An Introduction to Imaginary Geometry
UROP+ Final Paper, Summer 2016

Mark Sellke
Mentored by Andrew Ahn
Project suggested by Scott Sheffield

August 31, 2016

Abstract

This expository paper is based on the theory of imaginary geometry, due to Sheffield and Miller. We aim to explain enough of the theory for the reader to understand the main constructions and some basic interactions between them in a short time. We suppress the more technical aspects when possible to focus attention on building geometric intuition.

Contents

1 Introduction

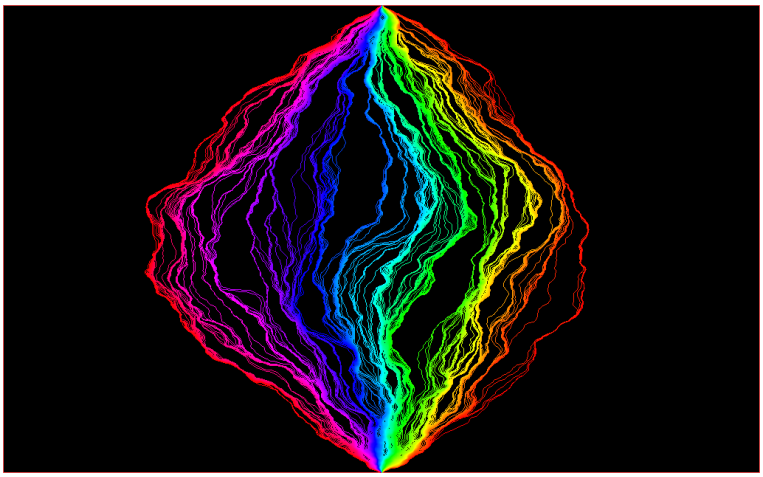

Gaussian free fields (GFF) and Schramm-Loewner Evolutions (SLE) are fundamental conformally invariant random objects in probability and statistical physics. In this paper we give a relatively quick and gentle introduction to the theory of imaginary geometry, in which SLEs are realized as straight lines in a geometry generated by a Gaussian free field.

The prerequisites for this paper are solid knowledge of complex analysis (especially the Riemann mapping theorem), probability (including basics of Ito calculus), and acquaintance with the Schramm-Loewner evolution. Familiarity with Gaussian free field is recommended, but we hope not strictly necessary. Good references for SLE include [BN16] and [Wer03], while good references for GFF include [She07] and the first chapter of [Ber16].

1.1 Initial Overview

Given a smooth function h on a planar domain $D \subseteq \mathbb{C}$, the vector field e^{ih} has well-defined, smooth integral curves. In fact, we may think of these curves as straight lines at angle 0, with straight lines at a general angle θ the integral curves of $e^{i(h+\theta)}$. One could define a smooth but imaginary analog of a Riemannian geometry in this way; since directions are globally defined, parallel transport along a closed loop fixes all angles, but may dilate lengths, in contrast to Riemannian geometry.

However, this paper is concerned not with the above situation, but with the problem of making sense of these integral curves, or **flow lines**, in the case that h is a Gaussian

free field (GFF). A Gaussian free field is a canonically random Schwartz distribution (or generalized function) on a domain D (always simply-connected for our purposes) which is invariant under conformal isomorphism. Unfortunately, GFFs cannot be realized as continuous functions, so it is highly unclear how to make sense of the "vector field" e^{ih}.

The flow lines of a GFF turn out to be another type of canonical, conformally invariant random object called an Schramm-Loewner evolution (SLE), or more precisely a variant known as $SLE_\kappa(\rho)$. These are highly non-smooth, and in fact have fractal dimension strictly between 1 and 2. One can also define *counterflow lines* which are slightly different and interact with flow lines in interesting ways. The main power of the theory comes from the ability to generate multiple flow and counterflow lines from a single instance h of a GFF.

An essentially degenerate case of imaginary geometry was studied in [SS09, SS13] in which *contour lines* or *level lines* of the GFF are studied; these correspond to sets on which the GFF is constant, and are also SLE curves. In this setting, it was shown that the level lines of successively finer discrete approximations to a continuum GFF converge to the continuum level lines. The analog of this for imaginary geometry is believed (and assumed in simulations) but has not been proved.

The theory of imaginary geometry was developed in detail, together with applications, in a series of paper by Miller and Sheffield ([MS12a, MS12b, MS12c, MS13]) . In particular, [MS12a] develops the theory of flow (and *counterflow*) lines started from a boundary point, [MS13] develops the theory for lines starting from interior points, and [MS12b, MS12c] apply the theory to reversibility of SLE processes. The interactions between flow (and counterflow) lines are carefully analyzed in a wide range of cases to great effect. In this paper we only present a subset of [MS12a].

The aim of this paper is to give the curious reader an expedited introduction to the basics of imaginary geometry. We hope to accomplish this by only outlining the technical points when they appear, allowing the geometric intuition to shine through. To the reader who completes this survey, the perhaps strange behavior of flow and counterflow lines should no longer be quite so mysterious.

1.2 Summary of Results

Although flow lines move monotonically to the right as the angle increases and do not self-intersect, there are many aspects of smooth differential topology which do not carry over to flow lines in imaginary geometry.

First, we can and will consider $e^{\frac{ih}{\chi}}$ as the vector field for any fixed χ, adding an extra degree of freedom. The flow lines look like SLE_κ for $\kappa < 4$ depending on χ. The level line case mentioned earlier corresponds to $\chi = 0, \kappa = 4$.

Flow lines of the same angle started from different boundary points can intersect, unlike the smooth case in which reversibility of integral curves immediately shows this is impossible. In fact, once two flow lines intersect, they immediately merge and stay together for all time.

Additionally, if we flow at angle θ and then at angle $\theta + \pi$, we do not retrace our previous path; indeed, this would be inconsistent with the above merging phenomenon. One might think this would be inconsistent with the previously mentioned conjecture that flow lines arise as scaling limits of well-behaved, classical flow lines, but this is not so because flow lines also do not vary continuously in the angle. In fact the set of points reachable in a straight line from a given boundary point, the **SLE fan**, has measure 0.

This interaction becomes more striking when we also consider *counterflow lines*. It turns out that given a GFF and $\kappa \in (0,4)$ such that the flow lines look like SLE_κ, we can also consider counterflow lines which look like $SLE_{\frac{16}{\kappa}}$. A single (zero-angle) counterflow line run from $b \to a$ contains a whole range of flow lines run from $a \to b$, namely those at angles $\theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$, and traverses each of them in reverse-order. The left and right boundaries of the counterflow line are precisely the extremal flow lines in the aforementioned range.

A single counterflow line actually is strictly larger than the union of the flow lines contained in it. However, we can reconstruct a counterflow line using *angle-varying flow lines*, in which we may change directions provided we are at all times travelling in a direction $\theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$; the set of points reachable in such a way is precisely the range of the counterflow line. In fact, only the extreme angles $\pm\frac{\pi}{2}$ are needed to generate the counterflow line.

1.3 Outline

In section 2 we begin by reviewing basics of GFF and SLE, and introducing the slightly more general $SLE_\kappa(\rho)$ processes which are what flow lines actually are. We discuss notation for assigning a certain type of boundary value to a GFF along a flow or counterflow line.

In section 3 we define the flow and counterflow lines. We remark now that the initial definition of flow and counterflow lines is rather indirect; it is not even clear that the GFF h determines them uniquely! We also state some useful technical results.

In sections 4 and 5 we explain the interactions between flow and counterflow lines, and use them to show that flow and counterflow lines are determined by the free field.

Finally, in section 6 we give a brief outline of further aspects.

2 Technical Preliminaries

2.1 Gaussian Free Field

Here we give the definition and basic background on the Gaussian free field. We briefly review basic notions in Liouville quantum gravity because of the close analogy with imaginary geometry.

The Gaussian Free Field h on a simply-connected planar domain D is a random Schwartz distribution or generalized function - recall that this means the pairing (h, f) is defined for all sufficiently smooth, compactly supported test functions f. We first define the 0-boundary GFF. To do this, we consider the Dirichlet energy $(f,f)_\nabla$ of a compactly supported, smooth function $f \in C_c^\infty(D)$. The Dirichlet energy is given by

$$(f,f)_\nabla = \int_{x \in D} |\nabla f(x)|^2 dx.$$

This defines an inner product (it actually assigns 0-energy to all constant functions, but we've assumed our functions are compactly supported so this is not an issue). Taking the Hilbert-space closure of $C_c^\infty(D)$ with this inner product gives the (Hilbert)-Sobolev space $H_0^1(D)$, roughly consisting of the functions f with gradient in $L^2(D)$ and with 0-boundary conditions. To define the Gaussian Free Field, we take a "canonical Gaussian random vector" for this Hilbert space H, by writing down an orthonormal basis (x_i),

generating i.i.d. standard Gaussians α_i and writing $h = \sum \alpha_i x_i$. It is easy to see that for any fixed $x \in H$, the inner product (x, h) should be a centered Gaussian with variance $|x|_H^2 = |x|_\nabla^2$. However, there is a caveat: h is not actually an element of H as we would clearly have $|h|_H = \infty$. Since Hilbert spaces are self-dual, this means that although $(h, f)_H$ exists for each fixed element $f \in H$ with probability 1, h almost surely does not pair with all elements simultaneously. However, we can still realize h as an element of the larger function space of Schwartz distributions.

One can also define a GFF with prescribed boundary data on D; to do this, one extends the boundary data to a harmonic function on D and deterministically adds this function to a 0-boundary GFF. In fact, we can do this for any harmonic function f on D, regardless of whether D has a nice boundary or if f extends sensibly to that boundary.

One very important property of the GFF is the **domain Markov property**, which states that if we condition on the complement U^c of a domain $U \subseteq D$, the law of $h|_U$ is a GFF with boundary data given by the harmonic extension of the boundary values. This justifies the definition of a GFF with non-zero boundary data. The proof requires a slightly technical ingredient, elliptic regularity, but is very intuitive.

Proposition 2.1. *Let $U \subseteq D$ open be a simply connected domain, and h a 0-boundary GFF on D. Then we may write $h = h_0 + \phi$ with h_0 a 0-boundary GFF on U and vanishing outisde U, and ϕ harmonic. Moreover h_0, ϕ are independent.*

Proof. The point is that we have an orthogonal decomposition of Hilbert spaces given by $H_0^1(D) = H_0^1(U) \oplus Harm(U)$, where $Harm(U)$ consists of functions harmonic in U. It is clear that $H_0^1(U)$ and $Harm(U)$ are disjoint except for 0 and orthogonal (by integrating by parts). To show they span, suppose $f \perp H_0^1(U)$. Then for any $\phi \in C_0^\infty(U)$, we have $0 = \langle f, \phi \rangle_\nabla = -(\Delta f, \phi)_{L^2(U)}$, and so $\Delta f = 0$ on U. Elliptic regularity implies that a harmonic distribution is actually a C^∞ function, and so $f \in Harm(U)$ as desired. The orthogonal decomposition makes the result clear when we recall the construction of h as a Gaussian random vector, since we can generate h using separate bases for $H_0^1(U)$ and $Harm(U)$.

