Schramm-Loewner Evolution
— and other —
conformally invariant objects

Vincent Beffara

Lectures given in Buzios (Brazil)
July 11 – August 7, 2010

Current version available at:

http://www.umpa.ens-lyon.fr/~vbeffara/files/Buzios.pdf

Contents

Foreword

These notes are not meant as a reference manual, but rather as an introduction combined with a kind of "user's guide" to existing bibliography. I plan to keep them mostly self-contained in the sense that the reader will need no additional information to understand the majority of the statements; but they contain essentially no detailed proof. In the case of very important results, I give indications about the main ideas of the demonstration, but of course that is hardly sufficient to a motivated student.

In each chapter, the most important section is therefore the extended bibliography at the end. I chose to gather all bibliographical references there, and to omit them from the main body of the text (in particular, the main results are only attributed to their respective authors, not to a particular publication). The point is to make reading through the text more natural; maybe it failed!

As a corollary, I chose to adopt an informal style. However, to distinguish vague statements of exact theorems from exact statements of conjectures, I will try to use the following adverbs consistently:

Essentially means that the statement is a simplification of an actual definition or theorem, but that the full level of precision is not necessary for the reader's understanding of the idea behind it — and in fact, that it might hide its core; all the gory details are published elsewhere and pointed at by the bibliographical notes at the end of each chapter.

Conjecturally means that the statement is widely believed to hold, and that it has at least some support from mathematically non-rigorous arguments, but that a formal proof is still missing. Physicists have an excellent track record here, so most of these should actually be true, although when they might be proved is hard to predict.

Hopefully means that the statement is a reasonable assumption to make about the actual behavior of the object under consideration, but whether it is actually true might depend on unstated (and often unknown) features.

Morally means that the statement corresponds to my own understanding of what happens; which means that it probably should not be trusted at all, and could be completely misleading. On the other hand, it might help the reader get an intuition of the topic.

The notes were started while I was giving a graduate course in Lyon during the spring preceding the school. As a result, they cover a certain quantity of material in addition to what will be / was discussed in Buzios (mostly the parts about random-cluster models and convergence to SLE, which correspond more to Smirnov's mini-course). These can constitute indications towards further reading, or can be ignored completely in a first reading.

For reference, here is a rough outline of the course schedule in Buzios. The contents of the exercise sessions will match these, and be added to the notes at the end of the school.

Course 1: Percolation and Cardy's formula.

Course 2: Loop-erased random walks and uniform spanning trees.

Course 3: Loewner chains in the radial case.

Course 4: Chordal Loewner chains, and definition of SLE.

Course 5: First properties of SLE.

Course 6: The locality property and SLE_6.

Course 7: The restriction property, $\mathrm{SLE}_{8/3}$ and restriction measures.

Course 8: More exotic objects: CLE, loop soups, Gaussian fields…

Chapter I

A few discrete models

An actual case suggested the following: An equal number of white and black balls of equal size are thrown into a rectangular box, what is the probability that there will be contiguous contact of white balls from one end of the box to the opposite end?

De Volson Wood, 1894, [22]

The problem is a pretty good one and if any one will furnish a complete solution to it, we will publish it in the next issue of the Monthly.

Editor's comment to [22]

CONTENTS

INTRODUCTION

The goal of these lectures is to provide a self-contained introduction to SLE and related objects, but some motivation is needed before introducing SLE as such; so it seems natural to start with a quick review of a few two-dimensional discrete models.

The focus of this chapter will be, for each model, to arrive to the question of scaling limits as quickly as possible, and to justify conformal invariance where it is known to hold in the limit. The proofs of actual convergence to SLE will of course have to be postponed (see Chapter III, Convergence to SLE) — but providing the key arguments is our main objective here.

1 LATTICE MODELS

We start with what we want to call *lattice models* — even though that might not exactly be the usual sense of that word. Essentially, given a (two-dimensional) lattice embedded in the plane, a *configuration* is a map from the set of vertices and/or edges of the lattice into a finite alphabet, and a probability measure on the set of configuration is constructed by taking a thermodynamical limit from measures in finite boxes derived from a Hamiltonian.

We choose to limit ourselves to models satisfying the *finite-energy condition* (*i.e.*, if two configurations differ only within a finite box, then their probabilities are comparable). This covers percolation, the Ising and Potts models, the random-cluster model, but excludes models such as spanning trees and domino tilings for which topological constraints are imposed on acceptable configurations. The uniform spanning tree (UST) is an important case because it was on of the first two-dimensional models for which convergence to SLE was proved; we will briefly come back to it in the next section in association to the loop-erased random-walk.

Besides, we will mostly be interested in models taken at their critical point, and defined on specific lattices for which more is understood about their asymptotic behavior (*e.g.*, we limit our description of percolation to the case of site percolation on the triangular lattice) — even though of course a lot is known in a more general setting.

1.1 PERCOLATION

The simplest lattice model to describe is *Bernoulli percolation*. Let $p \in (0,1)$ be a parameter; for each vertex of the triangular lattice $\mathcal{H}$, toss a coin and declare it to be *open* (resp. *closed*) with probability p (resp. $1-p$), independently of the others. Denote by P_p the corresponding probability measure on the set of configurations (it is simply a product measure). One can see a configuration as random subgraph of the underlying lattice, obtained by keeping the open vertices and all the edges connecting two open vertices.

— *Basic features of the model* —

The question of interest is that of the connectivity structure of this subgraph. Let

$$\theta(p) := P_p\left[0 \leftrightarrow \infty\right]$$

be the probability that the origin belongs to an infinite connected component, or *cluster* (*i.e.*, that it is "connected to infinity"). It is easy to show that the function θ is non-decreasing, and using a simple counting argument (known as a *Peierl's argument*), that for p small enough, $\theta(p)$ is equal to 0 and $\theta(1-p)$ is positive; in other words, defining

$$p_c := \inf\left\{p : \theta(p) > 0\right\} = \sup\left\{p : \theta(p) = 0\right\},$$

one has $0 < p_c < 1$. p_c is called the *critical point* of the model. Its value depends on the choice of the underlying lattice; in the case of the triangular lattice, it is equal to $1/2$.

The behavior of the system changes drastically across the critical point:

- If $p < p_c$, then almost surely all connected components are finite; moreover, they have finite expected volume, and the connection probabilities exhibit exponential decay: There exists $L(p) < \infty$ such that, for every $x, y \in \mathbb{Z}^2$,

$$P_p\left[x \leftrightarrow y\right] \leqslant C\, e^{-\|y-x\|/L(p)};$$

- If $p > p_c$, then almost surely there exists a *unique* infinite cluster and it has asymptotic density $\theta(p)$; but exponential decay still occurs for connectivity through finite clusters: There exists $L(p) < \infty$ such that for all $x, y \in \mathbb{Z}^2$,

$$P_p\left[x \leftrightarrow y; x \not\leftrightarrow \infty\right] \leqslant C\, e^{-\|y-x\|/L(p)};$$

- If $p = p_c$, there is no infinite cluster (*i.e.*, $\theta(p_c) = 0$) yet there is no finite characteristic length in the system; the two-point function has a power-law behavior, in the sense that for some $c > 0$ and for all $x, y \in \mathbb{Z}^2$,

$$c \, \|y - x\|^{-1/c} \leqslant P_p \left[x \leftrightarrow y\right] \leqslant c^{-1} \, \|y - x\|^{-c}.$$

The last statement is an instance of what is known as *Russo-Seymour-Welsh theory*: Essentially, the largest cluster within a large box of a given size has a diameter of the same order as the size of the box, and it crosses it horizontally with a positive probability, uniformly in the actual size of the box.

