

3 Statistical physics models

3.1 Percolation

In Werner's St. Flour lectures he discusses percolation in the first section and then in more detail in section 10. In the article by Kager and Nienhuis percolation is discussed in section 5. There is a vast literature on percolation. For the reader who wants more than we give here, there is an entire book: *Percolation*, by Geoffrey Grimmett.

Percolation can be defined on a variety of lattices. It also comes in two flavors - bond percolation and site percolation. We will only define one particular two-dimensional model. It is what is usually called site percolation on the triangular lattice. However, instead of using the triangular lattice we will define it in an equivalent way using the hexagonal lattice. We have a coin that has probability p of heads and $1 - p$ of tails. For each hexagon we flip the coin. If it is heads we color the hexagon gray; for tails we color it white. There are many questions one can ask about the resulting configuration. We start by looking at the connected components of the gray region.

If p is very small, the gray hexagons are very rare. So their connected components or "clusters" are typically very small. Large clusters are possible, but unlikely. If p is very close to 1, then most of the hexagons are gray and there should be an infinite cluster "percolating" throughout the lattice. When is there an infinite cluster?

Theorem 9 *There is a critical value of p , call it p_c , such that if $p < p_c$ then with probability one there is not an infinite cluster of gray hexagons and if $p > p_c$ then with probability one there is a unique infinite cluster of gray hexagons.*

For the hexagon model we are considering, the critical value of p is exactly $1/2$. For most versions of percolation the critical value is a non-trivial number which is not expected to be rational. There are many interesting questions one can ask about the clusters and there are a lot of rigorous results. We will focus on two questions - when do rectangles have left-right crossings and what do interfaces look like.

Fix $a > 0$ and $b > 0$ and consider a rectangle of sides a and b . We now introduce a hexagonal lattice whose bonds have length δ where δ is very small. We color the hexagons randomly as before. A left-right crossing of the rectangle is a path from the left edge to the right edge that only goes through gray hexagons. Let $p(a, b, \delta)$ be the probability of a left-right crossing. For $\delta > 0$, $p(a, b, \delta)$ is strictly between 0 and 1. We are interested in what happens to this number as $\delta \rightarrow 0$. Suppose that p is very small. So the vast majority of the hexagons are white. The probability of a left-right crossing is very small. It is not hard to prove for very small p that this probability converges to 0. Now suppose that p is very close to 1. Then the vast majority of the hexagons are gray. So the probability of a left-right crossing is almost 1. In fact, as δ goes to 0, this probability goes to 1. It is expected that for $p < p_c$, $p(a, b, \delta)$ will converge to 0, and for $p > p_c$, $p(a, b, \delta)$ will converge to 1. At $p = p_c$, $p(a, b, \delta)$ will converge to a number strictly between 0 and 1.

It can only depend on the ratio a/b . (This takes a little thought.) There is a famous formula for this function due to John Cardy. For the particular model we have defined, there is now a proof of Cardy's formula. The critical value of p depends on the lattice we consider, and whether we consider bond or site percolation. But it is believed that $F(b/a)$ is universal, i.e., independent of the choice of lattice and bond vs. site percolation. The question of whether or not there is a left-right crossing is a nice illustration of our probabilistic view of critical phenomena. When the model is not critical, whether or not there is such a crossing has a deterministic answer- below p_c there is not, above p_c there is. But when the model is critical, whether there is such a crossing is random.

We now study interfaces in our percolation model. We continue to work with a rectangle with sides of length a and b and a hexagonal lattice with edges of length δ . Fix two points z and w on the boundary of our rectangle. We will let z_δ and w_δ denote the points in the hexagonal lattice that are closest to z and w . We take a connected path of hexagons around the boundary of our rectangle, and refer to these hexagons as boundary hexagons. The two points z_δ and w_δ divide this path into two "arcs." We color the hexagons in one arc gray and in the other arc white. Then for the hexagons inside the rectangle that are not boundary hexagons, we color them by flipping the coin. The result is an interface. This is a curve that goes from z to w along bonds in the hexagonal lattice with the property that as we traverse the curve the hexagons on the left of the interface are all of the same color, and the hexagons on the right of the interface are all of the opposite color. (Which color depends on how we colored the two boundary arcs.) To be definite we will assume that the hexagons on the left of the interface as we traverse it from z to w are all white and those on the right of the interface are all gray.

To see that such an interface exists, we argue by induction. Suppose that we have defined an interface which begins at z and contains N bonds. So there is a path along N bonds in the hexagonal lattice such that as we traverse the path the hexagon on the left of the bond is white and the one to the right of the bond is gray. Now look at the last bond. Let x be the site it ends at. As we traverse the N th bond which ends at x , there is a white hexagon on the left and a gray one on the right. If the hexagon that we encounter at the end of the bond is white, then we turn right at x and if it is gray, we turn left at x . This extends the interface by another bond. It is not hard to see that this process cannot cross the perimeter of boundary hexagons (since a bond in the interface cannot have hexagons of the same color on both sides) except at the bond which goes to w . By drawing a few pictures you can convince yourself that this interface will never visit a site it has already visited.

