Percolation Theory

Arpon Basu

December 25, 2022

Contents

1 Notation and Convention

Throughout the report,

- We'll be representing the power-set of a set S by 2^S. Also, we'll always equip any power set with the inclusion partial order, ie:- if we say that for some $\omega_1, \omega_2 \in 2^S$ we have $\omega_2 \geq \omega_1$, then we mean $\omega_2 \supseteq \omega_1$.
- Given any set S, we will be using the terms 2^S and $\{0,1\}^S$ interchangeably. However, note a slight change of perspective in this interchange: When we say $\omega \in 2^S$, we are treating ω as a subset of S, while when we say $\omega \in \{0,1\}^S$, then we're treating ω as a **function** over S, where for any $s \in S$, $\omega(s) = 1$ if and only if $s \in \omega$.
- Let Ω be the sample space of a probability triple $(\Omega, \mathcal{F}, \mathbb{P})$. Further, assume that Ω is equipped with a partial order $\geq$. Then an event $\mathcal{E} \in \mathcal{F}$ is called *increasing*, if for any $\omega, \omega' \in \Omega$ we have

$$\omega \in \mathcal{E} \wedge \omega' \geq \omega \implies \omega' \in \mathcal{E}$$

- In the lattice $\mathbb{Z}^d$, we define the series of cubes $\Lambda_n := [-n, n]^d$ for every natural number $n \geq 1$. We also denote $x + \Lambda_n$ as $\Lambda_n(x)$ for any $x \in \mathbb{Z}^d$. Furthermore, let E_n be the set of edges of Λ_n.
- For any undirected graph $G := (V, E)$, and any subset of vertices $W \subset V$, we define the *boundary of* W to be

$$\partial W := \{w \in W : \exists w' \in V \setminus W, \{w, w'\} \in E\}$$

  Note that in our definition, $\partial W \subset W$ for any set W.
- Note that for a set of vertices W, ∂W is also a set of vertices. Similar to that definition, we can also define the boundary of edges to be

$$\triangle W := \{\{x, y\} \in E : x \in W, y \notin W\}$$

- Given an undirected graph $G = (V, E)$ and an edge $e \in E$, we call e a **bridge edge** if removing e from G increases the number of connected components in G.
- In this report, $\mathbb{E}_t[X]$ denotes the expectation of the random variable X as the parameter t *varies* over it's domain. In other words, $\mathbb{E}_t[X]$ is to be treated as a (deterministic) function of t.
- We use the $\sqcup$ symbol for *disjoint union*.

2 The Basic Setup

Our treatment of Bernoulli percolation follows the presentations given in [DC] and and [Ste].
For introducing percolation theory, we first need to establish a setting. Our setting is the (undirected) infinite lattice graph $\mathbb{L}_d = (\mathbb{V}, \mathbb{E})$ whose vertex set $\mathbb{V}$ is given by $\mathbb{Z}^d$, and whose edge set is given by the lattice itself, ie:-

$$\mathbb{E} := \{\{x, y\} : \|x - y\|_1 = 1, x, y \in \mathbb{Z}^d\}$$

Now, for every edge $e \in \mathbb{E}$, we independently assign a Bernoulli Random Variable X_e with parameter p. We call the edge e **open** if $X_e = 1$, and **closed** otherwise. The value of X_e shall henceforward be referred to as the *status* of the edge e.
We can now define the usual product measure on our space, with the probability triple $(\Omega, \mathcal{F}, \mathbb{P}_p)$, where the sample space Ω is $2^{\mathbb{E}}$, $\mathcal{F}$ is the σ-algebra generated by events dependent on only finitely many edges in $\mathbb{E}$, and $\mathbb{P}_p$ is the usual product measure on $\mathcal{F}$.

2.1 Interpretation Issues

For ease of interpretation, since every instance $\omega \in \Omega$ is a subset of $\mathbb{E}$, one can imagine that subset to be defining a subgraph of our lattice. Also note that when we say, for some $\omega, \omega' \in \Omega = 2^{\mathbb{E}}$, that $\omega' \geq \omega$, then in our graph-theoretic interpretation we can take this to mean that the graph of ω is a subgraph of ω', or equivalently, the graph of ω' can be obtained from that of ω by adding more edges into it. In this light, the definition of *increasing events* defined in Section 1 becomes much clearer: An event $\mathcal{E}$ (of graphs) is said to be increasing if, a certain graph ω belongs to $\mathcal{E}$, then any supergraph of ω must also belong to $\mathcal{E}$. Equivalently stated, adding edges to a graph can only increase the chance of an increasing event on it. Events like

- There exists an infinite cluster, ie:- our event $\mathcal{E}$ is the set of all infinite subgraphs of $\mathbb{L}$.
- There exists a path connecting **0** and **1**, ie:- our event $\mathcal{E}$ is the set of all subgraphs of $\mathbb{L}$ in which **0** and **1** belong to the same connected component.
- **0** belongs to an infinite cluster, ie:- our event $\mathcal{E}$ is the set of all subgraphs of $\mathbb{L}$ in which **0** belongs to an infinite connected component.

are standard examples of increasing events on graphs, in the context of percolation theory.
The entire investigation of percolation theory now focuses on "macro parameters" in this probability space, such as the existence of infinite clusters, the average size of a cluster containing the origin, and so on.
As we shall see below, what makes percolation theory so interesting is the natural occurrence of *phase transitions* in it: As we vary the parameter p, we shall see that many of the macroscopic quantities mentioned above change suddenly at a *critical probability*: The exact calculation of this critical probability, the equivalence of the critical probability arising from the phase transition of different quantities, and the probabilistic machinery needed to build all of this up will be the focus of the upcoming sections.

2.2 Definition of the infinite cluster probability

One can note that in our Bernoulli percolation process, the set of open edges forms various "clusters" on our lattice. For any point $x \in \mathbb{Z}^d$, one can formally define the cluster containing x to be the connected component containing x. In this light, we define

$$\theta(p,d) := \mathbb{P}_p(\mathbf{0} \text{ belongs to an infinite cluster})$$

Note that there is nothing special about **0**, it's just that our probability model is translationally invariant. At this stage, we also define our first critical probability

$$p_c(d) := \inf\{p \in [0,1] : \theta(p,d) > 0\}$$

Note that if $d = 1$, ie:- our lattice is 1-dimensional, and if $p < 1$, then almost surely there are no infinite clusters since infinite clusters in the $\mathbb{Z}$-lattice can only exist if there exists some $a \in \mathbb{Z}$ such that every edge in $(-\infty, a]$ or $[a, \infty)$ is open, which is an event of probability zero, and thus $p_c(1) = 1$.
Consequently, we now assume $d \geq 2$ hereon. We shall also suppress the dependence of θ and p_c on d for notational clarity.
Before we move on, we also define the **expected size of the cluster containing the origin**, ie:-

Definition 1. *Let C be the cluster containing the origin. Then the expected value of C is defined as*

$$\chi(p) := \sum_{n=0}^{\infty} n\mathbb{P}_p(|C| = n) + \infty \cdot \underbrace{\mathbb{P}_p(|C| = \infty)}_{=\theta(p)}$$

It is easy to see that χ is infinite if $\theta(p) > 0$, ie:- $p > p_c \implies \chi(p) = \infty$. One can also show that $p < p_c \implies \chi(p) < \infty$, but that is a much harder task, which we accomplish in Corollary 5.1.2.

It is very fruitful to express the χ function in terms of **2-point correlations**.

Lemma 2.1 (Two-point correlation)**.** *For any $p \in [0, 1]$, we have*

$$\chi(p) = \sum_{x \in \mathbb{Z}^d} \mathbb{P}_p(\mathbf{0} \longleftrightarrow x)$$

Proof. We know that

$$\chi(p) = \mathbb{E}[|C|] = \mathbb{E}\left[\sum_{\substack{C \subset \mathbb{Z}^d \\ \mathbf{0} \in C}} \sum_{c \in C} \mathbf{1}_{\mathbf{0} \longleftrightarrow c}\right] = \sum_{x \in \mathbb{Z}^d} \mathbb{P}_p(\mathbf{0} \longleftrightarrow x)$$

□

We also define

Definition 2. *For any $n \geq 1$ and $p \in [0, 1]$, define*

$$\theta_n(p) := \mathbb{P}_p(\mathbf{0} \longleftrightarrow \partial \Lambda_n)$$

Note that $\lim_{n \to \infty} \theta_n(p) = \theta(p)$.

The reason the definition of θ_n is so useful is that it allows us to study percolation in a finite, bounded, setting, and then many results about θ follow directly by passing to the limit $n \to \infty$. Just the definition of θ_n allows us to comment on the continuity properties of θ. Indeed,

Theorem 2.2. *$\theta : [0, 1] \mapsto [0, 1]$ is right continuous on $[0, 1]$.*

Proof. Observe that

1. $\theta_n(p)$ is a polynomial in p, and is thus continuous.
2. $\theta_n(p)$ is an increasing function of p [1].