Note that elliptic regularity isn't conceptually essential here; without it, we would obtain the same result with $Harm(U)$ defined to be the space of distributions harmonic in U. However, though it won't be apparent from this paper, it is legitimately useful to know that the conditional mean of $h|_U$ is a bona-fide function.

□

Though h is not a function, various properties of continuous functions may be extended to h; in this paper, as outlined before, we will investigate the flow lines of $e^{ih/\chi}$. One may also define the average values of h on circles (which is jointly continuous in center and radius) and use this to define $e^{\gamma h}$ for each $\gamma \in (0, 2)$, a random measure known as Liouville quantum gravity. Like the flow lines, Liouville quantum gravity has a fractal nature; it is supported on a set of dimension $2 - \frac{\gamma^2}{2}$.

The "definition" of Liouville quantum gravity above is similar to that of the flow lines, but LQG is real rather than imaginary. This results in a transformation rule for GFFs on domains $D, \tilde{D}$ related by a conformal isomorphism; we explain this here due to the analogous concept for imaginary geometry. Suppose $f : \tilde{D} \to D$ is a conformal isomorphism, h is a GFF on D and $\mu_h = e^{\gamma h}$ is the LQG measure on D. Then it turns out that the pullback $\mu_{\tilde{h}}$ of the measure μ_h to $\tilde{D}$ is given by LQG measure for $\tilde{h} = h \circ f + Q \log |f'|$, for $Q = \frac{\gamma}{2} + \frac{2}{\gamma}$. We say that f defines an isomorphism between the

random surfaces (D, h) and $(\tilde{D}, \tilde{h})$ when it preserves the LQG measures; that is, a *random LQG surface* is a pair (D, h) for a domain D and distribution h, up to isomorphism given by the above coordinate change formula. Note that if h were smooth, we would expect $Q = \frac{\gamma}{2}$. The reason that Q has another term comes from the precise definition of LQG, in which one approximates $e^{\gamma h}$ via small-ball circle averages; if one directly approximates h via circle averages, this will blow up badly as the circle radius goes to 0, and so the definition requires an explicit renormalization that explains the $\frac{2}{\gamma}$ term.

In the same vein, the imaginary surfaces we discuss in this paper have an analogous coordinate change formula which allows us to identify conformally isomorphic domains. In this case, though, pretending h is smooth actually does give the correct formula. One can make some intuitive sense of this by arguing that, whatever $e^{\frac{ih}{\chi}}$ means, it should not blow up and become infinite like $e^{\gamma h}$ might, so there's no reason that renormalization would come up. The transformation rule is thus $(D, h) \to (f^{-1}(D), h \circ f - \chi \arg f') = (\tilde{D}, \tilde{h})$; that is, under this transformation rule, flow lines of angle θ in $D, \tilde{D}$ are related by f. We will say more on this later.

Next, we make a few remarks on GFFs defined on finite graph. Let $G = (V, E)$ be a finite, simple graph and ∂ a non-empty distinguished set of vertices, which we view as the boundary. Let $D = V \backslash \partial$. Since $\int |\nabla f(x)|^2 dx = -\int f(x)(\Delta f(x))dx$, to define a GFF it suffices to define a discrete Laplacian for a graph, and this is given by $\Delta f(x) = \frac{1}{d(x)} \sum_{x \sim y} (f(y) - f(x))$ for $d(x)$ the degree. This then generates an inner product on the functions defined on the vertices of the graph, subjected to the constraint that they vanish on ∂. A (0-boundary) GFF, then, is simply a Gaussian vector in this finite-dimensional Hilbert space. Explicitly, we have $\langle f, f \rangle_\nabla = \sum_{x \sim y} |f(x) - f(y)|^2$, and so the values $f(x)$ are Gaussians biased to be close to their neighbors. We also retain a discrete domain Markov property in this discrete situation.

Figure 1: A Simulated Discrete GFF. Notice the roughness which prevents h from being defined as an honest function.

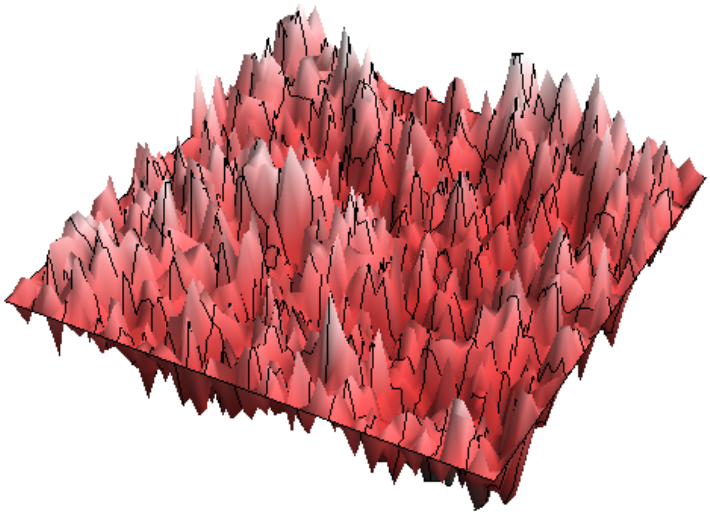

This construction allows one to represent a continuum GFF as a scaling limit of discrete GFFs on suitable lattices with mesh size decreasing to 0. More precisely, given a planar graph G with triangular faces, we may identify a function $f : G \to \mathbb{R}$ with the piece-wise linear interpolation of f to the faces of G. The piece-wise linear functions obtainable this way form a subspace onto which we may orthogonally project a given continuum GFF; taking finer meshes corresponds to projecting onto larger subspaces.

It is conjectured but has not been proven that the (piecewise smooth) flow lines of discrete approximations to a continuum GFF h converge to the continuum flow lines. [SS09, SS13] study what is essentially the degenerate case $\chi = 0$, in which the flow lines are just "level lines" on which h is "constant." In this case, convergence of discrete level lines to continuum level lines is proved, but proving an analogous result for flow lines has not been done.

As a final remark, GFFs are often defined with *free-boundary conditions*. In the discrete case this would amount to $\partial = \emptyset$, while in the continuum case we remove the compactly supported restriction on our test functions. Since constant functions have vanishing Dirichlet energy, we view a free-boundary GFF as being defined *modulo additive constant*. In practice, one sometimes normalizes a free-boundary GFF to have mean 0 on a suitable fixed set for convenience. One can also define *mixed boundary conditions* which are free on some but not all boundary arcs. In any case, free and mixed boundary conditions will not be relevant for this paper, and we will make no further mention of them.

2.2 SLE

Here we briefly review basics on the Schramm-Loewner Evolution (SLE_κ) processes, and discuss the generalized $SLE_\kappa(\rho)$ processes which appear as flow lines. For details on SLE we recommend the excellent survey [BN16].

Given an open, simply connected Jordan domain $D \subseteq \mathbb{C}$ with marked points $(a, b) \in \partial D$, the *Schramm-Loewner Evolution* (SLE) is a $\mathbb{R}^+$-parametrized family of random, non-self-crossing continuous paths γ from a to b. Since SLE is by construction invariant under conformal coordinate change, we typically take $(D, a, b) = (\mathbb{H}, 0, \infty)$, where $\mathbb{H} = \{x + iy : y > 0\}$ is the upper half-plane. (More precisely, one defines SLE first in $\mathbb{H}$, and then defines it in an arbitrary domain by conformal invariance.)

To define SLE, one considers not the path γ, but instead an $\mathbb{R}$-parametrized family of conformal isomorphisms $g_t : \mathbb{H}\backslash\gamma([0, t]) \to \mathbb{H}$. It is helpful to first consider a simple, deterministic model which is the degenerate case $SLE(0)$. Let

$$\gamma(t) = 2i\sqrt{t}.$$

By the Riemann mapping theorem, for each t there is a unique conformal isomorphism $g_t : \mathbb{H}\backslash\gamma([0, t]) \to \mathbb{H}$ such that $g_t(\infty) = \infty$ and $\partial_z g_t(\infty) = 1$. In fact in this case we have the exact formula

$$g_t(z) = \sqrt{z^2 + 4t}.$$

This means that we have

$$\partial_t g_t(z) = \frac{2}{g_t(z)}.$$

In other words, for $z \notin \gamma([0,t])$, the function $g_t(z)$ is the unique solution to the ODE $\partial_t g_t(z) = \frac{2}{g_t(z)}$ with initial condition $g_0(z) = z$, in other words the flow defined by the vector field $U(z) = \frac{2}{z}$. We may now define the set $K_t = \gamma([0,t])$ in terms of this flow, as the set of initial positions z_0 such that the flow $g_t(z_0)$ is not defined as a result of the flow curve hitting the real line. We note also that, having applied the conformal map g_t, we may view the remainder of this process as taking place in $g_t(\mathbb{H}\backslash\gamma([0,t])) = \mathbb{H}$; in other words, we may forget about the first t time units and start afresh. We might view this process as repeatedly eating away a small path starting at 0 and applying a conformal isomorphism back to $\mathbb{H}$ sending the tip of the path back to 0.

To define SLE_κ, we mimic the above construction, but instead of eating away only at 0 we eat from a continuously varying point W_t, called the **Loewner driving function**. In other words, after eating away each small amount of the upper half-plane and conformally mapping back to $\mathbb{H}$, we slightly shift the real point at which we eat away from $\mathbb{H}$. This also corresponds to horizontally translating the vector field U. To express this rigorously, we modify the ODE from before to incorporate the driving function, thus obtaining

$$\partial_t g_t(z) = \frac{2}{g_t(z) - W_t}$$

for a continuous function W_t with $W_0 = 0$. We then define the sets K_t as before, as the set of initial points z_0 for which $g_t(z_0)$ cannot be defined. In general, the sets K_t need not be described by a continuous path γ, but they do form a growth process of compact sets, and are parametrized by a natural quantity called the *half-plane capacity*.

The law of an SLE_κ curve is given by setting $W_t = \sqrt{\kappa}B_t$ for a standard Brownian motion B_t. Rohde and Schramm showed in [RS11] that in this case, K_t is the convex hull of the image $\eta([0,t])$ for a continuous function η, called the *SLE trace*.

