

Part III — Percolation and Random Walks on Graphs

Based on lectures by P. Sousi

Notes taken by Dexter Chua

Michaelmas 2017

These notes are not endorsed by the lecturers, and I have modified them (often significantly) after lectures. They are nowhere near accurate representations of what was actually lectured, and in particular, all errors are almost surely mine.

A phase transition means that a system undergoes a radical change when a continuous parameter passes through a critical value. We encounter such a transition every day when we boil water. The simplest mathematical model for phase transition is percolation. Percolation has a reputation as a source of beautiful mathematical problems that are simple to state but seem to require new techniques for a solution, and a number of such problems remain very much alive. Amongst connections of topical importance are the relationships to so-called Schramm–Loewner evolutions (SLE), and to other models from statistical physics. The basic theory of percolation will be described in this course with some emphasis on areas for future development.

Our other major topic includes random walks on graphs and their intimate connection to electrical networks; the resulting discrete potential theory has strong connections with classical potential theory. We will develop tools to determine transience and recurrence of random walks on infinite graphs. Other topics include the study of spanning trees of connected graphs. We will present two remarkable algorithms to generate a uniform spanning tree (UST) in a finite graph G via random walks, one due to Aldous–Broder and another due to Wilson. These algorithms can be used to prove an important property of uniform spanning trees discovered by Kirchhoff in the 19th century: the probability that an edge is contained in the UST of G , equals the effective resistance between the endpoints of that edge.

Pre-requisites

There are no essential pre-requisites beyond probability and analysis at undergraduate levels, but a familiarity with the measure-theoretic basis of probability will be helpful.


Contents

0	Introduction	3
1	Percolation	5
1.1	The critical probability	5
1.2	Correlation inequalities	12
1.3	Two dimensions	22
1.4	Conformal invariance and SLE in $d = 2$	31
2	Random walks	33
2.1	Random walks in finite graphs	33
2.2	Infinite graphs	41
3	Uniform spanning trees	44
3.1	Finite uniform spanning trees	44
3.2	Infinite uniform spanning trees and forests	48
	Index	51


0 Introduction

This course is naturally divided into two parts — percolation, and random walks on graphs. Percolation is one of the simplest models that experience phase transition — an abrupt change in quantitative feature due to a continuous change of a parameter. More sophisticated examples of percolation the boiling of water and the loss of long-range correlation in magnets when temperature increases.

But we are not physicists. So let's talk about percolation. For reasons that become clear later, consider an $n \times (n + 1)$ lattice connected by edges:




We now fix some $p \in [0, 1]$, and for each edge in the graph, we either keep it or remove it with probability p . There are many questions we may ask. For example, we may ask for the probability that there is a left-to-right crossing of open edges.



For example, we have $f_n(0) = 0$ and $f_n(1) = 1$.

An interesting choice of p to consider is $p = \frac{1}{2}$. We can argue that $f_n(\frac{1}{2})$ must be $\frac{1}{2}$, by symmetry. More precisely, consider the dual lattice:



Note that this lattice is isomorphic to the original lattice we were thinking about, by applying a rotation. Now each edge in the dual lattice crosses exactly one edge in the original lattice, and we can set the edge to be open iff the corresponding edge in the original lattice is open. This gives rise to a percolation on the dual lattice with $p = \frac{1}{2}$ as well.

Now notice that that there is a left-right crossing of open edges in the original lattice iff there is no top-bottom crossing in the dual graph. But since the dual and original lattices are isomorphic, it follows that the probability that there is a left-right crossing in the original lattice is equal to the probability that there is *no* left-right crossing in the original lattice. So both of these must be $\frac{1}{2}$.

The ability to talk about the dual graph is a very important property that is only true in 2 dimensions. In general, there are many things known for 2 dimensions via the dual, which do not generalize to higher dimensions.

The other topic we are going to discuss is random walks in graphs. In IB Markov chains, and maybe IA Probability, we considered random walks on the integer lattice \mathbb{Z}^d . Here we shall consider random walks on *any graph*. We shall mostly think about finite graphs, but we will also talk about how certain results can be extended to infinite graphs. It turns out a rather useful way of thinking about this is to think of the graph as representing an electrical network. Then many concepts familiar from high school physics translate to interesting properties about the graph, and importing well-known (and elementary) results from electrical networks helps us understand graphs better.

1 Percolation

1.1 The critical probability

There are two models of percolation — *bond percolation* and *site percolation*. In this course, we will focus on bond percolation, but we will look at site percolation in the example sheets.

The very basic set up of percolation theory involves picking a *graph* $G = (V, E)$, where V is the set of *vertices* and E is the set of *edges*. We also pick a *percolation probability* $p \in [0, 1]$. For each edge $e \in E$, we keep it with probability p and throw it with probability $1 - p$. In the first case, we say the edge is *open*, and in the latter, we say it is *closed*.

More precisely, we define the probability space to be $\Omega = \{0, 1\}^E$, where 0 denotes a closed edge and 1 denotes an open one (in the case of site percolation, we have $\Omega = \{0, 1\}^V$). We endow Ω with the σ -algebra generated by *cylinder sets*

$$\{\omega \in \Omega : \omega(e) = x_e \text{ for all } e \in A\},$$

where A is a finite set and $x_e \in \{0, 1\}$ for all e . In other words, this is the product σ -algebra. As probability measure, we take the product measure \mathbb{P}_p , i.e. every edge is 1 with probability p and 0 with probability $1 - p$. We will write $\eta_p \in \{0, 1\}^E$ for the state of the system.

Now what can we say about the graph resulting from this process? One question we may ask is whether we can connect two points in the graphs via the edges that remain. To further the discussion, we introduce some notation.

Notation. We write $x \leftrightarrow y$ if there is an open path of edges from x to y .

Notation. We write $\mathcal{C}(x) = \{y \in V : y \leftrightarrow x\}$, the *cluster* of x .

Notation. We write $x \leftrightarrow \infty$ if $|\mathcal{C}(x)| = \infty$.

From now on, we shall take $G = \mathbb{L}^d = (\mathbb{Z}^d, E(\mathbb{Z}^d))$, the d -dimensional integer lattice. Then by translation invariance, $|\mathcal{C}(x)|$ has the same distribution as $|\mathcal{C}(0)|$ for all x . We now introduce a key piece of notation:

Definition ($\theta(p)$). We define $\theta(p) = \mathbb{P}_p(|\mathcal{C}(0)| = \infty)$.

Most of the questions we ask surround this $\theta(p)$. We first make the most elementary observations:

Example. $\theta(0) = 0$ and $\theta(1) = 1$.

A natural question to ask is then if we can find $p \in (0, 1)$ such that $\theta(p) > 0$. But even before answering that question, we can ask a more elementary one — is θ an increasing function of p ?

Intuitively, it must be. And we can prove it. The proof strategy is known as *coupling*. We have already seen coupling in IB Markov Chains, where we used it to prove the convergence to the invariant distribution under suitable conditions. Here we are going to couple all percolation processes for different values of P .

Lemma. θ is an increasing function of p .

Proof. We let $(U(e))_{e \in E(\mathbb{Z}^d)}$ be iid $U[0, 1]$ random variables. For each $p \in [0, 1]$, we define

$$\eta_p(e) = \begin{cases} 1 & U(e) \leq p \\ 0 & \text{otherwise} \end{cases}$$

Then $\mathbb{P}(\eta_p(e) = 1) = \mathbb{P}(U(e) < p) = p$. Since the $U(e)$ are independent, so are η_p . Thus η_p has the law of bond percolation with probability p .

Moreover, if $p \leq q$, then $\eta_p(e) \leq \eta_q(e)$. So the result follows. \square

Note that this is not only useful as a theoretical tool. If we want to simulate percolation with different probabilities p , we can simply generate a set of $U[0, 1]$ variables, and use it to produce a percolation for all p .

If we wish, we can provide an abstract definition of what coupling is, but the detailed definition is not of much practical use:

Definition (Coupling). Let μ and ν be two probability measures on (potentially) different probability spaces. A *coupling* is a pair of random variables (X, Y) defined on the same probability space such that the marginal distribution of X is μ and the marginal distribution of Y is ν .

With the lemma, we can make the definition

Definition (Critical probability). We define $p_c(d) = \sup\{p \in [0, 1] : \theta(p) = 0\}$.

Recall we initially asked whether $\theta(p)$ can be non-zero for $p \in (0, 1)$. We can now rephrase and strengthen this question by asking for the value of $p_c(d)$. There are a lot more questions we can ask about p_c and $\theta(p)$.

For example, we know that $\theta(p)$ is a C^∞ function on $(p_c, 1]$. However, we do not know if θ is continuous at p_c in $d = 3$. We will see soon that $p_c = \frac{1}{2}$ in $d = 2$, but the exact value of p_c is not known in higher dimensions.

Let's start actually proving things about p_c . We previously noted that

Proposition. $p_c(1) = 1$.

The first actually interesting theorem is the following:

Theorem. For all $d \geq 2$, we have $p_c(d) \in (0, 1)$.

We shall break this up into two natural parts:

Lemma. For $d \geq 2$, $p_c(d) > 0$.

Proof. Write Σ_n for the number of open self-avoiding paths of length n starting at 0. We then note that

$$\mathbb{P}_p(|\mathcal{C}(0)| = \infty) = \mathbb{P}_p(\forall n \geq 1 : \Sigma_n \geq 1) = \lim_{n \rightarrow \infty} \mathbb{P}_p(\Sigma_n \geq 1) \leq \lim_{n \rightarrow \infty} \mathbb{E}_p[\Sigma_n].$$

We can now compute $\mathbb{E}_p[\Sigma_n]$. The point is that expectation is linear, which makes this much easier to compute. We let σ_n be the number of self-avoiding paths of length n from 0. Then we simply have

$$\mathbb{E}_p[\Sigma_n] = \sigma_n p^n.$$

We can bound σ_n by $2d \cdot (2d - 1)^{n-1}$, since we have $2d$ choices of the first step, and at most $2d - 1$ choices in each subsequent step. So we have

$$\mathbb{E}_p[\Sigma_n] \leq 2d(2d - 1)^{n-1}p^n = \frac{2d}{2d - 1}(p(2d - 1))^n.$$

So if $p(2d - 1) < 1$, then $\theta(p) = 0$. So we know that

$$p_c(d) \geq \frac{1}{2d - 1}. \quad \square$$

Before we move on to the other half of the theorem, we talk a bit more about *self-avoiding paths*.

Definition (σ_n). We write σ_n for the number of self-avoiding paths of length n starting from 0.

In the proof, we used the rather crude bound

$$\sigma_n \leq 2d \cdot (2d - 1)^{n-1}.$$

More generally, we can make the following bound:

Lemma. We have $\sigma_{n+m} \leq \sigma_n \sigma_m$.

Proof. A self-avoiding path of length $n + m$ can be written as a concatenation of self-avoiding paths of length n starting from 0 and another one of length m . \square

Taking the logarithm, we know that $\log \sigma_n$ is a *subadditive sequence*. It turns out this property alone is already quite useful. It is an exercise in analysis to prove the following lemma:

Lemma (Fekete's lemma). If (a_n) is a subadditive sequence of real numbers, then

$$\lim_{n \rightarrow \infty} \frac{a_n}{n} = \inf \left\{ \frac{a_k}{k} : k \geq 1 \right\} \in [-\infty, \infty).$$

In particular, the limit exists.

This allows us to define

Definition (λ and κ). We define

$$\lambda = \lim_{n \rightarrow \infty} \frac{\log \sigma_n}{n}, \quad \kappa = e^\lambda.$$

κ is known as the *connective constant*.

Then by definition, we have

$$\sigma_n = e^{n\lambda(1+o(1))} = \kappa^{n+o(n)}.$$

as $n \rightarrow \infty$.

Thus, asymptotically, the growth rate of σ_n is determined by κ . It is natural to then seek for the value of κ , but unfortunately we don't know the value of κ for the Euclidean lattice. A bit more on the positive side, the value for the hexagonal lattice has been found recently:

Theorem (Duminil-Copin, Smirnov, 2010). The hexagonal lattice has

$$\kappa_{\text{hex}} = \sqrt{2 + \sqrt{2}}.$$

We might want to be a bit more precise about how σ_n grows. For $d \geq 5$, we have the following theorem:

Theorem (Hara and Slade, 1991). For $d \geq 5$, there exists a constant A such that

$$\sigma_n = A\kappa^n(1 + O(n^{-\varepsilon}))$$

for any $\varepsilon < \frac{1}{2}$.

We don't really know what happens when $d < 5$, but we have the following conjecture:

Conjecture.

$$\sigma_n \approx \begin{cases} n^{11/32}\kappa^n & d = 2 \\ n^\gamma\kappa^n & d = 3 \\ (\log n)^{1/4}\kappa^n & d = 4 \end{cases}$$

One can also instead try to bound σ_n from above. We have the following classic theorem:

Theorem (Hammersley and Welsh, 1962). For all $d \geq 2$, we have

$$\sigma_n \leq C\kappa^n \exp(c'\sqrt{n})$$

for some constants C and c' .

In fact, a better bound was recently found:

Theorem (Hutchcroft, 2017). For $d \geq 2$, we have

$$\sigma_n \leq C\kappa^n \exp(o(\sqrt{n})).$$

This will be proved in the example sheet.

What would be a good way to understand self-avoiding walks? Fixing an n , there are only finitely many self-avoiding walks of length n . So we can sample such a self-avoiding walk uniformly at random. In general, we would expect the total displacement of the walk to be $\sim \sqrt{n}$. Thus, what we can try to do is to take $n \rightarrow \infty$ while simultaneously shrinking space by a factor of $\frac{1}{\sqrt{n}}$. We would then hope that the result converges toward some Brownian motion-like trajectory. If we can characterize the precise behaviour of this *scaling limit*, then we might be able to say something concrete about κ and the asymptotic behaviour of σ_n .


But we don't really know what the scaling limit is. In the case $d = 2$, it is conjectured to be $SLE(\frac{8}{3})$. Curiously, it was proven by Gwynne and Miller in 2016 that if we instead looked at self-avoiding walks on a *random surface*, then the scaling limit is $SLE(\frac{8}{3})$.

That's enough of a digression. Let's finish our proof and show that $p_c(d) < 1$. A first observation is that it suffices to show this for $d = 2$. Indeed, since \mathbb{Z}^d embeds into \mathbb{Z}^{d+1} for all d , if we can find an infinite cluster in \mathbb{Z}^d , then the same is true for \mathbb{Z}^{d+1} . Thus, it is natural to restrict to the case of $d = 2$, where *duality* will be prove to be an extremely useful tool.

Definition (Planar graph). A graph G is called planar if it can be embedded on the plane in such a way that no two edges cross.

Definition (Dual graph). Let G be a planar graph (which we call the *primal graph*). We define the *dual graph* by placing a vertex in each face of G , and connecting 2 vertices if their faces share a boundary edge.

Example. The dual of \mathbb{Z}^2 is isomorphic to \mathbb{Z}^2



The dual lattice will help us prove a lot of properties for percolation in \mathbb{Z}^2 .

Lemma. $p_c(d) < 1$ for all $d \geq 2$.

Proof. It suffices to show this for $d = 2$. Suppose we perform percolation on \mathbb{Z}^2 . Then this induces a percolation on the dual lattice by declaring an edge of the dual is open if it crosses an open edge of \mathbb{Z}^2 , and closed otherwise.

Suppose $|\mathcal{C}(0)| < \infty$ in the primal lattice. Then there is a closed circuit in the dual lattice, given by the “boundary” of $\mathcal{C}(0)$. Let D_n be the number of closed dual circuits of length n that surround 0. Then the union bound plus Markov’s inequality tells us

$$\mathbb{P}_p(|\mathcal{C}(0)| < \infty) = \mathbb{P}_p(\exists n \geq D_n \geq 1) \leq \sum_{n=4}^{\infty} \mathbb{E}_p[D_n],$$

using the union bound and Markov’s inequality.

It is a simple exercise to show that

Exercise. Show that the number of dual circuits of length n that contain 0 is at most $n \cdot 4^n$.

