Bipolar orientations on planar maps and SLE_{12}

Richard Kenyon   Jason Miller   Scott Sheffield   David B. Wilson

Brown   Cambridge   MIT   Microsoft

Abstract

We give bijections between bipolar-oriented (acyclic with unique source and sink) planar maps and certain random walks, which show that the uniformly random bipolar-oriented planar map, decorated by the "peano curve" surrounding the tree of left-most paths to the sink, converges in law with respect to the peanosphere topology to a $\sqrt{4/3}$-Liouville quantum gravity surface decorated by an independent Schramm-Loewner evolution with parameter $\kappa = 12$ (i.e., SLE_{12}). This result is universal in the sense that it holds for bipolar-oriented triangulations, quadrangulations, k-angulations, and maps in which face sizes are mixed.

1 Introduction

1.1 Planar maps

A *planar map* is a planar graph together with an embedding into $\mathbb{R}^2$ so that no two edges cross. More precisely, a planar map is an equivalence class of such embedded graphs, where two embedded graphs are said to be equivalent if there exists an orientation preserving homeomorphism $\mathbb{R}^2 \to \mathbb{R}^2$ which takes the first to the second. The enumeration of planar maps started in the 1960's in work of Tutte [Tut63], Mullin [Mul67], and others. In recent years, new combinatorial techniques for the analysis of random planar maps, notably via random matrices and tree bijections, have revitalized the field. Some of these techniques were motivated from physics, in particular from conformal field theory and string theory.

There has been significant mathematical progress on the enumeration and scaling limits of random planar maps chosen uniformly from the set of all rooted planar maps with a given number of edges, beginning with the bijections of Cori–Vauquelin [CV81] and Schaeffer [Sch98] and progressing to the existence of Gromov–Hausdorff metric space limits established by Le Gall [LG13] and Miermont [Mie13].

There has also emerged a large literature on planar maps that come equipped with additional structure, such as the instance of a model from statistical physics, e.g., a uniform spanning tree, or an Ising model configuration. These "decorated planar maps" are important in Euclidean 2D statistical physics. The reason is that it is often easier to compute "critical exponents" on planar maps than on deterministic

1

lattices. Given the planar map exponents, one can apply the KPZ formula to *predict* the analogous Euclidean exponents.[1] In this paper, we consider random planar maps equipped with bipolar orientations.

1.2   Bipolar and harmonic orientations

A *bipolar (acyclic) orientation* of a graph G with specified source and sink (the "poles") is an acyclic orientation of its edges with no source or sink except at the specified poles. (A *source* (resp. *sink*) is a vertex with no incoming (resp. outgoing) edges.) For any graph G with adjacent source and sink, bipolar orientations are counted by the coefficient of x in the Tutte polynomial $T_G(x, y)$, which also equals the coefficient of y in $T_G(x, y)$; see [dFdMR95] or the overview in [FPS09]. In particular, the number of bipolar orientations does not depend on the choice of source and sink as long as they are adjacent. If the source and sink are not adjacent, adjoining an edge between the source and sink does not affect the number of bipolar orientations, so bipolar orientations are counted by these Tutte coefficients in the augmented graph.

Let G be a finite connected planar map, with no self-loops but with multiple edges allowed, with a specified source and sink that are incident to the same face. It is convenient to embed G in the disk so that the source is at the bottom of the disk (and is denoted S, for south pole), the sink is at the top (and is denoted N, for north pole), all other vertices are in the interior of the disk (see Figure 1). Within the disk there are two faces that are boundary faces, which can be called W (the west pole) and E (the east pole). Endowing G with a bipolar orientation is a way to endow it and its dual map G^* with a coherent notion of "north, south, east, and west": one may define the directed edges to point *north*, while their opposites point *south*; a directed dual edge points *east* (resp. *west*) if the edge it crosses is oriented left (resp. right).

Given an orientation of a finite connected planar map G, its dual orientation of G^* is obtained by rotating directed edges counterclockwise. If an orientation has a sink or source at an interior vertex, its dual has a cycle around that vertex. Suppose an orientation has a cycle but has no source or sink at interior vertices. If this cycle surrounds more than one face, then one can find another cycle that surrounds fewer faces, so there is a cycle surrounding just one face, and the dual orientation has either a source or sink at that (interior) face. Thus an orientation of G is bipolar acyclic precisely when its dual orientation of G^* is bipolar acyclic. The east and west poles of G^* are the source and sink respectively of the dual orientation (see Figure 1).

One way to construct bipolar orientations is via electrical networks. Suppose every edge of G represents a conductor with some generic positive conductance, the

---

[1]This idea was used by Duplantier to derive the so-called Brownian intersection exponents [Dup98], whose values were subsequently verified mathematically by Lawler, Schramm, and Werner [LSW01b, LSW01c, LSW02] in an early triumph of Schramm's SLE theory [Sch00]. An overview with a long list of references can be found in [DS11].
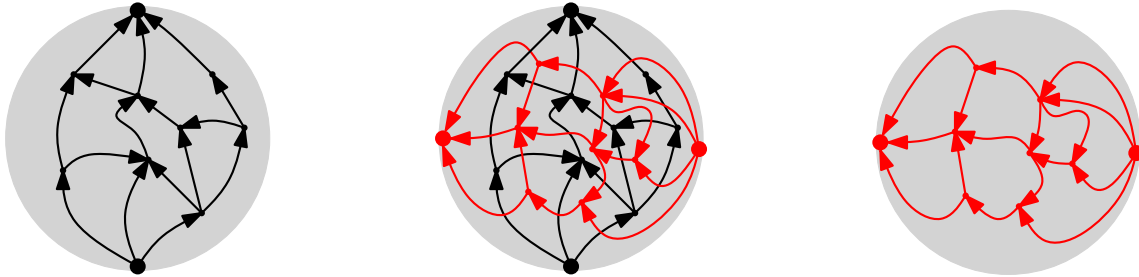
Figure 1: **Left:** A planar map embedded in a disk with two boundary vertices, with a north-going bipolar orientation. **Right:** The dual bipolar-oriented planar map, which has two boundary dual vertices on the disk. **Middle:** Primal and dual bipolar-oriented maps together. The dual orientations are obtained from the primal orientations by rotating the arrows left.

south pole is at 0 volts, and the north pole is at 1 volt. The voltages are harmonic except at the boundary vertices, and for generic conductances, provided every vertex is part of a simple path connecting the two poles, the interior voltages are all distinct. The *harmonic orientation* orients each edge towards its higher-voltage endpoint. The harmonic orientation is clearly acyclic, and by harmonicity, there are no sources or sinks at interior vertices. In fact, for any planar graph with source and sink incident to the same face, *any* bipolar orientation is the harmonic orientation for some suitable choice of conductances on the edges [AK15, Thm. 1], so for this class of graphs, bipolar orientations are equivalent to harmonic orientations.

Suppose that a bipolar-oriented planar map G has an interior vertex incident to at least four edges, which in cyclic order are oriented outwards, inwards, outwards, inwards. By the source-free sink-free acyclic property, these edges could be extended to oriented paths which reach the boundary, and by planarity and the acyclic property, the paths would terminate at four distinct boundary vertices. Since (in this paper) we are assuming that there are only two boundary vertices, no such interior vertices exist. Thus at any interior vertex, its incident edges in cyclic order consist of a single group of north-going edges followed by a single group of south-going edges, and dually, at each interior face the edges in cyclic order consist of a single group of clockwise edges followed by a single group of counterclockwise edges.

In particular, each vertex (other than the north pole) has a unique "west-most north-going edge," and the collection of such directed edges forms the *NW tree*. We define southwest, southeast, and northeast trees similarly.

We will exhibit (see Theorems 1 and 2) a bijection between bipolar-oriented planar maps (with given face-degree distribution) and certain types of random walks in the nonnegative quadrant $\mathbb{Z}_{\geq 0}^2$. This bijection leads to exact enumerative formulae as well as local information about the maps such as degree distributions. For previous enumerative work on this model, including bijections between bipolar-oriented planar maps and other objects, see e.g. [FPS09, BM11, BBMF11, FFNO11].

3

1.3 SLE and LQG

After the bijections our second main result is the identification of the scaling limit of the bipolar-oriented map with a *Liouville quantum gravity* (LQG) surface decorated by a *Schramm-Loewner evolution* (SLE) curve, see Theorem 9.

We will make use of the fact proved in [DMS14, MS15e, GHMS15] that an SLE-decorated LQG surface can be equivalently defined as a mating of a correlated pair of continuum random trees (a so-called *peanosphere*; see Section 4.2) where the correlation magnitude is determined by parameters that appear in the definition of LQG and SLE (namely γ and κ').

The scaling limit result can thus be formulated as the statement that a certain pair of discrete random trees determined by the bipolar orientation (namely the *northwest* and *southeast* trees, see Section 1.2) has, as a scaling limit, a certain correlated pair of continuum random trees. Although LQG and SLE play a major role in our motivation and intuition (see Sections 4.2 and 5), we stress that no prior knowledge about these objects is necessary to understand either the main scaling limit result in the current paper or the combinatorial bijections behind its proof (Sections 2 and 3).

