Lectures on Stochastic Loewner Evolution and Other Growth Processes in Two Dimensions*

John Cardy
Rudolf Peierls Centre for Theoretical Physics
1 Keble Road, Oxford OX1 3NP, U.K.†
and All Souls College, Oxford.

March 2006

Abstract

Random objects such as clusters in the plane can often be described in terms of the conformal mappings which take their boundaries into some standard shape. As the clusters grow, the mapping function changes in a well-defined manner, which is often easier to understand than the original problem. One of the simplest examples is Stochastic (Schramm) Loewner Evolution (SLE), which turns out to describe random curves in equilibrium statistical mechanics models. These lectures give an introduction to the use of such conformal mappings, and to SLE in particular, from the physicist's point of view.

*Lectures given at the school on *Non-Equilibrium Dynamics of Interacting Particle Systems*, March/April 2006, Isaac Newton Institute, Cambridge.

†Address for correspondence.

Contents

1 Introduction

These lectures sit rather uncomfortably in a workshop entitled 'Non-equilibrium Dynamics of Interacting Particle Systems'. They are going to be about dynamics, in the sense that we'll be discussing growth processes. But we are going to be much more interested in the *end result* of the process rather than the intermediate time dependence. Indeed, the notion of 'time' which we will end up using will turn out to be quite unrelated to the real time of a real or computer realisation of the growth process. And it will also turn out that in the simplest example of such a growth process, called SLE, the case when we can really calculate things analytically, the final pattern produced by the growth process corresponds to an *equilibrium* statistical mechanics problem.

The method we shall use involves describing the growing pattern not directly, but in terms of the evolving conformal mapping which sends the region outside to some standard region, like the upper half plane. It will turn out that the evolution of such mappings is determined, through the Loewner equation, in terms of a stochastic process on the real line, in one dimension. In many cases this can be realised by an 'Interacting Particle System'. But the simplest example, that of SLE, in fact corresponds to a *single* particle executing simple 1d Brownian motion, the most trivial random particle dynamics of all. Nevertheless, through the Loewner equation, it corresponds to something highly non-trivial happening in 2d, and that is what the main part of these lectures will be about.

I am very grateful to Wouter Kager for allowing me to use Figs. 1-3 which are from his thesis.

1.1 Bibliography

This is a very recent ($\geq$2000) subject, although the problems of 2d equilibrium lattice models and growth processes have been around for much longer. The original papers by G. Lawler, O. Schramm and W. Werner, are for pure mathematicians and are very difficult. Fortunately there are now several reviews written with physicists in mind:

- *2d growth processes: SLE and Loewner chains* by M. Bauer and D. Bernard (math-ph/0602049, to appear in Physics Reports, 172 pages) is the most complete and the approach is essentially along the lines of the present lectures, although I won't mention conformal field theory;
- *SLE for theoretical physicists* by J. Cardy (cond-mat/0503313, Annals Phys. 318 (2005) 81-118, 43 pages) is shorter and contains most of the ideas with, however, less introductory material;
- *A guide to stochastic Loewner evolution and its applications* by W. Kager and B. Nienhuis (math-ph/0312251, J. Stat. Phys. 115, 1149, 2004) gives a useful account of the mathematics for physicists;
- finally, there are mathematical reviews in *Random planar curves and Schramm-Loewner evolutions,* by W. Werner, (*Ecole d'Eté de Probabilités de Saint-Flour XXXII (2002)*, Springer Lecture Notes in Mathematics **1180**, 113, 2004 (math.PR/0303354)); *Conformally invariant processes in the plane*, by G. Lawler (AMS, 2005, ISBN: 0-8218-3677-3); *Stochastic Loewner Evolution,*, by G. Lawler (to appear in *Encyclopedia of Mathematical Physics*, J.-P. Françoise, G. Naber and T.S. Tsun, eds. (Elsevier, 2005.) (http://www.math.cornell.edu/~lawler/encyclopedia.ps).

2 Discrete Growth Processes in 2d

In this section we are going to describe the kind of growth processes on the lattice, whose continuum limit will be described by SLE. The simplest is

2.1 Percolation

For simplicity we consider a honeycomb lattice in some simply connected region of the plane. Fig. 1 shows a rectangular region. We describe a particular kind of kinetic self-avoiding walk on this lattice. It starts at some point on the boundary, which in the figure is chosen to be the top

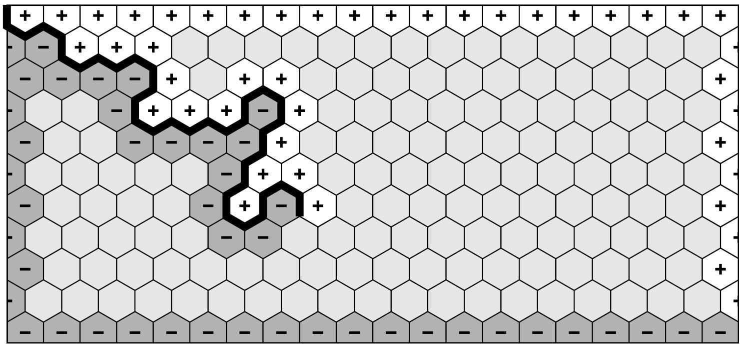

Figure 1: The percolation growth process.

left-hand corner. The walk always grows only at its tip, and can turn to the L or R with equal probability (it cannot retrace the previous step). This decided by choosing whether to mark the hexagon in front of the tip by $+$ or $-$. Thus, as the walk grows, all the hexagons to its immediate L are $+$ and those to its immediate R are $-$. Note that if the walk bends back on itself, it may encounter a hexagon which has already been marked. In that case it has no choice: it turns away from its previous trace. We also need to specify what happens at the boundary: we can do this by previously marking all the boundary hexagons between the starting point and some other point (in this case the bottom R corner) $+$, and the rest of the boundary hexagons $-$.

We end up with a growing self-avoiding walk on the lattice. It gets reflected from the boundary and will therefore always end up at the other point (in this case the bottom R corner). Note also that at any intermediate time there is always at least one path from the growing tip to this final point which does not intersect the existing part of the walk.