To be more specific, if $\mathcal{R}$ is a rectangle aligned with the axes of the lattice, denote by $LR(\mathcal{R})$ the probability that, within (the intersection between $\mathbb{Z}^2$ and) $\mathcal{R}$, there is a path of open edges connecting its two vertical sides. Then, RSW states that for every $\lambda > 0$, there exists $\eta(\lambda) \in (0, 1)$ such that, for every n large enough,

$$\eta(\lambda) \leqslant P_p \left[LR(\llbracket 0, \lambda n \rrbracket \times \llbracket 0, n \rrbracket)\right] \leqslant 1 - \eta(\lambda).$$

A natural question is then the following: Does the crossing probability above actually converge as $n \to \infty$? In fact, that question is still open in the general case, and in particular in the case of $\mathbb{Z}^2$, and it is not quite clear how many new ideas would be needed to prove convergence. But conjecturally, the limit does exist and does not depend on the choice of the underlying lattice, provided that it has enough symmetry — this is part of what is known as *universality*.

— *The Cardy-Smirnov formula* —

We now turn to the main result which we want to present in this section. In the case of site-percolation on the triangular lattice (which is the one we are considering here), in fact the above probability does converge, and the limit is known explicitly. This was first conjectured by Cardy using mathematically non-rigorous argument, and later proved by Smirnov.

Before stating the main theorem, we need some additional notation. Let Ω be a smooth, simply connected, bounded domain in the complex plane, and let a, b, c and d be four points on $\partial\Omega$, in this order if the boundary is oriented counterclockwise. Let $\delta > 0$, and consider the triangular lattice scaled by a factor of δ, which we will denote by $\mathcal{H}_\delta$; let $\mathcal{P}_\delta(\Omega, a, b, c, d)$ be the probability that, within percolation on $\Omega \cap \mathcal{H}_\delta$, there is an open path connecting the arc ab to the arc cd of the boundary.[1]

Theorem 1 (Cardy, Smirnov). *There exists a function f defined on the collection of all 5-tuples formed of a simply connected domain with four marked boundary points, satisfying the following:*

(i) *As $\delta \to 0$, $\mathcal{P}_\delta(\Omega, a, b, c, d)$ converges to $f(\Omega, a, b, c, d)$;*

(ii) *f is* conformally invariant, *in the following sense: If Ω and Ω' are two simply connected domains and if Φ maps Ω conformally to Ω', then*

$$f(\Omega, a, b, c, d) = f(\Omega', \Phi(a), \Phi(b), \Phi(c), \Phi(d));$$

(iii) *If $\mathcal{T}$ is an equilateral triangle, and if a, b and c are its vertices, then*

$$f(\mathcal{T}, a, b, c, d) = \frac{|cd|}{|ab|}$$

(which, together with the previous point, characterizes f uniquely).

[1] Technically, this definition would require constructing a discrete approximation of the domain; we choose to skip over such considerations here, and refer the avid reader to the literature for more detail.

A complete proof of this theorem can be found in several places, so it does not make much sense to produce yet another one here; instead, we briefly describe the main steps of Smirnov's general strategy in some detail. The same overall approach (though obviously with a few modifications) will be applied to other models below; the main point each time will be to find the correct *observable*, *i.e.* a quantity derived from the discrete model and which is computable enough that its asymptotic behavior can be obtained (and is non-trivial).

Step 1: Definition of the observable. Let Ω be as above, and let z be a vertex of the dual lattice $\mathcal{H}_\delta^*$ (or equivalently, a face of the triangular lattice); denote by $E_\delta^a(z)$ the event that there is a simple path of open vertices joining two points on the boundary of Ω and separating a and z on one side and b and c on the other side, and by $H_\delta^a(z)$ the probability of $E_\delta^a(z)$. Define H_δ^b and H_δ^c accordingly. Notice that if we choose $z = d$, we get exactly the crossing probability:

$$\mathcal{P}_\delta(\Omega, a, b, c, d) = H_\delta^a(d).$$

In fact, we will compute the limit of H_δ^a as $\delta \to 0$ in the whole domain; the existence of f will follow directly.

Step 2: Tightness of the observable. Let z and z' be two points within the domain, and let $\mathcal{A}$ be an annulus contained in Ω and surrounding both z and z'. If $\mathcal{A}$ contains an open circuit, then either both of the events $E_\delta^a(z)$ and $E_\delta^a(z')$ occur, or none of them does. The existence of such circuits in disjoint annuli are independent events, and if one fixes the modulus of the annuli, their probability is bounded below by RSW estimates. Besides, the number of such disjoint annuli which can be fit around $\{z, z'\}$ is of order $-\log|z'-z|$. This implies a bound of the form

$$\left|H_\delta^a(z') - H_\delta^a(z)\right| \leqslant C\,|z'-z|^c$$

for some C, $c > 0$ depending only on the domain Ω, but not on δ. In other words, the functions H_δ^a are uniformly Hölder of the same exponent and the same norm; and this implies, by Ascoli's theorem, that they form a relatively compact family of continuous maps from Ω to $[0,1]$. In particular, one can always choose a sequence (δ_k) going to 0 along which $H_{\delta_k}^a$ (as well as $H_{\delta_k}^b$ and $H_{\delta_k}^c$) converges to some continuous function h^a (resp. h^b, h^c) defined on $\bar\Omega$. Proving convergence of (H_δ^a) then amounts to proving the uniqueness of such a sub-sequential limit, *i.e.*, all that remains to be done is to identify the function h^a.

Step 3: Local behavior of the observable. This is essentially the only place in the proof where one uses the fact that the underlying model is site-percolation on a triangulation. Let z be a vertex of $\mathcal{H}_\delta^*$, and let z_1, z_2 and z_3 denote its three neighbors, ordered counter-clockwise; define $P_\delta^a(z, z_1)$ to be the probability that $E_\delta^a(z_1)$ occurs, but $E_\delta^a(z)$ does not. This is equivalent to the existence of three disjoint paths in $\mathcal{H}_\delta$, each joining one of the vertices of the triangle around z to one of the three boundary arcs delimited by a, b and c, and of appropriate state (two open, one closed — make a picture!). The core of Smirnov's proof is then a wonderful relation between these quantities; namely:

$$P_\delta^a(z, z_1) = P_\delta^b(z, z_2) = P_\delta^c(z, z_3).$$

The argument is very simple, but not easy to write down formally; it goes as follows: Assuming the existence of three arms as above, it is possible to discover two of them by exploring the percolation configuration starting from c (say), and always staying on the interface between open and closed vertices. The exploration path reaches z if and only if two of the above arms exist; besides, it gives us no information about the state of the vertices which are not along it, because the underlying measure is a product measure. The key remark is then the *color-swapping argument*: Changing the state of each of the vertices in the unexplored portion of Ω does not change the probability of the configuration (because we work at $p = p_c = 1/2$); but it does change the state of the third arm from open to closed. Swapping the colors of all the vertices in Ω (which still does not change probabilities) then arrives at a configuration with three arms of the appropriate colors, but where the role of a (resp. z_1) is now taken by b (resp. z_2).