Thus θ is a limit of continuous increasing functions, implying that it is right continuous.
It can also be shown that θ is continuous on $[0, 1] \setminus \{p_c\}$: Indeed, for $p \in [0, p_c)$, the result is trivial since θ is identically 0 on this interval. However, proving continuity on $(p_c, 1]$ requires some deep results, including showing that for $p > p_c$, there will be a unique infinite cluster almost surely. □

2.3 A First Introduction to Coupling Arguments

Coupling is a very useful technique in probability theory that allows us to compare two probabilities, especially when calculating them explicitly is difficult. Indeed, we shall see its power in the argument below, and the technique of coupling itself will be used many more times throughout our treatment of percolation theory.
By a *coupling argument*, we now claim that θ is a non-decreasing function of p: Indeed, let $p_1 < p_2$ be two parameters in $[0, 1]$. Then, for any instance $\omega \in 2^{\mathbb{E}}$ in the probability space generated by the parameter p_1, consider its set of closed

[1] this fact is intuitively clear: For a formal proof, refer Section 2.3

edges, ie:- $\mathbb{E} \setminus \omega =: \overline{\omega}$. For every edge $e \in \overline{\omega}$, with probability $\frac{p_2 - p_1}{1 - p_1}$, flip X_e from 0 to 1, and let $\tilde{\omega}$ be the set of open edges generated this way. Set $\omega' := \omega \sqcup \tilde{\omega}$. Then the probability space generated by ω' as ω varies over $2^{\mathbb{E}}$ is *exactly* the same as the one generated by the parameter p_2, since the probability that an edge $f \in \mathbb{E}$ belongs to ω' is

$$\Pr(f \in \omega') = \Pr(f \in \omega) + \Pr(f \in \tilde{\omega}) = p_1 + (1 - p_1)\frac{p_2 - p_1}{1 - p_1} = p_2$$

But since $\omega' \geq \omega$, we obtain that the probability of any *increasing event* is atleast as much in the p_2-space as in the p_1-space. Now we finish off by noting that the existence of an infinite cluster is an increasing event: Indeed, adding more edges doesn't change the infinitude of an infinite cluster.

2.4 Basic bounds on the critical probability

We first prove some basic lemmata.
We define a **self-avoiding path** (abbreviated as "SAP") to be a path of vertices in which no vertex is repeated.

Lemma 2.3. $\mathbf{0}$ *belonging to an infinite cluster (connected component) is equivalent to there being self-avoiding paths of arbitrarily large lengths containing* $\mathbf{0}$.

Proof. Suppose there are arbitrarily long SAPs originating from $\mathbf{0}$. Then the number of vertices connected to $\mathbf{0}$ can be arbitrarily large, and thus $\mathbf{0}$ is part of an infinite cluster.
Conversely, let $\mathbf{0}$ be part of an infinite cluster. Now, note that the degree of every vertex in our cluster is finite, ie:- our cluster is *locally finite* since $\mathbb{L}_d$ itself is locally finite with every vertex having degree $2d < \infty$. Thus our cluster is an infinite, connected, and locally finite graph, and consequently, by **König's lemma** in graph theory, there exists an infinite SAP in our cluster. If $\mathbf{0}$ is in this SAP, we're done. Otherwise, choose an arbitrary x in this SAP, and connect $\mathbf{0}$ via a path to x. Let y be the first member of the SAP on the path from $\mathbf{0}$ to x. Then $\mathbf{0} \longrightarrow y$, and then an infinite branch of the SAP originating at y forms an infinite SAP originating at $\mathbf{0}$. ◻

Let Ω_n be the set of SAPs in $\mathbb{Z}^d$ of length n originating at $\mathbf{0}$. Then it's easy to see that $|\Omega_n| \leq (2d)(2d-1)^{n-1}$.

Lemma 2.4. *For* $d \geq 2$, $p_c(d) \geq \frac{1}{2d-1}$.

Proof. Note that our result follows if we can show that $\theta(p) = 0$ on $[0, \frac{1}{2d-1})$.
Now, note that by Lemma 2.3

$$\theta(p) = \mathbb{P}_p(\mathbf{0} \text{ belongs to an infinite cluster}) = \mathbb{P}_p\left(\bigcap_{n \in \mathbb{N}} \{\exists \text{ open path in } \Omega_n\}\right)$$

But $\mathbb{P}_p\left(\exists \text{ open path in } \Omega_n\right) \leq |\Omega_n| p^n \leq 2d(2d-1)^{n-1}p^n \longrightarrow 0$ as $n \longrightarrow \infty$ for any $p < \frac{1}{2d-1}$. Consequently, $\theta(p) = 0$ on $[0, \frac{1}{2d-1})$. ◻

We can now state our first major theorem of percolation theory.

Theorem 2.5. *For* $d \geq 2$, $p_c(d) \in \left[\frac{1}{2d-1}, 0.9\right]$.

Proof. Note that if we can show that $p_c(d) \leq 0.9$, then we'll be done by Lemma 2.4. Also note that, by a coupling argument similar to the one given in Section 2.3, $\theta(p, d)$ is an increasing function of d for any given p: Indeed, for any two natural numbers $d < d'$, and any $\omega_{d'} \subset \mathbb{E}_{d'}$, one can obtain a $\omega_d \subset \mathbb{E}_d$ by setting $\omega_d := \omega_{d'} \cap \mathbb{Z}^d$.
Consequently, if $\theta(p, d) > 0$, then $\theta(p, d') > 0$ for all $d' > d$, and thus it suffices to prove our theorem for $d = 2$, ie:- show

that $\theta(0.9, 2) > 0$.[2]
We shall now construct the dual lattice of $\mathbb{Z}^2$ which we denote as $(\mathbb{Z}^2)^*$, where $(\mathbb{Z}^2)^*$ is defined as $\mathbb{Z}^2 + (\frac{1}{2}, \frac{1}{2})$, and let the dual lattice be $\mathbb{E}^*$. Note that there is a natural correspondence between elements of $\mathbb{E}$ and $\mathbb{E}^*$. Finally, an edge in $\mathbb{E}^*$ is set to be open if and only if the corresponding edge in $\mathbb{E}$ is closed. Effectively, the probability space on the dual lattice is $\mathbb{P}_{1-p} + (\frac{1}{2}, \frac{1}{2})$.
By a theorem of graph theory, it can be shown that $\mathbf{0}$ doesn't belong to an infinite cluster in $\mathbb{E}$ if and only if there is an open cycle C^* in $\mathbb{E}^*$ enclosing $\mathbf{0}$.
Thus

$$\begin{aligned} 1 - \theta(p) &\leq \sum_{n \geq 1} \Pr(\exists \text{ open cycle of length } n \text{ in } \mathbb{E}^* \text{ enclosing } \mathbf{0}) \\ &\leq \sum_{n \geq 1} |\text{cycles of length } n \text{ in } \mathbb{E}^*|(1-p)^n \end{aligned}$$

Similar to the calculation of $|\Omega_n|$, it's easy to see that the number of cycles of length n in a two-dimensional lattice is at most $n4 \cdot 3^{n-1}$. Thus

$$1 - \theta\,(0.9, 2) \leq \sum_{n \geq 1} n4 \cdot 3^{n-1}\,(1 - 0.9)^n < 1 \implies \theta(0.9, 2) > 0$$

as desired. □

Note:- The construction of the dual lattice and the subsequent argument is also known as a *Peierls argument*.

3 Crucial Inequalities

After establishing some basic results in the last section, we now need to develop some more technology to progress. This section has been modeled after [Gri99]. This section will require heavy use of coupling arguments, and also require one to recall the concept of increasing events.
Similar to the concept of increasing events, one can also define **increasing Random Variables**. Indeed, since a (real-valued) random variable X on a probability triple $(\Omega, \mathcal{F}, \mathbb{P})$ is just a function $X : \Omega \mapsto \mathbb{R}$, the R.V. X is called increasing if for any $\omega, \omega' \in \Omega$,

$$\omega' \geq \omega \implies X(\omega') \geq X(\omega)$$

Once we have defined increasing random variables, one may note that an increasing event is an event whose indicator function is an increasing random variable. Henceforward, we shall refer to events and their indicator random variables interchangeably. Again, similar to increasing events, we can use a coupling argument to present an inequality on increasing random variables.

Lemma 3.1. *Let X be an increasing random variable. Then*

$$p_1 \leq p_2 \implies (\mathbb{E}[X])_{p=p_1} \leq (\mathbb{E}[X])_{p=p_2}$$

Proof. Using the same coupling argument as Section 2.3, we can see that $X(\omega) \leq X(\omega')$, where ω and ω' are the same symbols as they were in Section 2.3. Taking an expectation of this inequality over Ω then yields our result. □

[2] Note that since $\theta(p, d)$ is an increasing function of d, $p_c(d)$ must be a decreasing function of d

3.1 The FKG Inequality

We now present the first of our inequalities, the FKG inequality [3], which confirms our intuition that the occurrence of an increasing event should positively influence the occurrence of another increasing event.