The rules we have given define a probability distribution, or measure, on the set of all such paths. (Note there is a finite number of such paths on a finite lattice.) However, there is another way in which we could generate paths according to this probability distribution. Suppose we start with an empty lattice, and mark all the boundary hexagons as before, then also mark all the interior hexagons $\pm$ independently with equal probability. This is the measure for independent *site percolation*, see Fig. 2. For each configuration these is exactly one path connecting the top L and bottom

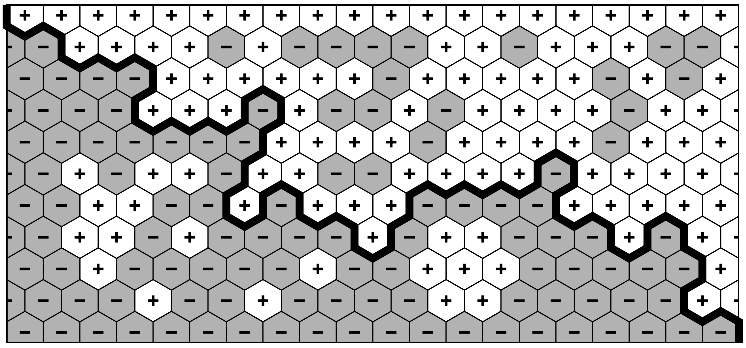

Figure 2: The completed process.

R corners, such that all the hexagons on one side are marked + and all those on the other are −. In addition, there will in general be other closed paths which are the boundaries of clusters of + and − hexagons. But we ignore these. A little thought should convince you that the measure on paths between the 2 boundary points we get this way is the same as that defined by the growth process. The path is simultaneously the boundary of the connected + cluster containing the upper R boundary of the whole domain, and that of the − cluster containing the lower L boundary.

Problem. What happens if the probabilities of turning L and R are different?

Problem. Consider a similar process on the square lattice, in which the walk can only turn L or R with equal probability (if this is possible) but cannot go straight on or backwards at each vertex. Show that in this case you get boundaries of (critical) bond percolation clusters on a (different) square lattice.

There is another important property of this measure we can see going back to Fig. 1: if we physically cut the lattice along the existing part of the path, then the measure on the rest of the path is just the same as if we started the process from the tip, and defined it in the cut domain, with the L and R sides of the existing part of the path forming part of the boundary. This is often called the *Markov* property. It is self-evident (at least for this model) on the lattice, but later on, when combined with conformal invariance, it will turn out to be very powerful.

Fig. 3 shows what happens if you simulate this process on a much bigger lattice. In this case, the whole domain is the upper half plane, and the

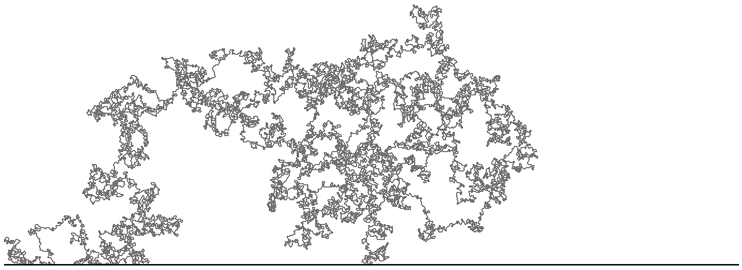

Figure 3: A much larger percolation growth process in the upper half plane.

boundary conditions are such that the path reflects from the real axis. You see that the walk continually does this, as well as reflecting off itself. It also appears to swallow whole regions of the plane.

2.2 Ising model

It is simple to modify the weights for this process and hence the measure on paths which results. For example, in Fig. 1, let us suppose that $\pm$ now represent the possible values of an Ising spin $s(r)=\pm 1$ on each hexagon. If R is the hexagon just in front of the tip, mark it $\pm$ with probability $\frac{1}{2}(1+m(R))$, where $m(R)$ is the mean magnetisation of an Ising spin at R in an Ising model defined in the whole of the cut domain, with the specified boundary conditions. That is

$$m(R)=\frac{\operatorname{Tr} s(R)e^{J\sum_{r,r'}s(r)s(r')}}{\operatorname{Tr} e^{J\sum_{r,r'}s(r)s(r')}} \qquad (1)$$

where $-J\sum_{r,r'}s(r)s(r')$ is the usual Ising energy function in units of kT.

Alternatively, imagine sampling the variable $s(R)$ in an equilibrium Monte Carlo simulation of the Ising model in the cut domain, with the specified fixed boundary conditions. If $s(R)=+1$, the walk turns to the R, if -1, it turns to the L. If $J=0$ (infinite temperature) we get back to the case of percolation. When $J\neq 0$ actually carrying out the algorithm would obviously be much harder. In fact it would be easier to simulate an equilibrium Ising model in the full rectangle with the same boundary conditions as before, sample one configuration, and draw the path connecting the top L and bottom R corners. Once again, the measure on curves we get this way is the same as that given by the growth process. This can be understood as follows: we start off with a configuration space with 2^N points, corresponding to all possible values of the N free spins in the lattice. The Ising weights define a measure on this space. As the path grows, we narrow this space down by repeatedly fixing the spin in front of the growing tip. This is called a *filtration*. At the same time, we condition the measure in such a way that the spins either side of the growing path become fixed. We can evaluate the expectation value of the spin $s(R)$ either in terms of this conditioned measure, or the original one. This guarantees that the measure on paths will satisfy the Markov property.

Looking at the example of the Ising model, we see that percolation has a special property, called *locality*: if we modify the domain by distorting its boundary or taking out a piece of the interior, the measure on paths which do not intersect this is unchanged: it terms of the growth problem, the growing path does not 'feel' the boundary of the region until it actually hits it. This is not true of the Ising model: the magnetisation of $s(R)$ depends on the shape of the whole region.

2.3 Harmonic explorer

There are many other ways we can similarly define measures on curves with the Markov property. We just have to give a rule whereby the relative probabilities of the 'spin' $s(R) = \pm 1$ in front of the tip are determined. The *harmonic explorer* is defined as follows: given the first part of the path, start a random walker off at the site R on the lattice of hexagons. This time it is a free random walk, moving from one hexagon to any of the 3 neighbouring hexagons with equal probability. Eventually it will hit the boundary, either that of the original domain or the L or R side of the existing path. If it hits it at a hexagon which is labelled +, the value of $s(R)$ is taken to be +1 and the walk turns R, and vice versa. A little thought shows that the space of paths we get this way is the same as in percolation or the Ising model, but the weights are different, although, by definition, they satisfy the Markov property.

Do these weights correspond to those of the paths in some equilibrium process? It is easy to show that the probability $\frac{1}{2}(1+m(r))$ that a free random walk starting at a given hexagon at r hits the boundary at a $+1$ site is simply the solution of the discrete Laplace equation $\Delta_{\rm lat}m(r)=0$ with the specified boundary conditions. Equivalently, if we define a real continuous variable $\phi(r)$ on every hexagon, then $m(R)$ is the expectation value of $\phi(R)$ with respect to the Gibbs measure $\exp\big(-\beta\sum_{r,r'}(\phi(r)-\phi(r'))^2\big)$, with the variables $\phi(r)$ being fixed to the values ±1 on the boundary – a gaussian model. Thus we could imagine doing a Monte Carlo simulation of this equilibrium model with $\phi=\pm1$ on the boundary and either side of the existing path, taking the value of $\phi(R)$, and saying that the path turns to the R or L according as $\phi(R)>0$ or <0. This would give the same measure on paths as defined by the harmonic explorer.

