2 Percolation

2.1 Definition of the model

Percolation can be defined on any lattice in any number of dimensions. It also comes in two flavors - bond percolation and site percolation. We will start with bond percolation and to be concrete think of the lattice $\mathbb{Z}^d$. We will denote the set of nearest neighbor bonds or edges in the lattice by $\mathbb{E}^d$.

We start with a hand-waving definition of the model and then give the precise definition. We fix a parameter $p \in [0,1]$. Each bond is the lattice is "open" with probability p and "closed" with probability $1-p$. The bonds are independent. So the model generates a random subset of $\mathbb{E}^d$.

Here is the more precise definition. The sample space is

$$\Omega = \prod_{e \in \mathbb{E}^d} \{0,1\}$$

with the convention that 1 represents open and 0 represents closed. $\mathcal{F}$ is the σ-field generated by the finite dimensional cylinders. On each $\{0,1\}$ we put the probability measure that gives 1 probability p and 0 probability $1-p$. Then we take P to be the product probability measure.

spell out what this means for an event that only depends on a finite number of bonds

We use ω to denote a point in the sample space. This corresponds to a function $\omega(e)$ from $\mathbb{E}^d$ into $\{0,1\}$. We can represent such a function by the set of edges e with $\omega(e) = 1$, which we will refer to as a bond configuration. Given such a bond configuration we let C denote the connected component which contains the origin. C can contain an infinite or finite number of bonds. Let $|C|$ denote the number of bonds in C. We define

$$\theta(p) = P(|C| = \infty)$$

So $\theta(p)$ is the probability that the origin belongs to an infinite cluster. It is trivial that

$$\theta(p) = 1 - \sum_{n=0}^{\infty} P(|C| = n)$$

Going beyond the question of whether there is an infinite cluster, we might ask what such a large scale structure looks like. We have defined the model on a lattice with unit lattice spacing. A very interesting thing to do is define it on a lattice with spacing a, e.g., $a\mathbb{Z}^d$ and studying the scaling limit or continuum limit in which $a \to 0$. We can do this in a finite volume. A particular interesting question is the following. Take a rectangle of size L by M. Put in a lattice of spacing a. Then we ask if there is a cluster of open bonds that connects the top edge to the bottom edge. Similarly, is there a cluster of open edges that connects the right edge to the left edge. In the finite volume these probabilities will lie in $(0, 1)$. The interesting question is what happens in the scaling limit.

So far we have been considering bond percolation. There is another model called site percolation. Again, we fix a parameter $p \in [0, 1]$. Each site is open with probability p and closed with probability $1-p$. Again we study connected cluster, but now the definition of connected is slightly different. A set of sites in the lattice is connected if we can get from any one site in the set to any other site in the set by a nearest neighbor path which only visits sites in the set. We refer to connected sets of sites as clusters and now C denotes the cluster containing the origin. $\theta(p)$ is defined as above, and the questions about crossings of a rectangle in the scaling limit are just as interesting.

2.2 Proof of a phase transition

It is not hard to show that $\theta(p) < 1$ if $p < 1$. (We leave this as an easy exercise.)

Proposition 1 *$\theta(p)$ is non-decreasing, $\theta(0) = 0$, and $\theta(1) = 1$.*

Proof: The fact that it is non-decreasing is intuitively obvious. To prove it we use a powerful idea known as "coupling." Explain what this is in general.

We define another probability space. We put i.i.d. random variables's X_e on the edges with each X_e uniformly distributed on $[0, 1]$. More precisely, let Ω' be the product

$$\Omega' = \prod_e [0, 1] \tag{1}$$

Now fix a p. Given an outcome in Ω', i.e., values of the X_e, we define a bond configuration $\omega_p(e)$ by declaring a bond to be open if $X_e \leq p$. This puts the same probability measure on the set of bond configurations as before.

Now suppose $p_1 < p_2$. Then $\omega_{p_1} \leq \omega_{p_2}$ is the sense that for every edge $\omega_{p_1}(e) \leq \omega_{p_2}(e)$.

If ω_{p_1} has an infinite cluster containing the origin, then so does ω_{p_2}. Hence the probability that there is an infinite cluster for p_1 is smaller than the probability there is an infinite cluster for p_2, i.e., $\theta(p_1) \leq \theta(p_2)$.

