BOND PERCOLATION ON $\mathbb{Z}^d$

ZIXIN YU

Abstract. This paper introduces the fundamental concepts of percolation theory, focusing on the bond percolation model on the d-dimensional cubic lattice $\mathbb{Z}^d$. We provide a rigorous mathematical framework for the model, define the lattice structure, the associated probability space, and key observables. The central focus is the the existence of a sharp phase transition in the macroscopic connectivity of the graph as the edge probability p varies. We will formally define the critical probability $p_c(d)$ and establish some of its elementary properties.

1. Introduction

Percolation theory originates from a simple physical question: consider a porous stone submerged in water. The pores within the stone form a network of fine channels, each of which may be randomly open or blocked. Can water permeate the stone and reach its center? Intuitively, if the probability of a single channel being open is high, it is more likely that water can find a path to the center. Percolation theory seeks to understand if there is a critical value for this probability, a threshold at which the global connectivity structure of the stone undergoes a fundamental change.

To study this problem, we simplify it into a model at the intersection of graph theory and probability theory. The simplest and most-studied case is bond percolation on the cubic lattice. We consider the integer grid $\mathbb{Z}^d$, where each vertex is connected to its nearest neighbors. Each connection, or edge, is declared "open" with a fixed probability p and "closed" with probability $1-p$, independently of all other edges.

While the "porous stone" provides the initial motivation, the applications of percolation theory are far-reaching. The model can be adapted to describe phenomena such as the spread of infectious diseases, the propagation of forest fires, and the conductivity of disordered materials. In a forest fire model, for instance, trees can be in one of three states: unburnt, burning, or burnt-out. A burning tree can ignite its unburnt neighbors. If we view the potential transmission of fire as an "open" edge, the spread of the fire becomes a percolation process. Such models are often more complex, as the states of edges may not be independent, but they share the same fundamental structure.

This paper will formalize the basic bond percolation model and explore its most striking feature: the existence of a critical probability.

2. The Mathematical Model

To analyze the percolation model with mathematical rigor, we must first precisely define the graph structure and the probability space.

2.1. **The Lattice $\mathbb{L}^d$.** We work in the d-dimensional Euclidean space $\mathbb{R}^d$.

Definition 2.1. The set of vertices of our graph is the integer lattice $\mathbb{Z}^d = \{x = (x_1, \dots, x_d) \mid x_i \in \mathbb{Z} \text{ for all } i = 1, \dots, d\}$.

To define the edges, we use the l_1-distance.

Definition 2.2. The l_1-distance between two vertices $x, y \in \mathbb{Z}^d$ is given by

$$\delta(x, y) = \sum_{i=1}^{d} |x_i - y_i|$$

Two vertices x, y are said to be adjacent if $\delta(x, y) = 1$.

Definition 2.3. The set of edges $\mathbb{E}^d$ consists of all unordered pairs $\langle x, y \rangle$ of adjacent vertices in $\mathbb{Z}^d$. The d-dimensional cubic lattice is the graph $\mathbb{L}^d = \langle \mathbb{Z}^d, \mathbb{E}^d \rangle$.

2.2. **The Probability Space.** The randomness in our model comes from the state of the edges. We formalize this by constructing a suitable probability space.

Definition 2.4. A bond configuration, denoted by ω, is an element of the space $\Omega = \{0, 1\}^{\mathbb{E}^d}$. For each edge $e \in \mathbb{E}^d$, its state in the configuration ω is given by $\omega(e) \in \{0, 1\}$. We say the edge e is open in ω if $\omega(e) = 1$ and closed if $\omega(e) = 0$.

Each $\omega \in \Omega$ represents a specific realization of the open and closed edges on the lattice $\mathbb{L}^d$. The sample space Ω is endowed with a σ-algebra $\mathcal{F}$ generated by the cylinder sets, making it a standard measurable space. We now define the probability measure.

Definition 2.5. For a given probability $p \in [0, 1]$, we define the product Bernoulli measure P_p on $(\Omega, \mathcal{F})$ by

$$P_p = \prod_{e \in \mathbb{E}^d} \mu_p^{(e)}$$

where for each edge e, $\mu_p^{(e)}$ is the Bernoulli measure on $\{0, 1\}$ satisfying $\mu_p^{(e)}(\{1\}) = p$ and $\mu_p^{(e)}(\{0\}) = 1 - p$.

In less formal terms, under the measure P_p, each edge $e \in \mathbb{E}^d$ is independently set to be open with probability p and closed with probability $1 - p$. The triple $(\Omega, \mathcal{F}, P_p)$ is the probability space for our percolation model. We use the subscript p, as in $P_p(A)$ or $E_p[X]$, to denote probabilities and expectations with respect to this measure.

2.3. **Open Clusters and the Percolation Probability.** With the model established, we can define the primary objects of interest.

Definition 2.6. An open path in a configuration ω is a sequence of vertices $x_0, x_1, \dots, x_n$ such that for all $i = 0, \dots, n - 1$, the edge $\langle x_i, x_{i+1} \rangle$ is open.

Definition 2.7. For a vertex $x \in \mathbb{Z}^d$, the open cluster containing x, denoted C_x, is the set of all vertices $y \in \mathbb{Z}^d$ that are connected to x by an open path. If the edge from x to all its neighbors is closed, then $C_x = \{x\}$. The size of the cluster is its cardinality, denoted $|C_x|$.

We are particularly interested in whether these clusters can be infinite. This leads to the central function of the theory.

Definition 2.8. The percolation probability, denoted $\theta(p)$, is the probability that the origin belongs to an infinite open cluster.

$$\theta(p) = P_p(|C_0| = \infty)$$

By the translational invariance of the lattice $\mathbb{L}^d$ and the i.i.d. nature of the edge states, the probability of any vertex x belonging to an infinite cluster is the same, i.e., $P_p(|C_x| = \infty) = \theta(p)$ for all $x \in \mathbb{Z}^d$.

A related quantity is the probability that an infinite cluster exists *anywhere* in the lattice.

Definition 2.9. Let $\psi(p)$ be the probability that there exists at least one infinite open cluster in the lattice.

$$\psi(p) = P_p\left(\bigcup_{x\in\mathbb{Z}^d} \{|C_x| = \infty\}\right)$$

Note that $\theta(p)$ describes a local property (at the origin), whereas $\psi(p)$ describes a global property. Their relationship is fundamental to the theory.

3. The Critical Phenomenon and its Properties

A remarkable feature of the percolation model is the emergence of a phase transition. The macroscopic connectivity of the graph changes abruptly as p crosses a certain threshold.

3.1. **The 0-1 Law for Percolation.**

Our first result shows that the global property $\psi(p)$ is trivial.

Theorem 3.1. *For any $p \in [0, 1]$, the probability of the existence of an infinite open cluster, $\psi(p)$, is either 0 or 1.*

Proof. Let S be the event that an infinite open cluster exists. The occurrence of S does not depend on the state of any finite collection of edges. To see this, consider any finite set of edges $E_{fin} \subset \mathbb{E}^d$. If a configuration ω contains an infinite open cluster, then changing the states of the edges in E_{fin} cannot destroy this property entirely, as an infinite path can always be rerouted to avoid a finite number of closed edges. Thus, S is a tail event with respect to the sequence of independent Bernoulli random variables corresponding to the states of the edges. By Kolmogorov's 0-1 Law, any such tail event must have a probability of either 0 or 1. Therefore, $\psi(p) = P_p(S) \in \{0, 1\}$. □

3.2. **Definition of the Critical Probability.**

The 0-1 law for $\psi(p)$ suggests a sharp transition. The following lemma connects this global property to the local percolation probability $\theta(p)$.