Unfortunately these weights do not satisfy the same simple conditional rules as in the Ising model: if we simulate the gaussian model in the full domain, we can identify a path between the end-points for which $\phi>0$ on the hexagons to its L and $\phi<0$ on those to its R, but this does not mean that they must take values ±1. It is a remarkable result (due to Sheffield and Schramm) that, for a particular value of β, the measure on paths we get by the two methods is the same in the continuum limit.

2.4 Self-avoiding walks

We have argued that some simple equilibrium models satisfy the Markov property and also have an interpretation as simple growth processes. As an example of an equilibrium model where the latter interpretation is absent, consider the classic problem of *self-avoiding walks* (SAWs). Of course, all the paths we have discussed so far have been self-avoiding, but in this case we shall weight them differently. Consider, for example, such walks in the half-plane which begin at a fixed point on the boundary, say the origin. It is straightforward to show that the total number of such walks of N steps grows as μ^N, where μ is lattice-dependent. Thus it is natural to weight each walk with a factor $\mu^{-\rm length}$, but the main point is that all walks of the *same* length are weighted *equally.* (This is supposed to model a polymer in a good solvent.)

Given that the walk chooses the first N' steps to be along a given path, there is no simple algorithm to determine the relative probability that it will turn R or L at the next step. To do this we would have to sample the entire population of walks. However, it does satisfy the Markov property: given the first part of the path, the measure on the remaining part weights all walks by the correct factor $\mu^{-(N-N')}$.

SAWs satisfy the nice property called *restriction*: if we restrict the domain as discussed for percolation, then the measure on SAWs in the whole domain, restricted to stay in the new domain, is that same as the measure on SAWs in the restricted domain. Once again, this apparently trivial observation, when combined with conformal invariance, will turn out to be very powerful.

2.5 Conformal invariance

So far, everything has been on a lattice. But if we look at pictures such as Fig. 3, we might ask the question as to whether it makes sense to discuss the properties of these curves in the continuum limit, when the lattice spacing a is taken to zero but the domain is kept fixed. For this we need some notions from the theory of critical behaviour in equilibrium statistical mechanics, in particular that of the correlation length ξ. This is the length scale over which correlations decay, typically exponentially. In the Ising model it these might be the spin-spin correlations, in percolation the probability that two points are in the same cluster. In general, ξ is of the order of a few lattice spacings. But at a continuous critical point, correlations do not decay exponentially, but as power laws. Renormalisation group (RG) theory tells us that if we take the limit $a \to 0$ and approach the critical point in such a way that ξ remains finite, then all correlation functions, if multiplicatively renormalised by suitable powers of a, have a finite limit, called the *scaling limit*. Moreover, this limit exists even at the critical point, when we have a stronger property: *scale invariance*. Because rescaling a is effectively the same as rescaling all distances, any such change can be absorbed into a renormalisation of the correlations.

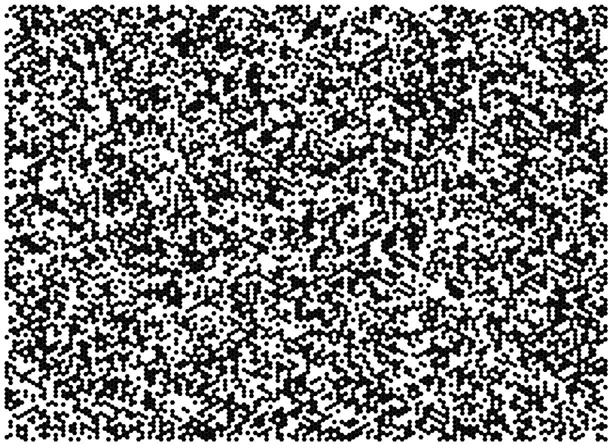

Figure 4: Typical configuration of site percolation in a large rectangle.

Scale invariance can be understood at the level of the measure on configurations on the lattice. Take percolation as an example. Fig. 4 shows a typical configuration of the site percolation model we defined before, with free boundary conditions. We see clusters of all sizes. Now, if we take, for example, the top left quadrant of this picture and blow it up so it is the same size as the whole, *statistically* we cannot tell it apart from the the original, as long as we look only a those features on scales $\gg a$ (i.e. those which survive in the scaling limit.)

Scale invariance can be generalised to a larger symmetry. Suppose we take the top L quadrant again and subject it to a conformal transformation, i.e. one which locally is equivalent to a rotation + a scale transformation, except that now the scaling factor can depend on position. Such transformations always preserve angles, however, see Fig. 5. Then the claim is that if we compare the conformally transformed picture with a similarly shaped region of the original, they are statistically indistinguishable (for features on scales $\gg a$.)

Of course, this is all conjecture for most critical statistical mechanics models, although many numerical studies support it. It does, however, suggest that the paths we have described before, in the case of critical models, should be conformally invariant in the scaling limit in a sense we will make precise.

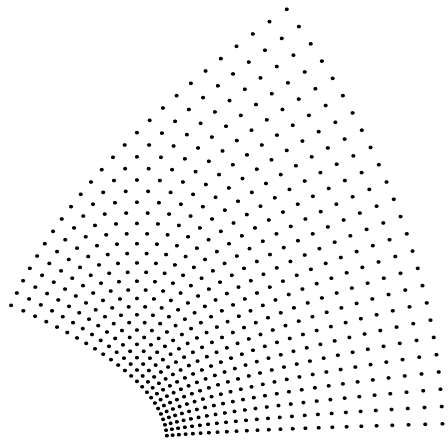

Figure 5: An example of a conformally transformed square.

2.6 Multiple curves

An obvious generalisation is to consider more than one curve, for example two. An important question is whether we get the same final joint measure by, for example, first growing one curve and then allowing the other to grow in the domain whose boundary is partly defined by the first, or by allowing them to grow at the same time. In the examples of percolation and the Ising model discussed above, the answer is obviously yes, because the final measure is just the equilibrium one.