Perhaps surprisingly, the qualitative behavior of η varies dramatically depending on the value of κ. η is always non-self-crossing, and for $\kappa \leq 4$ it is simple. For $\kappa \in (4,8)$ it is non-simple, and for $\kappa \geq 8$ it is space-filling. One can show in all cases that the trace $\eta(t)$ is transient, meaning that $|\eta(t)| \to \infty$ as $t \uparrow \infty$. Beffara showed that the Hausdorff dimensions satisfy $\dim_H(\eta([0,t])) = \min(1 + \frac{\kappa}{8}, 2)$.

2.3 $SLE_\kappa(\rho)$

The $SLE_\kappa(\rho)$ processes are variants of SLE with extra *force points* which affect the drift of the driving Brownian motion. In the absence of force points, one recovers ordinary SLE_κ. Given a vector

$$(x^L; x^R) = (x^{k,L} \leq x^{k-1,L} \leq ... \leq x^{1,L} \leq 0 \leq x^{1,R} \leq ... \leq x^{m,R})$$

of real positions and a vector

$$\rho = (\rho^L; \rho^R) = (\rho^{k,L}, ..., \rho^{1,L}, \rho^{1,R}, ..., \rho^{m,R})$$

of real weights, we define the $SLE_\kappa(\rho)$ process by the systems of stochastic differential equations

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

$$W_t = \sqrt{\kappa} B_t + \sum_{i,q} \int_0^t \frac{\rho^{i,q}}{W_s - V_s^{i,q}} ds.$$

The singularities in the SDEs result in a couple of technical points that we will freely gloss over, which can be found in full detail in the original paper [MS12a]. An important consequence of these technical points to keep in mind is that the force points $V^{i,L}$ to the left of W_t (say) will always stay to the left; upon colliding with W_t, they will bounce off but never cross over to the other side. Note that one or both of $x^{1,L}, x^{1,R}$ can be 0; in this case, we think of the force points as starting at 0^- and 0^+, so that they always bounce off on the left/right side of W_t.

Away from collisions, W_t is a Brownian motion with bounded drift. By Girsanov's theorem and absolute continuity, this fact combined with the corresponding result for ordinary SLE_κ means that on compact time-intervals away from collisions, K_t is generated by a continuous path. It turns out that $SLE_\kappa(\rho)$ is always continuous, even at boundary intersections. This results was first proved in full generality using imaginary geometry in [MS12a], and we will not use it in this paper. We mention it now for the psychological comfort it may provide.

Theorem 2.2. *$SLE_\kappa(\rho)$ is generated by an a.s. continuous curve η, at least up to the continuation threshold discussed below.*

The SDEs above surely look very strange and arbitrary to the reader. We add in a drift term to the driving function W_t, which is dependent on repelling force points which themselves drift away from W_t. The point, as will be explained in more detail later, is that we will view $SLE_\kappa(\rho)$ curves as flow lines for a GFF h with boundary data which is constant on each interval $(x^{i+1,L}, x^{i,L})$ or $(x^{\ell,R}, x^{\ell+1,R})$ with values depending on $(\rho^L; \rho^R)$ in a simple way. The drift of the force points is identical to the drift for the Loewner flow with driving function W_t, so the picture is that of a fixed domain, namely $\mathbb{H}$ with some stationary marked force points on the boundary, with a growing flow line $\eta([0,t]) = K_t$ exploring the domain.

2.4 The Continuation Threshold

A technical point we will discuss now is that of the *continuation threshold* for $SLE_\kappa(\rho)$. The continuation threshold isn't essential for this paper due to the restricted cases we consider, but it is certainly good to be aware of this notion. The story is that while ordinary SLE_κ is defined for all positive real times, $SLE_\kappa(\rho)$ may not be, depending on the force points and the evolution of the process. The *continuation threshold* for $SLE_\kappa(\rho)$ is thus the (possibly infinite) time after which the process stops being defined (i.e. we cannot continue past this time). Below we outline the situation.

To understand why we might run into trouble in continuing the $SLE_\kappa(\rho)$ SDEs, we need to think a bit about the SDEs from a technical standpoint. The force points are defined by the same ODE in W_t, and so away from W_t they will never merge. Conversely, if multiple force points on the same side of W_t simultaneously collide with W_t, they will stay together for all time by uniqueness, and we can think of them as having merged together (but note that the force points drift deterministically away from W_t, so collisions between opposite-sided force points will never occur at non-zero times). The continuation threshold comes from the technical problem of making sense of the SDEs near collision

points. If W_t collides with a single force point $V^{i,q}$, then $|W_t - V_t^{i,q}|$ will look like a rescaled *Bessel process* of dimension δ, i.e. a solution to the SDE

$$dX_t = dB_t + \frac{\delta - 1}{2X_t} dt.$$

To be precise, we have $\delta = 1 + \frac{2(\rho^{i,q}+2)}{\kappa}$. (The word "dimension" comes from the fact that when $\delta \in \mathbb{N}$, a Bessel process of dimension δ is simply the Euclidean norm of a Brownian motion in δ dimensions.) Bessel processes are well-understood, and it turns out that they behave just fine for $\delta > 1$, i.e. for $\rho > -2$, and so if $\sum_{\text{Force points colliding with } W_t} \rho^{i,q} > -2$ then we can safely continue the process past the collision. We therefore define the continuation threshold to be the first time at which this is not the case. This means that the continuation threshold will always happen at a collision of W_t with a new force point, i.e. a collision with the boundary.

2.5 Boundary Values and Notation

Here we discuss boundary values for GFFs along interior flow lines. In the next section we will use these ideas to define the flow line coupling. One might expect that this would be relevant because, as we said before, flow lines result from GFFs with piece-wise constant boundary conditions. If we stop a flow line at $\eta(\tau)$ and re-interpret our domain to be $D = \mathbb{H} \backslash \eta([0, \tau])$, a flow line started from the tip $\eta(\tau)$ at the same angle should just be the continuation of the original flow line. This means we should be able to say something about the boundary values of h on $\eta([0, \tau])$ so that we can define a flow line started from $\eta(\tau)$ in a nice way.

We first define an auxillary constant χ for each $\kappa > 0$ (though the reason for making this definition will not be clear until later). We set

$$\chi(\kappa) = \frac{2}{\sqrt{\kappa}} - \frac{\sqrt{\kappa}}{2}.$$

For an interior curve η (smooth for now), we want to consider boundary conditions on h that change linearly as a function of the winding angle, at least on each piece of η. In figure 2, we illustrate a convenient notation for expressing this.

In the case that η is non-smooth, it is not completely clear how one would define boundary values as above. The answer is by describing instead the harmonic extension to η^c; recall that this is what we really mean when we prescribe boundary values for a GFF. Now, if η were a simple flow-line of $e^{i(\frac{h}{\chi}+\theta)}$ for h smooth, with f_t the mapping-out function for $\eta([0, t])$, then we would have

$$h(x) \approx -\theta - \chi \arg f_t'(z)$$

for z just to the left side of $x \in \eta([0, t])$ and

$$h(x) \approx -\theta + \chi\pi - \chi \arg f_t'(z)$$

for z just to the right side, as one can easily see with a little thought.

Since $\arg f_t'$ is harmonic (this is just because f_t' is non-vanishing and holomorphic), it suffices to describe the harmonic extension of the angle-boundary values to η^c. We can do this canonically by mapping-out by f to reduce to the case of the whole upper half-plane.

Figure 2: For imaginary geometry, one constantly works with boundary conditions which change linearly with the winding angle of the boundary on each of several pieces of η. Here the different pieces are given by the arcs between adjacent black dots. More precisely, turning by θ to the left corresponds to an increase in boundary value (on both sides) by $\chi\theta$. When we label a straight boundary arc $\underset{\sim}{x}$, we mean that that arc has boundary value exactly x, and that the rest of that arc-piece is labelled according to the turning rule as shown. Thus the figures depicted side-by-side represent the same boundary conditions.

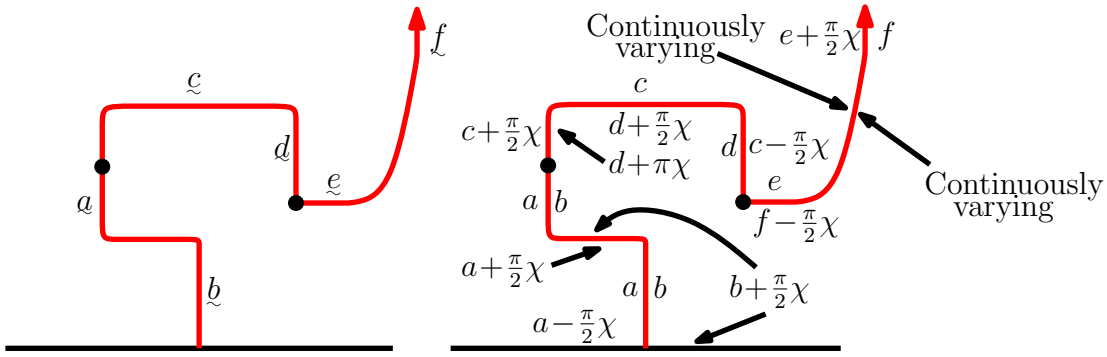

As a word of caution, imaginary geometry departs from this smooth analogy in an important way. Namely, the left and right boundary values on a flow line differ by constants depending on κ from what the above smooth story would suggest. For example, even in the degenerate case $\chi = 0$, in which flow lines correspond to "level lines" on which h is "constant," the boundary values for h differ by a constant amount on the left and right sides of the level lines ([SS09],[SS13]). Of course, if h were continuous the left and right boundary values would agree, so this discrepancy can be viewed as accounting for the "roughness" of h.

3 Foundational Results

3.1 The Coupling

Here we outline the construction of the flow lines which will be interpreted as $e^{ih/\chi}$. We will later show that they are determined a.s. by h, but this takes more work and for now we only construct couplings of $SLE_\kappa(\rho)$ processes with h. We do this by defining the conditional law of h given a stopped $SLE_\kappa(\rho)$ by defining the boundary values of h on the SLE. This gives a valid coupling of h with $SLE_\kappa(\rho)$, and the boundary values can be interpreted as a justification for our analogy with classical vector fields. We will also state the martingale characterization for $SLE_\kappa(\rho)$ which allows us to reverse the proof; this will be useful later because it allows us to show that flow lines in certain domains are also flow lines in different domains, allowing us to switch between domains.

In the coupling theorem, we use constants λ, λ' defined as follows:

$$\lambda = \lambda(\kappa) := \frac{\pi}{\sqrt{\kappa}},$$

$$\lambda' = \lambda(\kappa') = \lambda\left(\frac{16}{\kappa}\right) = \frac{\pi\sqrt{\kappa}}{4}.$$

The identity

$$\lambda' = \lambda - \frac{\pi\chi}{2}$$

will sometimes be useful.

