Networks out of Control:
Models and Methods for Random Networks

Matthias Grossglauser
Patrick Thiran
School of Computer and Communication Sciences
Ecole Polytechnique Fédérale de Lausanne (EPFL)

Contents

1

Introduction: Random Networks in Engineering

Scientists and engineers have to understand and reason about artifacts whose complexity and scale are often prohibitive. To make this task manageable, the reductionist approach calls for abstractions of reality that focus on the salient features of the problem. Such a *model* strips out all the details of the real artifact (an engineered system and its interactions with the environment) that are not crucial for understanding and reasoning. But beyond its use for solving problems, models impart on the scientist and engineer a "way of thinking", i.e., it shapes the way an engineer will approach a new problem, decompose it into manageable parts, and design a solution. In other words, the models taught to an engineering student constrains the design space that he will seek a solution in. The models are the engineer's toolbox.

The prevailing set of models in a field of engineering changes with technological progress, but is also influenced by academic tradition, intellectual elegance, etc. In the narrow disciplines of computer networking and telecommunications, the two main models and associated theories that students had to master are queueing theory and deterministic graph theory.

Queueing theory is the study of systems where clients try to obtain a service from a server. The key problem is that clients may have to await service if they request service when some other client is being served. In this case, a client is queued to await service. Queueing problems arise at various levels in networks. In packet-switched networks (such as the Internet), the arrival and departure of packets at intermediate nodes (routers) can be modeled as a queueing system to compute average packet delay and to dimension buffers. In the phone network, the calls established and terminated by users can be modeled as (another type of) queueing system to compute call block rates.

Graph theory is the study of structures consisting of nodes (vertices) that are connected by links (edges). A graph is a natural way to study global properties of a network. For example, delivering a packet or establishing a phone call involves finding an efficient path through the network to connect the originator to the destination. Graph theory provides elegant solutions to such shortest-path problems.

The goal of this class is to extend the toolbox of models of network engineers and researchers in the face of technological trends and new application scenarios that change the way we build, control, and use networks.

2

Bond percolation: setting and basic techniques

2.1 Introduction

Take a large block of material, such as a big porous stone, and poor water at one face of the block. Take a large cubic block of such a material, and poor water at one of its faces. How deeply will the water seep in it ? It turns out that there is a critical density of holes, below which the material will only get wet in surface, and above which the water will drain in the entire block, no matter how deep it is. This observation lead Broadbent and Hammersely to formulate it mathematically in 1957 [6], in a paper that gave birth in mathematics to the field of percolation theory. The first chapters of this course expose some of the main findings and techniques of percolation theory, mainly for the bond model. They follow very closely the seminal book by Grimmett [18] on percolation theory.

We will consider different models. The first one, and the easiest one, is the *bond model*, defined for the d-dimensional lattice $\mathbb{L}^d = (\mathbb{Z}^d, \mathbb{E}^d)$, where the set of edges $\mathbb{E}^d$ connects adjacent vertices. These edges represent the passages through which the water can flow. Each edge is randomly open (or maintained) with some probability p, and closed (or deleted) otherwise, independently of the others. The second one is the *site model*, defined for the d-dimensional lattice $\mathbb{L}^d$. Each vertex (or site) is now open with some probability p, and closed (or deleted) otherwise, independently of the others. The third model is the *Boolean model*, which is defined on the plane. Finally, we will move to a new model, closer to that used in wireless ad hoc/sensor networks.

2.2 Lattice bond model

In this section, we formalize the bond model that will be the basis of the following chapters. We consider the d-dimensional lattice $\mathbb{L}^d = (\mathbb{Z}^d, \mathbb{E}^d)$, where the set of edges $\mathbb{E}^d$ connects sites $(x, y) = ((x_1, \ldots, x_d), (y_1, \ldots, y_d))$ located at the vertices of $\mathbb{Z}^d$ for which the Manahattan distance, defined by

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

is no more than one: $\delta(x, y) \leq 1$. The edges of $\mathbb{E}^d$ connect thus adjacent vertices of $\mathbb{Z}^d$.

Let $0 \leq p \leq 1$. We declare an edge of $\mathbb{E}^d$ to be *open* with probability p, and closed otherwise, independently of all other edges. This amounts to work on the probability space $(\Omega, \mathcal{F}, \mathbb{P}_p)$ with the sample space $\Omega = \prod_{e \in \mathbb{E}} \{0, 1\}^e$ (its elements $\omega = (\omega(e) \mid e \in \mathbb{E}^d)$ are called *configurations*, with $\omega(e) = 0$ if the edge e is closed and $\omega(e) = 1$ if the edge e is open); where $\mathcal{F}$ is the associated σ-field of subsets of Ω and where $\mathbb{P}_p$ is the product measure

$$\mathbb{P}_p = \prod_{e \in \mathbb{E}} \mu_e$$

where μ_e is a Bernoulli measure given by

$$\mu_e(\omega(e) = 0) = 1 - p, \qquad \mu_e(\omega(e) = 1) = p.$$

We denote by $\mathbb{E}_p$ the corresponding expectation operator.

There is a partial order on the set Ω of configurations, given by $\omega \leq \omega'$ if and only if $\omega(e) \leq \omega'(e)$ for all edges $e \in \mathbb{E}^d$.

Let us introduce some notations and terminology that will be used throughout the course. A *path* of $\mathbb{L}^d$ from vertex x_0 to vertex x_n is an alternating sequence $x_0, e_1, x_2, \ldots, e_{n-1}, x_n$ of distinct vertices x_i and edges $e_i = \langle x_{i-1}, x_i \rangle$. The length of this path is n. If all edges of the path are open, the path is an *open path*. Conversely, if all edges are closed, the path is *closed*. A circuit is a path whose first and last vertices are identical ($x_0 = x_n$).

We denote by $C(x)$ the part of $\mathbb{L}^d$ containing the set of vertices connected by open paths to vertex x and the open edges of $\mathbb{E}^d$ connecting such vertices. By translation invariance of the lattice and the probability measure $\mathbb{P}_p$, the distribution of $C(x)$ does not depend on the vertex x. We therefore take in general $x = 0$, and denote by C the open cluster at the origin: $C = C(0)$. We denote by $|C(x)|$ the size (number of vertices) of $C(x)$.

If A and B are sets of vertices of $\mathbb{L}^d$, we write $A \leftrightarrow B$ to express the fact that there exists an open path connecting some vertex of A to some vertex of B. For example, $C(x) = \big\{y \in \mathbb{Z}^d \mid x \leftrightarrow y\big\}$. We write ∂A to denote the surface of A, which is the set of vertices of A which are adjacent to some vertex that does not belong to A. A typical subset of vertices is a *box*, defined as

$$B(n) = [-n, n]^d = \left\{ x \in \mathbb{Z}^d \mid \max_{1 \leq i \leq d} \{|x_i|\} \leq n \right\}$$

for some $n \in \mathbb{N}^* = \mathbb{N} \setminus \{0\}$. We write $B(n, x)$ for the box $x + B(n)$ having side-length $2n$ and center at x. We will also work often with "diamond" boxes

$$S(n) = \big\{x \in \mathbb{Z}^d \mid \delta(0, x) \leq n\big\}$$

or more general rectangular boxes. We also write $S(n, x)$ for the diamond box $x + S(n)$ centered in x.

2.3 Percolation probability

The main quantity of interest in percolation theory is the probability that the origin belongs to a cluster with an infinite number of vertices, which we denote by θ and call the *percolation probability*. With C denoting the cluster containing the origin, the percolation probability is thus defined as

$$\theta(p) = \mathbb{P}_p(|C| = \infty). \tag{2.1}$$

By space invariance, $\theta(p)$ is the probability that any node belongs to an infinite cluster.

Define the *critical* (or *percolation*) *threshold* as

$$p_c = \sup\{p \mid \theta(p) = 0\}. \tag{2.2}$$

In the one-dimensional case ($d = 1$), it is immediate to see that $p_c = 1$. Indeed, if $p_c < 1$, then walking along the lattice $\mathbb{L}$ in any direction, we will almost surely meet infinitely often an open edge, which yields that all clusters are almost surely finite. However, when $d \geq 2$, it is no longer the case. The main finding of percolation theory is that $0 < p_c < 1$, which implies that there are two phases: the *subcritical phase*, when $p < p_c$, where every vertex is almost surely in a finite open cluster, and the *supercritical phase*, when when $p > p_c$, where each vertex has a non zero probability of belonging to an infinite cluster. Computing the exact value of p_c is a challenge, and still remains an open problem for dimensions larger than 2. In this course, we will compute it for dimension $d = 2$, where Kesten closed the conjecture after more than two decades of research. We can already frame here the value of p_c within $1/3$ and $2/3$ thanks to the following theorem.

Theorem 2.1 (Non trivial phase transition). *The percolation threshold in $\mathbb{L}^2$ is such that $1/3 \leq p_c \leq 2/3$.*

The proof will make use of a technique that will prove to be quite powerful in $d = 2$ dimensions, but that does not generalize well to higher dimensions, which is to work with the *planar dual* graph. If G is a planar graph, drawn in the plane in such a way that edges intersect only at their common vertices, then the dual graph G_d is obtained by putting a vertex in every face of G, and by joining two such vertices by an edge whenever the corresponding faces of G share a common edge. When $G = \mathbb{L}^2$, its dual $G_d = \mathbb{L}^2_d$ is isomorphic to $\mathbb{L}$. The vertices of the dual lattice $\mathbb{L}^2_d$ are placed at the centers of the squares of $\mathbb{L}^2$, i.e. are the set $\left\{(i + 1/2, j + 1/2) \mid (i, j) \in \mathbb{Z}^2\right\}$, and its edges connect adjacent vertices. To every edge of $\mathbb{L}^2$ corresponds exactly one edge of L^2_d, and vice-versa. We declare an edge of the dual lattice L^2_d to be open (resp., close) if and only if its corresponding edge in the lattice $\mathbb{L}^2$ is open (resp., close), as shown in Figure 2.1. This results in a bond percolation process on the dual lattice with the same open edge probability p.

Proof: (i) We first prove that $p_c \geq 1/3$. Let $\sigma(n)$ be the number of distinct, loop free paths ("self-avoiding walks") of $\mathbb{L}^d$ having length n and beginning at the origin. The exact value of $\sigma(n)$ is very difficult to compute for already moderate values of n, but an upper bound on $\sigma(n)$ is $4 \cdot 3^{n-1}$. Indeed, walking from the origin, we have first 4 possible edges to take, and then, at each step, up to 3 different edges. Let $N(n)$ be the number of such paths that are open. Since each path is open with probability p^n,

$$\mathbb{E}_p[N(n)] = \sum_{s=1}^{\sigma(n)} \mathbb{E}_p\left[1_{\{\text{ path } s \text{ is open}\}}\right] = \sigma(n)p^n.$$

The origin belongs to an infinite open cluster if there are open paths of all possible lengths beginning

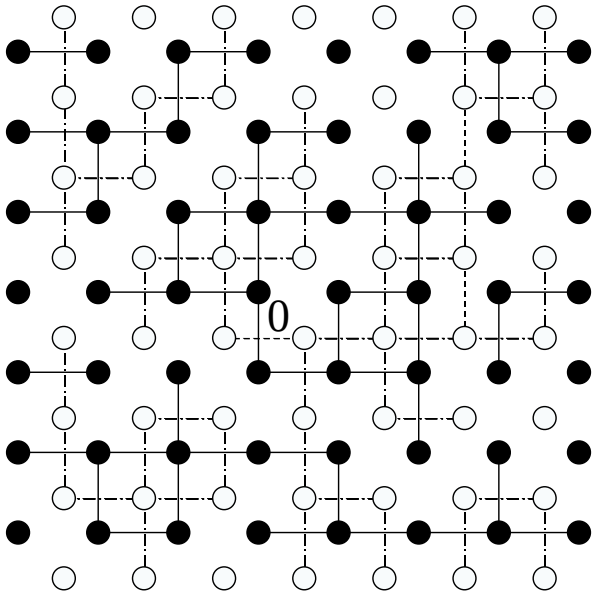

Figure 2.1: A portion of the lattice $\mathbb{L}^2$ (whose vertices are represented by full circles, open edges by plain lines) and its dual (whose vertices are represented by empty circles, and open edges by dashed lines).

at the origin, hence for all $n \in \mathbb{N}^*$

$$\begin{aligned}\theta(p) &\leq \mathbb{P}_p(N(n) \geq 1) = \sum_{s=1}^{\sigma(n)} \mathbb{P}_p(N(n) = s) \\ &\leq \sum_{s=1}^{\sigma(n)} s\mathbb{P}_p(N(n) = s) = \mathbb{E}_p[N(n)] = \sigma(n)p^n \\ &\leq \frac{4}{3}(3p)^n.\end{aligned}$$

Letting $n \to \infty$, we find that $\theta(p) = 0$ if $p < 1/3$. Hence $p_c \geq 1/3$.

(ii) We next prove $p_c \leq 2/3$. Let $m \in \mathbb{N}^*$, and let F_m be the event that there exists a closed circuit in the dual lattice $\mathbb{L}_d^2$ containing the box $B(m) = [-m, m] \times [-m, m]$ in its interior, and let G_m be the event that all edges of $B(m)$ are open. The origin belongs to an infinite cluster if F_m does not occur and G_m does occur, see Figure 2.2. Since these events are defined on disjoint sets of edges, they are independent and we have therefore that

$$\theta(p) \geq \mathbb{P}_p(\overline{F}_m \cap G_m) = \mathbb{P}_p(\overline{F}_m)\mathbb{P}_p(G_m). \tag{2.3}$$

Now, $\mathbb{P}_p(G_m) > 0$ and so all we need to do is to show that $\mathbb{P}_p(\overline{F}_m) > 0$ for $p \geq 2/3$.

Let $\gamma(n)$ be the number of self-avoiding circuits in the dual lattice L_d^2 surrounding the origin and of length n, and which consists of a single loop (In other words, the degree of every vertex of such a closed circuit is 2: we will speak of a "self-avoiding circuit"). Each such circuit must pass through a vertex of the form $(i + 1/2, 1/2)$ for some $0 \leq i \leq n - 1$, because (a) to surround the origin, it has to pass through a vertex $(i + 1/2, 1/2)$ for some $i \geq 0$, and (b) it cannot pass through a vertex $(i + 1/2, 1/2)$ for some $i \geq n$ since it would then be at least $2n$. Such a circuit contains a self-avoiding walk of length $n - 1$ starting from one of the n vertices $(i + 1/2, 1/2)$ for some

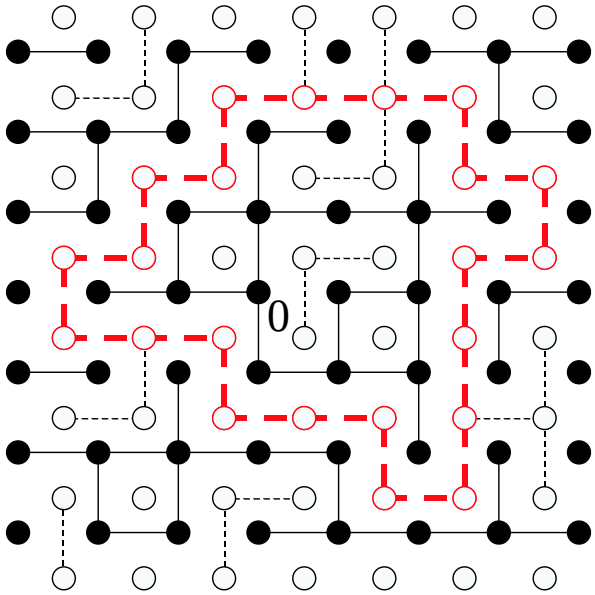

Figure 2.2: A portion of the lattice $\mathbb{L}^2$ (whose vertices are represented by full circles, open edges by plain lines) and its dual (whose vertices are represented by empty circles, and closed edges by dashed lines) Observe that there is a circuit of closed dual edges surrounding the origin (set in red bold on the figure), which therefore belongs to a finite open cluster.

$0 \leq i \leq n-1$. Therefore

$$\gamma(n) \leq n\sigma(n-1).$$

Now, the occurrence of the event F_m requires that there is at least one such closed circuit, with a length of at least $8m$ hops to contain $B(m)$:

$$\begin{aligned} F_m &\subseteq \{\text{there is at least one closed circuit of length } 8m \text{ surrounding } 0\} \\ &= \bigcup_{\text{circuit } g \text{ of length at least } 8m} \{g \text{ is closed}\}. \end{aligned}$$

Using the union bound, we get therefore that

$$\begin{aligned} \mathbb{P}_p(F_m) &\leq \sum_{\text{circuit } g \text{ of length at least } 8m} \mathbb{P}_p\left(g \text{ is closed}\right) \\ &= \sum_{n=8m}^{\infty} \sum_{\text{circuit } g \text{ of length } n} \mathbb{P}_p\left(g \text{ is closed}\right) \\ &\leq \sum_{n=8m}^{\infty} \gamma(n)\left(1-p\right)^n \\ &\leq \frac{4(1-p)}{3} \sum_{n=8m}^{\infty} n\left(3(1-p)\right)^{n-1}. \end{aligned} \tag{2.4}$$

If $p > 2/3$, this sum converges to some finite value, and we take m large enough so that it is less than $1/2$. Consequently, from (2.3), we get

$$\theta(p) \geq \mathbb{P}_p(\overline{F}_m)\mathbb{P}_p(G_m) \geq P_p(G_m)/2 > 0,$$

which proves the result.

■

For the 2-dim bond model, the exact value of p_c is known, and we will compute it in a few chapters, as it requires quite a lot of work. The simulations of Figure 2.3 show a 40×40 lattice. Although non infinite, the phase transition is already visible: all clusters are finite for $p = 0.3$ and $p = 0.49$, whereas one giant cluster is present for $p = 0.51$ and clearly for $p = 0.7$.

Figure 2.4 displays an estimate of the percolation probability $\theta(p)$, for a 5000×5000 lattice. Despite the finite size of the lattice, the phase transition, which stricto sensu only occurs for an infinite lattice, appears quite clearly on the figure.

In higher dimensions, the d-dim lattice $\mathbb{L}^d$ can always be embedded in a $(d+1)$-dim lattice $\mathbb{L}^{d+1}$, and therefore if the origin belongs to an infinite cluster in $\mathbb{L}^d$, it also belongs to an infinite cluster in $\mathbb{L}^{d+1}$. Therefore, the percolation threshold is a decreasing function of d: $p_c(d+1) \leq p_c(d)$.

A direct corollary of Theorem 2.1 is that the probability that there exists an infinite open cluster, which we denote by $\hat{\theta}(p)$ follows a zero-one law (see Appendix 2.10).

Corollary 2.1. *Existence of an open cluster The probability that there exists an open infinite cluster is*

$$\hat{\theta}(p) = \begin{cases} 0 & \text{if} \quad p < p_c \\ 1 & \text{if} \quad p > p_c. \end{cases}$$

We will however give a stronger result in the chapter on the super-critical phase.

2.4 Mean cluster size

The other quantities of interest in percolation theory are the *mean size of an open cluster*, which by translation invariance is the expected number of vertices in the open cluster at the origin, and which we denote

$$\chi(p) = \mathbb{E}_p[|C|]. \tag{2.5}$$

Expanding this expression, we have that

$$\chi(p) = \mathbb{E}_p[|C|] = \sum_{n=1}^{\infty} n\mathbb{P}_p(|C| = n) + \infty\mathbb{P}_p(|C| = \infty).$$

If $p > p_c$, then we see that $\chi(p) = \infty$. The converse is not obvious, and it will require quite some work to prove in the next chapter that if $p < p_c$ then $\chi(p) < \infty$. Figure 2.5 displays an estimate of the mean cluster size $\chi(p)$, for the 5000×5000 lattice.

In the supercritical phase, since the mean cluster size is infinite, one is more interested in the mean size of the finite clusters, which we denote $\chi^f(p)$ and which is defined as the mean of $|C|$ on the event that $|C|$ is finite:

$$\chi^f(p) = \mathbb{E}_p[|C|; |C| < \infty] = \mathbb{E}_p\left[|C|\, 1_{\{|C|<\infty\}}\right] = \mathbb{E}_p[|C| \mid |C| < \infty](1 - \theta(p)). \tag{2.6}$$

2.5 Increasing event

In the next three sections, we introduce three technical devices, which will be repeatedly used in the proofs of theorems in the following chapters. We need first the following definition.

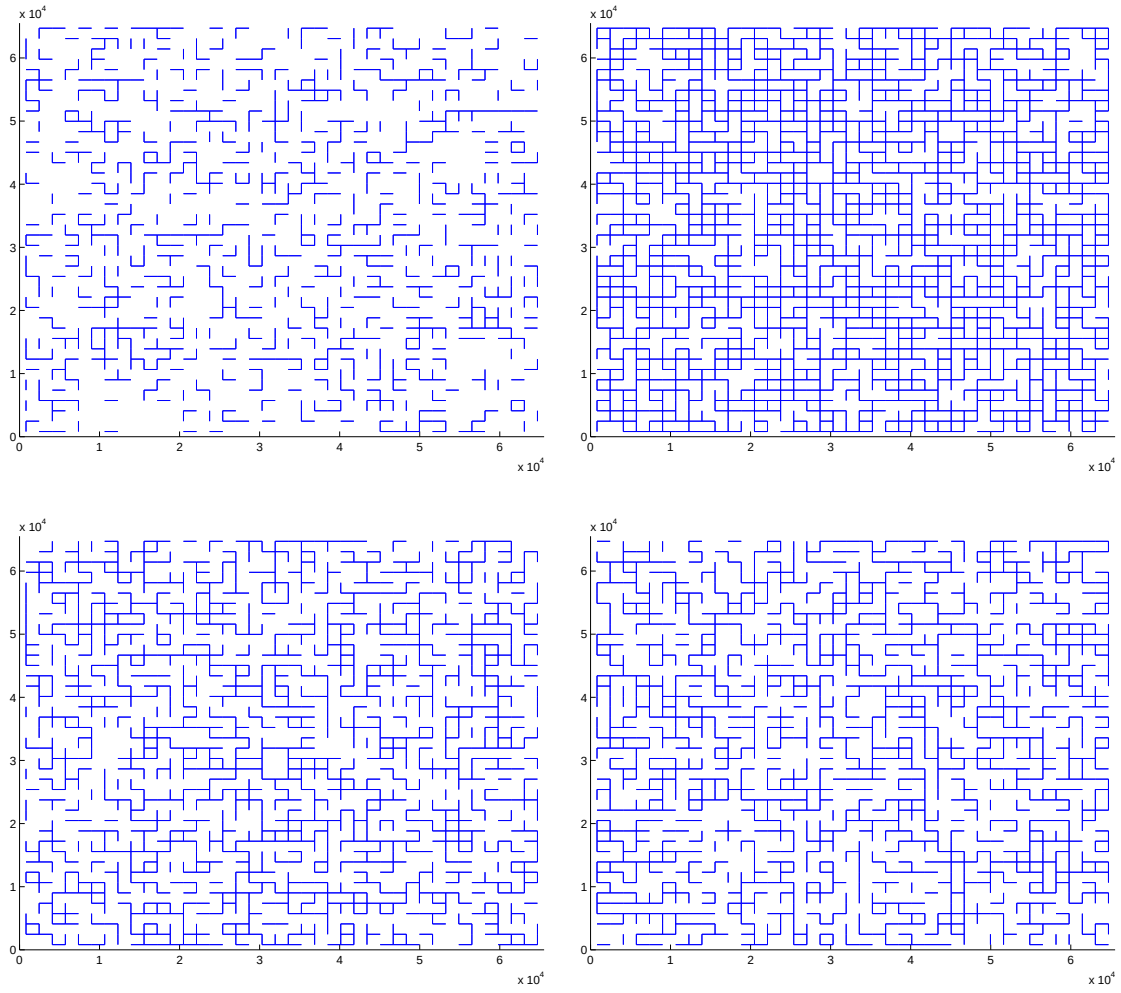

Figure 2.3: A simulation of bond percolation in a 40×40 lattice for different values of the open edge probability: $p = 0.3$ (upper left), $p = 0.49$ (bottom left), $p = 0.51$ (bottom right) and $p = 0.7$ (top right). Only the open edges are shown. A careful inspection of the two graphs at bottom reveals the emergence of a giant open cluster for $p \geq 0.51$, which was absent when $p \leq 0.49$.

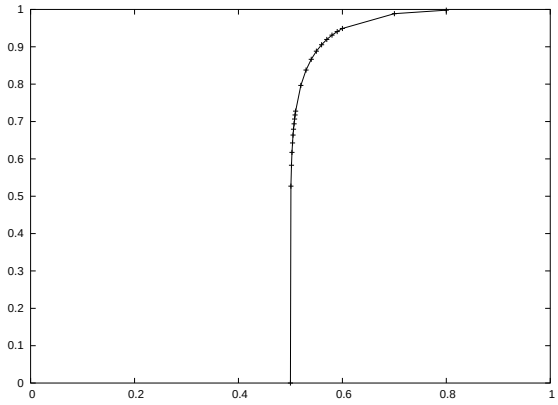

Figure 2.4: An estimation of the percolation probability $\theta(p)$, for a 5000×5000 lattice, as a function of p.

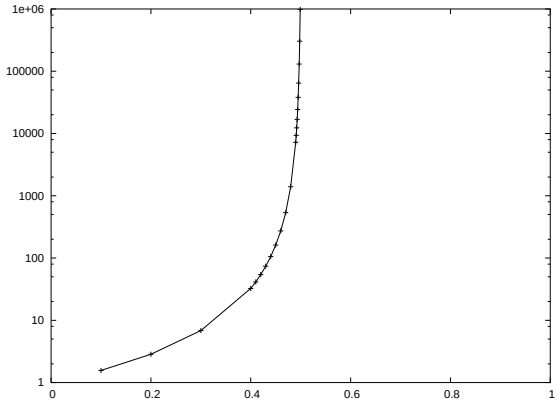

Figure 2.5: An estimation of the mean cluster size $\chi(p)$, for the 5000×5000 lattice, as a function of p.

Definition 2.1 (Increasing event). *A random variable X is* increasing *on $(\Omega, \mathcal{F})$ if $X(\omega) \leq X(\omega')$ whenever $\omega \leq \omega'$. It is decreasing if $-X$ is increasing. An event $A \in \mathcal{F}$ is increasing whenever its indicator function is an increasing variable, i.e. if $1_A(\omega) \leq 1_A(\omega')$ whenever $\omega \leq \omega'$.*

It is easy to show that if A is an increasing event, then $\mathbb{P}_p(A) \leq \mathbb{P}_{p'}(A)$ whenever $p \leq p'$.

2.6 FKG inequality

The FKG inequality (named after Fortuin Kasteleyn and Ginibre) was first shown by Harris in 1960. It expresses the fact that increasing events can only be positively correlated.

Theorem 2.2 (FKG inequality). *If A and B are two increasing events, then*

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

We establish the FKG inequality in the case where A and B are depend on finitely many edges. The proof of the FKG inequality when A and/or B depend on infinitely many edges is is found in [18]. The FKG inequality also holds when both A and B are two decreasing events.

Proof: Let $X = 1_A$ and $Y = 1_B$ be the indicators of the increasing events A and B, which are increasing random variables. We can then reformulate the FKG inequality as $\mathbb{E}_p[XY] \geq \mathbb{E}_p[X]\mathbb{E}_p[Y]$. Suppose that X and Y depend only on the state of edges $e_1, e_2, \ldots, e_n$ for some integer n. We prove the FKG inequality by induction.

Suppose first that $n = 1$, so that X and Y are only functions of the state $\omega(e_1)$ of the edge e_1. Pick any two states $\omega_1, \omega_2 \in \{0, 1\}$. Since both X and Y are increasing random variables,

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

with equality if $\omega_1 = \omega_2$. Therefore

$$\begin{aligned}
0 &\leq \sum_{\omega_1=0}^{1} \sum_{\omega_2=0}^{1} (X(\omega_1) - X(\omega_2))(Y(\omega_1) - Y(\omega_2))\mathbb{P}_p(\omega(e_1) = \omega_1)\mathbb{P}_p(\omega(e_1) = \omega_2) \\
&= \sum_{\omega_1=0}^{1} X(\omega_1)Y(\omega_1)\mathbb{P}_p(\omega(e_1) = \omega_1) + \sum_{\omega_2=0}^{1} X(\omega_2)Y(\omega_2)\mathbb{P}_p(\omega(e_1) = \omega_2) \\
&\quad - \sum_{\omega_1=0}^{1} \sum_{\omega_2=0}^{1} (X(\omega_1)Y(\omega_2) + X(\omega_2)Y(\omega_1)\mathbb{P}_p(\omega(e_1) = \omega_1)\mathbb{P}_p(\omega(e_1) = \omega_2) \\
&= 2\left(\mathbb{E}_p[XY] - \mathbb{E}_p[X]\mathbb{E}_p[Y]\right).
\end{aligned}$$

Let $1 < k \leq n$. Suppose now that the claim holds for all $m < k$, and that X and Y depend only on the states $\omega(e_1), \ldots, \omega(e_k)$ of the edges $e_1, \ldots, e_k$. Then, given $\omega(e_1), \ldots, \omega(e_{k-1})$, X and Y only depend on the state $\omega(e_k)$ of the edge e_k, and proceeding as above, we have that

$$\mathbb{E}_p[XY \mid \omega(e_1), \ldots, \omega(e_{k-1})] \geq \mathbb{E}_p[X \mid \omega(e_1), \ldots, \omega(e_{k-1})]\mathbb{E}_p[Y \mid \omega(e_1), \ldots, \omega(e_{k-1})]$$

and thus

$$\begin{aligned}
\mathbb{E}_p[XY] &= \mathbb{E}_p\left[\mathbb{E}_p[XY \mid \omega(e_1), \ldots, \omega(e_{k-1})]\right] \\
&\geq \mathbb{E}_p\left[\mathbb{E}_p[X \mid \omega(e_1), \ldots, \omega(e_{k-1})]\mathbb{E}_p[Y \mid \omega(e_1), \ldots, \omega(e_{k-1})]\right].
\end{aligned}$$

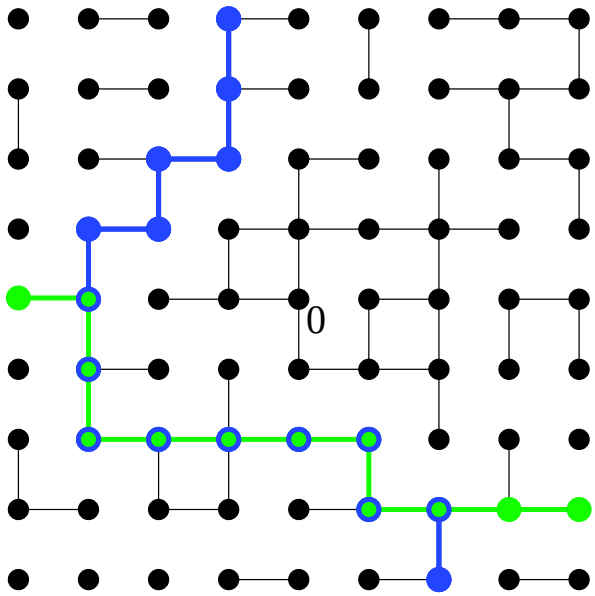

Figure 2.6: The box $B(5)$ with a LR and a TB open crossing.

Now, $\mathbb{E}_p[X \mid \omega(e_1), \ldots, \omega(e_{k-1})]$ and $\mathbb{E}_p[Y \mid \omega(e_1), \ldots, \omega(e_{k-1})]$ are increasing functions of the state of the $(k-1)$ edges $e_1, \ldots, e_{k-1}$. By induction, it implies that

$$\begin{aligned} \mathbb{E}_p[XY] &\geq \mathbb{E}_p\left[\mathbb{E}_p[X \mid \omega(e_1), \ldots, \omega(e_{k-1})]\right] \cdot \mathbb{E}_p\left[\mathbb{E}_p[Y \mid \omega(e_1), \ldots, \omega(e_{k-1})]\right] \\ &= \mathbb{E}_p[X]\mathbb{E}_p[Y]. \end{aligned}$$

■

As an example of application of the FKG inequality, consider the 2-dim. box $B(n)$, and let A be the event that there is an open path joining a vertex of the top face of $B(n)$ to the bottom face of $B(n)$ (we call such a path a TB (top-bottom) (open) crossing of $B(n)$, and B be the event that there is an open path joining a vertex of the left face of $B(n)$ to the right face of $B(n)$ (we call such a path a LR (left-right) (open) crossing of $B(n)$, as shown in Figure 2.6. Then the probability that there are both a TB and LR open crossings of $B(n)$ is at least the product of the probabilities that there is a TB open crossing and that there is a LR open crossing.

2.7 BK inequality

The BK inequality (named after van den Berg and Kesten, who proved it in 1985) can be regarded as the reverse of the FKG inequality, with one difference: it applies to the event $A \circ B$ that two increasing events A and B occur on *disjoint* sets of edges, and not to the larger event $A \cap B$ that events A and B occur on any sets of edges. $A \circ B$ is the set of configurations $\omega \in \Omega$ for which there are disjoint sets of open edges such that the first set guarantees the occurrence of A and the second set guarantees the occurrence of B. The formal definition is as follows.

Definition 2.2 (Disjoint occurrence). *Let A and B be two increasing events which depends on the states $\omega(e_1), \ldots, \omega(e_n)$ of n distinct edges $e_1, \ldots e_n$ of $\mathbb{L}^d$. Each such configuration is specified uniquely by the subset $K(\omega) = \{e_i \mid \omega(e_i) = 1\}$ of open edges among these n edges. Then $A \circ B$ is the*

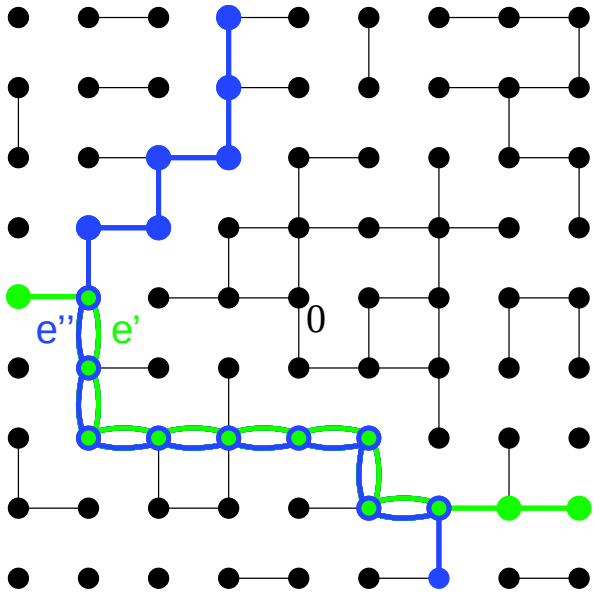

Figure 2.7: Construction of two independent copies of the lattice

set of ω for which there exists a subset $H \subset K(\omega)$ such that any ω' determined by $K(\omega') = H$ is in A and any ω'' determined by $K(\omega'') = K(\omega) \setminus H$ is in B.

Theorem 2.3 (BK inequality). *If A and B are two increasing events, then*

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

We only sketch the intuition behind the proof of van den Berg when A and B are the existence of two open paths between different sets of vertices. The full proof is given in [18]. The BK inequality also holds when both A and B are two decreasing events.

Proof: *(Sketch)* Let G be a finite subgraph of $\mathbb{L}^d$. Let A (respectively, B) be the event that there exists an open path between vertices u and v (respectively, x and y). $A \circ B$ is the event that there exist two disjoint open paths from u to v and from x to y. Let e be an edge of G. Replace e by two parallel edges e' and e'', having the same end vertices, each of which being open with the same probability p, independently of each other and of all other edges. The splitting of edge e in the two edges e' and e'' can only make our search for two disjoint open paths easier: indeed, if in graph G two paths from u to v and from x to y had to use the same edge e, they now can replace this common edge by the two distinct edges e' and e''. The probability of finding two disjoint open paths from from u to v and from x to y can therefore only increase or remain equal after this splitting. We continue this splitting process, as shown in Figure 2.7, replacing every edge f of G by two parallel edges f' and f''. At each stage, we look for two open paths, the first one avoiding all edges marked $''$ and the second one all edges marked $'$. The probability of finding two such paths can only increase or remain equal at each stage. When all edges of G have been split in two, we end up with two independent copies of G, in the first of which we look for an open path connecting u to v, and in the second of which we look for an open path connecting x to y. Since such paths occur independently in each copy of G, the probability that they both occur is $\mathbb{P}_p(A)\mathbb{P}_p(B)$.

■

As an example of application of the BK inequality, consider again the 2-dim. box $B(n)$, and let A be the event that there is an open TB crossing path of $B(n)$ to the bottom face of $B(n)$, and B be

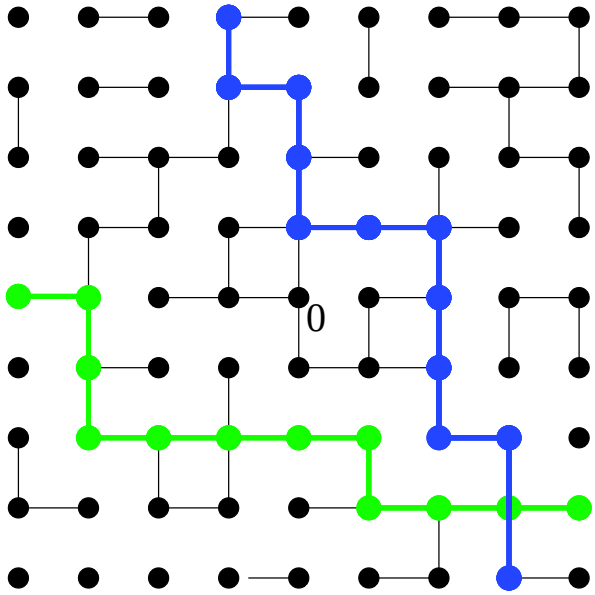

Figure 2.8: The box $B(5)$ with edge-disjoint LR and TB open crossings.

the event that there is an open LR crossing, which is edge-disjoint with A. This event does not occur in the example of Figure 2.6, but does occur for the example of Figure 2.8. Then the probability that there are edge disjoint TB and LR open crossings of $B(n)$ is no more than the product of the probabilities that there is a TB open crossing and that there is a LR open crossing.

2.8 Russo's formula

The third relation estimates the rate of change of the probability of occurrence of an increasing event A as p increases. We need first to introduce the definition of a pivotal edge. If A is increasing, an edge e is pivotal if and only if A occurs when e is open and does not occur is e is closed. A pivotal edge is thus a critical edge for the occurrence of A.

Definition 2.3 (Pivotal edge). *Let A be an event, and let ω be a configuration. The edge e is pivotal for the pair (A, ω) if the occurrence of A crucially depends on e, i.e., if $1_A(\omega) \neq 1_A(\omega')$ where ω' is the configuration such that $\omega'(e) = 1 - \omega(e)$ and $\omega'(f) = \omega(f)$ for all $f \in \mathbb{E}^d \setminus \{e\}$.*

The event "e is pivotal for A" is the set of all configurations ω for which e is pivotal for (A, ω). Observe that this event is independent from the state of e itself, but only depends on the state of the other edges.

For example, let A be the event that there is a LR open crossing of the 2-dim box $B(n)$. Any edge e of $B(n)$ is pivotal for A if, when it is removed from the graph, there is no more LR open crossing of $B(n)$, but one endvertex of e is joined to the left side of $B(n)$ by an open path, while the other endvertex of e is joined to the right side of $B(n)$ by another open path.

Figure 2.9 shows another example of edges that are pivotal for the event that the origin is connected by an open path to the boundary $\partial S(n)$ of a diamond box $S(n)$.

Theorem 2.4 (Russo's formula). *Let A be an increasing event, which depends on the state of*

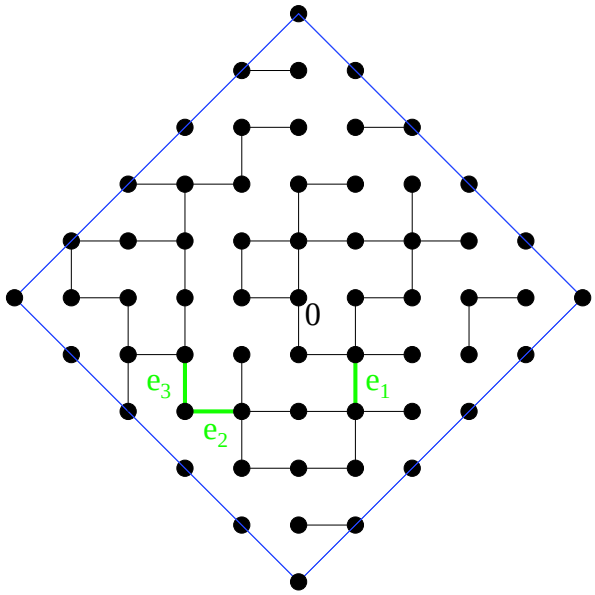

Figure 2.9: The three edges e_1, e_2 and e_3 are pivotal for the event $0 \leftrightarrow \partial S(5)$.

finitely many edges of $\mathbb{L}^d$*, and let* $N(A)$ *denote the number of edges that are pivotal for* A*. Then*

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

We give only a sketch of the proof, the full proof is not very long either and can be found in [18].

Proof: *(Sketch)*

Let $\{X(e), e \in \mathbb{E}^d\}$ be a collection of i.i.d. random variables indexed by the edge set $\mathbb{E}^d$, uniformly distributed on $[0,1]$. Let η_p be the configuration of edges defined by

$$\eta_p(e) = \begin{cases} 1 & \text{if} \quad X(e) < p \\ 0 & \text{if} \quad X(e) \geq p \end{cases}$$

for some $0 \leq p \leq 1$ and all $e \in \mathbb{E}^d$. Observe that

$$\begin{aligned} \mathbb{P}(\eta_p(e) = 0) &= \mathbb{P}(X(e) \geq p) = 1 - p \\ \mathbb{P}(\eta_p(e) = 1) &= \mathbb{P}(X(e) < p) = p. \end{aligned}$$

Hence $\mathbb{P}_p(A) = \mathbb{P}(\eta_p \in A)$. As A is an increasing event, we have that for $\delta > 0$

$$\begin{aligned} \mathbb{P}_{p+\delta}(A) &= \mathbb{P}(\eta_{p+\delta} \in A) \\ &= \mathbb{P}(\{\{\eta_{p+\delta} \in A\} \cap \{\eta_p \notin A\}\} \cup \{\eta_p \in A\}) \\ &= \mathbb{P}(\{\eta_{p+\delta} \in A\} \cap \{\eta_p \notin A\}) + \mathbb{P}(\eta_p \in A) \\ &= \mathbb{P}(\{\eta_{p+\delta} \in A\} \cap \{\eta_p \notin A\}) + \mathbb{P}_p(A). \end{aligned} \tag{2.7}$$

Now, if $\eta_p \notin A$ while $\eta_{p+\delta} \in A$, it means that there are some edges e on which A depends, and for which $\eta_p(e) = 0$ and $\eta_{p+\delta}(e) = 1$, or equivalently, $p \leq X(e) < p+\delta$. As A depends only on the state of finitely many edges, the probability that there are more than one edge e with $p \leq X(e) < p + \delta$ is negligible (of the order $o(\delta)$) in front of the probability that there is one such edge, when $\delta \downarrow 0$.