Step 4: Holomorphicity in the scaling limit. Now, we need to exhibit a holomorphic function built out of h^a, h^b and h^c; following the symmetry of order 3 in the setup, it is natural to define

$$H_\delta(z) := H^a_\delta(z) + \tau H^b_\delta(z) + \tau^2 H^c_\delta(z)$$

and $h := h^a + \tau h^b + \tau^2 h^c$ accordingly, where $\tau = e^{2\pi i/3}$. To prove that h is holomorphic, it is enough to show that, along every smooth curve γ contained in Ω, one has

$$\oint_\gamma h(z)\,\mathrm{d}z = 0$$

(by Morera's theorem); and to show that, it is enough to pick a sequence of suitable discretizations of γ and estimate the integral using H_δ, and to show that the discrete estimate vanishes as δ goes to 0. It is always possible to approach γ by a discrete path $\gamma_\delta = (z^\delta_0, z^\delta_1, \dots, z^\delta_{L_\delta} = z^\delta_0)$ on $\mathcal{H}^*_\delta$ in such a way that $L_\delta = \mathcal{O}\big(\delta^{-1}\big)$, and one then has

$$\oint_\gamma h(z)\mathrm{d}z = \sum_{j=0}^{L_\delta - 1} \frac{H_\delta(z^\delta_j) + H_\delta(z^\delta_{j+1})}{2}(z^\delta_{j+1} - z^\delta_j) + \mathcal{O}(\delta^c)$$

with $c > 0$ by the previous tightness estimate. One can then apply a discrete analog of Green's formula to make discrete derivatives of H_δ appear, and write these in terms of P^a_δ, P^b_δ and P^c_δ: after elementary calculus, one gets

$$\oint_\gamma h(z)\mathrm{d}z = \frac{i\delta\sqrt{3}}{2} \sum_{z\sim z'} \Big[P^a_\delta(z,z') + \tau P^b_\delta(z,z') + \tau^2 P^c_\delta(z,z')\Big]\,(z'-z) + \mathcal{O}(\delta^c)\,,$$

where the sum extends to all pairs of nearest neighbors in the interior of γ. Applying Smirnov's identity to write everything in terms of P^a_δ only then leads to

$$\oint_\gamma h(z)\mathrm{d}z = \frac{i\sqrt{3}}{2} \sum_{z\sim z'} \left[P^a_\delta(z,z') \sum_{j=0}^{2} \tau^j (z_j - z)\right] + \mathcal{O}(\delta^c)$$

(where the z_j are the neighbors of z, numbered counterclockwise in such a way that $z_0 = z'$). It is then easy to see that the inner sum is identically equal to 0 (because it is always proportional to $1 + \tau^2 + \tau^4$).

Step 5: Boundary conditions and identification. The same computation as above can be performed starting with $S_\delta := H^a_\delta + H^b_\delta + H^c_\delta$, and the conclusion is the same: The (subsequential) limit $s := h^a + h^b + h^c$ is holomorphic as well. But because it is constant, this leads to the conclusion that it is constant, equal to 1 by looking at the point $z = a$. This means that the triple (h^a, h^b, h^c) can be seen as the barycentric coordinates of $h(z)$ relative to the points 1, τ and τ^2, respectively, meaning that h maps Ω to the interior of the corresponding equilateral triangle $\mathcal{T}$. Since it sends boundary to boundary in a one-to-one way (the variations of h^a on the boundary are easy to determine), it has to be conformal, and so it has to be the unique conformal map from Ω to $\mathcal{T}$ mapping a (resp. b, c) to 1 (resp. τ, τ^2). Because the subsequential limit is thus identified uniquely, one obtains convergence of (H_δ) itself to h, and it is not difficult to conclude the proof.

This concludes the few features of percolation which we will need in the next chapters; we will come back to it (and say a little bit more about the exploration process) in Chapter III, Convergence to SLE. For now, the relevant piece of information to remember is that, at criticality, the scaling limit of percolation (in any reasonable sense) is non-trivial and exhibits conformal invariance.

1.2 THE RANDOM-CLUSTER MODEL

Percolation is very easy to describe, because the states of the vertices are independent from each other; but it is not very physically realistic. We now focus our attention on the *random-cluster model* (sometimes also referred to as *FK-percolation*, for the names of its inventors, Fortuin and Kasteleyn). It is a dependent variant of *bond* percolation.

— *Definitions and first properties* —

Let $G = (V, E)$ be a finite graph, and let $q \in [1, +\infty)$ and $p \in (0, 1)$ be two parameters. The random-cluster measure on G is defined on the set of subgraphs of G, seen as subsets of E, by

$$P_{p,q,G}[\{\omega\}] := \frac{p^{o(\omega)}(1-p)^{c(\omega)}q^{k(\omega)}}{Z_{p,q,G}},$$

where $o(\omega)$ is the number of open edges in ω, $c(\omega)$ the number of closed edges, and $k(\omega)$ the number of connected components of the subgraph (counting isolated vertices). The *partition function* $Z_{p,q,G}$ is chosen so as to make the measure a probability measure. Notice that the case $q = 1$ is exactly that of a product measure, in other words it is Bernoulli bond-percolation on G.

The definition in the case of an infinite graph needs a little more care — of course defining the measure as above makes little sense as the exponents would be infinite. The first step is to define *boundary conditions*, which in this case amounts to introducing additional edges whose state is fixed (either open or closed); if ξ denotes such a choice, then $P^{\xi}_{p,q,G}$ denotes the corresponding measure. Notice that the only effect ξ has is in the counting of connected components within G.

Now consider the square lattice, and a sequence of increasing boxes $\Lambda_n := [\![-n, n]\!]^2$. We will consider two types of boundary conditions for the random-cluster model on Λ_n: *free* (*i.e.*, ξ is empty) and *wired* (*i.e.*, all the vertices on the boundary of Λ_n are assumed to be connected). We denote these boundary conditions by f and w, respectively. A third boundary condition is known as the *Dobrushin* boundary condition, and consists in wiring the vertices of one boundary arc of the box (with prescribed endpoints) together while leaving the rest of the boundary free.

If $q \geqslant 1$, the model exhibits positive correlations (in the form of the FKG inequality), and that implies that, if $n < N$, the restriction of the wired (resp. free) measure on Λ_N to Λ_n is stochastically smaller (resp. larger) than the corresponding measure defined on Λ_n directly. As n goes to infinity, this allows for the definition of infinite-volume measures as monotonic limits of both sequences, which we will denote by $P^{w}_{p,q}$ and $P^{f}_{p,q}$.

For fixed q and either free of wired boundary conditions, these two measure families are stochastically ordered in p; this implies the existence of a critical point $p_c(q)$ (the same in both cases, as it turns out) such that, as in the case of Bernoulli percolation, there is a.s. no infinite cluster (resp. a unique infinite cluster) if $p < p_c$ (resp. $p > p_c$).