Theorem 3.2 (FKG Inequality)**.** *Let X and Y be increasing events in Ω such that $\mathbb{E}[X^2],\ \mathbb{E}[Y^2] < \infty$ then*

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

Proof. We first prove the theorem in the case that X and Y depend on finitely many $e \in \mathbb{E}$. To that end, for $n = 1$, we only have two states in the domain of X and Y, $\{(0), (1)\}$, and we note that for $\alpha, \beta \in \{(0), (1)\}$ (note that α, β aren't necessarily distinct)

$$\begin{aligned} &(X(\alpha) - X(\beta)) \cdot (Y(\alpha) - Y(\beta)) \geq 0 \\ \implies &\sum_{\alpha,\beta} (X(\alpha) - X(\beta)) \cdot (Y(\alpha) - Y(\beta)) \Pr(\alpha) \Pr(\beta) \geq 0 \\ \implies &2(\mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y]) \geq 0 \end{aligned}$$

Now assume for the sake of induction that the result is true for increasing random variables X and Y which are a function of $< k$ edges, and consider random variables X and Y which are functions of k edges, ie:-

$$X = X(\omega(1), \omega(2), \dots, \omega(k)), Y = Y(\omega(1), \omega(2), \dots, \omega(k))$$

$$X, Y : \{0, 1\}^k \mapsto \mathbb{R}$$

Then

$$\begin{aligned} \mathbb{E}[XY] &= \mathbb{E}_{\omega(1),\omega(2),\dots,\omega(k-1)} \left[\mathbb{E}_{\omega(k)}[XY|\omega(1), \omega(2), \dots, \omega(k-1)] \right] \\ &\geq \mathbb{E}_{\omega(1),\omega(2),\dots,\omega(k-1)} \left[\mathbb{E}_{\omega(k)}[X|\omega(1), \omega(2), \dots, \omega(k-1)] \cdot \mathbb{E}_{\omega(k)}[Y|\omega(1), \omega(2), \dots, \omega(k-1)] \right] \end{aligned}$$

where the inequality follows since $X_{\{\omega(i)\}_{1\leq i\leq k-1}}, Y_{\{\omega(i)\}_{1\leq i\leq k-1}}$ are increasing RVs of the single edge $\omega(k)$.
But note that $\mathbb{E}_{\omega(k)}[X|\omega(1), \omega(2), \dots, \omega(k-1)]$ is an increasing RV of $\{\omega(i)\}_{1\leq i\leq k-1}$, and thus

$$\begin{aligned} &\mathbb{E}_{\omega(1),\omega(2),\dots,\omega(k-1)} \left[\mathbb{E}_{\omega(k)}[X|\omega(1), \omega(2), \dots, \omega(k-1)] \cdot \mathbb{E}_{\omega(k)}[Y|\omega(1), \omega(2), \dots, \omega(k-1)] \right] \\ &\geq \mathbb{E}_{\omega(1),\omega(2),\dots,\omega(k-1)} \left[\mathbb{E}_{\omega(k)}[X|\omega(1), \omega(2), \dots, \omega(k-1)] \right] \cdot \mathbb{E}_{\omega(1),\omega(2),\dots,\omega(k-1)} \left[\mathbb{E}_{\omega(k)}[Y|\omega(1), \omega(2), \dots, \omega(k-1)] \right] \\ &= \mathbb{E}[X]\mathbb{E}[Y] \end{aligned}$$

Thus the theorem holds for increasing RVs of finitely many edges $e \in \mathbb{E}$. Now, since X and Y have finite second moments, $X_n \to X, Y_n \to Y$ as $n \to \infty$. Also, it can be seen that

$$\begin{aligned} \mathbb{E}\left[|X_nY_n - XY|\right] &\leq \mathbb{E}\left[|(X_n - X)Y_n| + |(Y_n - Y)X|\right] \\ &\leq \sqrt{\mathbb{E}\left[(X_n - X)^2\right]\mathbb{E}[Y_n^2]} + \sqrt{\mathbb{E}\left[(Y_n - Y)^2\right]\mathbb{E}[X^2]} \\ &\to 0 \end{aligned}$$

as $n \to \infty$. Thus $\mathbb{E}[X_nY_n] \geq \mathbb{E}[X_n]\mathbb{E}[Y_n]$ "passes" onto X and Y as $n \to \infty$, and hence the inequality holds. □

This result has many immediate corollaries, which we present below.

[3] the FKG inequality is named after its discoverers, Fortuin, Kasteleyn and Ginibre. It's alternatively also known as the Harris-FKG inequality

Corollary 3.2.1. *Let $\mathcal{A}$ and $\mathcal{B}$ be two increasing events. Then*

$$\Pr(\mathcal{A} \cap \mathcal{B}) \geq \Pr(\mathcal{A}) \Pr(\mathcal{B})$$

Corollary 3.2.2 (Square-Root Trick)**.** *Let $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ be increasing events. Then*

$$\max_{1 \leq i \leq n} \mathbb{P}_p(\mathcal{A}_i) \geq 1 - \left(1 - \mathbb{P}_p\left(\bigcup_{i=1}^{n} \mathcal{A}_i\right)\right)^{\frac{1}{n}}$$

Proof. Note that the FKG inequality holds for decreasing events too, where an event is defined to be decreasing if its complement is increasing.
Thus

$$\mathbb{P}_p\left(\bigcap_{1 \leq i \leq n} \overline{\mathcal{A}_i}\right) \geq \prod_{i=1}^{n} \mathbb{P}_p\left(\overline{\mathcal{A}_i}\right) \geq \left(\min_{1 \leq i \leq n} \mathbb{P}_p(\overline{\mathcal{A}_i})\right)^n = \left(1 - \max_{1 \leq i \leq n} \mathbb{P}_p(\mathcal{A}_i)\right)^n$$

and inequality follows. □

3.2 The BK Inequality

While the FKG inequality formalized our intuition that increasing events should be mutually positively correlated, the BK inequality [4] formalizes the intuition that the probability of two events required to happen "disjointly" is lesser than if no such conditions were imposed.
The following definitions will help formalize this intuition.

Definition 3. *Let $\mathcal{A}$ be an event (not necessarily increasing). For a given $\omega_0 \in A$, a subset $I = I_{\mathcal{A}}(\omega_0) \subset \mathbb{E}$ is said to be a* ***witness*** *of $\mathcal{A}$ for ω_0 if for any $\omega \in \Omega = 2^{\mathbb{E}}$, we have*

$$\omega \cap \omega_0 \supseteq I \implies \omega \in \mathcal{A}$$

Note that an event $\mathcal{A}$ is increasing if and only if every $\omega \in \mathcal{A}$ is a self-witness.

We also define

Definition 4. *Let $\mathcal{A}$ and $\mathcal{B}$ be events (not necessarily increasing). We then define $\mathcal{A} \circ \mathcal{B}$ to be*

$$\mathcal{A} \circ \mathcal{B} := \{\omega \in \mathcal{A} \cap \mathcal{B} : \text{there exist disjoint witnesses } I_{\mathcal{A}}(\omega), J_{\mathcal{B}}(\omega)\}$$

Note that $\mathcal{A} \circ \mathcal{B} \subseteq \mathcal{A} \cap \mathcal{B}$.

At the outset, we give some examples and clarifications to clear matters up. For example, if $\mathcal{A}$ and $\mathcal{B}$ depend on disjoint sets of edges, then $\mathcal{A} \circ \mathcal{B} = \mathcal{A} \cap \mathcal{B}$. Indeed, let $\mathcal{A}$ be the event that $\mathbf{0}$ is connected to some $x \in \partial\Lambda_5$ using edges only in Λ_5, and let $\mathcal{B}$ be the event that x is connected to some $y \in \mathbb{Z}^d \setminus \Lambda_5$ using only edges outside Λ_5 [5]. Then $\mathcal{A}$ and $\mathcal{B}$ depend on a disjoint set of edges, and thus they are bound to happen "disjointly" if they ever happen together, and thus $\mathcal{A} \circ \mathcal{B} = \mathcal{A} \cap \mathcal{B}$.
On the other hand, consider $x \in \mathbb{Z}^d \setminus \{\mathbf{0}\}$, and let $\mathcal{A}$ be the event that there is an open path between $\mathbf{0}$ and x. Also let $\mathcal{B} = \mathcal{A}$. Then $\mathcal{A} \cap \mathcal{B} = \mathcal{A}$, but $\mathcal{A} \circ \mathcal{A}$ is the event that there exist two **disjoint** paths between $\mathbf{0}$ and x, which is a proper subset of $\mathcal{A}$ in general.
We can now state the BK inequality.