We remark that the relationship between $\kappa, \chi, \lambda, \lambda'$ is designed precisely for the theorem below to work. Though we do not actually work through the computation here, we explain the condition one needs to check.

Theorem 3.1. *Fix $\kappa > 0$ and a vector $(\rho^L; \rho^R)$. Let K_t be the hull generated by the $SLE_\kappa(\rho)$ process for the parameters above, and let $\mathfrak{h}_t^0 : \mathbb{H} \to \mathbb{R}$ be the harmonic function with boundary values*

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

Here (for convenience of notation in the boundary formulas above) we assume $\rho^{0,L} = \rho^{0,R} = 0, x^{0,L} = 0^-, x^{0,R} = 0^+, x^{k+1,L} = -\infty, x^{\ell+1,R} = +\infty$.

Now with $\chi = \frac{2}{\sqrt{\kappa}} - \frac{\sqrt{\kappa}}{2}$, let

$$\mathfrak{h}_t(z) = \mathfrak{h}_t^0(f_t(z)) - \chi \arg f_t'(z).$$

Let $(\mathcal{F}_t)$ be the filtration generated by the $SLE_\kappa(\rho)$ processes $(W, V^{i,q})$. Then there exists a coupling (K, h) for $h = \tilde{h} + \mathfrak{h}_0$ and $\tilde{h}$ a zero boundary GFF on $\mathbb{H}$ such that for any stopping time τ, the conditional law of $h|_{\mathbb{H}\setminus K_\tau}$ given $\mathcal{F}_\tau$ is the law of $\mathfrak{h}_\tau + \tilde{h} \circ f_\tau$.

Remark. In the statement of the coupling theorem, we assume that h has piece-wise constant boundary conditions. One can easily extend this to general boundary conditions which are constant in a neighborhood of the initial point 0, so long as η stays away from non-piecewise-constant boundary spots. This is essentially because behavior of a GFF on separated sets A, B is mutually absolutely continuous with respect to the product measure; we say $h|_A, h|_B$ are *almost independent*. We won't use non-piecewise-constant boundary conditions in this paper, however.

Before discussing the proof of theorem 3.1, let's unpack what this (fundamental) theorem says. As mentioned before, we will construct flow (or counterflow) lines such that h has nicely described boundary conditions on the flow line $\eta([0, \tau])$; this theorem asserts that we can define a coupling of flow lines with a GFF h so that this actually holds. As stated in Theorem 3.2, the desired boundary conditions given a flow line are enough information to uniquely determine the joint law for the coupling. In the later sections, among other things, we will show that η is actually determined by h.

Figure 3: The function $\mathfrak{h}^0_\tau$ in Theorem 3.1 is the harmonic extension of the boundary values on the right panel. The function $\mathfrak{h}_\tau = \mathfrak{h}^0_\tau \circ f_\tau - \chi \arg f'_\tau$ is the harmonic extension on the left side. That the vertical boundary values on the left side are $\pm\lambda'$ follows from the identity $\lambda' = \lambda - \frac{\pi\chi}{2}$.

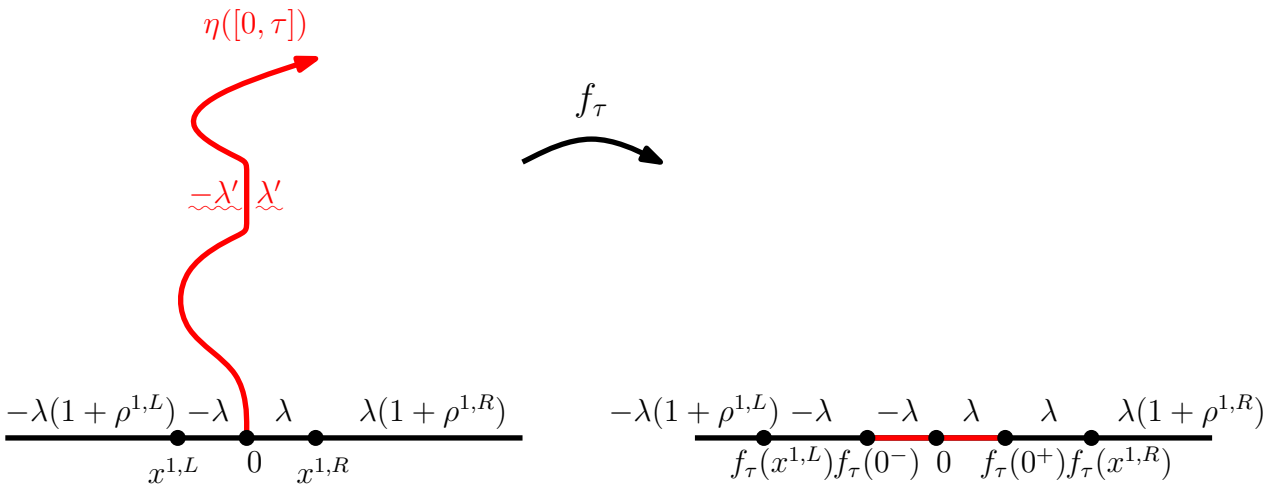

Proof. We only give an outline, explaining the key idea which is intuitive and rather pretty. Fix a function $F \in C_0^\infty(\mathbb{H})$ and assume that K_τ does not intersect the support of F. For the boundary conditions prescribed, the law of (h, F) is a Gaussian with mean and variance $\mu(F), \sigma^2(F)$, and knowing the (one dimensional) law for each of these Gaussians uniquely determines the law of h. Now $\mathfrak{h}_0$ is the conditional expectation of h, so

$$(\mathfrak{h}_0, F) = \mathbb{E}[(h, F)].$$

As t increases up to τ, it turns out that $(\mathfrak{h}_t, F)$ evolves as a time-changed Brownian motion, and we can describe explicitly the quadratic variation, i.e. the rate of passage of Brownian time. This means that $\mathfrak{h}_\tau(f_t(z))$ is a Gaussian with some variance. Since $\tilde{h}$ is chosen independently, for h to agree in law with $\mathfrak{h}_\tau + \tilde{h} \circ f_\tau$, what we want is for $(\tilde{h} \circ f_\tau, F)$ to be another Gaussian with complementary variance, so that we recover the correct law when paired with F.

In fact, the above is exactly what happens. Verifying all of this is a tedious but routine computation in Ito calculus which we won't carry out here. Making the numbers work in this argument is the reason for the strange looking boundary conditions and constants in the theorem.

There is a bit more work to be done to show such a coupling actually exists, and we won't go through it in detail; the above argument can be thought of as checking the main consistency condition. A summary is that one first finds a coupling that works for a given finite collection of stopping times, and then uses some abstract nonsense (in the form of Prohorov's theorem) to find a coupling that works for a given countable collection of stopping times. This countable collection can be taken to be dense, which suffices to give an arbitrary stopping time by approximating from above and using backwards-martingale convergence.

□

We now can define flow and counterflow lines. Note the differing sign conventions; the reason for this is that $\chi(\kappa) = -\chi(\frac{16}{\kappa})$, so the sign convention below allows us to naturally couple κ-flow lines with κ'-counter flow lines. See the next section for more on this point.

Definition 1. For $\kappa \in (0,4)$ we define a **flow line** η of h (at angle 0) to be an $SLE_\kappa(\rho)$ curve coupled with h as in Theorem 3.1. We define a flow line of angle θ to be a flow line of $h + \chi\theta$.

For $\kappa' \in (4,\infty)$ we define a **counterflow line** η' of h (at angle 0) to be an $SLE_{\kappa'}(\rho)$ curve coupled with the *negation* $-h$ as in Theorem 3.1. We define a counterflow line of angle θ to be a counterflow line of $h + \chi\theta$.

Remark. As a technical point, note that the ambiguity in the choice of branch for the arg function means that if we change θ by 2π, we can get the same flow line by deterministically shifting $\arg f_t'$. In practice this ambiguity will not be an issue because diagrams will be labelled with a precise boundary value on at least one boundary arc, so the normalization of arg will be implicit in the diagram.

Remark. How do we deal simultaneously with multiple flow lines? Once we show that flow lines are determined by h this will be a non-issue, but since flow lines are only constructed as a coupling for now the correct a priori definition might not be so clear. Figuring this out now is actually important, because we'll need to use flow/counterflow interactions to show that flow lines are determined by h.

The solution is to use a canonical way to combine multiple couplings with a fixed object (in this case h) by making the flow lines *conditionally independent* given h. That is, to sample multiple flow lines, we first sample h and then independently sample the flow lines from the conditional law given h. This convention is useful primarily due to local set results which we only mention in the epilogue, so don't worry about it too much for now.

3.2 Imaginary Random Surfaces

Here we expand on the point in the previous subsection on naturally coupling flow and counterflow lines together.

The coupling theorem motivates the notion of an **imaginary random surface**. Just as a Riemannian manifold (say) is a surface with a metric structure *up to isomorphism*, an imaginary random surface consists of a planar D together with a Schwartz distribution h on it, up to isomorphism. The equivalence relation is given by the coordinate change formula

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

Note that we implicitly fix χ here, so we really have a parameterized notion of χ-random surface for each χ value. The reader will note also that, by the chain rule, the coordinate change formula given above is functorial as one would hope; given $f : D \to D'$ and $f' : D' \to D''$, changing by f and f' separately or simultaneously give the same result.

The point here is that, as alluded to previously, $SLE_\kappa(\rho)$ processes satisfy the same conformal Markov property as SLE (so long as we account for the force points). Hence,

Figure 4: To extend the conformal Markov property for $SLE_\kappa(\rho)$ to the coupling with h, we must change the boundary data of h to account for the winding of f_τ. This leads to the notion of an imaginary random surface up to conformal equivalence.