If e is the only edge for which $p \le X(e) < p + \delta$, then e must be a pivotal edge for A, in the sense that $\eta_p \notin A$ but $\eta_{p'} \in A$ where $\eta_{p'}(e) = 1 = 1 - \eta_p(e)$ and $\eta_{p'}(e') = \eta_p(e')$ for all other edges $e' \neq e$. Therefore

$$\begin{aligned}
\mathbb{P}(\{\eta_{p+\delta} \in A\} \cap \{\eta_p \notin A\}) &= \sum_{e \in \mathbb{E}^d} \mathbb{P}(\{e \text{ is pivotal for } A\} \cap \{p \le X(e) < p + \delta\}) + o(\delta) \\
&= \sum_{e \in \mathbb{E}^d} \mathbb{P}(e \text{ is pivotal for } A)\mathbb{P}(p \le X(e) < p + \delta) + o(\delta) \\
&= \delta \sum_{e \in \mathbb{E}^d} \mathbb{P}(e \text{ is pivotal for } A) + o(\delta)
\end{aligned}$$

where the second equality follows from the independence of the state of an edge with the fact that is pivotal or not. Inserting this relation in (2.7), dividing by δ and taking the limit as $\delta \downarrow 0$, we get

$$\frac{d}{dp}\,\mathbb{P}_p(A) = \sum_{e \in \mathbb{E}^d} \mathbb{P}(e \text{ is pivotal for } A).$$

The right hand side of this last equation is $E_p[N(A)]$.

■

We can also recast Russo's formula in an integral form.

Corollary 2.2. *Let A be an increasing event, which depends on the state of finitely many edges of $\mathbb{L}^d$, and let $N(A)$ denote the number of edges that are pivotal for A. Then for any $0 \le p_1 < p_2 \le 1$*

$$\mathbb{P}_{p_2}(A) = \mathbb{P}_{p_1}(A) \exp\left(\int_{p_1}^{p_2} \frac{1}{p} \mathbb{E}_p[N(A) \mid A] dp \right).$$

Proof: From Russo's formula, as the state of an edge e is independent of the fact that e is pivotal for A or not,

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

Dividing both members of the last equality by $\mathbb{P}_p(A)$, and integrating from p_1 to p_2 gives the result.

■

2.9 Multiplicity of edge-disjoint paths

A last result, that uses the same coupling argument as above, is useful to relate the probability that r edge-disjoint paths cross a given portion of the lattice to the probability that at leat one such path exists. It follows directly from Theorem 2.45 in [18] and the remarks thereafter.

Lemma 2.1. *Let A_n be the event that there exists an open path between the left and right sides of $B(n)$ and $I_r(A_n)$ the event that there exist r edge-disjoint such LR crossings. We have*

$$1 - \mathbb{P}_p(I_r(A_n)) \leq \left(\frac{p}{p-p'}\right)^r [1 - \mathbb{P}_{p'}(A_n)]$$

for any $0 \leq p' < p \leq 1$.

Proof: Let $\{X(e), e \in \mathbb{E}^2\}$ be a collection of i.i.d. random variables indexed by the edge set $\mathbb{E}^2$, uniformly distributed on $[0,1]$. Let η_p be the configuration of edges defined by

$$\eta_p(e) = \begin{cases} 1 & \text{if} \quad X(e) < p \\ 0 & \text{if} \quad X(e) \geq p \end{cases}$$

for some $0 \leq p \leq 1$ and all $e \in \mathbb{E}^d$. Observe that

$$\begin{aligned} \mathbb{P}(\eta_p(e) = 0) &= \mathbb{P}(X(e) \geq p) = 1 - p \\ \mathbb{P}(\eta_p(e) = 1) &= \mathbb{P}(X(e) < p) = p. \end{aligned}$$

Hence $\mathbb{P}_p(A_n) = \mathbb{P}(\eta_p \in A_n)$.

Observe that the configuration of edges η_p does not have r LR edge-disjoint crossings of $B(n)$ (in other words, $\eta_p \notin I_r(A_n)$), if and only if there is a (possibly empty) collection C of edges, with (i) $|C| \leq r$, (ii) $\eta_p(e) = 1$ for all $e \in C$, and (iii) the configuration $(\eta_p \setminus C)$ obtained by declaring all edges in C closed does not have a LR open path that crosses $B(n)$, i.e. $(\eta_p \setminus C) \notin A_n$. Indeed, there are less than r edge-disjoints LR crossings if and only if we can find at most r edges that form the minimal cutset of the graph between the left and right sides of $B(n)$.

Suppose that $\eta_p \notin I_r(A_n)$. Then

$$\begin{aligned} \mathbb{P}(\eta_{p'} \notin A_n \mid \eta_p \notin I_r(A_n)) &= \mathbb{P}(\eta_{p'} \notin A_n \mid \text{there is a set } C \text{ verifying (i) - (iii) above}) \\ &= \mathbb{P}(\eta_{p'}(e) = 0 \text{ for all } e \in C \mid \text{there is a set } C \text{ verifying (i) - (iii) above}) \\ &= \frac{\mathbb{P}(\{\text{there is a set } C \text{ verifying (i) - (iii) above}\} \cap \{\eta_{p'}(e) = 0 \text{ for all } e \in C\})}{\mathbb{P}(\text{there is a set } C \text{ verifying (i) - (iii) above})} \\ &= \frac{\mathbb{P}(\{\text{there is a set } C \text{ verifying (i) - (iii) above}\} \cap \{p' \leq X(e) < p \text{ for all } e \in C\})}{\mathbb{P}(\{\text{there is a set } C \text{ verifying (i) - (iii) above}\} \cap \{X(e) < p \text{ for all } e \in C\})} \\ &\geq \left(\frac{p-p'}{p}\right)^r \end{aligned}$$

and

$$\mathbb{P}(\{\eta_{p'} \notin A_n\} \cap \{\eta_p \notin I_r(A_n)\}) \geq \left(\frac{p-p'}{p}\right)^r \mathbb{P}(\eta_p \notin I_r(A_n)),$$

from which we deduce that

$$\begin{aligned} 1 - \mathbb{P}_p(I_r(A_n)) &= \mathbb{P}_p(\eta_p \notin I_r(A_n)) \\ &\leq \left(\frac{p}{p-p'}\right)^r \mathbb{P}(\{\eta_{p'} \notin A_n\} \cap \{\eta_p \notin I_r(A_n)\}) \\ &\leq \left(\frac{p}{p-p'}\right)^r \mathbb{P}(\eta_{p'} \notin A_n) \\ &= \left(\frac{p}{p-p'}\right)^r [1 - \mathbb{P}_{p'}(A_n)]. \end{aligned}$$

■

This theorem is in fact much more general. First, it is not restricted to a portion of $\mathbb{L}^d$, which is a box $B(n)$. Second, it applies to any increasing event A, if we define $I_r(A)$ to be the *interior* of A with depth r, defined as the set of configuration in A, which remain in A even if the states of up to r edges is modified (see [18]).

2.10 Appendix: Kolmogorov's zero-one law and tail events

Let $\{X_n, n \in \mathbb{N}^*\}$ be a sequence of independent random variables. A *tail event* is an event whose occurrence or failure is determined by the values of these random variables, but which does not depend probabilistically of any finite subsequence of these random variables.

For example, the event $\{\sum_{n\in\mathbb{N}^*} X_n \text{ converges }\}$ is a tail event, because if we remove any finite subcollection of X_n, it does not change the convergence property of the series. Likewise, the event $\{\limsup_{n\to\infty} \frac{1}{n}\sum_{m=1}^{n} X_m \leq 2\}$ is a tail event. On the contrary, for if $\sum_{n\in\mathbb{N}^*} X_n$ converges, the event $\{\sum_{n\in\mathbb{N}^*} X_n \leq 2\}$ does change if we remove some finite subcollection of X_n, and thus is not a tail event.

In the case of Corollary 2.1, let X_n denotes the state of an edge and A be the existence of an infinite open cluster. Then A does not depend on any finite subcollection of variables X_n, and is therefore a tail event.

Tail events enjoy the following property.

Theorem 2.5 (Kolmogorov's zero-one law). *If $\{X_n, n \in \mathbb{N}^*\}$ is a sequence of independent variables, then any tail event A satisfies $\mathbb{P}(A) = 0$ or $\mathbb{P}(A) = 1$.*

The following corollary of the zero-one law will be useful later on (see [19]). Let Y be a random variable which is a function of the variables X_n. Then Y is a *tail function* if, roughly speaking, it does not depend crucially on any finite subcollection of X_n. More precisely, Y is a tail function if and only if the event $\{\omega \in \Omega \mid Y(\omega) \leq y\}$ is a tail event for all $y \in \mathbb{R}$.

For example, the random variable

$$Y = \limsup_{n\to\infty} \frac{1}{n} \sum_{m=1}^{n} X_n$$

is a tail function of the independent variables X_n.

Tail functions are almost surely constant. Indeed, since $\{\omega \in \Omega \mid Y(\omega) \leq y\}$ is a tail event for all $y \in \mathbb{R}$, $\mathbb{P}(Y \leq y)$ can only take the values 0 and 1. Let $k = \inf\{y \mid \mathbb{P}(Y \leq y) = 1\}$. Then for any $y \in \mathbb{R}$, $\mathbb{P}(Y \leq y) = 0$ when $y < k$ and $\mathbb{P}(Y \leq y) = 1$ when $y \geq k$.

Theorem 2.6 (Constant tail functions). *If Y is a tail function of the independent variables $X_n, n \in \mathbb{N}^*$, then there exists some $k \in \mathbb{Z} \cup \{-\infty, \infty\}$ such that $\mathbb{P}(Y = k) = 1$.*

3

Subcritical phase

In this chapter we study the situation in the subcritical phase, when $p < p_c$ and $d \geq 2$. In this case, we know that the open cluster C containing the origin is almost surely finite since $\theta(p) = \mathbb{P}_p(|C| = \infty) = 0$.

We will study the mean size of the open cluster containing the origin, i.e.

$$\chi(p) = \mathbb{E}_p[|C|] = \sum_{n=1}^{\infty} n\mathbb{P}_p(|C| = n) + \infty\mathbb{P}_p(|C| = \infty) = \sum_{n=1}^{\infty} n\mathbb{P}_p(|C| = n).$$

Because the process is space invariant, we can replace the origin by any vertex, so that $\chi(p)$ is the mean size of an open cluster.

The main result from this chapter, which is probably the most difficult we will have to prove in this course, is that the radius of the mean cluster size decreases exponentially when $p < p_c$, As a result, the mean cluster size the mean cluster size is finite in the subcritical phase: $\chi(p) < \infty$ when $p < p_c$.

3.1 Exponential decrease of the radius of the mean cluster size

Let $S(n)$ be the diamond of radius n (i.e., the ball of radius n with the Manhattan distance), that is the set of all vertices $x \in \mathbb{Z}^d$ for which $\delta(0,x) = |x| \leq n$. Let $A_n = \{0 \leftrightarrow \partial S(n)\}$ be the event that there exists a open path connecting the origin to any vertex lying on the surface of $S(n)$, which we denote by $\partial S(n)$. Defining the radius of C by $\mathrm{rad}(C) = \max_{x \in C}\{|x|\}$, we see that $A_n = \{\mathrm{rad}(C) \geq n\}$.

We follow the approach of Menshikov (1986), as exposed in [18], to prove that the radius of the cluster at the origin (and by space invariance, any cluster) has a tail that decreases at least exponentially in the subcritical phase.

Theorem 3.1 (Exponential decay of the radius of an open cluster). *If $p < p_c$, there exists $\psi(p) > 0$ such that*

$$\mathbb{P}_p(rad(C) \geq n) = \mathbb{P}_p(0 \leftrightarrow \partial S(n)) = \mathbb{P}_p(A_n) < \exp(-n\psi(p)).$$

The proof of this theorem is rather long, and we will need several lemmas to establish it.

The starting point is Russo's formula, expressed as in Corollary 2.2, which states that for any $0 \leq p_1 < p_2 \leq 1$, we have that

$$\mathbb{P}_{p_2}(A_n) = \mathbb{P}_{p_1}(A_n) \exp\left(\int_{p_1}^{p_2} \frac{1}{p}\mathbb{E}_p[N(A_n) \mid A_n]dp\right).$$

Denoting by $g_p(n) = \mathbb{P}_p(A_n)$, we can restate this inequality as

$$g_{p_1}(n) = g_{p_2}(n) \exp\left(-\int_{p_1}^{p_2} \frac{1}{p}\mathbb{E}_p[N(A_n) \mid A_n]dp\right) \tag{3.1}$$

We will choose $p_1 < p_c$ and we will show that the mean number of pivotal edges, given that A_n occurs, grows roughly linearly with n when $p < p_c$. The idea is that since $p < p_c$, then $\mathbb{P}_p(A_n) \to 0$ as $n \to \infty$, so that if A_n occurs, then it must depend critically on many edges, because there can only be very few open paths connecting 0 to $\partial S(n)$. As a result, one expects that the average number of pivotal edges for A_n linearly increases with n. We need thus to prove that $\mathbb{E}_p[N(A_n) \mid A_n]$ grows roughly linearly with n when $p < p_c$.

Before computing $\mathbb{E}_p[N(A_n) \mid A_n]$, we show the following lemma. Denote by $e_1, e_2, \ldots, e_N$ the (random) edges that are pivotal for A_n. Any path connecting the origin to $\partial S(n)$ uses one of these edges, as otherwise they would not be pivotal for A_n. We label the edges in the order of encountering when we move on a such a path from the origin to $\partial S(n)$, and we denote by x_i (respectively, y_i) the first (resp., second) endvertex of the pivotal edge e_i in the order of encountering from the origin to the surface, see the example shown in Figure 3.1. Hence $e_i = \langle x_i, y_i \rangle$. There are at least two edge disjoint paths from 0 (resp., any vertex y_{i-1}) to x_1 (resp,. and vertex x_i). Indeed, if this was not the case, then there would be a pivotal edge between 0 and x_1, a contradiction. The open cluster appears as a set of well meshed "islands" connected to each other by pivotal edges.

Let $R_i = \delta(y_{i-1}, x_i)$ for $1 \leq i \leq N$, with $y_0 = 0$. The random variables R_i, $1 \leq i \leq N$, are thus the (Manhattan) distances traversed by the shortest path connecting the origin to $\partial S(n)$, within the i-th "island" encountered as we walk along this path starting from the origin. The distribution of the random variables R_i is linked to $\mathbb{E}_p[N(A_n) \mid A_n]$ as follows. Knowing that A_n occurs, if $R_1 + R_2 + \ldots + R_k \leq n - k$, then the number $N(A_n)$ of pivotal edges for A_n must be at least k. As a result,

$$\mathbb{P}_p(R_1 + R_2 + \ldots + R_k \leq n - k \mid A_n) \leq \mathbb{P}_p(N(A_n) \geq k \mid A_n). \tag{3.2}$$

Therefore

$$\mathbb{E}_p[N(A_n) \mid A_n] = \sum_{k=1}^{\infty} \mathbb{P}_p(N(A_n) \geq k \mid A_n) \geq \sum_{k=1}^{\infty} \mathbb{P}(R_1 + R_2 + \ldots + R_k \leq n - k \mid A_n).$$

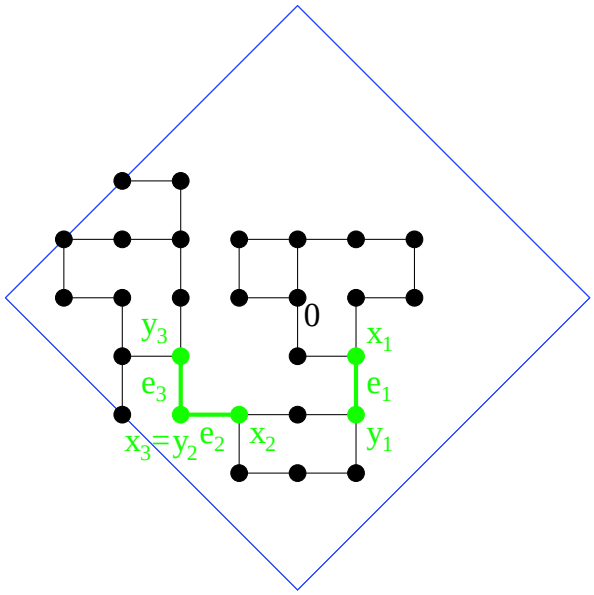

Figure 3.1: The three edges e_1, e_2 and e_3 are pivotal for A_5. For this example, $R_1 = 2$, $R_2 = 2$ and $R_3 = 0$.

So to compute $\mathbb{E}_p[N(A_n) \mid A_n]$, we need the distribution of the sum of the variables R_i. The first intermediate step will enable us to replace this sum by a sum of i.i.d. random variables, whose distribution is easier to compute.

Let $M = \max\{k \mid A_k \text{ occurs }\}$ be the radius of the largest ball whose surface is joined to the origin by an open path. The next lemma shows that, roughly speaking, the random variables $R_1, R_2, \ldots$ are jointly smaller in distribution than a sequence $M_1, M_2, \ldots$ of i.i.d. random variables distributed as M.

Lemma 3.1. *Let $k \in \mathbb{N}$, and let $r_1, r_2, \ldots r_k \in \mathbb{N}$, be such that $\sum_{i=1}^{k} r_i \leq n - k$. For $0 < p < 1$,*

$$\mathbb{P}_p(R_k \leq r_k, R_i = r_i \text{ for } 1 \leq i \leq k-1 \mid A_n) \geq \mathbb{P}_p(M \leq r_k)\mathbb{P}_p(R_i = r_i \text{ for } 1 \leq i \leq k-1 \mid A_n) \quad (3.3)$$

Proof: Let $k = 1$ and $r_1 \leq n-1$. If $R_1 > r_1$ then the first endvertex x_1 of the first pivotal edge e_1 lies either outside the ball $S(r_1 + 1)$, or on its surface $\partial S(r_1 + 1)$, as shown in the example of Figure 3.2. Hence $\{R_1 > r_1\} \subseteq A_{r_1+1}$. Since there are at least two edge disjoint paths between 0 and x_1, we have therefore that

$$\{R_1 > r_1\} \cap A_n \subseteq A_{r_1+1} \circ A_n$$

and since both A_{r_1+1} and A_n are increasing events, the BK inequality yields that

$$\mathbb{P}_p(\{R_1 > r_1\} \cap A_n) \leq \mathbb{P}_p(A_{r_1+1} \circ A_n) \leq \mathbb{P}_p(A_{r_1+1})\mathbb{P}_p(A_n).$$

Dividing by $\mathbb{P}_p(A_n)$, and noting that $\mathbb{P}_p(A_{r_1+1}) = \mathbb{P}_p(M \geq r_1 + 1)$, we get

$$\mathbb{P}_p(R_1 > r_1 \mid A_n) \leq \mathbb{P}_p(M \geq r_1 + 1),$$

and thus (3.3) for $k = 1$.

Suppose now that $k > 1$. For any edge $e =< u, v >$, let G_e be the set of vertices attainable from the origin along open paths not using e, together with all open edges between these vertices. Let

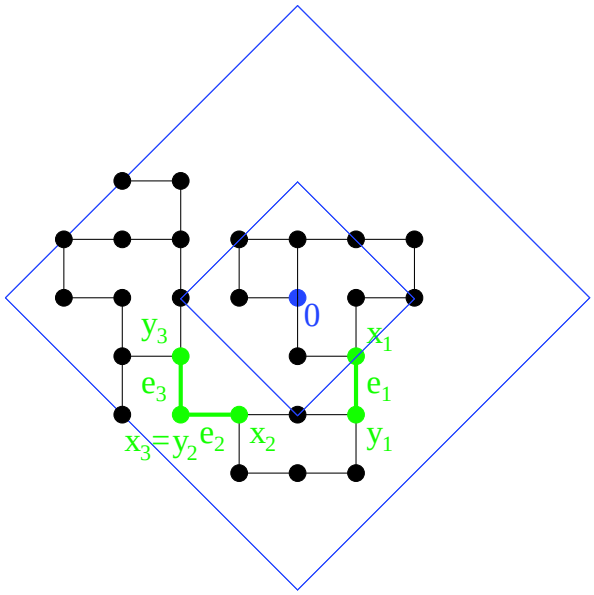

Figure 3.2: There are at least two edge-disjoint open paths connecting 0 to $\delta S(r_1+1)$ (here $r_1=1$).

B_e be the event that (i) e is open, (ii) $u \in G_e$ and $v \notin G_e$, (iii) G_e contains no vertex of $\partial S(n)$, and (iv) the pivotal edges for the event $\{0 \leftrightarrow \partial S(n)\}$ are $e_1 = \langle x_1, y_1\rangle, \ldots, e_{k-1} = \langle x_{k-1}, y_{k-1}\rangle = e$, where $\delta(y_{i-1}, x_i) = r_i$ for all $1 \le i \le k-1$. Let $B = \cup_e B_e$. For any $\omega \in A_n \cap B$, there is a unique edge $e = e(\omega)$ such that B_e, and thus B, occurs.

For $\omega \in B$, let $e = e(\omega)$ be an edge that verifies all conditions (i)-(iv) above. Let $G = G_e \cup \{e(\omega)\}$ be the set of vertices attainable from the origin along open paths not using this edge e, together with all open edges between these vertices, to which we the open pivotal edge $e = e(\omega)$ and its other endvertex y_{k-1}, we which also denote by $y(G)$. See Figure 3.3.

By conditioning on G, we obtain

$$\mathbb{P}_p(A_n \cap B) = \sum_{\Gamma} \mathbb{P}_p(B, G = \Gamma)\mathbb{P}_p(A_n \mid B, G = \Gamma),$$

where the sum is over all possible values Γ of G. Now, given the graph Γ, A_n occurs if and only if the vertex $y(\Gamma)$ is connected to $\partial S(n)$ by an open path which does not have any vertex other than $y(\Gamma)$ in common with Γ. Hence

$$\mathbb{P}_p(A_n \cap B) = \sum_{\Gamma} \mathbb{P}_p(B, G = \Gamma)\mathbb{P}_p(y(\Gamma) \leftrightarrow \partial S(n) \text{ off } \Gamma). \tag{3.4}$$

Similarly, given the graph Γ, if $\{R_k > r_k\}$, then the first endvertex x_k of the kth pivotal edge e_k lies either outside the ball $S(r_k+1, y(\Gamma))$ of radius r_k+1 centered on $y(\Gamma) = y_{k-1}$, or on its surface $\partial S(r_k+1, \Gamma)$. Hence $\{R_k > r_k\} \subseteq \{y(\Gamma) \leftrightarrow \partial S(r_k+1, \Gamma) \text{ off } \Gamma\}$. Since there are at least two edge disjoint open paths between $y(\Gamma) = y_{k-1}$, which moreover avoid any edge of Γ, we have therefore, conditionally to $G = \Gamma$, that

$$\{R_k > r_k\} \cap A_n \subseteq \{y(\Gamma) \leftrightarrow \partial S(r_k+1, y(\Gamma))\} \circ \{y(\Gamma) \leftrightarrow \partial S(n) \text{ off } \Gamma\}.$$

Now, by space invariance, $\mathbb{P}_p(y(\Gamma) \leftrightarrow \partial S(r_k+1, y(\Gamma))) = \mathbb{P}_p(A_{r_k+1})$. Therefore, by conditioning

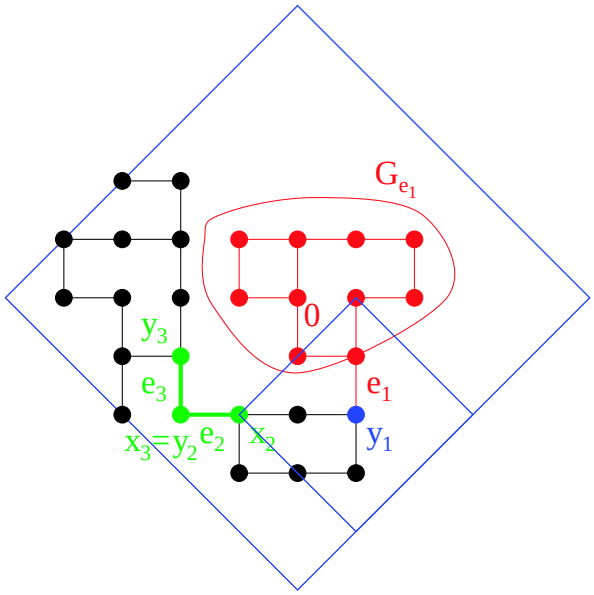

Figure 3.3: The circled set of vertices is the set G_{e_1}. At least two edge-disjoint open paths connect $y(G_{e_1}) = y_1$ to $\delta S(r_2 + 1, y_1)$ without passing by any vertex of G_{e_1} (here $k = 2$ and $r_k = r_2 = 1$).

on G and using the BK inequality, we get

$$\begin{aligned}
&\mathbb{P}_p(\{R_k > r_k\} \cap A_n \cap B) \\
&\quad = \sum_{\Gamma} \mathbb{P}_p(B, G = \Gamma)\mathbb{P}_p(\{R_k > r_k\} \cap A_n \mid B, G = \Gamma) \\
&\quad \leq \sum_{\Gamma} \mathbb{P}_p(B, G = \Gamma)\mathbb{P}_p(\{y(\Gamma) \leftrightarrow \partial S(r_k + 1, y(\Gamma))\} \circ \{y(\Gamma) \leftrightarrow \partial S(n) \text{ off } \Gamma\}) \\
&\quad \leq \sum_{\Gamma} \mathbb{P}_p(B, G = \Gamma)\mathbb{P}_p(y(\Gamma) \leftrightarrow \partial S(r_k + 1, y(\Gamma)))\mathbb{P}_p(y(\Gamma) \leftrightarrow \partial S(n) \text{ off } \Gamma) \\
&\quad = \sum_{\Gamma} \mathbb{P}_p(B, G = \Gamma)\mathbb{P}_p(A_{r_k+1})\mathbb{P}_p(y(\Gamma) \leftrightarrow \partial S(n) \text{ off } \Gamma) \\
&\quad = \mathbb{P}_p(A_{r_k+1})\mathbb{P}_p(A_n \cap B) \qquad (3.5)
\end{aligned}$$

where the latter equality follows from (3.4).

We finally divide both sides of (3.5) by $\mathbb{P}_p(A_n \cap B)$ to obtain

$$\mathbb{P}_p(R_k > r_k \mid A_n \cap B) \leq \mathbb{P}_p(A_{r_k+1})$$

and multiply the latter by $\mathbb{P}_p(B \mid A_n)$ to get

$$\mathbb{P}_p(\{R_k > r_k\} \cap B \mid A_n) \leq \mathbb{P}_p(A_{r_k+1})\mathbb{P}_p(B \mid A_n).$$

Now,

$$\begin{aligned}
\mathbb{P}_p(A_{r_k+1}) &= \mathbb{P}_p(M > r_k + 1) \\
\mathbb{P}_p(B \mid A_n) &= \mathbb{P}_p(R_i = r_i \text{ for } 1 \leq i \leq k-1 \mid A_n) \\
\mathbb{P}_p(\{R_k \leq r_k\} \cap B \mid A_n) &= \mathbb{P}_p(R_k \leq r_k, R_i = r_i \text{ for } 1 \leq i \leq k-1 \mid A_n),
\end{aligned}$$

from which we deduce (3.3).

■

From (3.3), we have that

$$\begin{aligned}
&\mathbb{P}_p(R_1 + R_2 + \ldots + R_k \leq n-k \mid A_n) \\
&= \sum_{i=0}^{n-k} \mathbb{P}_p(R_1 + R_2 + \ldots + R_{k-1} = i, R_k \leq n-k-i \mid A_n) \\
&\geq \sum_{i=0}^{n-k} \mathbb{P}_p(M \leq n-k-i)\mathbb{P}_p(R_1 + R_2 + \ldots + R_{k-1} = i \mid A_n) \\
&= \mathbb{P}_p(R_1 + R_2 + \ldots + R_{k-1} + M_k \leq n-k \mid A_n)
\end{aligned}$$

where M_k is a random variable independent from the state of all edges in $S(n)$, and distributed as M. Iterating $(k-1)$ more times this operation, we find that

$$\mathbb{P}_p(R_1 + R_2 + \ldots + R_k \leq n-k \mid A_n) \geq \mathbb{P}(M_1 + M_2 + \ldots + M_k \leq n-k) \tag{3.6}$$

where $M_1, \ldots, M_k$ is a sequence of i.i.d. random variables distributed as M. We can now find a lower bound on $\mathbb{E}_p[N(A_n) \mid A_n]$ thanks to (3.6); this is is our second intermediate result.

Lemma 3.2. *For $0 < p < 1$,*

$$\mathbb{E}_p[N(A_n) \mid A_n] \geq \frac{n}{\sum_{i=0}^{n} g_p(i)} - 1. \tag{3.7}$$

Proof: We come back to (3.2). Remember that if $R_1 + R_2 + \ldots + R_k \leq n-k$ and if A_n occurs, then the number $N(A_n)$ of pivotal edges for A_n must be at least k. Consequently,

$$\begin{aligned}
\mathbb{P}_p(N(A_n) \geq k \mid A_n) &\geq \mathbb{P}_p(R_1 + R_2 + \ldots + R_k \leq n-k \mid A_n) \\
&\geq \mathbb{P}(M_1 + M_2 + \ldots + M_k \leq n-k).
\end{aligned}$$

Now, as $\mathbb{P}(M_i \geq r) = \mathbb{P}_p(M \geq r) = g_p(r) \to \theta(p)$ for $r \to \infty$, we make a change of variable to avoid having $\mathbb{P}(M_i \geq r) > 0$ for $r \to \infty$ when $p_c < p$: let $M_i' = 1 + \min\{M_i, n\}$. Then

$$\mathbb{P}(M_1 + M_2 + \ldots + M_k \leq n-k) = \mathbb{P}(M_1' + M_2' + \ldots + M_k' \leq n),$$

and so we can continue with these truncated random variables. We have

$$\begin{aligned}
\mathbb{E}_p[N(A_n) \mid A_n] &= \sum_{k=1}^{\infty} \mathbb{P}_p(N(A_n) \geq k \mid A_n) \\
&\geq \sum_{k=1}^{\infty} \mathbb{P}(M_1' + M_2' + \ldots + M_k' \leq n) \\
&\geq \sum_{k=1}^{\infty} \mathbb{P}(K \geq k+1) \\
&= \mathbb{E}[K] - 1
\end{aligned}$$

where

$$K = \min\{k \mid M_1' + M_2' + \ldots + M_k' > n\}.$$

Since the M_i' are i.i.d. bounded random variables, Wald's equation (see Appendix) yields that

$$\mathbb{E}[M_1' + M_2' + \ldots + M_K'] = \mathbb{E}[M_i']\mathbb{E}[K]$$

As $M'_1 + M'_2 + \ldots + M'_K > n$ by definition of K, it follows that

$$\mathbb{E}[M'_1 + M'_2 + \ldots + M'_K] = \mathbb{E}[M'_i]\mathbb{E}[K] > n$$

whence

$$\begin{aligned}\mathbb{E}[K] &> \frac{n}{\mathbb{E}[M'_i]} = \frac{n}{1+\mathbb{E}[\min\{M_i, n\}]} = \frac{n}{1+\mathbb{E}[\min\{M, n\}]} \\ &= \frac{n}{1+\sum_{j=1}^{n}\mathbb{P}(M \geq j)} = \frac{n}{\sum_{j=0}^{n} g_p(j)}.\end{aligned}$$

■

Plugging (3.7) in (3.1), we find that

$$g_{p_1}(n) \leq g_{p_2}(n) \exp\left(-\int_{p_1}^{p_2} \frac{1}{p}\left(\frac{n}{\sum_{i=0}^{n} g_p(i)} - 1\right) dp\right).$$

This integral is difficult to compute as such, so we replace functions of p in the integrant as follows: $1/p \geq 1$ and $g_p(i) \leq g_{p_2}(i)$, and we obtain hat for any $n \in \mathbb{N}^*$

$$g_{p_1}(n) \leq g_{p_2}(n) \exp\left(-(p_2 - p_1)\left(\frac{n}{\sum_{i=0}^{n} g_{p_2}(i)} - 1\right)\right). \tag{3.8}$$

We still need one last intermediate result, namely that we need is that $\sum_{i=0}^{\infty} g_{p_2}(i)$ is finite.

Lemma 3.3. *For $0 < p < p_c$, there exists $\delta(p) < \infty$ such that*

$$g_p(n) \leq \delta(p)\sqrt{n}. \tag{3.9}$$

for $n \in \mathbb{N}^*$.

Proof: For any $n' \geq n$, (3.8) becomes

$$\begin{aligned}g_{p_1}(n') &\leq g_{p_2}(n') \exp\left((p_2 - p_1)\left(1 - \frac{n'}{\sum_{i=0}^{n'} g_{p_2}(i)}\right)\right) \\ &\leq g_{p_2}(n) \exp\left((p_2 - p_1)\left(1 - \frac{n'}{\sum_{i=0}^{n'} g_{p_2}(i)}\right)\right) \\ &\leq g_{p_2}(n) \exp\left(1 - \frac{n'(p_2 - p_1)}{\sum_{i=0}^{n'} g_{p_2}(i)}\right)\end{aligned}$$

because $g_{p_2}(n) = \mathbb{P}_{p_2}(A_n)$ is a decreasing function of n and $n' \geq n$. Now, we can decompose the summation

$$\begin{aligned}\frac{1}{n'}\sum_{i=0}^{n'} g_{p_2}(i) &= \frac{1}{n'}\left(\sum_{i=0}^{n-1} g_{p_2}(i) + \sum_{i=n}^{n'} g_{p_2}(i)\right) \\ &\leq \frac{1}{n'}\left(n g_{p_2}(0) + (n' - n + 1) g_{p_2}(n)\right) \\ &\leq \frac{1}{n'}\left(n + n' g_{p_2}(n)\right) \\ &\leq 3 g_{p_2}(n)\end{aligned}$$

by choosing $n' = n\lfloor 1/g_{p_2}(n)\rfloor$ (so that $n \leq 2n' g_{p_2}(n)$). Consequently,

$$g_{p_1}(n') \leq g_{p_2}(n) \exp\left(1 - \frac{(p_2 - p_1)}{3g_{p_2}(n)}\right). \tag{3.10}$$

Now we assume $p_2 < p_c$ and we choose p_1 by

$$p_1 = p_2 - 3g_{p_2}(n)(1 - \ln g_{p_2}(n)). \tag{3.11}$$

As $g_{p_2}(n)(1 - \ln g_{p_2}(n)) \to 0$ for $n \to \infty$, we pick n large enough to have the right hand side of (3.11) strictly positive. Plugging (3.11) in (3.10) finally yields

$$g_{p_1}(n') \leq \left(g_{p_2}(n)\right)^2. \tag{3.12}$$

Now we fix $0 < p < p_c$. We use the argument above to construct a subsequence $n_1, n_2, \ldots, n_i, \ldots$ along which $g_p(n_i)$ approaches 0 quickly. Pick q so that $p < q < p_c$, and construct two sequences. The first one is a sequence of probabilities p_i starting at $p_0 = q$ and the second one is a sequence of integers n_i starting at a value n_0 we will choose later, and defined as

$$n_{i+1} = n_i \gamma_i = n_i \lfloor 1/g_{p_i}(n_i)\rfloor \tag{3.13}$$

$$p_{i+1} = p_i - 3g_{p_i}(n_i)(1 - \ln g_{p_i}(n_i)) \tag{3.14}$$

with $\gamma_i = \lfloor 1/g_{p_i}(n_i)\rfloor$. Clearly, $n_{i+1} \geq n_i$ and that $p_{i+1} < p_i$. We still need to check that $p_i > 0$, and we will pick n_0 to ensure it. Because of the way we constructed the sequences (3.13) and (3.14), and because of the discussion leading to (3.12), we find that

$$g_{p_{j+1}}(n_{j+1}) \leq \left(g_{p_j}(n_j)\right)^2 \tag{3.15}$$

for $0 \leq j \leq i$. Now, any real sequence $\{x_j\}$ starting at a value $0 < x_0 < 1$ and defined by $x_{j+1} = x_j^2$ converges so quickly to 0 that the infinite sum $\sum_{j=0}^{\infty} 3x_j(1 - \ln x_j) < \infty$, and moreover converges to zero if $x_0 \to 0$. So we may pick x_0 sufficiently small that this iinite sum is smaller or equal to $q - p$, and next pick n_0 sufficiently large that $\mathbb{P}_q(A_n) = g_q(n_0) < x_0$. Using the fact that $3x(1 - \ln x)$ is an increasing function, we iterate (3.14) to find that

$$\begin{aligned}
p_{i+1} &= p_i - 3g_{p_i}(n_i)(1 - \ln g_{p_i}(n_i)) \\
&= p_0 - \sum_{j=0}^{i} 3g_{p_j}(n_j)(1 - \ln g_{p_j}(n_j)) \\
&\geq q - \sum_{j=0}^{\infty} 3g_{p_j}(n_j)(1 - \ln g_{p_j}(n_j)) \\
&\geq q - (q - p) = p.
\end{aligned}$$

Consequently, by suitably choosing n_0 large enough, we guarantee that $p_i > 0$ for all i. Even more, we get that $\lim_{i\to\infty} p_i \geq p$. Let us turn our attention now to the other sequence, and expand (3.13) to get

$$n_{i+1} = n_0 \gamma_0 \gamma_1 \ldots \gamma_i.$$

Next, expanding (3.15), we obtain that

$$\begin{aligned}
g_{p_i}^2(n_i) &= g_{p_i}(n_i) g_{p_i}(n_i) \leq g_{p_i}(n_i)\left(g_{p_{i-1}}(n_{i-1})\right)^2 \\
&\leq g_{p_i}(n_i) g_{p_{i-1}}(n_{i-1}) g_{p_{i-2}}^2(n_{i-2}) \leq \ldots \leq g_{p_i}(n_i) g_{p_{i-1}}(n_{i-1}) \ldots g_{p_1}(n_1) g_{p_0}^2(n_0) \\
&\leq \left(\gamma_i \gamma_{i-1} \ldots \gamma_1 \gamma_0\right)^{-1} g_{p_0}(n_0) \\
&= \delta^2 n_{i+1}^{-1}
\end{aligned}$$

with $\delta^2 = n_0 g_{p_0}(n_0)$.

Finally we fill in the gaps in the sequence $n_1, n_2, \ldots, n_i, \ldots$. Let $n > n_0$, and let i be such that $n_{i-1} \leq n < n_i$ (since $g_{p_i}(n_i) \to 0$ for $i \to \infty$, $n_{i-1} < n_i$ for large i). Then, since $p \leq p_{i-1}$,

$$g_p(n) \leq g_{p_{i-1}}(n_{i-1}) \leq \delta n_i^{-1/2} < \delta n^{-1/2}.$$

This is valid for $n < n_0$, adjusting δ we make a similar inequality valid for $n \geq 1$.

■

A consequence of this lemma, and more precisely of (3.9), is that there is some $\Delta(p) < \infty$ such that

$$\sum_{i=0}^{n} g_p(i) \leq \Delta(p) n^{1/2}$$

for $p < p_c$. Let $p_1 < p_c$, and pick $p_2 = p$ so that $p_1 < p_2 = p < p_c$. Inserting this sum in (3.8) yields that

$$\begin{aligned} g_{p_1}(n) &\leq g_{p_2}(n) \exp\left(-(p_2 - p_1)\left(\frac{n^{1/2}}{\Delta(p)} - 1\right)\right) \\ &\leq \exp\left(1 - \frac{p_2 - p_1}{\Delta(p)} n^{1/2}\right) \end{aligned}$$

and hence

$$\sum_{n=0}^{\infty} g_{p_1}(n) < \infty$$

for all $p_1 < p_c$. Consequently, since $p_2 < p_c$,

$$\mathbb{E}_{p_2}[M] = \sum_{n=0}^{\infty} g_{p_2}(n) < \infty.$$

Inserting this relation in (3.9), we get

$$\begin{aligned} g_{p_1}(n) &\leq g_{p_2}(n) \exp\left(-(p_2 - p_1)\left(\frac{n}{\mathbb{E}_{p_2}[M]} - 1\right)\right) \\ &\leq \exp\left(-\frac{p_2 - p_1}{\mathbb{E}_{p_2}[M]}\, n\right) = \exp(-\psi(p_1) n) \end{aligned}$$

with $\psi(p_1) = (p_2 - p_1)/\mathbb{E}_{p_2}[M] > 0$, which completes this long proof.

3.2 Correlation length and cluster size distribution

A number of easier results can be deduced from Theorem 3.1, which show that basically all metrics related to the size of the cluster containing the origin are exponentially decreasing with the size.

3.2.1 Connectivity function and correlation length

We begin with the *connectivity function* $\mathbb{P}_p(x \leftrightarrow y)$ which is defined as the probability that two vertices x and y are connected together by an open path. By translation invariance, we can take $y = 0$. Moreover, in the sake of simplicity and without loss of generality, we assume that x is positioned along the x-axis : $x = x_n$ where u_n is the d-dimensional vector $u_n = (n, 0, \ldots 0)$.

Theorem 3.2 (Exponential decay of connectivity function). *If $0 < p < p_c$, there exists $0 < \xi(p) < \infty$ and a constant $\kappa > 0$ independent of p, such that*

$$\kappa p n^{4(1-d)} \exp(-n/\xi(p)) \leq \mathbb{P}_p(0 \leftrightarrow u_n) \leq \exp(-n/\xi(p)). \tag{3.16}$$

This shows that

$$\mathbb{P}_p(0 \leftrightarrow u_n) \approx \exp(-n/\xi(p)).$$

The function $\xi(p)$ is called the *correlation length*, and one can show that $\xi(p) = 1/\psi(p)$, where $\psi(p)$ is the exponent in Theorem 3.1.