Before we move on to the combinatorics, let us highlight another point about the SLE connection. There are several special values of the parameters κ and $\kappa' = 16/\kappa$ that are related to discrete statistical physics models. (SLE_κ with $0 < \kappa \leq 4$ and $\mathrm{SLE}_{16/\kappa}$ are closely related [Zha08, Dub09, MS12a, MS13a], which is known as SLE-duality.) These special $\{\kappa, \kappa'\}$ pairs include $\{2, 8\}$ (for loop-erased random walk and the uniform spanning tree) [LSW04], $\{8/3, 6\}$ (for percolation and Brownian motion) [Smi01, LSW01a], $\{3, 16/3\}$ (for the Ising and FK-Ising model) [Smi10, CDCH⁺14], and $\{4, 4\}$ (for the Gaussian free field contours) [SS09, SS13]. The relationships between these special $\{\kappa, \kappa'\}$ values and the corresponding discrete models were all discovered or conjectured within a couple of years of Schramm's introduction of SLE, building on earlier arguments from the physics literature. We note that all of these relationships have random planar map analogs, and that they all correspond to $\{\kappa, \kappa'\} \subset [2, 8]$. This range is significant because the so-called *conformal loop ensembles* CLE_κ [She09, SW12] are only defined for $\kappa \in (8/3, 8)$, and the discrete models mentioned above are all related to random collections of loops in some way, and hence have either κ or κ' in the range $(8/3, 8)$.

In this paper the relevant $\{\kappa, \kappa'\}$ pair is $\{4/3, 12\}$. This special pair is interesting in part because it lies outside the range $[2, 8]$. It had been proposed, based on heuristic arguments and simulations, that "activity-weighted" spanning trees should have SLE scaling limits with κ anywhere in the range $[4/3, 4)$ and κ' anywhere in the range $(4, 12]$ [KW15]. In more recent work, subsequent to our work on bipolar orientations, using a generalization of the inventory accumulation model in [She11], the activity-weighted spanning trees on planar maps were shown to converge to SLE-decorated LQG in the peanosphere topology for this range of κ, κ' [GKMW15].

We will further observe that if one modifies the bipolar orientation model by

4

a weighting that makes the faces more or less balanced (in terms of their number of clockwise and counterclockwise oriented boundary edges), one can obtain any $\kappa \in (0, 2)$ and any $\kappa' \in (8, \infty)$. In a companion to the current paper [KMSW15], we discuss a different generalization of bipolar orientations that we conjecture gives SLE for $\kappa \in [12 - 8\sqrt{2}, 4)$ and $\kappa' \in (4, 12 + 8\sqrt{2}]$.

1.4 Outline

In Sections 2 and 3 we establish our combinatorial results and describe the scaling limits of the NW and SE trees in terms of a certain Brownian excursion. In Section 4 we explain how this implies that the uniformly random bipolar-oriented map with n edges, and fixed face-degree distribution, decorated by its NW tree, converges in law as $n \to \infty$ to a $\sqrt{4/3}$-Liouville quantum gravity sphere decorated by space-filling SLE_{12} from ∞ to ∞. This means that, following the curve which winds around the NW tree, the distances to the N and S poles scale to an appropriately correlated pair of Brownian excursions. We also prove a corresponding *universality* result: the above scaling limit holds for essentially any distribution on face degrees (or, dually, vertex degrees) of the random map.

In Section 5 we explain, using the so-called imaginary geometry theory, what is special about the value $\kappa' = 12$. These observations allow us to explain at a heuristic level why (even before doing any discrete calculations) one would *expect* $\kappa' = 12$ to arise as the scaling limit of bipolar orientations.

2 Bipolar-oriented maps and lattice paths

2.1 From bipolar maps to lattice paths

For the bipolar-oriented planar map in Figure 1, Figure 2 illustrates its *NW tree* (in red), *SE tree* (in blue), and the *interface path* (in green) which winds between them from the south pole to the north pole. The interface path has two types of steps:

1. Steps that traverse an edge (between red and blue sides).

2. Steps that traverse an interior face from its maximum to its minimum. Face steps can be subcategorized according to the number of edges on the west and east sides of the face, where the maximum and minimum vertex of a face
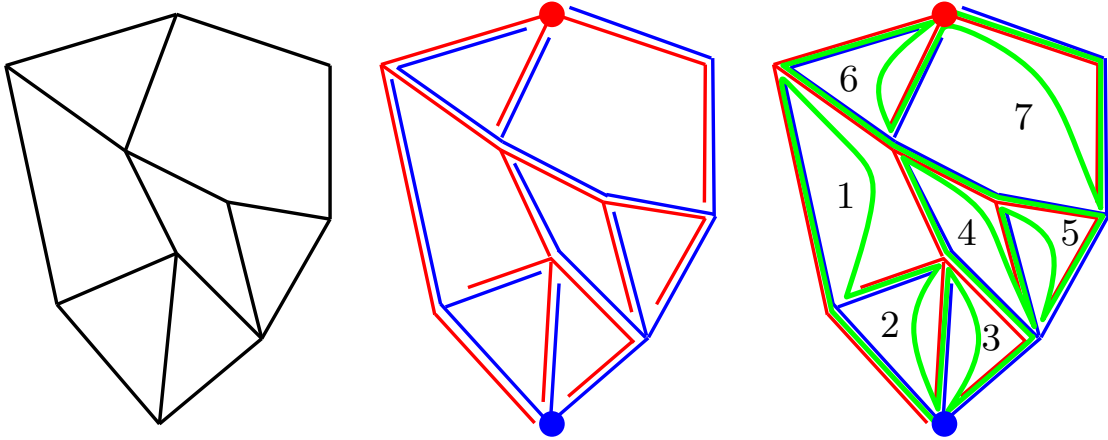
Figure 2: **Left:** A map with a bipolar orientation, embedded so each edge is oriented "upward" (i.e., in the direction along which the vertical coordinate increases). **Middle:** Set of oriented edges can be understood as a tree, the *northwest tree*, where the parent of each edge is the leftmost upward oriented edge it can merge into. If we reverse the orientations of all edges, we can define an analogous tree (blue) and embed both trees (using the British convention of driving on the left side) so that they don't cross each other. **Right:** We then add a green path tracing the interface between the two trees. Each edge of the interface moves along an edge of the map or across a face of the map. For illustration purposes, faces are numbered by the order they are traversed by the green path, but it is the traversals of the edges of the green path that correspond to steps of the lattice path.

separate its west from its east. If a face has $i + 1$ edges on its west and $j + 1$ edges on its east, we say that it is of type (i, j).

Observe that each face step has edge steps immediately before and after it.

Let E be the set of edges of the planar map, which we order $e_0, \ldots, e_{|E|-1}$ according to the green path going from the south pole S to the north pole N. For each edge e_t, let X_t be distance in the blue tree from the blue root (S) to the lower endpoint of e_t, and let Y_t be the distance in the red tree from the red root (N) to the upper endpoint of e_t. Suppose the west outer face has $m + 1$ edges and the east outer face has $n + 1$ edges. Then the sequence $\{(X_t, Y_t)\}_{0 \le t \le |E|-1}$ defines a walk or lattice path that starts at $(0, m)$ when $t = 0$ and ends at $(n, 0)$ when $t = |E| - 1$, and which remains in the nonnegative quadrant. If there is no face step between e_t and e_{t+1}, then the walk's increment $(X_{t+1}, Y_{t+1}) - (X_t, Y_t)$ is $(1, -1)$. Otherwise there is exactly one face step between e_t and e_{t+1}; if that face has $i + 1$ edges on its west and $j + 1$ edges on its east, then the walk's increment is $(-i, j)$, see Figure 3.

For the example in Figure 2, the walk starts at $(0, 2)$ and ends at $(3, 0)$.

2.2 From lattice paths to bipolar maps

The above construction can be reversed, constructing a bipolar-oriented planar map from a lattice path of the above type.
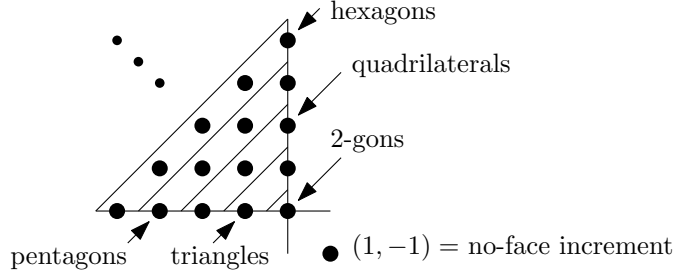
Figure 3: Lattice path increments.

We construct the bipolar-oriented planar map by sewing edges and oriented polygons according to the sequence of steps of the path. Let $m_{i,j}$ denote a step of $(-i, j)$ with $i, j \geq 0$, and m_e denote a step of $(1, -1)$. It is convenient to extend the bijection, so that it can be applied to any sequence of these steps. These steps give sewing instructions to augment the current "marked bipolar map".

A marked bipolar map is a bipolar-oriented planar map together with a "start vertex" on its western boundary which is not at the top, and an "active vertex" on its eastern boundary which is not at the bottom, such that the start vertex and every vertex below it on the western boundary has at most one downward edge, and the active vertex and every vertex above it on the eastern boundary has at most one upward edge. We think of the edges on the western boundary below the start vertex and on the eastern boundary above the active vertex as being "missing" from the marked bipolar map: they are boundaries of open faces that are part of the map, but are not themselves in the map.
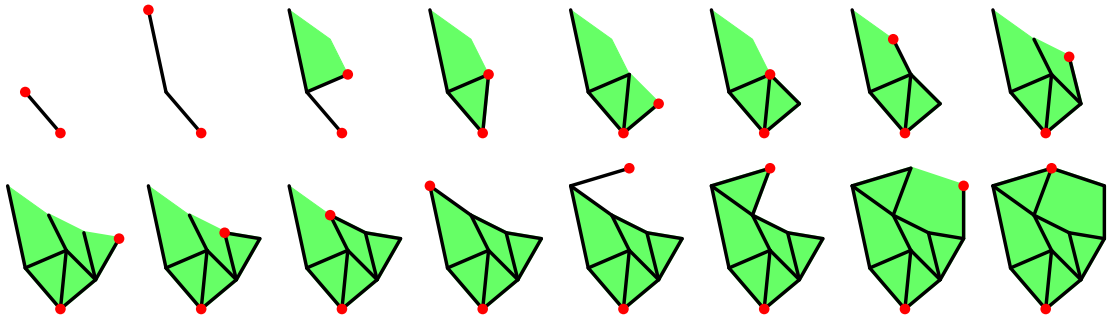

Figure 4: The process of sewing oriented polygons and edges to obtain a bipolar-oriented planar map. The intermediate structures are marked bipolar-oriented planar maps, which may have some edges missing on the boundaries.

