brownian motion), we therefore have $(h_t^n, \mathfrak{p}) - (\mathbf{C}_t^n, \mathfrak{p}) \rightarrow 0$ in law as $n \rightarrow \infty$ and thus $(h_t, \mathfrak{p}) - (\mathbf{C}_{K_t}, \mathfrak{p}) = 0$ almost surely. Since this is true for any \mathfrak{p} which is necessarily supported off of K_t , we have $h_t = \mathbf{C}_{K_t}$ on $\mathbb{H} \setminus K_t$ almost surely as desired. \square

We now give an alternate proof of the fact, proved in [SS09], that the $\widehat{\gamma}^n$ converge in law to SLE(4). Using the notation introduced in the previous proof, write

$$T_\epsilon^n(\mathfrak{p}) = \inf\{t : (h_t^n, \mathfrak{p}) - (\mathbf{C}_t^n, \mathfrak{p}) > \epsilon\}.$$

$$U_\epsilon^n(\mathfrak{p}) = \inf\{t : (h_t^n, \mathfrak{p}) - (\mathbf{C}_t^n, \mathfrak{p}) < -\epsilon\}.$$

Consider a subsequential limit along which the quadruple

$$(t, W_t^n, \mathbf{C}_t^n, \widehat{\gamma}^n([0, t])),$$

defined with $t = T_\epsilon^n(\mathfrak{p})$, converges in law to a limit

$$(T, W_T, \mathbf{C}_{K_T}, K_T).$$

We claim that in such a limit, for any fixed neighborhood of the support of \mathfrak{p} , T is almost surely large enough so that K_T intersects that neighborhood. The arguments in the previous proof would imply that—on the event that K_T does not intersect such a neighborhood—we have

$$(h_T, \mathfrak{p}) - (\mathbf{C}_{K_T}, \mathfrak{p}) = 0$$

almost surely. However, since $(h_t^n, \mathfrak{p}) - (\mathbf{C}_t^n, \mathfrak{p}) > \epsilon$ on the analogous event for each n —and the values $(h_t^n, \mathfrak{p}) - (\mathbf{C}_t^n, \mathfrak{p})$ converge to $(h_T, \mathfrak{p}) - (\mathbf{C}_{K_T}, \mathfrak{p})$ —we must have

$$(h_T, \mathfrak{p}) - (\mathbf{C}_{K_T}, \mathfrak{p}) > \epsilon$$

almost surely on this event, which implies that the event has probability zero. A similar argument holds with U in place of T .

The law of $(\mathbf{C}_t^n, \mathbf{p})$, up until the first time that $\widehat{\gamma}^n([0, t])$ intersects a fixed neighborhood of the support of \mathbf{p} , is a martingale whose largest increment size tends to zero in n . Thus, if we parameterized time by quadratic variation, then the limiting process $(\mathbf{C}_{K_T}, \mathbf{p})$ would be a Brownian motion up until the first time that K_T intersects the support of \mathbf{p} , and $(h_t^n, \mathbf{p}) - (\mathbf{C}_t^n, \mathbf{p})$ would converge in law to this limiting process with respect to the supremum norm on finite time intervals.

The height gap lemma implies that $(h_T, \mathbf{p}) = (\mathbf{C}_{K_T}, \mathbf{p})$ for any fixed stopping time (as discussed in the previous proof). We know that $(\mathbf{C}_{K_T}, \mathbf{p})$ is a Brownian motion when parameterized by $-E_t(\mathbf{p})$, as defined in Section 2.2, so the same must be true for (h_T, \mathbf{p}) , and the arguments in the proof of Lemma 2.8 then imply that the law of W is that of $\sqrt{4}$ times a Brownian motion (so that the Loewner evolution generated by W is SLE(4)) and $\lambda = \sqrt{\pi}/8$. Since the (h_t^n, \mathbf{p}) converge uniformly to (h_t, \mathbf{p}) , the same arguments imply that the W^n converge in law to W with respect to the supremum norm on compact intervals of time.

The fact that this driving parameter convergence holds for both forward and reverse parameterizations of the path implies that the convergence also holds in d_{STRONG} by the main result of [SS].