We prove only the upper bound, the proof for the lower bound is longer but builds essentially on a similar argument, based on the sub-additive limit theorem, which we recall here.

Lemma 3.4 (Sub-additive limit theorem). *Let $\{x_n, n \in \mathbb{N}^*\}$ be a sub-additive sequence of real non negative numbers, i.e. a sequence of real non negative numbers such that*

$$x_{m+n} \leq x_m + x_n \tag{3.17}$$

for all $m, n \in \mathbb{N}^*$, the limit

$$\overline{x} = \lim_{n\to\infty} x_n/n$$

exists and is finite. Moreover,

$$\overline{x} = \inf_{n\in\mathbb{N}^*} x_n/n$$

and thus $x_m \geq m\overline{x}$ for all $m \in \mathbb{N}^*$.

Proof: We prove only the upper inequality. The starting point is the observation (see Figure 3.4) that

$$\{0 \leftrightarrow u_{m+n}\} \supseteq \{0 \leftrightarrow u_m\} \cap \{u_m \leftrightarrow u_{m+n}\}$$

and thus, by first using the FKG inequality and next by translation invariance,

$$\mathbb{P}_p(0 \leftrightarrow u_{m+n}) \geq \mathbb{P}_p(0 \leftrightarrow u_m)\mathbb{P}_p(u_m \leftrightarrow u_{m+n}) = \mathbb{P}_p(0 \leftrightarrow u_m)\mathbb{P}_p(0 \leftrightarrow u_n).$$

Letting $x_n = -\ln \mathbb{P}_p(0 \leftrightarrow u_n)$, this inequality becomes (3.17), and therefore, by Lemma 3.4, the limit

$$\xi^{-1}(p) = \lim_{n\to\infty} \left(-\frac{\ln \mathbb{P}_p(0 \leftrightarrow u_n)}{n} \right)$$

exists. Moreover, $x_n \geq n/\xi(p)$ for all $n \in \mathbb{N}^*$, which yields the result.

■

3.2.2 Cluster size distribution

It follows from Theorem 3.1 that the distribution number $|C|$ of vertices contained in the open cluster at the origin has an exponentially decreasing tail. A more accurate bound is as follows.

Theorem 3.3 (Exponential decay of the cluster size distribution). *If $0 < p < p_c$, there exists $\lambda(p) > 0$ such that*

$$\mathbb{P}_p(|C| \geq n) \leq \exp(-n\lambda(p)) \tag{3.18}$$

and there exists $0 < \zeta(p) < \infty$ such that

$$\mathbb{P}_p(|C| = n) \leq \frac{(1-p)^2}{p} n \exp(-n\zeta(p)). \tag{3.19}$$

for $n \in \mathbb{N}^*$.

One can moreover show that

$$\mathbb{P}_p(|C| = n) \approx \exp(-n\zeta(p)).$$

The theorem is proven in [18].

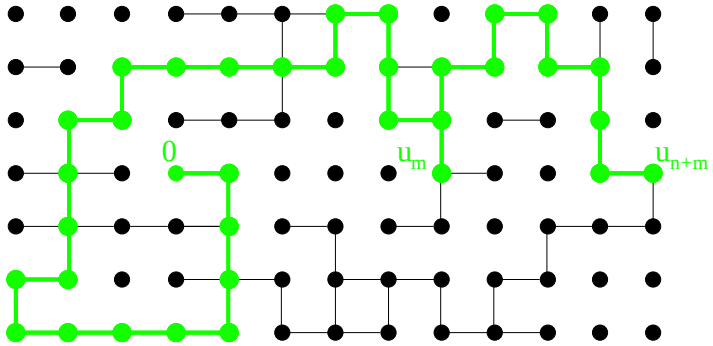

Figure 3.4: The event $\{0 \leftrightarrow u_{m+n}\}$ is more likely than the joint occurrence of events $\{0 \leftrightarrow u_m\}$ and $\{u_m \leftrightarrow u_{m+n}\}$.

Appendix: Wald's equation

Let $\{X_n, n \in \mathbb{N}^*\}$ be a sequence of i.i.d. variables with finite mean. Then N is a *stopping time* for this sequence if and only if for any $n \in \mathbb{N}^*$, the event $\{N = n\}$ is independent of X_i with $i \geq n+1$.

Theorem 3.4 (Wald's equation). *If $\{X_n, n \in \mathbb{N}^*\}$ is a sequence of i.i.d. non negative variables with finite mean $\mathbb{E}[X]$, and if N is a stopping time for this sequence, with $\mathbb{E}[N] < \infty$, then*

$$\mathbb{E}[X_1 + X_2 + \ldots + X_N] = \mathbb{E}[X]\mathbb{E}[N].$$

The assumption that the variables X_n are non negative is not needed, but to avoid using martingales, we give a proof, which follows instead [35], and which is always valid when variables X_n are non negative.

Proof: Since

$$\sum_{n=1}^{N} X_n = \sum_{n=1}^{\infty} X_n 1_{\{N \geq n\}}$$

taking expectations, we have that

$$\mathbb{E}\left[\sum_{n=1}^{N} X_n\right] = \mathbb{E}\left[\sum_{n=1}^{\infty} X_n 1_{\{N \geq n\}}\right] = \sum_{n=1}^{\infty} \mathbb{E}\left[X_n 1_{\{N \geq n\}}\right] \tag{3.20}$$

where the last interchange between expectation and summation is valid given that all X_n are non negative. Now, since N is a stopping time for the sequence $\{X_n, n \in \mathbb{N}^*\}$, $1_{\{N \geq n\}} = 1$ if and only if we have not stopped after having successively observed $X_1, X_2, \ldots, X_{n-1}$. The random variable $1_{\{N \geq n\}}$ is thus determined by $X_1, \ldots, X_{n-1}$ and is independent of X_n. As a result, (3.20) becomes

$$\begin{aligned}
\mathbb{E}\left[\sum_{n=1}^{N} X_n\right] &= \sum_{n=1}^{\infty} \mathbb{E}\left[X_n\right] \mathbb{E}\left[1_{\{N \geq n\}}\right] \\
&= \mathbb{E}\left[X\right] \sum_{n=1}^{\infty} \mathbb{P}(N \geq n) \\
&= \mathbb{E}\left[X\right] \mathbb{E}[N].
\end{aligned}$$

■

4

Supercritical phase

In this chapter we study the situation in the supercritical phase, when $p > p_c$ and $d \geq 2$. In this case, we know that there is almost surely an open cluster of infinite size. But how many are there ? We will first prove that there is exactly one such cluster. The next question will be to evaluate the size of the other, finite clusters. We will see in that they decrease sub-exponentially fast. We will prove the result only when $d = 2$, although it holds for $d \geq 3$ as well.

4.1 Uniqueness of the infinite open cluster

We follow the approach of Burton and Keane (1989), as exposed in [18], to prove that the infinite open cluster is unique in the supercritical phase.

Theorem 4.1 (Uniqueness of the infinite open cluster). *If $p > p_c$, then*

$$\mathbb{P}_p(\textit{there exists exactly one infinite open cluster}) = 1.$$

Proof: Let Y be the number of infinite open clusters. Because the sample space $\Omega = \prod_{e\in\mathbb{E}}\{0,1\}^e$ is a product space with a space invariant product measure $\mathbb{P}_p$, Y is a translation-invariant function on Ω. A property of translation-invariant functions under ergodic measures is to be almost surely constant. Consequently, there exists some $k \in \mathbb{N}\cup\{\infty\}$ such that $\mathbb{P}_p(Y = k) = 1$.

Since $p > p_c$, $k \neq 0$. We will prove by contradiction that (i) $k \notin [2,\infty[$ and (ii) $k \neq \infty$, which implies therefore that $k = 1$.

(i) Suppose first that $2 \leq k < \infty$. As in the previous chapter, denote by $S(n) = \{x \in \mathbb{Z}^d \mid \delta(0,x) = |x| \leq n\}$ the diamond of radius n (i.e., the ball of radius n with the Manhattan distance). Let $Y(0)$ be the number of infinite open clusters when all edges of $S(n)$ are closed. As the probability that all edges of $S(n)$ are closed is strictly positive,

$$\mathbb{P}_p(Y(0) = k) = \frac{\mathbb{P}_p(\{Y = k\}\cap\{\text{all edges of } S(n) \text{ are closed}\})}{P_p(\text{all edges of } S(n) \text{ are closed})} = 1.$$

Similarly, if $Y(1)$ denotes the number of infinite open clusters when all edges of $S(n)$ are open, $\mathbb{P}_p(Y(1) = k) = 1$, and therefore $\mathbb{P}_p(Y(0) = Y(1)) = 1$. We always have that $Y(0) \geq Y(1)$, but since there are only a finite number of open infinite clusters, we have $Y(0) = Y(1)$ if and and only if $S(n)$ intersects exactly one such cluster. So, if $M_{S(n)}$ is the number of infinite open clusters intersecting $S(n)$, $\mathbb{P}_p(M_{S(n)} \geq 2) = 0$ for all $n \in \mathbb{N}$. Letting $n \to \infty$, we have that the diamond $S(n)$ becomes the entire lattice $\mathbb{L}^d$ and therefore that

$$0 = \lim_{n\to\infty} \mathbb{P}_p(M_{S(n)} \geq 2) = \mathbb{P}_p(Y \geq 2), \tag{4.1}$$

a contradiction with $\mathbb{P}(Y = k) = 1$ for some $2 \leq k < \infty$.

(ii) Suppose next that $k = \infty$. We use a geometric argument to get a contradiction, which is based on the following object. We call a vertex x a *trifurcation* (see Figure 4.1) if

1. x belongs to an infinite open cluster;
2. there exist exactly three open edges incident to x; and
3. the deletion of x and of the three open edges incident to x splits the infinite open cluster containing x in exactly three disjoint infinite clusters and no finite cluster.

Because of the space invariance of $\mathbb{L}^d$, the probability that a vertex x is a trifurcation is independent of x, and therefore

$$\mathbb{P}_p(x \text{ is a trifurcation}) = \mathbb{P}_p(0 \text{ is a trifurcation}). \tag{4.2}$$

Let us show that this probability is non zero. Let $M_{S(n)}(0)$ be the number of infinite open clusters intersecting $S(n)$ when all edges of $S(n)$ are closed. Clearly, $M_{S(n)}(0) \geq M_{S(n)}$. Therefore

$$\mathbb{P}_p(M_{S(n)}(0) \geq 3) \geq \mathbb{P}_p(M_{S(n)} \geq 3) \to \mathbb{P}_p(Y \geq 3) = 1$$

as $n \to \infty$. Consequently, there is $n \in \mathbb{N}$ such that $\mathbb{P}_p(M_{S(n)}(0) \geq 3) \geq 1/2$, fix n to this value from now on until we have shown the probability of having a trifurcation at the origin is non zero.

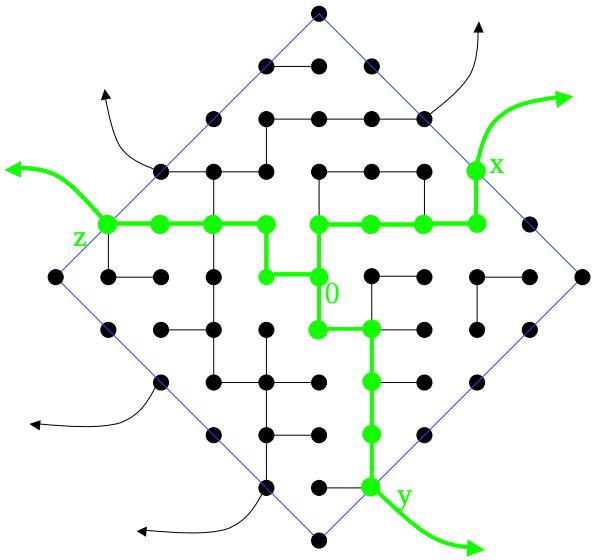

Figure 4.1: A sufficient condition for 0 to be a trifurcation if the three paths from x, y and z are open, and all all other edges in $S(n)$ are closed, and x, y and z belong to three distinct infinite open clusters. The arrows outside $\partial S(n)$ represent connectivity to distinct infinite clusters.

If $M_{S(n)}(0) \geq 3$, then there exists three vertices $x, y, z \in \partial S(n)$ lying in three distinct infinite open clusters. Moreover, there are three paths inside $S(n)$ joining the origin to respectively x, y, z, such that the origin is the unique vertex common to any two of them, and each touches exactly one vertex on $\partial S(n)$. For a configuration of edges $\omega \in \{M_{S(n)}(0) \geq 3\}$, we pick $x = x(\omega)$, $y = y(\omega)$ and $z = z(\omega)$ and the three paths as just described. Let $J_{x,y,z}$ be the event that all edges in these three paths are open and that all other edges in $S(n)$ are closed. Then

$$\mathbb{P}_p(J_{x,y,z} \mid M_B(0) \geq 3) \geq (\min\{p, 1-p\})^{R(n)}$$

where $R(n)$ is the total number of edges in $S(n)$. Now, if $M_{S(n)}(0) \geq 3$ and if $J_{x,y,z}$ occurs, then x is a trifurcation. Therefore

$$\begin{aligned} \mathbb{P}_p(0 \text{ is a trifurcation}) &\geq \mathbb{P}_p(J_{x,y,z} \mid M_B(0) \geq 3)\mathbb{P}_p(M_{S(n)}(0) \geq 3) \\ &\geq (\min\{p, 1-p\})^{R(n)}/2 > 0. \end{aligned}$$

Because of (4.2), we have therefore that $\mathbb{P}_p(x \text{ is a trifurcation}) > 0$ for all vertices $x \in \mathbb{Z}^d$. Let $T(m)$ denote the number of trifurcations in $S(m)$. As

$$\mathbb{E}_p[T(m)] = |S(m)|\mathbb{P}_p(0 \text{ is a trifurcation}),$$

it implies that $T(m)$ grows in the manner of $|S(m)|$ as $m \to \infty$. The contradiction is obtained by the following rough geometric argument (a more rigorous proof uses partitions, see [18]). Pick a trifurcation in $S(m)$, say t_1, and take a vertex $x_1 \in \partial S(m)$ that is connected to t_1 by an open path in $S(m)$. Pick a second trifurcation $t_2 \in S(m)$. By definition of a trifurcation, there must be a vertex $x_2 \in \partial S(m)$, distinct from x_1, such that $t_2 \leftrightarrow x_2$ in $S(m)$ (See Figure 4.2). Repeat this operation, at each stage picking a new trifurcation t_i and a new vertex $x_i \in \partial S(m)$, with $t_i \leftrightarrow x_i$ in $S(m)$. There are $T(m)$ trifurcations in $S(m)$, so we end up with $T(m)$ distinct vertices $x_i \in \partial S(m)$, which implies that $|\partial S(m)| \geq T(m)$. But $T(m)$ grows in the manner of $|S(m)|$ for large m, which would mean that $|\partial S(m)|$ would grow in the manner of $|S(m)|$ for large m as well.

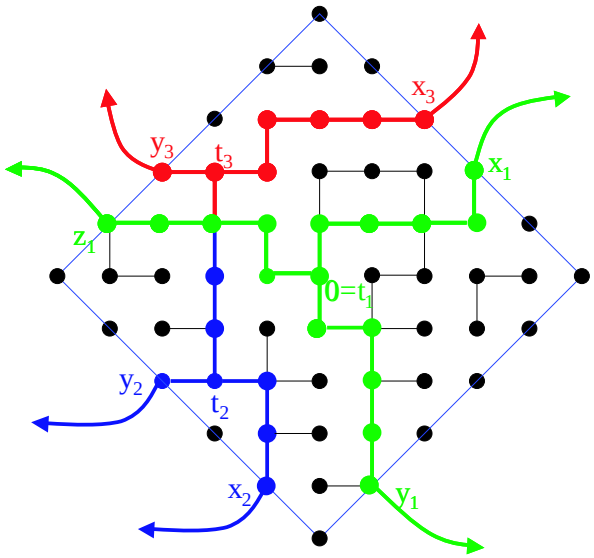

Figure 4.2: Finding trifurcations in $S(n)$.

We have reached a contradiction, because $|S(m)|$ grows in the manner of m^d while $|\partial S(m)|$ grows in the manner of m^{d-1}.

■

4.2 Finite cluster size distribution

We are now interested in the size of the finite clusters. which show that basically all metrics related to the size of the cluster containing the origin are exponentially decreasing with the size. We only consider the 2-dim case, where the proof is easier because of the use of duality. The theorem is however valid for

Theorem 4.2 (Sub-exponential decay of the finite cluster size distribution). *If $p_c < p < 1$, there exists $\eta(p) > 0$ such that*

$$\mathbb{P}_p(|C| = n) \leq \exp(-n^{(d-1)/d}\eta(p)) \tag{4.3}$$

for $n \in \mathbb{N}^*$.

We can also find a lower bound of the same form: there exists some $\gamma(p) < \infty$ such that

$$\mathbb{P}_p(|C| = n) \geq \exp(-n^{(d-1)/d}\gamma(p)).$$

Proof: We only prove a slightly weaker bound

$$\mathbb{P}_p(|C| = n) \leq n \exp(-\sqrt{n}\eta(p)),$$

and only for $d = 2$ and for $2/3 < p < 1$. Once we will have computed the exact value of p_c in the next chapter, the proof is directly extended for $p_c < p < 1$.

Suppose that that the origin belongs to a finite cluster of size n. Then there exists a closed circuit in the dual lattice $\mathbb{L}_d^2$, having the origin in its interior. Clearly, this circuit has less than n vertices.

Moreover, it can be shown using topological arguments (see Kesten 1982) that there is some value $\delta > 0$ such that this closed circuit contains at least $\delta\sqrt{n}$ vertices. For the same reason as in part (ii) of the proof of Theorem 2.1, it must pass through a vertex of the form $(i + 1/2, 1/2)$ for some $0 \leq i \leq n-1$, and therefore one of these n vertices must lie in a closed cluster of $\mathbb{L}_d^2$ of size at least $\delta\sqrt{n}$. Let us called 0_d this vertex, and C_d the closed cluster to which it belongs.

Now, each edge of $\mathbb{L}_d^2$ is closed with probability $(1-p)$, and $1-p < 1/3 \leq p_c$ because of Theorem 2.1. In other words, the process of closed edges of $\mathbb{L}_d^2$ is subcritical. Theorem 3.3 then yields that there exists $\lambda(p) > 0$ such that

$$\mathbb{P}_p(|C_d| \geq \delta\sqrt{n}) \leq \exp(-\lambda(p)\delta\sqrt{n}).$$

Since

$$\mathbb{P}_p\left((i + 1/2, 1/2) \text{ lies in a closed cluster of } \mathbb{L}_d^2 \text{ of size at least } \delta\sqrt{n}\right) = \mathbb{P}_p(|C_d| \geq \delta\sqrt{n}),$$

we have thus that

$$\begin{aligned}
\mathbb{P}_p(|C| = n) &\leq \sum_{i=0}^{n-1} \mathbb{P}_p\left((i + 1/2, 1/2) \text{ lies in a closed cluster of } \mathbb{L}_d^2 \text{ of size at least } \delta\sqrt{n}\right) \\
&= n\mathbb{P}_p(|C_d| \geq \delta\sqrt{n}) \\
&\leq n\exp(-\lambda(p)\delta\sqrt{n}).
\end{aligned}$$

Setting $\eta(p) = \lambda(p)\delta$ finishes the proof.

■

5

Near the critical threshold

5.1 Introduction

After having studied the key properties of the metrics associated to the cluster size distribution in the sub-critical phase, we now move to the critical point p_c. The first part of this chapter is devoted to the computation its value for the 2-dimensional case, where we will again make use of the self-duality of $\mathbb{L}^2$. We will next move to the behavior of the metrics of interest (percolation probability $\theta(p)$, mean cluster size $\chi(p)$, correlation length $\xi(p)$, finite cluster size distribution $\mathbb{P}_p(|C| = n)$) for values of p close to p_c, in dimensions $d \geq 2$. Very few rigorous results exist for the behavior of these quantities close to p_c, and we need to make conjectures, supported by techniques from statistical physics, such as scaling theory, which will be introduced in the second part of this chapter.

5.2 Critical threshold for bond percolation on the 2-dim. lattice

The previous chapters have equipped us with the necessary tools to eventually compute the value of p_c, which we will prove to be equal to $1/2$.

We begin by proving that in 2 dimensions, the percolation probability is zero when $p = 1/2$. An immediate consequence is that the critical percolation threshold $p_c \geq 1/2$.

The absence of infinite open cluster at the percolation threshold is also conjectured to hold in higher dimensions.

Lemma 5.1 (Absence of infinite open cluster for $p = 1/2$). *If $d = 2$, then $\theta(1/2) = 0$.*

Proof: We proceed by contradiction, and follow Zhang (1988) as exposed in [18]. Suppose that $\theta(1/2) > 0$. Consider the square $B(n) = [-n, n] \times [-n, n]$, and let $A^l(n)$ (respectively, $A^r(n)$, $A^t(n)$, $A^b(n)$) be the event that some vertex on the left (respectively, right, top, bottom) side of $B(n)$ belongs to an infinite open path of $\mathbb{L}^2$ that uses no other vertex of $B(n)$. Clearly, these are four increasing events that have equal probability (by symmetry) and whose union is the event that some vertex on $B(n)$ belongs to an infinite cluster. Since we assume that $\theta(1/2) > 0$, the Kolmogorov zero-one law implies that there is almost surely an infinite cluster, and therefore as $n \to \infty$,

$$\mathbb{P}_{1/2}\left(A^l(n) \cup A^r(n) \cup A^t(n) \cup A^b(n)\right) \to 1. \tag{5.1}$$

Now, using the "square root trick" (see homework 1), which states that if B_i, $1 \leq i \leq n$, are increasing events having the same probability,

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

we get

$$\mathbb{P}_{1/2}(A^l(n)) \geq 1 - \left(1 - \mathbb{P}_{1/2}\left(A^l(n) \cup A^r(n) \cup A^t(n) \cup A^b(n)\right)\right)^{1/4}.$$

It follows from (5.1) that, with $u = l, r, t, b$,

$$\mathbb{P}_{1/2}(A^u(n)) \to 1$$

as $n \to \infty$. Therefore there is n_0 large enough such that for $u = l, r, t, b$

$$\mathbb{P}_{1/2}(A^u(n_0)) > 7/8. \tag{5.2}$$

Let us next consider the dual box $B_d(n)$ defined as

$$B_d(n) = \{(i + 1/2, j + 1/2) \mid (i, j) \in B(n)\},$$

and let $A^l_d(n)$ (respectively, $A^r_d(n)$, $A^t_d(n)$, $A^b_d(n)$) be the event that some vertex on the left (respectively, right, top, bottom) side of $B_d(n)$ belongs to an infinite closed path of $\mathbb{L}^2_d$ that uses no other vertex of $B_d(n)$. Each edge of $\mathbb{L}^2_d$ is closed with a probability $1/2$, which is the same as the open edge probability in $\mathbb{L}^2$. Therefore $\mathbb{P}_{1/2}(A^u(n)) = \mathbb{P}_{1/2}(A^u_d(n))$ for $u = l, r, t, b$ and all $n \in \mathbb{N}^*$. In particular, for $n = n_0$,

$$\mathbb{P}_{1/2}(A^u_d(n_0)) > 7/8 \tag{5.3}$$

for $u = l, r, t, b$ because of (5.2).

We now consider the event $A = A^l(n_0) \cap A^r(n_0) \cap A^t_d(n_0) \cap A^b_d(n_0)$, that there exist infinite open paths of $\mathbb{L}^2$ connecting to some vertex on the left and right sides of $B(n)$, without using any other vertex of $B(n)$, and that there exists infinite closed paths connecting to some vertex on the top

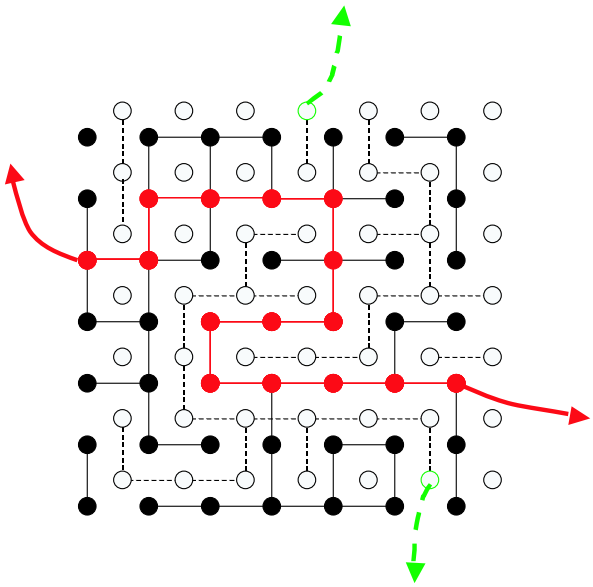

Figure 5.1: Infinite open paths of $\mathbb{L}^2 \setminus B(n_0)$ connecting to some vertex on the left and right sides of $B(n_0)$ and infinite closed paths of $\mathbb{L}_d^2 \setminus B_d(n_0)$ connecting to some vertex on the top and bottom sides of $B_d(n)$.

and bottom sides of $B_d(n)$, without using any other vertex of $B_d(n)$, as shown in Figure 5.1. Now, using the union bound,

$$\begin{aligned}
\mathbb{P}_{1/2}(A) &= 1 - \mathbb{P}_{1/2}\left(\overline{A}^l(n_0) \cup \overline{A}^r(n_0) \cup \overline{A}_d^t(n_0) \cup \overline{A}_d^b(n_0)\right) \\
&\geq 1 - \left(\mathbb{P}_{1/2}(\overline{A}^l(n_0)) + \mathbb{P}_{1/2}(\overline{A}^r(n_0)) + \mathbb{P}_{1/2}\overline{A}_d^t(n_0)) + \mathbb{P}_{1/2}(\overline{A}_d^b(n_0))\right) \\
&> 1/2
\end{aligned}$$

because of (5.2) and (5.3). If A occurs, then there must be two infinite open clusters in $\mathbb{L}^2 \setminus B(n_0)$, one containing the infinite open path connected to the left side of $B(n)$ and the other one containing the infinite open path connected to the right side of $B(n)$. Moreover, these two infinite open clusters must be disjoint, because they are separated by two infinite closed paths in $\mathbb{L}_d^2 \setminus B_d(n_0)$ connecting to some vertex on the top and bottom sides of $B_d(n)$. If there was an open path connecting the two infinite clusters of $\mathbb{L}^2 \setminus B(n_0)$ path, one of its (open) edges would cross a closed edge in $\mathbb{L}_d^2 \setminus B_d(n_0)$, which is impossible, as shown in Figure 5.1.

The same reasoning implies that there must be two disjoint infinite closed clusters in $\mathbb{L}_d^2 \setminus B_d(n_0)$, one containing the infinite closed path connected to the top side of $B_d(n)$ and the other one containing the infinite closed path connected to the bottom side of $B_d(n)$, and separated by the two infinite open paths of $\mathbb{L}^2 \setminus B(n_0)$. Now, as $\theta(1/2) > 0$, Theorem 4.1 yields that the infinite lattice $\mathbb{L}^2$ contains (almost surely) one and only one infinite open cluster. Therefore, there must be a left-right open crossing within $B(n)$, which forms a barrier to any top-bottom closed crossing of $B_d(n)$. As a result, there must be (almost surely) at least two disjoint infinite closed clusters in $\mathbb{L}_d^2$. But since $p = 1 - p = 1/2$, the probability that there are two infinite closed clusters in $\mathbb{L}_d^2$ is the same as the probability that there are two infinite open clusters in $\mathbb{L}^2$, which is zero. We have thus reached a contradiction, which means that $\mathbb{P}_{1/2}(A)$ cannot be nonzero. The initial assumption $\theta(1/2) > 0$ cannot be valid, which establishes the result.

■

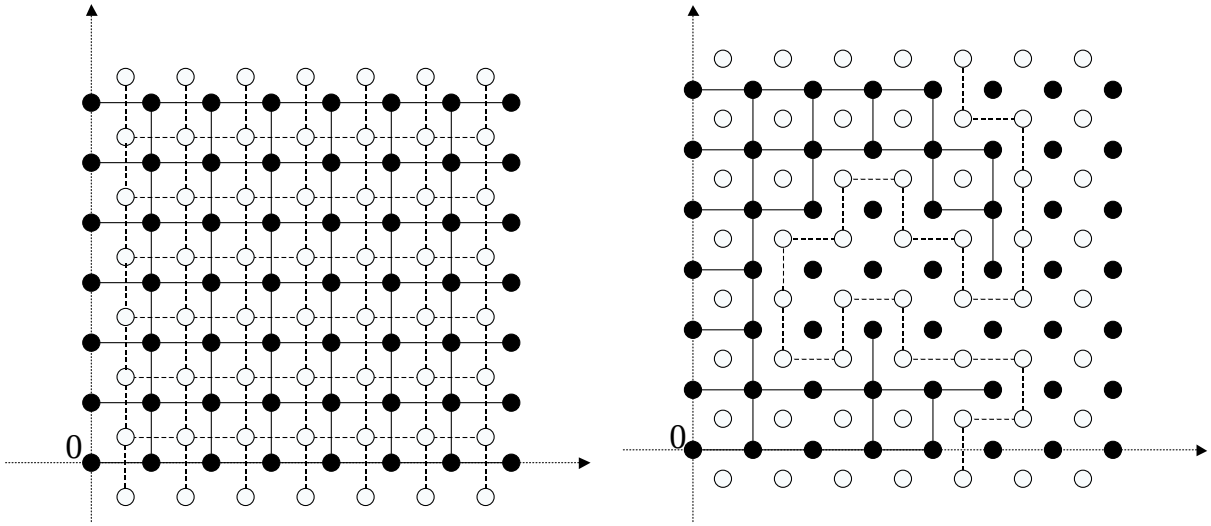

Figure 5.2: The box $R(n)$ and its dual $R_d(n)$ for $n = 6$ (left) and an illustration of the fact that there is no left-right open crossing of $R(n)$ is and only if there is a top-bottom closed crossing of $R_d(n)$ (right).

The previous theorem implies that $p_c \geq 1/2$. The following lemma is the main step in showing the converse, namely $p_c \leq 1/2$.

Lemma 5.2 (Crossing of a square for $p = 1/2$). *Let $LR(n)$ be the even that there is a left-right open crossing of the rectangle $R(n) = [0, n+1] \times [0, n]$ (that is, an open path connecting some vertex on the left side of $R(n)$ to some vertex on the right side of $R(n)$. Then $\mathbb{P}_{1/2}(LR(n)) = 1/2$ for all $n \in \mathbb{N}^{star}$.*

Proof: The rectangle $R(n)$ is the subgraph of $\mathbb{L}^2$ having vertex set $[0, n+1] \times [0, n]$ and edge set comprising all edges of $\mathbb{L}^2$ joining pairs of vertices in $S(n)$, except those joining pairs (i, j), (k, l) with either $i = k = 0$ or $i = k = n + 1$. Let $R_d(n)$ be the subgraph of $\mathbb{L}^2_d$ having vertex set $\{(i + 1/2, j + 1/2) \mid 0 \leq i \leq n, 1 \leq j \leq n\}$ and edge set all edges of $\mathbb{L}^2_d$ joining pairs of vertices in $R_d(n)$, except those joining pairs (i, j), (k, l) with either $i = k = -1/2$ or $i = k = n+1/2$. The two subgraphs can be obtained from each other by a 90 degrees rotation, which relocates the vertex labeled $(0, 0)$ at the point $(n + 1/2, -1/2)$, see Figure 5.2 (left).

Let us consider the two following events: $LR(n)$ is the event that there exists an open path of $R(n)$ joining a vertex on the left side of $R(n)$ to a vertex on its right side, and $TB_d(n)$ is the event that there exists a closed path of $R_d(n)$ joining a vertex on the top side of $R_d(n)$ to a vertex on its bottom side.

If $LR(n) \cap TB_d(n) \neq \emptyset$, there is a left-right open path in $R(n)$ crossing a top-bottom closed path in $S_d(n)$. But then, at the crossing of these two paths, there would be an open edge of $\mathbb{L}^2$ crossed by a closed edge of $\mathbb{L}^2_d$, which is impossible, see Figure 5.2 (right). Hence $LR(n) \cap TBd(n) = \emptyset$. On the other hand, either $LR(n)$ or $TBd(n)$ must occur. Let D be the set of vertices that are reachable from the left side of $R(n)$ by an open path. Suppose that $LR(n)$ does not occur. Then there exists a top-bottom closed path of $\mathbb{L}^2_d$ crossing only edges of $R(n)$ contained in the edge boundary of D, and so $TBd(n)$ occurs. Consequently $LR(n)$ and $TBd(n)$ form a partition of the sample space Ω, and

$$\mathbb{P}_p(LR(n)) + \mathbb{P}_p(TBd(n)) = 1. \tag{5.4}$$

Now, since $R(n)$ and $R_d(n)$ are isomorphic (they can be obtained from each other by a 90 degrees

rotation, which relocates the vertex labeled $(0,0)$ at the point $(n+1/2,-1/2))$, flipping the polarity of each edge of $\mathbb{L}_d^2$ yields that $\mathbb{P}_p(TBd(n)) = \mathbb{P}_{1-p}(LR(n))$. Plugging this equality in (5.4), the latter becomes

$$\mathbb{P}_p(LR(n)) + \mathbb{P}_{1-p}(LR(n)) = 1.$$

Taking $p = 1/2$ in this equation proves the lemma.

■

We now deduce directly one of the most famous theorems of percolation theory.

Theorem 5.1 ($p_c = 1/2$). *The percolation threshold in $\mathbb{L}^2$ is $p_c = 1/2$.*

Proof: We know from Lemma 5.1 that $p_c \geq 1/2$. Suppose that $p_c > 1/2$. Then the value $p = 1/2$ belongs to the supercritical phase, and we know from Theorem 3.1 that there exists $\psi(1/2) > 0$ such that for all $n \in \mathbb{N}^*$

$$\mathbb{P}_{1/2}(0 \leftrightarrow \partial^r R(n)) \leq \mathbb{P}_{1/2}(0 \leftrightarrow \partial S(n)) < \exp(-n\psi(1/2)),$$

where $\{0 \leftrightarrow \partial^r R(n)\}$ is the event that the origin is connected by an open path to a vertex lying on the right side of $R(n)$, defined as $\partial^r R(n) = \{(n+1,k) \in \mathbb{Z}^2 \mid 0 \leq k \leq n\}$, and where $\{0 \leftrightarrow \partial S(n)\}$ is the event that the origin is connected by an open path to a vertex lying on the perimeter of the ball of radius n centered in 0. Consequently, since $LR(n)$ is the event that there exists an open path of $R(n)$ joining a vertex on the left side of $R(n)$ to a vertex on its right side,

$$\begin{aligned} \mathbb{P}_{1/2}\left(LR(n)\right) &\leq \sum_{k=0}^{n} \mathbb{P}_{1/2}((0,k) \leftrightarrow \partial^r R(n)) \\ &\leq (n+1)\mathbb{P}_{1/2}(0 \leftrightarrow \partial^r R(n)) \\ &< (n+1)\exp(-n\psi(1/2)), \end{aligned}$$

which yields that $P_{1/2}\left(LR(n)\right) \to 0$ as $n \to \infty$, and therefore contradicts Lemma 5.2. Consequently $p_c \leq 1/2$, which completes the proof.

■

5.3 Near the critical threshold

5.3.1 Power laws in the 2-dim lattice

We know from Theorems 3.1 and 3.3 that the distributions of the radius and size of the cluster at the origin C decrease exponentially fast when $p < 1/2$. What happens when $p = 1/2$? Lemma 5.1 indicates that the cluster C is almost surely finite at the critical threshold, like in the sub-critical phase. The following theorem shows however that the distributions of the radius and size change radically of nature, and are follow no longer an exponential law, but a power law. A consequence is that the mean cluster size $\chi(1/2) = \infty$, contrary to the subcritical case.

Theorem 5.2 (Power law inequalities at the critical threshold). *In $\mathbb{L}^2$, for all $n \in \mathbb{N}^*$,*

$$\mathbb{P}_{1/2}(0 \leftrightarrow \partial B(n)) \geq \frac{1}{2\sqrt{n}} \tag{5.5}$$

$$\mathbb{P}_{1/2}(|C| \geq n) \geq \frac{1}{2\sqrt{n}}. \tag{5.6}$$

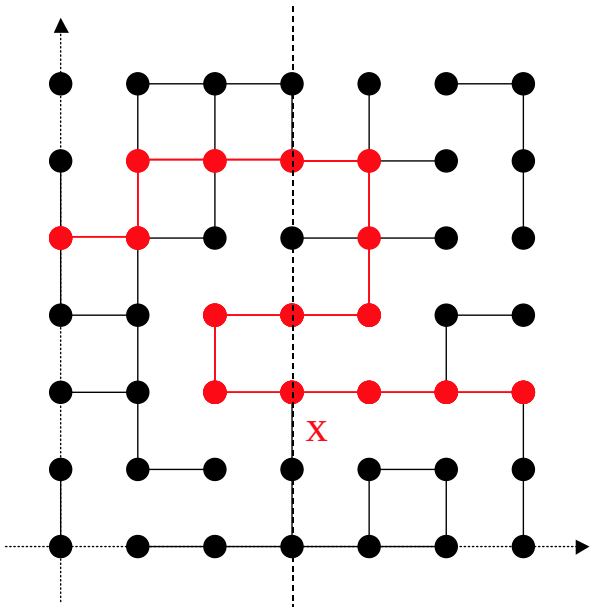

Figure 5.3: A left right open path crossing the box $R(2n-1)$ must hit the center vertical line at some vertex x, which is therefore joined by two disjoint paths to respectively the left and right sides of $R(2n-1)$.

Proof: Since any open path connecting the origin to the perimeter of $B(n)$ contains at least n vertices, $\mathbb{P}_{1/2}(|C| \geq n) \geq \mathbb{P}_{1/2}(0 \leftrightarrow \partial B(n))$, and so we only need to prove (5.5).

As before, let $LR(2n-1)$ be the event that is an open path in the rectangle $R(2n-1) = [0, 2n] \times [0, 2n-1]$ connecting some vertex on its left side to some vertex on its right side. This path must cross the center line $\{(n,k) \in \mathbb{Z}^2 \mid 0 \leq k \leq 2n-1\}$ in at least one vertex, which is therefore connected by two disjoint paths to respectively the left and right sides of $R(2n-1)$, as shown in Figure 5.3.

Denoting by $A_n(k)$ the event that the vertex (n,k) is joined by an open path to the surface $\partial B(n,(n,k))$ of the box $B(n,(n,k))$ having side-length $2n$ and centered at (n,k), we have therefore

$$\mathbb{P}_{1/2}(LR(2n-1)) \leq \sum_{k=0}^{2n-1} \mathbb{P}_{1/2}(A_n(k) \circ A_n(k))$$

and applying the BK inequality, we get

$$\begin{aligned} \mathbb{P}_{1/2}\left(LR(2n-1)\right) &\leq \sum_{k=0}^{2n-1} \mathbb{P}_{1/2}^2(A_n(k)) \\ &= 2n\mathbb{P}_{1/2}^2(A_n(0)) \\ &= 2n\mathbb{P}_{1/2}^2(0 \leftrightarrow \partial B(n)). \end{aligned}$$

Now, Lemma 5.2 states that $\mathbb{P}_{1/2}\left(LR(2n-1)\right) = 1/2$ for all $n \in \mathbb{N}^*$, from which we deduce (5.5).

■

We obtain directly that the tail of distribution of the radius $\mathrm{rad}(C) = \max_{x \in C}\{|x|\}$ of the cluster size C from (5.5) by noting that

$$\mathbb{P}_p(0 \leftrightarrow \partial B(n/2)) \leq \mathbb{P}_p(0 \leftrightarrow \partial S(n)) = \mathbb{P}_p(\mathrm{rad}(C) \geq n).$$

5.3.2 Scaling theory

Scaling theory has been used by mathematical physicists to study the behavior of quantities such as $\theta(p)$, $\xi(p)$, $\mathbb{P}_p(|C| = n)$ near the critical point p_c. Theorem 5.2 suggests, at least when $d = 2$, that they follow a power law distribution in that transition region, and indeed this is taken as starting assumption (ansatz in Physics) for such quantities. In other words, scaling theory assumes that

$$\begin{aligned}
\theta(p) &\approx (p - p_c)^{\beta} \qquad \text{as } p \downarrow p_c && (5.7)\\
\chi(p) &\approx (p - p_c)^{-\gamma} \qquad \text{as } p \uparrow p_c && (5.8)\\
\xi(p) &\approx (p - p_c)^{-\nu} \qquad \text{as } p \uparrow p_c && (5.9)\\
\mathbb{P}_{p_c}(|C| = n) &\approx n^{-1-1/\delta} \qquad \text{as } n \to \infty && (5.10)\\
\mathbb{P}_{p_c}(\mathrm{rad}(C) \geq n) &\approx n^{-1/\rho} \qquad \text{as } n \to \infty && (5.11)
\end{aligned}$$

where the "critical exponents" $\beta > 0$, $\gamma > 0$, $\nu > 0$, $\delta > 1$ and $\rho > 0$ depend on the dimension d. The notation $f(p) \approx g(p)$ as $p \to p_c$ means that $\lim_{p \to p_c} \ln f(p)/\ln g(p) = 1$.

Scaling theory predicts that these critical exponents are not independent from each other, but obey sets of relations called "scaling relations". We are going to derive the one linking β, γ and δ.

More precisely, the ansatz for the distribution of the cluster at the origin for values for $p \leq p_c$ is

$$\mathbb{P}_p(|C| = n) = n^{-(1+\delta^{-1})} f_- \left(n/\xi^{\tau}(p)\right) \tag{5.12}$$

where $\tau > 0$ is a constant, and $f_-(\cdot)$ is a smooth (differentiable) positive function on $\mathbb{R}^+$. Theorem 3.3 would suggest a function $f_-(x) \approx C \exp(-Ax)$ for some $A, C > 0$, but this is not very important here. We just assume that $f_-(x) \to 0$ faster than any power of $1/x$ as $x \to \infty$.

When $p \geq p_c$, we will take a similar ansatz for the distribution of the finite cluster at the origin

$$\mathbb{P}_p(|C| = n) = n^{-(1+\delta^{-1})} f_+ \left(n/\xi^{\tau}(p)\right) \tag{5.13}$$

where $f_+(\cdot)$ is a smooth (differentiable) positive function on $\mathbb{R}^+$. Here again, Theorem 4.2 would suggest to take $f_+(x) \approx C' \exp(-A' x^{(d-1)/d})$ for some $A', C' > 0$, but again we do not want to make this assumption here. We just assume that $f_+(x) \to 0$ faster than any power of $1/x$ as $x \to \infty$.