We now ask what the interface looks like for small p . First consider the extreme case of $p = 0$. So there are no gray hexagons except for the ones in one of the two boundary arcs. Thus the interface will go from z to w along the gray boundary hexagons. Now suppose $p > 0$, but p is very small. So the typical cluster of gray hexagons is small. Note that any cluster of gray hexagons that is not connected to the boundary will not be part of the interface. So only clusters that are attached to the boundary can affect the




Figure 1: The precolation exploration process below criticality, $p = 0.3$.

interface. They are typically small. So now the interface will follow the boundary gray hexagons with occasional excursions of a few hexagons away from the boundary. See figure 1. As $p \rightarrow p_c$, the average number of hexagons in a gray cluster will go to infinity. But for a fixed $p < p_c$ the average number of hexagons in a gray cluster will be finite. Since the hexagons are of size δ this means that the typical excursions of the interface away from the boundary will be of size $C(p)\delta$. So in the scaling limit, $\delta \rightarrow 0$, the interface will converge to the polygonal path that goes from z to w along the boundary of the rectangle that had the gray squares. Once again we see that for a non-critical system the scaling limit behaves deterministically. For $p > p_c$ the same reasoning suggest that in the scaling limit the interface will go from z to w along the boundary of the rectangle, but in the other direction.

Of course, the interesting interface is when $p = p_c$. In this case the interface will converge to a random curve that goes from z to w and stays inside the rectangle. See figure 2. For $\delta > 0$ the interface does not intersect itself. The random curves we get by letting $\delta \rightarrow 0$ cannot cross themselves, but they may touch themselves. This does in fact happen for the percolation interface, but there is no simple way to see this.

There is nothing special about rectangles. We could repeat the above for any reasonable region in the plane. When we make the connection with SLE we will be particularly interested in simply connected domains. Given such a domain D and two points on its boundary z and w , the the scaling limit of interfaces between z and w in critical percola-




Figure 2: The precolation exploration process at criticality, $p = 0.5$.

tion gives a probability measure on continuous curves that go between z and w and stay inside D .

Processes that give random curves between two boundary points of a domain will be called “chordal.” In other models, e.g., the loop erased random walk, we will have random curves from a boundary point to an interior point in the domain. Such processes will be called “radial”. There will be two corresponding types of SLE - chordal SLE and radial SLE.

The interface has a property, **locality**, that will be important when we determine which particular SLE will describe the scaling limit of the interface. Suppose we change the color of some of the hexagons which do not have an edge in the interface. Then the interface in the new configuration will be the same as the interface for the configuration we had to begin with. There is another way to think about this property which is useful if you want to simulate an interface in an efficient way. Rather than determine the colors of all the hexagons before we find the interface, we can instead determine them only as we need them. We begin with only the boundary hexagons colored. We construct the interface as above, and whenever we encounter a hexagon whose color has not yet been determined, we flip the coin and color it appropriately. Once a hexagon’s color has been chosen, we do not change it if the interface happens to encounter it later on. This dynamic way of generating the interface is often called an exploration process. We start at z follow the exploration process for n bonds, and let $\gamma(n)$ be the vertex we are at. So $\gamma(n)$ is a

random nearest neighbor walk on the hexagonal lattice. We give a precise statement of the locality property.

Theorem 10 (*Locality property for the percolation exploration process*) *Let D and D' be simply connected domains in the plane, each of which is a union of hexagons. Suppose that $D' \subset D$, and that the boundaries of D and D' have non-empty overlap. Let z be a vertex of one of the hexagons which is in both boundaries, w a vertex with $w \in \partial D$, and w' a vertex with $w' \in \partial D'$. Let $\gamma'(n)$ be the percolation exploration process in D' from z to w' , and $\gamma(n)$ the percolation exploration process in D from z to w . Then up until γ leaves D' , $\gamma(n)$ and $\gamma'(n)$ have the same distribution. More precisely, for every m , if we condition on the event that $\gamma(1), \gamma(2), \dots, \gamma(m)$ all lie in D' , then $\gamma(1), \gamma(2), \dots, \gamma(m)$ and $\gamma'(1), \gamma'(2), \dots, \gamma'(m)$ have the same joint distribution.*

The locality property should be true in the scaling limit. We will eventually prove that SLE_κ has the locality property only for $\kappa = 6$. Schramm argued that the percolation interface must be given by SLE_6 if the model is conformally invariant. Smirnov then proved the particular model we have defined is conformally invariant. Percolation is more than just a single interface, so one can ask if SLE_6 helps describe all of percolation. There is some fascinating work on this by Camia and Newman.