From this, it follows that

$$\mathbb{P}_p(|\mathcal{C}(0)| < \infty) \leq \sum_{n=4}^{\infty} n \cdot 4^n (1-p)^n.$$

Thus, if p is sufficiently near 1, then $\mathbb{P}_p(|\mathcal{C}(0)| < \infty)$ is bounded away from 1. \square

By definition, if $p < p_c(d)$, then 0 is almost surely not contained in an infinite cluster. If $p > p_c(d)$, then there is a positive probability that 0 is contained in an infinite cluster. However, it is of course not necessarily the case that 0 is connected to ∞ with probability 1. In fact, there is at least probability $(1-p)^{2d}$ that 0 is not connected to ∞ , since 0 cannot be connected to ∞ if all its neighbouring edges are closed. However, it is still possible that there is some infinite cluster somewhere. It’s just that it does not contain 0.

Proposition. Let A_∞ be the event that there is an infinite cluster.

- (i) If $\theta(p) = 0$, then $\mathbb{P}_p(A_\infty) = 0$.
- (ii) If $\theta(p) > 0$, then $\mathbb{P}_p(A_\infty) = 1$.

Proof.

- (i) We have

$$\mathbb{P}_p(A_\infty) = \mathbb{P}_p(\exists x : |\mathcal{C}(x)| = \infty) \leq \sum_{x \in \mathbb{Z}^d} \mathbb{P}_p(|\mathcal{C}(x)| = \infty) = \sum \theta(p) = 0.$$

- (ii) We need to apply the *Kolmogorov 0-1 law*. Recall that if X_1, X_2, \dots are independent random variables, and $\mathcal{F}_n = \sigma(X_k : k \geq n)$, $\mathcal{F}_\infty = \bigcap_{n \geq 0} \mathcal{F}_n$. Then \mathcal{F}_∞ is trivial, i.e. for all $A \in \mathcal{F}_\infty$, $\mathbb{P}(A) \in \{0, 1\}$.

So we order the edges of \mathbb{Z}^d as e_1, e_2, \dots and denote their states

$$w(e_1), w(e_2), \dots$$

These are iid random variables. We certainly have $\mathbb{P}_p(A_\infty) \geq \theta(p) > 0$. So if we can show that $A_\infty \in \mathcal{F}_\infty$, then we are done. But this is clear, since changing the states of a finite number of edges does not affect the occurrence of A_∞ . \square

The next follow up questions is how many infinite clusters do we expect to get?

Theorem (Burton and Keane). If $p > p_c$, then there exists a unique infinite cluster with probability 1.

This proof is considerably harder than the ones we have previously done. We might think we can use the Kolmogorov 0-1 law, but we can't, since changing a finite number of edges can break up or join together infinite clusters, so the event that there are k infinite clusters for $k > 0$ is not in \mathcal{F}_∞ . However, we can exploit the fact that N is translation invariant.

Exercise. Let A be an event that is translation invariant. Then $\mathbb{P}_p(A) = 0$ or 1 almost surely.

Proof. Let N be the number of infinite clusters. Then by the lemma, we know N is constant almost surely. So there is some $k \in \mathbb{N} \cup \{\infty\}$ such that $\mathbb{P}_p(N = k) = 1$. First of all, we know that $k \neq 0$, since $\theta(p) > 0$. We shall first exclude $2 \leq k < \infty$, and then exclude $k = \infty$.

Assume that $k < \infty$. We will show that $\mathbb{P}_p(n = 1) > 0$, and hence it must be the case that $\mathbb{P}_p(n = 1) = 1$.

To bound this probability, we let $B(n) = [-n, n]^d \cap \mathbb{Z}^d$ (which we will sometimes write as B_n), and let $\partial B(n)$ be its boundary. We know that

$$\mathbb{P}_p(\text{all infinite clusters intersect } \partial B(n)) \rightarrow 1$$

as $n \rightarrow \infty$. This is since with probability 1, there are only finitely many clusters by assumption, and for each of these configurations, all infinite clusters intersect $\partial B(n)$ for sufficiently large n .

In particular, we can take n large enough such that

$$\mathbb{P}_p(\text{all infinite clusters intersect } \partial B(n)) \geq \frac{1}{2}.$$

We can then bound

$$\mathbb{P}_p(N = 1) \geq \mathbb{P}_p(\text{all infinite clusters intersect } \partial B(n) \text{ and all edges in } B(n) \text{ are open}).$$

Finally, note that the two events in there are independent, since they involve different edges. But the probability that all edges in $B(n)$ are open is just $p^{E(B(n))}$. So

$$\mathbb{P}_p(N = 1) \geq \frac{1}{2} p^{E(B(n))} > 0.$$

So we are done.

We now have to show that $k \neq \infty$. This involves the notion of a *trifurcation*. The idea is that we will show that if $k = \infty$, then the probability that a vertex is a trifurcation is positive. This implies the expected number of trifurcations is $\sim n^d$. We will then show deterministically that the number of trifurcations inside $B(n)$ must be $\leq |\partial B(n)|$, and so there are $O(n^{d-1})$ trifurcations, which is a contradiction.

We say a vertex x is a *trifurcation* if the following three conditions hold:

- (i) x is in an infinite open cluster \mathcal{C}_∞ ;
- (ii) There exist exactly three open edges adjacent to x ;
- (iii) $\mathcal{C}_\infty \setminus \{x\}$ contains exactly three infinite clusters and no finite ones.


This is clearly a translation invariant notion. So

$$\mathbb{P}_p(0 \text{ is a trifurcation}) = \mathbb{P}_p(x \text{ is a trifurcation})$$

for all $x \in \mathbb{Z}^d$.

Claim. $\mathbb{P}_p(0 \text{ is a trifurcation}) > 0$.

We need to use something slightly different from $B(n)$. We define $S(n) = \{x \in \mathbb{Z}^d : \|x\|_1 \leq n\}$.



The crucial property of this is that for any $x_1, x_2, x_3 \in \partial S(n)$, there exist three disjoint self-avoiding paths joining x_i to 0 (exercise!). For each triple x_1, x_2, x_3 , we arbitrarily pick a set of three such paths, and define the event

$$J(x_1, x_2, x_3) = \{\text{all edges on these 3 paths are open} \\ \text{and everything else inside } S(n) \text{ is closed}\}.$$

Next, for every possible infinite cluster in $\mathbb{Z}^d \setminus S(n)$ that intersects $\partial S(n)$ at at least one point, we pick a designated point of intersection arbitrarily.

Then we can bound

$$\mathbb{P}_p(0 \text{ is a trifurcation}) \geq \mathbb{P}_p(\exists \mathcal{C}_\infty^1, \mathcal{C}_\infty^2, \mathcal{C}_\infty^3 \subseteq \mathbb{Z}^d \setminus S(n) \\ \text{infinite clusters which intersect } \partial S(n) \text{ at } x_1, x_2, x_3, \text{ and } J(x_1, x_2, x_3)).$$

Rewrite the right-hand probability as

$$\mathbb{P}_p(J(x_1, x_2, x_3) \mid \exists \mathcal{C}_\infty^1, \mathcal{C}_\infty^2, \mathcal{C}_\infty^3 \subseteq \mathbb{Z} \text{ intersecting } \partial S(n)) \\ \times \mathbb{P}_p(\exists \mathcal{C}_\infty^1, \mathcal{C}_\infty^2, \mathcal{C}_\infty^3 \subseteq \mathbb{Z}^d \setminus \partial S(n))$$

We can bound the first term by

$$\min(p, 1-p)^{E(S(n))}.$$

To bound the second probability, we have already assumed that $\mathbb{P}_p(N = \infty) = 1$. So $\mathbb{P}_p(\exists \mathcal{C}_\infty^1, \mathcal{C}_\infty^2, \mathcal{C}_\infty^3 \subseteq \mathbb{Z}^d \setminus S(n) \text{ intersecting } \partial S(n)) \rightarrow 1$ as $n \rightarrow \infty$. We can then take n large enough such that the probability is $\geq \frac{1}{2}$. So we have shown that $c \equiv \mathbb{P}_p(0 \text{ is a trifurcation}) > 0$.

Using the linearity of expectation, it follows that

$$\mathbb{E}_p[\text{number of trifurcations inside } B(n)] \geq c|B(n)| \sim n^d.$$

On the other hand, we can bound the number of trifurcations in $B(n)$ by $|\partial B(n)|$. To see this, suppose x_1 is a trifurcation in $B(n)$. By definition, there exists 3 open paths to $\partial B(n)$. Fix three such paths. Let x_2 be another trifurcation. It also has 3 open paths to the $\partial B(n)$, and its paths to the boundary could intersect those of x_1 . However, they cannot create a cycle, by definition of a trifurcation. For simplicity, we add the rule that when we produce the paths for x_2 , once we intersect the path of x_1 , we continue following the path of x_1 .

Exploring all trifurcations this way, we obtain a forest inside $B(n)$, and the boundary points will be the leaves of the forest. Now the trifurcations have degree 3 in this forest. The rest is just combinatorics.

Exercise. For any tree, the number of degree 3 vertices is always less than the number of leaves. \square

1.2 Correlation inequalities

In this section, we are going to prove some useful inequalities and equalities, and use them to prove some interesting results about $\theta(p)$ and the decay of $\mathbb{P}_p(0 \leftrightarrow \partial B_n)$.

To motivate our first inequality, suppose we have 4 points x, y, u, v , and we want to ask for the conditional probability

$$\mathbb{P}_p(x \leftrightarrow y \mid u \leftrightarrow v).$$

Intuitively, we expect this to be greater than $\mathbb{P}_p(x \leftrightarrow y)$, since $u \leftrightarrow v$ tells us there are some open edges around, which is potentially helpful. The key property underlying this intuition is that having more edges is beneficial to both of the events. To quantify this, we need the notion of an increasing random variable.

Again, let $G = (V, E)$ be a graph and $\Omega = \{0, 1\}^E$. We shall assume that E is countable.

Notation (\leq). Given $\omega, \omega' \in \Omega$, we write $\omega \leq \omega'$ if $\omega(e) \leq \omega'(e)$ for all $e \in E$.

This defines a partial order on Ω .

Definition (Increasing random variable). A random variable X is increasing if $X(\omega) \leq X(\omega')$ whenever $\omega \leq \omega'$, and is decreasing if $-X$ is increasing.

Definition (Increasing event). An event A is increasing (resp. decreasing) if the indicator $1(A)$ is increasing (resp. decreasing)

Example. $\{|\mathcal{C}(0)| = \infty\}$ is an increasing event.

An immediate consequence of the definition is that

Theorem. If N is an increasing random variable and $p_1 \leq p_2$, then

$$\mathbb{E}_{p_1}[N] \leq \mathbb{E}_{p_2}[N],$$

and if an event A is increasing, then

$$\mathbb{P}_{p_1}(A) \leq \mathbb{P}_{p_2}(A).$$

Proof. Immediate from coupling. □

What we want to prove is the following result, which will be extremely useful.

Theorem (Fortuin–Kasteleyn–Ginibre (FKG) inequality). Let X and Y be increasing random variables with $\mathbb{E}_p[X^2], \mathbb{E}_p[Y^2] < \infty$. Then

$$\mathbb{E}_p[XY] \geq \mathbb{E}_p[X]\mathbb{E}_p[Y].$$

In particular, if A and B are increasing events, then

$$\mathbb{P}_p(A \cap B) \geq \mathbb{P}_p(A)\mathbb{P}_p(B).$$

Equivalently,

$$\mathbb{P}_p(A \mid B) \geq \mathbb{P}_p(A).$$

Proof. The plan is to first prove this in the case where X and Y depend on a finite number of edges, and we do this by induction. Afterwards, the desired result can be obtained by an application of the martingale convergence theorem. In fact, the “real work” happens when we prove it for X and Y depending on a single edge. Everything else follows from messing around with conditional probabilities.

If X and Y depend only on a single edge e_1 , then for any $\omega_1, \omega_2 \in \{0, 1\}$, we claim that

$$(X(\omega_1) - X(\omega_2))(Y(\omega_1) - Y(\omega_2)) \geq 0.$$

Indeed, X and Y are both increasing. So if $\omega_1 > \omega_2$, then both terms are positive; if $\omega_1 < \omega_2$, then both terms are negative, and there is nothing to do if they are equal.

In particular, we have

$$\sum_{\omega_1, \omega_2 \in \{0, 1\}} (X(\omega_1) - X(\omega_2))(Y(\omega_1) - Y(\omega_2)) \mathbb{P}_p(\omega(e_1) = \omega_1) \mathbb{P}_p(\omega(e_1) = \omega_2) \geq 0.$$

Expanding this, we find that the LHS is $2(\mathbb{E}_p[XY] - \mathbb{E}_p[X]\mathbb{E}_p[Y])$, and so we are done.

Now suppose the claim holds for X and Y that depend on $n < k$ edges. We shall prove the result when they depend on k edges e_1, \dots, e_k . We have

$$\mathbb{E}_p[XY] = \mathbb{E}_p[\mathbb{E}_p[XY \mid \omega(e_1), \dots, \omega(e_{k-1})]].$$

Now after conditioning on $\omega(e_1), \dots, \omega(e_{k-1})$, the random variables X and Y become increasing random variables of $\omega(e_k)$. Applying the first step, we get

$$\begin{aligned} \mathbb{E}_p[XY \mid \omega(e_1), \dots, \omega(e_{k-1})] \\ \geq \mathbb{E}_p[X \mid \omega(e_1), \dots, \omega(e_{k-1})] \mathbb{E}_p[Y \mid \omega(e_1), \dots, \omega(e_{k-1})]. \quad (*) \end{aligned}$$

But $\mathbb{E}_p[X \mid \omega(e_1), \dots, \omega(e_{k-1})]$ is a random variable depending on the edges e_1, \dots, e_{k-1} , and moreover it is increasing. So the induction hypothesis tells us

$$\begin{aligned} \mathbb{E}_p[\mathbb{E}_p[X \mid \omega(e_1), \dots, \omega(e_{k-1})] \mathbb{E}_p[Y \mid \omega(e_1), \dots, \omega(e_{k-1})]] \\ \geq \mathbb{E}_p[\mathbb{E}_p[X \mid \omega(e_1), \dots, \omega(e_{k-1})]] \mathbb{E}_p[\mathbb{E}_p[Y \mid \omega(e_1), \dots, \omega(e_{k-1})]] \end{aligned}$$

Combining this with (the expectation of) $(*)$ then gives the desired result.

Finally, suppose X and Y depend on the states of a countable set of edges e_1, e_2, \dots . Let's define

$$\begin{aligned} X_n &= \mathbb{E}_p[X \mid \omega(e_1), \dots, \omega(e_n)] \\ Y_n &= \mathbb{E}_p[Y \mid \omega(e_1), \dots, \omega(e_n)] \end{aligned}$$

Then X_n and Y_n are martingales, and depend only on the states of only finitely many edges. So we know that

$$\mathbb{E}_p[X_n Y_n] \geq \mathbb{E}_p[X_n] \mathbb{E}_p[Y_n] = \mathbb{E}_p[X] \mathbb{E}_p[Y].$$

By the L^2 -martingale convergence theorem, $X_n \rightarrow X$, $Y_n \rightarrow Y$ in L^2 and almost surely. So taking the limit $n \rightarrow \infty$, we get

$$\mathbb{E}_p[XY] \geq \mathbb{E}_p[X] \mathbb{E}_p[Y]. \quad \square$$

What we want to consider next is the notion of disjoint occurrence. For example, we want to be able to ask the probability that there exists two disjoint paths connecting a pair of points.

To formulate this disjointness, suppose we have an event A , and let $\omega \in A$. To ask whether this occurrence of A depends only on some set $S \subseteq E$ of edges, we can look at the set

$$[\omega]_S = \{\omega' \in \Omega : \omega'(e) = \omega(e) \text{ for all } e \in S\}.$$

If $[\omega]_S \subseteq A$, then we can rightfully say this occurrence of A depends only on the edges in S . Note that this depends explicitly and importantly on ω , i.e. the “reason” A happened. For example, if $A = \{x \leftrightarrow y\}$, and $\omega \in A$, then we can take S to be the set of all edges in a chosen path from x to y in the configuration of ω . This choice will be different for different values of ω .

Using this, we can define what it means for two events to occur disjointly.