Now, we make the following approximate computations, first when $p < p_c$:

$$\begin{aligned}
\chi(p) &= \sum_{n \in \mathbb{N}^*} n \mathbb{P}_p(|C| = n) \simeq \sum_{n \in \mathbb{N}^*} n^{-\delta^{-1}} f_- \left(n/\xi^{\tau}(p)\right)\\
&\simeq \int_0^{\infty} n^{-\delta^{-1}} f_- \left(n/\xi^{\tau}(p)\right) dn\\
&= \xi^{\tau(1-\delta^{-1})}(p) \int_0^{\infty} u^{-\delta^{-1}} f_-(u) du.
\end{aligned}$$

Making the assumptions (5.8) and (5.9), the latter equation becomes

$$(p - p_c)^{-\gamma} \approx (p - p_c)^{-\nu\tau(1-\delta^{-1})} \int_0^{\infty} u^{-\delta^{-1}} f_-(u) du,$$

and since the integral converges because we assumed $\delta > 1$, we find that

$$\tau\nu = \frac{\gamma}{1 - \delta^{-1}} \tag{5.14}$$

We continue now with $p > p_c$, assuming that $\theta(p_c) = 0$ (we know it for sure for $d = 2$), so that

$$\begin{aligned}
\theta(p) &= 1 - \sum_{n\in\mathbb{N}^*} \mathbb{P}_p(|C| = n) \\
&= \sum_{n\in\mathbb{N}^*} [\mathbb{P}_{p_c}(|C| = n) - \mathbb{P}_p(|C| = n)] \\
&\simeq \sum_{n\in\mathbb{N}^*} n^{-1-\delta^{-1}} [f_+(0) - f_+(n/\xi^\tau(p))] \\
&\simeq \int_0^\infty n^{-1-\delta^{-1}} [f_+(0) - f_+(n/\xi^\tau(p))]\, dn \\
&= \xi^{\tau\delta^{-1}}(p) \int_0^\infty u^{-1-\delta^{-1}} [f_+(0) - f_+(u)]\, du
\end{aligned}$$

Plugging (5.7) and (5.9) in the latter equation, it becomes

$$(p - p_c)^\beta \approx (p - p_c)^{\nu\tau\delta^{-1}} \int_0^\infty u^{-1-\delta^{-1}} [f_+(0) - f_+(u)]\, du.$$

The integrand behaves like $u^{-1-\delta^{-1}} df_+(0)/du$ near $u = 0$, hence the integral converges. As a result, we find that

$$\tau\nu = \beta\delta.$$

Combining this relation with (5.14) gives the scaling relation

$$\gamma + \beta = \beta\delta. \tag{5.15}$$

This latter equation is one among many scaling relations. It shows that at least one among the three critical exponents, at most two are independent from each other. Interestingly, relation (5.15) does not depend on the dimension $d \geq 2$ of the lattice.

Another set of relations depend on the dimension d, and are believed to be valid only for dimensions $2 \leq d \leq d_c$ where the d_c is called the "critical dimension". These relations are called hyper-scaling relations, and read

$$d\nu = \gamma + 2\beta \tag{5.16}$$

$$d\rho = \delta + 1. \tag{5.17}$$

The scaling relations are widely accepted, but the hyper-scaling relations are more questionable.

The values of the scaling exponents obtained numerically for $d = 2$ are $\beta = 5/36$, $\gamma = 43/18$, $\delta = 91/5$, $\nu = 4/3$, $\rho = 48/5$.

6

Site and tree percolations

6.1 Introduction

In this chapter, we examine two other classical models of discrete percolation. The first one, tree percolation, is actually easier than bond percolation on the lattice. The major difference is that circuits are absent in a tree, contrary to a lattice, and this makes things considerably simpler: there is now a unique path between any two vertices of the tree. The second one, site percolation, is more difficult to handle, because it amounts to introduce dependencies between the edges, but is also more general than bond percolation, because to every bond model corresponds a site model, but not vice-versa. Nevertheless, most findings of bond percolation carry over at least qualitatively to site percolation.

6.2 Percolation on a tree

We replace the lattice $\mathbb{L}^d$ by a d-ary tree $\mathbb{T}$ (also called Bethe lattice), whose root is the origin. The root is connected to d nodes, called its "children", each of its children is in turn connected to d new "grand-children", and so forth. All nodes are thus directly connected to their common ancestor and to their d direct children. Note that the origin is only connected to its d children, hence to make the tree regular, one would need to add a new child only for the origin, so that each node would then be connected to exactly $d+1$ other nodes. However this has strictly no impact on critical exponents for the infinite tree, and so we do not need to care much about this detail. For the sake of simplification, we also consider only the binary tree ($d = 2$); there is no much difference with the general case.

Each edge of the tree $\mathbb{T}$ is open with probability p, and closed otherwise, independently of all other edges. Let us denote as usual by C the cluster at the origin (root of the tree).

6.2.1 Percolation probability

Let us first compute $\theta(p) = \mathbb{P}_p(|C| = \infty)$. Observe that C is the family of all descendants of the root, according to a Galton-Watson branching process. Let $X(n)$ be the number of vertices belonging to C at the nth layer of the tree (with the origin being at layer 0), they form the nth generation of the descendants of the root. Since $X(0) = 1$ (there is one node at the root of the tree), the probability that the branching process dies out for some finite n is

$$1 - \theta(p) = \mathbb{P}_p(X(n) = 0 \text{ for some } n \in \mathbb{N}^* \mid X(0) = 1).$$

Now, $X(n)$ is a homogeneous Markov chain, and we know from the theory of Markov chains that the above probability is the minimal solution of the set of equations, for all $i \in \mathbb{N}$,

$$\begin{array}{rcll} h_{i0} & = & 1 & \text{si} \quad i = 0 \\ h_{i0} & = & \sum_{j=0}^{\infty} p_{ij} h_{j0} & \text{if} \quad i \in \mathbb{N}^* \end{array} \tag{6.1}$$

where $h_{i0} = \mathbb{P}_p(X(n) = 0 \text{ for some } n \in \mathbb{N}^* \mid X(0) = i)$ and $p_{ij} = \mathbb{P}_p(X(n+1) = j \mid X(n) = i)$ are the transition probabilities of the chain. Here, $h_{i0} = h_{10}^i$ by independence between the i families born from the i ancestors, while for $j = 0, 1, 2$,

$$p_{1j} = \mathbb{P}_p(X(n+1) = j | X(n) = 1) = 2p^j(1-p)^{2-j},$$

so that $1 - \theta(p) = h_{10}$ is the minimal solution of

$$h_{10} = \sum_{j=0}^{\infty} p_{1j} h_{10}^j.$$

Denoting by $G(z) = \sum_j z^j p_{1j} = (1 - p + pz)^2$ the probability generating function of the number of children of a given vertex, we see that $1 - \theta(p) = h_{10}$ is the minimal solution of

$$z = G(z) = (1 - p + pz)^2,$$

which is 1 if $p \leq 1/2$ and $((1-p)/p)^2$ if $p > 1/2$. This shows that the critical threshold for bond percolation on $\mathbb{T}$ is $p_c = 1/2$, and

$$\theta(p) = \begin{cases} 0 & \text{if} \quad p \leq 1/2 \\ 1 - \left(\frac{1-p}{p}\right)^2 & \text{if} \quad p > 1/2. \end{cases} \tag{6.2}$$

Developing this expression in a Taylor expansion around $p = p_c = 1/2$, we find that

$$\theta(p) = 8(p - p_c) + O((p - p_c)^2)$$

for $p \downarrow p_c$. This indicates that the critical exponent β in the ansatz $\theta(p) \approx (p - p_c)^\beta$ is $\beta = 1$ for tree percolation.

6.2.2 Mean cluster size

It is as easy to compute the mean cluster size $\chi(p) = \mathbb{E}_p[|C|]$. Indeed, any vertex v of $\mathbb{T}$ belongs to C if and only if every edge in the (unique) path connecting v to 0 is open. If the vertex v is at the nth layer of the tree, this path if open with probability p^n, and as there are 2^n vertices at the nth layer, we have that for $p < 1/2$

$$\chi(p) = \sum_{n=0}^{\infty} \sum_{v=1}^{2^n} \mathbb{E}_p\left[1_{\{0 \leftrightarrow v\}}\right] = \sum_{n=0}^{\infty} 2^n p^n = \frac{1}{2}\left(\frac{1}{2} - p\right)^{-1},$$

which also shows that the critical exponent γ in the ansatz $\chi(p) \approx (p_c - p)^{-\gamma}$ for $p \uparrow p_c$ is $\gamma = 1$ for tree percolation.

6.2.3 Cluster size distribution

The computation of $\mathbb{P}_p(|C| = n)$ is not difficult either. Let $G_{|C|}(z)$ be the probability generating function of $|C|$. Since

$$\begin{aligned}
\mathbb{P}_p(|C| = 1) &= (1-p)^2 \\
\mathbb{P}_p(|C| = 2) &= 2p(1-p) \\
\mathbb{P}_p(|C| = n) &= 2p(1-p)\mathbb{P}_p(|C| = n-1) + p^2 \sum_{k=1}^{n-2} \mathbb{P}_p(|C| = k)\mathbb{P}_p(|C| = n-k-1) \text{ for } n \geq 3,
\end{aligned}$$

we obtain after some computations that

$$G_{|C|}(z) = \left(1 - 2p(1-p)z - \sqrt{1 - 4p(1-p)z}\right) / 2p^2 z$$

whose inverse z-transform yields that

$$\mathbb{P}_p(|C| = n) = \frac{1}{n} \begin{pmatrix} 2n \\ n-1 \end{pmatrix} p^{n-1}(1-p)^{n+1}. \tag{6.3}$$

Stirling's formula $n! \approx n^n e^{-n}\sqrt{2\pi n}$ yields that

$$\begin{pmatrix} 2n \\ n-1 \end{pmatrix} = \frac{n}{n+1} \begin{pmatrix} 2n \\ n \end{pmatrix} \approx \frac{4^n}{\sqrt{\pi n}}$$

and therefore (6.3) behaves like $\mathbb{P}_p(|C| = n) \approx 1/\sqrt{\pi n^3}$ for $n \to \infty$. This indicates that the critical exponent δ in the ansatz $\mathbb{P}_p(|C| = n) \approx n^{-1-1/\delta}$ is $\delta = 2$ for tree percolation.

Observe that the three values $\beta = 1$, $\gamma = 1$ and $\delta = 2$ satisfy the scaling relation (5.15)!

We could also compute that $\rho = 1/2$ and $\nu = 1/2$. Plugging these values in the hyperscaling relations (5.16) and (5.17), we find that $d = 6$. This suggests that the critical dimension $d_c = 6$. Indeed, we can embed $\mathbb{T}$ in $\mathbb{L}^\infty$, with each edge connecting a nth layer vertex of $\mathbb{T}$ to a $(n+1)$th layer vertex being parallel to the nth coordinate axis of $\mathbb{L}^\infty$. This would make percolation in $\mathbb{T}$ and $\mathbb{L}^\infty$ similar, roughly speaking. The computations for the tree suggests that the two processes are similar already for $\mathbb{L}^d$ with $d \geq 6$.

6.3 Site percolation

A much more important model in practice is to close vertices rather than edges in a lattice $\mathbb{L}^d$. The corresponding model is called site percolation, all definitions of percolation probability, critical probability, etc remain the same as in the bond model, the only difference is that vertices (and not edges) are open with probability p, and closed with probability $1-p$.

One can show that a phase transition occurs between a sub-critical and super-critical phases, and essentially most properties of bond percolation extend to site percolation. However, contrary to bond percolation, the percolation threshold p_c is not known mathematically. It is found numerically to be close to 0.59.

Site percolation is more general than bond percolation, in the sense that every bond model can be recast as a site model, but not the reverse. To recast a bond model as a site model, we make use of the notion of covering graph $\mathcal{G}_c$ of a graph $\mathcal{G}$, which is obtained as follows. Place a vertex of $\mathcal{G}_c$ on the middle of each edge of $\mathcal{G}$. Two vertices of $\mathcal{G}_c$ are declared to be adjacent if and only if the two corresponding edges of $\mathcal{G}$ share a common endvertex of $\mathcal{G}$. Defined now a bond percolation process

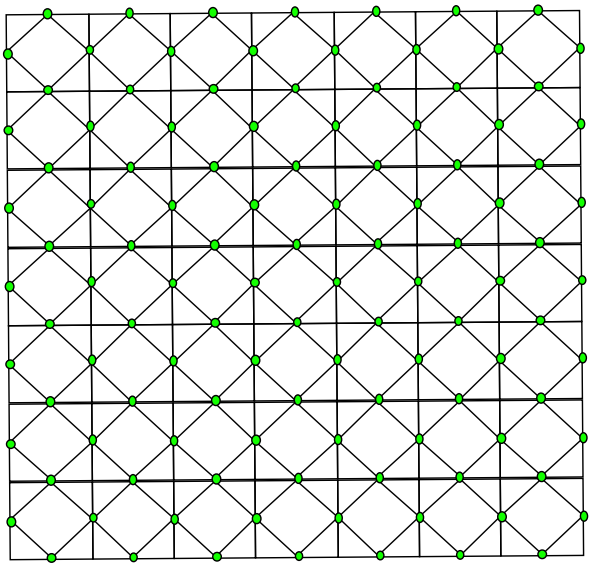

Figure 6.1: The covering graph $\mathbb{L}_c^2$ of the square lattice $\mathbb{L}^2$.

on $\mathcal{G}$, and declare a vertex of $\mathcal{G}_c$ to be open (resp., closed) if and only if the corresponding edges of $\mathcal{G}$ is open (resp., closed). This results in a site percolation process on $\mathcal{G}_c$. Any path of open edges of $\mathcal{G}$ corresponds to a path of open vertices of $\mathcal{G}_c$, and vice-versa. As a result,

$$p_c^{\text{bond}}(\mathcal{G}) = p_c^{\text{site}}(\mathcal{G}_c) \tag{6.4}$$

For example, if $\mathcal{G} = \mathbb{L}^2$, then we find that $\mathcal{G}_c = \mathbb{L}_c^2$ is the lattice shown in Figure 6.1, where each site has exactly six adjacent vertices. Because of (6.4), the site percolation threshold on this graph is $1/2$. We can show that the triangular lattice, where each vertex is also adjacent to six other vertices, has a also a site percolation threshold equal to $1/2$.

7

Full versus partial connectivity

Percolation is concerned with the emergence of a giant cluster in a large network, but it does not directly provide information on its full connectivity, which occurs when there is an open path between any pair of nodes of the network. Requiring a network to be fully connected instead of simply being super-critical is clearly much more stringent. What is the price to pay for full connectivity ? In this chapter, we return to the simplest, non-trivial setting of our bond percolation model to compare the two properties. The question is trivially answered if we take the infinite lattice $\mathbb{L}^d$. Indeed, in this case a direct application of the 0-1 law yields that

$$\mathbb{P}_p(\text{network fully connected}) = \begin{cases} 0 & \text{if} \quad p < 1 \\ 1 & \text{if} \quad p = 1. \end{cases}$$

Therefore the correct asymptotic is to take a finite graph of n vertices, and to study the full connectivity of the network as $n \to \infty$. At each step n, we will compute the value $p = p_n$ of the open edge probability for which the graph G_n obtained on a box of $\mathbb{L}_2$ containing n vertices is fully connected, or contains only isoalted nodes. The term "with high probability" (w.h.p.) means that the property holds asymptotically almost surely, i.e., with probability 1 as $n \to \infty$.

7.1 Poisson approximation using the Chen Stein method

Let $\{I_i\}_{1\leq i\leq n}$ be a sequence of Bernouilli random variables, with $\mathbb{P}(I_i = 1) = p_i = 1 - \mathbb{P}(I_i = 0)$. Let $\lambda = \sum_{i=1}^n p_i$, and $A \subset \mathbb{N}$. We are interested in computing $\mathbb{P}(W \in A)$, where $W = \sum_{i=1}^n I_i$. The Chen Stein method enables to bound the error when this probability is approximated by $\sum_{k\in A} e^{-\lambda}\lambda^k/k!$, which is the probability that a Poisson random variable of rate λ takes a value in A, and which we will denote by $Po_\lambda(A)$. In other words,

$$Po_\lambda(A) = \sum_{k\in A} e^{-\lambda}\frac{\lambda^k}{k!}.$$

We follow the textbooks [3, 35].

Given $A \subset \mathbb{N}$ and λ, the method starts by defining recursively a function g by $g(0) = 0$, and, for $j \geq 0$, by

$$g(j+1) = \frac{1}{\lambda}\left[1_{\{j\in A\}} - Po_\lambda(A) + jg(j)\right], \tag{7.1}$$

where $1_{\{j\in A\}} = 1$ if $j \in A$ and $1_{\{j\in A\}} = 0$ otherwise. We can recast the recursion as

$$\lambda g(j+1) - jg(j) = 1_{\{j\in A\}} - Po_\lambda(A) \tag{7.2}$$

from which we get, after multiplying by both sides by p_j and adding up for all $0 \leq j \leq n$ that

$$\mathbb{E}[\lambda g(W+1) - Wg(W)] = \mathbb{P}(W \in A) - Po_\lambda(A). \tag{7.3}$$

Now,

$$\begin{aligned}\lambda g(W+1) - Wg(W) &= \left(\sum_{i=1}^n p_i\right) g(W+1) - \left(\sum_{i=1}^n I_i\right) g(W)\\ &= \sum_{i=1}^n (p_i g(W+1) - I_i g(W))\end{aligned}$$

from which we deduce, by taking expectations, that

$$\begin{aligned}\mathbb{E}[\lambda g(W+1) - Wg(W)] &= \sum_{i=1}^n (p_i\mathbb{E}[g(W+1)] - \mathbb{E}[I_i g(W)])\\ &= \sum_{i=1}^n (p_i\mathbb{E}[g(W+1)] - p_i\mathbb{E}[g(W) \mid I_i = 1])\\ &= \sum_{i=1}^n p_i\left(\mathbb{E}[g(W+1)] - \mathbb{E}[g(W) \mid I_i = 1]\right)\\ &= \sum_{i=1}^n p_i\left(\mathbb{E}[g(W+1)] - \mathbb{E}[g(V_i+1)]\right)\\ &= \sum_{i=1}^n p_i\mathbb{E}[g(W+1) - g(V_i+1)],\end{aligned}$$

where V_i is any random variable whose distribution is the same as the conditional distribution of $\sum_{k\neq i} I_k$ given that $I_i = 1$. In other words, V_i is any random variable such that for all $v \in \mathbb{N}$,

$$\mathbb{P}(V_i = v) = \mathbb{P}\left(\sum_{k\neq i} I_k = v \mid I_i = 1\right).$$

We have therefore that

$$\mathbb{P}(W \in A) - Po_\lambda(A) = \sum_{i=1}^{n} p_i \mathbb{E}[g(W+1) - g(V_i+1)]. \tag{7.4}$$

We need the following bound.

Lemma 7.1 (Bound on delta g). *For any $A \subset \mathbb{N}$ and λ*

$$|g(j+1) - g(j)| \leq \frac{1-e^{-\lambda}}{\lambda}. \tag{7.5}$$

Proof: We can check that the solution of (7.1) is

$$g(j+1) = \frac{j!}{\lambda^{j+1}} e^{\lambda} \left[Po_\lambda(A \cap \{0,1,\ldots,j\}) - Po_\lambda(A) Po_\lambda(\{0,1,\ldots,j\}) \right]. \tag{7.6}$$

Take $A = \{i\}$, for some $i \in \mathbb{N}$. If $i \leq j$, (7.6) becomes

$$\begin{aligned} g(j+1) &= \frac{j!}{\lambda^{j+1}} e^{\lambda} Po_\lambda(\{i\}) \left[1 - Po_\lambda(\{0,1,\ldots,j\})\right] \\ &= Po_\lambda(\{i\}) \sum_{k=j+1}^{\infty} \frac{j!}{k!} \lambda^{k-j-1} \\ &= Po_\lambda(\{i\}) \sum_{l=0}^{\infty} \left(\begin{array}{c} l+j+1 \\ l+1 \end{array} \right)^{-1} \frac{\lambda^l}{(l+1)!}, \end{aligned}$$

which shows that, when $j \geq i$, $g(j+1)$ is positive and decreases with j.

Similarly, for $i \geq j+1$, (7.6) becomes

$$\begin{aligned} g(j+1) &= -\frac{j!}{\lambda^{j+1}} e^{\lambda} Po_\lambda(\{i\}) Po_\lambda(\{0,1,\ldots,j\}) \\ &= -Po_\lambda(\{i\}) \sum_{k=0}^{j} \frac{j!}{k!} \frac{1}{\lambda^{j+1-k}} \\ &= -Po_\lambda(\{i\}) \sum_{l=0}^{j} \left(\begin{array}{c} l \\ j \end{array} \right) \frac{l!}{\lambda^l}, \end{aligned}$$

which shows that, when $j \leq i-1$, $g(j+1)$ is negative and decreases with j.

As a result, given i, the only value of $j \geq 1$ for which the difference $g(j+1) - g(j)$ is positive is $j = i$. Let us write $g = g_A$ to make the dependence of g on A explicit. We can verify from (7.6) that $g_A = \sum_{\{i \in A\}} g_{\{i\}}$, so that

$$g_A(j+1) - g_A(j) = \sum_{\{i \in A\}} \left(g_{\{i\}}(j+1) - g_{\{i\}}(j) \right) \leq g_{\{j\}}(j+1) - g_{\{j\}}(j)$$

. Hence an upper bound is obtained by taking $A = \{j\}$. In this case (7.6) yields that

$$\begin{aligned} g(j+1) - g(j) &= \frac{j!}{\lambda^{j+1}} e^{\lambda} Po_\lambda(\{j\}) \left[1 - Po_\lambda(\{0,1,\ldots,j\})\right] + \frac{(j-1)!}{\lambda^j} e^{\lambda} Po_\lambda(\{j\}) \left[Po_\lambda(\{0,1,\ldots,j-1\})\right] \\ &= \frac{1}{\lambda} \left[1 - Po_\lambda(\{0,1,\ldots,j\}) + \frac{\lambda}{j} Po_\lambda(\{0,1,\ldots,j-1\}) \right] \\ &\leq \frac{1}{\lambda} \left[1 - Po_\lambda(\{0,1,\ldots,j\}) + Po_\lambda(\{1,\ldots,j\})\right] \\ &= \frac{1}{\lambda}(1 - Po_\lambda(\{0\})) = \frac{1-e^{-\lambda}}{\lambda}. \end{aligned}$$

■

For $i < j$, we have that

$$g(j) - g(i) = g(j) - g(j-1) + g(j-1) - g(j-2) + \ldots + g(i+1) - g(i),$$

from which we deduce from (7.5) and the triangle inequality that for any $i, j \in \mathbb{N}$

$$|g(j) - g(i)| \leq |j - i| \frac{1 - e^{-\lambda}}{\lambda}.$$

Therefore,

$$|g(W+1) - g(V_i+1)| \leq |W - V_i| \frac{1 - e^{-\lambda}}{\lambda}.$$

Combining this inequality with (7.4) and Jensen's inequality, we obtain that

$$\begin{aligned} |\mathbb{P}(W \in A) - Po_\lambda(A)| &\leq \sum_{i=1}^{n} p_i \left|\mathbb{E}[g(W+1) - g(V_i+1)]\right| \\ &\leq \sum_{i=1}^{n} p_i \mathbb{E}[|g(W+1) - g(V_i+1)|] \\ &\leq \frac{1 - e^{-\lambda}}{\lambda} \sum_{i=1}^{n} p_i \mathbb{E}[|W - V_i|]. \end{aligned}$$

We have therefore the following theorem.

Theorem 7.1 (Chen-Stein approximation). *Let $\{I_i\}_{1 \leq i \leq n}$ be a sequence of Bernoulli random variables, with $\mathbb{P}(I_i = 1) = p_i = 1 - \mathbb{P}(I_i = 0)$. Let $W = \sum_{i=1}^{n} I_i$, $\lambda = \sum_{i=1}^{n} p_i$. Let V_i be a random variable whose distribution is the same as the conditional distribution of $\sum_{k \neq i} I_k$ given that $I_i = 1$, i.e. for all $v \in \mathbb{N}$,*

$$\mathbb{P}(V_i = v) = \mathbb{P}\left(\sum_{k \neq i} I_k = v \mid I_i = 1\right).$$

Then for any $A \subset \mathbb{N}$,

$$\left|\mathbb{P}(W \in A) - \sum_{i \in A} e^{-\lambda} \frac{\lambda^i}{i!}\right| \leq \frac{1 - e^{-\lambda}}{\lambda} \sum_{i=1}^{n} p_i \mathbb{E}[|W - V_i|]. \tag{7.7}$$

7.2 Chen Stein method with coupled variables

The difficulty in applying directly Theorem 7.1 stems from the need to find the appropriate variable V_i whose distribution coincides with $W - 1$ given that $I_i = 1$.

If all variables I_i are independent, then one naturally takes $V_i = \sum_{j \neq i} I_j$, and we find that $\mathbb{E}[|W - V_i|] = \mathbb{E}[I_i] = p_i$, so that (7.7) becomes

$$\left|\mathbb{P}(W \in A) - \sum_{i \in A} e^{-\lambda} \frac{\lambda^i}{i!}\right| \leq \frac{1 - e^{-\lambda}}{\lambda} \sum_{i=1}^{n} p_i^2.$$

The power of the Chen-Stein method appears however when the random variables I_i are dependent. In this case, the theorem gets considerably simplified if we can find some natural coupling between

the variables I_i and another sequence of indicator variables J_{ij}, defined on the same probability space as I_i, and whose distribution given that $I_j = 1$ is identical to that I_i. That is,

$$\begin{aligned}\mathbb{P}(J_{ij} = 1) &= \mathbb{P}(I_i = 1 \mid I_j = 1)\\ \mathbb{P}(J_{ij} = 0) &= \mathbb{P}(I_i = 0 \mid I_j = 1).\end{aligned}$$

Such a coupling is

$$J_{ij} \geq I_i \tag{7.8}$$

for all $i \neq j$, which implies that the original Bernoulli variables I_i are positively correlated, as

$$COV[I_i, I_j] = \mathbb{E}[I_i I_j] - p_i p_j = \mathbb{E}[I_i \mid I_j = 1]p_j - \mathbb{E}[I_i]p_j = \mathbb{E}[J_{ij} - I_i]p_j \geq 0.$$

In this case indeed, we set

$$V_i = \sum_{k \neq i} J_{ki},$$

whence

$$\begin{aligned}p_i\mathbb{E}[|W - V_i|] &= p_i\mathbb{E}\left[\left|I_i + \sum_{k\neq i} I_k - J_{ki}\right|\right] \leq p_i\mathbb{E}\left[I_i + \left|\sum_{k\neq i} I_k - J_{ki}\right|\right]\\ &\leq p_i^2 + p_i\sum_{k\neq i}\mathbb{E}\left[|I_k - J_{ki}|\right] = p_i^2 + \sum_{k\neq i} p_i\mathbb{E}\left[J_{ki} - I_k\right] = p_i^2 + \sum_{k\neq i} COV[I_k, I_i].\end{aligned}$$

Since $VAR[W] = \sum_{k\neq i} COV[I_k, I_i] + \sum_{i=1}^n VAR[I_i] = \sum_{k\neq i} COV[I_k, I_i] + \sum_{i=1}^n p_i(1-p_i)$, (7.7) becomes

$$\left|\mathbb{P}(W \in A) - \sum_{i\in A} e^{-\lambda}\frac{\lambda^i}{i!}\right| \leq \frac{1-e^{-\lambda}}{\lambda}\left(VAR[W] - \lambda + 2\sum_{i=1}^n p_i^2\right).$$

The explicit coupling (7.8) is not always easy to find, and fortunately it can be replaced by a weaker condition, using associated random variables and an extension of the FKG inequality. The set of random variables $X_1, \dots X_n$ are said to be associated if for all increasing functions $f(\cdot)$ and $g(\cdot)$ (i.e. such that $f(x_1, \dots, x_n) \leq f(y_1, \dots, y_n)$ if $x_i \leq y_i$, $1 \leq i \leq n$),

$$\mathbb{E}[f(X_1, \dots, X_n)g(X_1, \dots, X_n)] \geq \mathbb{E}[f(X_1, \dots, X_n)]\mathbb{E}[g(X_1, \dots, X_n)].$$

Skipping the details [3], we can state the following theorem, that simplifies Theorem 7.1 when we know that the indicators I_i are increasing (or decreasing) functions of some independent random variables.

Theorem 7.2 (Chen-Stein approximation with monotone coupling). *Let $\{I_i\}_{1\leq i\leq n}$ be a sequence of Bernouilli random variables, with $\mathbb{P}(I_i = 1) = p_i = 1 - \mathbb{P}(I_i = 0)$, that are increasing (decreasing) functions of some independent random variables $X_1, \dots, X_m$. Let $W = \sum_{i=1}^n I_i$, $\lambda = \sum_{i=1}^n p_i$. Then for any $A \subset \mathbb{N}$,*

$$\left|\mathbb{P}(W \in A) - \sum_{i\in A} e^{-\lambda}\frac{\lambda^i}{i!}\right| \leq \frac{1-e^{-\lambda}}{\lambda}\left(VAR[W] - \lambda + 2\sum_{i=1}^n p_i^2\right). \tag{7.9}$$

7.3 Computation of the number of isolated nodes

We follow the approach of Franceschetti and Meester [16]. Let us consider a box $B(m) = [-m, m] \times [-m, m]$ of $\mathbb{L}^2$, on which we define an independent bond model, with open edge probability $p = p_n$

where $n = (2m+1)^2$ is the number of vertices in $B(m)$. The resulting graph is denoted by G_n. Let $i = (i_1, i_2)$ be a vertex of $B(m)$, let I_i be the indicator that vertex i is isolated (i.e. that all the incident edges in i are closed). The set of the four vertices at the corners of $B(m)$ is denoted by $\square B(m)$, the set of all other vertices on the boundaries is denoted by $\delta B(n)$, and the set of all interior vertices coincides with the set of vertices of $B(m-1)$. We have then that $\mathbb{P}(I_i = 1) = p_i$ is given by

$$\begin{aligned}\mathbb{P}(I_i = 1) &= (1-p_m)^4 \text{ if } i \in B(m-1)\\ \mathbb{P}(I_i = 1) &= (1-p_m)^3 \text{ if } i \in \delta B(m)\\ \mathbb{P}(I_i = 1) &= (1-p_m)^2 \text{ if } i \in \square B(m).\end{aligned}$$

We will also denote by $i \sim j$ the fact that vertices i and j are adjacent (i.e., are the two end-vertices of some edge), and by $i \nsim j$ the fact that vertices i and j are not adjacent. In this section, we want to compute the distance between the total number of isolated vertices $W = \sum_{i=1}^n I_i$ and a Poisson random variable of rate $\mu = n(1-p_n)^4$. We will keep μ fixed as $n \to \infty$, which amounts to take $(1-p_n) = O(n^{-1/4})$.

Now, it is clear that I_i is a decreasing function of the state of the four edges incident in i, which are independent random variables. Therefore we can apply Theorem 7.2. We need to compute $VAR[W]$ in order to evaluate the right hand side of (7.9).

We first get that

$$\sum_{i=1}^n p_i^2 = (2m-1)^2(1-p_n)^8 + (8m-4)(1-p_n)^6 + 4(1-p_n)^4 = O\left(n^{-1}\right).$$

Next, we compute

$$\begin{aligned}\mathbb{E}[W^2] &= \mathbb{E}\left[\sum_{i,j=1}^n I_i I_j\right] = \mathbb{E}\left[\sum_{i=1}^n I_i + \sum_{i=1}^n \sum_{j\sim i} I_i I_j + \sum_{i=1}^n \sum_{j\nsim i} I_i I_j\right]\\ &= \lambda + \sum_{i=1}^n \sum_{j\sim i} \mathbb{E}[I_i I_j] + \sum_{i=1}^n \sum_{j\nsim i} \mathbb{E}[I_i I_j]. \qquad (7.10)\end{aligned}$$

When $i \sim j$, $\mathbb{E}[I_i I_j] = \mathbb{P}(\{I_i = 1\} \cap \{I_j = 1\}) = (1-p_n)^x$ with $x = 7$ if $i, j \in B(m-1)$; $x = 6$ if either $i \in B(m-1)$ and $j \in \delta B(m)$, or vice versa; $x = 5$ if $i, j \in \delta B(m)$ and $x = 4$ if i or $j \in \square B(m)$. Therefore

$$\begin{aligned}\sum_{i=1}^n \sum_{j\sim i} \mathbb{E}[I_i I_j] &= O(n)(1-p_n)^7 + O(n^{1/2})\left[(1-p_n)^6 + (1-p_n)^5\right] + 4(1-p_n)^4\\ &= O\left(n^{-3/4}\right) \qquad (7.11)\end{aligned}$$

On the other hand, when $i \nsim j$, $\mathbb{E}[I_i I_j] = \mathbb{E}[I_i]\mathbb{E}[I_j] = p_i p_j$, from which we deduce that

$$\begin{aligned}\sum_{i=1}^n \sum_{j\nsim i} \mathbb{E}[I_i I_j] &= \sum_{i=1}^n p_i \sum_{j\nsim i} p_j = \sum_{i=1}^n p_i \left(\sum_{j=1}^n p_j - \sum_{j\sim i} p_j - p_i\right) = \left(\sum_{i=1}^n p_i\right)^2 - \sum_{i=1}^n p_i \sum_{j\sim i} p_j - \sum_{i=1}^n p_i^2\\ &= \lambda^2 + O(n)(1-p_n)^8 - \sum_{i=1}^n p_i^2 = \lambda^2 + O\left(n^{-1}\right) \qquad (7.12)\end{aligned}$$

Plugging (7.11) and (7.12) in (7.10), we find that $\mathbb{E}[W^2] = \lambda^2 + \lambda + O\left(n^{-3/4}\right)$, and therefore that

$$VAR[W] - \lambda + 2\sum_{i=1}^n p_i^2 = \lambda^2 + \lambda - \lambda^2 - \lambda + O\left(n^{-3/4}\right) = O\left(n^{-3/4}\right),$$

from which we deduce that

$$\left|\mathbb{P}(W \in A) - \sum_{i \in A} e^{-\lambda}\frac{\lambda^i}{i!}\right| \to 0$$

as $n \to \infty$.

Finally, as

$$\begin{aligned}\lambda &= \sum_{i=1}^{n} p_i = (2m-1)^2(1-p_n)^4 + (8m-4)(1-p_n)^3 + 4(1-p_n)^2 \\ &= n(1-p_n)^4 + O\left(n^{-1/2}\right) = \mu + O\left(n^{-1/2}\right),\end{aligned}$$

it shows that $\lambda \to \mu$ as $n \to \infty$.

We therefore shown the following:

Theorem 7.3 (Number of isolated vertices). *If for some $\mu > 0$, the open edge probability p_n is such that*

$$\mu = n(1-p_n)^4, \tag{7.13}$$

then as $n \to \infty$, the number W of isolated vertices in G_n converges in distribution to a Poisson random variable of rate μ.

A consequence of this theorem is as follows [16].

Corollary 7.1 (Probability of finding isolated vertices). *Let the open edge probability be $p_n = 1 - c_n n^{-1/4}$, where c_n is an arbitrary sequence of positive reals. Then, as $n \to \infty$,*

$$\mathbb{P}(\textit{no isolated vertices in} G_n) \to \exp(-c^4) \tag{7.14}$$

if and only if $c_n \to c$.

Proof: Let A_n be the event that there are no isolated vertices int he graph.

($\Leftarrow$) Suppose that $c_n \to c$. Then for any $\varepsilon > 0$, there exists $N_1 \in \mathbb{N}$ such that for all $n > N_1$, $c - \varepsilon \leq c_n \leq c + \varepsilon$.

Now, plugging $p_n = 1 - cn^{-1/4}$ in (7.13) and using Theorem 7.3, we find that

$$\lim_{n\to\infty} \mathbb{P}(A_n) = \exp(-c^4).$$

Because $\mathbb{P}(A_n)$ is decreasing in c for all n, there exists $N_2 \geq N_1$ such that for all $n > N_2$, $\exp(-(c+\varepsilon)^4 \leq \mathbb{P}(A_n) \leq exp(-(c-\varepsilon)^4$. As ε can be arbitrary small, it shows that if $c_n \to c$, then (7.14) holds.

($\Rightarrow$) Suppose that $c_n \nrightarrow c$. Then there exists a subsequence $\{c_{n_k}\}_{k\in\mathbb{N}}$ that converges to $c' \neq c$ as $k \to \infty$. The reasoning made in the first part of the proof then implies that the subsequence $\mathbb{P}(A_{n_k})$ converges to $\exp(-c'^4) \neq \exp(-c^4)$. As a result, $\mathbb{P}(A_n) \nrightarrow \exp(-c^4)$, because otherwise all subsequences of $\{\mathbb{P}(A_n)\}$ would also converge to $\exp(-c^4)$.

■

7.4 Condition for full connectivity

The following results, established by Franceschetti and Meester [16], show if the network is not fully connected, then it is very likely to have only isolated vertices. The term "with high probability" (w.h.p.) means that the property holds asymptotically almost surely, i.e., with probability 1 as $n \to \infty$.

Theorem 7.4 (Absence of larger finite clusters). *For any $0 < c < \infty$, if $p_n = 1 - cn^{-1/4}$, then w.h.p. either the graph G_n is fully connected, or it contains only isolated vertices.*

Proof: Suppose that $B(m)$ contains at least one disconnected cluster of at least two nodes. If this cluster is in the interior $B(m-1)$, then it must be surrounded by a circuit of length at least equal to 6. If one of the node of the cluster is on the boundary $\partial B(m)$, then this cluster must be separated from the rest of the network by a self-avoiding path in the dual lattice, starting and ending on $\partial B(m)$, and of length at least equal to 3.

We know from Chapter 2 the following bound on the probability of finding a self-avoiding path of length l in the dual lattice

$$\mathbb{P}(\exists \text{ path of length } l) \leq \frac{4}{3} \sum_{k=l}^{\infty} \left(3(1-p_n)\right)^k .$$

Taking $p_n = 1 - cn^{-1/4}$, this bound becomes

$$\mathbb{P}(\exists(\text{ path of length } l) \leq \frac{4}{3} \sum_{k=l}^{\infty} \left(3cn^{-1/4}\right)^k = \frac{4(3c)^l}{3} \frac{n^{-l/4}}{1-3cn^{-1/4}}. \tag{7.15}$$

For disconnected clusters whose vertices are all in the interior $B(m-1)$, we take $l = 6$ in (7.15). There are $(2m-1)^2 = n - 4\sqrt{n} + 4$ such vertices. Therefore, by the union bound, the probability of finding a disconnected cluster of at least two vertices is less than

$$\mathbb{P}(\exists \text{ disconnected cluster of size at least 2 in } B(m-1)) \leq \frac{4(3c)^6}{3} \frac{n^{-1/2} - 4n^{-1} + 4^{-3/2}}{1-3cn^{-1/4}} = O\left(n^{-1/2}\right).$$

Similarly, for disconnected clusters with at least one vertex on boundary $\partial B(m)$, we take $l = 3$ in (7.15), and we note that there are $4(2m+1) = 4\sqrt{n}$ such vertices, so that the probability of finding a disconnected cluster of at least two vertices is less than

$$\mathbb{P}(\exists(\text{ disconnected cluster of size at least 2 touching } \partial B(m)) \leq \frac{4(3c)^3}{3} \frac{4n^{-1/4}}{1-3cn^{-1/4}} = O\left(n^{-1/4}\right).$$

Consequently, as $n \to \infty$, the probability of finding a disconnected cluster of size at least equal to 2 becomes zero.

■

We combine Corollary 7.1 and Theorem 7.4 to obtain the main result of this chapter.

Theorem 7.5 (Threshold for full connectivity). *Let the open edge probability be $p_n = 1 - c_n n^{-1/4}$, where c_n is an arbitrary sequence of positive reals. The graph in $G(n)$ is fully connected w.h.p. if and only if $c_n \to 0$.*

8

Random Graphs

8.1 Introduction

The theory of random graphs began in the late 1950s with the seminal paper by Erdös and Rényi [31]. In contrast to percolation theory, which emerged from efforts to model physical phenomena such as the behavior of liquids in porous materials, random graph theory was devised originally as a mathematical tool in existence proofs of certain combinatorial objects. However, our goal is to study random graphs as models for networks, and this will govern our choice of results and insights we present here.

Both percolation theory and the theory of random graphs have permeated and enriched many fields beyond the initial focus of the pioneers: mathematics, statistical physics, sociology, computer science, biology. Another feature that random graph theory shares with percolation theory is that it remains a very active area of research, with many seemingly simple questions remaining open.

The key difference between percolation models and random graphs is that random graphs are not constrained by an underlying lattice. Rather, every pair of vertices (nodes, sites) can potentially be connected by an edge (bond). As such, RGs are clearly not an adequate model for networks that "live" in a geometric space, such as ad hoc wireless networks. However, they may be a reasonable model for certain features of "logical" networks that live in high or infinite-dimensional spaces, such as peer-to-peer overlay networks or sociological networks. In addition, the results we derive and the methods we develop in this chapter will also serve as a basis to define and study small-world and scale-free networks.

Despite this key difference, many of the phenomena we have studied in percolation theory occur also in random graphs. There are sharp thresholds that separate regimes where various graph properties jump from being very unlikely to very likely. One of these thresholds concerns the emergence of a giant component, just as in percolation.

As we shall see shortly, a random graph with constant p is not a very interesting object to study for our purposes, as the graph is so richly connected that every node is only two hops away from every other node. In fact, with constant p, the degree (i.e., the number of edges per vertex) grows linearly with n, while many real networks are much sparser. As we will see, interesting behavior (such as phase transitions from many small clusters to one dominating giant cluster) occurs within much sparser random graphs.

To focus on such graphs, it is necessary to let $p = p(n)$ depend on n; specifically, we will let $p(n)$ go to zero in different ways, which will give rise to several interesting regimes, separated by phase transitions. Contrast this with percolation theory, where the phase transition occurred for $p = p_c$ independent of the lattice size. For this reason, it was not necessary to work out results on a finite lattice of size n and then to study the limit $n \to \infty$; we could directly study the infinite lattice. In random graph theory, on the other hand, we need to perform the extra step of going to the limit, and we will be interested in properties of RGs whose probability goes to one when $n \to \infty$. Such a property Q is said to occur *asymptotically almost surely (a.a.s.)*, although many authors use the somewhat imprecise term *almost every* graph has property Q (a.e.), or also *property Q occurs with high probability (w.h.p.)*.