[4] which stands for van der Berg-Kesten

[5] an edge e is defined to be outside a set S if at least one of the endpoints of e is not in S

Theorem 3.3 (BK Inequality)**.** *If $\mathcal{A}$, $\mathcal{B}$ are events dependent on only finitely many edges in $\mathbb{E}$, then*

$$\Pr(\mathcal{A} \circ \mathcal{B}) \leq \Pr(\mathcal{A}) \Pr(\mathcal{B})$$

Proof. Suppose $\mathcal{A}$ and $\mathcal{B}$ together are dependent on atmost $m < \infty$ edges. Enumerate those edges from 1 to m, ie:- let $[m] := \{1, 2, \ldots, m\}$ be the universe of all edges in this proof. We consequently also choose to restrict our probability space ($2^{\mathbb{E}}$) to just $2^{[m]} \cong \{0,1\}^m =: \Gamma$.
Now, the proof sketch is simple [6]: Consider $\omega \in (\mathcal{A} \cap \mathcal{B}) \setminus (\mathcal{A} \circ \mathcal{B})$. Then $\exists i \in [m]$ such that $i \in \omega$, and $\omega' := \omega \setminus \{i\} \notin \mathcal{A} \cap \mathcal{B}$, ie:- "deleting" the i^{th} edge from ω destroys both $\mathcal{A}$ and $\mathcal{B}$ as properties [7]. Now, "extend" our probability space Γ into $\Gamma' \cong \{0,1\}^{m+1}$ such that Γ' has an extra "copy" of the i^{th} index, say i', with a stipulation that while dealing with the i^{th} indices of members of $\mathcal{A}$, we'll mean i, while dealing with the i^{th} indices of members of $\mathcal{B}$, we'll mean i' [8]. Thus, ω's which aren't in $\mathcal{A} \circ \mathcal{B}$, because they were "blocked" by the index i, can now potentially be included, and thus the probability of $\mathcal{A} \circ \mathcal{B}$ in Γ' is more than what it was in Γ.
Extending this argument to every index i, our new probability space becomes $\Gamma \times \Gamma$, we get that

$$\Pr_{\Gamma}(\mathcal{A} \circ \mathcal{B}) \leq \Pr_{\Gamma \times \Gamma}(\mathcal{A} \circ \mathcal{B}) = \Pr_{\Gamma}(\mathcal{A}) \Pr_{\Gamma}(\mathcal{B})$$

as desired. □

Corollary 3.3.1. *If $\{\mathcal{A}_i\}_{1 \leq i \leq k}$ are events in Γ, then*

$$\Pr(\mathcal{A}_1 \circ \mathcal{A}_2 \circ \ldots \circ \mathcal{A}_k) \leq \Pr(\mathcal{A}_1) \Pr(\mathcal{A}_2) \ldots \Pr(\mathcal{A}_k)$$

This result can be used to derive very important bounds. But before that, a definition.

Definition 5. *Given two points $x, y \in \mathbb{Z}^d$, we denote the event "x is connected to y" by $x \longleftrightarrow y$. To denote the event that x percolates to ∞, ie:- the connected component containing x is infinite, we use the notation $x \longleftrightarrow \infty$.*
Given a set $S \subset \mathbb{Z}^d$, we denote $x \overset{S}{\longleftrightarrow} y$ if there exists a path τ connecting x and y such that the endpoints of every edge in τ, except possibly x and y, are in S.
Finally, for any two sets $A, B \subset \mathbb{Z}^d$, we also define

$$\{A \longleftrightarrow B\} := \bigcup_{a \in A} \bigcup_{b \in B} \{a \longleftrightarrow b\}$$

Lemma 3.4. *Consider a finite set $S \subset \mathbb{Z}^d$ such that $\mathbf{0} \in S$, and also consider a finite $X \subset \mathbb{Z}^d \setminus S$. Then*

$$\mathbb{P}_p(\mathbf{0} \longleftrightarrow X) \leq \sum_{y \in \partial S} \mathbb{P}_p(\mathbf{0} \overset{S}{\longleftrightarrow} y) \mathbb{P}_p(y \longleftrightarrow X)$$

Proof. Let n be any natural number. Consider the event $\mathcal{E}_n$ that $\mathbf{0}$ is connected to x **inside** Λ_n, ie:- $\mathcal{E}_n := \mathbf{0} \overset{\Lambda_n}{\longleftrightarrow} X$. Let τ be a self-avoiding path inside Λ_n connecting $\mathbf{0}$ and X, and let y be the first vertex of ∂S on τ. Note that such a y always exists since τ begins at $\mathbf{0}$ inside S, but ends up in X, outside S, and thus τ must cross the boundary of S somewhere. Then note that

$$\{\mathbf{0} \overset{\Lambda_n}{\longleftrightarrow} X\} = \bigcup_{y \in \partial S} \{\mathbf{0} \overset{S}{\longleftrightarrow} y\} \circ \{y \overset{\Lambda_n}{\longleftrightarrow} X\}$$

[6] it captures the essential ideas, but is not completely formal

[7] Had ω been in $\mathcal{A} \circ \mathcal{B}$, then we'd have been able to put the i^{th} edge into either $\mathcal{A}$ or $\mathcal{B}$, ie:- we'd have had $\omega' = \omega \setminus \{i\} \in (\mathcal{A} \setminus \mathcal{B}) \sqcup (\mathcal{B} \setminus \mathcal{A})$

[8] $\omega_{\mathcal{A}} \in \mathcal{A} \subseteq \Gamma$ are transported to Γ' with their $(i')^{\text{th}}$ indices zero, while $\omega_{\mathcal{B}} \in \mathcal{B}$ are transported such that their i^{th} indices become 0, and their $(i')^{\text{th}}$ indices inherit whatever the i^{th} index of $\omega_{\mathcal{B}}$ originally was. This is basically to separate the elements of $\mathcal{A}$ and $\mathcal{B}$ into different probability spaces

Indeed, every (self-avoiding) path in $\{\mathbf{0} \xleftrightarrow{\Lambda_n} X\}$ has a "first vertex" of ∂S inside it, and similarly, if there are disjoint paths between $\mathbf{0}$ and y and then y and X, there is a self-avoiding path between $\mathbf{0}$ and X too.
Thus

$$\mathbb{P}_p(\mathbf{0} \xleftrightarrow{\Lambda_n} X) \leq \sum_{y \in \partial S} \mathbb{P}_p(\{\mathbf{0} \xleftrightarrow{S} y\} \circ \{y \xleftrightarrow{\Lambda_n} X\})$$

Since all events here depend on only finitely many edges, we can use the BK inequality to get

$$\sum_{y \in \partial S} \mathbb{P}_p(\{\mathbf{0} \xleftrightarrow{S} y\} \circ \{y \xleftrightarrow{\Lambda_n} X\}) \leq \sum_{y \in \partial S} \mathbb{P}_p(\{\mathbf{0} \xleftrightarrow{S} y\})\mathbb{P}_p(\{y \xleftrightarrow{\Lambda_n} X\})$$

Finally, letting $n \to \infty$ yields our desired result. □

Corollary 3.4.1. *Consider a finite set $S \subset \mathbb{Z}^d$ such that $\mathbf{0} \in S$, and also consider $x \in \mathbb{Z}^d \setminus S$. Then*

$$\mathbb{P}_p(\mathbf{0} \longleftrightarrow x) \leq \sum_{y \in \partial S} \mathbb{P}_p(\mathbf{0} \xleftrightarrow{S} y)\mathbb{P}_p(y \longleftrightarrow x)$$

A quick definition before we get to an exciting result.

Definition 6. *For a given set $S \subset \mathbb{Z}^d$ define*

$$\varphi_p(S) := p \sum_{\{x,y\} \in \Delta S, x \in S} \mathbb{P}_p(\mathbf{0} \xleftrightarrow{S} x)$$

We set $\varphi_p(S) = 0$ if $\mathbf{0} \notin S$.

We can now finally state the crown jewel of this section, given below.

Theorem 3.5. *For a given $p \in [0, 1]$, if there exists a finite set S containing $\mathbf{0}$ such that $\varphi_p(S) < 1$, then there exists a constant $c = c(p) > 0$ such that for every $n \geq 1$ we have $\theta_n(p) \leq \exp(-cn)$.*

Proof. Let n_0 be an integer such that $S \subset \Lambda_{n_0 - 1}$. Then a proof similar to that of Lemma 3.4 gives us

$$\theta_{kn_0}(p) \leq p \sum_{\{x,y\} \in \Delta S, x \in S} \mathbb{P}_p(\mathbf{0} \xleftrightarrow{S} x)\mathbb{P}_p(y \longleftrightarrow \partial \Lambda_{kn_0}) \leq \varphi_p(S)\theta_{(k-1)n_0}(p)$$

where $\mathbb{P}_p(y \longleftrightarrow \partial\Lambda_{kn_0}) \leq \theta_{(k-1)n_0}(p)$ since y is at a distance at least $(k-1)n_0$ from $\partial\Lambda_{kn_0}$. Thus

$$\theta_{kn_0}(p) \leq \varphi_p(S)^k \implies \theta_n(p) \leq \varphi_p(S)^{\lfloor n/n_0 \rfloor}$$

thus showing exponential decay. □

4 The Margulis-Russo Formula

We now introduce the extremely useful Margulis-Russo formula. As with the BK inequality, we will only deal with events dependent on finitely many edges in this section, and thus let $[m] := \{1, 2, \ldots, m\}$ be the universe of edges for our discussion below.