2.7 Other growth processes

Although they will not be the main subject of these lectures, we should mention other related problems in 2d. We can, for example, allow the path to branch: it could either do so at the growing tip, in which case we would need to give a rule for the probability of such branching to occur, or at some other point. In neither case is there is a simple generalisation of the Markov property. But there are other simple models. For example, given the existing set P (no longer necessarily a simple curve on the lattice) we reverse the harmonic explorer and allow a random walker to approach from infinity and ask for the point at which it first hits P: at this point a hexagon is added. Equivalently, we can solve the discrete Laplace equation

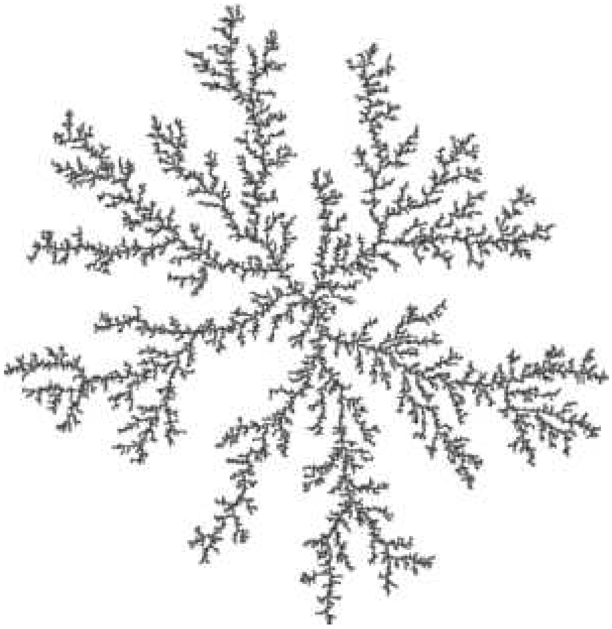

Figure 6: An example of a DLA cluster.

$\Delta\phi = 0$ with the boundary condition that $\phi \sim \log|z|$ as $|z| \to \infty$ and $\phi = 0$ on P. The value of ϕ on a hexagon neighbouring P determines the relative probability that this hexagon is added to P. This process defines (lattice) *diffusion limited aggregation* (DLA). When it is continued we get pictures that look like that in Fig. 6, which are supposed to represent clusters in smoke, or, perhaps, patterns in viscous fingering experiments.

The conformal mapping methods we shall introduce to describe non-branching processes can certainly be applied to these models, but unfortunately they do not appear to be analytically tractable in the same way.

3 Conformal mappings

3.1 Riemann's theorem

In two dimensions, the theory of conformal mappings is inseparable from that of complex analytic functions. Consider some general differentiable mapping $(x, y) \to (u(x, y), v(x, y))$ of some domain D to another D' in the plane, where (x, y) are cartesian coordinates. In particular suppose this is infinitesimal, i.e. $u-x$ and $v-y$ are small. We can think of this as being like the deformation of an elastic body. The local strain field has components $(u_x, u_y; v_x, v_y)$, which can be decomposed into a local rotation $u_y - v_x$, a local scale transformation u_x+v_y, and a local shear with components u_x-v_y and u_y+v_x. The requirement that the mapping is conformal, i.e. these last two vanish, gives the Cauchy-Riemann equations which imply that u and v are the real and imaginary parts of a differentiable function of $z = x+iy$:

$$u(x, y) + iv(x, y) = f(x + iy) \tag{2}$$

so that $u_x = v_y = \operatorname{Re} f'(z)$ and $u_y = -v_x = \operatorname{Im} f'(z)$. It is also necessary that $f'(z)$ does not vanish in the interior of D. This also guarantees that f is univalent: each point in D gets mapped to just one point in D' and vice versa.

A *simply connected* domain is one in which every closed loop can be continuously shrunk to a point. The most important theorem in conformal mapping is Riemann's: it states that for any two simply connected domains D and D' in the plane there is always an analytic function f which maps the interior of D to that of D'. This is true no matter how irregular the boundary of D, although in this case f is not necessarily analytic on the boundary.

If we take D' to be the upper half plane $\mathbf{H}$ and $h(z) = a + b/(z - c)$, with a, b and c all real, we see that $(h \circ f)(z) = h(f(z))$ does the job just as well as f, so that, in general, there are three arbitrary real parameters which can be fixed in Riemann's mapping.

3.2 Conformal mappings and growth problems

This is the central part of these lectures. Suppose we have some growing set K_t in a domain D, which is such that its complement in D, denoted by $D \setminus K_t$, is simply connected. Then Riemann's theorem assures us that there exists a conformal mapping g_t from $D \setminus K_t$ back to D, see Fig. 7. We can say that g_t eliminates K_t. So, rather than thinking of how K_t changes, we can think how the mapping g_t changes as K_t grows. This is often easier,

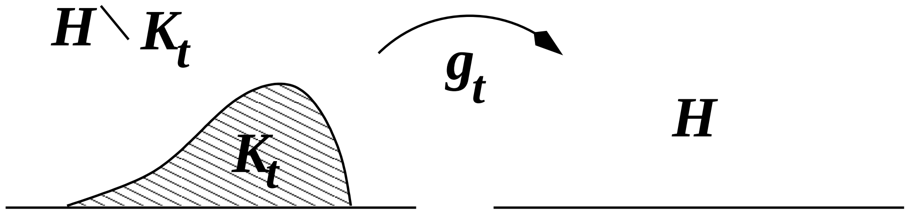

Figure 7: g_t maps $\mathbf{H} \setminus K_t$ to $\mathbf{H}$.

because for any point z of $D \setminus K_t$ not on its boundary, $g_t(z)$ changes much more smoothly.

In what follows we shall mostly assume that D is the upper half plane $\mathbf{H}$, so that K_t is a union of sets connected to the real axis. If we are trying to describe conformally invariant measures on sets, we can do this without loss of generality. However it should be stressed that in other situations, the form of the results depends on the choice of D.

Moreover, because of the freedom in Riemann's theorem, we can always choose that $g_t(z) \sim z$ as $z \to \infty$. We have the freedom of fixing one more real constant: this can be done by demanding that the $O(z^0)$ term in the expansion of g_t about infinity vanishes:

$$g_t = z + 0 + O(1/z)$$

A simple example

Suppose K_t is an interval $(0, i\ell)$, that is a straight stick of length ℓ perpendicular to the real axis. It is easy to check that

$$g_t(z) = (z^2 + \ell^2)^{1/2}$$

maps $\mathbf{H} \setminus (0, i\ell)$ to $\mathbf{H}$. Note that it maps the whole R half of the real axis to (ℓ, ∞), the R side of the stick to $(0, \ell)$, the L side to $(-\ell, 0)$ and the rest of the real axis to $(-\infty, -\ell)$.

It is useful to note the inverse function $f_t = g_t^{-1}$ which in this case takes the form

$$f_t(w) = (w^2 - \ell^2)^{1/2}$$

Note that this develops an imaginary part on the real axis at the image of the boundary of K_t.

3.3 Loewner's equation

In general, since $f_t(w)$ is analytic in the upper half plane and real on some part of the real axis, we can always write, by Cauchy's theorem

$$f_t(w) = g_t^{-1}(w) = w - \int \frac{\rho_t(x')dx'}{w - x'}$$

where $\rho_t \geq 0$.[1]

Problem. What is g_t and ρ_t if K_t is a half disc of radius r centred at the origin?