Definition 8.1 (Random graph). *Given n and p, a random graph $G(n,p)$ is a graph with labeled vertex set $[n] = \{1, \ldots, n\}$, where each pair of vertices has an edge independently with probability p.*

As the node degree has a binomial distribution $\text{Binom}(n-1, p)$, this random graph model is sometimes also referred to as the binomial model. We point out that various other types of random graphs have been studied in the literature; we will discuss *random regular graphs*, another class of random graphs, in the next chapter.

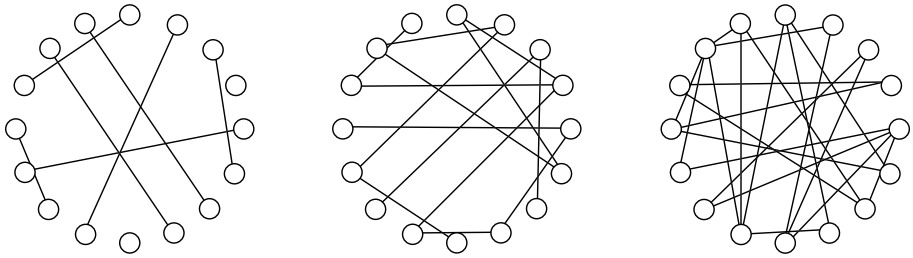

Figure 8.1: Three realizations of $G(16,p)$, with increasing p.

8.2 Preliminaries

Theorem 8.1 (Almost every $G(n,p)$ is connected). *For constant p, $G(n,p)$ is connected* a.a.s.

Proof: If G is disconnected, then there exists a bypartition of $V(G) = S \cup \bar{S}$ such that there are no edges between S and $\bar{S}$. We can union-bound the probability that there is no such partition $(S, \bar{S})$ by summing over all possible partitions.

Condition on $|S| = s$. There are $s(n-s)$ possible edges connecting a node in S to a node in $\bar{S}$, so $\mathbb{P}\left\{S \text{ and } \bar{S} \text{ are disconnected}\right\} = (1-p)^{s(n-s)}$.

The probability that $G(n,p)$ is disconnected is at most $\sum_{s=1}^{n/2} \left(\begin{array}{c} n \\ s \end{array} \right) (1-p)^{s(n-s)}$ (note that we do not need to sum beyond $n/2$). Using the bounds $\left(\begin{array}{c} n \\ s \end{array} \right) < n^s$ and $(1-p)^{n-s} \leq (1-p)^{n/2}$, we find $\mathbb{P}\left\{G(n,p) \text{disconnected}\right\} < \sum_{s=1}^{n/2} (n(1-p)^{n/2})^s$.

For n large enough, $x = n(1-p)^{n/2} < 1$, and the sum above is a convergent geometric series $\sum_{k=1}^{n/2} x^s < x/(1-x)$. Since $x \to 0$, the probability that $G(n,p)$ is disconnected $\to 0$ as well.

■

The above union bound is very loose, as graphs with many components are counted several times. We now illustrate two methods that are used frequently to prove results of the above type. Specifically, we often face the task of proving that a graph $G(n,p)$ has some property either with probability going to zero or to one. We assume here that X_n is an integer ≥ 0.

Theorem 8.2 (First moment method). *If* $\mathbb{E}[X_n] \to 0$, *then* $X_n = 0$ a.a.s.

Proof: Apply the Markov inequality $\mathbb{P}\{X \geq x\} \leq \mathbb{E}[X]/x$ with $x = 1$.

■

Theorem 8.3 (Second moment method). *If* $\mathbb{E}[X_n] > 0$ *for* n *large and* $Var[X_n]/(\mathbb{E}[X_n])^2 \to 0$, *then* $X(G) > 0$ a.a.s.

Proof: Chebyshev's inequality states that if $\mathrm{Var}[X]$ exists, then $\mathbb{P}\{|X - \mathbb{E}[X]| \geq x\} \leq \mathrm{Var}[X]/x^2, x > 0$. The result follows by setting $x = \mathbb{E}[X]$.

■

We now illustrate the use of this approach by deriving the following result that implies the preceding result, but is stronger because it also establishes that the diameter of $G(n,p)$ is very small.

Theorem 8.4 (Almost every $G(n,p)$ has diameter 2). *For constant* p, $G(n,p)$ *is connected and has diameter* 2 a.a.s.

Proof: Let X be the number of (unordered) vertex pairs with no common neighbor. To prove the theorem, we need to show that $X = 0$ *a.a.s.*

We apply the first-moment method of Theorem 8.2 above. Let $X_{u,v}$ an indicator variable with $X_{u,v} = 1$ if u and v do not have a common neighbor, and $X_{u,v} = 0$ if they do.

For a vertex pair u, v, if $X_{u,v} = 1$, then none of the other $n-2$ vertices is adjacent to both u and v. Therefore, $\mathbb{P}\{X_{u,v} = 1\} = (1-p^2)^{n-2}$, and therefore $\mathbb{E}[X] = \mathbb{E}\left[\sum_{u,v} X_{u,v}\right] = \binom{n}{2}(1-p^2)^{n-2}$. This expression goes to zero with n for fixed p, establishing $\mathbb{P}\{X = 0\} \to 1$.

■

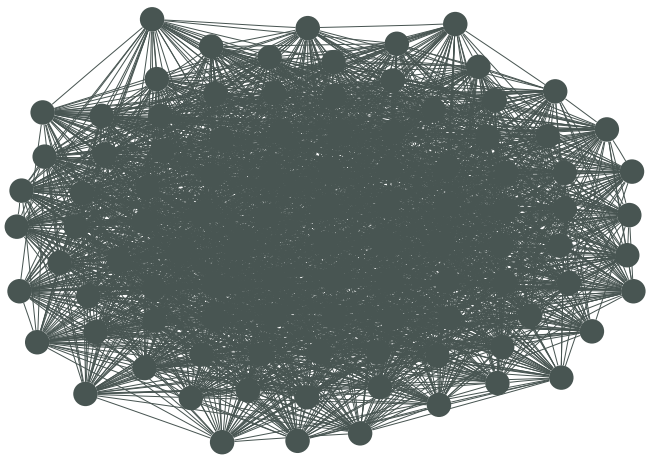

Figure 8.2: An instance of $G(100, 0.5)$, a very dense graph with diameter 2.

Definition 8.2 (Increasing property and threshold function). *An increasing property is a graph property conserved under the addition of edges. A function $t(n)$ is a threshold function for an increasing property if (a) $p(n)/t(n) \to 0$ implies that $G(n,p)$ does not possess the property* a.a.s.*, and if (b) $p(n)/t(n) \to \infty$ implies that it does* a.a.s.

Note that threshold functions are never unique; for example, if $t(n)$ is a threshold function, then so is $ct(n), c > 0$. Examples of increasing properties are:

1. A fixed graph H appears as a subgraph in G.
2. There exists a large component of size $\Theta(n)$ in G.
3. G is connected.
4. The diameter of G is at most d.

A counterexample is the appearance of H as an *induced* subgraph: indeed, the addition of edges in G can destroy this property.

Definition 8.3 (Balanced graph). *The ratio $2e(H)/|H|$ for a graph H is called its average vertex degree. A graph G is balanced if its average vertex degree is equal to the maximum average vertex degree over all its induced subgraphs.*

Note that trees, cycles, and complete graphs are all balanced.

Definition 8.4 (Automorphism group of a graph). *An automorphism of a graph G is an isomorphism from G to G, i.e., a permutation Π of its vertex set such that $(u,v) \in E(G)$ if and only if $(\Pi(u), \Pi(v)) \in E(G)$.*

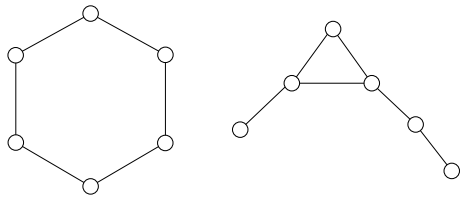

Figure 8.3: Two graphs with 6 vertices and 6 edges, the first with an automorphism group of size $a = 12$, the second with $a = 1$.

Lemma 8.1 (Chernoff bounds for Binomial RVs). *For $X \sim Binom(n,p)$,*

$$\begin{aligned} \mathbb{P}\left\{X \geq \mathbb{E}\left[X\right] + x\right\} &\leq \exp\left(-\frac{x^2}{2(np + x/3)}\right) \\ \mathbb{P}\left\{X \leq \mathbb{E}\left[X\right] - x\right\} &\leq \exp\left(-\frac{x^2}{2np}\right) \end{aligned} \tag{8.1}$$

We will also need the following theorem, which counts the number of different trees with n vertices.

Theorem 8.5 (Cayley's Formula). *There are n^{n-2} labeled trees of order n.*

8.3 Appearance of a subgraph

We now study the following problem: given an unlabeled graph H, what is the probability that this graph H is a subgraph of $G(n,p)$ when $n \to \infty$? This question has a surprisingly simple answer: we identify a threshold function for the appearance of H in $G(n,p)$ that only depends on the number of vertices and edges in H, with the caveat that H has to be balanced.

Theorem 8.6 (Threshold function for appearance of balanced subgraph). *For a balanced graph H with k vertices and l edges ($l \geq 1$), the function $t(n) = n^{-k/l}$ is a threshold function for the appearance of H as a subgraph of G.*

To be a bit more precise, a graph H appears in G if and only if there is at least one subgraph of G that is isomorphic to H.

Proof: The proof has two parts. In the first part, we show that when $p(n)/t(n) \to 0$, the H is not contained in $G(n,p)$ *a.a.s.*In the second part, we show that when $p(n)/t(n) \to \infty$, then the opposite is true.

Part 1: $p(n)/t(n) \to 0$.

Set $p(n) = o_n n^{-k/l}$, where o_n goes to zero arbitrarily slowly. Let X denote the number of subgraphs of G isomorphic to H.

We need to show that $\mathbb{P}\{X = 0\} \to 1$. Let A denote the set of labeled graphs H' isomorphic to H, whose vertex label set $[n]$ is the same as that of G.

$$|A| = \binom{n}{k} \frac{k!}{a} \leq n^k, \tag{8.2}$$

where a is the size of H's automorphism group.

$$\mathbb{E}[X] = \sum_{H' \in A} \mathbb{P}\{H' \subset G\}. \tag{8.3}$$

As H' is a labeled graph, the probability of it appearing in G is simply the probability that all edges of H' are present, i.e.,

$$\mathbb{P}\{H' \subset G\} = p^l. \tag{8.4}$$

$$\mathbb{E}[X] = |A|p^l \leq n^k p^l = n^k (o_n n^{-k/l})^l = o_n^l. \tag{8.5}$$

By the First Moment Method,

$$\mathbb{P}\{H \in G\} = \mathbb{P}\{X \geq 1\} \leq \mathbb{E}[X] \leq o_n^l \to 0. \tag{8.6}$$

Therefore, H does not appear in $G(n,p)$ *a.a.s.*

Part 2: $p(n)/t(n) \to \infty$.

Set $p(n) = \omega_n n^{-k/l}$, where ω_n goes to ∞ arbitrarily slowly.

We need to show that $\mathbb{P}\{X = 0\} \to 0$. For this, we bound the variance of X.

$$\mathbb{E}\left[X^2\right] = \mathbb{E}\left[\left(\sum_{H' \in A} \mathbf{1}_{\{H' \subset G\}}\right)^2\right] = \sum_{(H', H'') \in A^2} \mathbb{P}\{H' \cup H'' \subset G\}. \tag{8.7}$$

As the labeled graph $H' \cup H''$ has $2l - e(H' \cap H'')$ links, so $\mathbb{P}\{H' \cup H'' \subset G\} = p^{2l-e(H' \cap H'')}$.

As H is balanced, we know that any subgraph (induced or not) of H, including $H' \cap H''$, has $e(H' \cap H'')/|H' \cap H| \leq e(H)/|H| = l/k$. Therefore, if $|H' \cap H''| = i$, then $e(H' \cap H'') \leq il/k$. We partition the set A^2 into classes A_i^2 with identical order of the intersection, i.e.,

$$A_i^2 = \{(H', H'') \in A^2 : |H' \cap H''| = i\} \tag{8.8}$$

$$S_i = \sum_{(H', H'') \in A_i^2} \mathbb{P}\{H' \cup H'' \subset G\} \tag{8.9}$$

We will show that $\mathbb{E}[X]$ is dominated by $i = 0$, i.e., H' and H'' are disjoint. In this case, the events $\{H' \in G\}$ and $\{H'' \in G\}$ are independent, as they have no edges in common. Thus,

$$\begin{aligned} S_0 &= \sum_{(H', H'') \in A_0^2} \mathbb{P}\{H' \cup H'' \subset G\} \\ &= \sum_{(H', H'') \in A_0^2} \mathbb{P}\{H' \subset G\}\, \mathbb{P}\{H'' \subset G\} \qquad (H' \text{ and } H'' \text{ disjoint}) \\ &\leq \sum_{(H', H'') \in A^2} \mathbb{P}\{H' \subset G\}\, \mathbb{P}\{H'' \subset G\} \\ &= (\mathbb{E}[X])^2. \end{aligned} \tag{8.10}$$

We now examine the contribution to $\mathbb{E}\left[X^2\right]$ of the terms $i \geq 1$. For this, note that for a fixed graph H', the number of H'' such that $|H' \cap H''| = i$ is given by

$$\left(\begin{array}{c} k \\ i \end{array} \right) \left(\begin{array}{c} n-k \\ k-i \end{array} \right) \frac{k!}{a}, \tag{8.11}$$

as we need to select i nodes from H' to form the intersection, and $k - i$ nodes from the vertices outside H' for the rest; then there are $k!/a$ labelings for H'' isomorphic to H. Also note that it is easy to see that this expression is $O(n^{k-i})$.

We now use this to compute S_i.

$$\begin{aligned}
S_i &= \sum_{(H',H'')\in A_i^2} \mathbb{P}\{H' \cup H'' \subset G\} \\
&= \sum_{H'\in A} \sum_{H''\in A:|H'\cap H''|=i} \mathbb{P}\{H' \cup H'' \subset G\} \\
&\leq \sum_{H'} \binom{k}{i}\binom{n-k}{k-i} \frac{k!}{a} p^{2l} p^{-il/k} \quad (\text{as } e(H\cap H'') \leq il/k) \\
&= |A| \binom{k}{i}\binom{n-k}{k-i} \frac{k!}{a} p^{2l} (\omega_n n^{-k/l})^{-il/k} \\
&\leq |A| p^l c_1 n^{k-i} \frac{k!}{a} p^l \omega_n^{-il/k} n^i \quad (\text{because } \binom{k}{i}\binom{n-k}{k-i}\tfrac{k!}{a} = O(n^{k-i})) \\
&= \mathbb{E}[X]\, c_1 n^k \frac{k!}{a} p^l \omega_n^{-il/k} \quad (\text{using (8.5)}) \\
&\leq (\mathbb{E}[X]) c_2 \binom{n}{k} \frac{k!}{a} p^l \omega_n^{-il/k} \quad (\text{because } n^k = \Theta\binom{n}{k}) \\
&= (\mathbb{E}[X])^2 c_2 \omega_n^{-il/k} \quad (\text{using (8.2)}) \\
&\leq (\mathbb{E}[X])^2 c_2 \omega_n^{-l/k} \qquad (8.12)
\end{aligned}$$

for n large enough.

$$\mathbb{E}\left[X^2\right]/(\mathbb{E}[X])^2 = S_0/(\mathbb{E}[X])^2 + \sum_{i=1}^{k} S_i/(\mathbb{E}[X])^2 \leq 1 + kc_2\omega_n^{-l/k}, \qquad (8.13)$$

and therefore $\mathrm{Var}[X]/\mathbb{E}[X]^2 \to 0$. Therefore, by Theorem 8.3, $X > 0$ *a.a.s.*

■

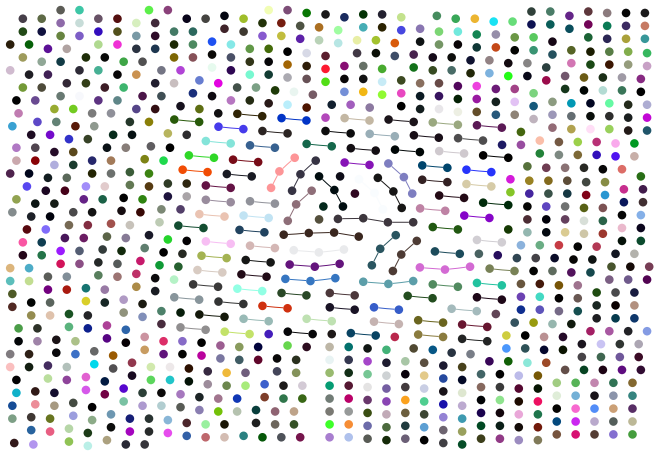

Figure 8.4: An instance of $G(1000, 0.2/1000)$. The graph consists only of small trees.

Corollary 8.1 (Appearance of trees of order k). *The function $t(n) = n^{-k/k-1}$ is a threshold function for the appearance of trees of order k.*

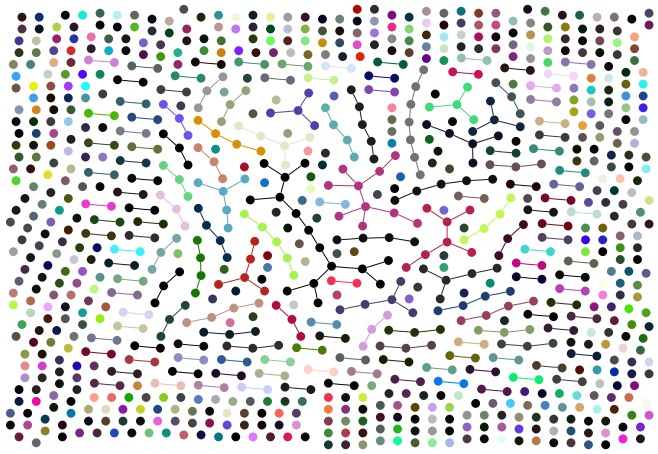

Figure 8.5: An instance of $G(1000, 0.5/1000)$. The graph still only consists of trees, but trees of higher order are appearing.

Proof: A tree of order k has k nodes and $k-1$ edges, and it is balanced. The result follows directly from Theorem 8.6 and the fact that there are finitely many trees of order k.

■

Corollary 8.2 (Appearance of cycles of all orders). *The function $t(n) = 1/n$ is a threshold function for the appearance of cycles of any fixed order.*

Corollary 8.3 (Appearance of complete graphs). *The function $t(n) = n^{-2/(k-1)}$ is a threshold function for the appearance of the complete graph K_k with fixed order k.*

8.4 The giant component

After studying the class of threshold functions of the form $n^{-k/l}$ for the appearance of subgraphs, we now focus in more detail on $p(n) = c/n$. Note that this is the threshold function for the appearance of cycles of all orders, which suggests that something special is happening for this function $p(n)$.

We set $p(n) = c/n$, and study the structure of $G(n, p)$ as a function of c. Specifically, we consider the set of components and their sizes that make up $G(n, p)$. As it turns out, a phase transition occurs at $c = 1$: when c goes from $c < 1$ to $c > 1$, the largest component jumps from $\Theta(\log n)$ to $\Theta(n)$ (this actually occurs as a "double jump", which we do not consider in more detail); the largest component for $c > 1$ is unique.

Let C_v denote the component that a vertex v belongs to.

Theorem 8.7 (Small components for $c < 1$). *If $c < 1$, then the largest component of $G(n, p)$ has at most*

$$\frac{3}{(1-c)^2} \log n \tag{8.14}$$

vertices a.a.s.

Proof: Let $c = 1 - \epsilon$. We consider a vertex v in $G(n, p)$ and study the following process to successively discover all the vertices of the component that v belongs to. Let the set A_i denote

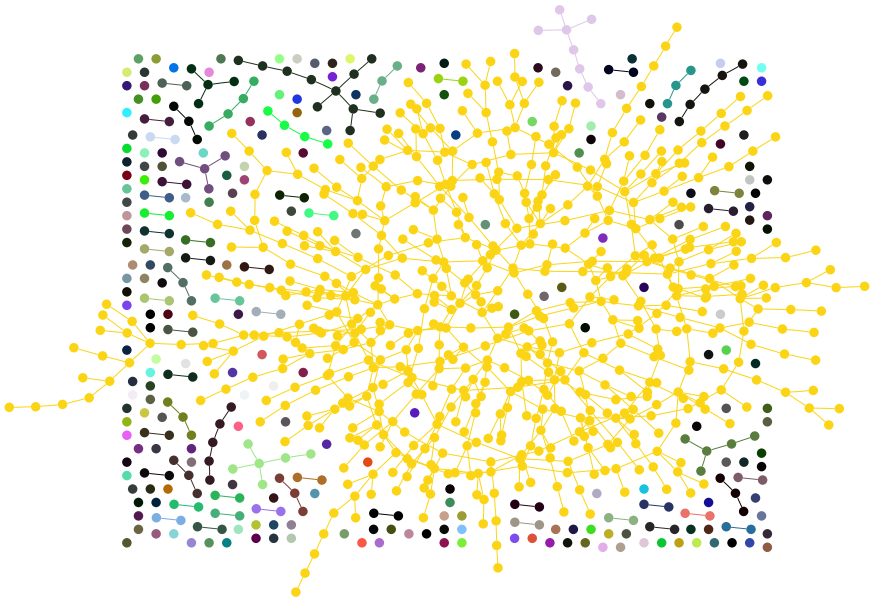

Figure 8.6: An instance of $G(1000, 1.5/1000)$, slightly above the threshold for the appearance of the giant cluster.

active vertices, and the set S_i denote *saturated* vertices, with $A_0 = \{v\}$, and $S_0 = \emptyset$. At the ith step, we select an arbitrary vertex u from A_i. We move u from the active to the saturated set, and mark all neighbors of u that have not been touched yet as active. In this manner, we touch all the vertices in v's component until $A_i = \emptyset$, which is equivalent to $|C_v| = |S_i| = i$.

Let $Y_i = |A_i \cup S_i|$ denote the total number of vertices visited by step i, and define $T = \min\{i : Y_i = i\}$, i.e, we have visited all nodes and $A_i = \emptyset$. Then Y_i is a Markov chain with $Y_{i+1} - Y_i \sim \text{Binom}(n - Y_i, p)$, because an edge exists from u to each of the $n - Y_i$ vertices not in $A_i \cup S_i$ independently with probability p, and T is a stopping time for this Markov chain.

Note that we can stochastically upper bound the process (Y_i) with a random walk (Y_i^+) with increments $X_i \sim \text{Binom}(n, p)$. The corresponding stopping time T^+ for the random walk stochastically dominates T.

We want to bound the probability that vertex v belongs to a component of size at least k. As $|C_v| \geq k \Leftrightarrow |A_k \cup S_k| \geq k$,

$$\mathbb{P}\{|C_v| \geq k\} = \mathbb{P}\{T \geq k\} \leq \mathbb{P}\{T^+ \geq k\} \leq \mathbb{P}\left\{\sum_{i=0}^{k} X_i \geq k\right\}. \tag{8.15}$$

The random walk has $Y_k^+ \sim B(kn, p)$. Using the Chernoff bound (Lemma 8.1) for the binomial distribution and setting $k = (3/\epsilon^2)\log n$, we find

$$\begin{aligned}
\mathbb{P}\left\{\max_v |C_v| \geq k\right\} &\leq n\mathbb{P}\left\{Y_k^+ \geq k-1\right\} \\
&= n\mathbb{P}\left\{Y_k^+ \geq ck + \epsilon k - 1\right\} \\
&\leq n\exp\left(-\frac{(\epsilon k-1)^2}{2(ck+\epsilon k/3)}\right) \\
&\leq n\exp\left(-\frac{\epsilon^2}{2}k\right) = n^{-1/2} = o(1). \qquad (8.16)
\end{aligned}$$

■

Theorem 8.8 (Unique giant component for $c > 1$). *If $c > 1$, then the largest component of $G(n,p)$ has $\Theta(n)$ vertices, and the second-largest component has $O(\log n)$ vertices* a.a.s.

Proof: We will study the same process as in the previous proof, starting at an arbitrary vertex v. The proof has three parts. In the first part, we show that it is very likely that the process either dies out early, i.e., $T \leq a_n$ *a.a.s.*, resulting in a small component, or continues for at least b_n steps, resulting in a large component. In the second part, we show that there is only one large component with $k \geq b_n$. In the third part, we confirm that the size of the largest component is of order n.

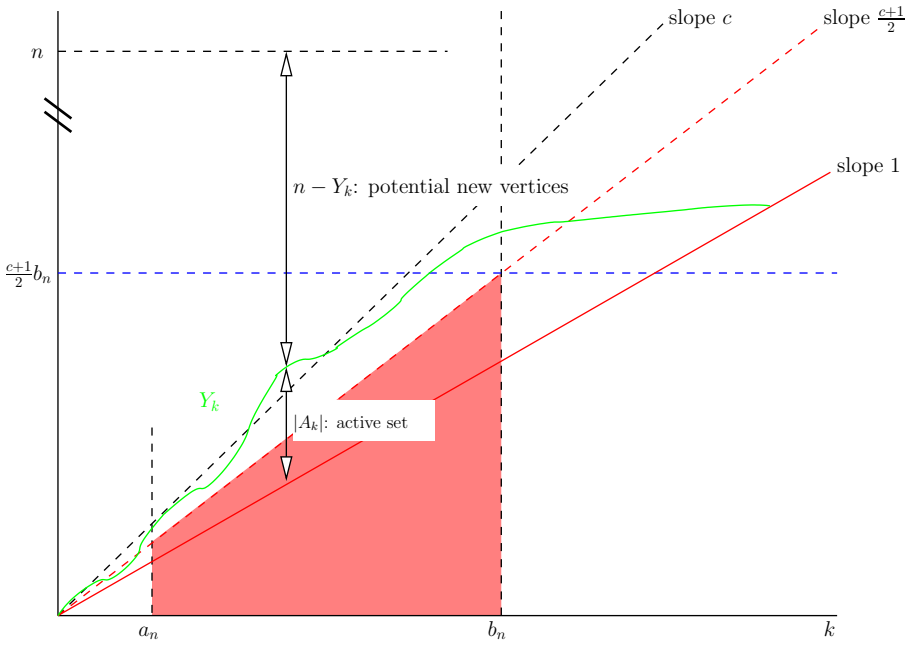

Figure 8.7: An illustration of the different variables involved in the proof for $c > 1$.

Part 1: each component is either small or large. Let $c = 1+\epsilon$, $a_n = \frac{16c}{\epsilon^2}\log n$, and $b_n = n^{2/3}$. We wish to show that the Markov chain Y_k either dies out before a_n (i.e., $T < a_n$), or that for any $a_n < k < b_n$, we have a large number of active nodes A_k left to continue the process, specifically $(\epsilon/2)k$ nodes.

The event that we have many vertices left at stage k satisfies

$$\left\{|A_k| \geq \frac{c-1}{2}k\right\} = \left\{Y_k \geq \frac{c+1}{2}k\right\}. \qquad (8.17)$$

Define the stopping time $T_b = \min\{i : Y_i \geq \frac{c+1}{2} b_n\}$. After time T_b, the condition on the size of the active set is satisfied until at least time b_n. Then

$$\begin{aligned}
\mathbb{P}\left\{|A_k| \geq \frac{c-1}{2}k\right\} &= \mathbb{P}\{T_b \leq k\} + \mathbb{P}\left\{Y_k \geq \frac{c+1}{2}k \middle| T_b > k\right\}\mathbb{P}\{T_b > k\} \\
&= \mathbb{P}\{T_b \leq k\}\left(1 - \mathbb{P}\left\{Y_k \geq \frac{c+1}{2}k \middle| T_b > k\right\}\right) + \mathbb{P}\left\{Y_k \geq \frac{c+1}{2}k \middle| T_b > k\right\} \\
&\geq \mathbb{P}\left\{Y_k \geq \frac{c+1}{2}k \middle| T_b > k\right\}. \qquad (8.18)
\end{aligned}$$

Fix $a_n \leq k \leq b_n$ and a starting vertex v. Conditional on $\{T_b > k\}$, there remain at least $n - \frac{c+1}{2}b_n$ untouched vertices. Therefore, to bound the second term in (8.18), we can stochastically lower-bound Y_i with a random walk Y_i^- with increments $X_i^- \sim \text{Binom}(n - \frac{c+1}{2}b_n, p)$. Therefore,

$$\mathbb{P}\left\{Y_k \geq \frac{c+1}{2}k\right\} \geq \mathbb{P}\left\{Y_k^- \geq \frac{c+1}{2}k\right\} \qquad (8.19)$$

Using this bound, we find

$$\begin{aligned}
\mathbb{P}\{\text{any component has size in } (a_n, b_n)\} &\leq n\sum_{k=a_n}^{b_n} \mathbb{P}\left\{Y_k^- < k + \frac{\epsilon}{2}k\right\} \\
&\leq n\sum_{k=a_n}^{b_n} \exp\left(-\frac{\epsilon^2 k^2}{9ck}\right) \\
&\leq nb_n \exp\left(-\frac{\epsilon^2}{9c}a_n\right) = o(1). \qquad (8.20)
\end{aligned}$$

Part 2: large component is unique. We now show that the largest component is unique, by considering two vertices u and v that both belong to a component of size larger than b_n, and showing that the probability that they lie in different components in asymptotically small.

Assume that we run the above process starting from u and from v. We had shown in Part 1 that starting at v, the set $A_{b_n}(v)$ will be of size at least $\epsilon b_n/2$. The same holds for the set of active vertices starting at u.

Now assume that the two processes have not "touched" yet, i.e., have no vertices in common. The probability that they touch at a later stage (after b_n) is larger than the probability that they touch in the next step, i.e., that there exists at least one vertex that is adjacent to both active sets $A_{b_n}(u), A_{b_n}(v)$.

$$\begin{aligned}
\mathbb{P}\{\text{processes do not touch in next step}\} &= (1-p)^{|A_{b_n}(u)||A_{b_n}(v)|} \\
&\leq (1-p)^{(\epsilon b_n/2)^2} \\
&\leq \exp\left(-\frac{\epsilon^2}{4}cn^{1/3}\right) = o(n^{-2}). \qquad (8.21)
\end{aligned}$$

Taking the union bound over all pairs of vertices (u, v) shows that the probability that any two vertices in giant components lie in different giant components goes to zero, i.e., the giant component is unique *a.a.s.*

Part 3: large component has size $\Theta(n)$. Recall that $b_n = n^{2/3}$. Therefore, to show that the unique giant component is of size $\Theta(n)$, we need to show that all the other vertices only make up at most

a constant fraction of all the vertices. For this, we consider a vertex v, and find an upper bound to the probability that it belongs to a small component.

Let N be the number of vertices in small components. Then the size of the giant component is $n - N$. By definition, each small component is smaller than a_n. The probability ρ that C_v small, i.e., that the process dies before a_n vertices have been reached, is smaller than $\rho^+ = \mathbb{P}\{\text{BP}(\text{Binom}(n - a_n, p)) \text{ dies}\}$, and larger than $\rho^- = \mathbb{P}\{\text{BP}(\text{Binom}(n, p)) \text{ dies}\} - o(1)$, where the $o(1)$ term corresponds to the probability that the process dies too late (after more than a_n vertices have been discovered).

Note that $\text{Binom}(n - o(n), c/n) \to \text{Poisson}(c)$ in distribution. The probability that the process dies out before a_n vertices have been reached is asymptotically equal to $\mathbb{P}\{\text{BP}(\text{Poisson}(c)) \text{ dies}\}$, which is given by $\rho = 1 - \beta$, with $\beta < 1$ the solution of $\beta + e^{-\beta c} = 1$.

Therefore,

$$\mathbb{P}\{v \text{ in small component}\} \to \rho, \tag{8.22}$$

and $\mathbb{E}[N]/n \to \rho$.

To show the result, we need to show that $\mathbb{E}\left[N^2\right] \to (\mathbb{E}[N])^2$ and invoke Lemma 8.3 to prove the result. For this, write

$$\begin{aligned}
\mathbb{E}\left[N^2\right] &= \mathbb{E}\left[\left(\sum_v \mathbf{1}_{\{v \text{ in small component}\}}\right)^2\right] \\
&= \sum_{u,v} \mathbb{P}\{C_u, C_v \text{ both small}\} \\
&= n\rho\mathbb{P}\{u \in C_v | C_v \text{ small}\} + \\
&\quad n(n-1)\rho\mathbb{P}\{C_u \text{ small}, C_u \neq C_v | C_v \text{ small}\} \\
&\leq n\rho^2 a_n + n^2\rho^2,
\end{aligned} \tag{8.23}$$

and therefore $\mathbb{E}\left[N^2\right] \leq (1 + o(1))n^2\rho^2$. This completes the proof.

■

8.5 Connectivity

We have seen that a unique giant cluster appears around $np = 1$. As in percolation, it is much harder to achieve full connectivity, such that the graph possesses a single component encompassing all vertices. We will now show that this happens for $p(n) = \log n/n$, i.e., when the average node degree hits $\log n$.

It is also interesting to understand what happens between the threshold probability for the giant cluster and the threshold for full connectivity. In fact, we will show that as we increase $p(n)$, the giant cluster consumes the remaining smaller clusters in descending order. The small clusters are in fact small trees, and there actually are threshold functions for the disappearance of trees of a given order between the thresholds for giant cluster and that for full connectivity.

Just before we hit full connectivity, the only remaining small components are isolated vertices. Theorem 8.10 below shows that $t(n) = \log n/n$ is a threshold function for $G(n, p)$ to be connected.

Theorem 8.9. *The function $t(n) = \log n/n$ is a threshold function for the disappearence of isolated vertices in $G(n, p)$.*

Proof: Let X_i denote the indicator for vertex i to be isolated, and X is the sum of all X_i. We have $\mathbb{E}[X] = n(1 - p)^{n-1}$.

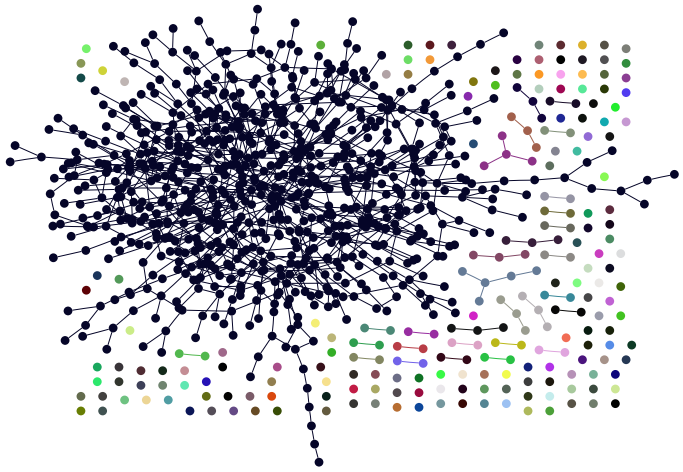

Figure 8.8: An instance of $G(1000, 2/1000)$, in between the critical average degree for the giant cluster and for full connectivity.

First, let $p(n) = \omega_n \log n/n$, with $\omega_n \to \infty$. We have

$$\begin{aligned} \mathbb{E}[X] &\leq ne^{-\omega_n \log n} \\ &= n^{1-\omega_n} \to 0. \end{aligned} \tag{8.24}$$

By the First Moment Method, there are *a.a.s.*no isolated vertices.

Second, let $p(n) = o_n \log n/n$, with $o_n \to 0$.

$$\begin{aligned} \mathbb{E}\left[X^2\right] &= \mathbb{E}\left[\sum_{i,j} X_i X_j\right] \\ &= \mathbb{E}[X] + \mathbb{E}\left[\sum_{i \neq j} X_i X_j\right] \\ &= \mathbb{E}[X] + n(n-1)(1-p)^{2(n-2)+1} \\ &\sim \mathbb{E}[X] + n^{2(1-o_n)} \sim \mathbb{E}[X] + \mathbb{E}[X]^2 . \end{aligned} \tag{8.25}$$

As $\mathbb{E}[X] \to \infty$, this shows that $\mathbb{E}\left[X^2\right] \sim \mathbb{E}[X]^2$, proving the result through the second-moment method.

■

Note that this result can actually easily be sharpened by setting $p(n) = c \log n/n$, and studying the phase transition as a function of c, in analogy to the way the giant component appeared.

Theorem 8.10. *The function $t(n) = \log n/n$ is a threshold function for connectivity in $G(n,p)$.*

Proof: Theorem 8.9 shows immediately that if $p(n) = o(t(n))$, then the graph is not connected, because it still has isolated vertices.

To show the converse, set $p(n) = \omega_n \log n/n$, with $\omega_n \to \infty$ arbitrarily. We know that the RG does not contain isolated vertices. We now show that it does not contain any other small components either, i.e., specifically components of size k at most $n/2$.

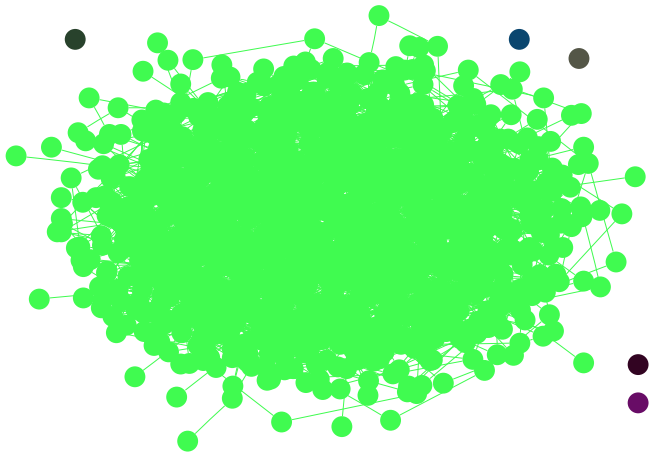

Figure 8.9: An instance of $G(1000, 5/1000)$, slightly below the threshold for full connectivity. The only remaining small components are isolated vertices.

We bound the probability that a small component of size between $2 \leq k \leq n/2$ appears. The case $k = 2$ is left as an exercise. For $k > 2$, we note that such a component contains necessarily a tree of order k, which means it contains at least $k - 1$ edges, and none of the the $k(n - k)$ possible edges to other vertices exists.

$$
\begin{aligned}
&\mathbb{P}\{G(n,p) \text{ contains component of order } k\} \\
&\leq \binom{n}{k} k^{k-2} p^{k-1} (1-p)^{k(n-k)} \\
&\leq k^{-2} \exp\left[k(\log n + 1) + (k-1)\left(\log(\omega_n \log n) - \log n\right) - k\omega_n \log n + \frac{k^2}{n}\omega_n \log n\right] \\
&\leq nk^{-2} \exp\left[k + k\log(\omega_n \log n) - \frac{1}{2}k\omega_n \log n\right] \\
&\leq nk^{-2} \exp\left[-(1/3)\omega_n k \log n\right] \qquad \text{(for } n \text{ large enough)} \\
&= k^{-2} n^{1-k\omega_n/3}, \qquad (8.26)
\end{aligned}
$$

using $\binom{n}{k} = (n)_k/(k)_k \leq n^k/(k/e)^k = (ne/k)^k$, $\log(1-p) \leq -p$.

$$
\begin{aligned}
&\mathbb{P}\{G(n,p) \text{ contains components of order } 3 \leq k \leq n/2\} \\
&\leq \sum_{k=2}^{n/2} k^{-2} n^{1-k\omega_n/3} = O(n^{2-2\omega_n/3}) = o(1). \qquad (8.27)
\end{aligned}
$$

Therefore, the RG contains no components smaller than $n/2$.

■

The following figures illustrate the evolution between $p(n) = 1/n$ and $p(n) = \log n/n$. Figure 8.8 shows an instance of $G(1000, 0.002)$, roughly halfway between the two thresholds.

Figures 8.9 and 8.10 show instances of $G(n,p)$ just below and just above the threshold (of approx. 6.9/1000) for full connectivity.

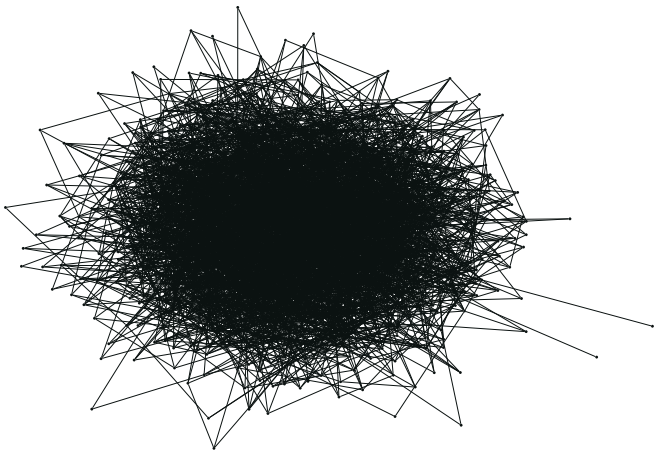

Figure 8.10: An instance of $G(1000, 8/1000)$, slightly above the threshold for full connectivity.

9

Random Regular Graphs

9.1 Introduction

Another model for random graphs is the random regular graph $G(n, r)$, in which every vertex has degree r. In contrast to the previous model $G(n, p)$, the existence of different edges is not independent, and this leads, not surprisingly, to some additional difficulties in the analysis. Even defining the probability space is not as straightforward as before: we would like to assign the same probability to every labelled graph over vertex set $[n]$ with degree r. Our first concern in this chapter will be to analyze a relaxation of this model, where we sample from a larger class of graphs $G^*(n, r)$, by allowing for the possibility of loops and multiple edges.

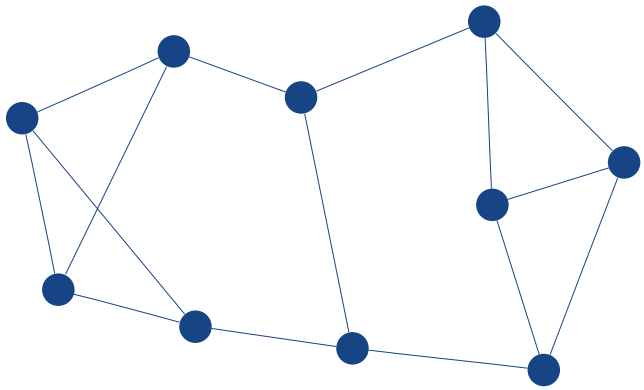

Figure 9.1: An instance of $G(10, 3)$.