3.2 Loop-erased random walk

We start with an ordinary random walk defined on \mathbb{Z}^2 which starts at the origin. So we let X_i be a *vector-valued* i.i.d. sequence with $X_i = (1, 0), (-1, 0), (0, 1), (0, -1)$ where each of the four values has probability $1/4$. The random walk is

$$\omega(n) = \sum_{i=1}^n X_i \tag{50}$$

These walks can, and often do, visit the same site more than once and so form loops. We will say that a walk is *simple* if it never visits a site more than once. We now define a construction that takes a walk and produces a simple walk. We take a random walk as above, but only consider the walk up to time n . Heuristically, we traverse the walk and whenever it forms a loop we erase the loop. The precise definition is as follows. We denote the simple walk we are going to define by γ . Let $\gamma(0) = \omega(0) = 0$. Assume inductively that $\gamma(i)$ is defined for $0 \leq i \leq j$. Let

$$m_j = \max\{i : i \leq n, \omega(i) = \gamma(j)\} \tag{51}$$

So m_j is the last time ($\leq n$) when ω visits $\gamma(j)$. Then we define $\gamma(j+1) = \omega(m_j+1)$. So $\gamma(j+1)$ is the next step in the original walk after its last visit to $\gamma(j)$. We stop when m_j is equal to $n-1$. We let the j for which this happens be $k-1$. So $\gamma(k) = \omega(n)$, i.e., the loop

erased random walk ends at the same site where the original random walk ended. Clearly, k is random. Note that many different ω 's can produce the same γ . So trying to compute the probability of a γ is quite nontrivial. We emphasize that the above construction is done for a finite random walk and the result depends on n . Suppose we take a single walk ω and carry out the loop erasure first for ω up to time n and then for ω up to time m . The two simple walks we get may have no relation to each other.

Now let A be a subset of \mathbb{Z}^2 which does not contain the origin. The definitions we are about to make need no further assumptions about A , but we will usually take A to be the lattice sites outside of some bounded domain containing the origin. In two dimensions the random walk is recurrent. With probability one it will eventually visit every site. So with probability one, the walk will enter A , i.e., there is an n such that $\omega(n) \in A$. Define τ_A to be this hitting time, i.e.,

$$\tau_A = \min\{n : \omega(n) \in A\} \tag{52}$$

The site $\omega(\tau_A)$ is in A and it is a boundary site in the sense that it has a nearest neighbor that is not in A . For every random walk ω , we take the walk up to time τ_A and carry out the loop-erasure defined above. The result γ is a simple random walk from 0 to a site in A such that γ does not hit A until its last step. This defines the loop-erased random walk or LERW.

We now fix a domain D containing the origin, and take the lattice to be $\delta\mathbb{Z}^2$. In the above construction we let A be the sites in the lattice that are outside of D . The LERW defined above gives a probability measure on the set of nearest neighbor simple walks on $\delta\mathbb{Z}^2$ that start at the origin, end at a site outside of D and do not leave D until their last step. Let y be a lattice site outside of D such that the probability the LERW ends at y is non-zero. (Note that y must have a least one neighbor in D in order for this probability to be non-zero.) Then we can condition the LERW to end at y . This just means that we take only the LERW's that end at y and renormalize the probability measure restricted to this set to make it a probability measure again. The result is a probability measure on the set of simple nearest-neighbor walks on $\delta\mathbb{Z}^2$ that start at the origin, end at y and remain in D until this final step to y . We will call this probability measure the LERW in D starting at the origin and ending at y . To take the scaling limit we send δ to zero. This gives a probability measure on continuous curves in the plane that start at the origin and end on the boundary of D . (This scaling limit has been proved to exist.)

Before we let δ go to zero, our loop-erased random walk is simple. However, this does not imply that in the limit the probability measure we get will be supported on simple curves in the plane. The curves we get cannot cross, but they may touch themselves. Recall that this happens in the scaling limit of percolation. For the loop-erased random walk this does not happen. The probability measure we get in the scaling limit on curves in \mathbb{R}^2 is supported on simple curves.

In the above we used the lattice \mathbb{Z}^2 . There are other two-dimensional lattices we could use - the triangular lattice or the hexagon lattice. The definition of the LERW on

these lattices is completely analogous. Universality is the statement that for each of these different lattices the LERW will have the same scaling limit. (This is a theorem I think.)

3.3 Self-avoiding random walk




Figure 3: Three self-avoiding walks in the full plane with 1K, 10K and 100K steps. Each walk has been scaled by $N^{-3/4}$.