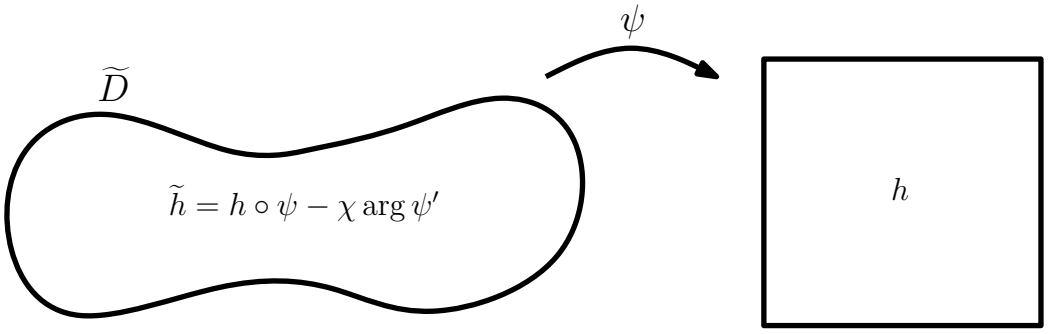

when we stop a flow line η at $\eta(\tau)$ and ask what it will do next, this is equivalent to beginning a flow line at $\eta(\tau)$ from scratch by viewing $\eta([0,\tau])$ as part of the boundary. That is, by taking the conformal isomorphism f_t from $\mathbb{H}\backslash K_t$ to $\mathbb{H}$ that preserves ∞ and takes $\eta(\tau) \to 0$ we should be able to define the continued flow lines starting from $\eta(\tau)$. At the level of the $SLE_\kappa(\rho)$ SDE, this is equivalent to stopping the SDE at some time and then pretending the state at that time is the initial condition.

Since we don't create any new force points as the SDE evolves, theorem 3.1 tells us what the mapped-out boundary data should be for the boundary segments of $\partial\mathbb{H}$ corresponding to $\eta(\tau)$, namely $\pm\lambda$. The conditional expectation formula

$$\mathfrak{h}_t(z) = \mathfrak{h}_t^0(f_t(z)) - \chi \arg f_t'(z)$$

from Theorem 3.1 now leads to the imaginary random surface coordinate change formula so that the coupling (h, η) is jointly conformally invariant.

As mentioned previously, for counterflows with $\kappa' > 4$ we always implicitly negate the sign of the GFF h. Because $\chi(\kappa) = -\chi(\kappa')$, this means that the values of χ corresponding to κ, κ' (counter)flow lines agree. Hence, the same "up to isomorphism" notion of a random surface transforms both the dual flow and counterflow lines functorially, so that we may sensibly couple them together on the same imaginary random surface (up to isomorphism). As we will soon see, a great deal of the power of imaginary geometry comes from exploiting this interaction.

3.3 Martingale Characterization of $SLE_\kappa(\rho)$

In proving the coupling theorem 3.1, the key fact was that $\mathfrak{h}_t(z)$ was a continuous local martingale; this allowed us to parametrize inner products with test functions by Brownian motions. The martingale characterization of $SLE_\kappa(\rho)$ asserts that this property characterizes $SLE_\kappa(\rho)$, and is useful for technical reasons.

Theorem 3.2. *Suppose η is a random, continuous curve in $\overline{\mathbb{H}}$ with continuous Loewner driving function W_t. For given $x^{i,q}, \rho^{i,q}$ and with $V_t^{i,q}$ the images of the force points $x^{i,q}$ under the Loewner maps, let $\mathfrak{h}_t$ be the harmonic function defined as in the statement of Theorem 3.1. Then $W_t, V_t^{i,q}$ evolve as an $SLE_\kappa(\rho)$ process iff $\mathfrak{h}_t(z)$ is a continuous local martingale in t for each fixed $z \in \mathbb{H}$ until z is absorbed by K_t.*

The proof is a bit technical and we don't discuss it. We do emphasize that it must be *assumed* that W_t is continuous. The main use of this theorem is that it allows us to relate flow lines in overlapping domains.

3.4 Flow Lines given Flow Lines are Flow Lines

The following result provides a logical foundation for the intuitive, geometric arguments of the next two sections. To make it completely rigorous one needs some results on local sets of the Gaussian free field, so we advise the reader not to worry too much about the details for now. This sort of result is true in more generality - see the epilogue for some commentary.

Theorem 3.3. *Let $a, b \in \partial\mathbb{H}$ be distinct boundary points and consider flow or counterflow lines η_a, η_b started at a, b. Suppose $\eta_a([0, \tau_a])$ a.s. does not accumulate at b, and K_{τ_a} the hull generated. Then η_b evolves as a (counter)flow line in $H_{\tau_a} = \mathbb{H}\backslash K_{\tau_a}$ with (counter)flow line boundary conditions on $\eta([0, \tau_a])$ and the original boundary conditions on $\partial\mathbb{H}$, at least until η_b hits η_a or $\partial\mathbb{H}$.*

Proof. It's easy to see that η_b has continuous Loewner driving function in H_{τ_a}. Furthermore, given $\eta_b([0, \tau_b])$, it is true that the conditional law of h on the remaining domain is given by a GFF with separate flow line boundary conditions on η_a, η_b (this is actually non-obvious without the machinery of local sets, and relies on the fact that η_a, η_b are separated by a positive distance and are coupled to be conditionally independent). So we know that

1. $h|_{H_{\tau_a}}$ has conditional law a GFF with flow line boundary conditions.
2. η_b has a continuous driving function.
3. Given $\eta_b([0, \tau_b])$, h has conditional law given by combining the flow line boundary conditions so long as η_b has not yet hit the boundary.
4. The conditional expectation of $h|_{H_{\tau_a}}$ given $\eta_b([0, t])$ is a continuous local martingale (essentially by definition a conditional expectation with respect to a filtration is a local martingale, and continuity is not hard to see).

The martingale characterization of $SLE_\kappa(\rho)$ now implies that the conditional law of $(h_{\mathbb{H}_{\tau_a}}, \eta_b([0, \tau_b])$ given K_{τ_a} is precisely that of a flow line in H_{τ_a}.

□

4 Flow and Counterflow Interactions 1

In this section and the next, we derive some properties of flow/counterflow lines in the case that the flow lines are not boundary intersecting. The results of this section are best stated on a strip $\mathcal{S} = \mathbb{R} \times [0, 1]$; we identify $\mathcal{S} \sim \mathbb{H}$ via a conformal isomorphism $\psi : \mathcal{S} \to \mathbb{H}$ with $\psi((0, 0)) = 0, \psi(+\infty) = \infty, \psi(-\infty) = -1$. We denote the upper and lower boundaries $\mathbb{R} \times \{0\}, \mathbb{R} \times \{1\}$ by $\partial_U\mathcal{S}, \partial_L\mathcal{S}$, respectively. We will repeatedly invoke continuity of the flow and counterflow lines, and we remark that this is not circular because in all cases considered here, because the $SLE_\kappa(\rho)$ processes avoid the boundary

and hence are absolutely continuous with respect to ordinary SLE_κ on compact time intervals.

As in figure 5, a putting a single force point of weight ρ at $-1 \in \partial\mathbb{H}$ corresponds to constant boundary data $a = -(1+\rho)\lambda - \pi\chi$, as follows from the imaginary random surface transformation rule.

Figure 5: For convenience, we identify $\mathcal{S} \sim \mathbb{H}$ via the conformal map $\psi^{-1} : \mathbb{H} \to \mathcal{S}$ depicted here, which sends $(0, -1, \infty) \to (0, -\infty, +\infty)$ and changes boundary data according to the transformation rule.

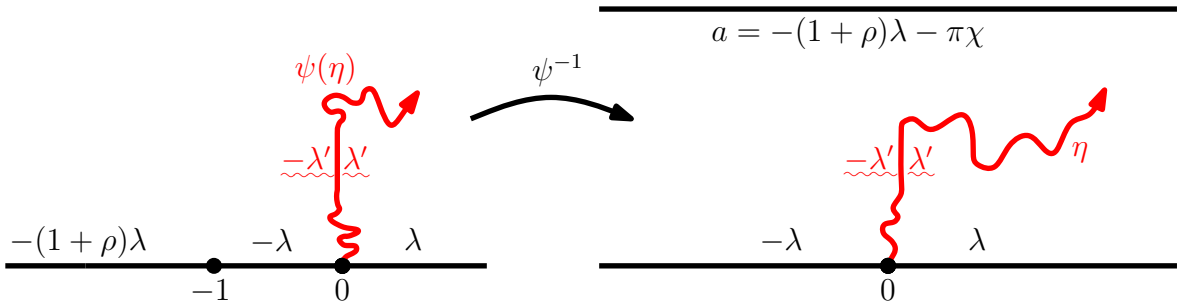

4.1 Boundary Value Lemmas

We will make heavy use of some lemmas which are essentially due to Dubédat in [Dub07] where a precise analysis of the $SLE_\kappa(\rho)$ SDE is undertaken. We will implicitly use Theorem 3.3 to apply these lemmas after conditioning on a flow line; the recipe for the proofs will be to condition on a flow or counterflow line $\eta_1([0, \tau_1])$, use a conformal map to make η_1 part of the boundary of $\mathcal{S}$, and apply one of the lemmas below to another flow or counterflow η_2 to glean information on the interaction of η_1, η_2. We encourage the reader not to worry to much about the precise boundary value calculations.

Figure 6: The boundary-hitting behavior of an $SLE_\kappa(\rho)$ flow line can be precisely determined in the special case depicted here by the boundary values. Depending on whether $a \geq \lambda$, $a \in (-\lambda, \lambda)$, or $a \leq -\lambda$, the flow line accumulates first at $-\infty$, $\partial_U\mathcal{S}$, or $+\infty$.

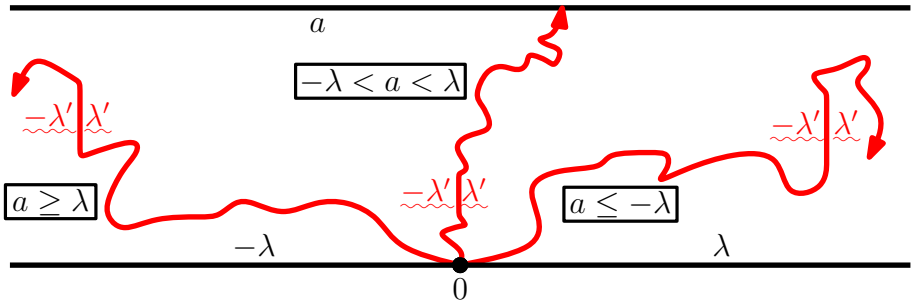

Lemma 4.1. *Suppose that h is a GFF on $\mathcal{S}$ with boundary data as depicted in Figure 6 and let η be the flow line of h starting from 0. If $a \geq \lambda$, then η a.s. accumulates at $-\infty$ and if $a \leq -\lambda$ then η a.s. accumulates at $+\infty$. In both cases, η a.s. does not intersect $\partial_U \mathcal{S}$. If, on the other hand, $a \in (-\lambda, \lambda)$, then η a.s. accumulates in $\partial_U \mathcal{S}$.*

The case $a \geq \lambda$ holds more generally if the boundary data is piece-wise constant with all values at least λ with finitely many changes, and similarly for the $a \leq \lambda$ case.

Finally, the analogous statements hold for counterflow lines $SLE_{\kappa'}(\rho')$ with λ replaced by λ'.