9.2 Preliminaries

Theorem 9.1 (Method of moments for Poisson RVs, factorial moments). *Let $(X_{n1}, \dots, X_{nl})$ be vectors of random variables, where $l \geq 1$ is fixed. If $\lambda_1, \dots, \lambda_l \geq 0$ are such that*

$$\mathbb{E}\left[(X_{n1})_{m_1} \dots (X_{nl})_{m_l}\right] \to \lambda_1^{m_1} \dots \lambda_l^{m_l} \tag{9.1}$$

for every $m_1, \dots, m_l \geq 0$, then $(X_{n1}, \dots, X_{nl}) \to (Z_1, \dots, Z_l)$, where $Z_i \sim Poisson(\lambda_i)$ are independent.

Definition 9.1 (Random regular graph). *Consider the set $\mathcal{G}(n, r)$ of all labeled r-regular graphs of order n, i.e., the set of labeled (simple) graphs with vertex label set $[n]$ and constant degree r. Then the random graph $G(n, r)$ is a uniform random element of $\mathcal{G}(n, r)$.*

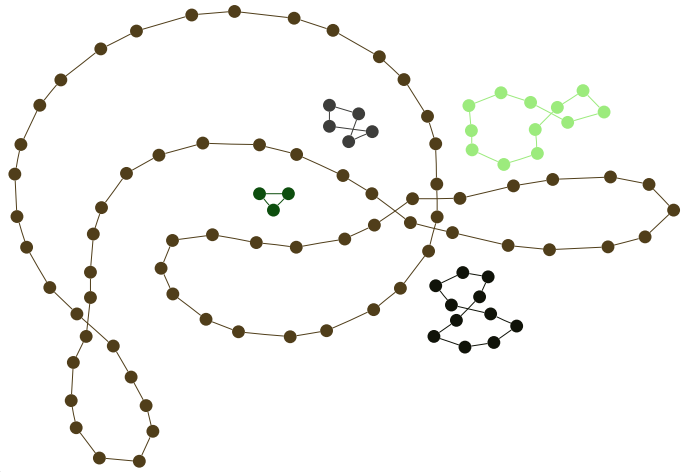

Figure 9.2: An instance of $G(100, 2)$. The graph $G(n, 2)$ is *a.a.s.*not connected (we do not prove this here), and is the union of several disjoint cycles.

9.3 The pairing model $G^*(n, r)$

Consider the set of *stubs* $[n] \times [r]$. Think of a stub as one endpoint of a potential edge. To generate a random regular multigraph $G^*(n, r)$, we generate a *pairing* (or matching) of the nr stubs, which results into $nr/2$ edges. In the pairing model, an edge should be thought of as labeled, i.e., a tuple $((u, i), (v, j))$, where $u, v \in [n]$, $i, j \in [r]$.

The random regular multigraph $G^*(n, r)$ is obtained by projecting a pairing, which simply corresponds to removing the stub labels. An edge $((u, i), (v, j))$ in the pairing corresponds to an edge (u, v) in the graph. This implies that $G^*(n, r)$ is not a simple graph, because it can have loops and multiple edges.

The pairing model is interesting because (a) if we condition on the projection of a random matching to be a simple graph, then that graph is sampled uniformly from $G(n, r)$; furthermore, the probability that it is simple is bounded away from zero; (b) is easier to handle to prove many properties of interest in $G^*(n, r)$, and (c) because a property that holds *a.a.s.*for $G^*(n, r)$ also holds *a.a.s.*for $G(n, r)$, as we will establish in the next two sections.

Let $S = [n] \times [r]$ denote the set of stubs. There are $(nr-1)!! = (nr-1)(nr-3)\cdots 3$ distinct pairings.

If we condition on G^* being a simple graph, then each element of $\mathcal{G}(n,r)$ is equally probable. This is because each element of $\mathcal{G}(n,r)$ corresponds to the same number $(r!)^n$ of distinct configurations. Note that this holds only conditional on G^* being simple; unconditionally, graphs with loops and/or multiple edges appear with smaller probability than a simple graph.

9.4 Appearance of a fixed subgraph in $G^*(n,r)$

The family of random variables $(Z_k)_{k=1}^{\infty}$ denotes the number of cycles of order k in $G^*(n,r)$.

Theorem 9.2 (Convergence in distribution of number of all cycles in $G^*(n,r)$). *The random variables $(Z_k), k \geq 1$ converges in distribution to a collection of independent random variables, with $Z_k \to Poisson((r-1)^k/2k)$.*

Proof: We view a sample of $G^*(n,r)$ as the projection of a random pairing. A labeled cycle of k labeled edges in the pairing corresponds to a cycle of order k in G^*. We will use this correpondence to compute the moments of Z_k, and to use the method of moments to establish convergence in distribution.

First we need the probability p_k that a set of k labeled edges is in a random pairing.

$$p_k = \frac{(rn-2k-1)!!}{(rn-1)!!} = \frac{1}{(rn-1)(rn-3)\ldots(rn-2k+1)}, \tag{9.2}$$

because each labeled edge blocks two stubs, which leaves $(rn-2k-1)!!$ configurations with the k labeled edges fixed.

Expectation of number of k-cycles. We count the number of ways a k-cycles can appear in G^*. As in Theorem 8.6, the number of distinct vertex-labeled k-cycles is

$$\left(\begin{array}{c} n \\ k \end{array} \right) \frac{k!}{2k}, \tag{9.3}$$

where $2k$ is the size of the automorphism group of the k-cycle.

Each edge so obtained has two distinct labels from $[r]$ in the pairing, for a total of

$$\left(\begin{array}{c} n \\ k \end{array} \right) \frac{k!}{2k}(r(r-1))^k \sim \frac{(nr(r-1))^k}{2k}. \tag{9.4}$$

For large n we also have $p_k \sim (rn)^{-k}$, and therefore $\mathbb{E}[Z_k] \sim (r-1)^k/2k$.

Expectation of number of other graphs H. Note that a similar argument shows that the expected number of copies of a graph is in general $\Theta(n^{v(H)-e(H)})$.

This is important for the following reason. We will study higher-order moments next, which amounts to counting the number of copies of graphs H where H is the union of several cycles. If all these cycles are disjoint, then $v(H) = e(H)$. Otherwise, H contains at least one component which is the union of several intersecting cycles; H then has $v(H) < e(H)$.

Second factorial moment. Before studying higher-order joint moments in their full generality, we consider the second factorial moment $\mathbb{E}[(Z_k)_2]$. Note that $(Z_k)_2$ is the number of ordered pairs of distinct k-cycles in G. We can express this number as a sum of two terms S_0 and $S_>$, where S_0 counts the number of ordered pairs of two distinct *disjoint* k-cycles, and where $S_>$ counts the number of ordered pairs of two intersecting k-cycles. We now show that S_0 asymptotically dominates.

Similar to Theorem 8.6, we can express $S_>$ as a sum of terms $S_{i,j}$ according to the number of vertices and edges (i,j) in the intersection between the two k-cycles. Obviously, the number of terms does not depend on n.

Each $S_{i,j}$ counts the number of copies of an unlabeled graph $H_{i,j}$, which is the union of two intersecting k-cycles, and thus $v(H) < e(H)$. Therefore, $S_>$ is $O(n^{v(H)-e(H)}) = o(1)$.

To compute S_0,

$$\left(\begin{array}{c} n \\ k \end{array} \right) \frac{k!}{2k} \left(\begin{array}{c} n-k \\ k \end{array} \right) \frac{k!}{2k} (r(r-1))^{2k} \left(\frac{n^k r^k (r-1)^k}{2k} \right)^2 . \tag{9.5}$$

Combining this with $p_k \sim (rn)^{-k}$, we obtain, as needed, $\mathbb{E}\left[(Z_k)_2\right] \to \lambda^2$.

Higher-order moments of number of k-cycles. We now generalize this argument to higher-order and joint moments, of the form

$$\mathbb{E}\left[(Z_1)_{m_1}(Z_2)_{m_2}\dots(Z_l)_{m_l}\right]. \tag{9.6}$$

Now H denotes an unlabeled graph resulting from the union of m_1 1-cycles, m_2 2-cycles, etc. A similar argument as before shows that all the terms corresponding to H where not all cycles are disjoint go to zero. The sum S_0 is easily shown to factor, so that

$$\mathbb{E}\left[(Z_1)_{m_1}(Z_2)_{m_2}\dots(Z_l)_{m_l}\right] \to \lambda_1^{m_1}\dots\lambda_l^{m_l}. \tag{9.7}$$

Theorem 9.1 shows the result.

■

9.5 The random regular graph $G(n,r)$

We can now go back to the original model $G(n,r)$. We first study the same random variables $(Z_k)_{k=1}^{\infty}$, the number of cycles of order k, but in $G(n,r)$ instead of in $G^*(n,r)$. Obviously, this forces $Z_1 = Z_2 = 0$. The following theorem is then a direct consequence of Theorem 9.2, by conditioning on $Z_1 = Z_2 = 0$.

Theorem 9.3 (Convergence in distribution of number of all cycles in $G(n,r)$). *The random variables $(Z_k), k \geq 3$ converge in distribution to a collection of independent random variables, with $Z_k \to$ Poisson$((r-1)^k/2k)$.*

We can now also show that the probability that $G^*(n,r)$ is simple is bounded away from zero. This suggests an efficient way of generating $G(n,r)$, by simply generating random pairings until a simple graph is found. Note however that the probability of success decreases quite quickly with r.

Theorem 9.4 (Probability that random regular multigraph is simple). $\mathbb{P}\{G^*(n,r)$ *is simple*$\} \to \exp\left(-\frac{r^2-1}{4}\right)$.

Proof: $\mathbb{P}\{G^*(n,r)$ is simple$\} = \mathbb{P}\{Z_1 = Z_2 = 0\}$.

■

Theorem 9.5 (Almost sure property of $G^*(n,r)$ carries over to $G(n,r)$). *Any property Q that holds* a.a.s. *for $G^*(n,r)$ also holds* a.a.s. *for $G(n,r)$.*

Figure 9.3: An instance of $G(100, 3)$.

Proof:

$$\begin{aligned} \mathbb{P}\{G \text{ does not have } Q\} &= \mathbb{P}\{G^* \text{ does not have } Q | G^* \text{ is simple}\} \\ &= \frac{\mathbb{P}\{G^* \text{ does not have } Q, G^* \text{ is simple}\}}{\mathbb{P}\{G^* \text{ is simple}\}} \\ &\leq \frac{\mathbb{P}\{G^* \text{ does not have } Q\}}{\mathbb{P}\{G^* \text{ is simple}\}} \to 0. \end{aligned} \tag{9.8}$$

■

This theorem allows us to try to prove properties of interest of the model G^*; the model G then has the same property. The converse is not true; for example, G has no cycles of order 1 and 2, but G^* does.

9.6 Connectivity of $G(n, r)$

A random regular graph $G(n, r)$ is connected *a.a.s.*for $r \geq 3$. While this might seem surprising at first when we compare with $G(n, p)$, where connectivity required about $\ln n$ average vertex degree, recall that the main challenge there was to eliminate isolated vertices; the majority of vertices have already been connected at a much lower average vertex degree of $c = np > 1$. In this sense, we should not be surprised that a constant r is enough to ensure connectivity in $G(n, r)$, as isolated vertices are a-priori impossible in this model.

We will in fact show a much stronger result, which is that $G(n, r)$ is *a.a.s.*r-connected, which means that there are (at least) k vertex-disjoint paths connecting any pair of vertices.

We note in passing (without proof) that $G(n, 2)$ is not connected *a.a.s.*, and consists of a collection of cycles, and $G(n, 1)$ is simply a random matching.

Theorem 9.6 (Connectivity of $G(n, r)$). *For $r \geq 3$, $G(n, r)$ is r-connected* a.a.s..

Proof: We partition the set of vertices into three sets A, S, and B. If there are no edges between A and B, we say that S *separates* A and B. The graph is r-connected if and only if the smallest

set S that separates the graph is of order at least r. We denote by T the subset of vertices of S adjacent to a vertex in A. Let H be the subgraph spanned by $A \cup T$.

Small component. Fix an arbitrarily large natural number a_0. We first consider a small component A, i.e., of fixed size $a = |A| < a_0$. For $a = 1$, the assertion is immediate. For $a = 2$, $A = \{u, v\}$, we distinguish two cases. If there is no edge (u, v), there are r edges incident to u that need to go to distinct vertices in T. If there is an edge (u, v), then there can be at most one vertex in T adjacent to both u and v, as otherwise $v(H) < e(H)$, which implies that H does not appear *a.a.s.*(c.f. the proof of Theorem 9.2).

For $a > 2$, we lower-bound the size t of T, and therefore the size of S. The subgraph H contains $a + t$ vertices and at least $(ra + t)/2$ edges, because there are by definition ar stubs in A and t stubs in T in the spanned subgraph H (recall that every vertex in T has at least one edge into A).

Therefore, to ensure $v(H) \geq e(H)$,

$$s \geq t \geq a(r - 2). \tag{9.9}$$

This shows that $s \geq r$ *a.a.s.*for a fixed $a \geq 3$ and $r \geq 3$, and therefore this holds over all $3 \leq a \leq a_0$.

Large component. We have shown the above result only for fixed a. The proof for large $a > a_0$ is left as an exercise.

■

In fact, it is possible to show much more than that: that the size of the separating set is in fact much larger than r for large a.

It is also worth pointing out that many of the results in this chapter allow for a degree $r = r(n)$ that grows slowly with n.

9.7 General degree distributions

We close this chapter with a brief discussion of an elegant result that generalizes our study of the emergence of the giant component in the models $G(n, p)$ and $G(n, r)$. In this model, the empirical degree distribution $\lambda = \{\lambda_i\}$ is given a-priori. A graph from $G(n, \lambda)$ has $n\lambda_i$ vertices of degree i. Clearly, $G(n, p) = G(n, \text{Binom}(n, p))$, and $G(n, r) = G(n, \{\lambda_r = 1\})$.

The model to sample from $G(n, \lambda)$ generalizes the pairing model we studied above for $G(n, r)$: we generate stubs for each vertex, then randomly connect the stubs to form a pairing, then project and condition on the graph being simple. What is different is that we generate classes of vertices with different degrees, to match the empirical distribution λ (i.e., we have roughly $n\lambda_i$ vertices of degree i).

Molloy and Reed [30] show the following simple criterion for the emergence of a giant component, subject to some technical conditions that we do not discuss here. Define $Q(\lambda) = \sum_{i \geq 1} i(i - 2)\lambda_i$. If $Q(\lambda) < 0$, then $G(n, \lambda)$ has only small components; if $Q(\lambda) > 0$, then it does have a giant component.

While proving this result is quite involved, we can easily develop an intuition of why the function $Q(\lambda)$ determines the appearance of a giant component. We consider for this the same component discovery process as in the discussion of $G(n, p)$, where A_k is the set of active vertices. We can then view $Q(\lambda)$ as the expected difference between A_{k+1} and A_k. Suppose at the ith step, we saturate a vertex of degree i; this means that we remove this vertex from the active set; there are $i - 1$ neighbors of this vertex that are added, for a total change of $(i - 2)$.

What is the probability of hitting a vertex of degree i? For this, it is important to note that this probability is not proportional to λ_i, but rather to $i\lambda_i$. This is because we sample edges, rather than vertices, which gives a bias towards higher-degree vertices. The function $Q(\lambda)$ is therefore proportional

to the expected change in the active set; if this change is positive, then the discovery process is likely to either die early, or to give rise to a giant component, in analogy to the proofs in Section 8.4.

10

Small Worlds

10.1 Introduction

The term *small world network* was coined by the sociologist Stanley Milgram in the 1960s in the context of his experiments on the structure of social networks. A social network has people as its vertices, and social connections (friendship, acquaintance, business relationship, etc.) as its edges. Milgram was interested in determining the distance in hops separating arbitrary persons. In an ingenious experiment, Milgram mailed letters to randomly chosen individuals in Nebraska. The task of these individuals was to send a letter to an target individual living in Boston, but the letter could be sent only through chains of social acquaintances.

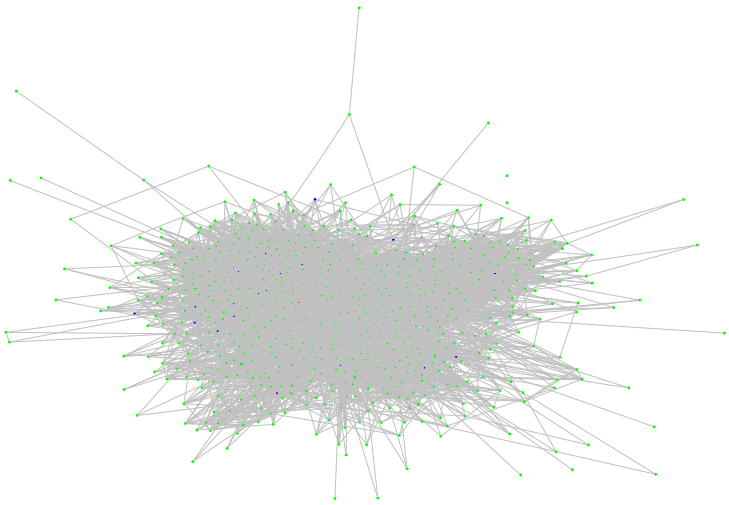

Figure 10.1: The graph spanned by the two-hop neighborhood of a starting vertex (in red), in a social network derived from email exchanges.

The outcome of these experiments was surprising: a relatively large fraction of these letters did indeed arrive at their target; furthermore, they did so after traversing only a small number (approx. five) of social edges. It is rather surprising to think that in a country with a population on the order of 10^8 people, two completely unrelated, arbitrarily chosen individuals, who might lead very different lives, belong to different social classes, a live thousands of kilometers apart, would nevertheless be so close in the social network. It's a small world.

10.2 Random graphs are small

The diameter of many randomly generated graphs, such as $G(n,p)$ and $G(n,r)$, are surprisingly small. At the risk of making an overly sweeping statement, we can say that randomness produces rapidly expanding, and hence small, networks. In this section, we study a slightly different model from the random graphs considered so far, to avoid certain technical difficulties. Specifically, we study a random network obtained by adding a random matching to an n-cycle. Note that an n-cycle alone has diameter $n/2$.

Theorem 10.1 (Cycle + random matching have small diameter[4]). *Let G be an undirected graph formed by adding a random matching to an n-cycle. Then G has diameter $diam(G)$ satisfying* a.a.s.

$$\log_2 n - c \leq diam(G) \leq \log_2 n + \log_2 \ln n + c, \tag{10.1}$$

where c is a constant (at most 10).

Proof: The idea of the proof is to show that most chords, i.e., edges in the random matching, lead to new vertices that are sufficiently far away from any previously visited vertices when we explore the graph starting from a fixed vertex v. In the proof, we have to proceed in two phases. In the first phase, we consider distances that are relatively short with respect to the diameter, and when most vertices have not been visited yet; in the second, we consider distances above that threshold l.

Let C denote the n-cycle, and M the random matching, so that $G = C \cup M$. Also, let $d_C(u,v)$ denote the distance between u and v in C; note that $d_C(u,v) \leq n/2$. We start at a vertex v, and define circles and balls around v as follows.

$$S_i = \{u : d(u,v) = i\}, B_i = \cup_{j \leq i} S_j = \{u : d(u,v) \leq i\} \tag{10.2}$$

Short distances $i \leq l = (1/5)\log_2 n$. Consider a chord (u,v) where $u \in S_i$ and $v \in S_{i+1}$. We call such a chord *local* if v is close on the cycle to at least one other vertex in B_{i+1}, i.e., if $d_C(v,v') \leq 3\log_2 n$ for any other $v' \in B_{i+1}$.

Note that $|B_i| \leq 3 \cdot 2^i$, because after the initial 3 neighbors of v, each node in the previous stage gives rise to at most 2 children.

The probability that a new chord after step i is local is at most

$$\frac{|B_{i+1}| 3\log_2 n}{n} \leq \frac{9 \cdot 2^{i+1} \log_2 n}{n}. \tag{10.3}$$

We now want to compute the probability that local chords are rare while $i \leq l$. Specifically, consider all the chords traversed in the first l steps, of which there are at most $|B_l|$. Therefore, the probability that there are two or more chords in this set is

$$\mathbb{P}\{\text{at least two local chords}\} \leq \binom{3 \cdot 2^l}{2} \left(\frac{9 \cdot 2^l \log_2 n}{n}\right)^2 = O(n^{-6/5}(\log_2 n)^2) = o(n^{-1}). \tag{10.4}$$

A union bound over all n starting vertices then shows that *a.a.s.*for every starting vertex v, there is at most one local chord up to step l. From now on, we condition on the event A that this is true.

Fresh neighbors on one or two sides. We have shown that most chords lead to vertices that are at least $3\log_2 n$ from other already discovered vertices. We now use this property to define two sets of vertices at step i. The set C_i contains vertices that have at least $\log_2 n$ untouched vertices on one side (of the cycle), and there is a unique path of length i to such a vertex; the set D_i contains vertices that have at least $\log_2 n$ vertices on both sides, and a unique path.

We now lower-bound the sizes of C_i and D_i, conditional on the event $A \cap B$, i.e., of having at most one local chord in the first phase, and at most $2^i n^{-1/10}$ local chords in each step of the second phase.

Consider a vertex $v \in C_i$. A neighbor u of v on the circle becomes an element of C_{i+1}, unless a local chord falls into the interval of free vertices next to u at step $i+1$ (or hits u itself). Also, because $v \in C_i$, there is a fresh chord (v,u), such that $u \in D_{i+1}$ unless (v,u) is a local chord.

Similarly, for $v \in D_i$, both neighbors of v become elements of C_{i+1}, unless a loal chord hits on either side of v.

Clearly, $C_i \cup D_i \subset S_i$. Conditional on A, the worst case (giving smallest C_i and D_i) is when the (unique) local chord goes to one of the neighbors of v. In that case, $C_1 = C_2 = 1, C_3 = 2$, and generally

$$\begin{aligned} |C_i| &\geq 2^{i-2} \\ |D_i| &\geq 2^{i-3}. \end{aligned} \tag{10.5}$$

Long distances $l < i \leq (3/5)\log_2 n$. The probability that a chord is local is

$$p_l = \mathbb{P}\{\text{chord local}\} \leq \frac{18 \cdot 2^i \log_2 n}{n} = O(n^{-1/6}). \tag{10.6}$$

There are at most 2^i chords leaving the set S_i. The probability that there are at least $2^i n^{-1/10}$ local chords leaving S_i is at most

$$\left(\begin{array}{c} 2^i \\ 2^i n^{-1/10} \end{array} \right) (p_l)^{2^i n^{-1/10}} \leq \left(\frac{e 2^i}{2^i n^{-1/10}} \right)^{2^i n^{-1/10}} = e^{2^i n^{-1/10}} n^{-(1/15) 2^i n^{-1/10}} = O(n^{-5}) \tag{10.7}$$

A union bound over all n vertices and all $(2/5)\log_2 n$ time steps i shows that *a.a.s.*, at most $2^i n^{-1/10}$ chords leave S_i. Call this event B.

A neighbor u of $v \in C_i$ is sure to be an element of C_{i+1}, except if a local chord falls into the interval of free vertices next to u at step $i+1$.

$$\begin{aligned} |C_{i+1}| &\geq |C_i| + 2|D_i| - 2^{i+1} n^{-1/10} \\ |D_{i+1}| &\geq |C_i| - 2^{i+1} n^{-1/10} \end{aligned} \tag{10.8}$$

Therefore, over the entire range of $i \in [3, (3/5)\log_2 n]$,

$$\begin{aligned} |C_i| &\geq 2^{i-3} \\ |D_i| &\geq 2^{i-4} \end{aligned} \tag{10.9}$$

All vertices are close. Set $i^* = (1/2)(\log_2 n + \log_2 \ln n + c) \leq (3/5)\log_2 n$. Suppose we go through the discovery process described above for two different starting vertices v' and v'', to generate two

sets $C_{i^*}(v')$ and $C_{i^*}(v'')$. Assuming that the two balls around v and v'' have not touched yet, we compute the probability that the set $C_{i^*}(v')$ and $C_{i^*}(v'')$ will be connected by one of the edges generated in the next step. Specifically, we can conservatively focus only on the chords generated in the next step by vertices in $C_{i^*}(v')$, and ask whether they will hit any vertex in $C_{i^*}(v'')$. The probability that none hits is upper-bounded by

$$\begin{aligned}
\mathbb{P}\{d(v,u) > 2i^*+1 | A \cap B\} &\leq \left(1 - \frac{2^{i^*-3}}{n}\right)^{2^{i^*-3}} \\
&\leq \exp\left(-2^{2i^*-7}/n\right) \quad \text{using } (p \leq -\log(1-p)) \\
&\leq \exp\left(-2^{c-7}\ln n\right) \\
&\leq n^{-4} = o(n^{-2}), \qquad (10.10)
\end{aligned}$$

for $c \geq 9$.

We can now bound the diameter of the entire graph.

$$\mathbb{P}\{\text{diam}(G) > 2i^*+1\} \leq \mathbb{P}\left\{\bar{A}\right\} + \mathbb{P}\left\{\bar{B}\right\} + \sum_{u,v} \mathbb{P}\{d(v,u) > 2i^*+1 | A \cap B\} = o(1). \qquad (10.11)$$

Therefore, G has diameter $2i^*+1 = \log_2 n + \log_2 \ln n + 10$ *a.a.s.*

■

Therefore, the addition of the random matching has decreased the diameter significantly, from $n/2$ for the n-cycle, to about $\log_2 n$. Similar results are known for other classes of random graphs; while the techniques are similar, the proofs are typically (even) more involved than the one above. We give two examples.

Theorem 10.2 (Diameter of giant component of $G(n,p)$ [7]). *The diameter of the giant component of $G(n,p)$ for $\ln n > np \to \infty$ satisfies*

$$diam(G(n,p)) = (1+o(1))\frac{\log n}{\log np} \text{ a.a.s..} \qquad (10.12)$$

Theorem 10.3 (Diameter of regular random graphs $G(n,r)$ [5]). *Let $r \geq 3$ and $\epsilon > 0$ be fixed. Then*

$$\frac{diam(G(n,r))}{\log n/\log(r-1)} \in (1-\epsilon, 1+\epsilon) \text{ a.a.s..} \qquad (10.13)$$

So random graphs do possess the small world property, in that their diameter behaves like $\log n$. In a sense, the absence of structure in random graphs, i.e., that edges are independent, ensure that the neighborhood of a vertex v grows quickly with distance, a fact we had explicitly used in the proof of the emergence of the giant component for $G(n, c/n)$.

Consider a graph G of maximum degree $\Delta(G)$ and of diameter D. We can establish the following inequality bounding the order of G. A vertex v can have at most Δ neighbors. Each of these neighbors in turn can have at most $\Delta - 1$ new neighbors, and so forth. Therefore

$$n \leq 1 + \Delta\left(\sum_{i=1}^{D}(\Delta-1)^{(i-1)}\right) = 1 + \Delta\frac{(\Delta-1)^D - 1}{\Delta - 2}. \qquad (10.14)$$

Graphs that have equality in (10.14) are called Moore graphs. Moore graphs exist only for particular sets of values of Δ and D. Note that trees are not Moore graphs, as their diameter is twice their height.

The above equation shows that $D = \Omega(\log n / \log(\Delta - 1))$. Informally, a small world graph has small diameter close to the above bound. (Sometimes, the small world property is also defined in terms of a definition of the average distance, rather than the diameter, i.e., the maximum distance). The key is that to achieve small diameter, it is necessary for the size of the i-hop neighborhood to grow exponentially, as in (10.14). This is not the case for regular lattices $\mathbb{L}^d$ of dimension d, whose i-hop neighborhood only grows polynomially as i^d, and whose diameter is therefore on the order of $n^{1/d}$.

10.3 Clustering

Another feature of real networks is their "cliquishness" or "transitivity", i.e., the tendency for two neighbors of a vertex v to be connected by an edge.

Two distinct definitions of the clustering coefficient are used in the literature. Both are based on a definition of the clustering coefficient C_v of a vertex, given by

$$\begin{aligned} C_v &= \frac{\text{number of edges between neighbors of } v}{\text{number of possible edges between neighbors of } v} \\ &= \frac{|\{(u,w) \in E(G) : (u,v) \in E(G), (w,v) \in E(G)\}|}{\binom{d(v)}{2}} \end{aligned} \tag{10.15}$$

Two different definitions of the clustering coefficient of the entire graph have been proposed in the literature.

$$C_1(G) = \frac{1}{n} \sum_{v \in V(G)} C_v(G) \tag{10.16}$$

$$\begin{aligned} C_2(G) &= \frac{\sum_{v \in V(G)} \binom{d(v)}{2} C_v(G)}{\sum_{v \in V(G)} \binom{d(v)}{2}} \\ &= \frac{3 \times \text{number of triangles}}{\text{number of pairs of adjacent edges}} \end{aligned} \tag{10.17}$$

Note that for the random graph $G(n,p)$, every possible edge exists independently of everything else with probability p. Therefore

$$\mathbb{E}\left[C_v(G(n,p))\right] = p. \tag{10.18}$$

Also, note that we had seen that for $G(n,r)$, the number Z_3 of triangles is asymptotically Poisson with mean independent of n, which shows that the clustering coefficient would decrease as $1/n$.

Note that the clustering coefficient is limited in that it only captures local connectivity within one hop. It is easy to construct graphs that have rich local connectivity over more than one hop, but that have $C(G) = 0$. It might therefore be desirable to develop more robust measures for clustering that would take into account more than the one-hop neighborhood of vertices. As an example, consider transforming a graph by breaking every edge in half and "inserting" an additional vertex. The transformed graph has $C = 0$, even though it has "almost" the same structure as the initial graph.

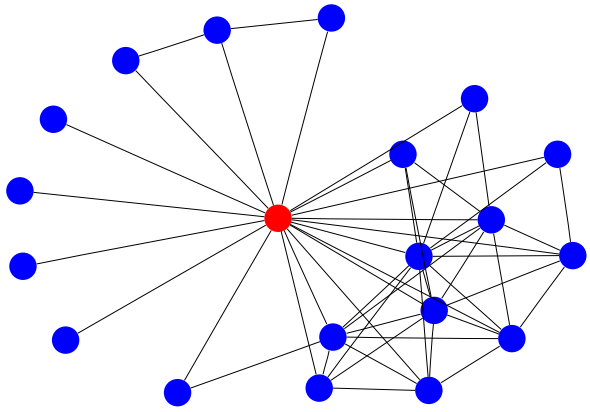

Figure 10.2: The graph spanned by the one-hop neighborhood of a starting vertex (in red), in a social network derived from email exchanges. Note that there are many triangles resulting from edges between neighbors of the red vertex.

10.4 Small worlds: small and transitive

Studies of real networks (social networks, the world wide web, the power grid, etc.) show that they usually possess the small world property like random graphs, while at the same time exhibiting a large clustering coefficient like lattices. This has motivated new models of real networks. The basic idea is to start with a regular lattice (e.g., a cycle of length n), and then to select each edge independently with probability p to "rewire" it, i.e., to reattribute randomly uniformly over all vertices one or both of the endpoints of such an edge. In a variation, edges are not rewired, but new random edges are added to the existing lattice.

This type of model can therefore be viewed as an interpolation between the lattice at $p = 0$, and a random graph at $p = 1$. A neighbor at distance i of v, $1 \leq i \leq k$, it is easy to see that this neighbor has $2k - 1 - i$ edges to other neighbors of v. Therefore, the number of edges among neighbors of v is $2\sum_{i=1}^{k} i = 3k(k-1)$. Hence, for $p = 0$, the clustering coefficient is $C(S(n,0)) = C_v = \frac{3(k-1)}{2(2k-1)}$ For $p > 1$, we can approximate the clustering coefficient by noting that a triangle survives rewiring with probability $(1-p)^3$, and

$$C(S(n,p)) \approx C(S(n,0))(1-p)^3. \tag{10.19}$$

This confirms our intuition that if the fraction of rewired edges is low, the impact on the clustering coefficient is quite small. On the other hand, even for small p, the fact that the subgraph of rewired edges resembles a random graph, the average distance and the diameter drop very quickly even for small p. Therefore, intermediate values of p model the two features of real networks: small distances (similar to random graphs), but large clustering coefficient (similar to lattices).

The Watts-Strogatz model starts out with a network that has high clustering, but a diameter much larger than $O(\log n)$, and then adds randomness to this network, which quickly brings the diameter down. One could argue that this model, while capturing the two key aspects of small world networks - high clustering, small distance - is rather artificial in that it embeds the network into a geometry (the initial lattice) which is unlikely to be a feature of many real networks. This becomes important, for example, when we study how such networks can be navigated, because the geometry provides important clues.

Another small world model results if we start out with a network that does not have high clustering, but small distances, and then adding additional edges to increase clustering. For example, one may

start with a random regular graph $G(n, r)$, and then add between each pair of vertices (u, v) for which $d(u, v) = 2$ an edge with probability q.

11

Continuum percolation for multi-hop wireless networks

11.1 Introduction

We now move to continuum percolation models, where nodes are no longer placed at the vertices of a lattice or tree, but are randomly scattered in $\mathbb{R}^d$. We will always make the assumption that the node distribution follows a homogeneous Poisson process of intensity $\lambda > 0$, although the results can be extended to many more general stationary and ergodic point processes. Percolation in this setting is referred to as *continuum percolation*, the main reference is the book by Meester and Roy [29]. We begin with the simplest model, where nodes connect with each other if their distance is less than some connectivity range r, which is the *Boolean model*. We then move to a model that captures better the physical layer of wireless networks, the *Signal to Interference Ratio model*.

There are many other models in continuum percolation; we should mention here an important one, which we will not study is in the course but which is addressed in [29], the *random connection model*, where two nodes located at points $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ are connected to each other with probability $g(\|\boldsymbol{x}_1 - \boldsymbol{x}_2\|)$, independently of all other points, for a given connection function g. The Boolean model with deterministic radius r is a particular case of the random connection model with $g(x) = 1$ if $x \leq r$ and $g(x) = 0$ if $x > r$. If we replace the Poisson point process by a deterministic point process placing a node at each vertex of $\mathbb{Z}^d$, and take $g(x) = p$ if $x \leq 1$ and $g(x) = 0$ otherwise, we find the lattice bond percolation model. The random connection model is thus a quite general model, but the proofs tend to be technically rather complex. Even for the Boolean model, although most results from the previous chapters are still valid, their proof is much more involved. This is why we will only see the main result (existence of a non trivial phase transition) in this chapter for continuum models, and make use of mapping techniques on the lattice. These techniques do not (always) give the tightest possible bounds, but they are among the most straightforward to use.

11.2 Boolean model

11.2.1 Model settings

In the Poisson Boolean model $\mathcal{B}(\lambda, r)$ (or Poisson blob model), the positions of the nodes are distributed according to a Poisson point process of constant, finite intensity λ in $\mathbb{R}^d$. We associate to each node a closed ball of fixed radius $r/2$, as shown in Figure 11.1. The plane is thus partitioned into two regions: the *occupied* region $\mathcal{W}$, which is the region covered by the balls, and the *vacant* region $\mathcal{V}$, which is the complement of the occupied region. The vacant region plays a similar role to the dual lattice in bon percolation, but is much more difficult to handle, see Chapter 4 in [29].

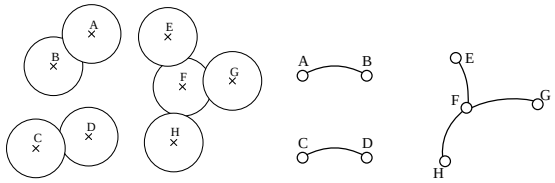

Figure 11.1: The Boolean model (left) and the associated graph (right).

Two nodes are *directly connected* or *immediate neighbors* if the intersection of their associated balls is non-empty. In other words, if this model represents a wireless network, two nodes are able to communicate together through a wireless channel if the distance between them is less than a characteristic *range* r. A *cluster* is a connected component of the occupied region. Finally, two nodes are said to be *connected* together if they belong to the same cluster.

Furthermore, one can associate with the random model $\mathcal{B}(\lambda, r)$ the graph $\mathcal{G}(\lambda, r)$ by associating a vertex to each node of $\mathcal{B}(\lambda, r)$ and an edge with each direct connection in $\mathcal{B}(\lambda, r)$. $\mathcal{G}(\lambda, r)$ is called the *associated graph* of $\mathcal{B}(\lambda, r)$.

We only consider the simple case where r is fixed (it would be the maximal radius allowed by power constraints).

The two models $\mathcal{B}(\lambda, r)$ and $\mathcal{B}(\lambda', r')$ lead to the same associated graph, namely $\mathcal{G}(\lambda, r) = \mathcal{G}(\lambda', r')$ if $\lambda' r'^d = \lambda r^d$. As result, the graph properties of $\mathcal{B}(\lambda, r)$ depend only on one parameter λr^d proportional to the the average node degree $\pi \lambda r^d$.

11.2.2 Percolation probability

The quantities of interest in continuum percolation are the same as in discrete percolation. The first one the probability that a given node, arbitrarily placed at the origin, belongs to a cluster with an infinite number of nodes, which we denote by θ and call the *percolation probability*. With C denoting the cluster containing the origin, the percolation probability is thus defined as before

$$\theta(\lambda, r) = \theta(\lambda r^d) = \mathbb{P}_{\lambda,r}(|C| = \infty). \tag{11.1}$$

By space invariance, $\theta(\lambda r^d)$ is the probability that any node belongs to an infinite cluster.

Define the *critical* (or *percolation*) *threshold* as

$$\left(\lambda r^d\right)_c = \sup\left\{\lambda r^2 \mid \theta(\lambda r^d) = 0\right\}. \tag{11.2}$$

In the one-dimensional case ($d = 1$), it is immediate to see that $\left(\lambda r^d\right)_c = \infty$. Indeed, in the 1-d case, the Poisson Boolean model with constant radius is an M/D/∞ queue, which is ergodic if $(\lambda r)_c < \infty$.

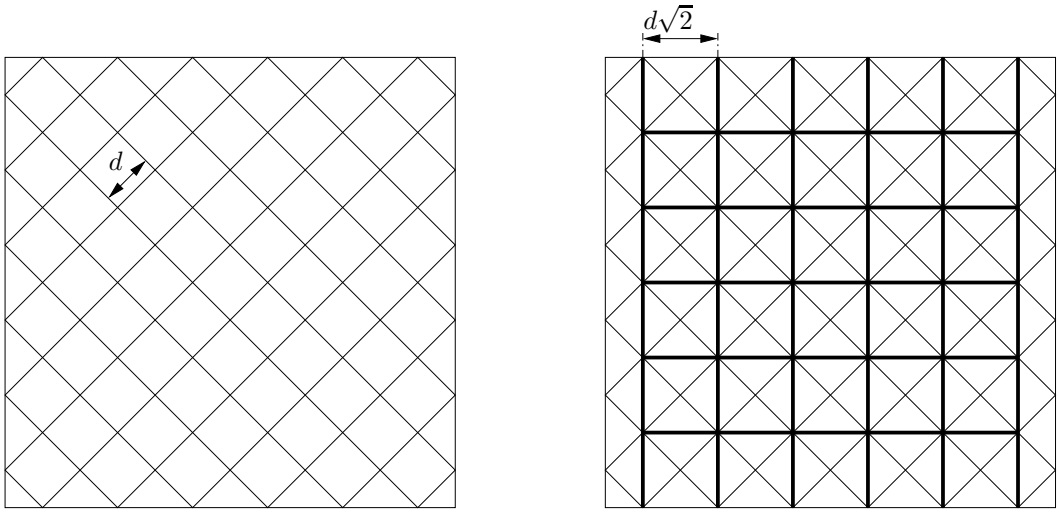

Figure 11.2: Construction of the bond percolation model. We declare each square on the left-hand side of the picture open, if there is at least a Poisson point inside it, closed otherwise. This corresponds to associate an edge to each square, traversing it diagonally, as depicted on the right-hand side of the figure, and declare the edge either open or closed according to the state of the corresponding square.

As a result, all clusters are almost surely finite. However, when $d \geq 2$, it is no longer the case. We will show that $0 < (\lambda r^d)_c < 1$, which implies that there are two phases: the *subcritical phase*, when $\lambda r^d < (\lambda r^d)_c$, where every vertex is almost surely in a finite open cluster, and the *supercritical phase*, when when $\lambda r^d > (\lambda r^d)_c$, where each node has a non zero probability of belonging to an infinite cluster. Computing the exact value of $(\lambda r^d)_c$ is still an open problem, numerical results show that it is close to 1.43 for $d = 2$.

Theorem 11.1 (Non trivial phase transition). *The percolation threshold in* $\mathbb{R}^2$ *is such that* $0 < (\lambda r^2)_c \leq 8 \ln 2$.

Proof:

(i) We first prove that $(\lambda r^2)_c \leq 8 \ln 2$. Let us divide the plane $\mathbb{R}^2$ in squares of size $d \times d$, as depicted in the left-hand of Fig. 11.2. Pick $d = r/\sqrt{8}$. Then any pair of nodes located in two squares having a common edge are connected in $\mathcal{G}(\lambda, r)$.

Let p denote the probability that a square contains at least one point:

$$\mathbb{P}_{\lambda,r}[\text{a square contains at least one point}] = 1 - e^{-\lambda d^2} := p. \tag{11.3}$$

We say that a square is *open* if it contains at least one point, and *closed* otherwise; note that the status of the squares is i.i.d.

In a second step, we construct a bond percolation model on $\mathbb{L}^2$. We draw an horizontal edge across half of the squares, and a vertical edge across the others, as shown on the right-hand side of Fig. 11.2. In this way we obtain a lattice of horizontal and vertical edges, each edge being open, independently of all other edges, with probability p. Whenever two edges sharing an endvertex are open, all nodes in both corresponding squares are connected to each other in $\mathcal{G}(\lambda, r)$.

If $\lambda r^2 > 8 \ln 2$, then $p = 1 - e^{-\lambda d^2} = 1 - e^{-\lambda r^2/8} > 1/2$, and the lattice $\mathbb{L}^2$ contains an infinite open cluster. As a result, the graph $\mathcal{G}(\lambda, r)$ also contains an infinite cluster, and $\lambda r^2 > 8 \ln 2$ is in the supercritical phase.