It is conjectured that for every $q \geqslant 1$,

$$p_c(q) = \frac{\sqrt{q}}{1+\sqrt{q}}.$$

This come from the following duality construction. Let for now G be a graph embedded into the 2-sphere, and let G^* be its dual graph. To any configuration ω of the random-cluster model on G, one can associate a configuration ω^* on G^* by declaring a dual bond to be open if and only if the corresponding primal bond is closed. As it turns out, if ω is distributed as $P_{p,q,G}$, then ω^* is distributed as P_{p^*,q^*,G^*}, *i.e.* it is a random-cluster model configuration, with

$$q^* = q \quad \text{and} \quad \frac{p^*}{1-p^*} = q\frac{1-p}{p}.$$

It is easy to see that there is a unique value p_{sd} of p satisfying $p_{sd} = (p_{sd})^*$, and it is then natural to hope that $p_c = p_{sd}$, leading to the value above.

Exercise I.1: *Prove the duality statement.*

— *The parafermionic observable* —

As before, we want to introduce an observable defined from the discrete model and giving us enough information to determine asymptotic properties in the scaling limit. From now on, let Ω be a smooth, simply connected domain in the complex plane, with two marked points a and b on its boundary. Let G_δ be the graph obtained as the intersection of Ω with a scaled copy of $\mathbb{Z}^2$ (with mesh $\delta > 0$); we will consider the random-cluster model with parameters $q \geqslant 1$ and $p = p_{sd}(q)$ on G_δ, with Dobrushin boundary conditions, wired on the (positively oriented) arc ab — and we will denote the corresponding measure simply by P (or later by P_δ when we insist on the scaling behavior as $\delta \to 0$).

As in the case of percolation, we briefly describe the main steps in Smirnov's proof of conformal invariance. A big difference is that the statement of convergence needs more notation, so we will have to postpone it a little bit.

Step 1: The loop representation. In addition to the graph G_δ and its dual G_δ^*, we need a third one known as the *medial graph* and denoted by $G_\delta^\diamond$; it is defined as follows. The vertices of the medial graph are in bijection with the bonds of either G_δ of G_δ^* (which is the same), and one can think of them as being at the intersection of each primal bond with its dual; there is an edges between two vertices of $G_\delta^\diamond$ if and only if the two corresponding primal edges share an endpoint and the two corresponding dual edges do as well.

One can encode a random-cluster configuration on G_δ using the medial graph, by following the boundary of each of its clusters (or equivalently, each of the clusters of the dual configuration). This leads to a covering of all the bonds of $G_\delta^\diamond$ by a family of edge-disjoint paths, one joining a to b (which we will call the *interface* and denote by γ), and the others being loops. If $l(\omega)$ denotes the number of loops obtained this way, it is possible to rewrite the probability of a configuration ω as

$$P[\{\omega\}] = \frac{x^{o(\omega)}(\sqrt{q})^{l(\omega)}}{Z_{x,q}} \quad \text{with} \quad x := \frac{p}{(1-p)\sqrt{q}}.$$

Since we work at the self-dual point, in fact we have $x = 1$ and the weight of a configuration is written as a function of only the number of loops in its loop representation.

Exercise I.2: *Prove the equivalence of the FK representation and the loop representation.*

Step 2: Definition of the parafermionic observable. Let e be an edge of the medial graph. We would like to be able to compute the probability that the interface passes through e; unfortunately, this seems to be out of the reach of current methods, so we need an alternative. Assume that γ does go through e. Then, it is possible to follow it from a to the mid-point of e, and to follow the variation of the angle of the tangent vector along the way: it increases (resp. decreases) by $\pi/2$ whenever γ turns left (resp. right). The *winding* of the curve at e is the value one gets when reaching e; it is in $(\pi/2)\mathbb{Z}$ and we denote it by $W(e)$. If γ does not pass through e, define $W(e)$ to be an arbitrary value, as it will not be relevant.

The parafermionic observable is then defined as

$$F_\delta(e) := E\left[e^{-i\sigma W(e)}\mathbb{1}_{e\in\gamma}\right] \quad \text{where } \sigma \text{ satisfies} \quad \cos\frac{\sigma\pi}{2} = \frac{\sqrt{q}}{2}.$$

Notice the difference here between the cases $q \leq 4$ (when σ is real) and $q > 4$ (when it is pure imaginary); we will come back to this distinction shortly. The parameter σ is known as the *spin* of the model. Morally, the main convergence result is that, at the self-dual point, $\delta^{-\sigma}F_\delta$ converges (to an explicit limit) as $\delta \to 0$; but giving a precise sense to that statement requires a little more preparation.

Step 3: Local behavior of the observable. Let ω be a configuration, and let e again be an edge of the medial lattice. We will denote by $F_\delta(e,\omega)$ the contribution of ω to the observable, so that $F_\delta(e) = \sum F_\delta(e,\omega)$. Besides, let ℓ be a bond of the primal lattice which is incident to e. There is a natural involution on the set of configurations given by changing the state of the bond ℓ without changing anything else — denote this involution by s_ℓ. It is easy to see how $F_\delta(e,\omega)$ and $F_\delta(e,s_\ell(\omega))$ differ: If both are non-zero, then the winding term is the same and their ratio is therefore either $x\sqrt{q}$ or $x/\sqrt{q}$ according to whether opening ℓ creates or destroys a loop.

Notice that the medial lattice can be oriented in a natural way, deciding that its faces corresponding to dual vertices are oriented positively. The definitions above ensure that γ always follows the orientation of the bonds it uses. We will use the notation $e \to \ell$ (resp. $\ell \to e$) to mean that the bond e is oriented towards (resp. away from) its intersection with ℓ. The above observations imply that, for every ℓ in the primal lattice and every $x > 0$,

$$\sum_{e\to\ell}\left[F_\delta(e,\omega) + F_\delta(e,s_\ell(\omega))\right] = \frac{e^{\sigma\pi i/2}+x}{1+xe^{\sigma\pi i/2}}\sum_{e\to\ell}\left[F_\delta(e,\omega) + F_\delta(e,s_\ell(\omega))\right].$$

The prefactor in the right-hand side is a complex number of modulus 1 if $q \leq 4$; it is real and positive if $q > 4$; and in both cases, it is equal to 1 if and only if $x = 1$ — in other words, exactly at the self-dual point. Summing the previous relation over ω, we get

$$\sum_{e\to\ell}F_\delta(e) = \frac{e^{\sigma\pi i/2}+x}{1+xe^{\sigma\pi i/2}}\sum_{e\to\ell}F_\delta(e)$$

and in particular, at the self-dual point, this relation boils down to the flow condition:

$$\sum_{e\to\ell}F_\delta(e) = \sum_{e\to\ell}F_\delta(e).$$

This is the base from which all the rest of the proof builds up: we now need to interpret this relation as the vanishing of a divergence, and in turn as the (discrete) holomorphicity of a well-chosen function.