Lemma 4.2. *Suppose a GFF h on $\mathcal{S}$ has boundary data as depicted in Figure 7 with $a \leq -\lambda, c \in (-\lambda, \lambda), b \geq \lambda$. Then a flow line started from 0 a.s. first accumulates in the interval (z_0, z_1).*

Similarly, in Figure 8 with boundary data as indicated, a flow line started from 0 a.s. first accumulates at z_0.

Figure 7: When the upper boundary has boundary data as shown in this diagram, the flow line will almost surely first accumulate on the boundary in the interval (z_0, z_1). This is intuitive given the previous figure because the outside upper boundaries should push η toward the middle.

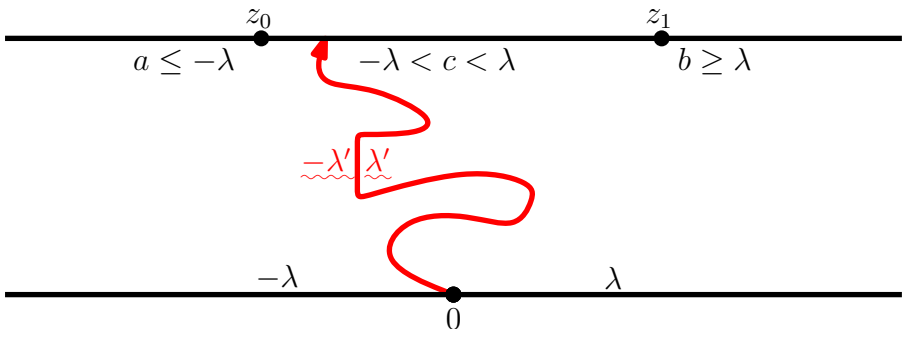

Figure 8: In this figure, we dispense with the middle intermediate boundary segment from the previous figure and only have very large and very small boundary data. In this case, the flow line will almost surely accumulate first exactly at z_0.

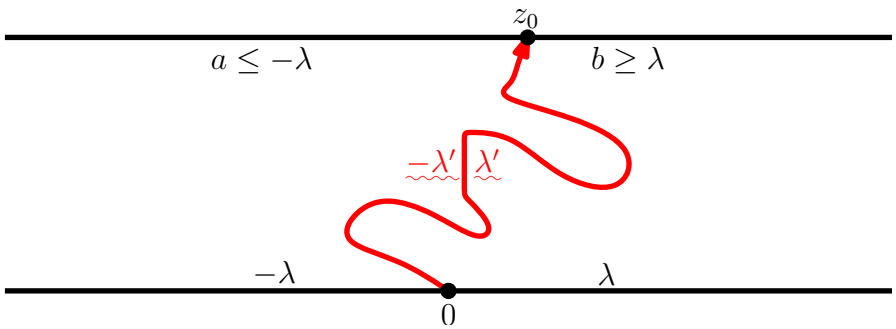

These lemmas are not hard to guess given the theory of Bessel processes. As explained earlier, when W_t is close to a force point $V^{i,q}$, the distance between them is essentially a

Bessel process, and for such processes it is known for which parameters it is possible to hit zero, i.e. for collisions to occur. Some easy arithmetic is all that is needed to translate this information into the present situation. However, the full proof is rather technical, because a first-hitting of a boundary interval corresponds to a merging of multiple force points, so the Bessel intuition is not obviously valid. We will take the lemmas on faith.

4.2 Basic Interactions

Finally we can prove some neat results on flow lines. We begin by showing that counterflow lines η' contain a whole range of flow lines.

Theorem 4.3. *As in Figure 8, fix $\kappa \in (0,4)$ let $\eta = \eta_\theta$ be the flow line at angle $\theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$ for h started at 0, and η' a conditionally independent counterflow line for h started at z_0. Let τ be a stopping time for η. Then η' a.s. first hits $\partial_L \mathcal{S} \cup \eta([0,\tau])$ at the tip $\eta(\tau)$. In particular, η' contains η and traverses it in reverse order; that is, for $s < t$, η' first hits $\eta(t)$ before $\eta(s)$.*

Proof. By Theorem 3.3, the conditional law of η' given $\eta([0,\tau])$ is a counterflow in the domain $\mathcal{S}\backslash\eta([0,\tau])$ with flow line boundary conditions on η, at least assuming η has not yet accumulated at the boundary. Mapping $\eta([0,\tau])$ to a boundary segment reduces this to an application of lemma 4.2.

To show the last result for all positive reals (s,t) it suffices to invoke continuity. □

Figure 9: When we stop a flow line $\eta = \eta_\theta$ for $\theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$ at $\eta(\tau)$, a conditionally independent counterflow line η' of $h - \frac{\pi}{2}\chi$ first meets $\partial_L \mathcal{S} \cup \eta([0,\tau])$ at $\eta(\tau)$ by lemma 4.2. This implies that η' contains all of η and traverses η in reverse chronological order. Here we depict the case $\theta = -\frac{\pi}{2}$. In this case, mapping out the right side of η to the boundary gives a boundary value of $\lambda - \pi\chi + \frac{\pi\chi}{2} = \lambda'$ by the transformation rule, so that lemma 4.2 barely guarantees that η' hits $\eta([0,\tau])$ first at $\eta(\tau)$. Symmetrically, in the case $\theta = \frac{\pi}{2}$, the left side barely works. This means that lemma 4.2 applies to all $\theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$.

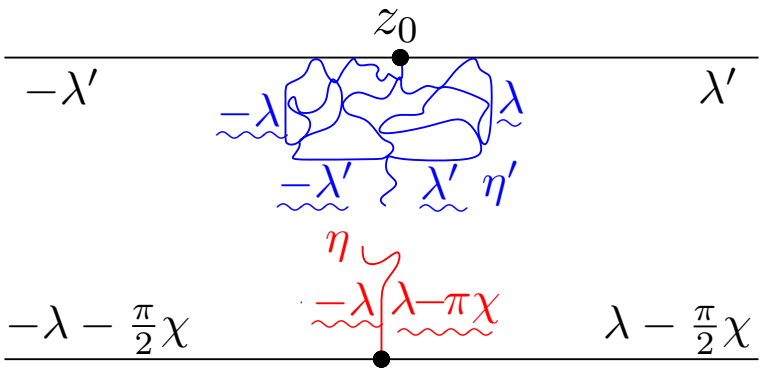

We now show that the left and right boundaries of η' are simply the extremal flow lines $\eta_{\frac{\pi}{2}}$ and $\eta_{-\frac{\pi}{2}}$, respectively. As a sanity check, one easily sees that the left boundary of $\eta_{\frac{\pi}{2}}$ has the same left-boundary conditions as η' and similarly for the other side. Nevertheless this result is still rather striking.

Theorem 4.4. *Fix $\kappa \in (0,4)$ and let h be a GFF on $\mathcal{S}$ with boundary data as in Figure 9. Let $\eta = \eta_{\frac{\pi}{2}}$ be a flow line of h started at 0 at angle $\frac{\pi}{2}$, and η' a conditionally independent counterflow line of h started at z_0, at angle 0. Almost surely, η is the left boundary of η'. Similarly, $\eta_{-\frac{\pi}{2}}$ is almost surely the right boundary of η'.*

Figure 10: In Theorem 4.4, we show that the left boundary of a counterflow line η' is almost surely a flow line η at angle $\frac{\pi}{2}$. It suffices to show that η' is never strictly to the left of η. Using our standard lemmas, one shows that for any stopping time τ' for η', the flow line η must first hit $\eta'([0,\tau'])$ on the left side of the boundary. Topology now shows that for $\eta'(\tau')$ to be on the left side of η, the flow line η must wrap under and around $\eta'(\tau')$ before hitting the right side of the counterflow. But this means that at time τ', the set of t such that η' has hit $\eta(t)$ is not an upward-closed interval, contradicting our earlier result.

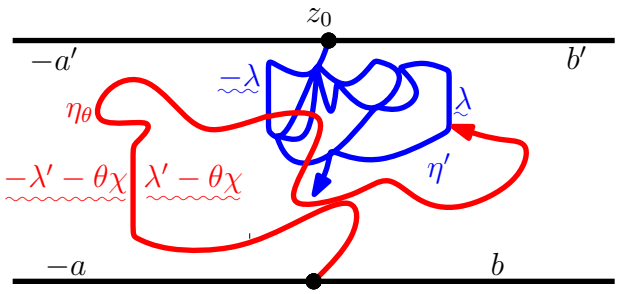

Proof. We just do the left-side case. We already know that η is contained in η', so it suffices to show that η' is never strictly to the left of η.

Let τ' be a stopping time for η'. We will show $\eta'(\tau')$ is on or to the right of η. Indeed, suppose not, and fix $\eta'([0,\tau'])$. An application of lemma 4.2 above shows that η first hits $\partial_U\mathcal{S} \cup \eta'([0,\tau'])$ either to the left of z_0 in $\partial_U\mathcal{S}$ or on the left side of $\eta'([0,\tau'])$. In the first case we are done. In the second case, for $\eta'(\tau')$ to be to the left of η, the curve η must wrap underneath and around $\eta'(\tau')$ and then hit η' on the right side as in Figure 10. But this means that the order-reversing traversal property cannot hold, since $\eta'([0,\tau'])$ has now hit two different points on η without having hit the points in between. We are done.

□

We can use these ideas to prove a result on monotonicity of flow lines.

Remark. Starting with the next two corollaries, we will begin to suppress the precise hypotheses on the boundary data in the statements of results. The book-keeping would become distracting otherwise for no tangible gain, and we feel that the previous results already give a clear enough idea of how lemmas 4.1 and 4.2 are applied. The reader may assure him/herself that assuming very negative boundary conditions on the left side of $\partial\mathcal{S}$ and very positive conditions on the right side of $\partial\mathcal{S}$ would suffice for all subsequent results.

As a tacit reminder that we are doing this, for the remainder we will simply hypothesize a "good setup" to mean that we work in $\mathcal{S}$ with flow lines $0 \to z_0$ and counterflows $z_0 \to 0$ with boundary data satisfying some appropriate hypotheses.

Corollary 4.5. *Suppose that h is a GFF on $\mathcal{S}$ in a good setup. Suppose $\theta_1 < \theta_2$.*

Let η_{θ_i} be the flow line of $h + \theta_i \chi$ starting from 0, stopped when it accumulates in $\partial_U \mathcal{S}$. Then a.s., η_{θ_2} lies on or to the left of η_{θ_1}.

Proof. Consider a counterflow line η' with left boundary η_{θ_2}. Since η' contains η_θ for all $\theta \in [\theta_2 - \pi, \theta_2]$, this shows the result for $\theta_2 - \theta_1 \leq \pi$. If the difference is greater than π, simply repeat the argument inductively with flow lines $\eta_{\theta_2 - k\pi}$ for integers k.