Lemma 3.2. *The relationship between $\theta(p)$ and $\psi(p)$ is as follows:*

1. *If $\theta(p) = 0$, then $\psi(p) = 0$.*
2. *If $\theta(p) > 0$, then $\psi(p) = 1$.*

Proof. (1) The event that an infinite cluster exists is the union of the events that each vertex belongs to an infinite cluster: {an infinite cluster exists} $= \bigcup_{x\in\mathbb{Z}^d}\{|C_x| = \infty\}$. Using the union bound (Boole's inequality) and the translational invariance of $\theta(p)$:

$$\psi(p) = P_p\left(\bigcup_{x\in\mathbb{Z}^d}\{|C_x| = \infty\}\right) \leq \sum_{x\in\mathbb{Z}^d} P_p(|C_x| = \infty) = \sum_{x\in\mathbb{Z}^d}\theta(p)$$

If $\theta(p) = 0$, the sum is 0, which implies $\psi(p) = 0$.

(2) If $\theta(p) > 0$, it is immediate that $\psi(p) \geq \theta(p) > 0$. Since we know from Theorem 3.1 that $\psi(p)$ can only be 0 or 1, we must have $\psi(p) = 1$.

□

This lemma shows that the transition point for $\psi(p)$ from 0 to 1 coincides exactly with the point where $\theta(p)$ becomes strictly positive. This motivates the central definition of the field.

Definition 3.3. The critical probability for the lattice $\mathbb{L}^d$, denoted $p_c(d)$, is defined as:

$$p_c(d) = \sup\{p \in [0, 1] \mid \theta(p) = 0\}$$

Equivalently, $p_c(d) = \inf\{p \in [0, 1] \mid \theta(p) > 0\}$, as $\theta(p)$ is a non-decreasing function of p.

The critical probability divides the behavior of the system into distinct phases:

- Subcritical Phase ($p < p_c(d)$): $\theta(p) = 0$ and $\psi(p) = 0$. Almost surely, all open clusters are finite.
- Supercritical Phase ($p > p_c(d)$): $\theta(p) > 0$ and $\psi(p) = 1$. There is a strictly positive probability for the origin to be in an infinite cluster, and an infinite cluster almost surely exists.
- Critical Phase ($p = p_c(d)$): The behavior is more complex and depends on the dimension d.

3.3. **Elementary Properties of the Critical Probability.** We conclude by proving two simple but important properties of $p_c(d)$.

Theorem 3.4. *For the one-dimensional lattice ($d = 1$), the critical probability is $p_c(1) = 1$.*

Proof. In $\mathbb{L}^1$, the vertices are the integers $\mathbb{Z}$. For the origin to belong to an infinite cluster, there must be an infinite open path extending in at least one direction. Let's consider the path to the right.

If $p < 1$, the probability that a given edge is closed is $1 - p > 0$. Let A_n be the event that the edge $\langle n, n + 1\rangle$ is closed for $n > 0$. The events $\{A_n\}_{n=1}^{\infty}$ are independent, and $P_p(A_n) = 1 - p$. The sum of these probabilities diverges:

$$\sum_{n=1}^{\infty} P_p(A_n) = \sum_{n=1}^{\infty}(1 - p) = \infty$$

By the second Borel-Cantelli Lemma, the event $\{A_n \text{ i.o.}\}$ occurs with probability 1. This means that almost surely, infinitely many edges to the right of the origin are closed. A symmetric argument shows that infinitely many edges to the left are also closed. Therefore, any open path starting from the origin must be finite. Thus, $|C_0| < \infty$ almost surely, which implies $\theta(p) = 0$ for all $p < 1$.

If $p = 1$, all edges are open with probability 1. The entire lattice forms a single infinite cluster, so $|C_0| = \infty$ and $\theta(1) = 1$.

From the definition $p_c(1) = \sup\{p \mid \theta(p) = 0\}$, we conclude that $p_c(1) = 1$. □

Theorem 3.5. *The sequence of critical probabilities is non-increasing in dimension, i.e., $p_c(d+1) \leq p_c(d)$ for all $d \geq 1$.*

Proof. We use a coupling argument. Let $\theta_d(p)$ denote the percolation probability on the lattice $\mathbb{L}^d$. Consider $\mathbb{L}^d$ as a subgraph of $\mathbb{L}^{d+1}$ by identifying $\mathbb{Z}^d$ with the set of vertices $\{x \in \mathbb{Z}^{d+1} \mid x_{d+1} = 0\}$.

Let there be a bond percolation process on $\mathbb{L}^{d+1}$ with edge probability p. The restriction of this process to the embedded subgraph $\mathbb{L}^d$ is a bond percolation process on $\mathbb{L}^d$ with the same parameter p.

If a configuration on this subgraph contains an infinite open cluster connected to the origin, this same cluster is also an infinite open cluster in the larger space $\mathbb{L}^{d+1}$. Therefore, the event $\{|C_0| = \infty \text{ on } \mathbb{L}^d\}$ implies the event $\{|C_0| = \infty \text{ on } \mathbb{L}^{d+1}\}$. This gives the inequality:

$$\theta_d(p) \leq \theta_{d+1}(p) \quad \text{for all } p \in [0, 1]$$

Recall the definition $p_c(d) = \sup\{p \mid \theta_d(p) = 0\}$. If $p > p_c(d)$, then by definition $\theta_d(p) > 0$. From the inequality above, it follows that $\theta_{d+1}(p) > 0$. This means that the set $\{p \mid \theta_{d+1}(p) > 0\}$ contains the set $\{p \mid \theta_d(p) > 0\}$. Taking the infimum over these sets gives:

$$p_c(d+1) = \inf\{p \mid \theta_{d+1}(p) > 0\} \leq \inf\{p \mid \theta_d(p) > 0\} = p_c(d)$$

The theorem is proved. □

The results here establish the existence of a phase transition. A crucial next step, which guarantees the existence of distinct subcritical and supercritical phases for higher dimensions, is to prove that for $d \geq 2$, the critical probability is non-degenerate, i.e., $0 < p_c(d) < 1$.

4. Non-Degeneracy of the Critical Probability for $d \geq 2$

In this section, we prove that for any dimension $d \geq 2$, the critical probability is non-degenerate. This result is fundamental, as it guarantees the existence of distinct subcritical and supercritical phases for percolation on higher-dimensional lattices.

Theorem 4.1. *For bond percolation on the cubic lattice $\mathbb{L}^d$ with $d \geq 2$, we have*

$$0 < p_c(d) < 1$$

Proof. The proof is divided into two main parts.

1. First, we prove that $p_c(d) > 0$ for all $d \geq 1$. This establishes a non-trivial subcritical phase.
2. Second, we prove that $p_c(2) < 1$ by using a duality argument specific to the planar case.