— *Interlude* —

Here we have to stop for a moment. All the preceding reasoning is perfectly general, and the intuition behind it is rather clear. However, a worrying remark is that we get one linear relation per bond of the primal lattice, but one unknown (the value of F_δ) per bond of the *medial* lattice, of which there are twice as many. That means that these relations cannot possibly characterize F_δ uniquely, and indicates that something more is needed; and in fact, the proof of conformal invariance is indeed not known in all generality.

The first "easy" case is that of $q > 4$. Here the observable is real-valued, and therefore lends itself to more analytic techniques, mostly inequalities. This is quite fruitful if one aims for the value of the critical point, and indeed in that case one can show that $p_c = p_{sd}$ using only elementary calculus. The use of inequalities is not optimal though, as it is not precise enough to derive any information at the critical point — much less to prove conformal invariance.

The second "easy" case is that of $q = 2$, for which $\sigma = 1/2$. Recall that γ always traverses edges in their positive direction; that means that $W(e)$ is known in advance up to a multiple of 2π. In turn, this means that the argument of $e^{i\sigma W(e)}$ is known up to an integer multiple of π; in other words, for every (oriented) edge e, the observable $F_\delta(e)$ takes its value on a line in the complex plane depending only on the direction of e (essentially, seeing e as a complex number, $F_\delta(e)/\sqrt{e}$ is a real number). The previous obstruction can then be bypassed: while we get a (complex) linear relation per primal bond, we only need to determine a *real* unknown per medial bond, which is the same quantity of information. That is morally why a complete derivation of conformal invariance is known only in that particular case, to which we restrict ourselves from now on.

— *The rest of the proof when* $q = 2$ *and* $x = 1$ —

Step 4: The integrated observable. We would like to prove that $\delta^{-1/2} F_\delta$ converges and be done with it. However, this is not very realistic, in particular because the argument of F_δ oscillates wildly (with a period of one lattice mesh). A solution to that is to consider an integrated version of it, or rather of its square. There is (up to an additive constant) a unique function H_δ, defined on the vertices of both the primal and dual lattices, satisfying the following condition: Whenever W (resp. B) is a vertex of the primal (resp. dual) lattice, and they are adjacent and separated by the bond e_{WB} of the primal lattice, then

$$H_\delta(B) - H_\delta(W) = |F_\delta(e_{WB})|^2 .$$

In some sense, one can think of H_δ as a discrete integral of $|F_\delta|^2$; proving its existence is a matter of checking that the sum of the prescribed increments around a vertex of the medial lattice vanishes; and this in turn is a direct consequence (via Pythagoras' theorem) of the flow condition, noticing that the values of $F_\delta(e)$ for $e \to \ell$ (resp. $\ell \to e$) are always orthogonal. Besides, it is easy to check that H_δ is constant along both boundary arcs of the domain, and that its discontinuity at a (and hence also at b) is exactly equal to 1 — because the interface *has* to pass through a with no winding. From now on, we will thus assume that H_δ is equal to 0 (resp. 1) on the arc ab (resp. ba).

Now, H_δ has two natural restrictions, H_δ^w to the primal vertices and H_δ^b to the dual ones. These two restrictions have nice properties: H_δ^w is super-harmonic (its discrete Laplacian is non-positive) wile H_δ^b is sub-harmonic (its Laplacian is non-negative). Besides, assuming that F_δ is small, they differ very little, so that any sub-sequential scaling limit of H_δ has to be harmonic — in fact, has to be the unique harmonic function h with boundary values 0 on ab and 1 on ba. That is already a non-empty statement, but extracting useful information from H_δ is not easy …

Exercise I.3: *Prove that H_δ^w is indeed super-harmonic.*

Step 5: Tying up the loose ends. The above arguments are rather convincing, but a lot is missing which would not fit comfortably in these notes. Following the scheme of the percolation argument, the main ingredient is relative compactness in the shape of uniform continuity; here, it follows from classical results about the Ising model (essentially, from the fact that the phase transition is of second order, or equivalently that the magnetization of the critical 2D Ising model vanishes) and it does imply the convergence of H_δ, as $\delta \to 0$, to h as defined above.

Then, one needs to come back down from H_δ to F_δ, which involves taking a derivative (the same way H_δ was obtained by integration). Notice in passing that $|F_\delta|^2$ is an increment of H_δ, so it is expected to be of the same order as the lattice spacing, and thus F_δ itself should be — and, indeed, is — of order $\sqrt{\delta}$, and related to the square root of the gradient of H_δ. The objection we raised before about the argument of F_δ still stands; we need to make the following adjustments.

Each vertex of the medial graph has four *corners* (one per adjacent face); by "rounding up" γ at each of its turns, it is possible to naturally extend F_δ to all such corners, with a winding at a corner defined to be the midpoint between that on the incident edge and that of the exiting one. Then, define $F_\delta^\diamond$ on each medial vertex to be the sum of F_δ over all adjacent corners. $F_\delta^\diamond$ is now a *bona fide* complex-valued function, and in fact F_δ can be recovered from it by appropriate projections.

We are now armed to state the main convergence result. Let Φ be a conformal map from Ω to the horizontal strip $\mathbb{R} \times (0,1)$, mapping a to $-\infty$ and b to $+\infty$. Such a map is unique up to a horizontal translation, which will not matter here since we will look at derivatives anyway; notice that h is simply the imaginary part of Φ.

Theorem 2 (Smirnov). *As $\delta \to 0$, and uniformly on compact subsets of Ω,*

$$\delta^{-1/2} F_\delta^\diamond \to \sqrt{2\Phi'}.$$

In particular, the scaling limit of $F_\delta^\diamond$ is conformally invariant.

2 PATH MODELS

Maybe the simplest model for which conformal invariance is well understood is that of the *simple random walk* on a periodic lattice, say $\mathbb{Z}^2$. Indeed, as the mesh of the lattice goes to 0, the random walk path converges in distribution to that of a Brownian motion, and this in turn is conformally invariant.

More precisely, let Ω be a (bounded, smooth, simply connected) domain of the complex plane, and let $z \in \Omega$; let (B_t) be a standard planar Brownian motion started from z, τ be its hitting time of the boundary of Ω. Besides, let Φ be a conformal map from Ω onto a simply connected domain Ω', and let (W_s) be a Brownian motion started from $\Phi(z)$ and σ its hitting time of $\partial\Omega'$.

It is not true that (W_s) and $(\Phi(B_t))$ have the same distribution in general, because their time parameterizations will be different, but in terms of the path considered as a subset of the plane, they do; the following statement is another instance of conformal invariance:

Theorem 3 (P. Lévy). *The random compact sets $\{\Phi(B_t) : t \in [0,\tau]\}$ and $\{W_s : s \in [0,\sigma]\}$ have the same distribution.*

A consequence of this and the study of SLE processes will be (among others) a very detailed description of the *Brownian frontier*, *i.e.* of the boundary of the connected component of infinity in the complement of $B_{[0,\tau]}$. However, the frontier is not visited by the Brownian path in chronological order, and that makes direct the use of planar Brownian motion problematic; it seems that things would be simpler if the random curve had no double point, and correspondingly if the underlying discrete path were self-avoiding.