Initially the marked bipolar map consists of an oriented edge whose lower endpoint is the start vertex and whose upper endpoint is the active vertex. The m_e moves will sew an edge to the current marked bipolar map upwards from the active vertex and

move the active vertex to the upper endpoint of the new edge. If the eastern boundary had a vertex above the active vertex, the new edge gets sewn to the southernmost missing edge on the eastern boundary, and otherwise there is a new vertex which becomes the current top vertex. The $m_{i,j}$ moves will sew a face with $i+1$ edges on its west and $j+1$ edges on its east, sewing the north of the face to the active vertex and the west of the face to the eastern boundary of the marked bipolar map, and then sew an edge to the southernmost east edge of the new face; the new active vertex is the upper vertex of this edge. If there are fewer than $i+1$ edges below the (old) active vertex, then the new face gets sewn to as many of them as there are, the start vertex is no longer at the bottom, and the remaining western edges of the face are missing from the map.

The final marked bipolar map is considered a valid bipolar-oriented planar map if the start vertex is at the south and the active vertex is at the north, or equivalently, if there are no missing edges.

Theorem 1. *The above mapping from sequences of moves from $\{m_e\} \cup \{m_{i,j} : i, j \geq 0\}$ to marked bipolar maps is a bijection.*

Proof. Consider a marked bipolar map obtained from a sequence of moves. The number of edges present in the structure determines the length of the sequence. If that length is positive, then the easternmost downward edge from the active vertex was the last edge adjoined to the structure. If this edge is the southernmost edge on the eastern boundary of a face, then the last move was one of the $m_{i,j}$'s, and otherwise it was m_e. Since the last move and preceding structure can be recovered, the mapping is an injection.

Starting from an arbitrary marked bipolar map, we can inductively define a sequence of moves as above by considering the easternmost downward edge from the active vertex. This sequence of moves yields the original marked bipolar map, so the mapping is a surjection. $\qquad\square$

Next we restrict this bijection to sequences of moves which give valid bipolar-oriented planar maps. A sequence of moves can of course be encoded as a path.

Theorem 2. *The above mapping gives a bijection from length-$(\ell - 1)$ paths from $(0, m)$ to $(n, 0)$ in the nonnegative quadrant having increments $(1, -1)$ and $(-i, j)$ with $i, j \geq 0$, to bipolar-oriented planar maps with ℓ total edges and $m + 1$ and $n + 1$ edges on the west and east boundaries respectively. A step of $(-i, j)$ in the walk corresponds to a face with degree $i + j + 2$ in the planar map.*

Proof. When we make a walk started from $(0, m)$ using these moves, the number of (non-missing) edges on the eastern boundary of the marked bipolar map is one plus the first coordinate of the walk. Thus the start vertex is at the south pole precisely when the first coordinate always remains nonnegative, and the final number of (nonmissing) edges on the eastern boundary is $n + 1$.

8

Observe that if we reverse a sequence of moves, then the structure obtained from the reversed sequence is the same as the structure obtained from the initial sequence rotated by 180°, with the roles of start and active vertices reversed.

Using this reversal symmetry with our previous observation, it follows that the active vertex is at the north pole precisely when the second coordinate achieves its minimum on the last step, and the number of (nonmissing) edges on the western boundary is $m + 1$. □

If we wish to restrict the face degrees, the bijection continues to hold simply by restricting the set of allowed steps of the paths.

Corollary 3. *Any bipolar-oriented planar map has a straight-line planar embedding (except that the edges of 2-gons overlap) such that edges are oriented upwards, i.e., in the direction of increasing y-coordinate, as in Figure 2.*

Proof. Given a bipolar map, we can convert it to walk, and then convert it back to a bipolar map. When converting it back to a bipolar map, each step of the process of sewing oriented polygons can be done with straight-line edges, as in Figure 4. □

2.3 Path scaling limit

What happens if we consider a random bipolar-oriented planar map such as the one in Figure 2, where we fix the left boundary length (3 in Figure 2), the right boundary length (4 in Figure 2), and the total number ℓ of edges (16 in Figure 2)? We consider the limiting case where the boundary lengths are fixed and $\ell \to \infty$. What can one say about the limiting joint law of the pair of trees in Figure 2 in this situation?

In light of Theorem 2, understanding this limiting law amounts to understanding the limiting law of its associated lattice path. For example, if the map is required to be a triangulation, then the lattice path is required to have increments of size $(1, -1)$, $(-1, 0)$, and $(0, 1)$. Since $\ell \to \infty$ with fixed endpoints, there are $\ell/3 + O(1)$ steps of each type. One can thus consider a random walk of length $\ell - 1$ with these increment sizes (each chosen with probability 1/3) conditioned to start and end at certain fixed values, and to stay in the nonnegative quadrant.

It is reasonable to expect that if a random walk on $\mathbb{Z}^2$ converges to Brownian motion with some non-degenerate diffusion matrix, then the same random walk conditioned to stay in a quadrant (starting and ending at fixed locations when the number of steps gets large) should scale to a form of the Brownian bridge constrained to stay in the same quadrant (starting and ending at 0). The recent work [DW15, Theorem 4] contains a precise theorem of this form, and Proposition 4 below is a special case of this theorem.

Recall that the period of a random walk on $\mathbb{Z}^2$ is the smallest integer $p \geq 1$ such that the random walk has a positive probability to return to zero after kp steps for all sufficiently large integers $k > 0$.

Proposition 4. *Let ν be a probability measure supported on $\mathbb{Z}^2$ with expectation zero and moments of all orders.[2] Let $p \geq 1$ denote the period of the random walk on $\mathbb{Z}^2$ with step distribution ν. Suppose that for given $z_{start}, z_{end} \in \mathbb{Z}_{\geq 0}^2$ there is a positive probability path from z_{start} to z_{end} with steps from ν of length ℓ. Suppose further that for any $R > 0$ there is a point $z \in \mathbb{Z}_{\geq 0}^2$ that is distance at least R from the boundary of the quadrant, such that there is a path from z_{start} to z to z_{end} with steps from ν that remains in the quadrant $\mathbb{Z}_{\geq 0}^2$. For sufficiently large n with $n \equiv \ell \bmod p$, consider a random walk $z_{start} = S_0, S_1, \ldots, S_n = z_{end}$ from z_{start} to z_{end} with increments chosen from ν, conditioned to remain in the quadrant $\mathbb{Z}_{\geq 0}^2$. Then the law of $S_{\lfloor nt \rfloor}/\sqrt{n}$ converges weakly w.r.t. the L^∞ norm on $[0,1]$ to that of a Brownian excursion (with diffusion matrix given by the second moments of ν) into the nonnegative quadrant, starting and ending at the origin, with unit time length.*

Now let us return to the study of random bipolar-oriented planar triangulations. By Theorem 2 these correspond to paths in the nonnegative quadrant from the y-axis to the x-axis which have increments of $(1, -1)$ and $(0, 1)$ and $(-1, 0)$. Fix the boundary lengths $m + 1$ and $n + 1$, that is, fix the start $(0, m)$ and end $(n, 0)$ of the walk, and let the length ℓ get large. Note that if ν is the uniform measure on the three values $(1, -1)$ and $(0, 1)$ and $(-1, 0)$, then the ν-expectation of an increment (X, Y) of the (unconstrained) walk is $(0, 0)$. Furthermore, (in the unconstrained walk) the variance of $X - Y$ is 2 while the variance of $X + Y$ is $2/3$, and the covariance of $X - Y$ and $X + Y$ is zero by symmetry. Thus the variance in the $(1, -1)$ direction is 3 times the variance in the $(1, 1)$ direction. The scaling limit of the random walk will thus be a Brownian motion with the corresponding covariance structure. We can summarize this information as follows:

Theorem 5. *Consider a uniformly random bipolar-oriented triangulation, sketched in the manner of Figure 2, with fixed boundary lengths $m + 1$ and $n + 1$ and with the total number of edges given by ℓ. Let $S_0, S_1, \ldots$ be the corresponding lattice walk. Then $S_{\lfloor \ell t \rfloor}/\sqrt{\ell}$ converges in law (weakly w.r.t. the L^∞ norm on $[0, 1]$), as $\ell \to \infty$ with $\ell \equiv n - m + 1 \bmod 3$, to the nonnegative quadrant constrained Brownian bridge starting and ending at the origin, with covariance matrix $\begin{pmatrix} 2/3 & -1/3 \\ -1/3 & 2/3 \end{pmatrix}$. (This is the covariance matrix such that if the Brownian motion were unconstrained, the difference and sum of the two coordinates at time 1 would be independent with respective variances 2 and $2/3$.)*

In particular, Theorem 5 holds when the lattice path starts and ends at the origin, so that the left and right sides of the planar map each have length 1. In this case,

---

[2]In fact it suffices that $|\cdot|^\alpha$ has ν-finite expectation, for a positive constant α defined in [DW15]. The constant α depends on the angle of the cone $L(\mathbb{R}_{\geq 0}^2)$, where $L : \mathbb{R}^2 \to \mathbb{R}^2$ is a linear map for which $L(S_n)$ scales to a constant multiple of standard two-dimensional Brownian motion. In the setting of Theorems 5 and 6 below, L can be the map that rescales the $(1, -1)$ direction by $1/\sqrt{3}$ and fixes the $(1, 1)$ direction. In this case, the cone angle is $\pi/3$ and $\alpha = 3$.

the two sides can be glued together and treated as a single edge in the sphere, and Theorem 5 can be understood as a statement about bipolar maps on the sphere with a distinguished south to north pole edge.