We take a lattice, e.g., in two dimensions the square, triangular or hexagonal lattice, and we fix a natural number N . We consider all walks with N steps which start at the origin, take only nearest neighbor steps and do not visit any site more than once. So a walk ω is a function ω from $\{0, 1, 2, \dots, N\}$ into the lattice such that

$$\begin{aligned}
 \omega(0) &= 0 \\
 |\omega(i) - \omega(i-1)| &= 1, \quad i = 1, 2, \dots, N \\
 \omega(i) &\neq \omega(j), \quad 0 \leq i < j \leq N
 \end{aligned}
 \tag{53}$$

There are a finite number of such walks for any fixed N , and we put a probability measure on this set by requiring that all such walks be equally probable.

The self-avoiding walk is of interest to physicists since it is model for polymers in dilute solution. More generally, it is of interest since it is a simple model that exhibits critical phenomena and universality. There are a variety of critical exponents that describe the behavior of the model. For example, the growth of the mean distance the walk travels as a function of N is described by a critical exponent ν . It is defined by

$$E[\omega(N)^2] \sim N^{2\nu} \tag{54}$$

The expected value E is with respect to the uniform probability measure described above and $\omega(N)^2$ is the square of the distance from the origin to the lattice point $\omega(N)$. Everyone believes that in two dimensions $\nu = 3/4$. There are essentially no rigorous results on ν . In fact, there is not even a proof in two (or three) dimensions that ν is bigger than $1/2$, the value for the ordinary random walk. Figure 3 shows three self-avoiding walks with $N = 1,000$, $N = 10,000$ and $N = 100,000$. Each walk has been scaled by $N^{-3/4}$ so that they are all on a scale of order one.




Figure 4: Distribution of the end to end distance and the end to midpoint distance for the SAW. The distances have been rescaled so their mean is one. The end to end distance for the ordinary random walk is also show.

The exponent ν does not tell us how the endpoint of the walk is distributed, so the next quantity to look at is the distribution of $\omega(N)$. It is natural to scale it by a factor of $N^{-\nu}$ and study this distribution in the limit that N goes to ∞ . Of course, for the ordinary random walk this would give a Gaussian distribution. The limiting distribution for the

self-avoiding walk is not expected to be Gaussian, but is still expected to be rotationally symmetric, so we will only look at the distribution of $N^{-\nu}|\omega(N)|$. This distribution for the self-avoiding walk is shown in figure 4. We also show the distribution for the distance from the endpoint of the walk to its midpoint. We have scaled both random variables so that they each have mean equal to 1. For comparison the analogous distribution for the ordinary random walk is also shown. This discussion of the behavior of the end to end distance was for your general education. Unfortunately, SLE has nothing to say (so far) about this distribution. We now turn to things SLE has a lot to say about.

The above definition was in the full plane. Now let D be a simply connected unbounded domain with $0 \in \partial D$. An important example is the upper half plane. Introduce the lattice $\delta\mathbb{Z}^2$. Fix an integer N and consider all the self-avoiding nearest neighbor walks that start at the origin and stay inside D . We put a probability measure on this finite set by making these walks equally probable. Now we take two limits. First we send $N \rightarrow \infty$. This should give a probability measure on infinite self-avoiding walks that stay in D . Everyone believes this limit exists, but this has only been proved for the half-plane. Next we let $\delta \rightarrow 0$. This should give a probability measure on continuous curves in D that start at 0 and presumably “end” at ∞ . This defines the scaling limit of the self-avoiding walk in an unbounded domain between a boundary point and ∞ . There are no rigorous results on the existence of this limit. Note that ∞ can be thought of as a boundary point for an unbounded domain. So this defines a chordal process.

Now consider a bounded domain D and let z and w be points on its boundary. We want to define the SAW in D between these two points. Again, we introduce the lattice $\delta\mathbb{Z}^2$. We take all self-avoiding walks that go from z to w and stay in D . We now consider walks with any number of steps. Define a probability measure on this finite set by requiring the probability of a walk with N steps to be proportional to β^{-N} where β is a parameter. Note that if we do this without the self-avoiding constraint and take $\beta = 4$ (for the square lattice), then we get the random walk in D starting at z and conditioned to exit D at w . This suggests that for the self-avoiding walk we should choose β as follows. The number of self-avoiding walks of length N is believed to grow like β_c^N for some constant β_c . (The precise claim is that if $C(N)$ is the number of walks, then $\lim_{N \rightarrow \infty} \ln(C(N))/N$ exists. β_c is the exponential of this limit.) We should take $\beta = \beta_c$. (This constant is dependent on the lattice.) To construct the scaling limit of the self-avoiding walk we let $\delta \rightarrow 0$. (The definition of the scaling limit for a bounded domain is rather different than for an unbounded domain. In particular it only involves a single limit.)