Definition (Disjoint occurrence). Let F be a set and $\Omega = \{0, 1\}^F$. If A and B are events, then the event that A and B occurs *disjointly* is

$$A \circ B = \{\omega \in \Omega : \exists S \subseteq F \text{ s.t. } [\omega]_S \subseteq A \text{ and } [\omega]_{F \setminus S} \subseteq B\}.$$

Theorem (BK inequality). Let F be a finite set and $\Omega = \{0, 1\}^F$. Let A and B be increasing events. Then

$$\mathbb{P}_p(A \circ B) \leq \mathbb{P}_p(A)\mathbb{P}_p(B).$$

This says if A and B are both events that “needs” edges to occur, then requiring that they occur disjointly is more difficult than them occurring individually.

The proof is completely magical. There exist saner proofs of the inequality, but they are rather longer.

Proof (Bollobás and Leader). We prove by induction on the size n of the set F . For $n = 0$, it is trivial.

Suppose it holds for $n - 1$. We want to show it holds for n . For $D \subseteq \{0, 1\}^F$ and $i = 0, 1$, set

$$D_i = \{(\omega_1, \dots, \omega_{n-1}) : (\omega_1, \dots, \omega_{n-1}, i) \in D\}.$$

Let $A, B \subseteq \{0, 1\}^F$, and $C = A \circ B$. We check that

$$C_0 = A_0 \circ B_0, \quad C_1 = (A_0 \circ B_1) \cup (A_1 \circ B_0).$$

Since A and B are increasing, $A_0 \subseteq A_1$ and $B_0 \subseteq B_1$, and A_i and B_i are also increasing events. So

$$\begin{aligned} C_0 &\subseteq (A_0 \circ B_1) \cap (A_1 \circ B_0) \\ C_1 &\subseteq A_1 \circ B_1. \end{aligned}$$

By the induction hypothesis, we have

$$\begin{aligned} \mathbb{P}_p(C_0) &= \mathbb{P}_p(A_0 \circ B_0) \leq \mathbb{P}_p(A_0)\mathbb{P}_p(B_0) \\ \mathbb{P}_p(C_1) &\leq \mathbb{P}_p(A_1 \circ B_1) \leq \mathbb{P}_p(A_1)\mathbb{P}_p(B_1) \\ \mathbb{P}_p(C_0) + \mathbb{P}_p(C_1) &\leq \mathbb{P}_p((A_0 \circ B_1) \cap (A_1 \circ B_0)) + \mathbb{P}_p((A_0 \circ B_1) \cup (A_1 \circ B_0)) \\ &= \mathbb{P}_p(A_0 \circ B_1) + \mathbb{P}_p(A_1 \circ B_0) \\ &\leq \mathbb{P}_p(A_0)\mathbb{P}_p(B_1) + \mathbb{P}_p(A_1)\mathbb{P}_p(B_0). \end{aligned}$$

Now note that for any D , we have

$$\mathbb{P}_p(D) = p\mathbb{P}_p(D_1) + (1-p)\mathbb{P}_p(D_0).$$

By some black magic, we multiply the first inequality by $(1-p)^2$, the second by p^2 and the third by $p(1-p)$. This gives

$$p\mathbb{P}_p(C_1) + (1-p)\mathbb{P}_p(C_0) \leq (p\mathbb{P}_p(A_1) + (1-p)\mathbb{P}_p(A_0))(p\mathbb{P}_p(B_1) + (1-p)\mathbb{P}_p(B_0)).$$

Expand and we are done. \square

It turns out the increasing hypothesis is not necessary:

Theorem (Reimer). For all events A, B depending on a finite set, we have $\mathbb{P}_p(A \circ B) \leq \mathbb{P}_p(A)\mathbb{P}_p(B)$.

But the proof is much harder, and in all the case where we want to apply this, the variables are increasing.

As an application of the BK inequality, we first prove a preliminary result about the decay of $\mathbb{P}_p(0 \leftrightarrow \partial B(n))$. To prove our result, we will need a stronger condition than $p < p_c$. Recall that we defined

$$\theta(p) = \mathbb{P}_p(|\mathcal{C}(0)| = \infty).$$

We also define

$$\chi(p) = \mathbb{E}_p[|\mathcal{C}(0)|].$$

If $\chi(p)$ is finite, then of course $\theta(p) = 0$. However, the converse need not hold.

Theorem. If $\chi(p) < \infty$, then there exists a positive constant c such that for all $n \geq 1$,

$$\mathbb{P}_p(0 \leftrightarrow \partial B(n)) \leq e^{-cn}.$$

Later, we will show that in fact this holds under the assumption that $p < p_c$. However, that requires a bit more technology, which we will develop after this proof.

The idea of the proof is that if we want a path from, say, 0 to $B(2n)$, then the path must hit a point on $\partial B(n)$. So there is a path from 0 to a point on $\partial B(n)$, and a path from that point to $\partial B(2n)$. Moreover, these two paths are disjoint, which allows us to apply the BK inequality.

Proof. Let

$$X_n = \sum_{x \in \partial B(n)} 1(0 \leftrightarrow x).$$

Now consider

$$\sum_{n=0}^{\infty} \mathbb{E}[X_n] = \sum_n \sum_{x \in \partial B(n)} \mathbb{P}_p(0 \leftrightarrow x) = \sum_{x \in \mathbb{Z}^d} \mathbb{P}_p(0 \leftrightarrow x) = \chi(p).$$

Since $\chi(p)$ is finite, we in particular have $\mathbb{E}_p[X_n] \rightarrow 0$ as $n \rightarrow \infty$. Take m large enough such that $\mathbb{E}_p[X_m] < \delta < 1$.

Now we have

$$\begin{aligned}
\mathbb{P}_p(0 \leftrightarrow \partial B(m+k)) &= \mathbb{P}_p(\exists x \in \partial B(m) : 0 \leftrightarrow x \text{ and } x \leftrightarrow \partial B(m+k) \text{ disjointly}) \\
&\leq \sum_{x \in \partial B(m)} \mathbb{P}_p(0 \leftrightarrow x) \mathbb{P}_p(x \leftrightarrow \partial B(m+k)) && \text{(BK)} \\
&\leq \sum_{x \in \partial B(m)} \mathbb{P}_p(0 \leftrightarrow x) \mathbb{P}_p(0 \leftrightarrow \partial B(k)) && \text{(trans. inv.)} \\
&\leq \mathbb{P}_p(0 \leftrightarrow B(k)) \mathbb{E}_p[X_m].
\end{aligned}$$

So for any $n > m$, write $n = qm + r$, where $r \in [0, m-1]$. Then iterating the above result, we have

$$\mathbb{P}_p(0 \leftrightarrow \partial B(n)) \leq \mathbb{P}_p(0 \leftrightarrow B(mq)) \leq \delta^q \leq \delta^{-1 + \frac{n}{m}} \leq e^{-cn}. \quad \square$$

To replace the condition with the weaker condition $p < p_c$, we need to understand how $\theta(p)$ changes with p . We know $\theta(p)$ is an increasing function in p . It would be great if it were differentiable, and even better if we could have an explicit formula for $\frac{d\theta}{dp}$.

To do so, suppose we again do coupling, and increase p by a really tiny bit. Then perhaps we would expect that exactly one of the edges switches from being closed to open. Thus, we want to know if the state of this edge is *pivotal* to the event $|\mathcal{C}(0)| = \infty$, and this should determine the rate of change of θ .

Definition (Pivotal edge). Let A be an event and ω a percolation configuration. The edge e is *pivotal* for (A, ω) if

$$1(\omega \in A) \neq 1(\omega' \in A),$$

where ω' is defined by

$$\omega'(f) = \begin{cases} \omega(f) & f \neq e \\ 1 - \omega(f) & f = e \end{cases}.$$

The event that e is pivotal for A is defined to be the set of all ω such that e is pivotal for (A, ω) .

Note that whether or not e is pivotal for (A, ω) is independent of $\omega(e)$.

Theorem (Russo's formula). Let A be an increasing event that depends on the states of a finite number of edges. Then

$$\frac{d}{dp} \mathbb{P}_p(A) = \mathbb{E}_p[N(A)],$$

where $N(A)$ is the number of pivotal edges for A .

Proof. Assume that A depends the states of m edges e_1, \dots, e_m . The idea is to let each e_i be open with probability p_i , where the $\{p_i\}$ may be distinct. We then vary the p_i one by one and see what happens.

Writing $\bar{p} = (p_1, \dots, p_m)$, we define

$$f(p_1, \dots, p_m) = \mathbb{P}_{\bar{p}}(A),$$

Now f is the sum of the probability of all configurations of $\{e_1, \dots, e - m\}$ for which A happens, and is hence a finite sum of polynomials. So in particular, it is differentiable.

We now couple all percolation process. Let $(X(e) : e \in \mathbb{L}^d)$ be iid $U[0, 1]$ random variables. For a vector $\bar{p} = (p(e) : e \in \mathbb{L}^d)$, we write

$$\eta_{\bar{p}}(e) = 1(X(e) \leq p(e)).$$

Then we have $\mathbb{P}_{\bar{p}}(A) = \mathbb{P}(\eta_{\bar{p}} \in A)$.

Fix an edge f and let $\bar{p}' = (p'(e))$ be such that $p'(e) = p(e)$ for all $e \neq f$, and $p'(f) = p(f) + \delta$ for some $\delta > 0$. Then

$$\begin{aligned} \mathbb{P}_{\bar{p}'}(A) - \mathbb{P}_{\bar{p}}(A) &= \mathbb{P}(\eta_{\bar{p}'} \in A) - \mathbb{P}(\eta_{\bar{p}} \in A) \\ &= \mathbb{P}(\eta_{\bar{p}'} \in A, \eta_{\bar{p}} \in A) + \mathbb{P}(\eta_{\bar{p}'} \in A, \eta_{\bar{p}} \notin A) - \mathbb{P}(\eta_{\bar{p}} \in A). \end{aligned}$$

But we know A is increasing, so $\mathbb{P}(\eta_{\bar{p}'} \in A, \eta_{\bar{p}} \in A) = \mathbb{P}(\eta_{\bar{p}} \in A)$. So the first and last terms cancel, and we have

$$\mathbb{P}_{\bar{p}'}(A) - \mathbb{P}_{\bar{p}}(A) = \mathbb{P}(\eta_{\bar{p}'} \in A, \eta_{\bar{p}} \notin A).$$

But we observe that we simply have

$$\mathbb{P}(\eta_{\bar{p}'} \in A, \eta_{\bar{p}} \notin A) = \delta \cdot \mathbb{P}_{\bar{p}}(f \text{ is pivotal for } A).$$

Indeed, we by definition of pivotal edges, we have

$$\mathbb{P}(\eta_{\bar{p}'} \in A, \eta_{\bar{p}} \notin A) = \mathbb{P}_{\bar{p}}(f \text{ is pivotal for } A, p(f) < X(f) \leq p(f) + \delta).$$

Since the event $\{f \text{ is pivotal for } A\}$ is independent of the state of the edge f , we obtain

$$\mathbb{P}_{\bar{p}}(f \text{ is pivotal}, p(f) < X(f) \leq p(f) + \delta) = \mathbb{P}_{\bar{p}}(f \text{ is pivotal}) \cdot \delta.$$

Therefore we have

$$\frac{\partial}{\partial p(f)} \mathbb{P}_{\bar{p}}(A) = \lim_{\delta \rightarrow 0} \frac{\mathbb{P}_{\bar{p}'}(A) - \mathbb{P}_{\bar{p}}(A)}{\delta} = \mathbb{P}_{\bar{p}}(f \text{ is pivotal for } A).$$

The desired result then follows from the chain rule:

$$\begin{aligned} \frac{d}{dp} \mathbb{P}_p(A) &= \sum_{i=1}^m \frac{\partial}{\partial p(e_i)} \mathbb{P}_{\bar{p}}(A) \Bigg|_{\bar{p}=(p, \dots, p)} \\ &= \sum_{i=1}^m \mathbb{P}_p(e_i \text{ is pivotal for } A) \\ &= \mathbb{E}_p[N(A)]. \end{aligned} \quad \square$$

If A depends on an infinite number of edges, then the best we can say is that

$$\liminf_{\delta \downarrow 0} \frac{\mathbb{P}_{p+\delta}(A) - \mathbb{P}_p(A)}{\delta} \geq \mathbb{E}_p[N(A)].$$

To see this, again set $B(n) = [-n, n]^d \cap \mathbb{Z}^d$. Define \bar{p}_n by

$$\bar{p}_n(e) = \begin{cases} p & e \notin B(n) \\ p + \delta & e \in B(n) \end{cases}.$$

Then since a is increasing, we know

$$\frac{\mathbb{P}_{p+\delta}(A) - \mathbb{P}_p(A)}{\delta} \geq \frac{\mathbb{P}_{\bar{p}_n}(A) - \mathbb{P}_p(A)}{\delta}.$$

We can then apply the previous claim, take successive differences, and take $n \rightarrow \infty$.

Corollary. Let A be an increasing event that depends on m edges. Let $p \leq q \in [0, 1]$. Then $\mathbb{P}_q(A) \leq \mathbb{P}_p(A) \left(\frac{q}{p}\right)^m$.

Proof. We know that $\{f \text{ is pivotal for } A\}$ is independent of the state of f , and so

$$\mathbb{P}_p(\omega(f) = 1, f \text{ is pivotal for } A) = p\mathbb{P}_p(f \text{ is pivotal for } A).$$

But since A is increasing, if $\omega(f) = 1$ and f is pivotal for A , then A occurs. Conversely, if f is pivotal and A occurs, then $\omega(f) = 1$.

Thus, by Russo's formula, we have

$$\begin{aligned} \frac{d}{dp}\mathbb{P}_p(A) &= \mathbb{E}_p[N(A)] \\ &= \sum_e \mathbb{P}_p(e \text{ is pivotal for } A) \\ &= \sum_e \frac{1}{p}\mathbb{P}_p(\omega(e) = 1, e \text{ is pivotal for } A) \\ &= \sum_e \frac{1}{p}\mathbb{P}_p(e \text{ is pivotal} \mid A)\mathbb{P}_p(A) \\ &= \mathbb{P}_p(A)\frac{1}{p}\mathbb{E}_p[N(A) \mid A]. \end{aligned}$$

So we have

$$\frac{\frac{d}{dp}\mathbb{P}_p(A)}{\mathbb{P}_p(A)} = \frac{1}{p}\mathbb{E}_p[N(A) \mid A].$$

Integrating, we find that

$$\log \frac{\mathbb{P}_q(A)}{\mathbb{P}_p(A)} = \int_p^q \frac{1}{u}\mathbb{E}_u[N(A) \mid A] du.$$

Bounding $\mathbb{E}_u[N(A) \mid A] \leq m$, we obtain the desired bound. \square

With Russo's formula, we can now prove the desired theorem.

Theorem. Let $d \geq 2$ and $B_n = [-n, n]^d \cap \mathbb{Z}^d$.

- (i) If $p < p_c$, then there exists a positive constant c for all $n \geq 1$, $\mathbb{P}_p(0 \leftrightarrow \partial B_n) \leq e^{-cn}$.

(ii) If $p > p_c$, then

$$\theta(p) = \mathbb{P}_p(0 \leftrightarrow \infty) \geq \frac{p - p_c}{p(1 - p_c)}.$$

This was first proved by Aizenman and Barsky who looked at the more general framework of long-range percolation. Menshikov gave an alternative proof by analyzing the geometry of pivotal edges. The proof we will see is by Duminil-Copin and Tassion. Recall that we defined

$$\chi(p) = \mathbb{E}_p[|\mathcal{C}(0)|].$$

We saw that if $\chi(p) < \infty$, then $\mathbb{P}_p(0 \leftrightarrow \partial B_n) \leq e^{-cn}$. We now see that $\chi(p) < \infty$ iff $p < p_c$.

The strategy of the proof is to define a new percolation probability \tilde{p}_c , whose definition makes it easier to prove the theorem. Once we have established the two claims, we see that (i) forces $\tilde{p}_c \leq p_c$, and (ii) forces $\tilde{p}_c \geq p_c$. So they must be equal.

Proof (Duminil-Copin and Tassion). If $S \subseteq V$ is finite, we write

$$\partial S = \{(x, y) \in E : x \in S, y \notin S\}.$$

We write $x \xrightarrow{S} y$ if there exists an open path of edges from x to y all of whose end points lie in S .