Combining the second result with the monotonicity property $p_c(d) \leq p_c(2)$ for $d > 2$ (established in Theorem 3.4), we conclude that $p_c(d) < 1$ for all $d \geq 2$. Together, these two parts prove the theorem.

4.1. **Proof of the Lower Bound:** $p_c(d) > 0$**.** To show that $p_c(d) > 0$, we will find a value $p > 0$ for which we can demonstrate that $\theta(p) = 0$. The argument relies on counting the number of possible open paths.

If the origin belongs to an infinite open cluster, there must exist at least one self-avoiding open path of infinite length starting from the origin. Consequently, for any finite length n, there must be at least one open self-avoiding path of length n starting at the origin.

Let us define $\sigma(n)$ as the total number of self-avoiding paths of length n starting at the origin in $\mathbb{L}^d$. Let $N(n)$ be the random variable representing the number of such paths that are open in a given configuration.

The event $\{|C_0| = \infty\}$ is a subset of the event $\{N(n) \geq 1\}$ for any $n \geq 1$. Therefore, we can bound the percolation probability:

$$\theta(p) = P_p(|C_0| = \infty) \leq P_p(N(n) \geq 1)$$

By Markov's inequality, $P_p(N(n) \geq 1) \leq E_p[N(n)]$. The expectation of $N(n)$ can be calculated by linearity. For any specific self-avoiding path of length n, the probability that all its n edges are open is p^n. Summing over all possible paths, we get:

$$E_p[N(n)] = \sum_{\pi:\text{self-avoiding}, |\pi|=n} P_p(\pi \text{ is open}) = \sigma(n)p^n$$

Combining these inequalities yields an upper bound for $\theta(p)$:

$$\theta(p) \leq \sigma(n)p^n$$

Our next task is to find an upper bound for $\sigma(n)$. A path of length n is a sequence of $n+1$ distinct vertices $v_0, v_1, \ldots, v_n$. Starting from the origin $v_0 = 0$, there are $2d$ choices for v_1. For any subsequent vertex v_{i+1}, there are at most $2d - 1$ choices, since the path cannot immediately return to v_i. This gives a simple combinatorial bound:

$$\sigma(n) \leq 2d(2d-1)^{n-1}$$

Substituting this into our bound for $\theta(p)$:

$$\theta(p) \leq 2d(2d-1)^{n-1}p^n = \frac{2d}{2d-1}\left[p(2d-1)\right]^n$$

This inequality must hold for all $n \geq 1$. If we choose p such that $p(2d-1) < 1$, or $p < \frac{1}{2d-1}$, then the term $[p(2d-1)]^n$ tends to 0 as $n \to \infty$. Since $\theta(p)$ is a non-negative constant with respect to n, this forces $\theta(p) = 0$.

Thus, for any $p < \frac{1}{2d-1}$, we have $\theta(p) = 0$. From the definition $p_c(d) = \sup\{p \mid \theta(p) = 0\}$, it follows that

$$p_c(d) \geq \frac{1}{2d-1} > 0$$

This completes the first part of the proof.

We can refine this bound by noting that the limit $\lambda(d) = \lim_{n\to\infty}[\sigma(n)]^{1/n}$, known as the connectivity constant of $\mathbb{L}^d$, exists. The argument above can be extended to show that $p_c(d) \geq 1/\lambda(d)$.

4.2. **Proof of the Upper Bound:** $p_c(2) < 1$**.** To prove that $p_c(2) < 1$, we need to find a value $p < 1$ for which $\theta(p) > 0$. We achieve this using a duality argument for the planar lattice $\mathbb{L}^2$.

Consider the dual lattice of $\mathbb{L}^2$, denoted $\mathbb{L}^{2,*}$. Its vertices are located at the centers of the faces of $\mathbb{L}^2$, at coordinates $(i + \frac{1}{2}, j + \frac{1}{2})$ for $i, j \in \mathbb{Z}$. Each edge in $\mathbb{L}^{2,*}$ crosses exactly one edge in $\mathbb{L}^2$.

We define the state of a dual edge to be the opposite of the primal edge it crosses: a dual edge is declared 'closed' if its corresponding primal edge is open (probability p), and 'open' if the primal edge is closed (probability $1 - p$).

A key topological fact in the plane is that the open cluster of the origin, C_0, is finite if and only if it is enclosed by a circuit of 'open' edges in the dual lattice. If no such circuit exists, C_0 must be infinite. This allows us to relate $\theta(p)$ to properties of the dual graph.

Let us bound the probability that C_0 is finite. This occurs if there is at least one 'open' dual circuit surrounding the origin.

$$1 - \theta(p) = P_p(|C_0| < \infty) = P_p(\text{exists an open dual circuit } \gamma^* \text{ surrounding the origin})$$

Using the union bound over all possible self-avoiding dual circuits γ^* that surround the origin:

$$1 - \theta(p) \leq \sum_{\gamma^*} P_p(\gamma^* \text{ is open})$$

The probability that a specific dual circuit γ^* of length n is open is $(1-p)^n$. Let $\rho(n)$ be the number of distinct self-avoiding dual circuits of length n that enclose the origin. The sum becomes:

$$1 - \theta(p) \leq \sum_{n=4}^{\infty} \rho(n)(1-p)^n$$

Now we must estimate $\rho(n)$. A simple bound can be obtained by relating circuits to paths. Any circuit of length n containing the origin can be cut at one of its n vertices to form a self-avoiding path of length $n - 1$. The number of self-avoiding paths of length k in the dual lattice (which is also a square lattice) is $\sigma(k)$. A loose but sufficient bound is $\rho(n) \leq n\sigma(n-1)$.

$$1 - \theta(p) \leq \sum_{n=4}^{\infty} n\sigma(n-1)(1-p)^n$$

Using the connectivity constant $\lambda(2)$ for $\mathbb{L}^2$, we know that $\sigma(k) \approx [\lambda(2)]^k$ for large k. The sum is thus bounded by a geometric-like series whose convergence is determined by the ratio $(1-p)\lambda(2)$. The series converges if $(1-p)\lambda(2) < 1$.

If we choose p such that $1 - p < 1/\lambda(2)$, or equivalently $p > 1 - 1/\lambda(2)$, the sum on the right-hand side converges. As $p \to 1$, the sum tends to 0. We can therefore choose a $p < 1$ (but sufficiently close to 1) such that the sum is strictly less than 1. For such a p:

$$1 - \theta(p) < 1 \implies \theta(p) > 0$$

By the definition $p_c(2) = \inf\{p \mid \theta(p) > 0\}$, we have established that

$$p_c(2) \leq 1 - \frac{1}{\lambda(2)} < 1$$

This completes the second part of the proof. □

4.3. **Conclusion.** We have shown that $p_c(d) > 0$ for $d \geq 2$ and $p_c(2) < 1$. Since $p_c(d) \leq p_c(2)$ for all $d > 2$, it follows that $p_c(d) < 1$ for all $d \geq 2$. This establishes the main theorem, confirming that for two or more dimensions, both the subcritical and supercritical phases exist and are non-trivial. The study of percolation can now proceed to analyze the distinct properties of these phases.

5. Correlation Inequalities and the Influence of p

We now introduce several important results that are frequently used in the study of probabilistic graph models. These tools allow us to understand how the probability of events changes with the edge density p and how different events relate to one another.