Lemma 4.1. *Let $\Gamma := \{0,1\}^m$ be our probability space. For any **boolean** function $\mathbf{f} : \Gamma \mapsto \{0,1\}$, define*

$$f(p) := \mathbb{E}[\mathbf{f}(\omega)]$$

for $p \in [0,1]$. Then

$$f'(p) := \frac{df}{dp} = \frac{1}{p(1-p)} \sum_{i=1}^{m} \mathbb{E}[\mathbf{f}(\omega)(\omega(i) - p)]$$

Proof. Define $|\omega| := \sum_{i=1}^{m} \omega(i)$. Then

$$f(p) = \sum_{\omega \in \Gamma} \mathbf{f}(\omega) p^{|\omega|}(1-p)^{m-|\omega|}$$

$$\implies \frac{df}{dp} = \frac{1}{p} \sum_{\omega \in \Gamma} \mathbf{f}(\omega)|\omega| p^{|\omega|}(1-p)^{m-|\omega|} - \frac{1}{1-p} \sum_{\omega \in \Gamma} \mathbf{f}(\omega)(m - |\omega|) p^{|\omega|}(1-p)^{m-|\omega|}$$

$$= \frac{1}{p}\mathbb{E}[\mathbf{f}(\omega)|\omega|] - \frac{1}{1-p}\mathbb{E}[\mathbf{f}(\omega)(m - |\omega|)] = \frac{1}{p(1-p)} \sum_{i=1}^{m} \mathbb{E}[\mathbf{f}(\omega)(\omega(i) - p)]$$

□

We now introduce the very important notion of a **pivotal edge**.

Definition 7. *Let $\mathcal{A}$ be an event on Γ, and let $e \in [m]$ be an edge. We define*

$$\text{Piv}_e(\mathcal{A}) := \{\omega \in \Gamma : \omega \cup \{e\} \in \mathcal{A}, \omega \setminus \{e\} \notin \mathcal{A}\}$$

*ie:- given an edge e, e is **pivotal** to $\mathcal{A}$ containing ω, if ω lies in $\mathcal{A}$ if and only if the edge e is included in it. In other words, if one deletes e from ω, then the property of ω belonging to $\mathcal{A}$ is destroyed. $\text{Piv}_e(\mathcal{A})$ then, is the set of all pivotal edges of $\mathcal{A}$ w.r.t e. Note that $\text{Piv}_e(\mathcal{A})$ is independent of whether e is in ω, ie:- independent of $\omega(e)$.*

A few examples are in order:

1. Consider the event $\mathcal{A} := \{\mathbf{0} \longleftrightarrow \partial\Lambda_n\}$, and consider an edge e. Then $\text{Piv}_e(\mathcal{A})$ is the set of all subgraphs ω such that $\mathbf{0}$ and $\partial\Lambda_n$ are connected in $\omega' := \omega \cup \{e\}$, and e is a **bridge edge** in ω', ie:- removing e from ω' increases the number of connected components of the graph, and $\mathbf{0}$ and the vertices of $\partial\Lambda_n$ lie in different connected components of $\omega' \setminus \{e\}$.

2. On similar lines as the example above, consider $x \in \mathbb{Z}^d \setminus \{\mathbf{0}\}$, and let $\mathcal{A} := \{\mathbf{0} \longleftrightarrow x\}$, and let e be an edge. Then note that $\text{Piv}_e(\mathcal{A}) \cap (\mathcal{A} \circ \mathcal{A}) = \emptyset$. Indeed, since $\mathcal{A} \circ \mathcal{A}$ contains only those ω for which there are two disjoint paths between $\mathbf{0}$ and x, disconnecting a single edge will not disconnect 0 and x, and thus e won't be a pivotal edge.

With this, we now state the Margulis-Russo formula.

Theorem 4.2. *Let $\mathcal{A}$ be an increasing event that depends on finitely many edges, let $\mathbf{f} := \mathbf{1}_{\mathcal{A}}$ be the characteristic function of $\mathcal{A}$, and let $f(p)$ be defined as in Lemma 4.1. Then*

$$f'(p) = \sum_{e \in [m]} \mathbb{P}_p(\text{Piv}_e(\mathcal{A})) = \sum_{e \in [m]} \mathbb{P}_p(e \text{ is pivotal for } \mathcal{A})$$

Proof. Note that by Lemma 4.1, we'll be done if we can show that

$$\mathbb{E}[\mathbf{f}(\omega)(\omega(e)-p)] = p(1-p)\mathbb{P}_p(\mathrm{Piv}_e(\mathcal{A}))$$

Now,

$$\mathbb{E}[\mathbf{f}(\omega)(\omega(e)-p)\mathbf{1}_{\overline{\mathrm{Piv}_e(\mathcal{A})}}] = \mathbb{E}[\mathbf{1}_{\mathcal{A}}(\omega(e)-p)\mathbf{1}_{\overline{\mathrm{Piv}_e(\mathcal{A})}}] = \mathbb{E}[\mathbf{1}_{\mathcal{A}\setminus\mathrm{Piv}_e(\mathcal{A})}(\omega(e)-p)]$$

But if $\omega \in \mathcal{A} \setminus \mathrm{Piv}_e(\mathcal{A})$, then

1. $\omega \in \mathcal{A} \setminus \mathrm{Piv}_e(\mathcal{A}) \implies \omega \in \mathcal{A} \implies \omega \cup \{e\} \in \mathcal{A}$, since $\mathcal{A}$ is an increasing event.
2. $\omega \in \mathcal{A} \setminus \mathrm{Piv}_e(\mathcal{A}) \implies \omega \notin \mathrm{Piv}_e(\mathcal{A})$, which means that e isn't pivotal for ω, implying, along with the first point, that $\omega \setminus \{e\} \in \mathcal{A}$.
3. The first and second points together imply that ω belonging to $\mathcal{A} \setminus \mathrm{Piv}_e(\mathcal{A})$ is independent of the status of $\omega(e)$.

Consequently,

$$\mathbb{E}[\mathbf{1}_{\mathcal{A}\setminus\mathrm{Piv}_e(\mathcal{A})}(\omega(e)-p)] = \mathbb{E}[\mathbf{1}_{\mathcal{A}\setminus\mathrm{Piv}_e(\mathcal{A})}]\mathbb{E}[(\omega(e)-p)] = \mathbb{E}[\mathbf{1}_{\mathcal{A}\setminus\mathrm{Piv}_e(\mathcal{A})}]\cdot 0 = 0$$

Thus

$$\mathbb{E}[\mathbf{f}(\omega)(\omega(e)-p)] = \mathbb{E}[\mathbf{f}(\omega)(\omega(e)-p)\mathbf{1}_{\mathrm{Piv}_e(\mathcal{A})}] = \mathbb{E}[\mathbf{1}_{\mathcal{A}\cap\mathrm{Piv}_e(\mathcal{A})}(\omega(e)-p)]$$

Now, consider some $\omega \in \mathcal{A} \cap \mathrm{Piv}_e(\mathcal{A})$. Since $\omega \in \mathrm{Piv}_e(\mathcal{A})$, $\omega \setminus \{e\} \notin \mathcal{A}$. But we also have $\omega \in \mathcal{A}$. Consequently, we must have $e \in \omega$, ie:- $\omega(e) = 1$. Thus, $\mathcal{A} \cap \mathrm{Piv}_e(\mathcal{A}) = \mathrm{Piv}_e(\mathcal{A}) \cap \{\omega(e) = 1\}$. Then

$$\mathbb{E}[\mathbf{1}_{\mathcal{A}\cap\mathrm{Piv}_e(\mathcal{A})}(\omega(e)-p)] = \mathbb{E}[\mathbf{1}_{\mathrm{Piv}_e(\mathcal{A})\cap\{\omega(e)=1\}}(\omega(e)-p)] = (1-p)\mathbb{E}[\mathbf{1}_{\mathrm{Piv}_e(\mathcal{A})\cap\{\omega(e)=1\}}]$$

$$= (1-p)\mathbb{P}_p(\mathrm{Piv}_e(\mathcal{A}) \cap \{\omega(e)=1\}) = (1-p)\mathbb{P}_p(\mathrm{Piv}_e(\mathcal{A}))\mathbb{P}_p(\omega(e)=1) = p(1-p)\mathbb{P}_p(\mathrm{Piv}_e(\mathcal{A}))$$

where the second last equality follows because $\mathrm{Piv}_e(\mathcal{A})$ and $\{\omega(e)=1\}$ are independent events. □

We can now prove a useful lemma.