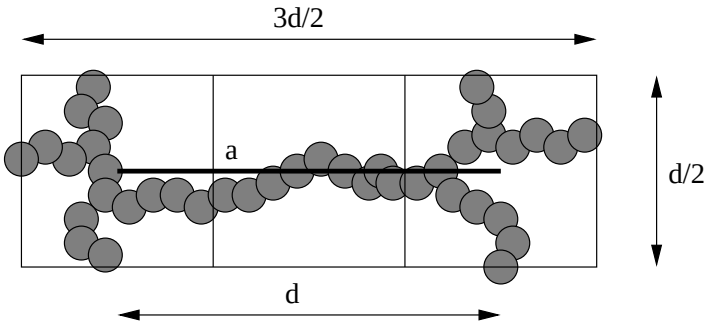

Figure 11.3: A horizontal edge a that fulfills the two conditions for having $A_a = 1$.

(ii) To prove $(\lambda r^2)_c > 0$, we divide again the plane $\mathbb{R}^2$ in squares of size $d \times d$, as depicted in the left-hand of Fig. 11.2. This time we pick $d = r$, so that a point in a square can only connect to points in the 8 adjacent squares.

We now construct a site percolation model by placing a vertex of $\mathbb{Z}^2$ inside each square, and declare the site open if there is at least at least one point in the square, and to closed otherwise. Denoting by p the open site probability, we have thus that

$$p^{site} = 1 - e^{-\lambda r^2}.$$

An open site is connected to every open site among its 8 adjacent sites. Consequently, if there is an infinite cluster in $\mathcal{B}(\lambda, r)$, then there is an infinite cluster in the site model as well. Let p_c^{site} be the site critical threshold, one can show that $0 < p_c^{site}$.

Suppose that $\lambda r^2 > (\lambda r^2)_c$. In this case, there is an infinite infinite cluster in $\mathcal{B}(\lambda, r)$ and thus also in the site model, whence $p^{site} > p_c^{site}$ and $\lambda r^2 > -\ln(1 - p_c^{site})$. In other words, $\lambda r^2 > (\lambda r^2)_c$ always implies $\lambda r^2 > -\ln(1 - p_c^{site})$, which means that $(\lambda r^2)_c \geq -\ln(1 - p_c^{site}) > 0$.

■

11.2.3 Another useful mapping

The main properties of the lattice bond percolation can be extended to the Boolean model, but the proofs are technically much more involved. The reader is referred to the textbook by Meester and Roy [29].

For $d > 0$, we denote by $\mathbb{L}^2$ the two-dimensional square lattice whose vertices are located at all points of the form (dx, dy) with $(x, y) \in \mathbb{Z}^2$. For each horizontal edge a of $\mathbb{L}^2$, we denote by $\boldsymbol{z}_a = (x_a, y_a)$ the point in the middle of the edge, and introduce the random field A_a, indexed by the edges of $\mathbb{L}^2$, that takes the value 1 if the following two events (illustrated in Figure 11.3) occur, and 0 otherwise:

1. the rectangle $[x_a - 3d/4, x_a + 3d/4] \times [y_a - d/4, y_a + d/4]$ is crossed from left to right by an occupied component in $\mathcal{B}(\lambda, r)$, and
2. both squares $[x_a - 3d/4, x_a - d/4] \times [y_a - d/4, y_a + d/4]$; $[x_a + d/4, x_a + 3d/4] \times [y_a - d/4, y_a + d/4]$ are crossed from top to bottom by an occupied component in $\mathcal{B}(\lambda, r)$.

We define A_a similarly for vertical edges, by rotating the above conditions by 90°.

According to [29, Corollary 4.1], the probability that $A_a = 1$ can be made as large as we like by choosing d large. The variables A_a are not independent in general. However, if edges a and b are not

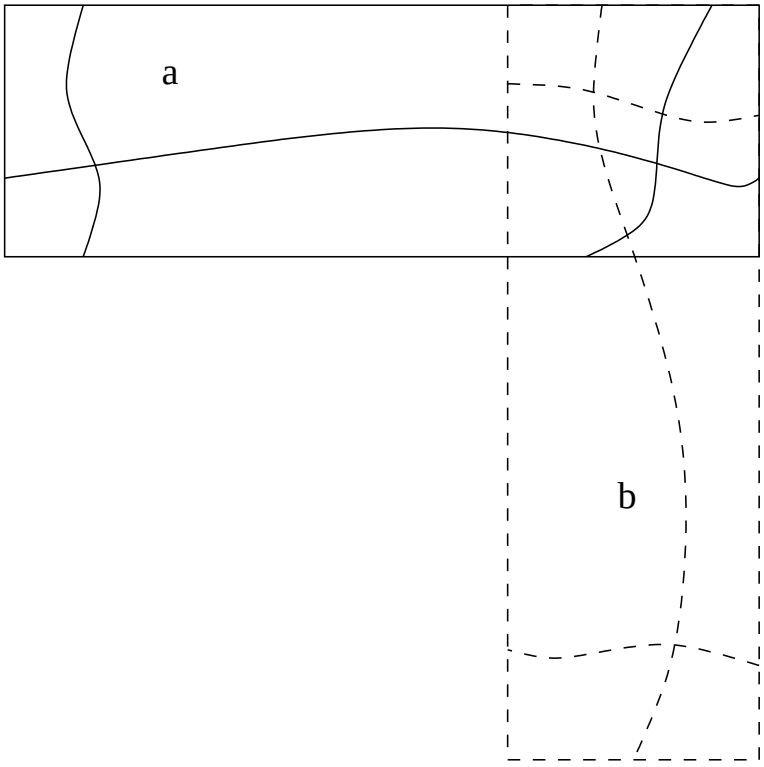

Figure 11.4: Two adjacent edges a (plain) and b (dashed) with $A_a = 1$ and $A_b = 1$. The crossings overlap, and form a connected component.

adjacent, then A_a and A_b are independent: these variables thus define a 1-dependent edge percolation process.

The reverse mapping follows from the observation that if $A_a = 1$, there exist crossings along edge a, as shown in Figure 11.3. These crossings are designed such that if for two adjacent edges a and b, $A_a = 1$ and $A_b = 1$, the crossings overlap, and they all belong to the same connected component (see Figure 11.4). Thus, an infinite cluster of such edges implies an infinite cluster in the Boolean model $\mathcal{B}(\lambda, r)$.

11.3 Signal to interference Ratio Model

The Boolean model is a very crude approximation of wireless multi-hop networks, and much research effort is currently devoted to obtain more realistic models. In particular, interferences need to be taken in account. We describe a model that incorporates interferences, and for which percolation holds under some assumptions [9, 10].

11.3.1 STIRG model

Nodes are distributed according to a Poisson point process of constant spatial intensity λ. Depending on its location, number of neighbors, and battery level, each node i will adjust its emitting power P_i within a given range $[0, P]$, where P is the maximal power of a node, which is finite. The power of the

signal emitted by Node i and received by Node j is $P_i L(\boldsymbol{x}_i - \boldsymbol{x}_j)$, where $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$ are the positions of Node i and j in the plane, respectively, and $L(\cdot)$ is the attenuation function in the wireless medium.

We assume that Node i can transmit data to Node j if the signal received by j is strong enough, compared to the thermal noise. Formally, this condition is written as

$$\frac{P_i L(\boldsymbol{x}_i - \boldsymbol{x}_j)}{N_0 + \gamma \sum_{k \neq i,j} P_k L(\boldsymbol{x}_k - \boldsymbol{x}_j)} \geq \beta, \tag{11.4}$$

where N_0 is the power of the thermal background noise and β is the signal to noise ratio required for successful decoding. The coefficient γ is the inverse of the processing gain of the system, it weights the effect of interferences, depending on the orthogonality between codes used during simultaneous transmissions. It is equal to 1 in a narrow band system, and is smaller than 1 in a broadband system that uses CDMA. The physical model of Gupta and Kumar [22] assumes $\gamma = 1$; other models [20] allow γ to be smaller than 1.

Similarly, Node j can transmit data to Node i if and only if

$$\frac{P_j L(\boldsymbol{x}_j - \boldsymbol{x}_i)}{N_0 + \gamma \sum_{k \neq i,j} P_k L(\boldsymbol{x}_k - \boldsymbol{x}_i)} \geq \beta. \tag{11.5}$$

From conditions (11.4) and (11.5), we can build an oriented graph that summarizes the available links between nodes. In order to define *connected components* (or *clusters*), we have to introduce a symmetric relation. In this paper, we choose to neglect unidirectional links, which are difficult to exploit in wireless networks [34]. In other words, we declare that Node i and Node j are *directly connected* if and only if both (11.4) and (11.5) are satisfied. This new relation leads to the definition of a non-oriented random graph associated with the Poisson point process, which we call *Poisson Signal To Interference Ratio Graph (STIRG)*.

As our model has many more parameters than degrees of freedom, we will focus on the node density λ and the orthogonality factor γ. The other parameters are supposed constant in the sequel. We will thus denote by $\mathcal{G}(\gamma, \lambda)$ the connectivity graph.

11.3.2 A Bound on the Degree of the Nodes

In the following theorem, we will prove that if $\gamma > 0$, the number of neighbors of each node is bounded from above (note that this is not the case in the Boolean Model with $\gamma = 0$).

Theorem 11.2. *Each node can have at most* $1 + 1/\gamma\beta$ *neighbors.*

Proof: Pick any node (called hereafter Node 0), and let N be the number of its neighbors (i.e. the number of nodes to which Node 0 is connected). If $N \leq 1$, the claim is trivially proven. Suppose next that $N > 1$, and denote by 1 the node whose signal power received by Node 0 is the smallest but is non zero, namely is such that

$$0 < P_1 L(\boldsymbol{x}_1 - \boldsymbol{x}_0) \leq P_i L(\boldsymbol{x}_i - \boldsymbol{x}_0), \; i = 2 \ldots N. \tag{11.6}$$

Since it is connected to Node 0, (11.4) imposes that

$$\frac{P_1 L(\boldsymbol{x}_1 - \boldsymbol{x}_0)}{N_0 + \gamma \sum_{i=2}^{\infty} P_i L(\boldsymbol{x}_i - \boldsymbol{x}_0)} \geq \beta. \tag{11.7}$$

Taking (11.6) into account, (11.7) implies that

$$\begin{aligned} P_1 L(\boldsymbol{x}_1 - \boldsymbol{x}_0) &\geq \beta N_0 + \beta\gamma \sum_{i=2}^{\infty} P_i L(\boldsymbol{x}_i - \boldsymbol{x}_0) \\ &\geq \beta N_0 + \beta\gamma(N-1) P_1 L(\boldsymbol{x}_1 - \boldsymbol{x}_0) + \beta\gamma \sum_{i=N+1}^{\infty} P_i L(\boldsymbol{x}_i - \boldsymbol{x}_0) \\ &\geq \beta\gamma(N-1) P_1 L(\boldsymbol{x}_1 - \boldsymbol{x}_0), \end{aligned}$$

from which we deduce that

$$N \leq 1 + \frac{1}{\beta\gamma}.$$

■

In CDMA cellular networks, this kind of bound is known under the name of pole capacity (see e.g. [37], [1]).

As a consequence of Theorem 11.2, we see that if $\gamma > 1/\beta$, each node has at most one neighbor. This is a very general and restrictive condition, that imposes the network to use efficient spread-spectrum encoding in order to keep γ small, or to introduce a scheduling between nodes to avoid having them emitting all the same time.

11.3.3 Shot-Noise

The sum in the denominator of (11.4) is a random variable that depends on the position of almost all nodes in the network. We can write it as $N_0 + \gamma I(\boldsymbol{x}_j) - \gamma P_i L(\boldsymbol{x}_i - \boldsymbol{x}_j)$ where

$$I(\boldsymbol{x}) = \sum_{i, \boldsymbol{x}_i \neq \boldsymbol{x}} P_i L(\boldsymbol{x}_i - \boldsymbol{x}) \tag{11.8}$$

is the *interference contribution*. This kind of variable is called a Poisson *shot-noise*. As it is an infinite sum, it may diverge to infinity, making connections impossible.

If we assume that the sequence $\{P_i\}$ is uniformly bounded from below by a strictly positive constant, and that the attenuation function is isotropic, i.e. that $L(\cdot)$ has the form $L(\boldsymbol{x}) = l(\|\boldsymbol{x}\|)$ where $l(t)$ is a non increasing function of t, the necessary and sufficient condition for the sum

$$\sum_i P_i L(\boldsymbol{x}_i - \boldsymbol{x})$$

to be a.s. finite is (see [8]):

$$\int_0^{\infty} l(t) t dt < \infty \tag{11.9}$$

This condition remains valid if $\inf_i\{P_i\} = 0$ but the sequence $\{1_{\{P_i < \varepsilon\}}\}$ is i.i.d and independent from the point process, for some $\varepsilon > 0$.

We note that for $l(t) = 1/t^2$, the integral in (11.9) diverges and thus no connection is possible in this case whenever $\gamma > 0$.

11.3.4 Attenuation function

We make the following assumptions on the attenuation function $L(\cdot)$:

1. Isotropic attenuation: $L(\cdot)$ has the form $L(\boldsymbol{x}) = l(\|\boldsymbol{x}\|)$ for some function $l : \mathbb{R}^+ \to \mathbb{R}^+$;

2. Physical energy dissipation: $l(x) \leq 1$;
3. Monotonicity: l is continuous and as long as it does not vanish, it is strictly decreasing;
4. Integrability: $\int_0^\infty xl(x)dx < \infty$.
5. Non vanishing power: $l(0) > \beta N_0/P$.

Among these assumptions, the second one is needed to avoid that the received power be larger than the emitted power, and the third one is quite natural given the physics of the problem. The fourth assumption has been shown to be necessary in the previous subsection. Finally, if the fifth assumption is not verified, Equation (11.4) never holds.

The most popular attenuation function is

$$l(x) = \min\left\{1, \frac{1}{x^\alpha}\right\}, \tag{11.10}$$

with α ranging from 3 to 6.

11.3.5 Existence of a percolation threshold for $\gamma = 0$

Let us first note that if we let $\gamma = 0$, the model described in Section 11.3.1 becomes equivalent to a generalized Boolean model, where two nodes are connected if and only if they are in a ball of radius r_b (which can be a deterministic or random value), independently from all the other nodes. Assuming all nodes emit at the maximum power P, this radius r_b is then constant and found from (11.4) to be equal to

$$r_b = l^{-1}\left(\frac{\beta N_0}{P}\right). \tag{11.11}$$

Since l is continuous, strictly monotone and larger than $\beta N_0/P$ at the origin, we have that $l^{-1}(\beta N_0/P)$ exists. For example, for the attenuation function (11.10), this radius reads $r_b = (P/(\beta N_0))^{1/\alpha}$.

We denote by λ_c the critical density of the model when $\gamma = 0$ (and all other parameters N_0, P and β fixed).

11.3.6 Some observations on the graph with $\gamma > 0$

If $\gamma > 0$, it is clear that for the same realization of the spatial point process giving the position of the nodes, the graph obtained with $\gamma > 0$ misses some edges in the graph obtained with $\gamma' = 0$. In other words, $\mathcal{G}(\gamma, \lambda) \subseteq \mathcal{G}(0, \lambda)$. As a result, it is not sure that percolation still occurs for nonzero values of γ. At least, for $\lambda < \lambda^*$, we are sure that $\mathcal{G}(\gamma, \lambda)$ is always sub-critical. However, for $\lambda > \lambda^*$, we know that:

1. For $\gamma = 0$, the network is super-critical
2. For $\gamma > 1/\beta$, the network is sub-critical

Therefore, there exists a critical value $0 \leq \gamma^*(\lambda) \leq 1/\beta$ at which one observes a phase transition. The next subsection will prove that $\gamma^*(\lambda)$ is strictly positive for sufficiently large values of λ.

We have computed by simulation the value of the percolation threshold $\gamma^*(\lambda)$, with $L(\boldsymbol{x}) = \min\{1, \|\boldsymbol{x}\|^{-3}\}$. The simulation results are shown in Figures 11.5, 11.6, 11.7 and 11.8. In the simulations, all nodes emit with the same power P. We observe in Figure 11.5 and 11.6 that $\mathcal{G}(0.02, \lambda) \subseteq \mathcal{G}(0, \lambda)$. We

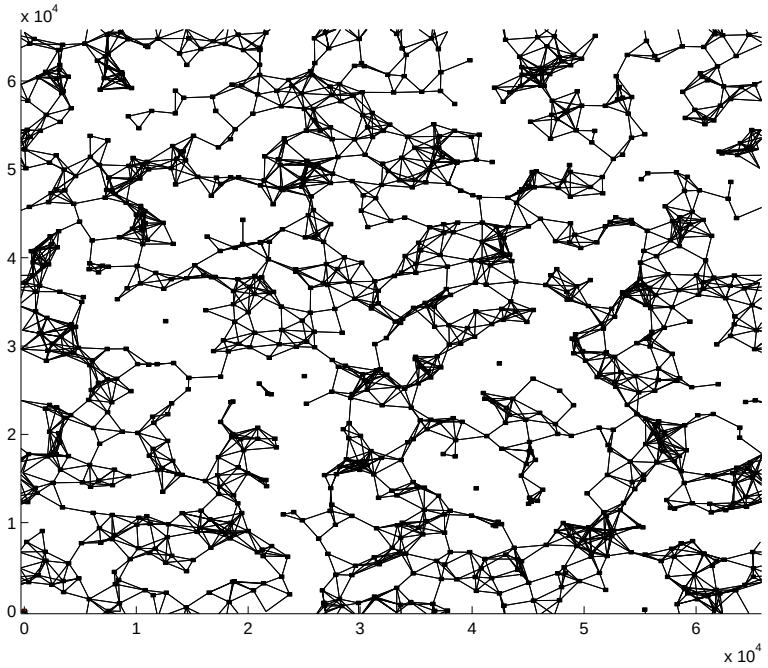

Figure 11.5: An example of graph $\mathcal{G}(0, \lambda)$ with no interference (Boolean Model). As the node density is super-critical ($\lambda > \lambda^*$), most of the nodes belong to the same connected component. [This simulation was run in a square of 65,536 $\times$ 65,536 pixels with parameters $\lambda = 9.31\dot{1}0^{-4}$, $\beta = 1$, $\gamma = 0$, $N_0 = 1$, $P_i = 100,000\ \forall i$.]

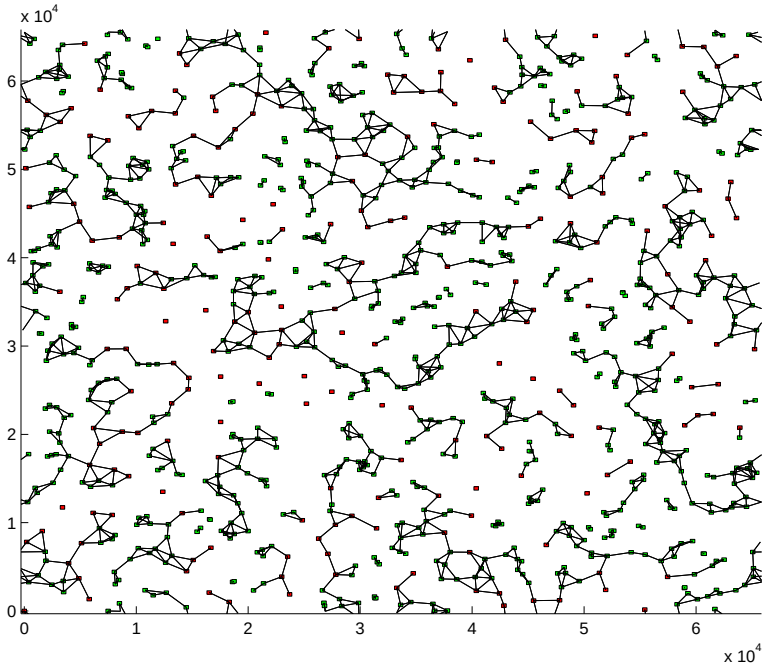

Figure 11.6: An example of graph $\mathcal{G}(\gamma, \lambda)$ with interferences ($\gamma = 0.02$). This simulation was run with the same parameters as in Figure 11.5, except γ that is now nonzero. Due to the interferences, the graph is split into many small components.

observe in Figure 11.7 that $\gamma^*(\lambda)$ exhibit a maximum at a certain density $\tilde{\lambda}$. Below $\tilde{\lambda}$, increasing the node density helps for connectivity, whereas after the maximum, the impact of interferences becomes preponderant, and $\gamma^*(\lambda)$ becomes decreasing. Figure 11.8 illustrates the percolation phenomenon with γ slightly smaller than $\gamma^*(\lambda)$.

11.3.7 Percolation for nonzero values of γ

We have shown above that if γ exceeds some finite, positive critical value, percolation does not occur. We want now to show that percolation can occur for nonzero values of γ. We make the simplifying assumption that every node emits at maximal power P: $P_i = P \; \forall i$.

We write $\boldsymbol{x} \leftrightarrow \boldsymbol{y}$ if there exists a sequence $\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_k$ of Poisson points such that $\boldsymbol{x}_1 = \boldsymbol{x}$, $\boldsymbol{x}_k = \boldsymbol{y}$, and $\boldsymbol{x}_l$ is connected by an edge to $\boldsymbol{x}_{l+1}$ for $1 \leq l < k$. A *(connected) component* or *cluster* is a set $\{\boldsymbol{x}_i : i \in J\}$ of points which is maximal with the property that $\boldsymbol{x}_i \leftrightarrow \boldsymbol{x}_j$ for all $i, j \in J$.

Here is the main result, which is proven in the next subsections.

Theorem 11.3. *Let λ_c be the critical node density when $\gamma = 0$ and assume that the attenuation function $\mathbb{L}(\cdot)$ satisfies the assumptions 1-5 stated in subsection 11.3.4. Then for any node density $\lambda > \lambda_c$, there exists $\gamma^*(\lambda) > 0$ such that for $\gamma \leq \gamma^*(\lambda)$, the SINR-model percolates.*

11.3.8 Proof of Theorem 12.1

The main strategy of the proof is by coupling the model to a discrete edge percolation model on the lattice, as described in Subsection 11.2.3. By doing so, we end up with a dependent discrete model,

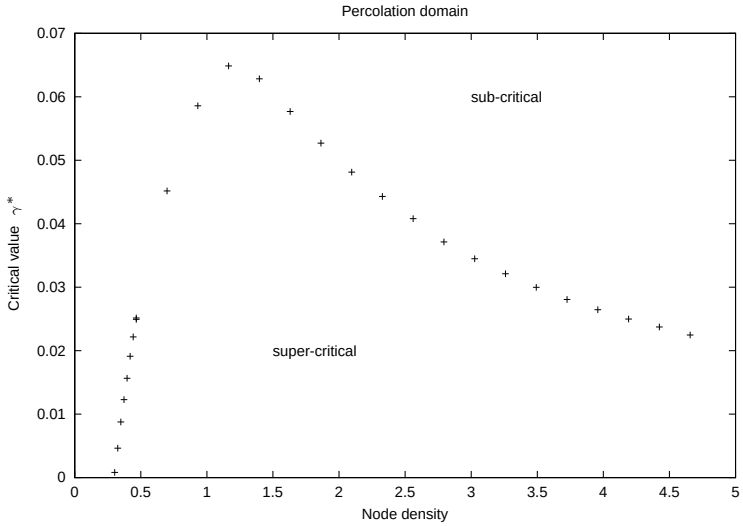

Figure 11.7: Critical value of γ as a function of the node density λ. The curve shows the critical value of γ below which the network percolates. [The parameters of this simulation are $\beta = 1$, $N_0 = 10^4$ and $P_i = 10^5\ \forall i$.]

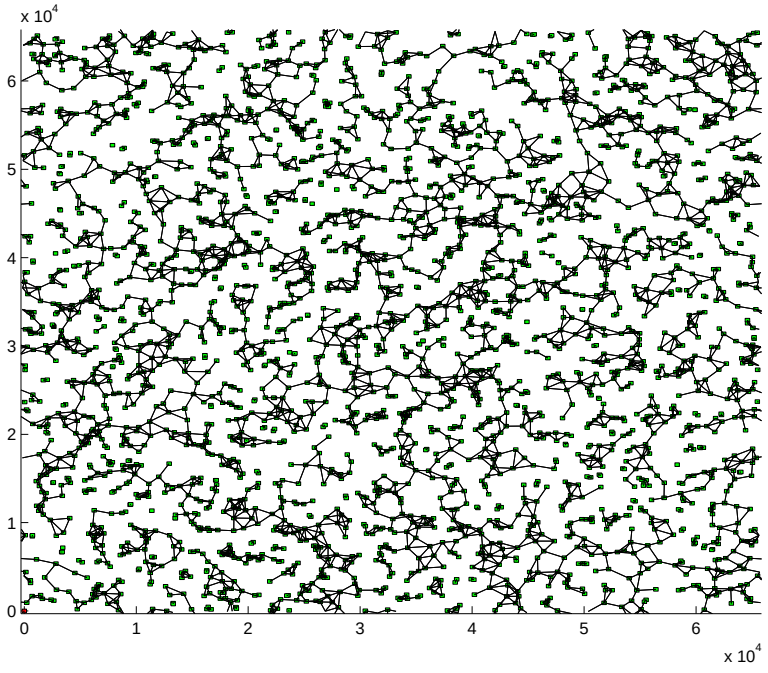

Figure 11.8: A barely supercritical graph with interferences. This simulation was run with the same parameters as in Figure 11.6, except that the node density is higher ($\lambda = 2.79\dot{1}0^{-3}$). The graph percolates despite the interferences because here $\gamma < \gamma^*(\lambda)$. One can observe that fewer edges are needed to achieve percolation than in Figure 11.5.

such that the existence of an infinite connected component in the edge percolation model implies the existence of an infinite connected component in the original graph. Although the edges of the discrete model are not finite-range dependent, we show that the probability of having a collection of n closed edges in the discrete model decreases exponentially as q^n, where q can be made arbitrarily small by appropriate choice of the parameters, and therefore the existence of an infinite connected component follows from a Peierls argument.

We describe the construction of the discrete model first, then we prove percolation of this model, and finally we show the final result by coupling it with the SINR model.

Mapping on a lattice

If we set $\gamma = 0$, we obtain a fixed radius Poisson Boolean model with radius r_b given by (11.11).

We consider next a supercritical Boolean model $\mathcal{B}(\lambda, r_b)$ with radius r_b where the node density λ is higher than the critical value λ_c. By rescaling the model, we can establish that the critical radius for a fixed density λ is

$$r^*(\lambda) = \sqrt{\frac{\lambda_c}{\lambda}} r_b < r_b.$$

Therefore, a Boolean model $\mathcal{B}(\lambda, r)$ with density λ and radius r satisfying $r^*(\lambda) < r < r_b$, is still supercritical.

We map this latter model on a discrete percolation model as described in Subsection 11.2.3, and use the same definition for the random field A_a indexed by the edges of $\mathbb{L}^2$.

We next define a second random field, B_a, indexed again by the edges in $\mathbb{L}^2$. We first define $\tilde{l}$, a shifted version of l, as follows:

$$\tilde{l}(x) = \begin{cases} l(0) & x \leq \frac{\sqrt{10}d}{4}, \\ l(x - \frac{\sqrt{10}d}{4}) & x > \frac{\sqrt{10}d}{4}. \end{cases}$$

We define the *shot-noise* processes I and $\tilde{I}$ as follows:

$$I(\boldsymbol{z}) = \sum_k l(\|\boldsymbol{z} - \boldsymbol{x}_k\|)$$

and

$$\tilde{I}(\boldsymbol{z}) = \sum_k \tilde{l}(\|\boldsymbol{z} - \boldsymbol{x}_k\|),$$

where $\boldsymbol{z} \in \mathbb{R}^2$ is an arbitrary point, and where the sum is over all points of the Poisson process X. Note that the shot-noises are random variables, since they depend on the random position of the points of X.

We define now the second random field B_a as taking the value 1 if the value of the shot-noise $\tilde{I}(\boldsymbol{z}_a)$ does not exceed a certain threshold M, and 0 otherwise. As the distance between any point $\boldsymbol{z}$ inside the rectangle $R(\boldsymbol{z}_a) = [x_a - 3d/4, x_a + 3d/4] \times [y_a - d/4, y_a + d/4]$ and its center $\boldsymbol{z}_a$ is at most $\sqrt{10}d/4$, the triangle inequality implies that $\|\boldsymbol{z}_a - \boldsymbol{x}_k\| \leq \sqrt{10}d/4 + \|\boldsymbol{z} - \boldsymbol{x}_k\|$, and thus that $I(\boldsymbol{z}) \leq \tilde{I}(\boldsymbol{z}_a)$ for all $\boldsymbol{z} \in R(\boldsymbol{z}_a)$. Therefore, $B_a = 1$ implies that $I(\boldsymbol{z}) \leq M$ for all $\boldsymbol{z} \in R(\boldsymbol{z}_a)$. Later, we will make an appropriate choice for first d and then M.

Percolation in the lattice

For any edge a of $\mathbb{L}^2$, we call the edge *open* if the product $C_a = A_a B_a = 1$, that is, if both of the following events occur: there exist crossings in the rectangle $R(\boldsymbol{z}_a)$ as described above, *and* the shot

noise is bounded by M for all points inside $R(\boldsymbol{z}_a)$. An edge a that is not open is *closed*. We want to show that for appropriate choice of the parameters M and d, there exists an infinite connected component of open edges at the origin, with positive probability.

To do this, we need an exponential bound on the probability of a collection of n closed edges. Most of the difficulty of obtaining this resides in the infinite range dependencies introduced by the random variables B_i's. A careful application of Campbell's theorem will take care of this.

Consider any collection of n edges $a_1, \ldots, a_n$. To keep the notation simple, we write $A_{a_i} = A_i$, $B_{a_i} = B_i$ and $C_{a_i} = C_i$, $i = 1, \ldots, n$. In Proposition 11.3, we will prove that the probability that all these edges are closed simultaneously, decreases exponentially with n. To do this, we first prove this for the fields A and B.

Proposition 11.1. *Let $\{a_i\}_{i=1}^n$ be a collection of n distinct edges, and let $\{A_i\}_{i=1}^n$ be the random variables of the field A associated with them. Then there exists a constant $q_A < 1$, independent of the particular collection, such that*

$$\mathbb{P}(A_1 = 0, A_2 = 0, \ldots, A_n = 0) \leq q_A^n.$$

Furthermore, for any $\varepsilon > 0$, one can choose d large enough so that $q_A \leq \varepsilon$.

Proof: This propositions follows directly from the observation that it is always possible to find a subset of indices $\{k_j\}_{j=1}^m$ with $1 \leq k_j \leq n$ for each j, such that the variables $\{A_{k_j}\}_{j=1}^m$ are independent, and such that $m \geq n/4$. Therefore we have

$$\begin{aligned}
\mathbb{P}(A_1 = 0, A_2 = 0, \ldots, A_n = 0) &\leq \mathbb{P}(A_{k_1} = 0, A_{k_2} = 0, \ldots, A_{k_m} = 0) \\
&= \mathbb{P}(A_1 = 0)^m \\
&\leq \mathbb{P}(A_1 = 0)^{\frac{n}{4}} \\
&\equiv q_A^n.
\end{aligned}$$

Furthermore, since $q_A = \mathbb{P}(A_1 = 0)^{1/4}$, it follows from [29, Corollary 4.1] that q_A tends to zero when d tends to infinity.

■

Proposition 11.2. *Let $\{a_i\}_{i=1}^n$ be a collection of n distinct edges, and let $\{B_i\}_{i=1}^n$ be the random variables of the field B associated with them. Then there exists a constant $q_B < 1$, independent of the particular collection, such that*

$$\mathbb{P}(B_1 = 0, B_2 = 0, \ldots, B_n = 0) \leq q_B^n.$$

Furthermore, for any $\varepsilon > 0$ and fixed d, one can choose M large enough so that $q_B \leq \varepsilon$.

Proof: To simplify notation, we denote by $\boldsymbol{z}_i$ the center $\boldsymbol{z}_{a_i}$ of the edge a_i. By Markov's inequality, we have for any $s \geq 0$,

$$\begin{aligned}
\mathbb{P}(B_1 = 0, B_2 = 0, \ldots, B_n = 0) &= \mathbb{P}\left(\tilde{I}(\boldsymbol{z}_1) > M, \tilde{I}(\boldsymbol{z}_2) > M, \ldots, \tilde{I}(\boldsymbol{z}_n) > M\right) \\
&\leq \mathbb{P}\left(\sum_{i=1}^n \tilde{I}(\boldsymbol{z}_i) > nM\right) \\
&\leq e^{-snM}\mathbb{E}\left[e^{s\sum_{i=1}^n \tilde{I}(\boldsymbol{z}_i)}\right].
\end{aligned}$$

Using Campbell's theorem (see e.g. [23, p. 28]) applied to the function

$$f(\boldsymbol{x}) = \sum_{i=1}^n \tilde{l}(\|\boldsymbol{z}_i - \boldsymbol{x}\|)$$

we obtain

$$\mathbb{E}\left[e^{s\sum_{i=1}^{n}\tilde{I}(\boldsymbol{z}_i)}\right] = \exp\left(\lambda\int_{\mathbb{R}^2}(e^{s\sum_{i=1}^{n}\tilde{l}(\|\boldsymbol{x}-\boldsymbol{z}_i\|)}-1)\boldsymbol{dx}\right). \tag{11.12}$$

We need to estimate the exponent $s\sum_{i=1}^{n}\tilde{l}(\|\boldsymbol{x}-\boldsymbol{z}_i\|)$. As $\{\boldsymbol{z}_i\}$ are centers of edges, they are located on a square lattice with edge length $d/\sqrt{2}$. So, if we consider the square in which $\boldsymbol{x}$ is located, the contribution to $\sum_{i=1}^{n}\tilde{l}(\|\boldsymbol{x}-\boldsymbol{z}_i\|)$ coming from the four corners of this square is at most equal to 4, since $\tilde{l}(x)\leq 1$. Around this square, there are 12 nodes, each located at distance at least $d/\sqrt{2}$ from $\boldsymbol{x}$. Further away, there are 20 other nodes at distance at least $2d/\sqrt{2}$, and so on. Consequently,

$$\begin{aligned}\sum_{i=1}^{n}\tilde{l}(\|\boldsymbol{x}-\boldsymbol{z}_i\|) &\leq \sum_{i=1}^{\infty}\tilde{l}(\|\boldsymbol{x}-\boldsymbol{z}_i\|)\\ &\leq 4+\sum_{k=1}^{\infty}(4+8k)\tilde{l}\left(\frac{kd}{\sqrt{2}}\right)\equiv K.\end{aligned}$$

Now Assumption 4 in Subsection 11.3.4 on l can easily be extended to $\tilde{l}$, and we clearly have

$$\int_{y}^{\infty}x\tilde{l}(x)dx<\infty \text{ for some } y>0. \tag{11.13}$$

Using the integral criterion and (11.13), we conclude that the sum converges and thus $K<\infty$.

The computation made above holds for any $s\geq 0$. We now take $s=1/K$, so that $s\sum_{i=1}^{n}\tilde{l}(\|\boldsymbol{x}-\boldsymbol{z}_i\|)\leq 1$, for all $\boldsymbol{x}$. Furthermore, since $e^x-1<2x$ for all $x\leq 1$ we have

$$e^{s\sum_{i=1}^{n}\tilde{l}(\|\boldsymbol{x}-\boldsymbol{z}_i\|)}-1<2s\sum_{i=1}^{n}\tilde{l}(\|\boldsymbol{x}-\boldsymbol{z}_i\|)=\frac{2}{K}\sum_{i=1}^{n}\tilde{l}(\|\boldsymbol{x}-\boldsymbol{z}_i\|).$$

Substituting this in (11.12), we obtain

$$\begin{aligned}\mathbb{E}\left[e^{\sum_{i=1}^{n}\tilde{I}(\boldsymbol{z}_i)/K}\right] &\leq \exp\left(\lambda\int_{\mathbb{R}^2}\frac{2}{K}\sum_{i=1}^{n}\tilde{l}(\|\boldsymbol{x}-\boldsymbol{z}_i\|)\boldsymbol{dx}\right)\\ &= \exp\left(\frac{2n\lambda}{K}\int_{\mathbb{R}^2}\tilde{l}(\|\boldsymbol{x}\|)\boldsymbol{dx}\right)\\ &= \left[\exp\left(\frac{2\lambda}{K}\int_{\mathbb{R}^2}\tilde{l}(\|\boldsymbol{x}\|)\boldsymbol{dx}\right)\right]^n.\end{aligned}$$

Putting things together, we have that

$$\begin{aligned}&\mathbb{P}\left(\tilde{I}(\boldsymbol{z}_1)>M,\tilde{I}(\boldsymbol{z}_2)>M,\ldots,\tilde{I}(\boldsymbol{z}_n)>M\right)\\ &\leq e^{-snM}\left[e^{s\sum_{i=1}^{n}\tilde{I}(\boldsymbol{z}_i)}\right]\\ &\leq e^{-nM/K}\left[\exp\left(\frac{2\lambda}{K}\int_{\mathbb{R}^2}\tilde{l}(\|\boldsymbol{x}\|)\boldsymbol{dx}\right)\right]^n\\ &= q_B^n,\end{aligned}$$

where q_B is defined as

$$q_B\equiv\exp\left(\frac{2\lambda}{K}\int\tilde{l}(\|\boldsymbol{x}\|)\boldsymbol{dx}-\frac{M}{K}\right). \tag{11.14}$$

Finally, it is easy to observe that this expression tends to zero when M tends to infinity (for fixed d and hence, since K depends only on d, for fixed K).

■

We next combine the two propositions, in order to obtain a similar result for the field C.

Proposition 11.3. *Let $\{a_i\}_{i=1}^n$ be a collection of n distinct edges, and let $\{C_i\}_{i=1}^n$ be the random variables of the field C associated to them. Then there exists a constant $q_C < 1$, independent of the particular collection, such that*

$$\mathbb{P}(C_1 = 0, C_2 = 0, \ldots, C_n = 0) \leq q_C^n.$$

Furthermore, for any $\varepsilon > 0$, one can choose d and M so that $q_C \leq \varepsilon$.

Proof: For convenience in the following calculations, we introduce the notation $\bar{A}_i = 1 - A_i$ and $\bar{B}_i = 1 - B_i$. First observe that

$$1 - C_i = 1 - A_i B_i \leq (1 - A_i) + (1 - B_i) = \bar{A}_i + \bar{B}_i.$$

Let us denote by $p(n)$ the probability that we want to bound, and let $(k_i)_{i=1}^n$ be a binary sequence (i.e. $k_i = 0$ or 1) of length n. We denote by $\mathcal{K}$ the set of the 2^n such sequences. Then we can write

$$\begin{aligned}
p(n) &= \mathbb{P}\left(C_1 = 0, C_2 = 0, \ldots, C_n = 0\right) \\
&= \mathbb{E}\left[(1 - C_1)(1 - C_2)\ldots(1 - C_n)\right] \\
&\leq \mathbb{E}\left[(\bar{A}_1 + \bar{B}_1)(\bar{A}_2 + \bar{B}_2)\ldots(\bar{A}_n + \bar{B}_n)\right] \\
&= \sum_{(k_i)\in\mathcal{K}} \mathbb{E}\left[\prod_{i:k_i=0} \bar{A}_i \prod_{i:k_i=1} \bar{B}_i\right] \\
&\leq \sum_{(k_i)\in\mathcal{K}} \sqrt{\mathbb{E}\left[\prod_{i:k_i=0} \bar{A}_i^2\right] \mathbb{E}\left[\prod_{i:k_i=1} \bar{B}_i^2\right]} \\
&= \sum_{(k_i)\in\mathcal{K}} \sqrt{\mathbb{E}\left[\prod_{i:k_i=0} \bar{A}_i\right] \mathbb{E}\left[\prod_{i:k_i=1} \bar{B}_i\right]},
\end{aligned}$$

where the two last inequalities follow respectively from Schwartz's inequality and from the observation that $\bar{A}_i^2 = \bar{A}_i$ and $\bar{B}_i^2 = \bar{B}_i$. Applying Propositions 11.1 and 11.2, we can bound each expectation in the sum. We have thus

$$\begin{aligned}
p(n) &\leq \sum_{(k_i)\in\mathcal{K}} \sqrt{\prod_{i:k_i=0} q_A \prod_{i:k_i=1} q_B} \\
&= \sum_{(k_i)\in\mathcal{K}} \prod_{i:k_i=0} \sqrt{q_A} \prod_{i:k_i=1} \sqrt{q_B} \\
&= \left(\sqrt{q_A} + \sqrt{q_B}\right)^n \\
&\equiv q_C^n.
\end{aligned}$$

Choosing first d large, and then M appropriately, we can make q_C is smaller than any given ε.

■

With Proposition 11.3, the existence of percolation in our dependent bond percolation model follows from standard arguments. Indeed, with our exponential bound in Proposition 11.3, we can apply the usual Peierls argument to establish the existence of percolation for appropriate M and d. We recall the Peierls argument here, it follows exactly the same line as the proof of Theorem 2.1) in Chapter 2.

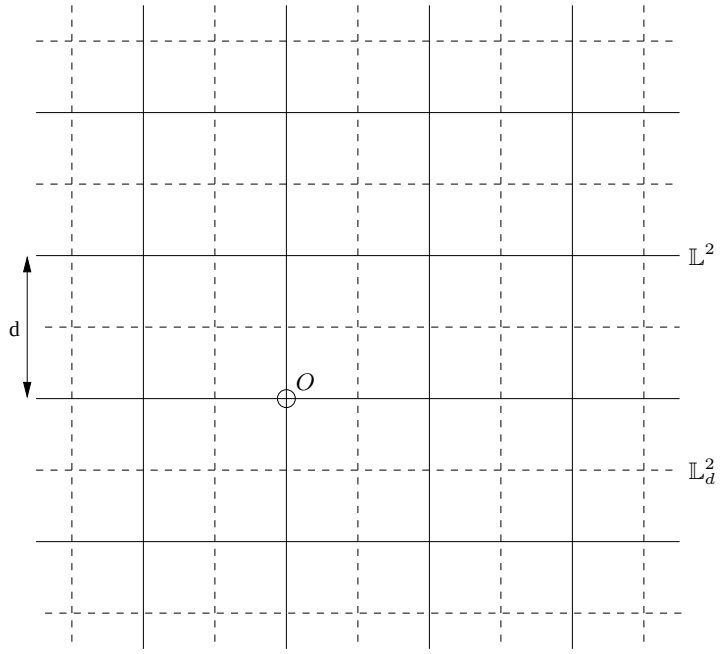

Figure 11.9: Lattice $\mathbb{L}^2$ (plain) and its dual $\mathbb{L}_d^2$ (dashed)

We consider the dual lattice $\mathbb{L}_d^2$ of $\mathbb{L}^2$, obtained by placing a vertex at the center of each face of $\mathbb{L}^2$, and an edge between two vertices whenever the two corresponding faces of $\mathbb{L}^2$ have a common edge (see Figure 11.9). We can set $\mathbb{L}_d^2 = \mathbb{L}^2 + (d/2, d/2) = \{\langle \boldsymbol{v}, \boldsymbol{w}\rangle \in \mathbb{R}^4 : \langle \boldsymbol{v} - (d/2, d/2), \boldsymbol{w} - (d/2, d/2)\rangle \in \mathbb{L}^2\}$.