5.1. **Monotone Events and their Properties.** The motivation for defining increasing events comes from the desire to understand how the global properties of the lattice change as the edge probability p varies. A natural question is whether the percolation probability, $\theta(p)$, is a monotonic function of p. Intuitively, as p increases, the model's connectivity should improve, making it more likely for the origin to belong to an infinite cluster. We formalize this intuition with the following definitions.

Definition 5.1. Let ω_1 and ω_2 be two bond configurations. We write $\omega_1 \leq \omega_2$ if $\omega_1(e) \leq \omega_2(e)$ for all edges $e \in \mathbb{E}^d$. An event $A \subseteq \Omega$ is called an **increasing event** if for any $\omega \in A$, the condition $\omega \leq \omega'$ implies $\omega' \in A$. Similarly, an event is **decreasing** if $\omega \in A$ and $\omega' \leq \omega$ implies $\omega' \in A$.

In essence, an increasing event is one whose occurrence is facilitated by having more open edges. A decreasing event is one whose occurrence is hindered by having more open edges.

Definition 5.2. A random variable $N : \Omega \to \mathbb{R}$ is called an **increasing random variable** if $N(\omega) \leq N(\omega')$ whenever $\omega \leq \omega'$.

With these definitions, we can rigorously show how the parameter p affects the model. We use the standard coupling construction where the state of each edge e is determined by comparing a uniform random variable $X(e) \sim U(0, 1)$ to p. Let η_p be the random configuration where $\eta_p(e) = 1$ if $X(e) < p$ and 0 otherwise. For a fixed realization of the vector $X = (X(e))_{e \in \mathbb{E}^d}$, if $p_1 < p_2$, it is clear that $\eta_{p_1} \leq \eta_{p_2}$.

Theorem 5.3. *Let A be an increasing event. Then its probability $P_p(A)$ is a non-decreasing function of p.*

Proof. Let I_A be the indicator function for the event A. Since A is an increasing event, I_A is an increasing random variable. For $p_1 < p_2$, using the coupling described above, we have $\eta_{p_1} \leq \eta_{p_2}$, which implies $I_A(\eta_{p_1}) \leq I_A(\eta_{p_2})$. Taking the expectation over all possible outcomes of the underlying random variables $X(e)$ preserves this inequality:

$$E[I_A(\eta_{p_1})] \leq E[I_A(\eta_{p_2})]$$

This is equivalent to $P_{p_1}(A) \leq P_{p_2}(A)$. □

As an immediate corollary, since the event $\{|C_0| = \infty\}$ is an increasing event, $\theta(p)$ is a non-decreasing function of p. Many other natural events, such as the existence of an open path between two vertices x and y, are also increasing.

5.2. **The FKG Inequality.** The FKG inequality, named after Fortuin, Kasteleyn, and Ginibre, formalizes the intuition that if two events are both increasing, they should be positively correlated. That is, the occurrence of one should make the other more likely.

Theorem 5.4 (FKG Inequality)**.** *If A and B are two increasing events, then*

$$P_p(A \cap B) \geq P_p(A)P_p(B)$$

More generally, if X and Y are two increasing random variables with finite second moments, then their covariance is non-negative:

$$E_p[XY] \geq E_p[X]E_p[Y]$$

Proof. We prove the version for random variables; the result for events follows by letting $X = I_A$ and $Y = I_B$. The proof proceeds by induction on the number of edges n on which the random variables depend.

Base Case ($n = 1$): Suppose X and Y depend only on the state of a single edge e_1. Let their values be $X(1), Y(1)$ if the edge is open and $X(0), Y(0)$ if it is closed.

$$\begin{aligned} E_p[XY] - E_p[X]E_p[Y] &= [pX(1)Y(1) + (1-p)X(0)Y(0)] \\ &\quad - [pX(1) + (1-p)X(0)][pY(1) + (1-p)Y(0)] \\ &= p(1-p)[X(1)Y(1) - X(1)Y(0) - X(0)Y(1) + X(0)Y(0)] \\ &= p(1-p)[X(1) - X(0)][Y(1) - Y(0)] \end{aligned}$$

Since X and Y are increasing, $X(1) \geq X(0)$ and $Y(1) \geq Y(0)$, so the expression is non-negative.

Inductive Step: Assume the inequality holds for random variables depending on $k-1$ edges. Now let X, Y depend on edges $e_1, \ldots, e_k$. We condition on the configuration $\omega' = (\omega(e_1), \ldots, \omega(e_{k-1}))$. By the law of total expectation, $E_p[XY] = E_p[E_p[XY \mid \omega']]$. Given ω', X and Y are random variables that depend only on the state of edge e_k. As functions of $\omega(e_k)$, they remain increasing. Thus, by the base case:

$$E_p[XY \mid \omega'] \geq E_p[X \mid \omega']E_p[Y \mid \omega']$$

Taking the expectation of both sides over the configurations of the first $k-1$ edges:

$$E_p[XY] \geq E_p\left[E_p[X \mid \omega']E_p[Y \mid \omega']\right]$$

Now, let $X'(\omega') = E_p[X \mid \omega']$ and $Y'(\omega') = E_p[Y \mid \omega']$. These are random variables that depend on the first $k-1$ edges. One can show they are also increasing. We can therefore apply the inductive hypothesis to X' and Y':

$$E_p[X'Y'] \geq E_p[X']E_p[Y']$$

Using the law of total expectation again, $E_p[X'] = E_p[X]$ and $E_p[Y'] = E_p[Y]$. Combining these steps gives the desired result:

$$E_p[XY] \geq E_p[X'Y'] \geq E_p[X']E_p[Y'] = E_p[X]E_p[Y]$$

The extension to variables depending on infinitely many edges can be done via a limiting argument using the martingale convergence theorem. The inequality can also be extended to decreasing events (the inequality direction is the same) or a mix of increasing and decreasing events (the inequality is reversed). ☐

5.3. **The BK Inequality.** In contrast to the FKG inequality, the BK inequality (due to van den Berg and Kesten) provides an upper bound on the probability of two increasing events occurring "disjointly".

Definition 5.5. For two events A and B, the event $A \circ B$ (read as "A and B occur disjointly") is the event that there exist two disjoint sets of open edges, E_1 and E_2, such that the configuration with open edges E_1 is in A and the configuration with open edges E_2 is in B.

For example, if A is the event "there is an open path from x to y" and B is the event "there is an open path from u to v", then $A \circ B$ is the event "there exist edge-disjoint open paths from x to y and from u to v". Intuitively, the requirement that the paths be disjoint makes the event harder to satisfy.

Theorem 5.6 (BK Inequality)**.** *If A and B are increasing events that depend on only a finite set of edges, then*

$$P_p(A \circ B) \leq P_p(A)P_p(B)$$

The proof is more involved and is omitted here, but its core idea is to "decouple" the events. One imagines replacing each edge in the graph with two parallel edges, one designated for event A and the other for event B. In this new graph, A and B become independent. The probability $P_p(A \circ B)$ in the original graph can be shown to be no more than the probability $P_p(A \cap B)$ in the new, decoupled graph, which equals $P_p(A)P_p(B)$.