Lemma 4.3. *Fix a $n \geq 1$. Define the set $\mathcal{S}_n$ to be:*

$$\mathcal{S}_n := \{s \in \Lambda_n : s \not\longleftrightarrow \partial\Lambda_n\}$$

Then for any $p \in (0,1)$ we have

$$\theta_n'(p) := \frac{d\theta_n(p)}{dp} = \frac{1}{p(1-p)}\mathbb{E}[\varphi_p(\mathcal{S}_n)]$$

where $\theta_n(p)$ and $\varphi_p(\cdot)$ are as they were defined in Definition 2 and Definition 6, respectively.

Proof. Let E_n be the set of edges between vertices in Λ_n. Then by the Margulis-Russo formula

$$\theta_n'(p) = \sum_{e\in E_n} \mathbb{P}_p(e \text{ is pivotal for } \mathbf{0} \longleftrightarrow \partial\Lambda_n)$$

$$= \frac{1}{1-p}\sum_{e\in E_n} \mathbb{P}_p(e \text{ is pivotal for } \mathbf{0} \longleftrightarrow \partial\Lambda_n \text{ and } \omega(e) = 0)$$

where the equality in the second line follows since the status of $\omega(e)$ is independent of ω belonging to $\mathrm{Piv}_e(\mathcal{A})$.
Now, since $e = \{x, y\}$ is pivotal for $\mathbf{0} \longleftrightarrow \partial\Lambda_n$, it is **a bridge edge** for ω. Thus, deleting e from ω increases the number

of connected components of ω, and the component (let's call it S) containing $\mathbf{0}$, is not connected to $\partial\Lambda_n$, and thus S is a valid candidate for $\mathcal{S}_n$ [9]. Moreover, note also that $x \in S$, and $y \notin S$. Thus the above expression is equal to

$$= \frac{1}{1-p} \sum_{S \subset \Lambda_n, \{x,y\} \in \triangle S} \mathbb{P}_p(\mathbf{0} \overset{S}{\longleftrightarrow} x, \mathcal{S}_n = S)$$

Now note that the event $\{\mathbf{0} \overset{S}{\longleftrightarrow} x\}$ depends on edges inside S, while the event $\{\mathcal{S}_n = S\}$ depends on edges outside S [10]. Consequently, they're independent events, and the above expression equals

$$= \frac{1}{1-p} \sum_{S \subset \Lambda_n, \{x,y\} \in \triangle S} \mathbb{P}_p(\mathbf{0} \overset{S}{\longleftrightarrow} x)\mathbb{P}_p(\mathcal{S}_n = S)$$

$$= \frac{1}{1-p} \sum_{S \subset \Lambda_n} \left(\sum_{\{x,y\} \in \triangle S} \mathbb{P}_p(\mathbf{0} \overset{S}{\longleftrightarrow} x) \right) \mathbb{P}_p(\mathcal{S}_n = S)$$

$$= \frac{1}{1-p} \mathbb{E}\left[\frac{1}{p}\varphi_p(\mathcal{S}_n)\right]$$

as desired. □

5 Behavior in non-critical Régimes

In this section, we shall investigate the behavior of various macroscopic quantities under 2 "régimes": When our parameter p is lesser than, and greater than the critical probability p_c.

5.1 Sub-critical Régime ($p < p_c$)

We shall begin right away with one of the most important results of this section.

Theorem 5.1 (Exponential Decay of the Diameter)**.** *Fix a dimension $d \geq 2$. Then for every $p < p_c$, there exists a constant $c_1 = c(p) > 0$ such that for every $n \geq 1$, we have*

$$\theta_n(p) = \mathbb{P}_p(\mathbf{0} \longleftrightarrow \partial\Lambda_n) \leq \exp(-c_1 n)$$

There also exists a constant $c_2 > 0$ such that for all $p \geq p_c$, we have $\theta(p) \geq c_2(p - p_c)$.

Proof. Define

$$\tilde{p_c} := \sup\{p \in [0,1] : \exists \text{ finite set } S \ni \mathbf{0} \text{ such that } \varphi_p(S) < 1\}$$

We claim that $\tilde{p_c} = p_c$: Indeed, if $p < \tilde{p_c}$, then $\theta_n(p)$ decays exponentially as a function of n by Theorem 3.5, and consequently $\theta(p) = 0$, implying $p < p_c$, which further implies that $\tilde{p_c} \leq p_c$.
Conversely, fix an arbitrary $p \in (\tilde{p_c}, 1]$. Then for every finite set $S \ni \mathbf{0}$, we have $\varphi_p(S) \geq 1$. Thus by Lemma 4.3 we have

$$\theta_n'(p) = \frac{1}{p(1-p)}\mathbb{E}[\varphi_p(\mathcal{S}_n)]$$

[9] conversely, every $\mathcal{S}_n$ containing $\mathbf{0}$ and not connected to $\partial\Lambda_n$ can be a 'S' of some ω
[10] since $\mathcal{S}_n = S$, edges outside S must conspire to prevent $\mathcal{S}_n$ from being connected to $\partial\Lambda_n$

But note that

$$\mathbb{E}[\varphi_p(\mathcal{S}_n)] = \sum_{S\subset\Lambda_n\setminus\partial\Lambda_n} \varphi_p(S)\mathbb{P}_p(\mathcal{S}_n = S) \geq \sum_{\substack{S\subset\Lambda_n\setminus\partial\Lambda_n\\ S\ni\mathbf{0}}} \varphi_p(S)\mathbb{P}_p(\mathcal{S}_n = S)$$

$$\geq \sum_{\substack{S\subset\Lambda_n\setminus\partial\Lambda_n\\ S\ni\mathbf{0}}} \mathbb{P}_p(\mathcal{S}_n = S) \geq \mathbb{P}_p(\mathcal{S}_n \ni \mathbf{0}) = 1 - \theta_n(p)$$

Thus

$$\theta_n'(p) \geq \frac{1-\theta_n(p)}{p(1-p)} \implies \left[\log\left(\frac{1}{1-\theta_n}\right)\right]' \geq \left[\log\left(\frac{p}{1-p}\right)\right]'$$

Integrating the inequality between $\tilde{p}_c$ and p yields

$$\theta_n(p) \geq \frac{p-\tilde{p}_c}{p(1-\tilde{p}_c)} \implies \theta(p) \geq \frac{p-\tilde{p}_c}{p(1-\tilde{p}_c)}$$

Consequently, since $p > \tilde{p}_c$, $\theta(p) > 0 \implies p > p_c$. Since p was arbitrary, we have $\tilde{p}_c \geq p_c \implies \tilde{p}_c = p_c$. Consequently, by Theorem 3.5 we have $\theta_n(p) \leq \exp(-c_1 n)$. We also have, by the derivation above, that for $p > p_c = \tilde{p}_c$, $\theta(p) \geq \frac{p-p_c}{p(1-p_c)} \geq \frac{p-p_c}{1-p_c}$, and thus the second assertion of our statement is also proved, with $c_2 = \frac{1}{1-p_c}$. □

Corollary 5.1.1. *For every $p < p_c$, there exists a constant $c_p > 0$ such that the probability that there exists a cluster of radius larger than $c_p \log n$ in the box of size Λ_n tends to 0 as n tends to infinity.*

Corollary 5.1.2. *If $p < p_c$, then $\chi(p) < \infty$, ie:- the expected size of a cluster is finite in the sub-critical régime. This settles a question we had raised in Definition 1.*

Proof. Set $M := \max\{n : \mathbf{0} \longleftrightarrow \partial\Lambda_n \text{ happens}\}$. Since $p < p_c$, $\mathbb{P}_p(M < \infty) = 1$. Also, let C be the connected component containing $\mathbf{0}$. Then

$$\chi(p) = \sum_{n=0}^{\infty} n\mathbb{P}_p(|C| = n) = \sum_{k=0}^{\infty} \mathbb{P}_p(M = k) \sum_{n=0}^{\infty} n\mathbb{P}_p(|C| = n \text{ and } M = k)$$

$$\leq \sum_{k=0}^{\infty} \mathbb{P}_p(M = k) \sum_{n=0}^{\infty} n \frac{\mathbb{P}_p(|C| = n \text{ and } M = k)}{\mathbb{P}_p(M = k)} = \sum_{k=0}^{\infty} \mathbb{P}_p(M = k)\mathbb{E}\left[|C|\middle|M = k\right]$$

$$\leq \sum_{k=0}^{\infty} \mathbb{P}_p(M = k)|B_k| \leq \sum_{k=0}^{\infty} \exp(-ck)O(k^d) < \infty$$

where B_k is the set of all points whose ℓ_1-distance from $\mathbf{0}$ is at most k. □

Corollary 5.1.3 (Sub-exponential decay of volume)**.** *Fix a $p < p_c$. Let C be the connected component containing $\mathbf{0}$. Then there exists a constant $c_p > 0$ such that $\mathbb{P}_p(|C| > n) \leq \exp(-c_p n^{1/d})$.*

Proof. Suppose $|C| > n$. Then it's not too difficult to see that $C \cap \partial B_k \neq \emptyset$ for some $k = \Theta(n^{1/d})$. Then

$$\mathbb{P}_p(|C| > n) \leq \theta_k(p) = \exp(-\Theta(n^{1/d}))$$

as desired. □

5.2 Super-critical Régime ($p > p_c$)

From our discussions so far, we already know quite a few things about percolation in the super-critical régime. For example,

1. The probability of the existence of an infinite cluster is non-zero.
2. The expected size of the cluster containing **0** is infinite.
3. $\theta(p)$ grows super-linearly in p: Indeed, $\theta(p) \geq \frac{p-p_c}{p(1-p_c)}$ for $p > p_c$.