In percolation we get critical behavior only when the parameter p is equal to a specific value. Similarly, we get critical behavior for the self-avoiding walk in a bounded domain only if $\beta = \beta_c$. What happens if $\beta \neq \beta_c$? My guess is that for $\beta < \beta_c$, the scaling limit will give a curve that is just a straight line from z to w (if such a line lies in D .) For $\beta > \beta_c$, I have no idea. Maybe the scaling limit does not exist.

The self-avoiding walk has an important property called the **restriction property**. At first glance it is easy to confuse with the locality property we saw before.

Theorem 11 (*Restriction property for the self-avoiding walk*) *Let z, w be sites in the lattice δZ^2 . Let D and D' be simply connected domains whose boundaries have non-empty overlap and such that $D' \subset D$. and such that $z, w \in \partial D \cap \partial D'$. Then the self-avoiding walk in D between w and z conditioned on the event that the walk stays in D' has the same distribution as the self-avoiding walk in D' between w and z .*

The proof of the theorem is completely trivial. The consequences are far from trivial. This property should hold in the scaling limit as well, assuming it exists. We will prove that SLE_κ has this restriction property only if $\kappa = 8/3$. So if the self-avoiding has a scaling limit that is SLE_κ , then κ must be $8/3$.

For the experts we note that the weakly self-avoiding walk also has the restriction property.

In the above we have defined scaling limits that give chordal processes - the self-avoiding walks go between two boundary points. One can also define radial processes in which the self-avoiding walk goes between an interior point and a boundary point.

3.4 Ising model

This model was suggested to Ising by his thesis adviser, Lenz. Ising solved the one-dimensional model, ..., and on the basis of the fact that the one-dimensional model had no phase transition, he asserted that there was no phase transition in any dimension. As we shall see, this is false. It is ironic that on the basis of an elementary calculation and erroneous conclusion, Ising's name has become among the most commonly mentioned in the theoretical physics literature. But history has had its revenge. Ising's name, which is correctly pronounced "E-zing," is almost universally mispronounced "I-zing."

Barry Simon

Like all our discrete models, the Ising model can be defined on a variety of lattices. I will use the hexagonal lattice as we did for percolation since it simplifies the discussion of interfaces. Let Λ be a large but bounded set made up of hexagons. We use variables like i and j to denote hexagons. For each hexagon i we have a variable (the spin) σ_i which only takes on the values $+1$ and -1 . (Normally one assigns these spins to the sites in the lattice. What I am doing is equivalent to the usual definition of the Ising model on a triangular lattice.) We use σ to denote the set of σ_i for all hexagons $i \in \Lambda$. We will refer to a choice of σ as a spin configuration or just a configuration. There are a finite number of such configuration, although the number grows exponentially with the number of hexagons in Λ . We are going to define a probability measure μ on this set of configurations. It depends on an energy function or "Hamiltonian" H . We let $\partial\Lambda$ denote the hexagons which are outside Λ but share an edge with a hexagon in Λ . We let σ^∂

denote an assignment of a spin σ_i to each hexagon $i \in \partial\Lambda$. Then we define

$$H(\sigma, \sigma^\partial) = - \sum_{\langle i,j \rangle: i,j \in \Lambda} \sigma_i \sigma_j - \sum_{\langle i,j \rangle: i \in \Lambda, j \in \partial\Lambda} \sigma_i \sigma_j^\partial \quad (55)$$

The first sum is over all pairs of hexagons in Λ which share an edge. The second sum is over pairs of hexagons which share an edge and have one hexagon in Λ and the other outside it.

Given a choice of boundary spins, σ^∂ , we define a probability measure by

$$\mu(\sigma) = \exp(-\beta H(\sigma, \sigma^\partial)) / Z \quad (56)$$

where Z is a constant chosen to make this a probability measure. Explicitly,

$$Z = \sum_{\sigma} \exp(-\beta H(\sigma, \sigma^\partial)) \quad (57)$$

where the sum is over all the spin configurations on Λ . Both Z and the probability measure depend on the boundary spins, but we have not made this explicit.

Two important choices of boundary spins are to take all the boundary spins to be $+1$ or take them all to be -1 . These are typically called $+$ boundary conditions and $-$ boundary conditions. We let μ^+ and μ^- denote the probability measures we get using these two boundary conditions.

Having defined a probability measure we can compute expectations with respect to it. Of course, since the probability space is finite these expectations are nothing more than finite sums. Physicists denote this expectation or sum by $\langle \ \rangle$. For example,

$$\langle \sigma_0 \rangle = \frac{1}{Z} \sum_{\sigma} \sigma_0 \exp(-\beta H(\sigma)) \quad (58)$$

but we will continue to use E for expectations and write the above as $E\sigma_0$. We will use E^+ or E^- to denote expectations with respect to the probability measure using $+$ or $-$ boundary conditions.