The BK and FKG inequalities can be combined. For example, if A_i is the event that a path set π_i contains an open path, then $A_1 \circ \cdots \circ A_k$ is the event that there exist edge-disjoint open paths, one from each set π_i. We have:

$$P_p(A_1 \circ \cdots \circ A_k) \leq \prod_{i=1}^{k} P_p(A_i) \leq P_p(A_1 \cap \cdots \cap A_k)$$

5.4. **Russo's Formula.** Russo's formula provides an explicit expression for the derivative of the probability of an increasing event with respect to p.

Definition 5.7. An edge e is said to be **pivotal** for an event A in a configuration ω if changing the state of e (and nothing else) changes whether A occurs.

Theorem 5.8 (Russo's Formula)**.** *Let A be an increasing event that depends on only a finite set of edges. Then $P_p(A)$ is differentiable, and*

$$\frac{d}{dp}P_p(A) = \sum_e P_p(e \textit{ is pivotal for } A) = E_p[N(A)]$$

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

Proof Sketch. Using the definition of the derivative and the coupling construction, we find that

$$P_{p+\delta}(A) - P_p(A) = P(\eta_p \notin A, \eta_{p+\delta} \in A)$$

For small δ, the event on the right happens almost exclusively when exactly one edge e has its state flipped from closed to open because $p \leq X(e) < p + \delta$. The probability of two or more edges flipping is $O(\delta^2)$ and does not contribute to the derivative. The event occurs if this single flipped edge is pivotal.

$$P_{p+\delta}(A) - P_p(A) \approx \sum_e P(p \leq X(e) < p + \delta \text{ and } e \text{ is pivotal})$$

Since the pivotal property of e is independent of its own state, this is approximately

$$\sum_e P(p \leq X(e) < p + \delta) \cdot P_p(e \text{ is pivotal}) = \delta \sum_e P_p(e \text{ is pivotal})$$

Dividing by δ and taking the limit gives the result. □

A useful alternative form of Russo's formula is derived as follows:

$$\frac{dP_p(A)}{dp} = \sum_e \frac{1}{p} P_p(e \text{ is open and pivotal}) = \frac{1}{p} \sum_e P_p(A \cap \{e \text{ is pivotal}\})$$

The second equality holds because if an edge is pivotal for an increasing event, the event must occur if the edge is open.

$$\frac{dP_p(A)}{dp} = \frac{1}{p} \sum_e P_p(A) P_p(e \text{ is pivotal} \mid A) = \frac{P_p(A)}{p} E_p[N(A) \mid A]$$

Separating variables, $\frac{dP_p(A)}{P_p(A)} = \frac{1}{p} E_p[N(A) \mid A] dp$, and integrating from p_1 to p_2 yields:

$$P_{p_2}(A) = P_{p_1}(A) \exp\left(\int_{p_1}^{p_2} \frac{E_p[N(A) \mid A]}{p} dp \right)$$

If A depends on m edges, then $N(A) \leq m$, which gives the bound:

$$\frac{P_{p_2}(A)}{P_{p_1}(A)} \leq \left(\frac{p_2}{p_1} \right)^m$$

This shows that the probability of an event depending on a small number of edges cannot change too rapidly with p.

5.5. **Inequalities from Reliability Theory.** Finally, we present an identity that relates the derivative of $P_p(A)$ to the covariance between the event and the number of open edges. Assume A depends on a finite set of m edges. Let N be the number of open edges in this set.

Theorem 5.9. *For any event A depending on a finite set of m edges,*

$$\frac{d}{dp} P_p(A) = \frac{1}{p(1-p)} cov_p(N, I_A)$$

Proof. $P_p(A) = \sum_\omega p^{N(\omega)}(1-p)^{m-N(\omega)} I_A(\omega)$. Differentiating with respect to p gives

$$\begin{aligned}
\frac{dP_p(A)}{dp} &= \sum_\omega I_A(\omega)\left[Np^{N-1}(1-p)^{m-N} - (m-N)p^N(1-p)^{m-N-1}\right] \\
&= \sum_\omega I_A(\omega)p^N(1-p)^{m-N}\left[\frac{N}{p} - \frac{m-N}{1-p}\right] \\
&= \frac{1}{p(1-p)}\sum_\omega I_A(\omega)p^N(1-p)^{m-N}(N-mp) \\
&= \frac{1}{p(1-p)}\left(E_p[NI_A] - mpE_p[I_A]\right) = \frac{\operatorname{cov}_p(N, I_A)}{p(1-p)}
\end{aligned}$$

since $E_p[N] = mp$. □

If A is an increasing event, we can use the FKG inequality to get a lower bound. Note that $N - I_A$ is an increasing random variable. By FKG, $\operatorname{cov}_p(N - I_A, I_A) \geq 0$.

$$\operatorname{cov}_p(N, I_A) = \operatorname{cov}_p(N - I_A, I_A) + \operatorname{Var}_p(I_A) \geq \operatorname{Var}_p(I_A)$$

This leads to the inequality:

$$\frac{dP_p(A)}{dp} \geq \frac{\operatorname{Var}_p(I_A)}{p(1-p)} = \frac{P_p(A)(1-P_p(A))}{p(1-p)}$$

This provides a useful lower bound on the rate of change of the probability of an increasing event. A tighter bound, stated without proof, is

$$\frac{dP_p(A)}{dp} \geq \frac{P_p(A)\log P_p(A)}{p\log p}$$

6. Uniqueness of the Critical Probability

In this section, we introduce the theorem on the uniqueness of the critical probability. We will first state the theorem and explore some of its powerful consequences. The full proof, which is lengthy, will be methodically broken down into a series of lemmas in subsequent discussions.

6.1. New Observables and an Alternative Critical Probability.

6.1. New Observables and an Alternative Critical Probability. Before proceeding, we introduce several key quantities that will be used throughout our analysis.

Definition 6.1. Let C_0 be the open cluster containing the origin. We define the following:

- The mean cluster size: $\chi(p) = E_p[|C_0|]$.
- The mean finite cluster size: $\chi^f(p) = E_p[|C_0| \cdot I(|C_0| < \infty)]$.
- The expected reciprocal cluster size: $\kappa(p) = E_p[1/|C_0|]$.

These four quantities, $(\theta, \chi, \chi^f, \kappa)$, are central to the study of percolation. Some of their properties are immediate from the definitions.

- In the subcritical phase ($p < p_c$), infinite clusters almost surely do not exist, so $|C_0| < \infty$ a.s. This implies $\chi(p) = \chi^f(p)$.

- In the supercritical phase ($p > p_c$), we have $\theta(p) = P_p(|C_0| = \infty) > 0$. The expectation for the cluster size becomes:

$$\chi(p) = \infty \cdot P_p(|C_0| = \infty) + \sum_{n=1}^{\infty} nP_p(|C_0| = n) = \infty$$

A direct conclusion is that the mean cluster size $\chi(p)$ is infinite in the supercritical phase. But does the converse hold? Is $\chi(p)$ necessarily finite in the subcritical phase? This depends on the convergence of the series $\sum nP_p(|C_0| = n)$ for $p < p_c$. Answering this is a core goal of this section.

This question motivates an alternative definition for the critical probability, based on the mean cluster size.