Since the super-critical phase is when infinite clusters exist in our lattice, the natural question to ask is: How many infinite clusters do we have? Answering this will require us to use some foundational results of probability theory, such as Kolmogorov's $0-1$ law. But before that, a short observation on our probability space $(2^{\mathbb{E}}, \mathcal{F}, \mathbb{P}_p)$: Note that the measure $\mathbb{P}_p$ is translation invariant: In other words, let $\mathcal{A}$ be any event in $\mathcal{F}$, and let $\tau_x(\mathcal{A})$ be a translation of $\mathcal{A}$, ie:- for every $\omega \in \mathcal{A}$, translate ω by x to get the corresponding element of $\tau_x(\mathcal{A})$. Then observe that $\mathbb{P}_p(\mathcal{A}) = \mathbb{P}_p(\tau_x(\mathcal{A}))$.
One can also show that $\mathbb{P}_p$ is **ergodic**, ie:- **every translation invariant event** [11] **has probability** 0 **or** 1. Consequently, we have

Lemma 5.2. *For $p > p_c$, there exists an infinite cluster **almost surely, ie:- with probability 1**.*

Proof. Note that the existence of an infinite cluster is a translation invariant event. Since the probability of the existence of an infinite cluster is > 0 for $p > p_c$, it can only be 1. □

Lemma 5.3. *Fix $p > p_c$. Let X be the random variable denoting the number of infinite clusters. Then $X = 1$ or ∞ almost surely.*

Proof. Note that $\mathcal{E}_k := \{X = k\}$, where $k \in \mathbb{N} \cup \{\infty\}$, is a translation invariant event, and thus has probability 0 or 1. Thus there exists a specific $k = k_0$ for which $\mathbb{P}_p(\mathcal{E}_{k_0}) = 1$, and for all other $k \neq k_0$, we have $\mathbb{P}_p(\mathcal{E}_k) = 0$.
Thus we must show that $k_0 = 1$ or ∞. Assume for the sake of contradiction that k_0 is some finite integer greater than 1. Thus, $\mathbb{P}_p(\mathcal{E}_1) = 0$. Now, consider the sequence of events $\mathcal{L}_1 \subseteq \mathcal{L}_2 \subseteq \ldots$, where $\mathcal{L}_i$ denotes the event that Λ_i intersects all k_0 clusters. Note that $\mathbb{P}_p(\mathcal{L}_i)$ is an increasing sequence, and also note that since $\bigcup_{i \in \mathbb{N}} \mathcal{L}_i = \mathcal{E}_{k_0}$, we have $\lim_{i \to \infty} \mathbb{P}_p(\mathcal{L}_i) = 1$. Consequently, there exists $i = i_0$ such that $\mathbb{P}_p(\mathcal{L}_{i_0}) > \frac{1}{2}\mathbb{P}_p(\mathcal{E}_{k_0}) = \frac{1}{2}$.
Now, let $\mathcal{E}$ be the event that all edges in Λ_{i_0} are open. Then note that $\mathcal{E}$ and $\mathcal{L}_{i_0}$ are independent: Indeed, changing the status of edges inside Λ_{i_0} may change how deep an infinite cluster "penetrated it", but it won't change the fact that the infinite cluster intersected it at the first place, because that is dependent on edges outside Λ_{i_0}. Consequently,

$$\mathbb{P}_p(\mathcal{E} \cap \mathcal{L}_{i_0}) = \mathbb{P}_p(\mathcal{E})\mathbb{P}_p(\mathcal{L}_{i_0}) > \frac{p^{|E_n|}}{2} > 0$$

But note that since all clusters intersected Λ_{i_0} in the event $\mathcal{L}_{i_0}$, in the event $\mathcal{E} \cap \mathcal{L}_{i_0}$ we have that all clusters amalgamate into one since all edges in our cube are opened, implying that $\mathbb{P}_p(\mathcal{E}_1) > 0$, which is a contradiction. □

We now show that $X = 1$ almost surely, ie:- almost surely we don't have infinitely many infinite clusters. However, to do that, we need to define the notion of a *trifurcation point*, also known as an *encounter point* in some sources. To motivate the definition, we first state an elementary lemma.

Lemma 5.4. *Let T be a tree. Let k_i be the number of vertices of T with degree i. Then $k_3 + 2 \leq k_1$.* A fortiori, *we have $k_3 < k_1$.*

[11] ie:- events $\mathcal{A}$ such that $\mathcal{A} = \tau_x(\mathcal{A})$ for every $x \in \mathbb{Z}^d$

Proof. Note that $\sum_{i=1}^{n-1} ik_i$ is the sum of degrees of all vertices of T. But we know that quantity to be $2(n-1)$ since a tree has $n-1$ edges. Thus

$$\sum_{i=1}^{n}(i-2)k_i = 2(n-1) - 2\sum_{i=1}^{n-1} k_i = 2(n-1) - 2n = -2$$

But we also have

$$\sum_{i=1}^{n}(i-2)k_i = -k_1 + k_3 + 2k_4 + \ldots \geq k_3 - k_1$$

□

Note:- The above result immediately implies $k_3 < k_1$ for forests too.
As promised, we now define trifurcation points.

Definition 8. *Let C be a (connected) infinite cluster. A point $x \in C$ is called a* ***trifurcation point*** *if $C \setminus \{x\}$ has exactly 3 connected components, all of them infinite. Note that the degree of a trifurcation point is exactly 3.*

We now prove the most important part of our proof that $X \neq \infty$ almost surely.

Lemma 5.5. *Let T be the maximum number of trifurcation points contained in Λ_n, under any configuration $\omega \in 2^{\mathbb{E}}$. Then $T < |\partial\Lambda_n|$.*

Proof. Let ω be any configuration. Consider the graph $\omega_0 := \omega \cap \Lambda_n$. Delete the minimum number of edges possible from ω_0 so that it becomes a forest. Call this forest ω'. Enumerate the edges in ω' as $\{e_1, e_2, \ldots, e_s\}$. Now, run the iteration below, where our iterating index i starts from 1 and goes till s:

1. Let our forest at this stage be ω'_i. We set $\omega'_1 = \omega'$.
2. If $\omega'_i \setminus \{e_i\}$ contains any connected component α such that $\alpha \cap \partial\Lambda_n = \emptyset$, then delete e_i and α from ω'_i. Otherwise, leave ω'_i as it is.

Now, notice that all the leaves (ie:- vertices of degree 1) of the forest left after this iteration lie in $\partial\Lambda_n$, and also notice that any trifurcation points inside Λ_n must be vertices of this forest. Since trifurcation points have a degree of 3, they must be lesser in number than the leaves, which can be at most $|\partial\Lambda_n|$, as desired. □

Theorem 5.6. *Let $p > p_c$. Then we have a unique infinite cluster almost surely.*

Proof. By the preceding discussion, the only thing left to prove is that the probability of having infinitely many infinite clusters is 0.
To that end, let $\mathcal{T}_0$ be the event that $\mathbf{0}$ is a trifurcation point. Let $k = k(d)$ be a large enough integer such that if k infinite clusters intersect Λ_n, then we have 3 points $x, y, z \in \partial\Lambda_n$ such that x, y, z are far away enough from each other to ensure that there exist 3 edge-disjoint paths from $\mathbf{0}$ to them. By an argument similar to that given in the proof of Lemma 5.3, we can find a large enough n such that at least k infinite clusters intersect Λ_n with some positive probability $\delta > 0$, and consequently the probability of the aforementioned x, y, z existing is also positive.
Now, choose 3 edge-disjoint paths π_x, π_y, π_z between $\mathbf{0}$ and x, y, z respectively, and change the status of edges within Λ_n so that only the edges on π_x, π_y, π_z remain open, and all other edges are closed. Then

$$\mathbb{P}_p(\mathcal{T}_0) \geq \delta \cdot \big(p(1-p)\big)^{|E_n|} = \eta > 0$$

Indeed, the inequality follows because the event that "more than k infinite clusters intersect Λ_n", and independently [12], "exactly a certain configuration of edges are open in Λ_n", is a subset of the event $\mathcal{T}_0$.
Now, note that by translational invariance, the probability that some $x \in \Lambda_n$ is a trifurcation point is equal to the probability that $\mathbf{0}$ is a trifurcation point. Consequently, if T is a random variable denoting the number of trifurcation points in Λ_n, then we have

$$\mathbb{E}[T] = \mathbb{P}_p(\mathcal{T}_0) \cdot |\Lambda_n| = \eta|\Lambda_n|$$

But on the other hand we also have $T \leq |\partial\Lambda_n|$, and consequently, we have that for all n large enough, $\eta|\Lambda_n| < |\partial\Lambda_n|$, which is a contradiction since $\lim_{n\to\infty} \frac{|\partial\Lambda_n|}{|\Lambda_n|} = 0$. □

6 Behavior at Criticality and the Computation of p_c

Unlike our study of behavior under non-critical régimes, behavior under criticality is much harder to determine. However, we can still say a few things.
This section will also mark the end of our exciting journey through Bernoulli percolation: We will finally determine the exact value of the critical probability for dimension $d = 2$. So let's begin!