Next we consider more general bipolar-oriented planar maps. Suppose we allow not just triangles, but all face sizes between 2 and some value r. Then instead of taking ν to be uniform on the triple $\{(1,-1),(0,1),(-1,0)\}$ we may take ν to be uniform on the set

$$\{(1,-1)\} \cup \{(-i,j) : i \geq 0, j \geq 0, i+j \leq r-2\}.$$

The ν defined in this is way does not have expectation $(0,0)$. However, note that if we *weight* the law of ν on (x,y) pairs by $a^x b^y$ (where a and b are any fixed positive constants) this does not change the probability of any lattice path *conditioned* on its endpoints. Since ν is already symmetric w.r.t. reflection across the line $y = -x$, one can find an a such that the measure ν weighted by $a^{(x-y)}$ (normalized to be a probability measure) has ν-expectation zero.

Indeed, in the limiting case $r = \infty$, taking $a = 1/2$ yields a probability measure with expectation zero; precisely, define

$$\nu_\infty\{(-i,j)\} = \begin{cases} 2^{-i-j-3} & i,j \geq 0 \ \text{ or } \ i = j = -1 \\ 0 & \text{otherwise.} \end{cases}$$

Using the fact that $\sum_{i=0}^{\infty}(1/2)^i \sum_{j=0}^{\infty}(1/2)^j j = \sum_{i=0}^{\infty}(1/2)^i \cdot 2 = 4$ (and the analog with i and j reversed) it is easy to check that ν_∞ has expectation $(0,0)$ and satisfies the requirements of Proposition 4.[3]

To generalize still further, we may consider situations where one weights the probability of a configuration by $\prod_{k=1}^{\infty} a_k^{n_k}$ where a_k are some constants and n_k is the number of faces with k edges. (Taking $a_k = 0$ means that faces with k edges are forbidden.) The ν one obtains in this general setting will assign probabilities $p_0, p_{i,j}$ to steps m_e and $m_{i,j}$ respectively, where $p_{i,j}$ only depends on $i+j$, and thus each of the $k-1$ possible steps $m_{i,k-2-i}$ (corresponding to adding k-gons) is assigned equal probability p_k. (Recall that we allow faces of degree 2, with corresponding step $m_{1,1}$.) Then

$$p_0 + \sum_{k \geq 2}(k-1)p_k = 1 \tag{1}$$

and the probability that in a large random map a randomly chosen face has degree k is then

$$\mathbb{P}(\text{face has degree } k) = \frac{(k-1)p_k}{1-p_0}.$$

---

[3]It is also possible to derive the distribution ν_∞ for uniformly random bipolar-oriented planar maps using a different bijection, one to noncrossing triples of lattice paths [FPS09].

The zero-drift condition on the increment (X, Y) is

$$p_0 - \sum_{k \geq 2} \frac{(k-2)(k-1)}{2} p_k = 0. \tag{2}$$

The variances of $X - Y$ and $X + Y$ are respectively

$$\mathrm{Var}[X - Y] = 4p_0 + \sum_{k \geq 2}(k-2)^2(k-1)p_k,$$

$$\mathrm{Var}[X + Y] = \sum_{k \geq 2} p_k((k-2)^2 + (k-4)^2 + \cdots + (-k+2)^2) = \sum_{k \geq 2} p_k \times 2\binom{k}{3},$$

$$\mathrm{Var}[X - Y] = \sum_{k \geq 2}(k-2)(k-1)k\, p_k = \sum_{k \geq 2} p_k\, 6\binom{k}{3} = 3\,\mathrm{Var}[X + Y].$$

Thus, the argument used to prove Theorem 5 also implies the following.

Theorem 6. *The conclusion of Theorem 5 still holds if triangulations are replaced by quadrangulations or k-angulations for any fixed k (up to a possible multiplication of the covariance matrix by a constant). More generally, the conclusions still hold in the "mixed face size" models described just above as long as the corresponding ν (as described above) has moments of all orders. In particular, they apply whenever one allows only a finite set of face sizes. They also apply (considering the case $\nu = \nu_\infty$) when face size is completely unrestricted.*

Remark 1. *If one relaxes the requirement that the probabilities assigned by ν be the same for all increments corresponding to a given face size, one can find a ν with the same support as ν_∞ such that the expectation is still $(0,0)$ and when (X, Y) is sampled from ν, the law is still symmetric w.r.t. reflection about the line $y = -x$ but the variance ratio $\mathrm{Var}[X - Y]/\mathrm{Var}[X + Y]$ assumes any value strictly between 1 and ∞. Indeed, one approaches one extreme by letting (X, Y) be (close to being) supported on the $y = -x$ antidiagonal, and the other extreme by letting (X, Y) be (close to being) supported on the x- and y-axes far from the origin (together with the point $(1, -1)$). The former corresponds to a preference for nearly balanced faces (in terms of the number of clockwise and counterclockwise oriented edges) while the latter corresponds to a preference for unbalanced faces.*

Remark 2. *In each of the models treated above, it is natural to consider an "infinite-volume limit" in which lattice path increments indexed by $\mathbb{Z}$ are chosen i.i.d. from ν. The standard central limit theorem then implies that the walks have scaling limits given by a Brownian motion with the appropriate covariance matrix.*

3   Bipolar-oriented triangulations

3.1   Enumeration

Corollary 7. *The number of bipolar-oriented triangulations of the sphere with ℓ edges in which S and N are adjacent and marked is (with $\ell = 3n$)*

$$B_\ell = \frac{2\,(3n)!}{(n+2)!\,(n+1)!\,n!}$$

(and zero if ℓ is not a multiple of 3).

Proof. In a triangulation $2E = 3F$ so the number of edges is a multiple of 3. Since S and N are adjacent, there is a unique embedding in the disk so that the west boundary has length 1 and the east boundary has length 2. The lattice walks as discussed there go from $(0,0)$ to $(1,0)$. It is convenient to concatenate the walk with a final $m_{1,0}$ step, so that the walks are from $(0,0)$ to $(0,0)$ of length ℓ and remain in the first quadrant. Applying a shear $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$, the walks with steps $m_e, m_{0,1}, m_{1,0}$ become walks with steps $(1,0), (0,1), (-1,-1)$ which remain in the domain $y \geq x \geq 0$. Equivalently this is the number of walks from $(0,0,0)$ to (n,n,n) with steps $(1,0,0), (0,1,0), (0,0,1)$ remaining in the domain $y \geq x \geq z$. These are the so-called 3D Catalan numbers, see A005789 in the OEIS. $\qquad\square$

There are other formulas for bipolar-oriented "near-triangulations" [BM11, Prop. 5.3].

3.2   Vertex degree

Using the bijection between paths and bipolar-oriented maps, we can easily get the distribution of vertex degrees of a large bipolar-oriented triangulation.

Proposition 8. *In a large bipolar-oriented planar triangulation, the limiting in-degree and out-degree distributions of a random vertex are independent and geometrically distributed with mean 3.*

Proof. We examine the construction of bipolar-oriented planar maps when the steps give triangles. Any new vertex or new edge is adjoined to the marked bipolar map on its eastern boundary, which we also call the *frontier*.

A new vertex is created by an $m_{0,1}$ move, or an m_e move if there are currently no frontier vertices above the active vertex, and when a vertex is created it is the active vertex. Each subsequent move moves frontier vertices relative to the active vertex, so let us record their position with respect to the active vertex by integers, with positive integers recording the position below the active vertex and negative integers recording the position above it.

The following facts are easily verified.

1. A vertex moves off the frontier exactly when it is at position 1 and an $m_{1,0}$ move takes place.

2. m_e moves increase the index of vertices by 1.

3. $m_{0,1}$ moves decrease the index of a vertex by 1 if it is non-positive, else leave it fixed.

4. $m_{1,0}$ moves decrease the index by 1 if it is ≥ 2, else leave it fixed (if the index is 1 it is moved off of the frontier).

5. The up-degree of a vertex increases by 1 each time it visits position 0, the down-degree increases each time it visits position 1.

The transition diagram is summarized here:

off frontier
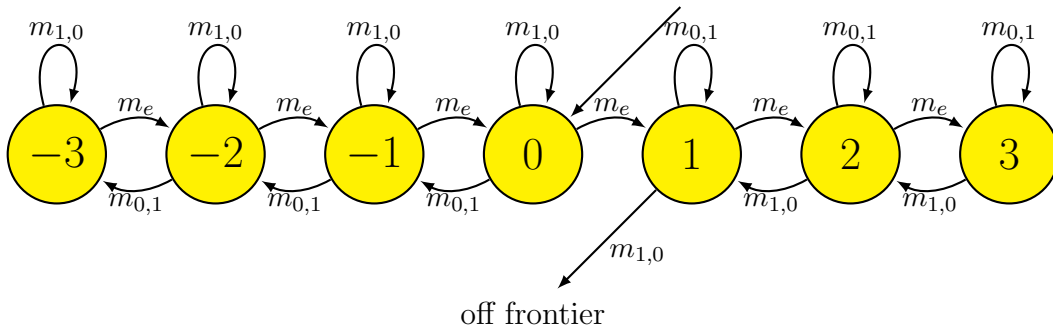
For the purposes of computing the final degree of a vertex we can simply count the number of visits to 0 before its index becomes positive, then count the number of visits to 1 before it is absorbed. Since $m_e, m_{0,1}, m_{1,0}$ occur with probabilities $1/3, 1/3, 1/3$, both of these counts are geometric with mean 3. □

4 Scaling limit

4.1 Statement

The proof of the following theorem is an easy computation upon application of the infinite-volume tree-mating theory introduced in [DMS14], a derivation of the relationship between the SLE/LQG parameters and a certain variance ratio in [DMS14, GHMS15], and a finite volume elaboration in [MS15e]. For clarity and motivational purposes we will reverse the standard conventions and give the proof first, explaining the relevant background in the following subsection.