Definition 6.2. We define a new critical probability, p_T, as

$$p_T = \sup\{p \mid \chi(p) < \infty\}$$

This definition is physically meaningful, as p_T marks the threshold where the average cluster size diverges. A natural and fundamental question arises: do these two definitions of the critical point coincide? That is, does $p_c = p_T$? This is the problem of the uniqueness of the critical probability. Proving that they are equal is equivalent to proving that the mean cluster size $\chi(p)$ is finite if and only if $p < p_c$.

6.2. **An Overview of the Proof Strategies.** We will establish that $p_c = p_T$ by proving that $\chi(p) < \infty$ for all $p < p_c$. There are two primary approaches to this proof.

Strategy 1: Exponential Decay of the Cluster Radius. The first approach is to show that in the subcritical phase, the size of open clusters is limited by an exponential decay law. The key theorem to be proven is:

Theorem 6.3. *For any $p < p_c(d)$, there exists a function $h(p) > 0$ such that for all $n \geq 1$,*

$$P_p(\textit{the origin is connected to the boundary of a box of radius } n) \leq e^{-nh(p)}$$

We will show that this exponential decay is sufficient to ensure the convergence of the sum for $\chi(p)$, thus proving it is finite.

Strategy 2: The Ghost-Site Method. The second approach involves augmenting the percolation model to derive new relationships between the fundamental quantities. This method will be used to prove the following theorem:

Theorem 6.4. *If p is such that $\chi^f(p) = \infty$, then one of two conditions must hold:*

(1) $\theta(p) > 0$*, or*
(2) $\theta(p) = 0$ *and for all $p' > p$, $\theta(p') \geq \frac{p'-p}{2p'}$.*

While the first approach is perhaps more direct, the second approach is rich with consequences and reveals deep connections between θ, p, and χ. Both are essential to a full understanding of the theory.

6.3. **Consequences of the Uniqueness Theorem.** Assuming the theorems stated in the two strategies are true, we can now demonstrate that $p_c = p_T$.

Proof of $p_c = p_T$ using Strategy 2: We prove by contradiction. Assume there exists a $p < p_c$ for which $\chi(p) = \infty$. Since we are in the subcritical phase, $|C_0| < \infty$ a.s., so $\chi(p) = \chi^f(p) = \infty$. By the theorem of Strategy 2, since $\theta(p) = 0$ (as $p < p_c$), the second condition must apply: for any $p' > p$, we must have $\theta(p') \geq \frac{p'-p}{2p'}$. However, because $p < p_c$, we can choose a p' such that $p < p' < p_c$. For such a p', we know by the definition of p_c that $\theta(p') = 0$. But the inequality requires $\theta(p') > 0$, a clear contradiction. Therefore, our initial assumption must be false. It must be that for all $p < p_c$, $\chi(p) < \infty$. This implies $p_c \leq p_T$. Since we already established that $\chi(p) = \infty$ for $p > p_c$, we have $p_T \leq p_c$. We conclude that $p_c = p_T$.

Proof of $p_c = p_T$ using Strategy 1: To formalize the argument, we introduce some notation. Let $S(n) = \{x \in \mathbb{Z}^d \mid \delta(0,x) \leq n\}$ be the ball of radius n around the origin, and let $\partial S_n = \{x \in \mathbb{Z}^d \mid \delta(0,x) = n\}$ be its boundary. Let A_n be the event that there exists an open path from the origin to a point on ∂S_n. The theorem of Strategy 1 states that for $p < p_c$, $P_p(A_n) \leq e^{-nh(p)}$.

Let $M = \max\{n \mid A_n \text{ occurs}\}$ be the radius of the origin's cluster. For $p < p_c$, $P_p(M < \infty) = 1$. If $M = n$, the cluster C_0 is contained within the ball $S(n)$. The number of vertices in this ball is bounded by a polynomial in n, i.e., $|S(n)| \leq \pi(d)(n+1)^d$ for some constant $\pi(d)$. We can now bound the mean cluster size:

$$\chi(p) = E_p[|C_0|] = \sum_{n=0}^{\infty} E_p[|C_0| \mid M = n] P_p(M = n)$$

Since $P_p(M = n) \leq P_p(M \geq n) = P_p(A_n)$ and $E_p[|C_0| \mid M = n] \leq |S(n)|$, we have:

$$\chi(p) \leq \sum_{n=0}^{\infty} |S(n)| P_p(A_n) \leq \sum_{n=0}^{\infty} \pi(d)(n+1)^d e^{-nh(p)}$$

This series converges because the exponential term $e^{-nh(p)}$ decays faster than any polynomial term $(n+1)^d$ grows. Thus, $\chi(p) < \infty$ for all $p < p_c$, which again proves $p_c = p_T$.

A powerful corollary of the exponential decay in Strategy 1 is a bound on the tail probability of the cluster size distribution. The core idea is to relate the size of a cluster (a number of vertices) to its spatial extent (a radius).

Theorem 6.5. *For $p < p_c$, there exists a constant $\beta(p) > 0$ such that for all $n \geq 1$:*

$$P_p(|C_0| \geq n) \leq e^{-\beta(p) n^{1/d}}$$

Proof Sketch. Let us assume the result that for $p < p_c$, there exists a function $h(p) > 0$ such that for any integer $k \geq 1$, the probability of the origin connecting to the boundary of a ball of radius k is bounded by:

$$P_p(A_k) = P_p(C_0 \text{ connects to } \partial S_k) \leq e^{-kh(p)}$$

where A_k is the event that there is an open path from the origin to a vertex in $\partial S_k = \{x \in \mathbb{Z}^d \mid \delta(0,x) = k\}$. Now, consider the event that the origin's cluster C_0 has size at least n, i.e., $\{|C_0| \geq n\}$.

A cluster is a set of vertices. If this set is to contain at least n vertices, it must occupy a certain volume in the lattice. Let us consider a ball of radius m centered at the origin, $S(m) = \{x \in \mathbb{Z}^d \mid \delta(0, x) \leq m\}$. The number of vertices in this ball, $|S(m)|$, is bounded by a polynomial in m. Specifically, there exists a constant $\pi(d)$ (depending on the dimension d) such that $|S(m)| \leq \pi(d)(m + 1)^d$.

If the entire cluster C_0 were contained within the ball $S(m)$, its size would be at most the volume of the ball: $|C_0| \leq |S(m)|$. Suppose the event $\{|C_0| \geq n\}$ occurs. This implies that the cluster cannot be contained within any ball $S(m)$ whose volume is less than n. Let us find an integer radius k such that if the cluster's radius is less than k, its size must be less than n. We need to find k such that if C_0 is contained in $S(k - 1)$, then $|C_0| < n$. The maximum size of a cluster contained in $S(k - 1)$ is its volume, $|S(k - 1)| \leq \pi(d)k^d$. We want to find k such that $\pi(d)k^d < n$. This is true if $k < (n/\pi(d))^{1/d}$.