Now suppose that $0 \in S$. We define

$$\varphi_p(S) = p \sum_{(x,y) \in \partial S} \mathbb{P}_p(0 \xrightarrow{S} x).$$

Define

$$\tilde{p}_c = \sup\{p \in [0, 1] : \text{exists a finite set } S \text{ with } 0 \in S \text{ and } \varphi_p(S) < 1\}.$$

Claim. It suffices to prove (i) and (ii) with p_c replaced by \tilde{p}_c .

Indeed, from (i), if $p < \tilde{p}_c$, then $\mathbb{P}_p(0 \leftrightarrow \partial B_n) \leq e^{-cn}$. So taking the limit $n \rightarrow \infty$, we see $\theta(p) = 0$. So $\tilde{p}_c \leq p_c$. From (ii), if $p > \tilde{p}_c$, then $\theta(p) > 0$. So $p_c \leq \tilde{p}_c$. So $p_c = \tilde{p}_c$.

We now prove (i) and (ii):

- (i) Let $p < \tilde{p}_c$. Then there exists a finite set S containing 0 with $\varphi_p(S) < 1$. Since S is finite, we can pick L large enough so that $S \subseteq B_{L-1}$. We will prove that $\mathbb{P}_p(0 \leftrightarrow \partial B_{kL}) \leq (\varphi_p(S))^{k-1}$ for $k \geq 1$.

Define $\mathcal{C} = \{x \in S : 0 \xrightarrow{S} x\}$. Since $S \subseteq B_{L-1}$, we know $S \cap \partial B_{kL} = \emptyset$.

Now if we have an open path from 0 to ∂B_{kL} , we let x be the last element on the path that lies in \mathcal{C} . We can then replace the path up to x by a path that lies entirely in S , by assumption. This is then a path that lies in \mathcal{C} up to x , then takes an edge on ∂S , and then lies entirely outside of \mathcal{C}^c . Thus,

$$\mathbb{P}_p(0 \leftrightarrow \partial B_{kL}) \leq \sum_{\substack{A \subseteq S \\ 0 \in A}} \sum_{(x,y) \in \partial A} \mathbb{P}_p(0 \xrightarrow{A} x, (x,y) \text{ open}, \mathcal{C} = A, y \xrightarrow{\mathcal{C}^c} \partial B_{kL}).$$

Now observe that the events $\{\mathcal{C} = A, 0 \xrightarrow{S} x\}$, $\{(x, y) \text{ is open}\}$ and $\{y \xrightarrow{A^c} \partial B_{kL}\}$ are independent. So we obtain

$$\mathbb{P}_p(0 \leftrightarrow \partial B_{kL}) \leq \sum_{A \subseteq S, 0 \in A} \sum_{(x, y) \in \partial S} p \mathbb{P}_p(0 \xrightarrow{S} x, \mathcal{C} = A) \mathbb{P}_p(y \xrightarrow{A^c} \partial B_{kL}).$$

Since we know that $y \in B_L$, we can bound

$$\mathbb{P}_p(y \xrightarrow{A^c} \partial B_{kL}) \leq \mathbb{P}_p(0 \leftrightarrow \partial B_{(k-1)L}).$$

So we have

$$\begin{aligned} \mathbb{P}_p(0 \leftrightarrow \partial B_{kL}) &\leq p \mathbb{P}_p(0 \leftrightarrow \partial B_{(k-1)L}) \sum_{A \subseteq S, 0 \in A} \sum_{(x, y) \in \partial S} \mathbb{P}_p(0 \xrightarrow{S} x, \mathcal{C} = A) \\ &= \mathbb{P}_p(0 \leftrightarrow \partial B_{(k-1)L}) p \sum_{(x, y) \in \partial S} \mathbb{P}_p(0 \xrightarrow{S} x) \\ &= \mathbb{P}_p(0 \leftrightarrow \partial B_{(k-1)L}) \varphi_p(S). \end{aligned}$$

Iterating, we obtain the desired result.

(ii) We want to use Russo's formula. We claim that it suffices to prove that

$$\frac{d}{dp} \mathbb{P}_p(0 \leftrightarrow \partial B_n) \geq \frac{1}{p(1-p)} \inf_{S \subseteq B_n, 0 \in S} \varphi_p(S) (1 - \mathbb{P}_p(0 \leftrightarrow \partial B_n)).$$

Indeed, if $p > \tilde{p}_c$, we integrate from \tilde{p}_c to p , use in this range $\varphi_p(S) \geq 1$, and then take the limit as $n \rightarrow \infty$.

The event $\{0 \leftrightarrow \partial B_n\}$ is increasing and only depends on a finite number of edges. So we can apply Russo's formula

$$\frac{d}{dp} \mathbb{P}_p(0 \leftrightarrow \partial B_n) = \sum_{e \in B_n} \mathbb{P}_p(e \text{ is pivotal for } \{0 \leftrightarrow \partial B_n\})$$

Since being pivotal and being open/closed are independent, we can write this as

$$\begin{aligned} &= \sum_{e \in B_n} \frac{1}{1-p} \mathbb{P}_p(e \text{ is pivotal for } \{0 \leftrightarrow \partial B_n\}, e \text{ is closed}) \\ &= \sum_{e \in B_n} \frac{1}{1-p} \mathbb{P}_p(e \text{ is pivotal for } \{0 \leftrightarrow \partial B_n\}, 0 \not\leftrightarrow \partial B_n) \end{aligned}$$

Define $\mathcal{S} = \{x \in B_n : x \not\leftrightarrow \partial B_n\}$. Then $\{0 \not\leftrightarrow \partial B_n\}$ implies $0 \in \mathcal{S}$. So

$$\frac{d}{dp} \mathbb{P}_p(0 \leftrightarrow \partial B_n) = \frac{1}{1-p} \sum_{e \in B_n} \sum_{A \subseteq B_n, 0 \in A} \mathbb{P}_p(e \text{ is pivotal}, \mathcal{S} = A)$$

Given that $\mathcal{S} = A$, an edge $e = (x, y)$ is pivotal iff $e \in \partial A$ and $0 \xrightarrow{A} x$. So we know

$$\frac{d}{dp} \mathbb{P}_p(0 \leftrightarrow \partial B_n) = \frac{1}{1-p} \sum_{A \subseteq B_n, 0 \in A} \sum_{(x, y) \in \partial A} \mathbb{P}_p(0 \xrightarrow{A} x, \mathcal{S} = A).$$

Observe that $\{0 \stackrel{A}{\leftrightarrow} x\}$ and $\{\mathcal{S} = A\}$ are independent, since to determine if $\mathcal{S} = A$, we only look at the edges on the boundary of A . So the above is equal to

$$\begin{aligned} & \frac{1}{1-p} \sum_{A \subseteq B_n, 0 \in A} \sum_{(x,y) \in \partial A} \mathbb{P}_p(0 \stackrel{A}{\leftrightarrow} x) \mathbb{P}_p(\mathcal{S} = A) \\ &= \frac{1}{p(1-p)} \sum_{A \subseteq B_n, 0 \in A} \varphi_p(A) \mathbb{P}_p(\mathcal{S} = A) \\ &\geq \frac{1}{p(1-p)} \inf_{S \subseteq B_n, 0 \in S} \varphi_P(S) \mathbb{P}_p(0 \not\leftrightarrow \partial B_n), \end{aligned}$$

as desired. \square

We might ask if $\mathbb{P}_p(0 \leftrightarrow \partial B_n) \leq e^{-cn}$ is the best convergence rate when $p < p_c$, but we cannot do better, since $\mathbb{P}_p(0 \leftrightarrow \partial B_n) \geq p^n$.

Also, if $p < p_c$, then we can easily bound

$$\mathbb{P}_p(|\mathcal{C}(0)| \geq n) \leq \mathbb{P}_p(0 \leftrightarrow \partial B_{n^{1/d}}) \leq \exp(-cn^{1/d}).$$

However, this is not a good bound. In fact, $n^{1/d}$ can be replaced by n , but we will not prove it here. This tells us the largest cluster in B_n will have size of order $\log n$ with high probability.

1.3 Two dimensions

We now focus on 2 dimensions. As discussed previously, we can exploit duality to prove a lot of things specific to two dimensions. In particular, we will show that, at $p = p_c = \frac{1}{2}$, certain probabilities such as $\mathbb{P}_{\frac{1}{2}}(0 \leftrightarrow \partial B(n))$ exhibit a power law decay. This is in contrast to the exponential decay for subcritical percolation (and being bounded away from zero for supercritical percolation).

First we establish that p_c is actually $\frac{1}{2}$.

Theorem. In \mathbb{Z}^2 , we have $\theta(\frac{1}{2}) = 0$ and $p_c = \frac{1}{2}$.

It is conjectured that for all $d \geq 2$, we have $\theta(p_c(d)) = 0$. It is known to be true only in $d = 2$ and $d \geq 11$.


This was proved first by Harris, Kesten, Russo, Seymour, Welsh in several iterations.

Proof. First we prove that $\theta(\frac{1}{2}) = 0$. This will imply that $p_c \geq \frac{1}{2}$.

Suppose not, and $\theta(\frac{1}{2}) > 0$. Recall that $B(n) = [-n, n]^2$, and we define

$$C(n) = [-(n-1), (n-1)]^2 + (\frac{1}{2}, \frac{1}{2})$$

in the dual lattice. The appearance of the -1 is just a minor technical inconvenience. For the same n , our $B(n)$ and $C(n)$ look like



We claim that for large n , there is a positive probability that there are open paths from the left and right edges of $B(n)$ to ∞ , and also there are *closed* paths from the top and bottom edges of $C(n)$ to ∞ . But we know that with probability 1, there is a *unique* infinite cluster in both the primal lattice and the dual lattice. To connect up the two infinite open paths starting from the left and right edges of $B(n)$, there must be an open left-right crossing of $B(n)$. To connect up the two infinite closed paths starting from the top and bottom of $C(n)$, there must be a closed top-bottom crossing. But these cannot both happen, since this would require an open primal edge crossing a closed dual edge, which is impossible.

To make this an actual proof, we need to show that these events do happen with positive probability. We shall always take $p = \frac{1}{2}$, and will not keep repeating it.

First note that since there is, in particular, an infinite cluster with probability 1, we have

$$\mathbb{P}(\partial B(n) \leftrightarrow \infty) \rightarrow 1.$$

So we can pick n large enough such that

$$\mathbb{P}(\partial B(n) \leftrightarrow \infty), \mathbb{P}(\partial C(n) \leftrightarrow \infty) \geq 1 - \frac{1}{8^4}.$$

Let $A_\ell/A_r/A_t/A_b$ be the events that the left/right/top/bottom side of $B(n)$ is connected to ∞ via an open path of edges. Similarly, let D_ℓ be the event that the left of $C(n)$ is connected to ∞ via a closed path, and same for D_r, D_t, D_b .

Of course, by symmetry, for $i, j \in \{\ell, r, t, b\}$, we have $\mathbb{P}(A_i) = \mathbb{P}(A_j)$. Using FKG, we can bound

$$\mathbb{P}(\partial S_n \not\leftrightarrow \infty) = \mathbb{P}(A_\ell^c \cap A_r^c \cap A_t^c \cap A_b^c) \geq (\mathbb{P}(A_\ell^c))^4 = (1 - \mathbb{P}(A_\ell))^4.$$

Thus, by assumption on n , we have

$$(1 - \mathbb{P}(A_\ell))^4 \leq \frac{1}{8^4},$$

hence

$$\mathbb{P}(A_\ell) \geq \frac{7}{8}.$$

Of course, the same is true for other A_i and D_j .

Then if $G = A_\ell \cap A_r \cap D_t \cap D_b$, which is the desired event, then we have

$$\mathbb{P}(G^c) \leq \mathbb{P}(A_\ell^c) + \mathbb{P}(A_r^c) + \mathbb{P}(D_t^c) + \mathbb{P}(D_b^c) \leq \frac{1}{2}.$$

So it follows that

$$\mathbb{P}(G) \geq \frac{1}{2},$$


which, as argued, leads to a contradiction.

So we have $p_c \geq \frac{1}{2}$. It remains to prove that $p_c \leq \frac{1}{2}$. Suppose for contradiction that $p_c > \frac{1}{2}$. Then $p = \frac{1}{2}$ is in the subcritical regime, and we expect exponential decay. Thus, (again with $p = \frac{1}{2}$ fixed) there exists a $c > 0$ such that for all $n \geq 1$,

$$\mathbb{P}(0 \leftrightarrow \partial B(n)) \leq e^{-cn}.$$

Consider $C_n = [0, n+1] \times [0, n]$, and define A_n to be the event that there exists a left-right crossing of C_n by open edges.

Again consider the dual box $D_n = [0, n] \times [-1, n] + (\frac{1}{2}, \frac{1}{2})$.



Define B_n to be the event that there is a top-bottom crossing of D_n by closed edges of the dual.

As before, it cannot be the case that A_n and B_n both occur. In fact, A_n and B_n partition the whole space, since if A_n does not hold, then every path from left to right of C_n is blocked by a closed path of the dual. Thus, we know

$$\mathbb{P}(A_n) + \mathbb{P}(B_n) = 1.$$

But also by symmetry, we have $\mathbb{P}(A_n) = \mathbb{P}(B_n)$. So

$$\mathbb{P}(A_n) = \frac{1}{2}.$$

On the other hand, for any point on the left edge, the probability of it reaching the right edge decays exponentially with n . Thus,


$$\mathbb{P}(A_n) \leq n(n+1)\mathbb{P}(0 \leftrightarrow \partial B_n) \leq (n+1)e^{-cn}$$

which is a contradiction. So we are done. \square

So we now know that $p_c = \frac{1}{2}$. We now want to prove that


$$\mathbb{P}_{\frac{1}{2}}(0 \leftrightarrow \partial B(n)) \leq An^{-\alpha}$$

for some A, α . To prove this, we again consider the dual lattice. Observe that if we have a closed dual circuit around the origin, then this prevents the existence of an open path from the origin to a point far far away. What we are going to do is to construct “zones” in the dual lattice like this:




The idea is to choose these zones in a way such that the probability that each zone contains a closed circuit around the origin is $\geq \zeta$ for some fixed ζ . Thus, if $B(n)$ contains m many of these zones, then the probability that 0 is connected to $\partial B(n)$ is bounded above by $(1 - \zeta)^m$. We would then be done if we can show that $m \sim \log n$.

The main strategy to bounding these probabilities is to use FKG. For example, if we want to bound the probability that there is a closed circuit in a region



then we cut it up into the pieces



If we can bound the probability of there being a left-to-right crossing of a horizontal piece, and hence by symmetry the probability of there being a top-to-bottom crossing of a vertical piece, then FKG gives us a bound on the probability of there being a closed circuit.

For convenience of notation, we prove these bounds for open paths in the primal lattice. We make the following definitions:

$$B(k\ell, \ell) = [-\ell, (2k - 1)\ell] \times [-\ell, \ell]$$

$$B(\ell) = B(\ell, \ell) = [-\ell, \ell]^2$$

$$A(\ell) = B(3\ell) \setminus B(\ell)$$

$$LR(k\ell, \ell) = \{\text{there exists left-right crossing of } B(k\ell, \ell) \text{ of open edges}\}$$

$$LR(\ell) = LR(\ell, \ell)$$

$$O(\ell) = \{\text{there exists open circuit in } A(\ell) \text{ that contains } 0 \text{ in its interior}\}.$$

We first note that we have already proven the following:

Proposition. $\mathbb{P}_{\frac{1}{2}}(LR(\ell)) \geq \frac{1}{2}$ for all ℓ .