Theorem 9. *The scaling limit of the bipolar-oriented planar map with its interface curve, with fixed boundary lengths $m + 1$ and $n + 1$, and number of edges $\ell \to \infty$ (with a possible congruence restriction on ℓ, m, and n to ensure such maps exist), with respect to the peanosphere topology, is a $\sqrt{4/3}$-LQG sphere decorated by an independent SLE_{12} curve.*

Proof. In Section 2.3 it was shown that the contour function (X_t, Y_t) for the bipolar-oriented random planar map converges as $\ell \to \infty$ to a Brownian excursion in the nonnegative quadrant with increments (X, Y) having covariance matrix (up to scale) $\begin{pmatrix} 2/3 & -1/3 \\ -1/3 & 2/3 \end{pmatrix}$, that is $X - Y$ and $X + Y$ are independent, and $\mathrm{Var}[X - Y] = 3\mathrm{Var}[X + Y]$.

The fact that the limit is a Brownian excursion implies, by [DMS14, Theorem 1.13] and the finite volume variant in [MS15e] and [GHMS15, Theorem 1.1], that the scaling limit in the peanosphere topology is a peanosphere, that is, a γ-LQG sphere decorated by an independent space-filling $\mathrm{SLE}_{\kappa'}$. The values γ, κ' are determined by the covariance structure of the limiting Brownian excursion. The ratio of variances $\mathrm{Var}[X - Y]/\mathrm{Var}[X + Y]$ takes the form

$$(1 + \cos[4\pi/\kappa'])/(1 - \cos[4\pi/\kappa']). \tag{3}$$

This relation was established for $\kappa \in [2, 4)$ and $\kappa' \in (4, 8]$ in [DMS14], and more generally for $\kappa \in (0, 4)$ and $\kappa' \in (4, \infty)$ in [GHMS15].[4] Setting it equal to 3 and solving we find $\kappa' = 12$. For this value of κ' we have $\gamma = \sqrt{\kappa} = \sqrt{16/\kappa'} = \sqrt{4/3}$. □

Remark 3. *If the covariance ratios vary as in Remark 1, then the κ' values varies between 8 and ∞. In other words, one may obtain any $\kappa' \in (8, \infty)$, and corresponding $\gamma = \sqrt{16/\kappa'}$, by introducing weightings that favor faces more or less balanced.*

Remark 4. *The infinite-volume variant described in Remark 2 corresponds to the mated pair of infinite-diameter trees first described in [DMS14], which in turn corresponds to the so-called γ-quantum cone described in the next subsection.*

4.2 Peanosphere background

The purpose of this section is to give a brief description of how Liouville quantum gravity (LQG) surfaces [DS11] decorated by independent SLE processes can be viewed as matings of random trees which are related to Aldous' continuum random tree (CRT) [Ald91a, Ald91b, Ald93]. The results that underly this perspective are established in [DMS14, MS15e], building on prior results from [DS11, She10, She09, MS12a, MS12b, MS12c, MS13a].

Recall that if h is an instance of the Gaussian free field (GFF) on a planar domain D with zero-boundary conditions and $\gamma \in (0, 2)$, then the γ-LQG surface associated with h of parameter γ is described by the measure μ_h on D which formally

---

[4]There is as yet no analogous construction corresponding to the limiting case $\kappa = \kappa' = 4$, where (3) is zero so that $\mathrm{Var}(X - Y) = 0$ and $X = Y$ a.s. It is not clear what such a construction would look like, given that space-filling $\mathrm{SLE}_{\kappa'}$ has only been defined for $\kappa' > 4$, not for $\kappa' = 4$, and the peanosphere construction in Section 4.2 is trivial when the limiting Brownian excursion is supported on the diagonal $x = y$.
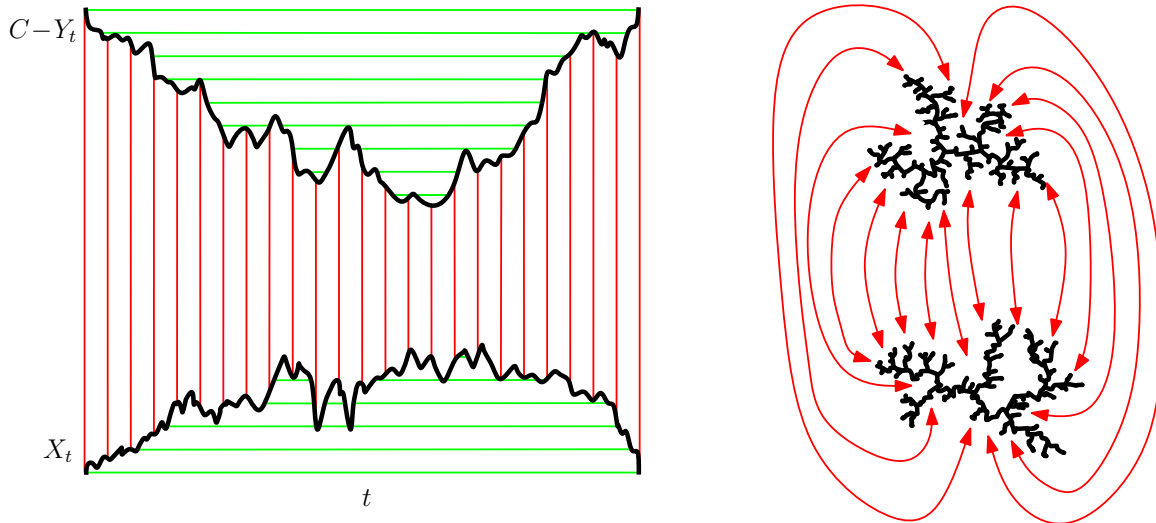
Figure 5: Gluing together a pair of CRTs to obtain a topological sphere. Illustration of the peanosphere construction. (This figure first appeared in [DMS14].)

has density $e^{\gamma h}$ with respect to Lebesgue measure. As h is a distribution and does not take values at points, this expression requires interpretation. One can construct this measure rigorously by considering approximations h_ε to h (by averaging the field on circles of radius ϵ) and then take μ_h to be the weak limit as $\varepsilon \to 0$ of $\varepsilon^{\gamma^2/4} e^{h_\varepsilon(z)} dz$ where dz denotes Lebesgue measure on D; see [DS11]. If one has two planar domains D_1, D_2, a conformal transformation $\varphi \colon D_1 \to D_2$, an instance of the GFF h_2 on D_2, and lets

$$ h_1 = h_2 \circ \varphi + Q \log |\varphi'| \quad \text{where} \quad Q = \frac{2}{\gamma} + \frac{\gamma}{2} \tag{4} $$

then the γ-LQG measure μ_{h_2} associated with h_2 is a.s. the image under φ of the γ-LQG measure μ_{h_1} associated with h_1. A *quantum surface* is an equivalence class of fields h where we say that two fields are equivalent if they are related as in (4).

This construction generalizes to any law on fields h which is absolutely continuous with respect to the GFF. The results in this article will be related to two such laws [She10, DMS14]:

1. The γ-quantum cone (an infinite-volume surface).

2. The γ-LQG sphere (a finite-volume surface).

We explain how they can both be constructed with the ordinary GFF h as the starting point.

The γ-quantum cone can be constructed by the following limiting procedure starting with an instance of the GFF h as above. Fix a constant $C > 0$ and note that adding C to h has the effect of multiplying areas as measured by μ by the factor $e^{\gamma C}$.

16

If one samples $z \in D$ according to μ and then rescales the domain so that the mass assigned by μ_{h+C} to $B(z,1)$ is equal to 1 then the law one obtains in the $C \to \infty$ limit is that of a γ-quantum cone. (The construction given in [She10, DMS14] is more direct in the sense that a precise recipe is given for sampling from the law of the limiting field.) That is, a γ-quantum cone is the infinite-volume γ-LQG surface which describes the local behavior of an γ-LQG surface near a μ_h-typical point.

The (unit area) γ-LQG sphere can also be constructed using a limiting procedure using the ordinary GFF h as above as the starting point. This construction works by first fixing $C > 0$ large, $\epsilon > 0$ small, and then *conditioning* on the event that the amount of mass that μ assigns to D is in $[e^{\gamma C}, e^{\gamma(C+\epsilon)}]$, so that the amount mass assigned to D by μ_{h-C} is in $[1, e^{\gamma\epsilon}]$, then sends first $C \to \infty$ and then $\epsilon \to 0$. (The constructions given in [DMS14, MS15e] are more direct because they involve precise recipes for sampling from the law of the limiting h.) One can visualize this construction by imagining that conditioning the area to be large (while keeping the boundary values of h constrained to be 0) leads to the formation of large a bubble. In the $C \to \infty$ limit, the opening of the bubble (which is the boundary of the domain) collapses to a single point, and it turns out that this point is typical (i.e., conditioned on the rest of the surface its law is given by that of the associated γ-LQG measure).