A very important question is how much $E\sigma_0$ depends on the choice of the boundary conditions. If Λ is very large, does the spin at the origin feel the effect of the boundary spins? The answer depends on β . Intuitively, one expects that if β is very small (very high temperature) then the spins are almost independent and the effect of the boundary spins should decay as we move away from the boundary. If β is very large, the weighting factor $e^{-\beta H}$ heavily favors configurations in which most of the adjacent spins have the same sign. So with $+$ boundary conditions we might expect it to favor configurations which are predominantly $+1$. What is not obvious intuitively is that there is a sharp transition between these two cases.

Theorem 12 *If the number of dimensions is at least two, then there is a positive number β_c such that for $\beta < \beta_c$ the limit $\lim_{\Lambda \rightarrow \infty} E^+ \sigma_0$ exists and is zero while for $\beta > \beta_c$ the limit exists and is strictly greater than zero.*

One should say exactly what $\Lambda \rightarrow \infty$ means. Intuitively, it means that Λ converges to all of the plane, but I won't be more precise than that. Note that symmetry implies that $E^+ \sigma_0 = -E^- \sigma_0$.

Now consider how we can force an interface into our model. We do this by introducing boundary conditions just as we did for percolation. We take two sites z, w in the hexagonal lattice which are on the boundary of Λ . These sites divide the boundary hexagons into two "arcs." We take the boundary spins to be $+1$ in one arc and -1 in the other arc. Just as for percolation, there will then be a unique path between z and w such that as we traverse the path there is always a $+$ spin on the left and a $-$ spin on the right.

Now we consider the scaling limit. We take a domain D and two points z, w in its boundary. We introduce a hexagonal lattice of spacing δ . We let Λ be the hexagons in D and let z_δ and w_δ be lattice sites in the boundary of Λ closest to w and z . The above construction then gives a probability measure on continuous paths from z_δ to w_δ along bonds in the hexagonal lattice which stay in D . The scaling limit is given by letting $\delta \rightarrow 0$. We expect that this limit will give a probability measure on continuous paths from z to w in D .

If the Ising model is not critical, i.e., $\beta \neq \beta_c$ then we expect that the interface will converge to a straight line from z to w . (I am not sure about this!) At the critical point, we should get a random curve between the two boundary points. I don't understand how this random curve should be related to SLE. Bauer, Bernard and Houdayer have introduced something they call "dipolar SLE" and conjectured that it gives the scaling limit of the interface. They have some modest simulations to test their conjecture. Their preprint is at <http://arxiv.org/abs/math-ph/0411038>.

We can define the Ising model on other lattices. It is believed that you should get the same scaling limit. Note that if we use a different lattice, for example assign spins to the squares in a square lattice, then there is some ambiguity in the definition of the interface. I am not sure what you do about this. One can also consider variations on the above simple nearest neighbor Hamiltonian. For example, one could add a term that couples sites that are two bonds apart. There are variations on the Hamiltonian that will "change the universality class," i.e., produce a different scaling limit. But a large class of Hamiltonians should have the same scaling limit as the one we defined above.

3.5 FK percolation and Potts models

There is a beautiful representation of the Ising model that relates it to a correlated form of percolation discovered by Fortuin and Kasteleyn. This correlated percolation is often called a random cluster model. One important feature of this FK representation is that

it is the basis for the fastest algorithm for Monte Carlo simulations of the Ising model : the Swensen-Wang algorithm. We give a glimpse of this representation. For more details on the FK representation see Grimmett's St. Flour lectures:

www.statslab.cam.ac.uk/~grg/papers/USstflour.ps

Before we used a hexagonal lattice and associated a spin with each hexagon. If you put a site at the center of each hexagon and draw a bond between sites whose hexagons share an edge, then you get the triangular lattice. So what we did was equivalent to using the triangular lattice and assigning spins to each site in this lattice. This is the standard way of defining the Ising model and we will now use it. So i and j will denote sites in the lattice.

Let $\delta_{\sigma_i, \sigma_j}$ be the function that is 1 if $\sigma_i = \sigma_j$ and 0 otherwise. Then

$$\exp(\beta\sigma_i\sigma_j) = e^{\beta}[e^{-2\beta} + \delta_{\sigma_i, \sigma_j}(1 - e^{-2\beta})] \quad (59)$$

$$= e^{\beta}[(1 - p) + \delta_{\sigma_i, \sigma_j}p] \quad (60)$$

where $p = 1 - e^{-2\beta}$. Thus the partition function can be written as

$$Z = \sum_{\sigma} \exp(\beta\sigma_i\sigma_j) = e^{\beta|E|} \sum_{\sigma} \prod_{\langle i, j \rangle \in E} [(1 - p) + \delta_{\sigma_i, \sigma_j}p] \quad (61)$$

where E is the set of bonds in the lattice and $|E|$ is the number of such bonds. We now expand out the product over bonds in E . We represent a term in the expansion by letting B be the set of bonds for which we take the factor $\delta_{\sigma_i, \sigma_j}p$ rather than the factor $(1 - p)$. So

$$Z = e^{\beta|E|} \sum_{\sigma} \sum_{B \subset E} p^{|B|} (1 - p)^{|B^c|} \prod_{\langle i, j \rangle \in B} \delta_{\sigma_i, \sigma_j} \quad (62)$$