Let's choose $k = \lfloor (n/\pi(d))^{1/d} \rfloor$. Now, consider the contrapositive argument. If the event A_k does *not* occur, it means the cluster C_0 has a radius less than k, so it must be entirely contained within the ball $S(k - 1)$. In this case, its size is bounded by:

$$|C_0| \leq |S(k-1)| \leq \pi(d)k^d = \pi(d) \left\lfloor \left(\frac{n}{\pi(d)}\right)^{1/d} \right\rfloor^d \leq \pi(d) \left(\left(\frac{n}{\pi(d)}\right)^{1/d}\right)^d = n$$

So, if A_k does not occur, then $|C_0| \leq n$. The contrapositive statement is that if $|C_0| > n$ (or $|C_0| \geq n$, as the argument holds), then the event A_k must have occurred. This gives us the crucial event inclusion:

$$\{|C_0| \geq n\} \subseteq A_k, \quad \text{where } k = \left\lfloor \left(\frac{n}{\pi(d)}\right)^{1/d} \right\rfloor$$

This inclusion implies the probability inequality:

$$P_p(|C_0| \geq n) \leq P_p(A_k)$$

Now we apply the assumed exponential decay for $P_p(A_k)$:

$$P_p(|C_0| \geq n) \leq e^{-kh(p)} = \exp\left(-\left\lfloor \left(\frac{n}{\pi(d)}\right)^{1/d} \right\rfloor h(p)\right)$$

For large n, the floor function $\lfloor x \rfloor$ is very close to x. We can bound it from below, for instance by $\lfloor x \rfloor \geq x - 1$.

$$P_p(|C_0| \geq n) \leq \exp\left(-\left(\left(\frac{n}{\pi(d)}\right)^{1/d} - 1\right) h(p)\right) = e^{h(p)} \exp\left(-\frac{h(p)}{(\pi(d))^{1/d}} n^{1/d}\right)$$

This expression is of the form $C \cdot e^{-\beta' n^{1/d}}$. By choosing a slightly smaller constant $\beta(p) < \beta'$ to absorb the leading factor $e^{h(p)}$ over all $n \geq 1$, we arrive at the stated result: there exists a $\beta(p) > 0$ such that for all $n \geq 1$,

$$P_p(|C_0| \geq n) \leq e^{-\beta(p) n^{1/d}}$$

□

This result shows that large clusters are extremely rare in the subcritical phase. Furthermore, these methods can be extended to show that for p slightly above p_c, $\theta(p)$ grows at least linearly, i.e., $\theta(p) \geq a(p - p_c)$ for some $a > 0$, indicating a sharp transition at the critical point.

Theorem 6.6 (Uniqueness of the Infinite Cluster)**.** *For bond percolation on the lattice $\mathbb{L}^d$ with $d \geq 2$, if the edge probability p is in the supercritical phase, $p > p_c(d)$, then there is almost surely exactly one infinite open cluster.*

Proof. Let $N(\omega)$ be the random variable for the number of distinct infinite open clusters (IOCs) in a configuration ω. We will prove that for $p > p_c(d)$, $P_p(N = 1) = 1$.

Let $\mathcal{F}_n$ be the σ-algebra generated by the states of edges within the box $B(n) = \{x \in \mathbb{Z}^d : ||x||_\infty \leq n\}$. The tail σ-algebra is $\mathcal{T} = \bigcap_{n=1}^\infty \sigma(\bigcup_{k=n}^\infty \mathcal{F}_k)$. The number of IOCs, N, is $\mathcal{T}$-measurable. By Kolmogorov's 0-1 Law, any $\mathcal{T}$-measurable random variable must be almost surely constant. Thus, there exists a $k \in \{0, 1, \dots, \infty\}$ such that $P_p(N = k) = 1$. In the supercritical phase, the probability of at least one IOC existing is $\psi(p) = P_p(N \geq 1) = 1$. This implies $k \geq 1$.

The proof proceeds by demonstrating that $P_p(N \geq 2) = 0$, which leaves $k = 1$ as the only possibility.

First, we rule out a finite plurality, $2 \leq k < \infty$. Let M_B be the number of distinct IOCs intersecting a finite set $B \subset \mathbb{Z}^d$. It is a foundational result of percolation theory that two distinct IOCs cannot enter the same finite region.

Lemma 6.7. *For any finite set $B \subset \mathbb{Z}^d$, $P_p(M_B \geq 2) = 0$.*

Let $\{B_n\}_{n=1}^\infty$ be an increasing sequence of finite sets with $\bigcup B_n = \mathbb{Z}^d$. The sequence of random variables $M_{B_n}(\omega)$ is non-decreasing in n for any ω, and converges pointwise to $N(\omega)$:

$$\lim_{n\to\infty} M_{B_n}(\omega) = N(\omega) \quad \forall \omega \in \Omega$$

The event $\{N \geq 2\}$ is the increasing limit of the events $\{M_{B_n} \geq 2\}$. By the continuity of probability measure for increasing sequences of events:

$$P_p(N \geq 2) = P_p\left(\bigcup_{n=1}^\infty \{M_{B_n} \geq 2\}\right) = \lim_{n\to\infty} P_p(M_{B_n} \geq 2)$$

By Lemma 6.7, $P_p(M_{B_n} \geq 2) = 0$ for all n. Thus, $P_p(N \geq 2) = 0$. This eliminates the possibility that $k \in \{2, 3, \dots\}$.

It remains to show that $P_p(N = \infty) = 0$. We prove this by contradiction. Assume $P_p(N = \infty) = 1$. We show this implies a positive density of "trifurcation points," which in turn leads to a contradiction derived from geometric constraints.

Definition 6.8. A vertex x is a **trifurcation point** if it belongs to an IOC, and there are exactly three open edges incident to it, such that removing these three edges fractures the cluster into exactly three distinct IOCs. Let T_x be the event that x is a tri-point.

The contradiction is built upon the following two lemmas, which we now prove.

Lemma 6.9. *If $P_p(N = \infty) = 1$, then $P_p(T_0) > 0$.*

Proof. Let $M_{B(n)}(0)$ be the number of IOCs existing entirely outside the box $B(n) = \{x : ||x||_\infty \leq n\}$. Since $N = \infty$ a.s., for any integer k, $P_p(N \geq k) = 1$. This implies $\lim_{n\to\infty} P_p(M_{B(n)}(0) \geq 3) = P_p(N \geq 3) = 1$. We can therefore choose a box radius n_0 large enough such that $P_p(M_{B(n_0)}(0) \geq 3) \geq 1/2$.

Let $B_0 = B(n_0)$. Define an event J which depends only on edges inside B_0: let J be the event that there exist three open, edge-disjoint paths from the origin to the boundary ∂B_0. For any $p > 0$, such a configuration has positive probability, $P_p(J) > 0$.