Lemma 6.1. *For any dimension $d \geq 2$, $\chi(p_c) = \infty$, ie:- the expected size of the cluster containing the origin is infinite at the critical probability.*

Proof. Consider any finite set $S \subset \mathbb{Z}^d$ such that $\mathbf{0} \in S$: Then $\varphi_p(S)$ is a (non-trivial) polynomial in p, and thus is a continuous function, from $(0,1)$ to $(0,\infty)$. Thus, by the topological definition of continuity, we have that $\varphi_p^{-1}((0,1))$ must be an open set $I \subset (0,1)$. Also, observe from the proof of Theorem 5.1 that for $p > p_c$, we have $\varphi_p(S) \geq 1$. Thus $I \subset (0, p_c]$. But since I is open, it can not contain p_c without also containing some $p > p_c$, and consequently, I can't contain p_c, implying that $\varphi_{p_c}(S) \geq 1$.
In particular, $\varphi_{p_c}(\Lambda_n) \geq 1$ for every $n \geq 1$. Now, note that

$$\chi(p_c) = \mathbb{E}[|C|] = \sum_{x\in\mathbb{Z}^d} \mathbb{P}_{p_c}(\mathbf{0} \longleftrightarrow x) \geq \frac{1}{dp_c} \sum_{n\geq 1} \varphi_{p_c}(\Lambda_n) = \infty$$

□

We now arrive at our moment of truth: We can finally calculate $p_c(2)$. But before we do that, a word about higher dimensions: Calculating the critical probability for dimensions greater than 2 is an extremely difficult task. For large enough dimensions, however, we do know that $p_c(d) \sim \frac{1}{d}$, which is an approximation we obtain from **Bethe lattices**. However, for intermediate dimensions of $3, 4, 5$, it is much harder to make an accurate comment.
While calculating the value of $p_c(2)$, ie:- the critical probability for 2 dimensions, we shall also show that $\theta(p_c(2), 2) = 0$: ie:- even though the expected size of the cluster containing the origin is infinite at p_c, that cluster is almost surely not infinite! Since θ is 0 at the critical probability, we can then also show that $\theta(\cdot, 2)$ **is a continuous function on** $[0,1]$: Indeed, note that the continuity of θ on $[0, p_c)$ is obvious. Since θ was right continuous everywhere, combined with the fact that $\theta(p_c(2), 2) = 0$, we get that $\theta(\cdot, 2)$ is continuous at $p_c(2)$ too. Showing the continuity of $\theta(\cdot, 2)$ on $(p_c(2), 1]$ requires the fact that the infinite cluster in our lattice is almost surely unique. However, the argument is nuanced and we shall not present it here. Since the value of $\theta(p_c(d), d)$ is not known [13], **for dimensions $d \geq 3$, the strongest statement we can make, as of today, is that θ is continuous on $[0,1] \setminus \{p_c\}$, and θ will be continuous at p_c if and only if it vanishes on p_c.**
So without any further ado, let's get to it.

[12] the independence argument is the same as in Lemma 5.3
[13] but suspected to be 0

Lemma 6.2. *Let us fix our dimension to $d = 2$.*
Let $\mathcal{H}_n$ be the event that there exists an open path going from the left vertical edge to the right vertical edge of the rectangle $R_n := [0, n] \times [0, n-1]$. Also, fix a $p < p_c$. Then

$$\lim_{n\to\infty} \mathbb{P}_p(\mathcal{H}_n) = 0$$

Proof. By Theorem 5.1, there exists a constant $c > 0$ such that

$$\mathbb{P}_p(\mathbf{0} \longleftrightarrow \partial\Lambda_n) \le e^{-cn}$$

But then

$$\begin{aligned}\mathbb{P}_p(\mathcal{H}_n) &\le \sum_{i=0}^{n-1} \mathbb{P}_p((0,i) \text{ is connected to the right vertical edge of } R_n) \\ &\le n\mathbb{P}_p(\mathbf{0} \longleftrightarrow \partial\Lambda_n) \le ne^{-cn} \to 0\end{aligned}$$

as $n \to \infty$. □

Lemma 6.3. *Let $\mathcal{H}_n$ be as defined above. Then $\mathbb{P}_{1/2}(\mathcal{H}_n) = 1/2$ for all $n \in \mathbb{N}$.*

Proof. Construct the dual lattice, as in the Peierls argument. Then there is a left-right path in R_n if and only if the event $\mathcal{B}_n$ happens, where $\mathcal{B}_n$ is the event that there is no bottom to top path in R_n^*. Consequently, for any $p \in [0, 1]$, we have $\mathbb{P}_p(\mathcal{H}_n) = 1 - \mathbb{P}_p(\mathcal{B}_n)$. But for $p = 1/2$, by symmetry, we have that $\mathbb{P}_{1/2}(\mathcal{H}_n) = \mathbb{P}_{1/2}(\mathcal{B}_n)$, and consequently $\mathbb{P}_{1/2}(\mathcal{H}_n) = 1/2$ for all $n \in \mathbb{N}$. □

Lemma 6.4. *Fix a p such that $\theta(p) > 0$* [14]*. Let $\mathcal{H}_n$ be as defined above. Then*

$$\lim_{n\to\infty} \mathbb{P}_p(\mathcal{H}_n) = 1$$

Proof. Fix integers $n \gg k \gg 1$. Since a path from Λ_k to ∞ intersects $\partial\Lambda_n$ in one of its four sides (left, right, bottom, top), we have, by Corollary 3.2.2,

$$\mathbb{P}_p(\Lambda_k \text{ is connected to the left boundary of } \Lambda_n) \ge 1 - \mathbb{P}_p(\Lambda_k \not\longleftrightarrow \infty)^{1/4}$$

Set $n' := \lfloor (n-1)/2 \rfloor$, and let $\mathcal{A}_n$ be the event that there exists a path from $\Lambda_k + (n', n')$ to the left side of $\partial\Lambda_n$, and there also exists a path from $\Lambda_k + (n'+2, n')$ to the right side of $\partial\Lambda_n$. By Corollary 3.2.2, we then have

$$\mathbb{P}_p(\mathcal{A}_n) \ge 1 - 2\mathbb{P}_p(\Lambda_k \not\longleftrightarrow \infty)^{1/4}$$

Now, note that the event $\mathcal{A}_n \setminus \mathcal{H}_n$ is contained in the event $\mathcal{A}$, that there are two different clusters, each intersecting Λ_k, but not each other, and $\mathcal{A}$ further is included in the event that there are two different clusters, which has zero probability. Consequently,

$$\liminf_{n\to\infty} \mathbb{P}_p(\mathcal{H}_n) = \liminf_{n\to\infty} \mathbb{P}_p(\mathcal{A}_n) \ge 1 - 2\mathbb{P}_p(\Lambda_k \not\longleftrightarrow \infty)^{1/4}$$

Since there exists an infinite cluster, the right-hand side of the above inequality $\to 1$ when $k \to \infty$, and consequently, $\lim_{n\to\infty} \mathbb{P}_p(\mathcal{H}_n) = 1$. □

Corollary 6.4.1. *$p_c(2) = 1/2$, ie:- the critical probability for 2 dimensions is $1/2$.*

Corollary 6.4.2. *$\theta(p_c(2), 2) = 0$, ie:- there is no percolation at criticality, in two dimensions.*

[14] note how we didn't instead require $p > p_c$, as one would expect from the previous lemmata. This will help us show later that $\theta(p_c, 2) = 0$

7 Acknowledgements

I thank Amitava Bhattacharya for pushing me to write this report, as well as pointing out to me the right resources, in the correct order, for this topic. This report would not have fructified but for his infinite patience and sage guidance.

References

[DC] Hugo Duminil-Copin. Introduction to Bernoulli percolation. URL: https://www.ihes.fr/~duminil/publi/2017percolation.pdf.

[Gri99] Geoffrey Grimmett. *Percolation*. Springer, 1999.

[Ste] Jeffrey Steif. A mini course on percolation theory. URL: http://www.math.chalmers.se/~steif/perc.pdf.