The most natural way to generate a self-avoiding path in a discretized simply connected domain, from an inside point z to the boundary, would be to notice that there are finitely many such paths and to define a probability measure on the set of paths (morally, uniform given the length of the path). This leads to the definition of the *self-avoiding walk*, but unfortunately not much is known about its scaling limit, so we turn our attention to a different object which is a bit more difficult to define but much easier to study.

2.1 LOOP-ERASED RANDOM WALK

Let again Ω be a simply connected domain in the plane, and let $\delta > 0$; let Ω_δ be an appropriate discretization of Ω by $\delta\mathbb{Z}^2$ (say, the largest component of their intersection), and let z_δ be a vertex in Ω_δ. In addition, let (X_n) be a discrete-time random walk on $\delta\mathbb{Z}^2$, starting from z_δ, and let

$$\tau := \inf\{n : X_n \notin \Omega_\delta\}$$

its exit time from Ω_δ. The loop-erasure $LE(X)$ is defined, as the name indicates, by removing the loops from (X_n) as they are created. Formally, define the (n_i) inductively by letting $n_0 = 0$ and, as long as $n_i < \tau$,

$$n_{i+1} := \max\{n \leqslant \tau : X_n = X_{n_i}\} + 1.$$

Then, $LE(X)_i := X_{n_i}$.

Clearly, the loop-erasure of a discrete path is a self-avoiding path, as the same vertex cannot appear twice in $LE(X)$; when as above X is a simple random walk, $LE(X)$ is known as the *loop-erased random walk* (from z_δ to $\partial\Omega_\delta$ in Ω_δ). If b is a boundary point of Ω_δ, one can condition X to leave Ω_δ at b and the loop-erasure of that conditioned random walk is called the loop-erased random walk from a to b in Ω_δ.

The profound link between the loop-erased random walk and the simple random walk itself will be instrumental in the study of its asymptotic properties as δ goes to 0. For instance, the distribution of the exit point is the same for both (it follows the discrete harmonic measure from z_δ).

The counterpart of RSW for loop-erased walks (in the sense that it is one of the basic building blocks in proofs of convergence) will be a statement that $LE(X)$ does not "almost close a loop" — so that in particular, if it does have a scaling limit, the limit will be supported on simple curves. We defer the exact statement to a later section, but essentially what happens is the following: for $LE(X)$ to form a fjord, without closing it, X itself needs to approach its past path and then proceed to the boundary of the domain without actually closing the loop (as this would vanish in $LE(X)$); the escape itself is very unlikely to happen as a consequence of Beurling's estimate.

2.2 UNIFORM SPANNING TREES AND WILSON'S ALGORITHM

Let $G = (V, E)$ be a finite graph; let v_∂ be a vertex of G (the "boundary" of the graph). A *spanning tree* of G is a connected sub-graph of G containing all its vertices and no loop (a sub-graph with all the vertices and no loop is called a *spanning forest*, and a tree is a connected forest). The set of spanning trees of G is finite; a *uniform spanning tree* is a random tree with the uniform distribution on that set.

Given a vertex $v \neq v_\partial$, we now have two ways of constructing a random self-avoiding path from v to v_∂:

- The loop-erased random walk in G from v to v_∂ (defined exactly as in the case of the square lattice above);
- The (unique) branch of a uniform spanning tree joining v to v_∂.

As it turns out, these two random paths have the same distribution. In particular, because in the second definition the roles of v and v_∂ are symmetric, we get an extremely non-obvious feature of loop-erased random walks: the time-reversal of the loop-erased walk from v to v_∂ is exactly the lop-erased walk from v_∂ to v. This is instrumental in the proof of convergence of the loop-erased walk to SLE_2 in the scaling limit.

As an aside, loop-erased walks provide a very efficient method for sampling a uniform spanning tree, which is due to David Wilson. Essentially: pick a point v_1, and run a loop-erased walk γ_1 from it to v_∂; then, pick a vertex v_2 which is not on γ_1 (if there is such a vertex) and run a loop-erased walk γ_2 from v_2 to γ_1; proceed until all the vertices of V are exhausted, each time building a loop-erased walk from a vertex to the union of all the previous walks. When the construction stops, one is facing a random spanning tree of G; and as it happens, the distribution of this tree is that of a uniform spanning tree.

3 BIBLIOGRAPHICAL NOTES

Section 1.1, Percolation. A very complete review of percolation theory is Grimmett's book [7]; it contains everything mentioned in these notes except for Cardy's formula. Its bibliography section is far more complete than I could hope to gather here, so I will just list a few key papers. An alternative, which is a bit hard to find but well worth reading, is the book of Kesten [10]. For more recent progress and conformal invariance (and more exercises), one can *e.g.* consult the lecture notes for Werner's lectures [20].

Besides the anecdotal quotation from [22], the first proper introduction of percolation as a mathematical model is the article of Broadbent and Hammersley [5]. Exponential decay (up to the critical point) was derived in a very general setting by Menshikov [13]. The first derivation of the value of a critical parameter was obtained (for bond-percolation on the square lattice) by Kesten [9]; RSW estimates were obtained independently by Russo [15] and by Seymour and Welsh [17].

Cardy's formula was first conjectured by — well, Cardy [6], and then proved on the triangular lattice by Smirnov [18]. A slightly simplified exposition of the proof (which is the one we followed here) can be found in [1], and a very (very!) detailed one in the book of Bollobás and Riordan [4].

Section 1.2, The random-cluster model. Here again, the reader is advised to refer to the (other) book of Grimmett [8] (and references therein) for a general introduction to random-cluster models, including most of the results which are mentioned in this section. The proof of conformal invariance for the $q = 2$ critical FK model was first obtained by Smirnov [19]; the approach we follow here is very close to the original, but some notation is borrowed from [2] (and technically, the notation $F_\delta^\diamond$ is only used here).

The equality $p_c = p_{sd}$ is related to the so-called Kramers-Wannier duality [11]; while still open in the general case, it is known to hold in the case $q = 1$ (where it is exactly Kesten's result on the percolation critical point in [9]); in the case $q = 2$ (where it is related to the derivation of the critical temperature of the two-dimensional Ising model by Onsager [14] — see also [2]); and in the case $q > 4$ as proved in [3] using Smirnov's observable.

Section 2.1, Loop-erased random walk. For the contents of this section, and an introduction (possibly the best introduction), one can have a look at Schramm's original paper on LERWs and USTs [16]. Wilson's article [21] complements it nicely; and for more quantitative results, parts of the paper by Lawler, Schramm and Werner [12] can be read without any prior knowledge of SLE.

Chapter II

Schramm-Loewner Evolution

CONTENTS

INTRODUCTION

The previous chapter introduced a few discrete models, and for each of them we saw that, in the scaling limit as the lattice mesh goes to zero, a particular *observable* converges to a conformally invariant limit. It is natural to hope that convergence will actually occur in a much stronger sense, and in particular that the *interfaces* of the discrete model will have a continuous counterpart described as random curves in a planar domain.

Schramm's insight was to realize that, under mild (and reasonable) assumptions in addition to conformal invariance, the limit has to be distributed as one of a one-parameter family of measures on curves, which he named Stochastic Loewner Evolutions and are now universally known as *Schramm-Loewner Evolutions*. The aim of this chapter is to define these random curves and give a few of their fundamental properties.