An important property of g_t, or equivalently of K_t, is the *half-plane capacity*

$$C_t \equiv \int \rho_t(x')dx'$$

so that, as $w \to \infty$

$$f_t(w) \sim w - \frac{C_t}{w} + \cdots$$

and

$$g_t(z) = z + \frac{C_t}{z} + O(1/z^2)$$

Problem. To what 2d electrostatics problem is C_t the answer?

Note that if the set grows, that is ρ_t increases, then so does C_t. Since we haven't in general defined what we mean by 'time' t (for a fractal, the area or length of K_t is not useful), we might as well reparametrise it so that

$$t \equiv C_t/2$$

(the factor of 2 is historical). This defines *Loewner time.*

It has the nice property that if we compose two such mappings, g_{t_1} ad g_{t_2}, their times are additive:

$$g_{t_2}(g_{t_1}(z)) \sim g_{t_1}(z) + \frac{2t_2}{g_{t_1}(z)} + \cdots \sim z + \frac{2t_1}{z} + \frac{2t_2}{z} + \cdots$$

[1] Actually we should write $\rho_t(x')dx' \to d\mu_t(x')$ where μ_t is a measure on the real axis.

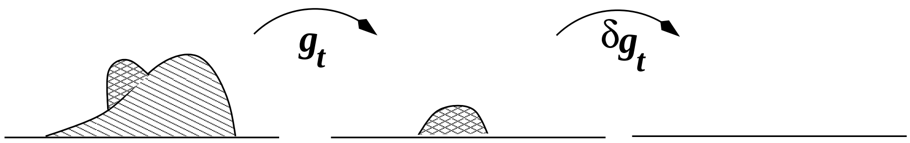

Figure 8: The sequence of maps g_t and δg_t.

Now suppose we have some increasing sequence of sets, that is $K_t \subset K_{t'}$ if $t < t'$. Suppose that $t' = t + \delta t$ where δt is small, and let $\delta K_t \equiv g_t(K_{t+\delta t} \setminus K_t)$, eliminated by (say) δg_t. See Fig. 8. Then $g_{t+\delta t} = \delta g_t \circ g_t$, or, equivalently, $g_t = \delta g_t^{-1} \circ g_{t+\delta t}$. Using the dispersion representation for δg_t^{-1} we therefore have

$$g_t(z) = g_{t+\delta t}(z) - \int \frac{\delta\rho_t(x')dx'}{g_{t+\delta t}(z) - x'}$$

Now let $\delta t \to 0$, with $\delta\rho_t \sim \delta t \nu_t$.

$$\frac{d}{dt} g_t(z) = \int \frac{\nu_t(x')dx'}{g_t(z) - x'}$$

The measure $\nu_t(x')dx'$ determines precisely how and where K_t grows. If we use Loewner time, then $\int \nu_t(x')dx' = 2$.

The physical interpretation is more clear for the evolution of the inverse mapping $f_t \equiv g_t^{-1}$. This satisfies

$$\frac{d}{dt} f_t(w) = -f_t'(w) \int \frac{\nu_t(x')dx'}{w - x'}$$

This is the equation determining the analytic function $\dot{f}_t(w)/f_t'(w)$ given its imaginary part $\pi\nu_t(x)$ on the real axis.

The boundary of K_t is the curve $z = f_t(x)$, $x \in \mathbf{R}$. The above equation then says that its normal velocity is

$$v_n(z) = \pi |f_t'(x)| \nu_t(x)$$

Thus $\nu_t(x)$ tells us amount of matter per unit time which is being added to K_t at $z = f_t(x)$; the factor $|f_t'(x)|$ takes account that this is being expressed in terms of x rather than z. Thus if we add matter at a uniform rate and density in the z coordinate, this corresponds to a nonuniform density $\nu_t \sim |f_t'|^{-1}$ in x.

So far this has been a way of describing any suitable 2d growth problem using conformal mappings.

Problem. Use this to formulate equations for diffusion-limited aggregation in the half-plane.

Now let us specialise to the case where the growth occurs only at one point, the tip, at any given time. Then we can write $\nu_t(x) = 2\delta(x - a_t)$. a_t is the image of the growing tip under g_t. Note that if we use Loewner time there can be no implicit dependence on g_t in this. This gives the standard Loewner equation

$$\frac{d}{dt}g_t(z) = \frac{2}{g_t(z) - a_t}$$

Moreover if the growing tip traces out a curve, we expect that a_t is a *continuous* real function. We can think of this equation in two ways. Given a suitable growing curve (basically one where the growing tip is always connected to infinity), we can define a sequence of conformal mappings g_t, and hence a real continuous function $a_t \equiv g_t(\text{tip})$, which will automatically satisfy the Loewner equation. Conversely, given a (suitable) continuous real function a_t (a sufficient condition is that it have Hölder exponent $> \frac{1}{2}$), we can integrate the Loewner equation to get a sequence of functions g_t and hence a sequence of tips $g_t^{-1}(a_t)$ which will trace out a curve.

Thus we have reduced the problem of classifying (random) curves in $\mathbf{H}$ to a similar problem of (random) continuous real functions. Putting a measure on such curves is equivalent to putting a measure on such functions.

Problem. What happens in general if a_t has discontinuities?

4 Stochastic Loewner Evolution

4.1 Schramm's Theorem

Now we are going to combine the idea of Loewner evolution with that of conformal invariance. The fact that Loewner uses conformal mappings of

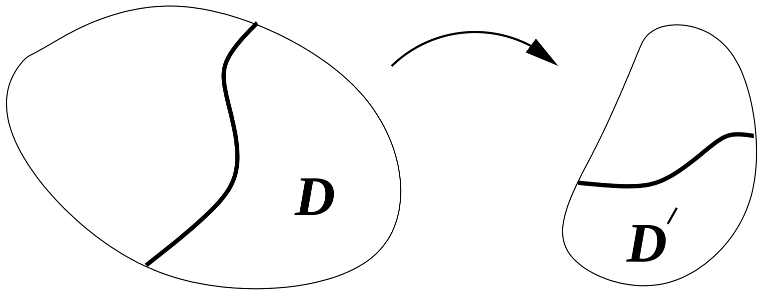

Figure 9: The image of a curve under a conformal transformation.

course makes this natural. We first need a precise definition of what it means for a measure on curves to be conformally invariant.