Proof. We have already seen that the probability of there being a left-right crossing of $[0, n + 1] \times [0, n]$ is at least $\frac{1}{2}$. But if there is a left-right crossing of $[0, n + 1] \times [0, n]$, then there is also a left-right crossing of $[0, n] \times [0, n]$! \square

For a general p , Russo–Seymour–Welsh (RSW) lets us bound $\mathbb{P}_p(O(\ell))$ by $\mathbb{P}_p(LR(\ell))$:

Theorem (Russo–Seymour–Welsh (RSW) theorem). If $\mathbb{P}_p(LR(\ell)) = \alpha$, then

$$\mathbb{P}_p(O(\ell)) \geq \left(\alpha (1 - \sqrt{1 - \alpha})^4 \right)^{12}.$$

A large part of the proof is done by the cut-and-paste argument we sketched above. However, to successfully do cut-and-paste, it turns out we need bounds on the probability of a left-right crossing on something that is *not* a square. The key, non-trivial estimate is the following:

Lemma. If $\mathbb{P}_p(LR(\ell)) = \alpha$, then

$$\mathbb{P}_p\left(LR\left(\frac{3}{2}\ell, \ell\right)\right) \geq (1 - \sqrt{1 - \alpha})^3.$$

To prove this, we need a result from the first example sheet:

Lemma (*n*th root trick). If A_1, \dots, A_n are increasing events all having the same probability, then

$$\mathbb{P}_p(A_1) \geq 1 - \left(1 - \mathbb{P}_p\left(\bigcup_{i=1}^n A_i\right)\right)^{1/n}.$$


Proof sketch. Observe that the proof of FKG works for *decreasing* events as well, and then apply FKG to A_i^c . \square

We can now prove our initial lemma.

Proof sketch. Let \mathcal{A} be the set of left-right crossings of $B(\ell) = [-\ell, \ell]^2$. Define a partial order on \mathcal{A} by $\pi_1 \leq \pi_2$ if π_1 is contained in the closed bounded region of $B(\ell)$ below π_2 .

Note that given any configuration, if the set of open left-right crossings is non-empty, then there exists a lowest one. Indeed, since \mathcal{A} must be finite, it suffices to show that meets exist in this partial order, which is clear.

For a left-right crossing π , let $(0, y_\pi)$ be the last vertex on the vertical axis where π intersects, and let π_r be the path of the path that connects $(0, y_\pi)$ to the right.




Let

$$\mathcal{A}_- = \{\pi \in \mathcal{A} : y_\pi \leq 0\}$$

$$\mathcal{A}_+ = \{\pi \in \mathcal{A} : y_\pi \geq 0\}$$

Letting $B(\ell)' = [0, 2\ell] \times [-\ell, \ell]$, our goal is to find a left-right crossing of the form



More precisely, we want the following paths:

- (i) Some $\pi \in \mathcal{A}_-$
- (ii) Some top-bottom crossing of $B(\ell')$ that crosses π_r .
- (iii) Some left-right crossing of $B(\ell')$ that starts at the positive (i.e. non-negative) y axis.

To understand the probabilities of these events happening, we consider the “mirror” events and then apply the square root trick.

Let π'_r be the reflection of π_r on $\{(\ell, k) : k \in \mathbb{Z}\}$. For any $\pi \in \mathcal{A}$, we define

$$\begin{aligned} V_\pi &= \{\text{all edges of } \pi \text{ are open}\} \\ M_\pi &= \{\text{exists open crossing from top of } B(\ell)' \text{ to } \pi_r \cup \pi'_r\} \\ M_\pi^- &= \{\text{exists open crossing from top of } B(\ell)' \text{ to } \pi_r\} \\ M_\pi^+ &= \{\text{exists open crossing from top of } B(\ell)' \text{ to } \pi'_r\} \\ L^+ &= \{\text{exists open path in } \mathcal{A}_+\} \\ L^- &= \{\text{exists open path in } \mathcal{A}_-\} \\ L_\pi &= \{\pi \text{ is the lowest open LR crossing of } B(\ell)\} \\ N &= \{\text{exists open LR crossing of } B(\ell)'\} \\ N^+ &= \{\text{exists open LR crossing in } B(\ell)' \text{ starting from positive vertical axis}\} \\ N^- &= \{\text{exists open LR crossing in } B(\ell)' \text{ starting from negative vertical axis}\} \end{aligned}$$

In this notation, our previous observation was that

$$\underbrace{\bigcup_{\pi \in \mathcal{A}_-} (V_\pi \cap M_\pi^-)}_G \cap N^+ \subseteq LR\left(\frac{3}{2}\ell, \ell\right)$$

So we know that

$$\mathbb{P}_p\left(LR\left(\frac{3}{2}\ell, \ell\right)\right) \geq \mathbb{P}_p(G \cap N^+) \geq \mathbb{P}_p(G)\mathbb{P}_p(N^+),$$

using FKG.

Now by the “square root trick”, we know

$$\mathbb{P}_p(N^+) \geq 1 - \sqrt{1 - \mathbb{P}_p(N^+ \cup N^-)}.$$

Of course, we have $\mathbb{P}_p(N^+ \cup N^-) = \mathbb{P}_p(LR(\ell)) = \alpha$. So we know that

$$\mathbb{P}_p(N^+) \geq 1 - \sqrt{1 - \alpha}.$$

We now have to understand G . To bound its probability, we try to bound it by the union of some disjoint events. We have

$$\begin{aligned} \mathbb{P}_p(G) &= \mathbb{P}_p\left(\bigcup_{\pi \in \mathcal{A}_-} (V_\pi \cap M_\pi^-)\right) \\ &\geq \mathbb{P}_p\left(\bigcup_{\pi \in \mathcal{A}_-} (M_\pi^- \cap L_\pi)\right) \\ &= \sum_{\pi \in \mathcal{A}_-} \mathbb{P}_p(M_\pi^- \mid L_\pi)\mathbb{P}_p(L_\pi). \end{aligned}$$

Claim.

$$\mathbb{P}_p(M_\pi^- \mid L_\pi) \geq 1 - \sqrt{1 - \alpha}.$$

Note that if π intersects the vertical axis in one point, then $\mathbb{P}_p(M_\pi^- | L_\pi) = \mathbb{P}_p(M_\pi^- | V_\pi)$, since L_π^- tells us what happens below π , and this does not affect the occurrence of M_π^- .

Since M_π^- and V_π are increasing events, by FKG, we have

$$\mathbb{P}_p(M_\pi^- | V_\pi) \geq \mathbb{P}_p(M_\pi^-) \geq 1 - \sqrt{1 - \mathbb{P}_p(M_\pi^- \cup M_\pi^+)} = 1 - \sqrt{1 - \mathbb{P}_p(M_\pi)}.$$

Since $\mathbb{P}_p(M_\pi) \geq \mathbb{P}_p(LR(\ell)) = \alpha$, the claim follows.

In the case where π is more complicated, we will need an extra argument, which we will not provide.

Finally, we have

$$\mathbb{P}_p(G) \geq \sum_{\pi \in \mathcal{A}_-} \mathbb{P}_p(L_\pi)(1 - \sqrt{1 - \alpha}) = (1 - \sqrt{1 - \alpha})\mathbb{P}_p(L^-).$$

But again by the square root trick,

$$\mathbb{P}_p(L^-) \geq 1 - \sqrt{1 - \mathbb{P}_p(L^+ \cup L^-)} = 1 - \sqrt{1 - \alpha},$$


and we are done. □

We now do the easy bit to finish off the theorem:

Lemma.

$$\begin{aligned} \mathbb{P}_p(LR(2\ell, \ell)) &\geq \mathbb{P}_p(LR(\ell)) \left(\mathbb{P}_p \left(LR \left(\frac{3}{2}\ell, \ell \right) \right) \right)^2 \\ \mathbb{P}_p(LR(3\ell, \ell)) &\geq \mathbb{P}_p(LR(\ell)) (\mathbb{P}_p(LR(2\ell, \ell)))^2 \\ \mathbb{P}_p(O(\ell)) &\geq \mathbb{P}_p(LR(3\ell, \ell))^4 \end{aligned}$$

Proof. To prove the first inequality, consider the box $[0, 4\ell] \times [-\ell, \ell]$.



We let

$$\begin{aligned} LR_1 &= \{\text{exists left-right crossing of } [0, 3\ell] \times [-\ell, \ell]\} \\ LR_2 &= \{\text{exists left-right crossing of } [\ell, 4\ell] \times [-\ell, \ell]\} \\ TB_1 &= \{\text{exists top-bottom crossing of } [\ell, 3\ell] \times [-\ell, \ell]\}. \end{aligned}$$

Then by FKG, we find that

$$\mathbb{P}_p(LR(2\ell, \ell)) \geq \mathbb{P}_p(LR_1 \cap LR_2 \cap TB_1) \geq \mathbb{P}_p(LR_1)\mathbb{P}_p(LR_2)\mathbb{P}_p(TB_1).$$

The others are similar. □

Theorem. There exists positive constants $\alpha_1, \alpha_2, \alpha_3, \alpha_4, A_1, A_2, A_4$ such that

$$\begin{aligned}\mathbb{P}_{\frac{1}{2}}(0 \leftrightarrow \partial B(n)) &\leq A_1 n^{-\alpha_1} \\ \mathbb{P}_{\frac{1}{2}}(|\mathcal{C}(0)| \geq n) &\leq A_2 n^{-\alpha_2} \\ \mathbb{E}(|\mathcal{C}(0)|^{\alpha_3}) &\leq \infty\end{aligned}$$

Moreover, for $p > p_c = \frac{1}{2}$, we have

$$\theta(p) \leq A_4 \left(p - \frac{1}{2}\right)^{\alpha_4}.$$

It is an exercise on the example sheet to prove that $\mathbb{P}_{\frac{1}{2}}(0 \leftrightarrow B(n)) \geq \frac{1}{2\sqrt{n}}$ using the BK inequality. So the true decay of $\mathbb{P}_{\frac{1}{2}}(0 \leftrightarrow \partial B(n))$ is indeed a power law.

Proof.

(i) We first prove the first inequality. Define dual boxes

$$B(k)_d = B(k) + \left(\frac{1}{2}, \frac{1}{2}\right).$$

The dual annuli $A(\ell)_d$ are defined by

$$A(\ell)_d = B(3\ell)_d \setminus B(\ell)_d.$$

We let $O(\ell)_d$ be the event that there is a closed dual circuit in $A(\ell)_d$ containing $(\frac{1}{2}, \frac{1}{2})$. Then RSW tells us there is some $\zeta \in (0, 1)$ such that

$$\mathbb{P}_{\frac{1}{2}}(O(\ell)_d) \geq \zeta,$$

independent of ℓ . Then observe that

$$\mathbb{P}_{\frac{1}{2}}(0 \leftrightarrow \partial B(3^k + 1)) \leq \mathbb{P}_p(O(3^r)_d \text{ does not occur for all } r < k).$$

Since the annuli $(A(3^r)_d)$ are disjoint, the events above are independent. So

$$\mathbb{P}_{\frac{1}{2}}(0 \leftrightarrow \partial B(3^k + 1)) \leq (1 - \zeta)^k,$$

and this proves the first inequality.

- (ii) The second inequality follows from the first inequality plus the fact that $|\mathcal{C}(0)| \geq n$ implies $0 \leftrightarrow \partial B(g(n))$, for some function $g(n) \sim \sqrt{n}$.
- (iii) To show that $\mathbb{E}_{\frac{1}{2}}[|\mathcal{C}(0)|^{\alpha_3}] < \infty$ for some α_3 , we observe that this expectation is just

$$\sum_n \mathbb{P}_{\frac{1}{2}}(|\mathcal{C}(0)|^{\alpha_3} \geq n).$$

(iv) To prove the last part, note that

$$\theta(p) = \mathbb{P}_p(|\mathcal{C}(0)| = \infty) \leq \mathbb{P}_p(0 \leftrightarrow \partial B_n)$$

for all n . By the corollary of Russo's formula, and since $\{0 \leftrightarrow \partial B_n\}$ only depends on the edges in B_n , which are $\leq 18n^2$, we get that

$$\mathbb{P}_{\frac{1}{2}}(0 \leftrightarrow \partial B_n) \geq \left(\frac{1}{2p}\right)^{18n^2} \mathbb{P}_p(0 \leftrightarrow \partial B_n).$$

So

$$\theta(p) \leq (2p)^{18n^2} \mathbb{P}_{\frac{1}{2}}(0 \leftrightarrow \partial B_n) \leq A_1 (2p)^{18n^2} n^{-\alpha_1}.$$

Now take $n = \lfloor (\log 2p)^{-1/2} \rfloor$. Then as $p \searrow \frac{1}{2}$, we have

$$n \sim \frac{1}{(2p-1)^{\frac{1}{2}}}.$$

Substituting this in, we get

$$\theta(p) \leq C \left(p - \frac{1}{2}\right)^{\alpha_1/2}. \quad \square$$

By similar methods, we can prove that

Theorem. When $d = 2$ and $p > p_c$, there exists a positive constant c such that

$$\mathbb{P}_p(0 \leftrightarrow \partial B(n), |\mathcal{C}(0)| < \infty) \leq e^{-cn}.$$

It is natural to ask if we have a similar result in higher dimensions. In higher dimensions, all the techniques in \mathbb{Z}^2 don't work.

Higher dimensions

In $d \geq 3$, define the *slab*

$$S_k = \mathbb{Z}^2 \times \{0, 1, \dots, k\}^{d-2}.$$

Then $S_k \subseteq S_{k+1} \subseteq \mathbb{Z}^d$.

In general, for a graph G , we let $p_c(G)$ be the critical probability of bond percolation in G . Then we have

$$p_c(S_k) \geq p_c(S_{k+1}) \geq p_c.$$

So the sequence $(p_c(S_k))$ must converge to a limit. Call this limit

$$p_c^{slab} = \lim_{k \rightarrow \infty} p_c(S_k).$$

We know that $p_c^{slab} \geq p_c$.

A lot of the results we can prove for \mathbb{Z}^2 can be proven for p_c^{slab} instead of p_c . So the natural question is how p_c^{slab} is related to p_c , and in particular, if they are equal. This has been an open question for a long time, until Grimmett and Marstrand proved it.

Theorem (Grimmett–Marstrand). Let F be an infinite-connected subset of \mathbb{Z}^d with $p_c(F) < 1$. Then for all $\eta > 0$, there exists $k \in \mathbb{N}$ such that

$$p_c(2kF + B_k) \leq p_c + \eta.$$

In particular, for all $d \geq 3$, $p_c^{slab} = p_c$.

We shall not prove the theorem, but we shall indicate how the “in particular” part works.

We take $F = \mathbb{Z}^2 \times \{0\}^{d-2}$. Then

$$2kF + B_k = \mathbb{Z}^2 \times ([-k, k]^{d-2} \cap \mathbb{Z}^{d-2}),$$

a translate of S_{2k} . So $p_c(S_{2k}) = p_c(2kF + B_k) \rightarrow p_c$ as $k \rightarrow \infty$.

A consequence of Grimmett-Marstrand is that

Theorem. If $d \geq 3$ and $p > p_c$, then there exists $c > 0$ such that

$$\mathbb{P}_p(0 \leftrightarrow \partial B(n), |\mathcal{C}(0)| < \infty) \leq e^{-cn}.$$

1.4 Conformal invariance and SLE in $d = 2$

Instead of working on \mathbb{Z}^2 , we will now work on the triangular lattice \mathbb{T} . We consider *site percolation*. So every vertex is open (black) with probability p and closed (white) with probability $1 - p$, independently for different vertices.

Like for \mathbb{Z}^2 , we can show that $p = p_c(\mathbb{T}) = \frac{1}{2}$.

Let D be an open simply connected domain in \mathbb{R}^2 with ∂D a Jordan curve. Let $a, b, c \in \partial D$ be 3 points labeled in anti-clockwise order. Consider the triangle T with vertices $A = 0$, $B = 1$ and $C = e^{i\pi/3}$. By the Riemann mapping theorem, there exists a conformal map $\varphi : D \rightarrow T$ that maps $a \mapsto A$, $b \mapsto B$, $c \mapsto C$.

Moreover, this can be extended to ∂D such that $\varphi : \bar{D} \rightarrow \bar{T}$ is a homeomorphism. If x is in the arc bc , then it will be mapped to $X = \varphi(x)$ on the edge BC of T .

Focus on the case $p = \frac{1}{2}$. We can again put a triangular lattice inside D , with mesh size δ . We let $ac \leftrightarrow bx$ to mean the event that there is an open path in D joining the arc ac to bx .

Then by RSW (for \mathbb{T}), we get that

$$\mathbb{P}_\delta(ac \leftrightarrow bx) \geq c > 0,$$

independent of δ . We might ask what happens when $\delta \rightarrow 0$. In particular, does it converge, and what does it converge to?

Cardy, a physicist studying conformal field theories, conjectured that

$$\lim_{\delta \rightarrow 0} \mathbb{P}_\delta(ac \leftrightarrow bx) = |BX|.$$

He didn't really write it this way, but expressed it in terms of hypergeometric functions. It was Carleson who expressed it in this form.