In [DMS14, MS15e], it is shown that it is possible to represent various types of γ-LQG surfaces (cones, spheres, and disks) decorated by an independent SLE as a gluing of a pair of continuous trees. We first explain a version of this construction in which $\gamma = \sqrt{2}$ and the surface is a unit-area LQG sphere decorated with an independent SLE_8. Let X and Y be independent one-dimensional Brownian excursions parametrized by $[0,1]$. Let C be large enough so that the graphs of X and $C - Y$ are disjoint, as illustrated in Figure 5. We define an equivalence relation $\sim$ on the rectangle $R = [0,1] \times [0,C]$ by declaring to be equivalent points which lie on either:

1. horizontal chords either entirely below the graph of X or entirely above graph of $C - Y$ (green lines in Figure 5), or

2. vertical chords between the graphs of X and $C - Y$ (red lines in Figure 5).

We note that under $\sim$, all of ∂R is equivalent so we may think of $\sim$ as an equivalence relation on the two-dimensional sphere $\mathbb{S}^2$. It is elementary to check using Moore's theorem [Moo25] (as explained in [DMS14, Section 1.1]) that almost surely the topological structure associated with $R/\sim$ is homeomorphic to $\mathbb{S}^2$. This sphere comes with additional structure, namely:

1. a space-filling path[5] η' (corresponding to the projection of the path which follows the red lines in Figure 5 from left to right), and

2. a measure μ (corresponding to the projection of Lebesgue measure on $[0,1]$).

---

[5]As explained just below, η' is related to an $\text{SLE}_{\kappa'}$ curve with $\kappa' > 4$. We use the convention here from [MS12a, MS12b, MS12c, MS13a], which is to use a prime whenever $\kappa' > 4$.

We refer to this type of structure as a *peanosphere*, as it is a topological sphere decorated with a path which is the peano curve associated with a space-filling tree.

The peanosphere associated with the pair (X, Y) does not *a priori* come with an embedding into the Euclidean sphere $\mathbb{S}^2$. However, it is shown in [DMS14, MS15e] that there is a canonical embedding (up to Möbius transformations) of the peanosphere associated with (X, Y) into $\mathbb{S}^2$, which is measurable with respect to (X, Y). This embedding equips the peanosphere with a conformal structure. The image of μ under this embedding is a $\sqrt{2}$-LQG sphere, see [DMS14, MS15e] as well as [DKRV14, AHS15]), and the law of the space-filling path η' is the following natural version of SLE_8 in this context [MS13a]: If we parametrize the $\sqrt{2}$-LQG sphere by the Riemann sphere $\widehat{\mathbb{C}}$, then η' is equal to the weak limit of the law of an SLE_8 on $B(0, n)$ from $-in$ to in with respect to the topology of local uniform convergence when parametrized by Lebesgue measure. (The construction given in [MS13a] is different and is based on the GFF.) The random path η' and the random measure μ are coupled together in a simple way. Namely, given μ, one samples from the law of the path by first sampling an SLE_8 (modulo time parametrization) independently of μ and then reparametrizing it according to μ-area (so that in t units of time it fills t units of μ-area).

This construction generalizes to all values of $\kappa' \in (4, \infty)$. In the more general setting, we have that $\gamma = \sqrt{\kappa}$ where $\kappa = 16/\kappa' \in (0, 4)$, and the pair of independent Brownian excursions is replaced with a continuous process (X, Y) from $[0, 1]$ into $\mathbb{R}_{\geq 0}^2$ which is given by the linear image of a two-dimensional Brownian excursion from the origin to the origin in the Euclidean wedge of opening angle

$$\theta = \frac{\pi \gamma^2}{4} = \frac{\pi \kappa}{4} = \frac{4\pi}{\kappa'}$$

see [DMS14, MS15e, GHMS15]. (In the infinite-volume version of the peanosphere construction, the Brownian excursions (X, Y) are replaced with Brownian motions, and the corresponding underlying quantum surface is a γ-quantum cone [DMS14].)

The main results of [DMS14, MS15e] imply that the information contained in the pair (X, Y) is a.s. *equivalent* to that of the associated $\mathrm{SLE}_{\kappa'}$-decorated γ-LQG surface. More precisely, the map f from $\mathrm{SLE}_{\kappa'}$-decorated γ-LQG surfaces to Brownian excursions is almost everywhere well-defined and almost everywhere invertible, and both f and f^{-1} are measurable.

The peanosphere construction leads to a natural topology on surfaces decorated by a space-filling tree and dual tree. Namely, we say that two such tree-decorated surfaces are close if the interface functions (i.e., the function which records the distance of a point on the tree to the root when one traces its boundary with unit speed) of the corresponding tree/dual-tree pair are uniformly close. The topology induced by sup-norm metric on interface functions is called the peanosphere topology.

The peanosphere approach to SLE/LQG convergence (i.e., identifying a natural pair of trees in the discrete model and proving convergence in the topology where two configurations are close if their tree contour functions are close) was introduced in

[She11, DMS14] to deal with infinite-volume limits of FK-cluster-decorated random planar maps, which correspond to $\kappa \in [2, 4)$ and $\kappa' \in (4, 8]$. Extensions to the finite volume case and a "loop structure" topology appear in [GMS15, GS15a, GS15b, GM15].

Since bipolar-oriented planar maps converge in the peanosphere topology to SLE_{12}-decorated $\sqrt{4/3}$-LQG, we conjecture that they also converge in other natural topologies, such as

- The *conformal path topology* defined as follows. Assume we have selected a method of "conformally embedding" discrete planar maps in the sphere. (This might involve circle packing, Riemann uniformization, or some other method.) Then the green path in Figure 2 becomes an actual path: a function η_n from $[0, 1]$ to the unit sphere (where n is the number of lattice steps) parameterized so that at time k/n the path finishes traversing its kth edge. An SLE_{12}-decorated $\sqrt{4/3}$-LQG sphere can be described similarly by letting η be the SLE path parameterized so that a t fraction of LQG volume is traversed between times 0 and t. (Note that the parameterized path η encodes both the LQG measure *and* the SLE path.) The conformal path topology is the uniform topology on the set of paths from $[0, 1]$ to the sphere. The conjecture is that η_n converges to η weakly w.r.t. the uniform topology on paths. See [DS11, She10] for other conjectures of this type.

- The Gromov–Hausdorff–Prokorov topology on metric measure spaces. So far, this problem has only been solved in the setting of uniformly random planar maps and $\sqrt{8/3}$-LQG in the works [LG13, Mie13] (Gromov–Hausdorff–Prokorhov convergence of the planar maps to a limit, the Brownian map) and [MS13b, MS15a, MS15b, MS15c, MS15d, MS15e] (construction of the metric space structure on $\sqrt{8/3}$-LQG). It is still an open problem to endow γ-LQG with a canonical metric space structure for $\gamma \neq \sqrt{8/3}$.

An interesting problem which illustrates some of the convergence issues that arise is the following: In the discrete setting, the interface functions between the NW and SE trees determine the bipolar map which in turn determine the interface functions between the NE and SW trees. Likewise, in the continuous setting, the interface functions (a Brownian excursion) between the NW and SE trees a.s. determine the SLE-decorated LQG which in turn a.s. determine the interface function (another Brownian excursion) between the NE and SW trees.

Conjecture 1. *The joint law of both NW/SE and NE/SW interface functions of a random bipolar-oriented planar map converges to the joint law of both NW/SE and NE/SW interface functions of SLE_{12}-decorated $\sqrt{4/3}$-LQG.*

One might expect to be able to approximate the discrete NW/SE interface function with a continuous function, obtain the corresponding continuous NE/SW

function, and hope that this approximates the discrete NE/SW function. One problem with this approach is that while the maps f^{-1} and f are measurable, they are (presumably almost everywhere) discontinuous, so that even if two interface functions are close, it does not follow that the corresponding measures and paths are close. However, since Brownian excursions are random perturbations rather than "worst case" perturbations of random walk excursions, we expect the joint laws to converge despite the discontinuities of f and f^{-1}.

5 Imaginary geometry: why $\kappa' = 12$ is special

When proving that a family of discrete random curves has SLE_κ as a scaling limit, it is sometimes possible to figure out in advance what κ should be by proving that there is only one κ for which SLE_κ has some special symmetry. For example, it is by now well known that SLE_6 is the only SLE curve with a certain *locality* property (expected of any critical percolation interface scaling limit) and that $SLE_{8/3}$ is the only SLE curve with with a certain *restriction* property (expected of any self-avoiding-walk scaling limit). The purpose of this section is to use the imaginary geometry theory of [MS12a, MS13a] to explain what is special about the values $\kappa = 4/3$ and $\kappa' = 12$.

5.1 Winding height gap for uniform spanning trees

The connections between winding height functions, statistical mechanics models, and height gaps are nicely illustrated in the discrete setting by the uniform spanning tree (UST). Temperley showed that spanning trees on the square grid are in bijective correspondence with dimer configurations (perfect matchings) on a larger square grid. Dimer configurations have a height function which is known to converge to the Gaussian free field [Ken00]. Under the Temperley correspondence, this dimer height function is related to the "winding" of the spanning tree, where the winding of a given edge in the tree is defined to be the number of right turns minus the number of left turns taken by the tree path from that edge to the root [KPW00]. If we multiply the dimer heights by $\pi/2$, then this function describes the accumulated amount of angle by which the path has turned right on its journey toward the root.