We interchange the sums over σ and B . Consider

$$\sum_{\sigma} \prod_{\langle i, j \rangle \in B} \delta_{\sigma_i, \sigma_j} \quad (63)$$

The delta functions force the spins to be constant on each connected component or cluster of B . For each cluster the spin has two possible values. So the sum over σ equals $2^{C(B)}$ where $C(B)$ denotes the number of clusters in B . So

$$Z = e^{\beta|E|} \sum_{B \subset E} p^{|B|} (1 - p)^{|B^c|} q^{C(B)} \quad (64)$$

where $q = 2$ for the Ising model.

In the above derivation q turns out to be 2. But we can now forget about the Ising model and use (64) as the definition of a model of bond percolation with q as a parameter in addition to the parameter p . If we take $q = 1$ in (64) we obtain ordinary bond percolation. For $q = 3, 4, \dots$ the random cluster model is closely related to the q state

Potts model. This model is similar to the Ising model but with q possible values for the spins instead of just 2. In the random cluster model q need not be an integer.

The connection between the Ising/Potts model and the FK random cluster model is deeper than the above relation between partition functions. Correlation functions in the Ising/Potts model are related to connectivity functions in the random cluster model. See the Grimmett article above for details.

The random cluster model depends on two parameters p and q . For a given value of q there will be a value of p which makes it critical. There is a way to introduce an interface in the model. This is bit involved. A brief discussion is in the Kager and Nienhuis article. It is conjectured that for $0 \leq \kappa < 4$, the scaling limit of this interface is SLE_κ where κ is given by

$$\cos\left(\frac{4\pi}{\kappa}\right) = \frac{-\sqrt{q}}{2} \tag{65}$$

3.6 Conformal invariance

Definition 9 *By a chordal processes we will mean a collection of probability measures, one for each choice of (D, z, w) where D is a simply connected domain (not equal to the entire plane), and z and w are points on its boundary. The probability measures are defined on the set of continuous curves from z to w that remain inside the closure of D .*

Definition 10 *By a radial processes we will mean a collection of probability measures, one for each choice of (D, z, w) where D is a simply connected domain (not equal to the entire plane), z is a point on its boundary, and w a point in its interior. The probability measures are defined on the set of continuous curves from z to w that remain inside the closure of D .*

We have defined several models whose scaling limits (assuming they exist) should give chordal and/or radial processes. Let D and D' be simply connected domains (not equal to the whole plane.) A conformal map from D onto D' is a one to one, onto, differentiable function that preserves angles. Locally, a conformal map must consist of the composition of a dilation and a rotation. We will see in the next chapter that there is a close relation between conformal maps and analytic functions. Here we will just need the non-trivial fact that given (D, z, w) and (D', z', w') as above, there is there is a conformal map of D onto D' that takes z to z' and w to w' . In the radial case, the map is unique, whereas in the chordal case there is a one paramter family of such maps.

Definition 11 *A chordal or radial process is conformally invariant if the following is true. Let D, z, w and D', z', w' be two choices as in the definition. Let ϕ be a conformal map of D onto D' with $\phi(z) = z'$ and $\phi(w) = w'$. Let ω denote the random curve with respect to the D, z, w process and ω' the random curve with respect to the D', z', w' process. Then $\phi \circ \omega$ is a random curve in D' from z' to w' . Conformal invariance is the requirement that the distribution of $\phi \circ \omega$ is the same as that of ω' .*

Brownian motion gives an example of a radial process. We start a Brownian motion at z and condition it to exit D at w . A theorem of Levy says that this is a conformally invariant. (I hope to prove this later.) It is believed that all of the models we have defined are conformally invariant. For percolation and the loop-erased random walk this is a theorem. For the self-avoiding walk it is a conjecture. There is no good heuristic reason it should be true, although there is numerical evidence it is.

3.7 Markov property

The random curves we have defined by the scaling limits of our models are continuous time stochastic processes. Are any of them Markov processes? No. However, there is another more subtle type of property that these models have which plays an important role in their relation to SLE and is a type of Markov property. To define it we first make the following observations. Let D be a simply connected domain, $z \in \partial D$. Let $\gamma : [0, t] \rightarrow D$ be a simple, continuous curve starting at z which lies entirely in D . Then $D \setminus \gamma$ is a simply connected domain and the terminal point $\gamma(t)$ is a boundary point of $D \setminus \gamma$. So we can consider the process in $D \setminus \gamma$ going from $\gamma(t)$ to w .