1 CHORDAL SLE

1.1 LOEWNER EVOLUTION IN THE HALF-PLANE

Let $\mathbb{H}$ denote the open upper half-plane, seen as a subset of the complex plane, and (for now) let $\gamma : [0, \infty) \to \bar{H}$ be a continuous, simple curve. To keep with probabilistic tradition, we will denote the position of γ at time t by γ_t instead of $\gamma(t)$; besides, we will assume that γ satisfies the following conditions:

- $\gamma_0 = 0$;
- For every $t > 0$, $\gamma_t \in \mathbb{H}$ (or in other words, $\gamma_t \notin \mathbb{R}$);
- $|\gamma_t| \to \infty$ as $t \to \infty$.

The results we will state in this section are actually valid in much more generality, but the intuition is not fundamentally different in the general case.

Let $H_t := \mathbb{H} \backslash \gamma_{[0,t]}$ be the complement of the path up to time t. Our assumptions ensure that H_t is a simply connected domain, and therefore Riemann's mapping theorem can be applied to show that there exists a conformal map

$$g_t : H_t \to \mathbb{H}$$

(we refer the reader to Section 2.1 for a refresher on complex analysis, if needed). g_t is uniquely determined if one imposes the *hydrodynamic normalization*, which amounts to fixing the following asymptotic behavior at infinity:

$$g_t(z) = z + \frac{a(t)}{z} + \mathcal{O}\left(\frac{1}{z^2}\right).$$

With this notation, it is not hard to prove that a is a strictly increasing, continuous function; it need not go to infinity with t, but we will add this as an assumption on the curve. a can therefore be used to define a "natural" time parametrization of the curve: Up to reparameterization, it is always possible to ensure that $a(t) = 2t$ for all $t > 0$. From now on, we shall assume that γ is indeed parametrized that way.

Exercise II.1: *Prove the statements made so far in the section, and in particular prove that a is indeed continuous and strictly increasing. Give an example of a curve going to infinity, but for which a is bounded.*

The normalizations of g and t are chosen in such a way that the behavior of $g_t(z)$ as a function of t is then easy to describe:

Theorem 4 (Loewner). *There exists a continuous function $\beta : [0, \infty) \to \mathbb{R}$ such that, for every $t \geqslant 0$ and every $z \in H_t$,*

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

This differential equation is known as *Loewner's equation (in the half-plane)*. The gain is substantial: We were able to encode the whole geometry of γ, up to reparameterization, in terms of a single real-valued function. Indeed, it is not difficult to show that the construction up to now is essentially reversible: Given β, one can solve Loewner's equation to recover (g_t), and hence (H_t) and (γ_t) as well[1].

1.2 DEFINITION OF SLE

Consider, say, critical site-percolation on the triangular lattice in the upper half-plane, with boundary conditions open to the right of the origin and close to the left — this corresponds to the case in Section 1.1 with one of the boundary points at infinity. This creates an interface starting from the origin, which is the path of the exploration process and satisfies the previous hypotheses on the curve γ.

Now, assume that, as the lattice mesh goes to 0, the exploration curve converges in distribution to a (still random) curve in the upper half-plane. This scaling limit can then be encoded into a real-valued process β using Loewner's equation; of course, β will be random as well. The question is now whether we can use the results of the previous chapter to identify β.

[1]Well, some care is needed here: It is *not* true that one can plug any function β into Loewner's equation and obtain a Jordan curve γ out of it. It is true if β is Hölder with exponent $1/2$ and small enough norm, but a sharp condition is not known. Obviously everything works out fine if β comes from the above construction in the first place!

Let $R > 0$, and stop the exploration process at the first time τ_R when it reaches the circle of radius R centered at 0. Conditionally on its path so far, the next steps are exactly the exploration process of percolation in a new domain H_{τ_R}, namely the unbounded connected component of the complement of the current path: This is known as the *domain Markov property*, and it parallels the DLR conditions of statistical mechanics.

Cardy's formula being conformally invariant, it is natural to expect the scaling limit of the exploration process will be as well, or in other words, that the path after time τ_R is distributed as the conformal image of the path in $\mathbb{H}$ by a map sending $\mathbb{H}$ to the appropriate domain.

We almost have such a map at our disposal, from Loewner's equation: The map g_{τ_R} sends H_{τ_R} to $\mathbb{H}$, so its inverse map looks like what we are looking for. The only difference is that $g_{\tau_R}(\gamma_{\tau_R})$ is equal to β_{τ_R} instead of 0. Taking this into account, we get the following property (assuming of course the existence of the scaling limit): **The image of $(\gamma_t)_{t\geqslant\tau_R}$ by $g_{\tau_R} - \beta_{\tau_R}$ has the same distribution as $(\gamma_t)_{t\geqslant 0}$.**

Besides, all information coming from the path up to time τ_R is forgotten in this map: Only the shape of H_{τ_R} is relevant (because such is the case at the discrete level), and the dependence of that shape vanishes by conformal invariance. **The image of $(\gamma_t)_{t\geqslant\tau_R}$ by $g_{\tau_R} - \beta_{\tau_R}$ is independent of $(\gamma_t)_{0\leqslant t\leqslant\tau_R}$.**

It remains to investigate what these two properties translate to in terms of the process (β_t). First, notice that the coefficient in $1/z$ in the asymptotic expansion at infinity which we are using is additive under composition; and let $s, t > 0$. The previous reasoning, applied at time t, then leads to the following:

$$g_{t+s} = \beta_t + \tilde{g}_s \circ (g_t - \beta_t),$$

where equality holds in distribution and where $\tilde{g}$ is an independent copy of g_s; the addition of β_t takes care of the normalization at infinity.

Differentiating in s and using Loewner's equation, this leads to

$$\beta_{t+s} = \beta_t + \tilde{\beta}_s,$$

where again equality holds in distribution and $\tilde{\beta}$ is an independent copy of β. **The process $(\beta_t)_{t\geqslant 0}$ has independent and stationary increments.**

Besides, the distribution of γ is certainly invariant under vertical reflexion (because this holds at the discrete level), so β and $-\beta$ have the same distribution. So, we arrive at the following characterization: **Under the hypotheses of conformal invariance and domain Markov property, there exist a constant $\kappa \geqslant 0$ and a standard Brownian motion (B_t) such that**

$$(\beta_t)_{t\geqslant 0} = (\sqrt{\kappa} B_t)_{t\geqslant 0}.$$

2 THE RADIAL CASE

We just say a few words here about the case of radial Loewner chains, since not much needs to be changed from the chordal setup. Here, we are given a continuous, Jordan curve γ in the unit disk $\mathbb{D}$, satisfying $\gamma_0 = 1$, $\gamma_t \neq 0$ for all $t > 0$ and $\gamma_t \to 0$ as $t \to 0$. In other words, the reference domain is not the upper-half plane with two marked boundary points, but the unit disk with one marked boundary point and one marked inside point.

Let D_t be the complement of $\gamma_{[0,t]}$ in the unit disk; notice that 0 is in the interior of D_t, so there exists a conformal map g_t from D_t onto $\mathbb{D}$ fixing 0; this map is unique if one requires in addition that $g_t'(0) \in \mathbb{R}_+$, which we will do from now on.