6 Remark on other contour lines

A more general problem than the one dealt with in this paper is to try to identify the collection of *all* chordal contour lines and contour loops of an instance of the GFF with *arbitrary* boundary conditions—and to show that they are limits of the chordal contour lines and contour loops of the piecewise linear approximations of the field. The second author is currently collaborating with Jason Miller on some aspects of this general problem.

For now, we only briefly mention one reason to expect Theorems 1.2 and 1.3 to also provide information about the general problem. Let h be a GFF with boundary conditions h_∂ as in Theorem 1.2. When $\psi \in H_s(D)$, it is not hard to see that the law of h is absolutely continuous with respect to the law of $h + \psi$.

Corollary 6.1. *In the context of Theorem 1.2, if $\psi \in H_s(\mathbb{H})$ and h is replaced by $h + \psi$, then as $r_D \rightarrow \infty$, the random paths $\widehat{\gamma}_D = \phi_D \circ \gamma_D$, viewed as Λ_C -*

valued random variables on (Ω, \mathcal{F}) , converge in probability (with respect to the metric d_{STRONG} on Λ_C) to an (Ω, \mathcal{F}) -measurable random path $\gamma^\psi \in \Lambda_C$ whose law is absolutely continuous with respect to the law of SLE(4).

We may interpret γ^ψ as the **zero contour line of $h + \psi$** . Note that if ψ is equal to a constant $-C$ on an open set, then the intersection of γ^ψ with this set can be viewed as (a collection of arcs of) a C contour line of h . By choosing different ψ functions and patching together these arcs, one might hope to obtain a family of height C contour loops. It requires work to make all this precise, however, and we will not discuss this here.

References

- [DS08] B. Duplantier and S. Sheffield. Liouville Quantum Gravity and KPZ. *ArXiv e-prints*, August 2008, 0808.1560.
- [Dub09] Julien Dubédat. SLE and the free field: partition functions and couplings. *J. Amer. Math. Soc.*, 22(4):995–1054, 2009.
- [Dud02] R. M. Dudley. *Real analysis and probability*, volume 74 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2002. Revised reprint of the 1989 original.
- [Gro67] Leonard Gross. Abstract Wiener spaces. In *Proc. Fifth Berkeley Sympos. Math. Statist. and Probability (Berkeley, Calif., 1965/66)*, Vol. II: *Contributions to Probability Theory, Part 1*, pages 31–42. Univ. California Press, Berkeley, Calif., 1967.
- [HBB10] C. Hagendorf, D. Bernard, and M. Bauer. The Gaussian Free Field and SLE₄ on Doubly Connected Domains. *Journal of Statistical Physics*, 140:1–26, July 2010, 1001.4501.
- [IK10] K. Izyurov and K. Kytölä. Hadamard’s formula and couplings of SLEs with free field. *ArXiv e-prints*, June 2010, 1006.1853.
- [Jan97] Svante Janson. *Gaussian Hilbert spaces*, volume 129 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1997.

- [Pom92] Ch. Pommerenke. *Boundary behaviour of conformal maps*. Springer-Verlag, Berlin, 1992.
- [Pro90] Philip Protter. *Stochastic integration and differential equations*, volume 21 of *Applications of Mathematics (New York)*. Springer-Verlag, Berlin, 1990. A new approach.
- [RS05] Steffen Rohde and Oded Schramm. Basic properties of SLE. *Ann. of Math. (2)*, 161(2):883–924, 2005.
- [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.
- [She03] Scott Sheffield. Gaussian free fields for mathematicians, 2003, arXiv:math.Pr/0312099.
- [She05] S. 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.
- [SS] S. Sheffield and N. Sun. Strong path convergence from loewner driving convergence. arXiv:1003.4675.
- [SS09] Oded Schramm and Scott Sheffield. Contour lines of the two-dimensional discrete Gaussian free field. *Acta Math.*, 202(1):21–137, 2009.
- [SW05] Oded Schramm and David B. Wilson. SLE coordinate changes. *New York J. Math.*, 11:659–669, 2005, <http://nyjm.albany.edu:8000/j/2005/11-31.html>.
- [Vai] N. Vaillant. Probability tutorials: Tutorial 13, Regular Measure. www.probability.net.