Suppose we start with some lattice model in a simple connected domain D, with boundary conditions prescribed, or otherwise, so that we have a well-defined measure on lattice curves γ connecting two boundary points z_1 and z_2. Assume that, in the scaling limit when the lattice spacing $\to 0$, this defined a measure $\mu(\gamma; z_1, z_2; D)$ on continuous curves. Now consider another domain related to D by the conformal mapping $\Phi : D \to D'$. This induces a mapping $\Phi \circ \mu$ on curves γ' from $\Phi(z_1)$ to $\Phi(z_2)$ in D' (Fig. 9: we can imagine computing the expectation values of any observable of γ' by conformally transporting the the problem back to D. Conformal invariance states that the measure $\Phi \circ \mu$ is the *same* as the measure we get by taking the scaling limit of lattice curves in D'. Note that the lattice is not supposed to be transformed (otherwise this would be a tautology): we take the scaling limit of (e.g.) the same square lattice in both D and D'.

Conformal invariance is a property of the scaling limit on curves in an arbitrary domain which may or may not be true. We are going to assume it holds for the lattice models we discussed, and explore its consequences: in certain cases, however, it has actually been proved.

We are now going to apply this principle to the Loewner mapping $g_t : \mathbf{H} \setminus K_t \to \mathbf{H}$. Suppose we evolve for a time t and then for a further time $s - t$ (see Fig. 10). The Markov property tells us that the conditional measure on $K_s \setminus K_t$, given K_t, in $\mathbf{H}$, is the same as the unconditional measure on $K_s \setminus K_t$ in $\mathbf{H} \setminus K_t$. Under g_t, this is the same as the measure on sets K_{s-t} in $\mathbf{H}$, except that these will correspond to curves starting at

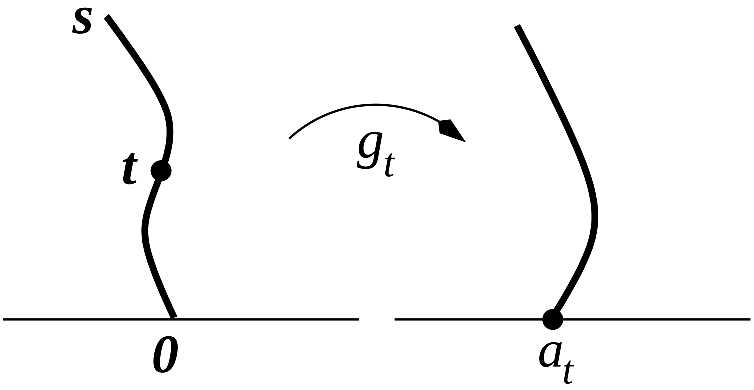

Figure 10: The curve evolved for a time t and a further time $s-t$, and its image under g_t.

a_t rather than a_0.

Translated into a statement about the measure on a_t, this says that the law of $a_s - a_t$, given a_t, is the same as the law of a_{s-t}, given a_0. In particular, this implies that the increments $a_{(n+1)\delta t} - a_{n\delta t}$ are independent identically distributed random variables, for all $\delta t > 0$. The only process satisfying this is 1d Brownian motion, with a possible drift term:

$$a_t = \sqrt{\kappa} B_t + \alpha t$$

where $\langle B_t \rangle = 0$, $\langle (B_s - B_t)^2 \rangle = |s-t|$, and κ, α are constants. If the measure on the curve in $\mathbf{H}$ is symmetric under $x \to -x$ (which is the case for the examples we discussed), then the drift term $\alpha = 0$.

This defines the family of conformally invariant measures SLE_κ on curves from a_0 to ∞ in $\mathbf{H}$. Different values of κ correspond to different models, e.g.:

- $\kappa = 6$: boundaries of percolation clusters
- $\kappa = 3$: boundaries of Ising spin clusters
- $\kappa = 4$: the harmonic explorer, and level lines of a gaussian field
- $\kappa = \frac{8}{3}$: self-avoiding walks

Some of these are proven and some (so far) conjectured.

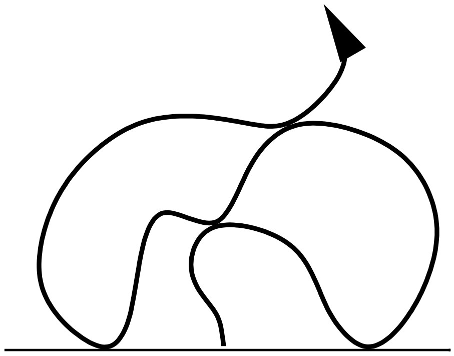

Figure 11: A curve with many double points, and its hull.

4.2 Properties of SLE

The first question to ask is whether the SLE trace is really a simple curve or whether it looks more like Fig. 11. Note that although the lattice curves we discussed are simple, they may approach themselves arbitrarily closely as the lattice spacing $a \to 0$, so the question is whether such double points appear with finite probability.[2] It turns out the answer depends on κ.

To see this it is useful to consider the shifted mapping function $\hat{g}_t(z) \equiv g_t(z) - a_t$, which satisfies

$$d\hat{g}_t(z) = \frac{2dt}{\hat{g}_t(z)} - da_t$$

(Note that we must now write this in differential form as a stochastic equation, since a_t is not differentiable.) This always maps the growing tip back to the origin. A point x_0 on the real axis gets mapped into $x_t = g_t(x_0)$ where

$$dx_t = \frac{2dt}{x_t} - \sqrt{\kappa}\, dB_t$$

This real process is much studied in the literature, and is called the Bessel process. It describes the motion of a particle repelled from the origin and also subject to Brownian noise. If κ is small, the repulsive force always wins, and the particle goes off to infinity (almost surely), while if it is

[2]If this is the case, we should apply Loewner's mapping not to the curve itself, but to its *hull* K_t: that is the set of points both on the curve, and enclosed by it (and the real axis). However this does not substantially change any of the earlier arguments.

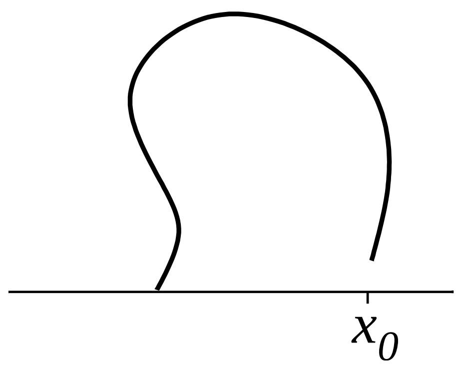

Figure 12: A curve about to swallow a whole region and hit the axis.

large, eventually random noise will cause the particle to reach the origin. The critical value of κ can be found by the crude approximation of either ignoring the noise, in which case $x_t^2 \sim 4t$, or ignoring the repulsion, when $\langle x_t^2 \rangle \sim \kappa t$. So for $\kappa < 4$ we expect the particle to go off to infinity, and for $\kappa > 4$ it will be absorbed at the origin. This turns out to be correct.