The natural parameterization of the curve is still whatever is additive under composition of conformal maps; here, the only choice (up to a multiplicative constant) is the logarithm of $g_t'(0)$: up to reparameterization, we can ensure that for every $t > 0$, $g_t'(0) = e^t$. With this normalization, we have the following:

Theorem 5 (Loewner). *There exists a continuous function $\theta : [0, \infty) \to \mathbb{R}$ such that, for every $t \geqslant 0$ and every $z \in D_t$,*

$$\partial_t g_t(z) = \frac{g_t(z) + e^{i\theta_t}}{g_t(z) - e^{i\theta_t}} g_t(z).$$

This is known as Loewner's equation in the disk.

Everything we just saw in the chordal case extends to the radial case. In particular, if the curve is related to a conformally invariant model (say, if it is the scaling limit of the loop-erased random walk), then under the same hypothesis of domain Markov property, one gets that there must exist $\kappa > 0$ such that

$$(\theta_t)_{t \geqslant 0} = (\sqrt{\kappa} B_t)_{t \geqslant 0}$$

(where again (B_t) is a standard real-valued Brownian motion). Solving Loewner's equation in the disk with such a driving function defines *radial* SLE_κ.

Remark. The local behavior of this equation around the singularity at $z = e^{i\theta_t}$ involves a numerator of norm 2; it is the same 2 as in Loewner's equation in the upper half-plane, in the sense that the local behavior of the solution for the same value of κ will then be the same on both sides.

Chapter III

Convergence to SLE

Understanding the connections between grid-based models and continuous processes is a project of fundamental importance, and so far has only limited success. As mathematicians, we should not content ourselves with the vague notion that the discrete and continuous models behave "essentially the same", but strive to make the relations concrete and precise.

Oded Schramm

Appendix A

Mathematical toolbox

1 PROBABILISTIC TOOLS

1.1 CORRELATION INEQUALITIES: FKG, BK . . .

1.2 STOCHASTIC CALCULUS

2 ANALYTIC TOOLS

2.1 COMPLEX ANALYSIS

Appendix B

Answers to the exercises

Exercise I.1: Use Euler's formula to relate the number of open and closed bonds in the primal and the dual configurations with their numbers of faces and clusters. It helps to rewrite the weight of a configuration as

$$\left(\frac{p}{1-p}\right)^{o(\omega)} q^{k(\omega)}$$

and to notice that $o(\omega) + o(\omega^*)$ does not depend on the configuration.

Exercise I.2: It works exactly the same way as the previous exercise, use Euler's formula in the natural way and it will work.

Exercise I.3: That actually takes some doing (the last, very tedious 4 pages of Smirnov's article [19]), but it is completely elementary. Simply expand the discrete Laplacian in terms of F_δ, and use the flow relation repeatedly to eliminate terms (it allows to express each value of F_δ in terms of its values at 3 neighboring edges, but projecting on lines brings this down to 2).

Exercise II.1: It is enough to show that $a(t)$ is strictly positive for every $t > 0$ — look at what happens under composition. One can then use Schwarz' Lemma to conclude. a will remain bounded if γ remains close enough to the real line.

Bibliography

[1] V. BEFFARA, *Cardy's formula on the triangular lattice, the easy way*, in Universality and Renormalization, I. Binder and D. Kreimer, eds., vol. 50 of Fields Institute Communications, The Fields Institute, 2007, pp. 39–45.

[2] V. BEFFARA AND H. DUMINIL-COPIN, *The critical temperature of the Ising model on the square lattice, the easy way*, In preparation, (2010).

[3] V. BEFFARA, H. DUMINIL-COPIN, AND S. SMIRNOV, *The self-dual point of the two-dimensional random-cluster model is critical above* $q = 4$, In preparation, (2010).

[4] B. BOLLOBÁS AND O. RIORDAN, *Percolation*, Cambridge University Press, 2006.

[5] S. R. BROADBENT AND J. M. HAMMERSLEY, *Percolation processes, I and II*, Proc. Cambridge Philos. Soc., 53 (1957), pp. 629–645.

[6] J. CARDY, *Critical percolation in finite geometries*, J. Phys. A, 25 (1992), pp. L201–L206.

[7] G. R. GRIMMETT, *Percolation*, vol. 321 of Grundlehren Math. Wiss., Springer-Verlag, Berlin, second ed., 1999.

[8] ——, *The random-cluster Model*, vol. 333 of Grundlehren Math. Wiss., Springer-Verlag, Berlin, 2006.

[9] H. KESTEN, *The critical probability of bond percolation on the square lattice equals* $1/2$, Communications in Mathematical Physics, 74 (1980), pp. 41–59.

[10] ——, *Percolation theory for mathematicians*, vol. 2 of Progress in Probability and Statistics, Birkhäuser, Boston, Mass., 1982.

[11] H. A. KRAMERS AND G. H. WANNIER, *Statistics of the two-dimensional ferromagnet I*, Phys. Rev., 60 (1941), pp. 252–262.

[12] G. F. LAWLER, O. SCHRAMM, AND W. WERNER, *Conformal invariance of planar loop-erased random walks and uniform spanning trees*, Ann. Probab., 32 (2004), pp. 939–995.

[13] M. V. MENSHIKOV, *Coincidence of critical points in percolation problems*, Soviet Mathematics Doklady, 33 (1986), pp. 856–859.

[14] L. ONSAGER, *Crystal statistics. I. A two-dimensional model with an order-disorder transition*, Phys. Rev. (2), 65 (1944), pp. 117–149.

[15] L. RUSSO, *A note on percolation*, Z. Wahrscheinlichkeitstheorie und Verw. Gebiete, 43 (1978), pp. 39–48.

[16] O. SCHRAMM, *Scaling limits of loop-erased random walks and uniform spanning trees*, Israel Journal of Mathematics, 118 (2000), pp. 221–288.

[17] P. D. SEYMOUR AND D. J. A. WELSH, *Percolation probabilities on the square lattice*, Ann. Discrete Math., 3 (1978), pp. 227–245. Advances in graph theory (Cambridge Combinatorial Conf., Trinity College, Cambridge, 1977).

[18] S. SMIRNOV, *Critical percolation in the plane: Conformal invariance, Cardy's formula, scaling limits*, Comptes Rendues de l'Académie des Sciences, Paris, Série I : Mathématiques, 333 (2001), pp. 239–244.

[19] ——, *Conformal invariance in random cluster models. i. holomorphic fermions in the ising model*, Annals of Mathematics, (2007). To appear, arXiv:0708.0039.

[20] W. WERNER, *Lectures on two-dimensional critical percolation*, Institute for Advanced Studies, Lectures given at the Park City summer school, 2007. arXiv:0710.0856.

[21] D. B. WILSON, *Generating random spanning trees more quickly than the cover time*, in Proceedings of the Twenty-eighth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996), New York, 1996, ACM, pp. 296–303.

[22] D. V. WOOD, *Average and probability, problem 5*, The American Mathematical Monthly, 1 (1894), pp. 211–212.

Index