■

We define a critical value by $p_c = \sup\{p : \theta(p) = 0\}$. The monotonicity of $\theta(p)$ then implies $\theta(p) = 0$ for $p < p_c$ and $\theta(p) > 0$ for $p > p_c$. But we can't say whether or not $\theta(p_c)$ is zero. In one dimension $p_c = 1$ (exercise). In two and higher dimensions there is a phase transition, as we will now prove.

Theorem 1 *For $d \geq 2$, $0 < p_c < 1$.*

Proof: We first prove that $p_c > 0$. So we want to show that for sufficiently small p, $\theta(p) = 0$. For a finite connected set D of bonds which contains the origin, we define E_D to be the event that all the bonds in D are open. Clearly $P(E_D) = p^{|D|}$. We showed in the previous chapter that there is lattice dependent constant c such that the number of such D with n bonds is bounded by c^n. The key idea is that if the connected component containing the origin is infinite then for any n it contains a finite connected subset with n bonds that contains the origin. So for any n,

$$\{|C| = \infty\} \subset \cup_{connected\, D\, containing\, 0,\, |D|=n} E_D \tag{2}$$

So

$$P(\{|C| = \infty\}) \leq \sum_{connected\, D\, containing\, 0,\, |D|=n} P(E_D) \leq c^n p^n \tag{3}$$

If $p < 1/c$ this bound goes to 0 as $n \to \infty$. So $P(\{|C| = \infty\}) = 0$.

Now we show that $p_c < 1$. So we must show that if p is close to 1 there is a nonzero probability that there is an infinite cluster. We use a Peierls argument. We define the dual lattice as before. We then define a edge configuration on the dual lattice as follows Note that each edge in the original lattice is bisected by one edge in the dual lattice. This sets up a one to one correpondence between edges in the original lattice and edges in the dual lattice. Given an edge configuration for the original lattice we define a edge in the dual lattice to be open if and only if the edge in the original lattice that it bisects is open. Clearly the resulting configurations of edges in the dual lattice is a percolation process on the dual with the same p.

Now suppose the cluster containing the origin is finite. Then we can find a loop of closed edges in the dual lattice which encloses the origin. (It is highly nontrivial to write out an honest proof of this. We prove it by drawing pictures.)

Let γ be a loop of edges in the dual lattice, and let E_γ be the event that all the edges in γ are closed. Then the event that the cluster in the original lattice containing the origin is finite is contained in $\cup_\gamma E_\gamma$. So

$$P(|C| < \infty) \leq \sum_\gamma P(E_\gamma) = \sum_\gamma (1-p)^{|\gamma|}$$

As in the last chapter there is a lattice dependent constant c such that the number of γ containing the origin with n edges is bounded by c^n. So the above is

$$\leq \sum_{n=1}^{\infty} c^n (1-p)^n$$

If p is sufficiently close to 1 this series converges and is less than 1. So for p is sufficiently close to 1, $P(|C| = \infty) > 0$.

■

We leave it to the reader to modify these proofs for the case of site percolation. One certainly expects the critical p to depend on the lattice and there is no reason to expect it to be the same for bond and site percolation on the same lattice. For the square lattice you can prove $p_c = 1/2$ for bond percolation. For site percolation on the square lattice the value of p_c is not know exactly but is around 0.59.

Theorem 2 *The probability there is an infinite open cluster somewhere in the lattice is* 0 *if* $\theta(p) = 0$ *and* 1 *if* $\theta(p) > 0$.

Proof

review tail events and the 0-1 law

Let E be the event that there is an infinite open cluster. Note that whether or not there is an infinite open cluster does not change if we change the bond configuration on a finite number of bonds. So E is a tail event. By the 0-1 law it has probability 0 or 1.

If $\theta(p) = 0$ then the probability the origin belongs to an infinite cluster is 0. Let C_x be the event that there is an infinite cluster containing x. By translation invariance, for all x, $P(C_x) = 0$. The event that there is an infinite cluster somewhere is contained in $\cup_x C_x$. So it has probability 0 since there are a countable number of terms in this union.