In 2001, Smirnov proved this conjecture.

Theorem (Smirnov, 2001). Suppose (Ω, a, b, c, d) and $(\Omega', a', b', c', d')$ are conformally equivalent. Then

$$\mathbb{P}(ac \leftrightarrow bd \text{ in } \Omega) = \mathbb{P}(a'c' \leftrightarrow b'd' \text{ in } \Omega').$$

This says percolation at criticality on the triangular lattice is conformally invariant.

We may also take the dual of the triangular lattice, which is the hexagonal lattice. Colour a hexagon black if the center is black, and similarly for whites.

Suppose we impose the boundary condition on the upper half plane that the hexagons are black when $x > 0, y = 0$, and white when $x < 0, y = 0$.

Starting at $(x, y) = (0, 0)$, we explore the interface between the black and white by always keeping a black to our right and white on our right. We can again take the limit $\delta \rightarrow 0$. What is this exploration path going to look like in the limit $\delta \rightarrow 0$? It turns out this is an SLE(6) curve.

To prove this, we use the fact that the exploration curve satisfies the *locality property*, namely that if we want to know where we will be in the next step, we only need to know where we currently are.

2 Random walks

2.1 Random walks in finite graphs

We shall consider weighted random walks on finite graphs. Let $G = (V, E)$ be a finite connected graph. We assign a conductances/weights to the edges $(c(e))_{e \in E}$. Here c is a function of the edges, so we require $c(xy) = c(yx)$.

In the weighted random walk on G , the probability of going from x to y (assuming $x \sim y$) is

$$\mathbb{P}(x, y) = \frac{c(x, y)}{c(x)}, \quad c(x) = \sum_{z \sim x} c(x, z).$$

If we put $c(e) = 1$ for all e , then this is just a simple random walk, with

$$\mathbb{P}(x, y) = \begin{cases} \frac{1}{\deg x} & y \sim x \\ 0 & \text{otherwise} \end{cases}.$$

The weighted random walk is reversible with respect to $\pi(x) = \frac{c(x)}{c_G}$, where $c_G = \sum_x c(x)$, because

$$\pi(x)P(x, y) = \frac{c(x)}{c_G} \cdot \frac{c(x, y)}{c(x)} = \pi(y)\mathbb{P}(y, x).$$

Conversely, every reversible Markov chain can be represented as a random walk on a weighted graph — place an edge $\{x, y\}$ if $\mathbb{P}(x, y) > 0$. By reversibility, $\mathbb{P}(x, y) > 0$ iff $\mathbb{P}(y, x) > 0$. Define weights

$$c(x, y) = \pi(x)\mathbb{P}(x, y)$$

for all $x \sim y$.

One way to understand random walks on weighted graphs is to think of them as currents in electrical networks. We set the resistance on the edges by

$$r(e) = \frac{1}{c(e)}.$$

Naturally, we would like to talk about currents flowing through the circuit.

Definition (Flow). A *flow* θ on G is a function defined on oriented edges which is anti-symmetric, i.e. $\theta(x, y) = -\theta(y, x)$.

Not every flow can be thought of as a current. We fix two distinguished vertices a and z , called the *source* and the *sink* respectively. We think of current as entering the network through a and exiting through z . Through any other vertex, the amount of current going in should equal the amount of current going out.

Definition (Divergence). The *divergence* of a flow θ is

$$\operatorname{div} \theta(x) = \sum_{y \sim x} \theta(x, y).$$

By the antisymmetry property, we have

$$\sum_x \operatorname{div} \theta(x) = \sum_x \sum_{y \sim x} \theta(x, y) = \sum_{x \sim y} \theta(x, y) = 0.$$

Definition (Flow from a to z). A flow θ from a to z is a flow such that

- (i) $\operatorname{div} \theta(x) = 0$ for all $x \notin \{a, z\}$. (Kirchhoff's node law)
- (ii) $\operatorname{div} \theta(a) \geq 0$.

The *strength* of the flow from a to z is $\|\theta\| = \operatorname{div} \theta(a)$. We say this is a *unit flow* if $\|\theta\| = 1$.

Since $\sum \operatorname{div} \theta(x) = 0$, a flow from a to z satisfies $\operatorname{div} \theta(z) = -\operatorname{div} \theta(a)$.

Now if we believe in Ohm's law, then the voltage difference across each edge is $I(x, y)r(x, y)$. So we have

$$W(x) = W(y) + r(x, y)I(x, y)$$

for all y . We know that if we sum $I(x, y)$ over all y , then the result vanishes. So we rewrite this as

$$c(x, y)W(x) = c(x, y)W(y) + I(x, y).$$

Summing over all y , we find that the voltage satisfies

$$W(x) = \sum_{x \sim y} \frac{c(x, y)}{c(x)} W(y) = \sum_{x \sim y} \mathbb{P}(x, y) W(y).$$

Definition (Harmonic function). Let P be a transition matrix on Ω . We call h *harmonic* for P at the vertex x if

$$h(x) = \sum_y \mathbb{P}(x, y) h(y).$$

We now take this as a *definition* of what a voltage is.

Definition (Voltage). A *voltage* W is a function on Ω that is harmonic on $\Omega \setminus \{a, z\}$.

The first theorem to prove is that given any two values $W(a)$ and $W(z)$, we can find a unique voltage with the prescribed values at a and z .

More generally, if $B \subseteq \Omega$, and X is a Markov chain with matrix P , then we write

$$\tau_B = \min\{t \geq 0 : X_t \in B\},$$

Proposition. Let P be an irreducible matrix on Ω and $B \subseteq \Omega$, $f : B \rightarrow \mathbb{R}$ a function. Then

$$h(x) = \mathbb{E}_x[f(X_{\tau_B})]$$

is the unique extension of f which is harmonic on $\Omega \setminus B$.

Proof. It is obvious that $h(x) = f(x)$ for $x \in B$. Let $x \notin B$. Then

$$\begin{aligned} h(x) &= \mathbb{E}_x[f(X_{\tau_B})] \\ &= \sum_y \mathbb{P}(x, y) \mathbb{E}_x[f(X_{\tau_B}) \mid X_1 = y] \\ &= \sum_y \mathbb{P}(x, y) \mathbb{E}_y[f(X_{\tau_B})] \\ &= \sum_Y \mathbb{P}(x, y) h(y) \end{aligned}$$

So h is harmonic.

To show uniqueness, suppose h' be another harmonic extension. Take $g = h - h'$. Then $g = 0$ on B . Set

$$A = \left\{ x : g(x) = \max_{y \in \Omega \setminus B} g(y) \right\},$$

and let $x \in A \setminus B$. Now since $g(x)$ is the weighted average of its neighbours, and $g(y) \leq g(x)$ for all neighbours y of x , it must be the case that $g(x) = g(y)$ for all neighbours of x .

Since we assumed G is connected, we can construct a path from x to the boundary, where g vanishes. So $g(x) = 0$. In other words, $\max g = 0$. Similarly, $\min g = 0$. So $g = 0$. \square

Given a voltage, Ohm's law defines a natural flow, called the *current flow*.

Definition (Current flow). The *current flow* associated to the voltage W is

$$I(x, y) = \frac{W(x) - W(y)}{r(x, y)} = c(x, y)(W(x) - W(y)).$$

The current flow has an extra property, namely that it satisfies the *cycle law*. For any cycle e_1, e_2, \dots, e_n , we have

$$\sum_{i=1}^n r(e_i) I(e_i) = 0.$$

Proposition. Let θ be a flow from a to z satisfying the cycle law for any cycle. Let I the current flow associated to a voltage W . If $\|\theta\| = \|I\|$, then $\theta = I$.

Proof. Take $f = \theta - I$. Then f is a flow which satisfies Kirchoff's node law at all vertices and the cycle law for any cycle. We want to show that $f = 0$. Suppose not. Then we can find some e_1 such that $f(e_1) > 0$. But

$$\sum_{y \sim x} f(x, y) = 0$$

for all x , so there must exist an edge e_2 that e_1 leads to such that $f(e_2) > 0$. Continuing this way, we will get a cycle of ordered edges where $f > 0$, which violates the cycle law. \square

Let W_0 be a voltage with $W_0(a) = 1$ and $W_0(z) = 0$, and let I_0 be the current associated to W_0 . Then any other voltage W can be written as

$$W(x) = (W(a) - W(z))W_0(x) + W(z).$$

So, noting that $W_0(a) = 1$, we have

$$W(a) = W(x) = (W(a) - W(z))(W_0(a) - W_0(x)).$$

Thus, if I is the current associated to W , then

$$\|I\| = \sum_{x \sim a} I(a, x) = \sum_{x \sim a} \frac{W(a) - W(x)}{r(a, x)} = (W(a) - W(z))\|I_0\|.$$

So we know that

$$\frac{W(a) - W(z)}{\|I\|} = \frac{1}{\|I_0\|},$$

and in particular does not depend on W .

Definition (Effective resistance). The *effective resistance* $R_{\text{eff}}(a, z)$ of an electric network is defined to be the ratio

$$R_{\text{eff}}(a, z) = \frac{W(a) - W(z)}{\|I\|}$$

for any voltage W with associated current I . The *effective conductance* is $C_{\text{eff}}(a, z) = R_{\text{eff}}(a, z)^{-1}$.

Proposition. Take a weighted random walk on G . Then

$$\mathbb{P}_a(\tau_z < \tau_a^+) = \frac{1}{c(a)R_{\text{eff}}(a, z)},$$

where $\tau_a^+ = \min\{t \geq 1 : X_t = a\}$.

Proof. Let

$$f(x) = \mathbb{P}_x(\tau_z < \tau_a).$$

Then $f(a) = 0$ and $f(z) = 1$. Moreover, f is harmonic on $\Omega \setminus \{a, z\}$. Let W be a voltage. By uniqueness, we know

$$f(x) = \frac{W(a) - W(x)}{W(a) - W(z)}.$$

So we know

$$\begin{aligned} \mathbb{P}_a(\tau_z < \tau_a^+) &= \sum_{x \sim a} \mathbb{P}(a, x) f(x) \\ &= \sum_{x \sim a} \frac{c(a, x)}{c(a)} \frac{W(a) - W(x)}{W(a) - W(z)} \\ &= \frac{1}{c(a)(W(a) - W(z))} \sum_{x \sim a} I(a, x) \\ &= \frac{\|I\|}{c(a)(W(a) - W(z))} \\ &= \frac{1}{c(a)R_{\text{eff}}(a, z)}. \end{aligned}$$

□

Definition (Green kernel). Let τ be a stopping time. We define the *Green kernel* to be

$$G_\tau(a, x) = \mathbb{E}_a \left[\sum_{t=0}^{\infty} \mathbf{1}(X_t = x, t < \tau) \right].$$

This is the expected number of times we hit x before τ occurs.

Corollary. For any reversible chain and all a, z , we have

$$G_{\tau_z}(a, a) = c(a)R_{\text{eff}}(a, z).$$

Proof. By the Markov property, the number of visits to a starting from a until τ_z is the geometric distribution $\text{Geo}(\mathbb{P}_a(\tau_z < \tau_a^+))$. So

$$G_{\tau_z}(a, a) = \frac{1}{\mathbb{P}_a(\tau_z < \tau_a^+)} = c(a)R_{\text{eff}}(a, z). \quad \square$$

Practically, to compute the effective resistance, it is useful to know some high school physics. For example, we know that conductances in parallel add:




Thus, if e_1 and e_2 are two edges with the same endvertices, then we can replace them by a single edge of conductance the sum of the conductances. We can prove that this actually works by observing that the voltages and currents remain unchanged outside of e_1, e_2 , and we check that Kirchoff's and Ohm's law are satisfied with

$$I(e) = I(e_1) + I(e_2).$$

Similarly, resistances in series add. Let v be a node of degree 2 and let v_1 and v_2 be its 2 neighbours. Then we can replace these 2 edges by a single edge of resistance $r_1 + r_2$. Again, to justify this, we check Kirchoff's and Ohm's law with

$$I(v_1 v_2) = I(v_1 v) + I(v v_2).$$



We can also glue 2 vertices of the same potential, keeping all other edges. Current and potential are unchanged since current doesn't flow if there is the same potential.

Exercise. Let T be a finite connected tree, and a, z two vertices. Then $R_{\text{eff}}(a, z)$ is the graph distance between a and z .

It turns out it is not convenient to work with our definition of effective resistance. Instead, we can use the well known result from high school physics $P = IV = I^2 R$. For any flow, we define

Definition (Energy). Let θ be a flow on G with conductances $(c(e))$. Then the energy is

$$\mathcal{E}(\theta) = \sum_e (\theta(e))^2 r(e).$$

Here we sum over unoriented edges.

Note that given any flow, we can always increase the energy by pumping more and more current along a cycle. Thus, we might expect the lowest energy configuration to be given by the flows satisfying the cycle law, and if we focus on unit flows, the following should not be surprising:

Theorem (Thomson's principle). Let G be a finite connected graph with conductances $(c(e))$. Then for any a, z , we have

$$R_{\text{eff}}(a, z) = \inf\{\mathcal{E}(\theta) : \theta \text{ is a unit flow from } a \text{ to } z\}.$$

Moreover, the unit current flow from a to z is the unique minimizer.

Proof. Let i be the unit current flow associated to potential φ . Let j be another unit flow $a \rightarrow z$. We need to show that $\mathcal{E}(j) \geq \mathcal{E}(i)$ with equality iff $j = i$.

Let $k = j - i$. Then k is a flow of 0 strength. We have

$$\mathcal{E}(j) = \mathcal{E}(i + k) = \sum_e (i(e) + k(e))^2 r(e) = \mathcal{E}(i) + \mathcal{E}(k) + 2 \sum_e i(e)k(e)r(e).$$

It suffices to show that the last sum is zero. By definition, we have

$$i(x, y) = \frac{\varphi(x) - \varphi(y)}{r(x, y)}.$$

So we know

$$\begin{aligned} \sum_e i(e)k(e)r(e) &= \frac{1}{2} \sum_x \sum_{y \sim x} (\varphi(x) - \varphi(y))k(x, y) \\ &= \frac{1}{2} \sum_x \sum_{y \sim x} \varphi(x)k(x, y) + \frac{1}{2} \sum_x \sum_{y \sim x} \varphi(x)k(x, y). \end{aligned}$$

Now note that $\sum_{y \sim x} k(x, y) = 0$ for any x since k has zero strength. So this vanishes.

Thus, we get

$$\mathcal{E}(j) = \mathcal{E}(i) + \mathcal{E}(k),$$

and so $\mathcal{E}(j) \geq \mathcal{E}(i)$ with equality iff $\mathcal{E}(k) = 0$, i.e. $k(e) = 0$ for all e . \square

Using this characterization, it is then immediate that

Theorem (Rayleigh's monotonicity principle). Let G be a finite connected graph and $(r(e))_e$ and $(r'(e))_e$ two sets of resistances on the edges such that $r(e) \leq r'(e)$ for all e . Then

$$R_{\text{eff}}(a, z; r) \leq R_{\text{eff}}(a, z; r').$$

for all $a, z \in G$.

Proof. By definition of energy, for any current i , we have

$$\mathcal{E}(i; r) \leq \mathcal{E}(i; r').$$

Then take the infimum and conclude by Thompson. \square

Corollary. Suppose we add an edge to G which is not adjacent to a . This increases the escape probability

$$\mathbb{P}_a(\tau_z < \tau_a^+).$$

Proof. We have

$$\mathbb{P}_a(\tau_z < \tau_a^+) = \frac{1}{c(a)R_{\text{eff}}(a, z)}.$$

We can think of the new edge as an edge having infinite resistance in the old graph. By decreasing it to a finite number, we decrease R_{eff} . \square

Definition (Edge cutset). A set of edges Π is an edge-cutset separating a from z if every path from a to z uses an edge of Π .

Theorem (Nash–Williams inequality). Let (Π_k) be disjoint edge-cutsets separating a from z . Then

$$R_{\text{eff}}(a, z) \geq \sum_k \left(\sum_{e \in \Pi_k} c(e) \right)^{-1}.$$

The idea is that if we combine all paths in each Π_k to a single path with conductance $\sum_{e \in \Pi_k} c(e)$, then every path has to use all of these edges, so we can bound the effective resistance.