Notice that if v is a vertex along a branch of a spanning tree, and there are two tree edges off of that branch that merge into the vertex v from opposite directions, then the winding height at the edge just to the right of the branch is larger by π than at the edge just to the left, see Figure 6. Because of this, it is intuitively natural to expect that there will typically be a "winding height gap" across the long tree branch of magnitude π, i.e., the winding just right of the long tree branch should be (on average) larger by π than the winding just left of the tree branch.
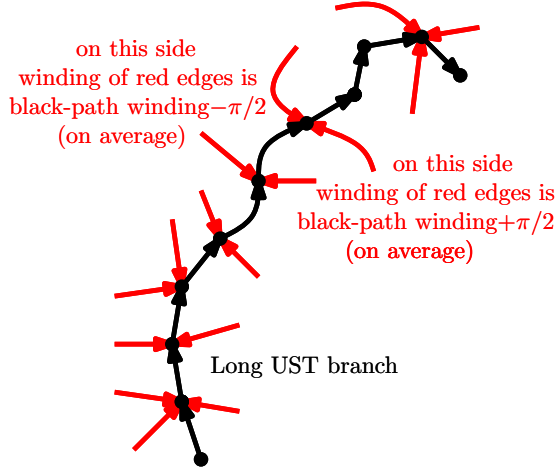
20

Figure 6: Every edge in a uniform spanning tree is assigned a "winding" — a real number that indicates the total amount of right turning (minus left turning) that takes place as one moves along the tree from the midpoint of that edge to the root. The black directed edges are those of a long branch in the UST directed towards the root, and adjacent spanning tree edges are shown in red. Given the long black path, the conditional law of the configuration on the left side of the path does not depend on the orientation of the long black path. Thus, by symmetry, one expects (on average) a $\pi/2$ angle gap between red edges and their black neighbors. This accounts for a total "average winding gap" of π between the left and right side.

5.2   Winding height gap for SLE

Imaginary geometry extends these notions of winding height function and the height gap to SLE. The starting point is an instance h of the GFF, which we divide by a parameter $\chi > 0$ to convert into units of radians. SLE_κ can be constructed as a flow line of the vector field in which z is assigned the complex unit vector $e^{ih/\chi}$, where

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

Although this vector field does not make literal sense, as h is a distribution and not a function, one can still construct the flow lines in a natural way [MS12a, MS13a]. While it has been conjectured that these GFF flow lines are limits of flow lines of mollified versions of the GFF, their rigorous construction follows a different route. One first proves that they are the unique paths coupled with the GFF that satisfy certain axiomatic properties (regarding the conditional law of the field given the path), and then establishes *a posteriori* that the paths are uniquely determined by the GFF.

We interpret flow lines of $e^{ih/\chi}$ as "east going" rays in an "imaginary geometry". A ray of a different angle θ is a flow line of $e^{i[h/\chi+\theta]}$. In contrast to Euclidean geometry, the rays of different angles started from the same point may intersect each other.

21

There is a critical angle θ_c given by

$$\theta_c = \frac{\pi\kappa}{4 - \kappa} = \frac{4\pi}{\kappa' - 4},$$

such that the flow lines of $e^{ih/\chi}$ and $e^{i[h/\chi+\theta]}$ started from a common point a.s. intersect when $\theta < \theta_c$, and a.s. do not intersect when $\theta \geq \theta_c$. If we condition on a flow line η, there is a winding height gap in the GFF, in the sense that $\mathbb{E}[h/\chi \mid \eta]$ just to the right of η is larger by θ_c than the value just to the left of η.

The east-going flow lines of $e^{ih/\chi}$ started from different points can intersect, at which point they merge. The collection of east-going rays started from all points together form a continuum spanning tree, whose branches are SLE_κ's. There is a space-filling curve which is the analog of the UST peano curve, which traces the boundary of this spanning tree, and is a space-filling version of $\mathrm{SLE}_{\kappa'}$ with $\kappa' = 16/\kappa$.

When $\kappa = 4n/(n+1)$ the critical angle is $\theta_c = n\pi$. The well-known "special values" of κ have in the past corresponded to integer values of n, together with the limiting case $n \to \infty$. For example, $n \in \{1, 2, 3, 5, \infty\}$ gives $\kappa \in \{2, 8/3, 3, 10/3, 4\}$ and $\kappa' \in \{8, 6, 16/3, 24/5, 4\}$. In the case $\kappa' = 12$ and $\kappa = 4/3$, we have $n = 1/2$.

5.3 Bipolar winding height gap should be $\pi/2$

To make sense of winding angle in the context of a planar map, one may view the map as a Riemannian surface obtained by interpreting the faces as regular unit polygons glued together, and then conformally map that surface, as in Figure 7. If a random planar map is decorated by a bipolar orientation, we can assign a winding to every edge that indicates the total amount of right turning (minus left turning) that takes
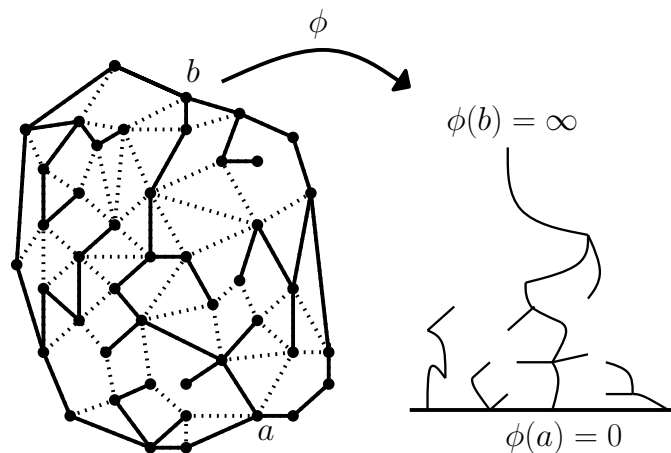

Figure 7: Planar map with a distinguished outer-boundary-plus-one-chord-rooted spanning tree (solid black edges), with chord joining marked boundary points a and b, plus image of tree under conformally uniformizing map ϕ to $\mathbb{H}$ (sketch). (This figure first appeared in [She10].)

22

place as one moves along *any* north-going path (it doesn't matter which one) from the midpoint of that edge to the north pole.

Consider a NW path started from a vertex incident to the eastern face continued to a vertex incident to the western face. The portion of the bipolar-oriented map that is south of this path may be east-west reflected to obtain a new bipolar-oriented planar map. If this NW path is suitably chosen, so that it is determined by the portion of the map north of it, then reflecting the portion of the map south of it is a bijection.

This east-west symmetry suggests that the edges to the left of a long NW path have a winding which is, on average, $\pi/2$ less than the winding of the NW path (as in Figure 8).

The bipolar-oriented map to the right of the NW path does not have this same reflection symmetry. Indeed, to the right of the NW path (but not the left), there can be other north-going paths that split off the NW path and rejoin the NW path at another vertex. Reflecting the bipolar map on the right side of the NW path would then create a cycle.

However, if we reflect the map to the right of the NW path and then reverse the orientations of the edges, no cycles are created, and no new sources or sinks are created except at the endpoints of the NW path. Thus for a long NW path, we expect the bipolar-oriented map to the right of the path to be approximately reflection-reversal symmetric.

Edges to the right of the NW path may be oriented either toward or away from the path; one expects those oriented away from the path to have smaller winding on average and those oriented toward the path to have larger winding on average. However, by the approximate reflection-reversal symmetry, these two effects should cancel, so that overall there is no expected angle gap between the black path and the red edges to its right (as in Figure 8).
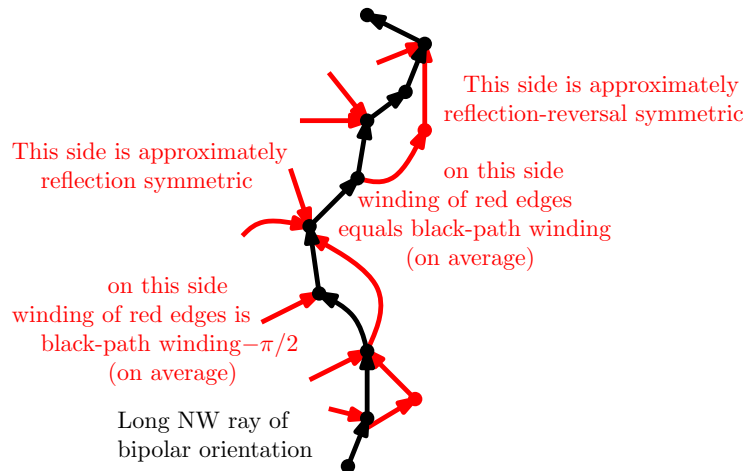

Figure 8: The average height gap on either side of a long NW ray (black) of a bipolar orientation. On the western side the average height gap is $-\pi/2$, on the eastern side it is 0.

23

6 Open questions

In addition to questions regarding strengthening the topology of convergence, which are discussed at the end of Section 4.2, it would be interesting to extend the theory to other surface graphs, such as the torus, or a disk with four boundary vertices which are alternately source, sink, source, sink.

References

[AHS15]   J. Aru, Y. Haung, and X. Sun. Two perspectives of the unit area quantum sphere and their equivalence. 2015. In preparation.

[AK15]   A. Abrams and R. Kenyon. Fixed-energy harmonic functions. 2015. arXiv:1505.05785.