Now suppose that $\theta(p) > 0$. So the probability there is an infinite cluster containing the origin is not zero. Hence the probabilty there is an infinite cluster somewhere is not zero. But this probability is 0 or 1, so it must be 1. ∎

2.3 Critical exponents, scaling limit and universality

Just as for the Ising model the behavior of percolation at and near the critical point is described by several critical exponents.

It is believed that $\theta(p)$ goes to zero as $p \to p_c$ and in the manner

$$\theta(p) \asymp (p - p_c)^\beta, p \to p_c^+ \tag{4}$$

This statement should be taken to mean that the following limit exists and is finite and nonzero:

$$\lim_{p \to p_c^+} \frac{\log(\theta(p))}{\log(p - p_c)} \tag{5}$$

If we are in the *subcritical* phase, $p < p_c$ then the cluster containing the origin is finite a.s. Define

$$\chi(p) = E[|C|] = \infty P(|C| = \infty) + \sum_{n=0}^{\infty} n P(|C| = n) \tag{6}$$

In the subcritical phase the first term on the right side is zero, but it is unclear whether the second term is finite. In fact it is, but it diverges as $p \to p_c$.

$$\chi(p) \asymp (p_c - p)^{-\gamma}, p \to p_c^- \tag{7}$$

In the supercritical phase, $p > p_c$, $\chi(p)$ is simply infinity. But we can look at a "truncated" version:

$$\chi^f(p) = E[|C|1(|C| < \infty)] = \sum_{n=0}^{\infty} nP(|C| = n) \tag{8}$$

It is also believed to behave as

$$\chi^f(p) \asymp (p - p_c)^{-\gamma}, p \to p_c^+ \tag{9}$$

with the same exponent.

In the subcritical phase we have already asserted that $P(|C| = n)$ decays fast enough that $E[|C|]$ is finite. In fact it decays exponentially fast:

$$P(|C| = n) \asymp \exp^{-\alpha(p)n} \tag{10}$$

At the critical point it decays as a power. The power is usually defined using the cumulative distribution ($|C| \geq n$ rather than $|C| = n$) :

$$P(|C| \geq n) \asymp n^{-1/\delta} \tag{11}$$

which should be taken to mean

$$-\frac{1}{\delta} = \lim n \to \infty \frac{log(P(|C| \geq n))}{\log(n)} \tag{12}$$

Instead of looking at the number of edges in the cluster containing the origin, one can look at the size of this cluster. We let $r(C)$ be the radius of the cluster, the distance to the site in the cluster farthest from the origin. Then

$$P(n \leq r(C) < \infty) \asymp n^{-1/\delta_r} \tag{13}$$

One can define a correlation length by

$$\xi^2(p) = \frac{1}{\chi^2(p)} \sum_x |x|^2 \, P(\{0 \to x\} \cap \{|C| < \infty\}) \tag{14}$$

and it should diverge as

$$\xi(p) \asymp |p - p_c|^{-\nu} \tag{15}$$

At the critical point,

$$P(0 \to x) \asymp |x|^{2-d-\eta} \tag{16}$$

There are scaling relations that are believed to hold in any number of dimensions

$$\begin{aligned} \gamma + 2\beta &= \beta(\delta + 1) \\ \gamma &= \nu(2 - \eta) \end{aligned} \tag{17}$$

and hyperscaling relations that should only hold for $2 \leq d \leq 6$.

$$\begin{aligned} d\delta_r &= \delta + 1 \\ 2 - \eta &= d\frac{\delta - 1}{\delta + 1} \end{aligned} \tag{18}$$

Theorem 3 *For site percolation on the triangular lattice,*

$$\begin{aligned} &\beta = 5/36, \quad \gamma = 43/18, \quad \nu = 4/3 \\ &\eta = 5/24, \quad \delta_r = 5/48, \quad \delta = 91/5 \end{aligned} \tag{19}$$

Conjecture 1 *(Universality) For a given dimension, the critical exponents are the same for all lattices and for bond or sites percolation*

Universality is much more that just the fact that critical exponents do not depend on the lattice or site vs. bond. It is expected that the scaling limit of the model should also be independent of the lattice and of site vs. bond. What is meant by the scaling limit? Crudely speaking we take the lattice spacing to be a parameter δ and take the scaling limit by letting $\delta \to 0$. But what exactly are we taking the limit of ? This is a hard question, but there are some quantities for which we can make sense of this. We focus on two of them - crossing probabilities and the "percolation exploration process."