□

Corollary 4.6. *On a good setup, the flow line η is determined by h.*

Proof. Take a counterflow line η' with left boundary flow line η. Recall that η', η were implicitly coupled to be conditionally independent throughout this section; this means that η has the same conditional law given h and given (h, η'). Since η is determined by η', it must be determined by (h, η') and hence by h alone.

□

5 Counterflow = Lightcone

Here we give another striking result, proving that counterflows are equal to *lightcones* of the flow lines. This implies that counterflow lines η' are determined by h.

We introduce angle-varying flow lines, which are the imaginary geometric analogs of polygonal line segments. The idea for the definition is that at the tip of a flow line $\eta(\tau)$, we may start a new flow line at a different angle by interpreting $\eta([0, \tau])$ as part of the boundary of a new domain, and in fact we may iterate this construction any finite number of times.

Definition 2. Fix angles $(\theta_1, ..., \theta_\ell)$. Let τ_1 be a stopping time for η_{θ_1}. After stopping at $\eta_{\theta_1}(\tau_1)$, let $\eta_{\theta_1,\theta_2}^{\tau_1,\tau_2}$ be the flow line at angle θ_2 started at $\eta_{\theta_1}(\tau_1)$, and stopped at a stopping time τ_2. Inductively, define

$$\eta_{\theta_1,...,\theta_j}^{\tau_1,...,\tau_j}$$

by starting at the tip of

$$\eta_{\theta_1,...,\theta_{j-1}}^{\tau_1,...,\tau_{j-1}}$$

and starting a flow line of angle θ_j at the tip. We call $\eta_{\theta_1,...,\theta_\ell}^{\tau_1,...,\tau_\ell}$ the **angle-varying flow line** with angles (θ_i) with respect to the stopping times (τ_i).

We now define the light-cone $\mathbf{L}$ of points reachable from 0 with angles in the range $[-\frac{\pi}{2}, \frac{\pi}{2}]$, and more generally the light-cones at the tip of any angle-varying flow line to be the set of points reachable from that tip. Because we don't have a simultaneous definition of all the flow lines, we cannot reason about an uncountable number of angles simultaneously. However, we can easily circumvent this technical hurdle by working with a countable dense set of angles and stopping times and taking $\mathbf{L}$ to be the closure of the resulting set.

Definition 3. The **light cone L** of h started at 0 is the closure of the union of all angle-varying flow lines with rational angles $\theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$ with all angle-changes time rational.

Similarly, for any angle-varying flow line $\eta^{\sigma_1,\ldots,\sigma_k}_{\phi_1,\ldots,\phi_k}$, the light cone $\mathbf{L}(\eta^{\sigma_1,\ldots,\sigma_k}_{\phi_1,\ldots,\phi_k})$ is the closure of the set of points reachable bu angle-varying flow lines starting at the tip of $\eta^{\sigma_1,\ldots,\sigma_k}_{\phi_1,\ldots,\phi_k}$ with angles in $[-\frac{\pi}{2}, \frac{\pi}{2}]$ and all stopping times rational.

Remark. It is not difficult to show that the light cones are independent of the countable dense set of angles used (above we took the rationals). This is because if $\theta_i \to \theta$ are deterministic angles, then we almost surely have convergence of the flow lines $\eta_{\theta_i} \to \eta_\theta$ (in, say, the Hausdorff topology). However (see the epilogue section on the SLE fan) we caution the reader that the flow lines do NOT in any deterministic sense vary continuously in the angle for a fixed GFF h.

The results on ordinary flow lines extend to angle-varying flow lines, which we record here.

Proposition 5.1. *If $\theta_i \in [-\frac{\pi}{2}, \frac{\pi}{2}]$, then any angle-varying flow line on a good setup with angles (θ_i) is simple and continuous, at least until hitting $\partial\mathcal{S}$, and is traversed in reverse order by η' and a.s. determined by h.*

Proof. The first claim is proved by induction on the number of straight parts in the angle-varying flow line using lemma 4.1 as usual to show that no self-intersections occur. The second and third claims are also easy inductions.

□

Theorem 5.2. *Let $\eta^{\sigma_1,\ldots,\sigma_k}_{\phi_1,\ldots,\phi_k}$ be an angle-varying flow line of h with angles $\phi_i \in [\phi_0, \phi_0+\pi]$ on a good setup. Then the light cone $\mathbf{L}(\eta^{\sigma_1,\ldots,\sigma_k}_{\phi_1,\ldots,\phi_k})$ of points reachable from the tip with angles in $[-\frac{\pi}{2}, \frac{\pi}{2}]$ almost surely coincides with the range of the counterflow η' of h started at z_0, stopping at the first hitting time for $\eta^{\sigma_1,\ldots,\sigma_k}_{\phi_1,\ldots,\phi_k}$. In fact, this is true even if we only allow the extreme angles $\pm\frac{\pi}{2}$ in the lightcone.*

Proof. We only give an outline of the proof. It suffices to show that angle-varying flow lines approach $\eta'(\tau')$ for any stopping time τ' such that $\eta'([0,\tau'])$ has not yet hit $\eta^{\sigma_1,\ldots,\sigma_k}_{\phi_1,\ldots,\phi_k}$. We can conformally map $\eta^{\sigma_1,\ldots,\sigma_k}_{\phi_1,\ldots,\phi_k}$ away to the boundary, essentially reducing this to the case of **L**. We only consider this case for notational simplicity.

The idea is to initially go straight at angle $\frac{\pi}{2}$ until nearly hitting $\eta'([0,\tau'])$ (say, until reaching within distance ε), and then repeatedly change direction by negating the angle after nearly hitting the opposite side. This generates an angle-varying flow line $\eta^{\tau_1,\ldots,\tau_\ell}_{\theta_1,\theta_2,\ldots,\theta_\ell}$ with $\theta_j = (-1)^{j+1}\frac{\pi}{2}$. This is depicted in Figure 11.

The same arguments as in the previous section show that the segments at angle $\frac{\pi}{2}$ can only approach the left boundary, while the segments at angle $-\frac{\pi}{2}$ can only approach the right boundary. This means that the angle-varying flow line will go back-and-forth between the left and right boundaries of $\eta'([0,\tau'])$, and we need to show that we can actually approach the tip $\eta'(\tau')$ in this way.

We conformally map away the hull $K_{\tau'}$ generated by the counterflow up to time τ', sending $\eta'(\tau') \to \infty$, $0 \to 0$. We are now reduced to the case of an angle-varying flow line oscillating back and forth between the left and right sides of the boundary, as in Figure 12. The point is to show that this angle-varying flow line (depicted at $\tilde{\eta}_3$) is not

contained in a bounded region after iterating this back-and-forth algorithm infinitely many times; this means that $\tilde{\eta}_3$ accumulates at ∞.

To show this, it suffices to show that $\tilde{\eta}_3$ has unbounded half-plane capacity, i.e. that the defining SDE runs for infinite time (recall that SLE_κ is parametrized by half-plane capacity; the same is true for $\tilde{\eta}_3$ by absolute continuity, since it never actually hits the boundary). To show this, the idea is to check that the $SLE_\kappa(\rho)$ SDE for each left-right iteration of the algorithm looks approximately identical. Indeed, one ends up obtaining an SDE with a single nearby force point together with many far-away force points which alternate in both sign and position. This means that the SDE is determined up to some minor changes such as a bounded drift term, and one can show from this that the time taken to complete a single left-right iteration of the algorithm is uniformly stochastically bounded from below by a nontrivial, a.s. positive random variable. This immediately implies that the time, hence half-plane capacity, of $\tilde{\eta}_3$ goes to infinity as desired. □

Figure 11: To prove that $\eta'(\tau')$ is in the light-cone $\mathbf{L}$, we go back-and-forth at angles $\pm\frac{\pi}{2}$, changing direction once we approach the boundary of $\eta'([0,\tau'])$.

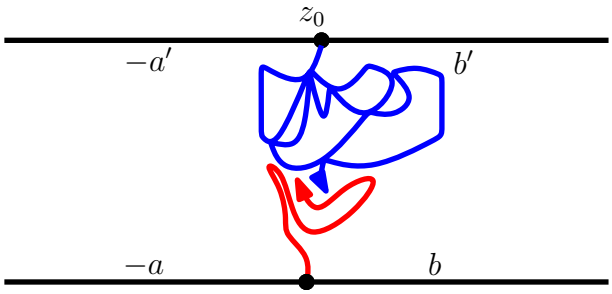

Figure 12: After conformally mapping $\eta'([0,\tau'])$ to the boundary so that $\eta'(\tau')$ is sent to infinity, it suffices to show that the angle-varying flow line $\tilde{\eta}_3$ has infinite half-plane capacity, since this implies that it is not contained in any bounded region. To do this, we show that each left-right iteration looks approximately the same on the SDE level, so that the amount of half-plane capacity parametrized time taken must go to infinity.

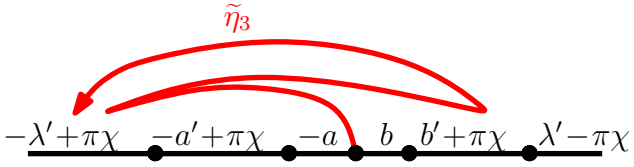

Corollary 5.3. *The counterflow line η' is determined almost surely by h on a good setup.*

Proof. It is now immediate that the trace of η' is determined as a compact set from the angle-varying flow lines, and hence from h, and we can show with only slightly more work that the curve η' itself is determined. Indeed, we can fill η' with a countable dense

set of points on our angle-varying flow lines, which are totally ordered by their η'-hitting times. By continuity this totally orders all the points of the lightcone by first-hitting time, and since η' is parametrized deterministically by halfplane-capacity (after applying an appropriate conformal isomorphism), this shows that η' is uniquely determined by h. □

Figure 13: Here we depict a counterflow line being filled by back-and-forth movement along angle-varying flow lines, as in Theorem 5.2. The first figure shows a partial filling of the counterflow line in which only a single angle-change is allowed, while in the second, limiting figure the counterflow line is filled.