[Ald91a] D. Aldous. The continuum random tree. I. *Ann. Probab.*, 19(1):1–28, 1991. MR1085326

[Ald91b] D. Aldous. The continuum random tree. II. An overview. In *Stochastic analysis*, London Math. Soc. Lecture Note Ser. #167, pages 23–70. Cambridge Univ. Press, 1991. MR1166406

[Ald93]   D. Aldous. The continuum random tree. III. *Ann. Probab.*, 21(1):248–289, 1993. MR1207226

[BBMF11] N. Bonichon, M. Bousquet-Mélou, and É. Fusy. Baxter permutations and plane bipolar orientations. *Sém. Lothar. Combin.*, 61A:Article B61Ah, 2009/11. MR2734180

[BM11]   M. Bousquet-Mélou. Counting planar maps, coloured or uncoloured. In *Surveys in combinatorics 2011*, London Math. Soc. Lecture Note Ser. #392, pages 1–49. Cambridge Univ. Press, 2011. MR2866730

[CDCH+14] D. Chelkak, H. Duminil-Copin, C. Hongler, A. Kemppainen, and S. Smirnov. Convergence of Ising interfaces to Schramm's SLE curves. *C. R. Math. Acad. Sci. Paris*, 352(2):157–161, 2014. MR3151886

[CV81]   R. Cori and B. Vauquelin. Planar maps are well labeled trees. *Canad. J. Math.*, 33(5):1023–1042, 1981. MR638363

[dFdMR95] H. de Fraysseix, P. O. de Mendez, and P. Rosenstiehl. Bipolar orientations revisited. *Discrete Appl. Math.*, 56(2-3):157–179, 1995. MR1318743

[DKRV14] F. David, A. Kupiainen, R. Rhodes, and V. Vargas. Liouville quantum gravity on the Riemann sphere. 2014. arXiv:1410.7318.

[DMS14]   B. Duplantier, J. Miller, and S. Sheffield. Liouville quantum gravity as a mating of trees. 2014. arXiv:1409.7055.

[DS11]   B. Duplantier and S. Sheffield. Liouville quantum gravity and KPZ. *Invent. Math.*, 185(2):333–393, 2011. MR2819163

[Dub09]   J. Dubédat. Duality of Schramm-Loewner evolutions. *Ann. Sci. Éc. Norm. Supér. (4)*, 42(5):697–724, 2009. MR2571956

[Dup98]   B. Duplantier. Random walks and quantum gravity in two dimensions. *Phys. Rev. Lett.*, 81(25):5489–5492, 1998. MR1666816

[DW15]   J. Duraj and V. Wachtel. Invariance principles for random walks in cones. 2015. arXiv:1508.07966.

[FFNO11]   S. Felsner, É. Fusy, M. Noy, and D. Orden. Bijections for Baxter families and related objects. *J. Combin. Theory Ser. A*, 118(3):993–1020, 2011. MR2763051

[FPS09]   É. Fusy, D. Poulalhon, and G. Schaeffer. Bijective counting of plane bipolar orientations and Schnyder woods. *European J. Combin.*, 30(7):1646–1658, 2009. MR2548656

[GHMS15]   E. Gwynne, N. Holden, J. Miller, and X. Sun. Brownian motion correlation in the peanosphere for $\kappa > 8$. 2015, arXiv:1510.04687.

[GKMW15]   E. Gwynne, A. Kassel, J. Miller, and D. B. Wilson. Topologically-weighted spanning trees on planar maps and SLE-decorated Liouville quantum gravity for $\kappa \geq 8$. 2015. In preparation.

[GM15]   E. Gwynne and J. Miller. Convergence of the topology of critical Fortuin-Kasteleyn planar maps to that of CLE_κ on a Liouville quantum surface. 2015. In preparation.

[GMS15]   E. Gwynne, C. Mao, and X. Sun. Scaling limits for the critical Fortuin-Kasteleyn model on a random planar map I: cone times. 2015. arXiv:1502.00546.

[GS15a]   E. Gwynne and X. Sun. Scaling limits for the critical Fortuin-Kastelyn model on a random planar map II: local estimates and empty reduced word exponent. 2015. arXiv:1505.03375.

[GS15b]   E. Gwynne and X. Sun. Scaling limits for the critical Fortuin-Kastelyn model on a random planar map III: finite volume case. 2015. In preparation.

[Ken00]   R. Kenyon. Conformal invariance of domino tiling. *Ann. Probab.*, 28(2):759–795, 2000. MR1782431

[KMSW15]   R. W. Kenyon, J. Miller, S. Sheffield, and D. B. Wilson. Square ice and space-filling SLE. 2015. In preparation.

[KPW00]   R. W. Kenyon, J. G. Propp, and D. B. Wilson. Trees and matchings. *Electron. J. Combin.*, 7:Research Paper 25, 34 pp., 2000. MR1756162

[KW15]   A. Kassel and D. B. Wilson. Activity-weighted trees. 2015. In preparation.

[LG13]   J.-F. Le Gall. Uniqueness and universality of the Brownian map. *Ann. Probab.*, 41(4):2880–2960, 2013. MR3112934

[LSW01a]   G. F. Lawler, O. Schramm, and W. Werner. The dimension of the planar Brownian frontier is 4/3. *Math. Res. Lett.*, 8(4):401–411, 2001. MR1849257

[LSW01b]   G. F. Lawler, O. Schramm, and W. Werner. Values of Brownian intersection exponents. I. Half-plane exponents. *Acta Math.*, 187(2):237–273, 2001. MR1879850

[LSW01c]   G. F. Lawler, O. Schramm, and W. Werner. Values of Brownian intersection exponents. II. Plane exponents. *Acta Math.*, 187(2):275–308, 2001. MR1879851

[LSW02]   G. F. Lawler, O. Schramm, and W. Werner. Values of Brownian intersection exponents. III. Two-sided exponents. *Ann. Inst. H. Poincaré Probab. Statist.*, 38(1):109–123, 2002. MR1899232

[LSW04]   G. F. Lawler, O. Schramm, and W. Werner. Conformal invariance of planar loop-erased random walks and uniform spanning trees. *Ann. Probab.*, 32(1B):939–995, 2004. MR2044671

[Mie13]   G. Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.*, 210(2):319–401, 2013. MR3070569

[Moo25]   R. L. Moore. Concerning upper semi-continuous collections of continua. *Trans. Amer. Math. Soc.*, 27(4):416–428, 1925. MR1501320

[MS12a]   J. Miller and S. Sheffield. Imaginary geometry I: Interacting SLEs. 2012. arXiv:1201.1496.

[MS12b]   J. Miller and S. Sheffield. Imaginary geometry II: reversibility of $\mathrm{SLE}_\kappa(\rho_1; \rho_2)$ for $\kappa \in (0, 4)$. 2012. arXiv:1201.1497.

[MS12c]   J. Miller and S. Sheffield. Imaginary geometry III: reversibility of SLE_κ for $\kappa \in (4, 8)$. 2012. arXiv:1201.1498.

[MS13a]   J. Miller and S. Sheffield. Imaginary geometry IV: interior rays, whole-plane reversibility, and space-filling trees. 2013. arXiv:1302.4738.

[MS13b]   J. Miller and S. Sheffield. Quantum Loewner evolution. 2013. To appear in *Duke Math. J.* arXiv:1312.5745.

[MS15a]   J. Miller and S. Sheffield. An axiomatic characterization of the Brownian map. 2015. arXiv:1506.03806.

[MS15b]   J. Miller and S. Sheffield. Liouville quantum gravity and the Brownian map I: The QLE(8/3,0) metric. 2015. arXiv:1507.00719.

[MS15c]   J. Miller and S. Sheffield. Liouville quantum gravity and the Brownian map II: geodesics and continuity of the embedding. 2015. In preparation.

[MS15d]   J. Miller and S. Sheffield. Liouville quantum gravity and the Brownian map III: the conformal structure is determined. 2015. In preparation.

[MS15e]   J. Miller and S. Sheffield. Liouville quantum gravity spheres as matings of finite-diameter trees. 2015. arXiv:1506.03804.

[Mul67]   R. C. Mullin. On the enumeration of tree-rooted maps. *Canad. J. Math.*, 19:174–183, 1967. MR0205882

[Sch98]   G. Schaeffer. *Conjugaison d'arbres et cartes combinatoires aléatoires*. PhD thesis, Université Bordeaux I, 1998.

[Sch00]   O. Schramm. Scaling limits of loop-erased random walks and uniform spanning trees. *Israel J. Math.*, 118:221–288, 2000. math/9904022. MR1776084

[She09]   S. Sheffield. Exploration trees and conformal loop ensembles. *Duke Math. J.*, 147(1):79–129, 2009. MR2494457

[She10]   S. Sheffield. Conformal weldings of random surfaces: SLE and the quantum gravity zipper. 2010. To appear in *Ann. Probab.* arXiv:1012.4797.

[She11]   S. Sheffield. Quantum gravity and inventory accumulation. 2011. To appear in *Ann. Probab.* arXiv:1108.2241.

[Smi01]   S. Smirnov. Critical percolation in the plane: conformal invariance, Cardy's formula, scaling limits. *C. R. Acad. Sci. Paris Sér. I Math.*, 333(3):239–244, 2001. MR1851632

[Smi10]   S. Smirnov. Conformal invariance in random cluster models. I. Holomorphic fermions in the Ising model. *Ann. of Math. (2)*, 172(2):1435–1467, 2010. MR2680496

[SS09]   O. Schramm and S. Sheffield. Contour lines of the two-dimensional discrete Gaussian free field. *Acta Math.*, 202(1):21–137, 2009. MR2486487

[SS13]   O. Schramm and S. Sheffield. A contour line of the continuum Gaussian free field. *Probab. Theory Related Fields*, 157(1-2):47–80, 2013. MR3101840

[SW12]   S. Sheffield and W. Werner. Conformal loop ensembles: the Markovian characterization and the loop-soup construction. *Ann. of Math. (2)*, 176(3):1827–1917, 2012. MR2979861

[Tut63]   W. T. Tutte. A census of planar maps. *Canad. J. Math.*, 15:249–271, 1963. MR0146823

[Zha08]   D. Zhan. Duality of chordal SLE. *Invent. Math.*, 174(2):309–353, 2008. MR2439609