We take a rectangle with width L and height M, and introduce a lattice with spacing δ. We consider the probability that there is a crossing between the top and bottom sides. For $0 < p < 1$ this probability is always strictly between 0 and 1. The interesting question is what happens when $a \to 0$. If $p < p_c$ it can be proved that the probability goes to 0, and if $p > p_c$ it can be proved that it goes to 1. At $p = p_c$ it converes to a value strictly between 0 and 1. In fact there is an explicit formula for this probability in terms of L/M. (It is left as an exercise to show that it can only depend on L and M through their ratio L/M.) This formula was first found by John Cardy using conformal field theory and then proved by Stanislav Smirnov for site percolation on the triangular lattice. The formula should be true for bond or site percolation on any two dimensonal lattice, but only this one case has been proved.

In fact Cardy's formula is much more general. Take any domain and take two "arcs" along its boundary. Cardy's formula gives the probability there is a crossing between the two boundary arcs.

Gap - Explain conformal invariance and how the crossing probabilities are conformally invariant.

Universality of the crossing probabilities is the statement that you get the same crossing probabilities in a given dimension for any choice of lattice and choice of sites vs. bond. Cardy's formula is believed to hold in two dimensions for any lattice and site or bond percolation, but it has only been proved for site precolation on the triangular lattice. Here the universality conjecture is not just that a few numbers are the same but rather that a *function* is universal!

The second object that is part of the scaling limit that we will consider is a random curve that can be defined in percolation which is called the percolation exploration process. This is the same as a domain wall. There have been spectacular results on the scaling limit of this random process using a process known as Schramm-Loewner Evolution (SLE). Stanislav Smirnov was awarded the Fields medal in 2010 for his work proving that SLE is the scaling limit of site percolation on the triangular lattice.

Consider site percolation on the triangular lattice. The dual lattice to the triangular lattice is the hexagonal lattice, and there is an equivalent way to define site percolation on the triangular lattice 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.

A remark on connectedness on the hexagonal lattice. If we did something similar for the square lattice we would need to decide whether two square that share a vertex but not an edge should be defined to be connected. On the hexagonal lattice we do not have this issue. The coordination number of the hexagonal lattice is three. At each site three hexagons have that site as a vertex. It follows that if two hexagons share a lattice site, then they must share an edge.

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" througout the lattice. In fact there is a critical p_c such that for $p < p_c$ the gray clusters are all finite and there is an infinite white cluster. And for $p > p_c$ the white clusters are all finite and there is an infinite gray cluster.

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.

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 Nth 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 interface. They are typically small. So now the interface will follow the boundary gray hexagons with occasional excusions of a few hexagons away from the boundary. See figure 1. As $p \to p_c$, 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, $\delta \to 0$, 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 $\delta \to 0$ 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 particularilarly 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 percolation gives a probability measure on continuous curves that go between z and w and stay inside D.

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 4 *(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 $\gamma(1), \gamma(2), \cdots \gamma(m)$ all lie in D', then $\gamma(1), \gamma(2), \cdots \gamma(m)$ and $\gamma'(1), \gamma'(2), \cdots \gamma'(m)$ have the same joint distribution.*

Another equivalent way to define the explorer(i.e. the interface) is that it is a myopic self-avoiding walk. At each step looks at its three neighbours (on the honeycomb lattice) and chooes at random one of the sites that it has no visited yet (there are one or two such sites since one site is anyway forbidden because it was the previous location of the walk). A priori it looks like this walk could get stuck - but the fact that we are doing this on the hexagonal lattice prevents such trapping. This way of thinking about the explorer process shows us how to simulate it in an efficient way. The slow way would be to generate a random number for every hexagon to determine if the hexagon is white or gray and then find the interface. Just generating the coloring take a time $O(L^2)$ if the rectangle is an L by L square. Using the myopic self-avoiding walk point of view we see that we can generate the explorer in time $O(L)$.

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.