Proof. By Thompson, it suffices to prove that for any unit flow θ from a to z , we have

$$\sum_e (\theta(e))^2 r(e) \geq \sum_k \left(\sum_{e \in \Pi_k} c(e) \right)^{-1}.$$

Certainly, we have

$$\sum_e (\theta(e))^2 r(e) \geq \sum_k \sum_{e \in \Pi_k} (\theta(e))^2 r(e).$$

We now use Cauchy–Schwarz to say

$$\begin{aligned} \left(\sum_{e \in \Pi_k} (\theta(e))^2 r(e) \right) \left(\sum_{e \in \Pi_k} c(e) \right) &\geq \left(\sum_{e \in \Pi_k} |\theta(e)| \sqrt{r(e)c(e)} \right)^2 \\ &= \left(\sum_{e \in \Pi_k} |\theta(e)| \right)^2. \end{aligned}$$

The result follows if we show that the final sum is ≥ 1 .

Let F be the component of $G \setminus \Pi_k$ containing a . Let F' be the set of all edges that start in F and land outside F . Then we have

$$1 = \sum_{x \in F} \text{div } \theta(x) \leq \sum_{e \in F'} |\theta(e)| \leq \sum_{e \in \Pi_k} |\theta(e)|. \quad \square$$

Corollary. Consider $B_n = [1, n]^2 \cap \mathbb{Z}^2$. Then

$$R_{\text{eff}}(a, z) \geq \frac{1}{2} \log(n-1).$$

There is also an upper bound of the form $R_{\text{eff}}(a, z) \leq c \log n$. So this is indeed the correct order. The proof involves constructing an explicit flow with this energy.

Proof. Take

$$\Pi_k = \{(v, u) \in B_n : \|v\|_\infty = k-1, \|u\|_\infty = k\}.$$

Then

$$|\Pi_k| = 2(k-1),$$

and so

$$R_{\text{eff}}(a, z) \geq \sum_{k=1}^{-1} \frac{1}{|\Pi_k|} \geq \frac{1}{2} \log(n-1). \quad \square$$

We shall need the following result:

Proposition. Let X be an irreducible Markov chain on a finite state space. Let τ be a stopping time such that $\mathbb{P}_a(X_\tau = a) = 1$ and $\mathbb{E}_a[\tau] < \infty$ for some a in the state space. Then

$$G_\tau(a, x) = \pi(x) \mathbb{E}_a[\tau].$$

A special case of this result was proven in IB Markov Chain, where τ was taken to be the first return time to a . The proof of this is very similar, and we shall not repeat it.

Using this, we can prove the commute-time identity.

Theorem (Commute time identity). Let X be a reversible Markov chain on a finite state space. Then for all a, b , we have

$$\mathbb{E}_a[\tau_b] + \mathbb{E}_b[\tau_a] = c(G) R_{\text{eff}}(a, b),$$

where

$$c(G) = 2 \sum_e c(e).$$

Proof. Let $\tau_{a,b}$ be the first time we hit a after we hit b . Then $\mathbb{P}_a(X_{\tau_{a,b}} = a) = 1$. By the proposition, we have

$$G_{\tau_{a,b}}(a, a) = \pi(a) \mathbb{E}_a[\tau_{a,b}].$$

Since starting from a , the number of times we come back to A by $\tau_{a,b}$ is the same as up to time τ_b , we get

$$G_{\tau_{a,b}}(a, a) = G_{\tau_b}(a, a) = c(a) R_{\text{eff}}(a, b).$$

Combining the two, we get

$$\pi(a) \mathbb{E}_a[\tau_{a,b}] = c(a) R_{\text{eff}}(a, b).$$

Thus, the result follows. □

In particular, from the commute-time identity, we see that the effective resistance satisfies the triangle inequality, and hence defines a metric on any graph.

2.2 Infinite graphs

Finally, let's think a bit about how we can develop the same theory for infinite graphs.

Let G be an infinite connected weighted graph with conductances $(c(e))_e$. Let 0 be a distinguished vertex. Let $(G_k = (V_k, E_k))_k$ be an *exhaustion* of G by finite graphs, i.e.

- $V_k \subseteq V_{k+1}$ for all k ;
- $0 \in V_k$ for all k ;
- $\bigcup_k V_k = V$;
- E_k is the set of edges of E with both end points in V_k .

Let G_n^* be the graph obtained by gluing all the vertices of $V \setminus V_n$ into a single point z_n and setting

$$c(x, z_n) = \sum_{z \in V \setminus V_n} c(x, z)$$

for all $x \in V_n$. Now observe that $R_{\text{eff}}(0, z_n; G_n^*)$ is a non-decreasing function of n (Rayleigh). So the limit as $n \rightarrow \infty$ exists. We set

$$R_{\text{eff}}(0, \infty) = \lim_{n \rightarrow \infty} R_{\text{eff}}(0, z_n; G_n^*).$$

This is independent of the choice of exhaustion by Rayleigh.

Now we have

$$\mathbb{P}_0(\tau_0^+ = \infty) = \lim_{n \rightarrow \infty} \mathbb{P}_0(\tau_{z_n} < \tau_0^+) = \frac{1}{c(0)R_{\text{eff}}(0, z_n; G_n^*)} = \frac{1}{c(0)R_{\text{eff}}(0, \infty)}.$$

Definition (Flow). Let G be an infinite graph. A *flow* θ from 0 to ∞ is an anti-symmetric function on the edges such that $\text{div } \theta(x) = 0$ for all $x \neq 0$.

Theorem. Let G be an infinite connected graph with conductances $(c(e))_e$. Then

- (i) Random walk on G is recurrent iff $R_{\text{eff}}(0, \infty) = \infty$.
- (ii) The random walk is transient iff there exists a unit flow i from 0 to ∞ of finite energy

$$\mathcal{E}(i) = \sum_e (i(e))^2 r(e).$$

Proof. The first part is immediate since

$$\mathbb{P}_0(\tau_0^+ = \infty) = \frac{1}{c(0)R_{\text{eff}}(0, \infty)}.$$

To prove the second part, let θ be a unit flow from 0 to ∞ of finite energy. We want to show that the effective resistance is finite, and it suffices to extend Thompson's principle to infinite graphs.

Let i_n be the unit current flow from 0 to z_n in G_n^* . Let $v_n(x)$ be the associated potential. Then by Thompson,

$$R_{\text{eff}}(0, z_n; G_n^*) = \mathcal{E}(i_n).$$

Let θ_n be the restriction of θ to G_n^* . Then this is a unit flow on G_n^* from 0 to z_n . By Thompson, we have

$$\mathcal{E}(i_n) \leq \mathcal{E}(\theta_n) \leq \mathcal{E}(\theta) < \infty.$$

So

$$R_{\text{eff}}(0, \infty) = \lim_{n \rightarrow \infty} \mathcal{E}(i_n).$$

So by the first part, the flow is transient.

Conversely, if the random walk is transient, we want to construct a unit flow from 0 to ∞ of finite energy. The idea is to define it to be the limit of $(i_n)(x, y)$.

By the first part, we know $R_{\text{eff}}(0, \infty) = \lim \mathcal{E}(i_n) < \infty$. So there exists $M > 0$ such that $\mathcal{E}(i_n) \leq M$ for all n . Starting from 0, let $Y_n(x)$ be the number of visits to x up to time τ_{z_n} . Let $Y(x)$ be the total number of visits to x . Then $Y_n(x) \nearrow Y(x)$ as $n \rightarrow \infty$ almost surely.

By monotone convergence, we get that

$$\mathbb{E}_0[Y_n(x)] \nearrow \mathbb{E}_0[Y(x)] < \infty,$$

since the walk is transient. On the other hand, we know

$$E_0[Y_n(x)] = G_{\tau_{z_n}}(0, x).$$

It is then easy to check that $\frac{G_{\tau_{z_n}}(0, x)}{c(x)}$ is a harmonic function outside of 0 and z_n with value 0 at z_n . So it has to be equal to the voltage $v_n(x)$. So

$$v_n(x) = \frac{G_{\tau_{z_n}}(0, x)}{c(x)}.$$

Therefore there exists a function v such that

$$\lim_{n \rightarrow \infty} c(x)v_n(x) = c(x)v(x).$$

We define

$$i(x, y) = c(x, y)(v(x) - v(y)) = \lim_{n \rightarrow \infty} c(x, y)(v_n(x) - v_n(y)) = \lim_{n \rightarrow \infty} i_n(x, y).$$

Then by dominated convergence, we know $\mathcal{E}(i) \leq M$, and also i is a flow from 0 to ∞ . \square

Note that this connection with electrical networks only works for reversible Markov chains.

Corollary. Let $G' \subseteq G$ be connected graphs.

- (i) If a random walk on G is recurrent, then so is random walk on G' .
- (ii) If random walk on G' is transient, so is random walk on G .

Theorem (Polya's theorem). Random walk on \mathbb{Z}^2 is recurrent and transient on \mathbb{Z}^d for $d \geq 3$.

Proof sketch. For $d = 2$, if we glue all vertices at distance n from 0, then

$$R_{\text{eff}}(0, \infty) \geq \sum_{i=1}^{n-1} \frac{1}{8i-4} \geq c \cdot \log n$$

using the parallel and series laws. So $R_{\text{eff}}(0, \infty) = \infty$. So we are recurrent.

For $d = 3$, we can construct a flow as follows — let S be the sphere of radius $\frac{1}{4}$ centered at the origin. Given any edge \mathbf{e} , take the unit square centered at the midpoint \mathbf{m}_e of \mathbf{e} and has normal \mathbf{e} . We define $|\theta(\mathbf{e})|$ to be the area of the radial projection on the sphere, with positive sign iff $\langle \mathbf{m}_e, \mathbf{e} \rangle > 0$.

One checks that θ satisfies Kirchoff's node law outside of 0. Then we find that

$$\mathcal{E}(\theta) \leq C \sum_n n^2 \cdot \left(\frac{1}{n^2}\right)^2 < \infty,$$

since there are $\sim n^2$ edges at distance n , and the flow has magnitude $\sim \frac{1}{n^2}$. \square

Another proof sketch. We consider a binary tree T_ρ with edges joining generation to $n + 1$ having resistance ρ^n for ρ to be determined. Then

$$R_{\text{eff}}(T_\rho) = \sum_{n=1}^{\infty} \left(\frac{\rho}{2}\right)^n,$$

and this is finite when $\rho < 2$.

Now we want to embed this in \mathbb{Z}^3 in such a way that neighbours in T_ρ of generation n and $n + 1$ are separated by a path of length of order ρ^n . In generation n of T_ρ , there are 2^n vertices. On the other hand, the number of vertices at distance n in \mathbb{Z}^3 will be of order $(\rho^n)^2$. So we need $(\rho^n)^2 \geq 2^n$. So we need $\rho > \sqrt{2}$. We can then check that

$$R_{\text{eff}}(0, \infty; \mathbb{Z}^3) \leq c \cdot R_{\text{eff}}(T_\rho) < \infty. \quad \square$$

Note that this method is rather robust. We only need to be able to embed our T_ρ into \mathbb{Z}^3 . Using a similar method, Grimmett, Kesten and Zhang (1993) showed that simple random walk on a supercritical bond percolation is transient for $d \geq 3$.

3 Uniform spanning trees

3.1 Finite uniform spanning trees

The relation between uniform spanning trees and electrical networks was discovered by Kirchoff in 1850's. Of course, he was not interested in uniform spanning trees, but in electrical networks themselves. One of the theorems we will prove is

Theorem (Foster's theorem). Let $G = (V, E)$ be a finite weighted graph on n vertices. Then

$$\sum_{e \in E} R_{\text{eff}}(e) = n - 1.$$

We can prove this using the commute-time formula, but there is a one-line proof using this connection between electrical networks and uniform spanning trees.

Definition (Spanning tree). Let $G = (V, E)$ be a finite connected graph. A *spanning tree* T of G is a connected subgraph of G which is a tree (i.e. there are no cycles) and contains all the vertices in G .

Let \mathcal{T} be the set of all spanning trees of G . Pick $T \in \mathcal{T}$ uniformly at random. We call it the *uniform spanning tree*. We shall prove the following theorem:

Theorem. Let $e \neq f \in E$. Then

$$\mathbb{P}(e \in T \mid f \in T) \leq \mathbb{P}(e \in T).$$

It is tempting to ask a very similar question — if we instead considered all spanning *forests*, do we get the same result? Intuitively, we should, but it turns out this is an open question. So let's think about trees instead.

In order to prove this, we need the following result:

Theorem (Kirchoff). Let T be a uniform spanning tree, e an edge. Then

$$\mathbb{P}(e \in T) = R_{\text{eff}}(e)$$

Of course, this immediately implies Foster's theorem, since every tree has $n - 1$ edges.

Notation. Fix two vertices s, t of G . For all every edge $e = (a, b)$, define $\mathcal{N}(s, a, b, t)$ to be the set of spanning trees of G whose unique path from s to t passes along the edge (a, b) in the direction from a to b . Write

$$N(s, a, b, t) = |\mathcal{N}(s, a, b, t)|,$$

and N the total number of spanning trees.

Theorem. Define, for every edge $e = (a, b)$,

$$i(a, b) = \frac{N(s, a, b, t) - N(s, b, a, t)}{N}.$$

Then i is a unit flow from s to t satisfying Kirchoff's node law and the cycle law.

Note that from the expression for i , we know that

$$i(a, b) = \mathbb{P}(T \in \mathcal{N}(s, a, b, t)) - \mathbb{P}(T \in \mathcal{N}(s, b, a, t)).$$

Proof. We first show that i is a flow from s to t . It is clear that i is anti-symmetric. To check Kirchoff's node law, pick $a \notin \{s, t\}$. We need to show that

$$\sum_{x \sim a} i(a, x) = 0.$$

To show this, we count how much each spanning tree contributes to the sum. In each spanning tree, the unique path from s to t may or may not contain a . If it does not contain a , it contributes 0. If it contains A , then there is one edge entering a and one edge leaving a , and it also contributes 0. So every spanning tree contributes exactly zero to the sum.

So we have to prove that i satisfies the cycle law. Let $C = (v_1, v_2, \dots, v_{n+1} = v_1)$ be a cycle. We need to show that

$$\sum_{i=1}^n i(v_i, v_{i+1}) = 0.$$

To prove this, it is easier to work with "bushes" instead of trees. An s/t bush is a forest that consists of exactly 2 trees, T_s and T_t , such that $s \in T_s$ and $t \in T_t$. Let $e = (a, b)$ be an edge. Define $\mathcal{B}(s, a, b, t)$ to be the set of s/t bushes such that $a \in T_s$ and $b \in T_t$.

We claim that $|\mathcal{B}(s, a, b, t)| = N(s, a, b, t)$. Indeed, given a bush in $\mathcal{B}(s, a, b, t)$, we can add the edge $e = (a, b)$ to get a spanning tree whose unique path from s to t passes through e , and vice versa.

Instead of considering the contribution of each tree to the sum, we count the contribution of each bush to the set. Then this is easy. Let

$$\begin{aligned} F_+ &= |\{(v_j, v_{j+1}) : B \in \mathcal{B}(s, v_j, v_{j+1}, t)\}| \\ F_- &= |\{(v_j, v_{j+1}) : B \in \mathcal{B}(s, v_{j+1}, v_j, t)\}| \end{aligned}$$

Then the contribution of B is

$$\frac{F_+ - F_-}{N}.$$

By staring it at long enough, since we have a cycle, we realize that we must have $F_+ = F_-$. So we are done.

Finally, we need to show that i is a unit flow. In other words,

$$\sum_{x \sim s} i(s, x) = 1.$$

But this is clear, since each spanning tree contributes $\frac{1}{N}$ to $i(s, x)$, and there are N spanning trees. \square

We can now prove the theorem we wanted to prove.