Problem. Show that the linear distance of a free Brownian particle from the origin in d dimensions evolves according to a Bessel process. How does the effective value of κ depend on d?

What does this correspond to in terms of $\hat{g}_t$? In Fig. 12 we show the curve about to hit the real axis at x_0. The whole real axis and the R and L sides of the curve all get mapped by $\hat{g}_t$ to the real axis, but the R side of the curve and the section of the real axis about to be swallowed get mapped into a very small segment, which gets smaller as the tip approaches x_0. This means that, as this happens, both x_0 and the tip, as well as the whole region which is swallowed, get mapped in to the origin, that is $x_t \to 0$.

Thus for $\kappa < 4$ (and in fact $\kappa = 4$) the curve is simple: there are no double points (almost surely), while for $\kappa > 4$ the curve has double points (and in fact, since it is self-similar) it has infinitely many such points in any finite region containing the curve. Another consequence is that eventually every point in the upper half plane gets swallowed.

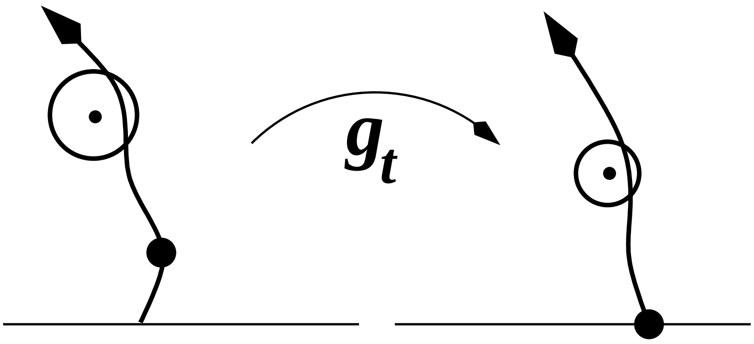

Figure 13: The curve intersecting a small disc, and its image under $g_{\delta t}$.

5 Calculating with SLE

Most SLE results come down to some very simple calculations of stochastic analysis. We consider a couple of examples.

5.1 Fractal dimension

The physicist's definition of the fractal dimension of an object is roughly as follows: suppose we try to cover it with a large number of (overlapping) discs of radius ϵ. Let $N(\epsilon)$ be the minimum such number. Then if $N(\epsilon) \sim \epsilon^{-d_f}$ as $\epsilon \to 0$, d_f is the fractal dimension.

An equivalent definition, for random fractals, is to ask for the probability $P(r,\epsilon)$ that the object intersects a disc of radius ϵ centred on r. If

$$P(r,\epsilon) \sim \epsilon^{D-d_f} f(r)$$

as $\epsilon \to 0$, where D is the embedding dimension (2 in our case), then d_f is again the fractal dimension. We shall use this last definition and get an equation for $P(r,\epsilon)$.

Suppose the curve γ intersects the disc, see Fig. 13. Look at the image of this picture under $\hat{g}_{\delta t}$. This maps the first short section of the curve to the real axis, and the rest of the curve to a new curve also starting from the origin (because we use $\hat{g}_t$ rather than g_t.) The point r gets moved to $g_{\delta t}(r)$ and the radius of the disc changes slightly to $|g'_{\delta t}(r)|\epsilon$. But the measure on the image is the same as that on the original curve, by conformal invariance.

So we can write an equation

$$P(x,y,\epsilon) = \langle P\left(x + \frac{2x\delta t}{x^2+y^2} - \sqrt{\kappa}\delta B_t, y - \frac{2y\delta t}{x^2+y^2}, (1 - \frac{2(x^2-y^2)\delta t}{(x^2+y^2)^2})\epsilon\right)\rangle$$

where the average is over the Brownian motion δB_t, i.e. the possible configurations of the first of the curve which has been erased by $g_{\delta t}$.

Expanding to first order in δt, and remembering that $\langle(\delta B_t)^2\rangle = \delta t$ (so we have to expand this part to second order), we get the PDE

$$\left(\frac{2x}{x^2+y^2}\frac{\partial}{\partial x} - \frac{2y}{x^2+y^2}\frac{\partial}{\partial y} + \frac{\kappa}{2}\frac{\partial^2}{\partial x^2} - \frac{2(x^2-y^2)}{(x^2+y^2)^2}\epsilon\frac{\partial}{\partial \epsilon}\right)P = 0$$

Thus if $P \sim \epsilon^{2-d_f}$, we see that $2 - d_f$ is an eigenvalue of a certain differential operators. This is typical of SLE calculations: the various critical exponents all turn out to be given by eigenvalue problems. In this case it turns out that the eigenfunction can be guessed by inspection: we find

$$P \sim \epsilon^{1-\kappa/8}\, y^{(\kappa-8)^2/8\kappa}\, (x^2+y^2)^{(\kappa-8)/2\kappa}$$

Thus we see that $d_f = 1 + \kappa/8$. This is correct for $\kappa \le 8$: when $\kappa > 8$ there is another solution with $d_f = 2$.

5.2 Crossing probability

Looking at Fig. 4 we can ask the question whether there is a crossing e.g. from the L edge to the R edge only on white hexagons. The probability this happens will depend on the shape of the rectangle.

Because of conformal invariance we can transform this problem to the upper half plane $\mathbf{H}$. It is always possible to make a fractional linear conformal mapping which takes the L edge into $(-\infty, x_1)$ and the R edge into $(0, x_2)$, where $x_1 < 0$ and $x_2 > 0$. However, within SLE, it still takes a certain amount of ingenuity to relate this problem to a question about a single curve. Now go back to the lattice picture and consider critical site percolation on the triangular lattice in the upper half plane, so that each site is independently coloured black or white with equal probabilities $\frac{1}{2}$. Choose

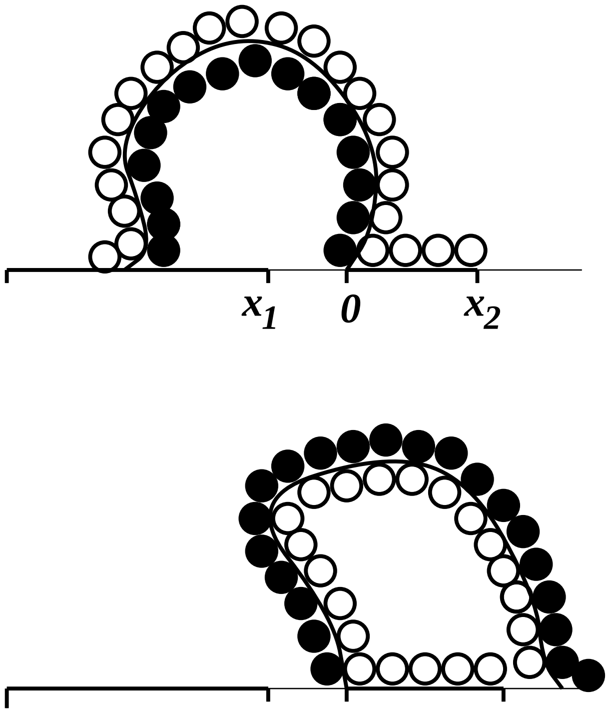

Figure 14: Is there a crossing on the white discs from $(0, x_2)$ to $(-\infty, x_1)$? This happens if and only if x_1 gets swallowed by the SLE before x_2.