We need to figure out what to do if γ is not simple or if the curve γ hits the boundary of D , i.e., γ is in \bar{D} but not D . If γ is not simple, then $D \setminus \gamma$ need not be simply connected because the loops in γ will form holes in the set. We can make it simply connected by deleting the regions enclosed by the loops in γ . When γ hits the boundary it is not clear what region we should remove to get our simply connected set. The process must eventually reach z , so we make the following definition. We consider the connected components of $D \setminus \gamma$. We remove all of them from $D \setminus \gamma$, except for the component containing w . We define what is left to be the ‘‘hull’’ of $D \setminus \gamma$, and denote it by $\text{hull}(D \setminus \gamma)$. (Note that if γ is simply connected and does not touch the boundary, then there is only one component and so $\text{hull}(D \setminus \gamma) = D \setminus \gamma$.)

We first give a general but somewhat vague definition of the property.

Definition 12 *Let D be a simply connected domain, z a point on its boundary and w a point on its boundary (for a chordal processes) or in its interior (for a radial process). Let $\gamma_0 : [0, t] \rightarrow D$ be a continuous curve in D beginning at z which does not hit w . We condition the chordal or radial process γ on the event that it agrees with γ_0 up to time t . (This is an event with probability zero, so some limiting procedure is needed to make sense of the conditioning.) We say the process has the Markov property if this process after time t conditioned on γ_0 has the same distribution as the process in the simply connected domain $\text{hull}(D \setminus \gamma_0)$ from $\gamma_0(t)$ to w .*

We now consider this property for our models. First consider percolation. Let γ_0 be the first n steps of some exploration process starting from z . If we condition on the event that the exploration process γ agrees with γ_0 for the first n steps, then the color of each hexagon which has an edge in γ_0 is determined by whether it is on the left or right of

γ_0 . Let Γ denote the set of hexagons which share an edge with γ_0 , and let z' be $\gamma_0(n)$. The conditioning amounts to imposing the boundary conditions we would use to define an exploration process in $D \setminus \Gamma$ using the boundary points z' and w . So the percolation exploration process has a discrete form of the Markov property.

Next we consider the self-avoiding walk in a bounded domain D between two boundary points z and w . Let γ_0 be a self avoiding walk which begins at z and ends in the interior of D . Recall that the probability of a SAW γ from z to w is $\beta^{-|\gamma|}$ where $|\gamma|$ is the number of steps in γ . Now condition on the event that γ agrees with γ_0 for the first n steps. Let γ' be the part of γ after the first n steps. (So $\gamma = \gamma_0 \cup \gamma'$.) The probability of γ is proportional to $\beta^{-|\gamma_0|}\beta^{-|\gamma'|}$. So under the conditioning the probability of γ' is proportional to $\beta^{-|\gamma'|}$. Thus the conditional measure is just the probability measure for the SAW in $D \setminus \gamma_0$ between $\gamma_0(n)$ and w .

Next we consider the loop-erased random walk. It also has a Markov property but unlike the above models the proof is far from trivial.

Theorem 13 (*Markov property of LERW*) *Consider any two dimensional lattice. Let D be a finite set of lattice sites containing w and let z be a site outside of D which is connected by a bond to a site in D . We consider the loop erased random walk γ in D from w to z . We let τ denote the number of steps in γ , so that $\gamma(\tau) = z$. Let γ_0 be a simple nearest neighbor walk with n steps which begins at z , stays inside D and does not hit y . We condition γ on the event that the last τ steps of γ are γ_0 traversed backwards. More precisely, we condition on the event*

$$\{\gamma(\tau - 1) = \gamma_0(1), \gamma(\tau - 2) = \gamma_0(2), \dots, \gamma(\tau - n) = \gamma_0(n)\} \quad (66)$$

With this conditioning, γ restricted to the indices $0, 1, \dots, \tau - n$ is a simple walk from w to $\gamma_0(n)$ which does not hit γ_0 , except for its last site. The conclusion is that under the conditioning this restricted γ has the same distribution as the loop-erased random walk from w to $\gamma_0(n)$ in $D \setminus \{\gamma_0(1), \dots, \gamma_0(n)\}$.

A proof of the theorem (with some details left to the reader) can be found in Werner's St. Flour notes. The theorem suggest that we should traverse the loop erased random walk in the opposite direction so that it goes from the boundary point to the interior point. This is a radial process. The theorem says it satisfies the Markov property.

The formulation and trivial proof of the Markov property for the interface in the Ising model we defined with spins assigned to hexagons is left to the reader.

3.8 Missing

The frontier of a two-dimension Brownian motion

The uniform spanning tree