Theorem. Let $e \neq f \in E$. Then

$$\mathbb{P}(e \in T \mid f \in T) \leq \mathbb{P}(e \in T).$$

Proof. Define the new graph $G.f$ to be the graph obtained by gluing both endpoints of f to a single vertex (keeping multiple edges if necessary). This gives a correspondence between spanning trees of G containing f and spanning trees of $G.f$. But

$$\mathbb{P}(e \in T \mid f \in T) = \frac{\text{number of spanning trees of } G.f \text{ containing } e}{\text{number of spanning trees of } G.f}.$$

But this is just $\mathbb{P}(e \in \text{UST of } G.f)$, and this is just $R_{\text{eff}}(e; G.f)$. So it suffices to show that

$$R_{\text{eff}}(e; G.f) \leq R_{\text{eff}}(e; G).$$

But this is clear by Rayleigh's monotone principle, since contracting f is the same as setting the resistance of the edge to 0. \square

In practice, how can we generate a uniform spanning tree? One way to do so is *Wilson's method*.

Definition (Loop erasure). Let $x = \langle x_1, \dots, x_n \rangle$ be a finite path in the graph G . We define the *loop erasure* as follows: for any pair $i < j$ such that $x_i = x_j$, remove $x_{i+1}, x_{i+2}, \dots, x_j$, and keep repeating until no such pairs exist.

To implement Wilson's algorithm, distinguish a root vertex of G , called r , and take an ordering of V . Set $T_0 = \{r\}$. Define T_i inductively as follows: take the first vertex in the ordering is not in T_i . We start a simple random walk from this vertex, and run until it hits T_i , and erase the loops. Then set T_{i+1} to be the union of T_i with this (loop-erased) path. Repeat until all vertices are used.

Theorem (Wilson). The resulting tree is a uniform spanning tree.

Note that in particular, the choice of ordering is irrelevant to the resulting distribution.

To prove this theorem, it is convenient to have a good model of how the simple random walk on G is generated. Under any vertex G , place an infinite number of cards that contain iid instructions, i.e. each card tells us to which neighbour we jump. Different vertices have independent stacks. Whenever we are at a vertex x , look at the top card underneath it and go to the neighbour it instructs. Then throw away the top card and continue.

Given stacks of cards under the vertices of the graph, by revealing the top card of each vertex, we get a directed graph, where directed edges are of the form (x, y) , where y is the instruction on the top card under x .

If there is no directed cycle, then we stop, and we have obtained a spanning tree. If there is a directed cycle, then we remove the top cards from all vertices on the cycle. We call this procedure *cycle popping*. Then we look at the top cards again and remove cycles and cards used.

We first prove a *deterministic* lemma about cycle popping.

Lemma. The order in which cycles are popped is irrelevant, in the sense that either the popping will never stop, or the same set of cycles will be popped, thus leaving the same spanning tree lying underneath.

Proof. Given an edge (x, S_x^i) , where S_x^i is the i th instruction under x , colour (x, S_x^i) with colour i . A colour is now coloured, but not necessarily with the same colour.

Suppose C is a cycle that can be popped in the order C_1, C_2, \dots, C_n , with $C_n = C$. Let C' be any cycle in the original directed graph. We claim that either C' does not intersect C_1, \dots, C_n , or $C' = C_k$ for some k , and C' does not intersect C_1, \dots, C_{k-1} . Indeed, if they intersect, let $x \in C' \cap C_k$, where k is the smallest index where they intersect. Then the edge coming out of x will have colour 1. Then S_x^1 is also in the intersection of C' and C_k , and so the same is true for the edge coming out of S_x^1 . Continuing this, we see that we must have $C_k = C'$. So popping $C_k, C_1, \dots, C_{k-1}, C_{k+1}, \dots, C_n$ gives the same result as popping C_1, \dots, C_n .

Thus, by induction, if C is a cycle that is popped, then after performing a finite number of pops, C is still a cycle that can be popped. So either there is an infinite number of cycles that can be popped, so popping can never stop, or every cycle that can be popped will be popped, thus in this case giving the same spanning tree. \square

Using this, we can prove that Wilson's method gives the correct distribution. In Wilson's algorithm, we pop cycles by erasing loops in the order of cycle creation. This procedure will stop with probability 1. With this procedure, we will reveal a finite set of coloured cycles O lying over a spanning tree. Let X be the set of all (O, T) , where O is the set of coloured cycles lying on top of T , which arise from running the random walks.

Now if we can get to (O, T) , then we could also get to (O, T') for any other spanning tree T' , since there is no restriction on what could be under the stacks. So we can write

$$X = X_1 \times X_2,$$

where X_1 is a certain set of finite sets of coloured cycles, and X_2 is the set of all spanning trees.

Define a function Ψ on the set of cycles by

$$\Psi(C) = \prod_{e \in C} p(e), \quad p(e) = \frac{c(e)}{c(e_-)},$$

where $e = (e_-, e_+)$, and $c(e_-)$ is the sum of the conductances of all edges with vertex e_- .

More generally, if O is a set of cycles, then we write

$$\Psi(O) = \prod_{C \in O} \Psi(C).$$

Then

$$\mathbb{P}(\text{get } (O, T) \text{ using Wilson's algorithm}) = \prod_{e \in \bigcup_{C \in O} C \cup T} p(e) = \Psi(O) \cdot \Psi(T).$$

Since this probability factorizes, it follows that the distribution of O is independent of T , and the probability of getting T is proportional to $\Psi(T)$. So T has the distribution of a uniform spanning tree.

Corollary (Cayley's formula). The number of labeled unrooted trees on n -vertices is equal to n^{n-2} .

Proof. Exercise. \square

How can we generalize Wilson's algorithm to infinite graphs? If the graph is recurrent, then we can still use the same algorithm, and with probability one, the algorithm terminates. If we have a transient graph, then we can still perform the same algorithm, but it is not clear what the resulting object will be.

The *Aldous–Broder algorithm* is another way of generating a uniform spanning tree. Run a random walk that visits all vertices at least once. For every vertex other than the root, take the edge the walk crossed the first time it visited the vertex, and keep the edge in the spanning tree. It is an exercise to prove that this does give a uniform spanning tree. Again, this algorithm works only on a recurrent graph. A generalization to transient graphs was found by Hutchcroft.

3.2 Infinite uniform spanning trees and forests

Since recurrent graphs are boring, let G be an infinite transient graph. Then a “naive” way to define a uniform spanning tree is by exhaustions.

Let G_n be an exhaustion of G by finite graphs. Define G_n^W (where W stands for “wired”) to be the graph G_n , where all of the vertices of $G \setminus G_n$ are glued to a single vertex z_n . The “free” G_n is just G_n considered as a subgraph. It turns out the wired one is much easier to study.

Let μ_n be the uniform spanning tree measure on G_n^W . Let B be a finite set of edges $B = \{e_1, \dots, e_n\}$. Then for T a UST of G_n^W , we have

$$\begin{aligned} \mu_n(B \subseteq T) &= \prod_{k=1}^m \mu(e_k \in T \mid e_j \in T \text{ for all } j < k) \\ &= \prod_{k=1}^n R_{\text{eff}}(e_k; G_n^W / \{e_j : j < k\}), \end{aligned}$$

where $G_n^W / \{e_j : j < k\}$ is the graph obtained from G_n^W by contracting the edges $\{e_j : j < k\}$.

We want to show that μ_n as a sequence of measures converges. So we want to understand how $R_{\text{eff}}(e_k; G_n^W / \{e_j : j < k\})$ changes when we replace n with $n + 1$. By Rayleigh's monotonicity principle, since G_n^W can be obtained from G_{n+1}^W by contracting more edges, we know the effective resistance increases when we replace n with $n + 1$. So we know that

$$\mu_n(B \subseteq T) \leq \mu_{n+1}(B \subseteq T).$$

So $(\mu_n(B \subseteq T))$ is an increasing sequence. So we can define

$$\mu(B \subseteq \mathcal{F}) = \lim_{n \rightarrow \infty} \mu_n(B \subseteq T),$$

and \mathcal{F} is our uniform forest. We can then extend the definition of μ to cylinder events

$$\{\mathcal{F} : B_1 \subseteq \mathcal{F}, B_2 \cap \mathcal{F} = \emptyset\}$$

with B_1, B_2 finite sets of edges, using inclusion-exclusion:

$$\mu(B_1 \subseteq \mathcal{F}, B_2 \cap \mathcal{F} = \emptyset) = \sum_{S \subseteq B_2} \mu(B_1 \cup S \subseteq \mathcal{F}) (-1)^{|S|}.$$

Then by commuting limits, we find that in fact

$$\mu(B_1 \subseteq \mathcal{F}, B_2 \cap \mathcal{F} = \emptyset) = \lim_{n \rightarrow \infty} \mu_n(B_1 \subseteq T, B_2 \cap T = \emptyset).$$

If \mathcal{A} is a finite union or intersection of cylinder sets of this form, then we have

$$\mu(\mathcal{A}) = \lim_{n \rightarrow \infty} \mu_n(\mathcal{A}).$$

So this gives a consistent family of measures.

By Kolmogorov's theorem, we get a unique probability measure μ supported on the forests of G . μ is called the *wired uniform spanning forest*. One can also consider the *free uniform spanning forest*, where we do the same thing but without gluing, however, this is more complicated and we will not study it.

We can adapt Wilson's algorithm to generate uniform spanning forests. This is called *Wilson's method rooted at infinity* (BKPS, 2001).

Let G be a transient graph, and start with $\mathcal{F}_0 = \emptyset$. Pick an ordering of the vertices of $V = \{x_1, x_2, \dots\}$. Inductively, from x_n , we start a simple random walk until the first time it hits \mathcal{F}_{n-1} if it does. If it doesn't, then run indefinitely. Call this (possibly infinite) path \mathcal{P}_n . Since G is transient, \mathcal{P}_n will visit every vertex finitely many times with probability 1. So it makes sense to define the loop erasure of \mathcal{P}_n . Call it $LE(\mathcal{P}_n)$. We set

$$\mathcal{F}_n = \mathcal{F}_{n-1} \cup LE(\mathcal{P}_n),$$

and take

$$\mathcal{F} = \bigcup_n \mathcal{F}_n.$$

As before, the order of V that we take does not affect the final distribution.

Proposition. Let G be a transient graph. The wired uniform spanning forest is the same as the spanning forest generated using Wilson's method rooted at infinity.

Proof. Let e_1, \dots, e_M be a finite set of edges. Let $T(n)$ be the uniform spanning tree on G_n^W . Let \mathcal{F} be the limiting law of $T(n)$. Look at G_n^W and generate $T(n)$ using Wilson's method rooted at z_n . Start the random walks from u_1, u_2, \dots, u_L in this order, where (u_i) are all the end points of the e_1, \dots, e_M (L could be less than $2M$ if the edges share end points).

Start the first walk from u_1 and wait until it hits z_n ; Then start from the next vertex and wait until it hits the previous path. We use the same infinite path to generate all the walks, i.e. for all i , let $(X_k(u_i))_{k \geq 0}$ be a simple random walk on G started from u_i . When considering G_n^W , we stop these walks when they exit G_n , .e. if they hit z_n . In this way, we couple all walks together and all the spanning trees.

Let τ_i^n be the first time the i th walk hits the tree the previous $i - 1$ walks have generated in G_n^W . Now

$$\mathbb{P}(e_1, \dots, e_M \in T(n)) = \mathbb{P} \left(e_1, \dots, e_M \in \bigcup_{j=1}^L LE(X_k(u_j) : k \leq \tau_j^n) \right).$$

Let τ_j be the stopping times corresponding to Wilson's method rooted at infinity. By induction on j , we have $\tau_j^n \rightarrow \tau_j$ as $n \rightarrow \infty$, and by transience, we have $LE(X_k(u_j) : k \leq \tau_j^n) \rightarrow LE(X_k(u_j) : k \leq \tau_j)$. So we are done. \square

Theorem (Pemantle, 1991). The uniform spanning forest on \mathbb{Z}^d is a single tree almost surely if and only if $d \leq 4$.

Note that the above proposition tells us

Proposition (Pemantle). The uniform spanning forest is a single tree iff starting from every vertex, a simple random walk intersects an independent loop-erased random walk infinitely many times with probability 1. Moreover, the probability that x and y are in the same tree of the uniform spanning forest is equal to the probability that simple random walk started from x intersects an independent loop-erased random walk started from y .

To further simplify the theorem, we use the following theorem of Lyons, Peres and Schramm:

Theorem (Lyons, Peres, Schramm). Two independent simple random walks intersect infinitely often with probability 1 if one walks intersects the loop-erasure of the other one infinitely often with probability 1.

Using these, we can now prove the theorem we wanted. Note that proving things about intersection, rather than collisions, tends to be higher, because any attempts to, say, define the first intersection time would inherently introduce an asymmetry.

We prove half of Pemantle's theorem.

Theorem. The uniform spanning forest is not a tree for $d \geq 5$ with probability 1.

Proof of Pemantle's theorem. Let X, Y be two independent simple random walks in \mathbb{Z}^d . Write

$$I = \sum_{t=0}^{\infty} \sum_{s=0}^{\infty} \mathbf{1}(X_t = Y_s).$$

Then we have

$$\mathbb{E}_{x,y}[I] = \sum_t \sum_s \mathbb{P}_{x-y}(X_{t+s} = 0) \approx \sum_{t=\|x-y\|} t \mathbb{P}_{x-y}(X_t = 0).$$

It is an elementary exercise to show that

$$\mathbb{P}_x(X_t = 0) \leq \frac{c}{t^{d/2}},$$

so

$$\mathbb{P}_x(X_t = 0) \leq \sum_{t=\|x-y\|} \frac{1}{t^{d/2-1}}.$$

For $d \geq 5$, for all $\varepsilon > 0$, we take x, y such that

$$\mathbb{E}_{x,y}[I] < \varepsilon.$$

Then

$$\mathbb{P}(\text{USF is connected}) \leq \mathbb{P}_{x,y}(I > 0) \leq \mathbb{E}_{x,y}[I] < \varepsilon.$$

Since this is true for every ε , it follows that $\mathbb{P}(\text{USF is connected}) = 0$. \square

Index

- $B(n)$, 10
- B_n , 10
- E , 5
- $R_{\text{eff}}(0, \infty)$, 41
- V , 5
- \mathbb{L}^d , 5
- $\chi(p)$, 16
- η_p , 5
- κ , 7
- λ , 7
- \leq , 13
- $\mathcal{C}(x)$, 5
- σ_n , 7
- θ_p , 5
- n th root trick, 26
- p_c^{slab} , 30
- s/t bush, 45
- $x \leftrightarrow y$, 5

- Aldous–Broder algorithm, 48

- BK inequality, 15
- bond percolation, 5

- Cayley’s formula, 47
- closed edge, 5
- cluster, 5
- commute time identity, 40
- connective constant, 7
- coupling, 5, 6
- Critical probability, 6
- current flow, 35
- cutset, 39
- cycle law, 35
- cycle popping, 46
- cylinder sets, 5

- decreasing event, 13
- decreasing random variable, 13
- disjoint occurrence, 15
- divergence, 33
- dual graph, 9

- edge
 - closed, 5
 - open, 5
- edge cutset, 39
- edges, 5

- effective conductance, 36
- Effective resistance, 36
- energy, 38
- event
 - decreasing, 13
 - increasing, 13
- exhaustion, 41

- Fekete’s lemma, 7
- FKG inequality, 13
- flow, 33, 34, 41
- Foster’s theorem, 44
- free uniform spanning forest, 49

- graph, 5
 - dual, 9
 - planar, 9
 - primal, 9
- Green kernel, 37
- Grimmett–Marstrand theorem, 30

- harmonic function, 34

- increasing event, 13
- increasing random variable, 13

- Kirchhoff’s node law, 34
- Kolmogorov 0-1 law, 10

- loop erasure, 46

- monotonicity principle, 38

- Nash–Williams inequality, 39

- Ohm’s law, 35
- open edge, 5

- percolation probability, 5
- pivotal edge, 17
- planar graph, 9
- Polya’s theorem, 42
- primal graph, 9

- random variable
 - decreasing, 13
 - increasing, 13
- Rayleigh’s monotonicity principle, 38
- RSW, 25
- Russo–Symour–Welsh, 25

- scaling limit, 8
- self-avoiding paths, 6, 7
- sink, 33
- site percolation, 5, 31
- slab percolation, 30
- source, 33
- spanning tree, 44
- strength, 34
- subadditive sequence, 7
- Thomson's principle, 38
- tree, 44
- trifurcation, 11
- uniform spanning tree, 44
- unit flow, 34
- vertices, 5
- voltage, 34
- Wilson's method, 46
- Wilson's method rooted at infinity, 49
- wired uniform spanning forest, 49