all the boundary sites on the positive real axis to be white, all those on the negative real axis to be black (see Fig. 14). There is a cluster boundary starting at the origin, which, in the continuum limit, will be described by SLE_6. Since $\kappa > 4$, it repeatedly hits the real axis, both to the L and R of the origin. Eventually every point on the real axis is swallowed. Either x_1 is swallowed before x_2, or vice versa.

Note that every site on the L of the curve is black, and every site on its R is white. Suppose that x_1 is swallowed before x_2. Then, at the moment it is swallowed, there exists a continuous path on the white sites, just to the R of the curve, which connects $(0, x_2)$ to the row just above $(-\infty, x_1)$. On the other hand, if x_2 is swallowed before x_1, there exists a continuous path on the black sites, just to the L of the curve, connecting $0-$ to a point on the real axis to the R of x_2. This path forms a barrier (as in the game of Hex) to the possibility of a white crossing from $(0, x_2)$ to $(-\infty, x_1)$. Hence there is such a crossing if and only if x_1 is swallowed before x_2 by the SLE curve.

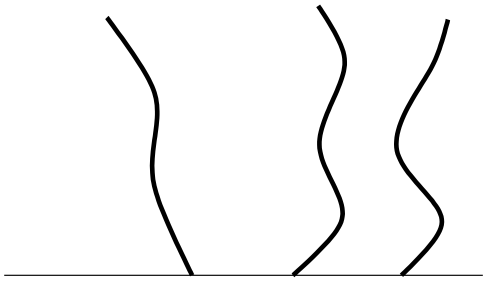

Figure 15: Multiple curves.

Recall that in Sec. **??** we related the swallowing of a point x_0 on the real axis to the vanishing of $x_t = g_t(x_t) - a_t$, which undergoes a Bessel process on the real line. Therefore

$$\Pr(\text{crossing from } (0, x_2) \text{ to } (-\infty, x_1)) = \Pr(x_{1t} \text{ vanishes before } x_{2t})\,.$$

Denote this by $P(x_1, x_2)$. By generalising the SLE to start at a_0 rather than 0, we can write a differential equation for this in similar manner to before:

$$\left(\frac{2}{x_1 - a_0}\frac{\partial}{\partial x_1} + \frac{2}{x_2 - a_0}\frac{\partial}{\partial x_2} + \frac{\kappa}{2}\frac{\partial^2}{\partial a_0^2}\right) P(x_1, x_2; a_0) = 0\,.$$

Translational invariance implies that we can replace ∂_{a_0} by $-(\partial_{x_1} + \partial_{x_2})$. Finally, P can in fact depend only on the ratio $\eta = (x_2 - a_0)/(a_0 - x_1)$, which again reduces the equation to hypergeometric form. The solution is (specialising to $\kappa = 6$ for percolation)

$$P = \frac{\Gamma(\frac{2}{3})}{\Gamma(\frac{4}{3})\Gamma(\frac{1}{3})}\eta^{1/3}{}_2F_1(\tfrac{1}{3}, \tfrac{2}{3}, \tfrac{4}{3}; \eta)\,.$$

5.3 Multiple curves

We can generalise Loewner's equation to describe the growth of not just one curve, but many. Suppose the set K_t consists of the union N curves (more generally their hulls), growing in the half plane in such a way that they do not trap each each other: that is $\mathbf{H} \setminus K_t$ is simply connected, see Fig. 15. Then we know there is a mapping g_t which sends this back to

H. If a_{jt} is the image of the tip of the jth curve, the generalised Loewner equation reads

$$\frac{d}{dt}g_t(z) = \sum_{j=1}^{N} \frac{2b_j}{g_t(z) - a_{jt}}$$

Note that if we choose t to be Loewner time we only know that $\sum_j b_j = 1$. In order to go further we need to specify the dynamics of the a_{jt} and the b_j.

In the case when we are trying to describe multiple curves in the conformally invariant scaling limit of 2d lattice models, we can argue as follows: let us consider the mapping $g_{t+\delta t} \circ g_t$, which evolves for a short time δt. We can imagine growing each curve j in turn for a short time δt_j. As this happens, a_{jt} will change by an amount $\sqrt{\kappa} B_{\delta t_j}$, while the other points with $k \neq j$ will move according to the usual Loewner mapping $a_{kt} \to a_{kt} + 2\delta t_j/(a_{kt} - a_{jt})$. Thus the total change in a_{jt} is

$$\delta a_{jt} = \sqrt{\kappa} B_{\delta t_j} + \sum_{k \neq j} \frac{2\delta t_k}{a_{jt} - a_{kt}}$$

We get consistency with the general form for g_t above if we choose $\delta t_j = b_j \delta t$. Using the fact that $B_{\delta t_j}$ then is the same as $\sqrt{b_j} \delta B_{\delta t}$ we get the following stochastic dynamics for the a_{jt}:

$$da_{jt} = \sqrt{b_j \kappa} dB_{jt} + \sum_{k \neq j} \frac{2b_j dt}{a_{jt} - a_{kt}}$$

where the B_{jt} are independent standard Brownian motions.

Thus the image points repel each other (corresponding to the entropic repulsion of the curves) and are also subject to Brownian noise. When the b_j are all equal, this process is well-known: it is called *Dyson's Brownian motion.* It describes, for particular values of κ, the evolution of the eigenvalues of random Hermitian matrices as their elements themselves undergo independent Brownian motions. If we rescale each a_{jt} by $\sqrt{t}$, the ensemble reaches a stationary distribution.

The interesting feature of the application to describing conformally invariant curves is that the final measure on the N curves is independent of the particular choice of the $\{b_j\}$. Taking all of then equal correspond to growing the curves at 'equal' rates (not in terms of length, however). But we could equally well choose one $b_j = 1$ and all the others zero. This would give us a measure on the jth curve, in the presence of the others.

Of course, one could take one's own favorite stochastic particle dynamics in 1d (as discussed in this school, the ASEP or SSEP, for example) and use the generalised Loewner equation to generate all kinds of measures on multiple curves. To what would these correspond?