The event T_0 (origin is a tri-point) is implied by the joint occurrence of J and $\{M_{B_0}(0) \geq 3\}$, as the three internal arms can connect to the three distinct external IOCs. Thus, $J \cap \{M_{B_0}(0) \geq 3\} \subseteq T_0$. As J depends on edges inside B_0 and $\{M_{B_0}(0) \geq 3\}$ depends on edges outside B_0, they are independent. This gives:

$$P_p(T_0) \geq P_p(J \cap \{M_{B_0}(0) \geq 3\}) = P_p(J) \cdot P_p(M_{B_0}(0) \geq 3)$$

Since $P_p(J) > 0$ and $P_p(M_{B_0}(0) \geq 3) \geq 1/2$, their product is strictly positive. We conclude that $P_p(T_0) > 0$. □

Lemma 6.10. *Let $\mathcal{P} = \{\pi_1, \ldots, \pi_k\}$ be a collection of 3-partitions of a finite set Y. If all pairs in $\mathcal{P}$ are compatible (a part of one partition is a subset of a part of another), then $k \leq |Y| - 2$.*

Proof. We proceed by induction on $|Y|$. The base case $|Y| = 3$ holds trivially, as $k = 1 \leq 3 - 2 = 1$. Assume the lemma holds for all sets of size up to m. Consider a set Y with $|Y| = m + 1$, and let $\mathcal{P}$ be a collection of k pairwise compatible 3-partitions. Choose an element $y \in Y$ and let $Z = Y \setminus \{y\}$. We split $\mathcal{P}$ into two sub-collections:

1. $\mathcal{P}_A = \{\pi \in \mathcal{P} \mid \{y\}$ is not a standalone part of $\pi\}$. Each $\pi \in \mathcal{P}_A$ induces a compatible 3-partition on Z. By the inductive hypothesis, $|\mathcal{P}_A| \leq |Z| - 2 = m - 2$.
2. $\mathcal{P}_B = \{\pi \in \mathcal{P} \mid \{y\}$ is a standalone part of $\pi\}$. Any two distinct, compatible partitions in this set can be shown to be impossible. Thus, $|\mathcal{P}_B| \leq 1$.

The total size is $k = |\mathcal{P}_A| + |\mathcal{P}_B| \leq (m - 2) + 1 = m - 1 = |Y| - 2$. □

Assuming these lemmas, we derive the contradiction. Let $\tau_C(B(n))$ be the number of tri-points in the box $B(n)$ on a specific IOC, C. Each such tri-point induces a 3-partition on the set of boundary points $C \cap \partial B(n)$. As all tri-points on the same IOC induce compatible partitions, Lemma 6.10 implies:

$$\tau_C(B(n)) \leq |C \cap \partial B(n)| - 2$$

Summing over all IOCs gives the total number of tri-points in $B(n)$:

$$\sum_{x \in B(n)} I_{T_x} = \sum_{\text{IOCs } C} \tau_C(B(n)) \leq \sum_C (|C \cap \partial B(n)| - 2) \leq |\partial B(n)|$$

where I_{T_x} is the indicator for event T_x. The last inequality follows because the sets $C \cap \partial B(n)$ are disjoint for distinct IOCs. Taking the expectation and using translation invariance:

$$E_p\left[\sum_{x \in B(n)} I_{T_x}\right] = |B(n)| \cdot P_p(T_0)$$

This yields the inequality $|B(n)| \cdot P_p(T_0) \leq |\partial B(n)|$, or

$$P_p(T_0) \leq \frac{|\partial B(n)|}{|B(n)|}$$

For the box $B(n)$, the volume is $|B(n)| = (2n+1)^d$. The boundary size is $|\partial B(n)| = (2n+1)^d - (2n-1)^d$. By the mean value theorem, for some $\xi \in (2n-1, 2n+1)$,

$$|\partial B(n)| = d(\xi)^{d-1} \cdot 2 = O(n^{d-1})$$

The inequality becomes:

$$P_p(T_0) \leq \frac{O(n^{d-1})}{O(n^d)} = O\left(\frac{1}{n}\right)$$

As this must hold for all n, we take the limit as $n \to \infty$:

$$P_p(T_0) \leq \lim_{n\to\infty} O\left(\frac{1}{n}\right) = 0$$

This forces $P_p(T_0) = 0$, which directly contradicts the result of Lemma **??**. The assumption $P_p(N = \infty) = 1$ must be false.

Since k is an almost-sure constant, $k \geq 1$, and we have shown $P_p(N \geq 2) = 0$ (which implies $k < 2$) and $P_p(N = \infty) = 0$ (which implies $k \neq \infty$), the only remaining possibility is $k = 1$. Therefore, $P_p(N = 1) = 1$. □

For bond percolation on the lattice $\mathbb{L}^d$ with $d \geq 2$, if the edge probability p is in the supercritical phase, $p > p_c(d)$, then there is almost surely exactly one infinite open cluster.

7. Conclusion and Further Topics

This paper introduced bond percolation on $\mathbb{Z}^d$, defining the lattice $\mathbb{L}^d = \langle \mathbb{Z}^d, \mathbb{E}^d \rangle$ and the product Bernoulli measure P_p on edge configurations $\Omega = \{0,1\}^{\mathbb{E}^d}$. We established the fundamental properties of the critical probability $p_c(d) = \sup\{p \in [0,1] \mid \theta(p) = 0\}$, including $p_c(1) = 1$, $0 < p_c(d) < 1$ for $d \geq 2$, and $p_c(d+1) \leq p_c(d)$. The uniqueness of the infinite open cluster for $p > p_c(d)$ ($d \geq 2$) was proven using techniques involving FKG and BK inequalities. We also demonstrated that $p_c(d) = \sup\{p \mid E_p[|C_0|] < \infty\}$, and established exponential decay for cluster properties in the subcritical phase: $P_p(\text{origin connects to } \partial B(n)) \leq e^{-nh(p)}$ and $P_p(|C_0| \geq n) \leq e^{-\beta(p) n^{1/d}}$ for $h(p), \beta(p) > 0$.

7.1. Further Directions.

Key areas for further study involve the precise characterization of the critical behavior. While $p_c(2) = 1/2$ is known for the square lattice, the determination of $p_c(d)$ for $d \geq 3$ remains an open problem. Near $p_c(d)$, various quantities exhibit power-law scaling defined by critical exponents. For example, the percolation probability $\theta(p)$ behaves as $\theta(p) \sim (p - p_c)^\beta$ for $p \downarrow p_c$, and the mean cluster size $\chi(p) = E_p[|C_0|]$ diverges as $\chi(p) \sim |p - p_c|^{-\gamma}$ for $p \neq p_c$. The **correlation length** $\xi(p)$, representing the typical diameter of finite clusters, diverges as $\xi(p) \sim |p - p_c|^{-\nu}$. These exponents $(\beta, \gamma, \nu, \dots)$ are believed to be universal, depending only on dimension d and not on lattice specifics. Rigorous proofs and exact values for these exponents are known only for $d = 2$ (e.g., $\beta_{2D} = 5/36$, $\gamma_{2D} = 43/18$, $\nu_{2D} = 4/3$) and for $d \geq 6$ (where mean-field exponents apply).

Further research also include percolation on general graphs and other variants, such as

- **Site percolation**, where vertices, rather than edges, are open.
- **Oriented percolation**, where edges permit directed paths.
- **Bootstrap percolation**, where nodes are activated based on the number of active neighbors and exhibits discontinuous phase transitions.
- **Percolation on random graphs** (e.g., Erdős-Rényi graphs or scale-free networks), where the critical phenomena and cluster structures differ significantly from lattice models. This also has connections with network robustness.

References

[1] G. Grimmett, *Probability on Graphs: Random Processes on Graphs and Lattices*. Statistical Laboratory, University of Cambridge, 2012.

[2] G. F. Lawler and V. Limic, *Random Walk: A Modern Introduction*. Cambridge University Press, 2010.

[3] P. Sousi, *Percolation and Random Walks on Graphs*. Lecture notes taken by Dexter Chua, Michaelmas 2017.