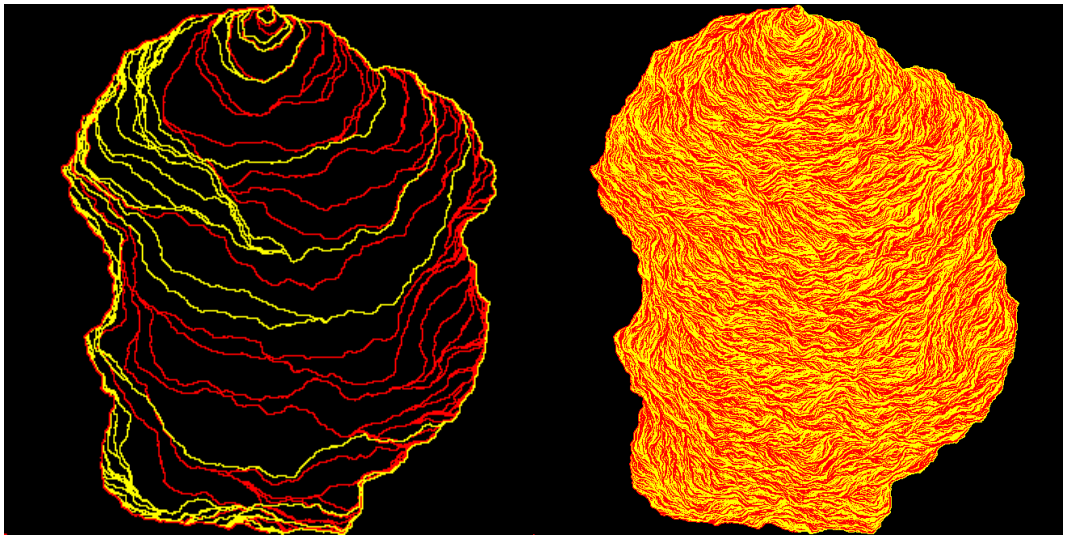

6 Epilogue

Here we give brief overviews of other aspects of imaginary geometry. We begin with some technical points we de-emphasized in the bulk of this paper and outline how they are used. We then discuss some other important topics. We will be more free here in making arguments which need more technical justification but should be intuitive and believable.

6.1 Local Sets of the Gaussian Free Field

A *local set* A of a GFF h on a domain D is a coupling (h, A) of a GFF with a closed set $\partial D \subseteq A \subseteq D$ such that conditioned on the harmonic extension of $h|_A$ to A^c, knowledge of the set A itself gives no more information about $h|_{A^c}$. Locality has a couple of other useful characterizations; for example, A is local iff for any (deterministic) open set $U \subseteq D$, the conditional probability given h that $A \subseteq U^c$ depends only on the projection of h onto $Harm(U)$.

Flow and counterflow lines are easily seen to be local from the conditional law of h in the coupling construction, and various properties of local sets turn out to be quite useful

in the study of imaginary geometry. For example, conditionally independent unions of local sets are local, which was implicitly used in the proof of Theorem 3.3. We have already used a couple of technical results which rely on locality of flow and counterflow lines, and we will illustrate a few more uses below.

6.2 Counterflow Determines Flow Lines for $\kappa \in (2, 4)$

One application of local sets comes in showing that for $\kappa \in (2, 4)$, the counterflow path η' determines all the flow lines η_θ which are contained inside it. The proof goes in the following steps, which are justified using local set results:

1. Since the range A of η' has Minkowski dimension $1 + \frac{\kappa'}{8} < 2$, $h|_{A^c}$ determines h almost surely.

2. Hence, the law of h given η' does not change if we condition further on $h|_{\eta'}$, meaning the intersection of the σ-algebras generated by $h|_U$ over all U containing A.

3. Since η_θ is local and contained in A, conditioning on η_θ does nothing if we have already conditioned on $h|_{\eta'}$.

4. Therefore, conditioning further on η' gives no more information about h if we have already conditioned on η. This implies that η' determines η.

In fact the same argument shows that η' determines all of the *angle-restricted light-cones* that it contains, i.e. the light-cones generated by a subinterval $[\theta_1, \theta_2] \subseteq [-\frac{\pi}{2}, \frac{\pi}{2}]$ of angles.

6.3 Conditional Laws for Multiple Flow Lines

We can show that boundary conditions given multiple flow lines are what you would expect, namely one just takes the boundary conditions for each flow line and combines them. Slightly technical GFF results show that this holds away from the intersections, using absolute continuity of GFF on separated sets. To show nothing weird happens at the intersections, it suffices to show there is no "jump" in the conditional expectation of (h, f) when intersections happen, for test functions f with support not intersecting the flow lines. This is fine because this conditional expectation is a time-changed Brownian motion, hence continuous.

Given some flow lines, we want to show that flow lines in between can be viewed as flow lines in the restricted domain (with flow line boundary values). This follows from the martingale characterization, but we first must establish that the intermediate flow line has a continuous Loewner driving function. This is actually quite simple; we essentially just need to know that no "tunnelling" happens, i.e. that the flow lines do not overlap for positive time-intervals. This is true because if tunnelling happened, the boundary values for h on the overlap would have two different values, and this turns out to be impossible.

With this fact in hand, one can use the martingale characterization to establish that, conditional on a flow line on the right side, a new flow line is just a flow line in the left-hand component, and similarly for flow lines between two other flow lines as explain in Figure 14.

Figure 14: Conditioned on flow lines $\eta_{\theta_1}, \eta_{\theta_2}$, the law of a flow line η between them is that of an independent flow line in each component in between, with boundary conditions coming from η_{θ_i}.

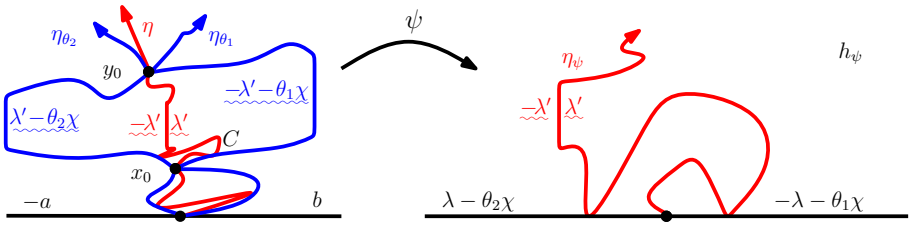

While interesting in itself, this idea of embedding flow lines between each other is also important for extending the interaction results we proved earlier to domains with arbitrary boundary conditions. For example, because with appropriate boundary values a, b as in Figure 14 we know that η is boundary-avoiding and hence continuous, we immediately conclude that any flow line is continuous when the only force points are at $0^-, 0^+$ since any such setup can be realized as a connected component between $\eta_{\theta_1}, \eta_{\theta_2}$ and ψ is a homeomorphism.

6.4 The SLE Fan

The SLE fan is the closure of the union of flow lines in an imaginary geometry started from a given boundary point, with angles ranging over a countable dense set. As with the light cones, a countable dense set is needed since we can only work simultaneously with countably many flow lines, but it is again independent of the choice of countable dense set.

Figure 15: Unlike in a classical geometry, the fan of points accessible from a given boundary point via straight lines always has measure 0 in imaginary geometry.

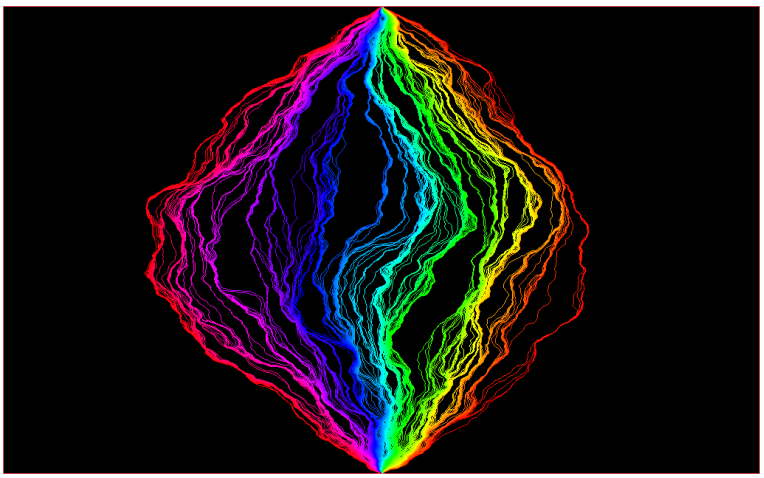

For $\kappa \in (2, 4)$ we know that the SLE fan is contained in the lightcone, hence has

dimension at most $1 + \frac{\kappa'}{8} = 1 + \frac{2}{\kappa} < 2$, meaning that almost every point is not in the fan. In fact, the dimension of the SLE fan is $1 + \frac{\kappa}{8}$, the same as that of a single flow line. A consequence of the sparcity of the fan is that flow lines must vary discontinuously in the angle; any sensible sort of continuous variation would have to sweep out every point in between the extreme flow lines.

7 Acknowledgements

I thank Andrew Ahn for his mentorship and direction. I thank Professor Scott Sheffield for answering my questions and explaining lots of exciting math, and granting permission to use the pictures from [MS12a], which was the source of every picture used here. Finally I thank the MIT UROP+ program for the opportunity to study this incredible material.

References

[Ber16] Nathaniel Berestycki. Introduction to the gaussian free field and liouville quantum gravity. Jul 2016.

[BN16] Nathaniel Berestycki and James Norris. Lectures on schrammloewner evolution. Jan 2016.

[Dub07] Julien Dubédat. Duality of schramm-loewner evolutions. *arXiv preprint arXiv:0711.1884*, 2007.

[MS12a] Jason Miller and Scott Sheffield. Imaginary geometry i: interacting sles. *arXiv preprint arXiv:1201.1496*, 2012.

[MS12b] Jason Miller and Scott Sheffield. Imaginary geometry ii: reversibility of sle_\ kappa (\ rho_1;\ rho_2) for\ kappa\ in (0, 4). *arXiv preprint arXiv:1201.1497*, 2012.

[MS12c] Jason Miller and Scott Sheffield. Imaginary geometry iii: reversibility of sle_\ kappa\ for\ kappa\ in (4, 8). *arXiv preprint arXiv:1201.1498*, 2012.

[MS13] Jason Miller and Scott Sheffield. Imaginary geometry iv: interior rays, whole-plane reversibility, and space-filling trees. *arXiv preprint arXiv:1302.4738*, 2013.

[RS11] Steffen Rohde and Oded Schramm. Basic properties of sle. In *Selected Works of Oded Schramm*, pages 989–1030. Springer, 2011.

[She07] Scott Sheffield. Gaussian free fields for mathematicians. *Probability theory and related fields*, 139(3-4):521–541, 2007.

[SS09] Oded Schramm and Scott Sheffield. Contour lines of the two-dimensional discrete Gaussian free field. *Acta mathematica*, 202(1):21–137, 2009.

[SS13] Oded Schramm and Scott Sheffield. A contour line of the continuum gaussian free field. *Probability Theory and Related Fields*, 157(1-2):47–80, 2013.

[Wer03] Wendelin Werner. Random planar curves and schramm-loewner evolutions. *arXiv preprint math/0303354*, 2003.