Let $\gamma(n)$ be the number of self-avoiding circuits in the dual lattice L_d^2 surrounding the origin and of length n. As we have seen in Chapter 2, $\gamma(n) \leq 4n3^{n-1}$.

We can now bound the probability to find a closed circuit surrounding the origin:

$$\begin{aligned}
\mathbb{P}\left(\exists \text{ closed circ.}\right) &\leq \sum_{\text{circuit } g \text{ of length at least } 4} \mathbb{P}_p\left(g \text{ is closed}\right) \\
&= \sum_{n=4}^{\infty} \sum_{\text{circuit } g \text{ of length } n} \mathbb{P}_p\left(g \text{ is closed}\right) \\
&\leq \sum_{n=4}^{\infty} \gamma(n) q_C^n \\
&\leq \frac{4q_C}{3} \sum_{n=4}^{\infty} n \left(3q_C\right)^{n-1} \\
&\leq \frac{4q_C}{3} \left(\frac{1}{1-3q_C}\right)^2.
\end{aligned}$$

The latter expression is smaller than 1 if $q_C < (11 - 2\sqrt{10})/27$. It is thus sufficient to have $q_A, q_B < ((11 - 2\sqrt{10})/27)^2/2$. We must thus choose d and M such that these conditions are verified. First, we know that we can choose d sufficiently large so that $q_A < (11 - 2\sqrt{10})/27)^2/2$. Second, it is easy to see in Equation (11.14) that one can choose M sufficiently large so that $q_B < (11 - 2\sqrt{10})/27)^2/2$. It is thus possible to find appropriate parameters M and d to have percolation in the bond percolation model.

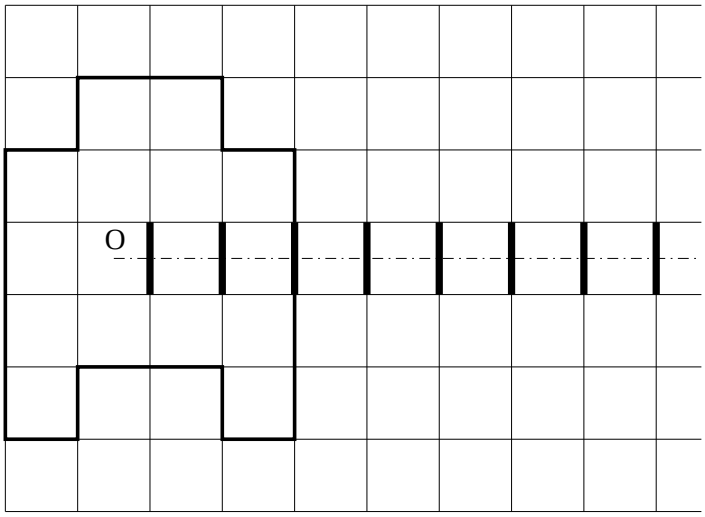

Figure 11.10: Circuits in the dual lattice that surround the origin. All such circuits must contain at least one of the bold edges.

Percolation of the SINR model

To conclude the proof we need to show that percolation of C_a implies percolation in the SINR model, with appropriate γ. If $B_a = 1$, the interference level in the rectangle $R(\boldsymbol{z}_a)$ is at most equal to M. Therefore, for two nodes $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$ in $R(\boldsymbol{z}_a)$ such that $\|\boldsymbol{x}_i - \boldsymbol{x}_j\| \leq r$, we have

$$\begin{aligned} \frac{Pl(\|\boldsymbol{x}_i - \boldsymbol{x}_j\|)}{N_0 + \gamma \sum_{k \neq i,j} Pl(\|\boldsymbol{x}_k - \boldsymbol{x}_j\|)} &\geq \frac{Pl(\|\boldsymbol{x}_i - \boldsymbol{x}_j\|)}{N_0 + \gamma PM} \\ &\geq \frac{Pl(r)}{N_0 + \gamma PM}. \end{aligned} \tag{11.15}$$

As $r < r_b$ and as l is strictly decreasing, we choose

$$\gamma = \frac{N_0}{PM} \left(\frac{l(r)}{l(r_b)} - 1 \right) > 0, \tag{11.16}$$

yielding

$$\frac{Pl(r)}{N_0 + \gamma PM} = \frac{Pl(r_b)}{N_0} = \beta. \tag{11.17}$$

Therefore, there exists a positive value of γ such that any two nodes separated by a distance less than r are connected in the SINR model. This means that in the rectangle $R(\boldsymbol{z}_a)$ all connections of $\mathcal{B}(\lambda, r)$ also exist in the SINR model.

Finally, since all edges a of the infinite cluster of the discrete model are such that $A_a = 1$ and $B_a = 1$, and since an an infinite cluster of open edges in the $\mathbb{L}^2$ implies an infinite cluster in the Boolean model $\mathcal{B}(\lambda, r)$ as we saw in Subsection 11.2.3, this means that the crossings also exist in the SINR model, and thus form an infinite connected component.

Here follows a summary of what is known about the set of couples (λ, γ) for which percolation occurs:

- no percolation occurs when $\lambda < \lambda_c$,
- no percolation occurs when $\gamma > 1/\beta$,

- when $\lambda > \lambda_c$, there exists $\gamma^*(\lambda) > 0$ such that percolation occurs whenever $\gamma < \gamma^*(\lambda)$, and
- there exists $c_1 < \infty$ and $\lambda' < \infty$ such that $\gamma^*(\lambda) \leq c_1/\lambda$ for all $\lambda > \lambda'$.

The last property follows from [9].

12

Transport capacity in wireless multihop networks

12.1 Introduction

Chapter 11 concluded with a model of multi-hop wireless networks (STIRG) where, for a given value β of the signal to noise ratio threshold at the receiver at the and node spatial distribution in $\mathbb{R}^2$, one deduced a network topology. Shannon's information theory relates the signal to noise ratio SNR on a channel to its capacity C (in bits per second) as

$$C = \frac{W}{2} \log_2 (1 + SNR) \tag{12.1}$$

where W is the bandwidth (in Hertz) and SNR is the signal to noise ratio. Consequently, imposing a minimal value of β on a given link amounts to require a minimal capacity $C(\beta)$ where SNR is replaced by β in (12.1). Now, (12.1) only holds for a single-user channel (i.e. only one source and one destination), and is not applicable as such to wireless multi-hop networks, at least in an information theoretic sense. Indeed, if we write

$$SNR = \frac{P_i L(\boldsymbol{x}_i - \boldsymbol{x}_j)}{N_0 + \sum_{k \neq i,j} P_k L(\boldsymbol{x}_k - \boldsymbol{x}_j)} \tag{12.2}$$

for the signal to noise ratio from Node i to Node j, with the same notations as in the previous chapter, it amounts to treat interferences are as noise. However a clever decoding may make a much better usage of that information, and the capacity given by (12.1) and (12.2) is thus a lower bound on the actual, truly information theoretic value. However, the determination of the information theoretic capacity of a multi-user channel remains an open problem, even with as few as three nodes, whereas using the rather crude and pessimistic assumption that noise is nothing but additional noise enables to have results for very large networks.

In this chapter, we will look at the dual of the problem tackled at the end of Chapter 11. That is, if instead of requiring a given capacity (in the "non-information-theoretic sense" described above) on any possible link, and looking at the resulting topology, we will first impose a topology and then determine the capacity. Contrary to the single user case, we now need to impose a traffic matrix

expressing the origin-destination pairs for each communication. Moreover, we will explicitly make the dependence of the capacity on the distance between the source and destination. This leads to the notion of *transport capacity*, introduced by Gupta and Kumar [22], and which is the amount of bits transported by time unit and distance unit by the network: it is thus expressed in bits/(sec $\times$ meter).

As a benchmark, and following [22], we will assume that the network is made of n nodes, and assume a uniform traffic matrix. Each node is a source and picks a destination at random, with the same probability, irrespectively of the distance between them. Each communication must receive the same minimal throughput; a lower bound on the transport capacity of the network is thus n times the individual throughput enjoyed by each connection when $n \to infty$. This definition of minimal throughput that *each* pair origin-destination must have forces the network to be fully connected.

Typically, there are two ways of letting the number of nodes n tend to infinity. One can either keep the area on which the network is deployed constant, and make the node density λ tend to infinity (*dense* networks); or one can keep the node density λ constant, and increase the area to infinity (*extended* networks). The first bound was found by [22], for dense networks. However, as the node density becomes larger and larger, the singularity at the origin of the attenuation function needs to be accounted for [12].

For extended networks, near field effects do not play a fundamental role and the power received by every node can be bounded without affecting the final result. In this case, the work in [38] presents an information theoretic bound of $\Theta(1/\sqrt{n})$ bit/sec per node[1], for arbitrarily located nodes satisfying a minimum distance constraint, and a power attenuation function that exhibits a power law behavior with exponent $\alpha > 6$, or an exponential attenuation. When nodes are randomly located, the work in [28] shows an upper bound of $O(1/n^{1/2-1/\alpha})$ that holds for $\alpha > 2$.

Constructive strategies proposed for networks of randomly located nodes [17, 22, 26, 36] achieve only $\Omega(1/\sqrt{n \log n})$ per-node bit rate, somehow suggesting that at least a $\sqrt{\log n}$ factor is the price to pay for randomness. We will see indeed that it is the case. Instead, using percolation theory [15], it is possible to achieve a per-node throughput capacity of $\Omega(1/\sqrt{n})$ in random networks. Hence, nodes in a random network can transmit at the same rate than nodes in an arbitrary network and there is no price to pay —at least asymptotically— for the additional randomness present in the system. The latter result holds in the same setting of [22], i.e., for dense networks, as long as near field effects are negligible. It also holds in the more general case of extended networks with bounded transmitted and received power, assuming only a power decay law whose tail exhibits a power law behavior with exponent $\alpha > 2$, or an exponential attenuation (which is typical if there is absorption in the medium [14]).

A major difference in the approaches to compute the throughput scaling between is whether one should first assume to be fully connected or not, since eventually *each* pair origin-destination must be guaranteed the same minimal rate. We begin by a brief digression on the extra cost required to fully connect the network, instead of just requiring it to be supercritical.

12.2 Full connectivity vs percolation

Let us return to the Poisson Boolean model, but with n nodes over a finite area. We will start with an intuitive – but not rigorous – reasoning following [11], that enables to link the probability that the network is fully connected with the function $\theta(\pi\lambda r^2)$. A rigorous and detailed version of this reasoning can be found in [21]. It builds upon the results of [33] and [32].

We assume without loss of generality that nodes are distributed according to a Poisson point process

[1] We use the following standard notation: $f = O(g)$ if $\limsup_{n\to\infty} \frac{f(n)}{g(n)} < +\infty$; $f = \Omega(g)$ if $g = O(f)$; $f = \Theta(g)$ if $f = O(g)$ and $g = O(f)$. Thus all $O(\cdot)$ results are upper bounds, $\Omega(\cdot)$ results are lower bounds and $\Theta(\cdot)$ results are sharp scaling estimates.

of spatial density $\lambda = 1$, and that the network area is equal to n. The total number of nodes in the network is therefore approximately equal to n.

We start with a result from Penrose and Pisztora [33], who proved that for a large but finite area, the fraction of connected nodes is always close to the deterministic function $\theta(\lambda r^2)$. Furthermore, when the number of nodes becomes large, the events that different nodes are connected become almost independent. Therefore

$$\mathbb{P}[\text{the network is connected}] \simeq \theta^n(\lambda r^2). \tag{12.3}$$

In order to have the above probability tend to one when n tends to infinity, we must have $\theta(\lambda r^2) \to 1$. This is only possible if $\lambda r^2 = \lambda r^2(n)$ increases with n.

To compute how fast $r(n)$ should grow, we derive an approximation of the function $\theta(\lambda r^2)$ when λr^2 is large. According to Theorem 6.4 in [29], p. 173, when the average node degree is large, the ratio between the probability that a node is disconnected, and the probability that a node is isolated (i.e. has degree zero) tends to one. Therefore, asymptotically, all disconnected nodes are isolated nodes. The probability that a node is isolated is easy to compute:

$$\mathbb{P}[\text{a given node is isolated}] = \exp(-\pi\lambda r^2).$$

Therefore, when $\pi\lambda r^2$ is large, $\theta(\lambda r^2) \simeq 1 - \exp(-\pi\lambda r^2)$.

Using this approximation in (12.3), we obtain

$$\mathbb{P}[\text{the network is connected}] \simeq [1 - \exp(-\pi\lambda r^2(n))]^n.$$

The latter expression tends to one provided $n \exp(-\pi\lambda r^2(n))$ tends to zero. Taking the logarithm of this expression leads to the conclusion that the network is asymptotically connected if and only if

$$\pi\lambda r^2(n) - \log n \to \infty.$$

This means that the average node degree must grow approximately like $\log n$, when the number of nodes in the network increases. However, it appears in this intuitive derivation of the critical range for full connectivity that the *most isolated node* is determining the result. In fact, the network becomes connected when the last node joins the network. This means that full connectivity is not really a global property of the network; it just answers the question "how isolated is the most isolated node". As we just saw above, the distance to the first neighbor of the most isolated node increases with the number of nodes. But this is a pure statistical effect: we are taking a set of randomly distributed distances, and pick the largest one. As we increase the sample set, the largest element becomes longer and longer. This explains why the range of the nodes has to increase, even though the node density remains constant.

12.3 Upper bound on transport capacity in dense networks

We return to the problem of determining the capacity of the model where each node picks a destination at random, and transmit data to it (uniform traffic matrix). The computation of the transport capacity in an information theoretic sense is a challenging open problem, as we saw in Section 12.1.

However, in the Gupta and Kumar model [22], an upper bound can easily be derived as follows.

We consider a network made of n nodes distributed on a unit square $[0,1] \times [0,1]$, with a power law attenuation function $L(\boldsymbol{x}) = l(\|x\|) = x^{-\alpha}$. For simplicity, we take $P_i = P$ for all nodes i and that the condition for having a direct transmission between an emitter i and a receiver j is that the Signal to Noise Ratio given by (12.2) is larger than or equal to $\beta > 1$.

What is the maximal throughput $T(n)$ that can be reached ? We follow the same reasoning as Gupta and Kumar [22].

Suppose bit b traverses $\overline{h}(b)$ hops from source to destination, with δ_b^h denoting the distance traversed during the hth hop. Suppose that $nT(n)$ bits are served all together during 1 second. Let L be the average distance of a source-destination pair: since the area is a unit square, $L = O(1)$ for unit area. The sum of the distances covered by all bits served in one second is therefore such that

$$nT(n)L \leq \sum_{b=1}^{nT(n)} \sum_{h=1}^{\overline{h}(b)} \delta_b^h \tag{12.4}$$

With $r(i) = j$ denoting the receiver of node i, we get from condition (12.2) that for any $k \neq i, r(i)$,

$$L(\boldsymbol{x}_i - \boldsymbol{x}_{r(i)}) \geq \beta L(\boldsymbol{x}_k - \boldsymbol{x}_{r(i)}),$$

whence

$$\|\boldsymbol{x}_k - \boldsymbol{x}_{r(i)}\| \geq \beta^{1/\alpha} \|\boldsymbol{x}_i - \boldsymbol{x}_{r(i)}\|. \tag{12.5}$$

The triangle inequality yields therefore that

$$\begin{aligned} \|\boldsymbol{x}_{r(i)} - \boldsymbol{x}_{r(k)}\| &\geq \beta^{1/\alpha} \|\boldsymbol{x}_i - \boldsymbol{x}_{r(i)}\| - \|\boldsymbol{x}_{r(k)} - \boldsymbol{x}_k\| \\ \|\boldsymbol{x}_{r(k)} - \boldsymbol{x}_{r(i)}\| &\geq \beta^{1/\alpha} \|\boldsymbol{x}_k - \boldsymbol{x}_{r(k)}\| - \|\boldsymbol{x}_{r(i)} - \boldsymbol{x}_i\|, \end{aligned}$$

which add up to give

$$\|\boldsymbol{x}_{r(i)} - \boldsymbol{x}_{r(k)}\| \geq \frac{\beta^{1/\alpha} - 1}{2} \left(\|\boldsymbol{x}_i - \boldsymbol{x}_{r(i)}\| + \|\boldsymbol{x}_k - \boldsymbol{x}_{r(k)}\| \right).$$

This shows that the distance between two receivers must be at least $(\beta^{1/\alpha} - 1)$ the average distance between each transmitter and receiver. Each transmission consumes therefore a circular footprint of radius $(\beta^{1/\alpha} - 1)$ times the distance travelled in one hop. Now the total surface covered by all these footprints cannot exceed to the total area on which the network is deployed, which is 1. Neglecting boundary effects, and summing over all the $nT(n)$ bits sent in one second, we get therefore the following packing constraint:

$$\sum_{b=1}^{nT(n)} \sum_{h=1}^{\overline{h}(b)} \pi \left((\beta^{1/\alpha} - 1) \delta_b^h \right)^2 \leq 1. \tag{12.6}$$

Let $H = \sum_{b=1}^{nT(n)} \overline{h}(b)$ be the total number of hops traversed by all $nT(n)$ bits in 1 second. Because no more than $n/2$ nodes can transmit simultaneously at any given time, $H \leq n/2$.

Now, by convexity of the quadratic function,

$$\left(\sum_{b=1}^{nT(n)} \sum_{h=1}^{\overline{h}(b)} \frac{1}{H} \delta_b^h \right)^2 \leq \sum_{b=1}^{nT(n)} \sum_{h=1}^{\overline{h}(b)} \frac{1}{H} \left(\delta_b^h \right)^2,$$

which, combined with (12.6) yields that

$$\left(\sum_{b=1}^{nT(n)} \sum_{h=1}^{\overline{h}(b)} \frac{1}{H} \delta_b^h \right)^2 \leq \frac{1}{H\pi(\beta^{1/\alpha} - 1)^2},$$

and finally, with (12.4), that

$$T(n) \leq \frac{\sqrt{H}}{\sqrt{\pi}(\beta^{1/\alpha} - 1)nL} \leq \frac{1}{\sqrt{\pi}(\beta^{1/\alpha} - 1)L} \frac{1}{\sqrt{n}}.$$

This establishes an upper bound on the throughput per source-destination pair $T(n) = O(1/\sqrt{n})$.

At this point, we should mention that the shape of the attenuation function is particularly important. Indeed, in a dense network, near field effects dominate, and if we replace an unrealistic pure power law function $l(\|x\|) = x^{-\alpha}$ by a bounded power law function $l(\|x\|) = \min\{1, x^{-\alpha}\}$, the he throughput per source-destination pair is worse, as $T(n) = O(1/n)$ [12]. However, in extended networks, both attenuation functions lead to $T(n) = O(1/\sqrt{n})$ [38].

12.4 Lower bound on transport capacity in dense networks: the full connectivity approach

We now to the problem of determining the capacity of the model where each node picks a destination at random, and transmit data to it (uniform traffic matrix). The node density is kept constant, whereas the network area increases. n denotes the average number of nodes.

This situation has been first studied in [22], where a constructive scheme is found, that achieves a throughput of order $1/\sqrt{n \log n}$. However, if the nodes are placed in a more regular (non-random) fashion, a throughput of order $1/\sqrt{n}$ is achievable. The key for this gap is that we require that *each node* benefits from the same throughput. Of course, this rules out the possibility of having a few disconnected nodes (which would have zero throughput). In this section, we follow an alternative approach, similar to [26], and show how the requirement of full connectivity is responsible for this $\sqrt{\log n}$ factor.

We assume that nodes are distributed according to a Poisson point process of intensity $\lambda = n$ over a square of unit area $[0,1] \times [0,1]$. Hence the average number of nodes in the network is n. Furthermore, we take an unbounded attenuation function l of the form

$$l(d) = d^{-\alpha} e^{-\gamma d}.$$

We divide that network into squares of size $c_n = c\sqrt{\log n / n}$. By adjusting $c > 1$, we can adjust the probability that a square contains at least one point:

$$\mathbb{P}[\text{a square contains at least one point}] = 1 - e^{-n c_n^2} = 1 - n^{-c^2}.$$

As a result, the probability that all squares contain at least one node is

$$\mathbb{P}[\text{all squares contain at least one point}] = \left(1 - n^{-c^2}\right)^{n/c^2 \log n},$$

which tends to 1 as $n \to \infty$. Consequently, *all* squares are populated w.h.p.

Next, we adjust the power P and find TDMA scheduling, such that one node per square can transmit to any destination located within a radius of d squares (in Manhattan distance) with a fixed rate independent of n, for any given integer $d > 0$ and when n is sufficiently large. The proof of this theorem is very similar to that of Theorem 12.4 at the end of this chapter. This TDMA schemes defines therefore $(2d+1)^2$ classes of equivalence, containing each a set of nodes that can transmit simultaneously.

Finally, we can upper bound the number of nodes in a horizontal slab of length 1 and height $c_n = c\sqrt{\log n / n}$, using Chernoff's bound. We find that each such slab contains no more than $2c\sqrt{n \log n}$ nodes w.h.p. The same is true of course for vertical slab of height 1 and width $c_n = c\sqrt{\log n / n}$.

These three properties give us the following simple routing and scheduling scheme. The route from a source to a destination follows first a path connecting nodes in neighboring squares with the horizontal slab containing the source, until it has reached the squares lying in the same vertical slab as the

destination. Then the route switches to a path connecting nodes in neighboring squares with this vertical slab containing, until it has reached the destination.

Now let us consider the number of slots we need to schedule. Because of the TDMA schedule, we need $(2d+1)^2$ times slots to cover all the squares. As there are at most $2c\sqrt{n \log n}$ nodes in the slab whose traffic needs to be relayed, we need at most $(2d+1)^2 \times 2c\sqrt{n \log n}$ slots for the horizontal phase, and the same amount of slots for the vertical phase, which gives a maximum of $4c(2d+1)^2\sqrt{n \log n}$ time slots to serve every node. From this, we can conclude that a per-node throughput capacity of

$$T(n) = \Omega(1/\sqrt{n \log n}) \text{ bit/sec}$$

is achievable in a random dense network. We will see in Section 12.6 that a better lower bound can be achieved by a percolation approach.

12.5 Lower bound on transport capacity in extended networks: the percolation approach

An alternative to the scaling adopted in the previous section is to allow for some empty squarelets, and to keep c constant. This lead to a *site percolation* model on a square grid, where nodes are occupied with probability $p = 1 - \exp(-c^2)$ and empty with probability $1 - p$.

With a constant squarelet size c, there are still disconnected nodes w.h.p. To cope with them, we use a separate time slot, for draining the traffic from disconnected nodes to the connected part. This is done by using longer hops. As this traffic is only local (from remote nodes to the closest connected nodes), it can be shown that the throughput per node during this special time slot decreases slower than $1/\sqrt{n}$. Therefore, as the overall throughput is determined by the minimum of the throughput during the two time slots, introducing this draining phase does not affect the asymptotic behavior of the result. The difference between this strategy and the previous one is that here we treat the problems of transport capacity and connectivity *separately* (in two distinct time slots), whereas by letting c tend to infinity, we solve the capacity *and* connectivity issues simultaneously, with high cost on the throughput.

We detail here this solution built along these lines, following [15], where the main result is the following.

Theorem 12.1. *Assuming a power attenuation function of the type* $l(d) = \min\{1, e^{-\gamma d}/d^{\alpha}\}$ *with* $\alpha > 0, \gamma > 0$ *or* $\alpha > 2, \gamma = 0$*, a per-node throughput capacity of*

$$T(n) = \Omega(1/\sqrt{n}) \textit{ bit/sec}$$

is achievable in a random extended network.

12.5.1 Overview of the solution

The main idea of the scheme achieving the result in Theorem 12.1 is to have a wireless backbone of nodes that carry packets across the network at constant rate, using short hops, and to drain the rest of the traffic to the wireless backbone using single hops of longer length. See Fig. 12.1 for a schematic representation.

Since all nodes transmit at the same power, we expect the longer hops that access the backbone to have a lower bit-rate, due to higher power loss across longer distances; however, one needs to take into account other components that influence the bit-rate: namely, interference, and cost of relay packets from other nodes. It turns out that when all these components are accounted for, the bottleneck is due to the traffic on the backbone.

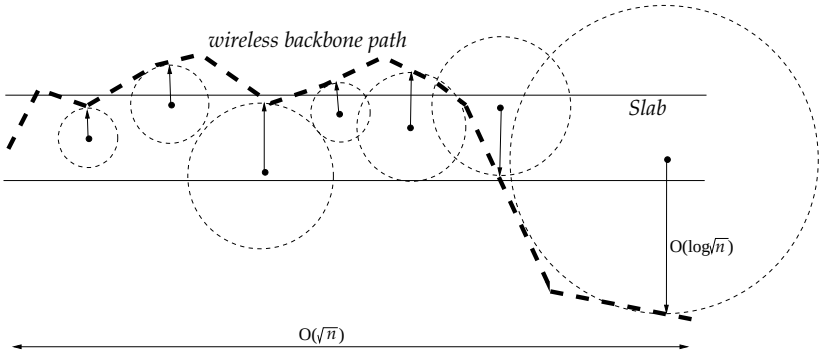

Figure 12.1: Nodes inside a slab of constant width access a path of the wireless backbone in single hops of length at most proportional to $\log\sqrt{n}$. Packets are carried across a distance $\sqrt{n}$ along the wireless backbone in constant-length hops.

The wireless backbone consists of paths of constant length hops. We will show that the interference caused by nodes sufficiently far away along the path is small enough to allow a constant transmission rate along the path. Each path, however, needs to relay packets coming from other nodes that access the backbone in single hops. This traffic is at most proportional to $\sqrt{n}$, if we associate to each path only nodes that are within a slab of constant width that crosses the network area, see Fig. 12.1. Hence, the per-node throughput on the backbone can be only of order $1/\sqrt{n}$.

Now, let us look at the throughput of the nodes that access the wireless backbone in single hops. We will show that these single hops are of length at most proportional to $\log\sqrt{n}$ and can sustain a rate higher than $1/\sqrt{n}$, even in presence of exponential power attenuation. Furthermore, there is no relay burden for nodes accessing the backbone in single hops. It follows that the bottleneck is represented by the nodes on the backbone that transmit at a rate of $1/\sqrt{n}$.

There are three key points in our reasonings: i) there exist paths of constant hop length that cross the entire network forming the wireless backbone, ii) these paths can be put into a one to one correspondence with $\sqrt{n}$ slabs of constant width, each containing at most a constant times $\sqrt{n}$ number of nodes, and iii) these paths are somehow regularly spaced so that there is always one within a $\log\sqrt{n}$ distance factor from any node in the network.

In the following, Theorem 12.2 ensures the existence of many paths using percolation theory arguments. Theorem 12.3 shows that each path in the wireless backbone can transport packets at a constant rate, and that packets can be drained to the backbone at a rate higher than $\sqrt{n}$. Finally, Lemmas 12.3 and 12.4 are needed to bound the number of nodes that access any given path.

12.5.2 Percolation results

In this section we establish the percolation results that are needed to show existence of a cluster of nodes forming the wireless backbone. The objective is to formally construct a mesh of paths that can simultaneously carry information across the network at a constant rate, independent of the number of nodes n. We call this mesh the *highway system*, and will use it to carry packets over most of the distance.

To begin our construction, we divide the area into squares of constant side length c, as depicted in the left-hand of Fig. 12.2. By adjusting c, we can adjust the probability that a square contains at least one point:

$$\mathbb{P}[\text{a square contains at least one point}] = 1 - e^{-c^2} := p. \tag{12.7}$$

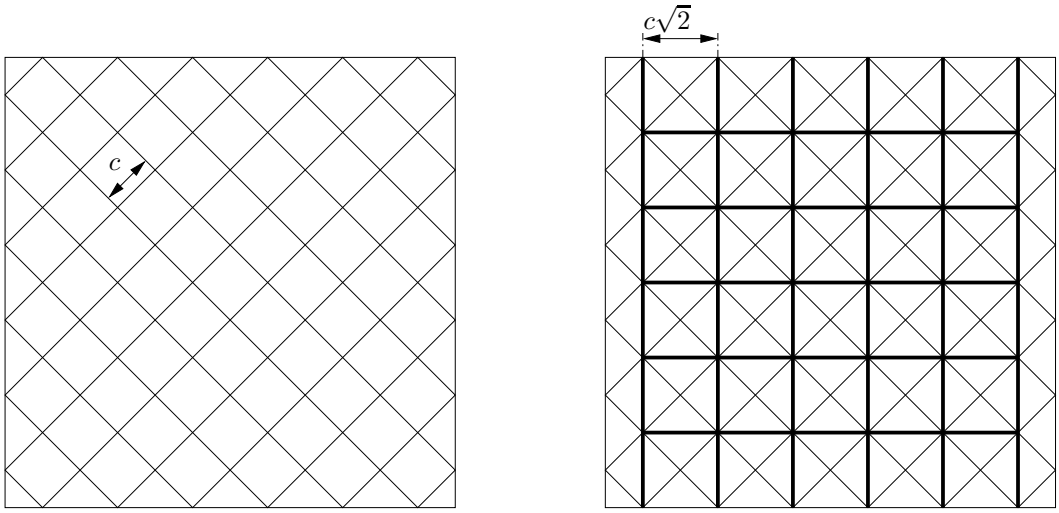

Figure 12.2: Construction of the bond percolation model. We declare each square on the left-hand side of the picture open, if there is at least a Poisson point inside it, closed otherwise. This corresponds to associate an edge to each square, traversing it diagonally, as depicted on the right-hand side of the figure, and declare the edge either open or closed according to the state of the corresponding square.

Object	Size	Lattice size	Nb. of objects
Whole network area	$\sqrt{n} \times \sqrt{n}$	$m \times m$	1×1
Rectangle $\overline{R}_n$	$\sqrt{n} \times \sqrt{2}c\kappa \log \frac{\sqrt{n}}{\sqrt{2}c}$	$m \times \kappa \log m$	$1 \times \frac{m}{\kappa \log m}$
Slab (Lemma 12.4)	$\sqrt{n} \times \frac{\sqrt{2}c}{\beta}$	$m \times \frac{1}{\beta}$	$1 \times \beta m$

Table 12.1: Divisions of the network. Recall that $m = \sqrt{n}/\sqrt{2}c$.

We say that a square is *open* if it contains at least one point, and *closed* otherwise; note that the status of the squares is i.i.d.

We now map our construction to a bond percolation model. We draw an horizontal edge across half of the squares, and a vertical edge across the others, as shown on the right-hand side of Fig. 12.2. In this way we obtain a grid of horizontal and vertical edges, each edge being open, independently of all other edges, with probability p. We call a path *open* (resp. *closed*) if it contains only open (resp. closed) edges. Note that, for c large enough, our construction produces winding open paths that cross the network area, see Fig. 12.3. Next, we turn to the question of how many of these paths there are.

Throughout the proof, we will introduce several subdivisions of the networks and other objects of different sizes. We summarize them in Table 12.1 to ease the reading of the proof.

Number of disjoint crossing paths

Let us divide the network area into horizontal rectangles $\overline{R}_n$, of size $\sqrt{n} \times \sqrt{2}c\kappa \log \frac{\sqrt{n}}{\sqrt{2}c}$, for some constant $\kappa > 0$, see Fig. 12.4. Each of these rectangles has thus lattice size $m \times \kappa \log m$ in the bond percolation model, with $m = \frac{\sqrt{n}}{\sqrt{2}c}$ (as the edges have length $\sqrt{2}c$). We want to show that there exist many disjoint open paths from left to right inside such rectangles.

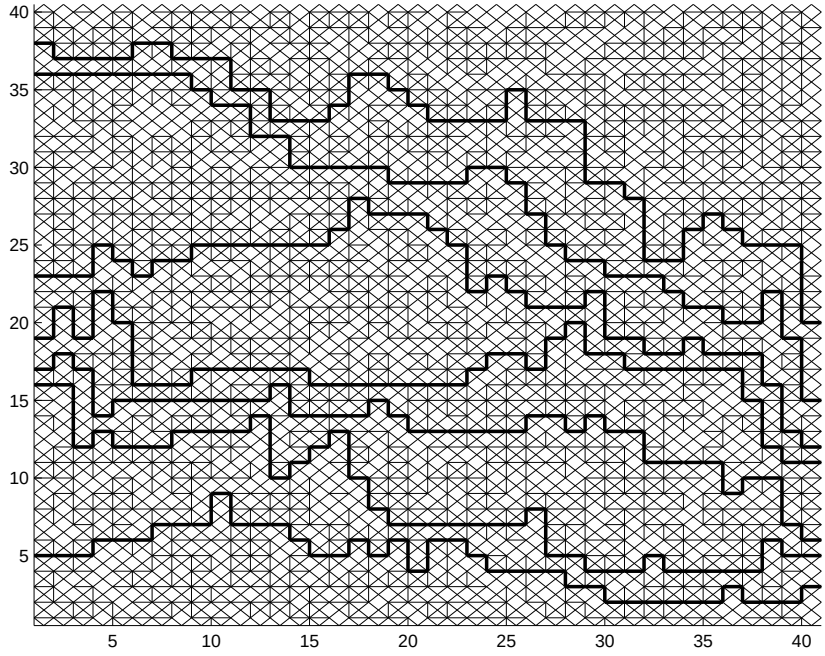

Figure 12.3: Horizontal paths in a 40×40 bond percolation model obtained by computer simulation. Each square is traversed by an open edge with probability p ($p = 0.7$ here). Closed edges are not depicted. We find seven disjoint open path crossing the area from left to right.

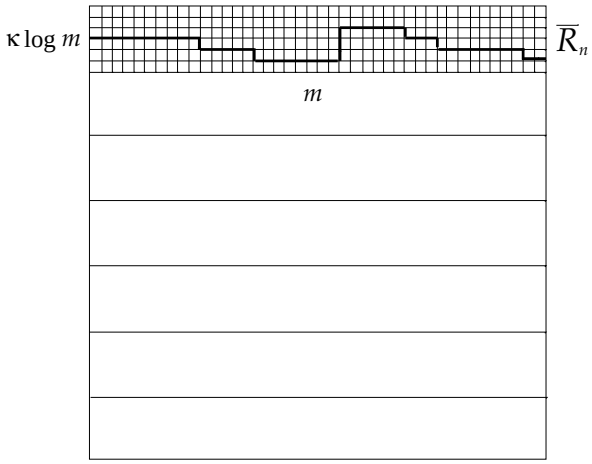

Figure 12.4: The network area is divided into $m/\kappa \log m$ horizontal rectangles of lattice size $m \times \kappa \log m$. A left to right crossing of rectangle $\overline{R}_n$ is shown.

Theorem 12.2. *For any constant $\kappa > 0$ and if c is so large that*

$$c^2 > \log 6 + \frac{2}{\kappa}, \tag{12.8}$$

then there exists a strictly positive constant $\beta = \beta(c, \kappa)$ such that w.h.p. there exist $\beta\kappa \log m = \beta\kappa \log \frac{\sqrt{n}}{\sqrt{2}c}$ disjoint open paths inside each rectangle $\overline{R}_n$, that cross it from left to right.

In order to prove this theorem we recall two preliminary results expressed by the lemmas below. The first lemma was proven in the first part of the proof of Theorem 2.1. The second lemma is simply Lemma 2.1.

Lemma 12.1. *Let S_n be a square lattice of size $n \times n$ and $0 < p \leq 1$. The probability that there exists an open path from the center 0 of S_n to its boundary ∂S_n is upper bounded by*

$$\mathbb{P}_p(0 \leftrightarrow \partial S_n) \leq \frac{4}{3}(3p)^n.$$

Lemma 12.2. *Let R_n be a rectangle embedded in the square lattice $\mathbb{L}^2$. Let A_n be the event that there exists an open path between the left and right sides of R_n and $I_r(A_n)$ the event that there exist r edge-disjoint such LR crossings. We have*

$$1 - \mathbb{P}_p(I_r(A_n)) \leq \left(\frac{p}{p - p'}\right)^r [1 - \mathbb{P}_{p'}(A_n)]$$

for any $0 \leq p' < p \leq 1$.

Proof: We consider bond percolation in the rectangle $\overline{R}_n$, with each edge having probability p' to be open, independently of all other edges. Let the *dual graph* of $\overline{R}_n$ be obtained by placing a vertex in each square of the percolation lattice, and joining two such vertices by an edge whenever the corresponding squares share a side, see Fig. 12.5. An edge of the dual is open if it crosses an open edge of the original lattice, it is closed otherwise. Let A_n be the event of having at least one open path inside $\overline{R}_n$ that crosses it from left to right, and let B_n be the event that a closed path crosses the rectangle vertically in the dual lattice. We have $A_n \bigcap B_n = \emptyset$, because if both A_n and B_n occur, then there must be an intersection between an open edge of $\overline{R}_n$ and a closed edge of its dual, which is impossible. Moreover, whenever A_n does not occur then B_n occurs (one can be convinced of this again by looking at Fig. 12.5). It follows that A_n and B_n are disjoint events which partition the sample space, and hence $\mathbb{P}_{p'}(A_n) + \mathbb{P}_{p'}(B_n) = 1$.

We index the nodes at the base of the dual graph by index i, and denote by $i \leftrightarrow \square\overline{R}_n$ the existence of a closed path in the dual from node i to the opposite side $\square\overline{R}_n$ at the top of the dual graph. As $\overline{R}_n$ has width $\kappa \log m$, we have for any $0 \leq q \leq 1$:

$$\mathbb{P}_q(i \leftrightarrow \square\overline{R}_n) \leq \mathbb{P}_q(0 \leftrightarrow \partial S_{\kappa \log m}).$$

Therefore, as edges are closed with probability $1 - p'$,

$$\begin{aligned}
\mathbb{P}_{p'}(B_n) &\leq \sum_{i=1}^{m} \mathbb{P}_{1-p'}(i \leftrightarrow \square\overline{R}_n) \\
&\leq \sum_{i=1}^{m} \mathbb{P}_{1-p'}(0 \leftrightarrow \partial S_{\kappa \log m}) \\
&\leq \frac{4m}{3}(3(1-p'))^{\kappa \log m}
\end{aligned}$$

where the first inequality is a union bound, and the third inequality follows from Lemma 12.1.

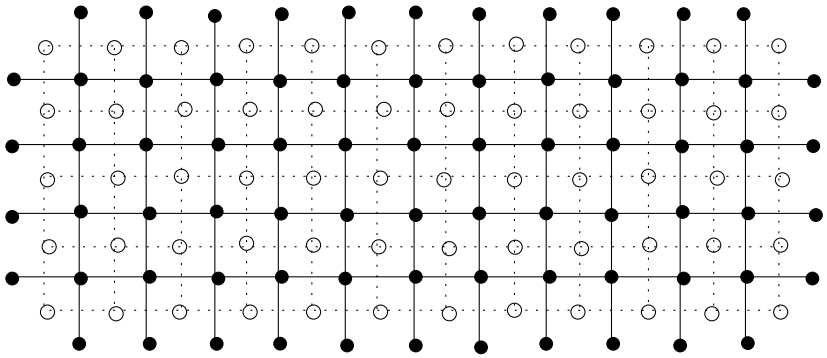

Figure 12.5: A picture of $\overline{R}_n$ (solid line) and its dual graph (dotted line) . Note that if there is no open path traversing the rectangle from left to right, then there is a closed path in the dual graph traversing it from top to bottom.

Now we look at the event that $\beta\kappa\log m$ disjoint paths exist. We apply Lemma 2.1 to rectangle $\overline{R}_n$ with $r = \beta\kappa\log m$, obtaining

$$1 - \mathbb{P}_p(I_{\beta\kappa\log m}(A_n)) \leq \left(\frac{p}{p-p'}\right)^{\beta\kappa\log m} \mathbb{P}_{p'}(B_n), \tag{12.9}$$

for any $p' < p$. Let us choose $p' = 2p - 1$. We have thus

$$\frac{p}{p-p'} = e^{c^2} - 1 < e^{c^2},$$

and

$$1 - p' = 2(1-p) = 2e^{-c^2}.$$

Hence, Equation (12.9) becomes

$$\begin{aligned} 1 - \mathbb{P}_p(I_{\beta\kappa\log m}(A_n)) &\leq \left(e^{c^2}\right)^{\beta\kappa\log m} \mathbb{P}_{p'}(B_n) \\ &\leq m^{\beta\kappa c^2}\frac{4m}{3}(6e^{-c^2})^{\kappa\log m} \\ &= m^{\beta\kappa c^2}\frac{4m}{3}m^{-\kappa c^2+\kappa\log 6} \\ &= \frac{4}{3}m^{(\beta-1)\kappa c^2+\kappa\log 6+1}. \end{aligned}$$

The probability to find at least $\beta\log m$ paths in $\overline{R}_n$ is thus

$$\mathbb{P}_p(I_{\beta\kappa\log m}(A_n)) \geq 1 - \frac{4}{3}m^{(\beta-1)\kappa c^2+\kappa\log 6+1}.$$

As this happens independently in each of the $\frac{m}{\kappa\log m}$ rectangles, the probability of having $\beta\kappa\log m$ disjoint paths in each rectangle is

$$\mathbb{P}_p(I_{\beta\kappa\log m}(A_n))^{\frac{m}{\kappa\log m}} \geq \left(1 - \frac{4}{3}m^{(\beta-1)\kappa c^2+\kappa\log 6+1}\right)^{\frac{m}{\kappa\log m}}.$$

Finally note that if $(\beta-1)\kappa c^2 + \kappa\log 6 + 1 \leq -1$, the above expression tends to 1 when m goes to infinity. Thus, if $c^2 > \log 6 + 2/\kappa$, one can choose

$$\beta(c,\kappa) = 1 - \frac{\kappa\log 6 + 2}{\kappa c^2} > 0$$

such that the above condition is fulfilled.

■

Joining all the rectangles together, we obtain $\beta\sqrt{n}$ paths in the whole network. The same is true of course if we divide the area into vertical rectangles and look for paths crossing the area from bottom to top. Using a simple union bound argument, we conclude that there exist $\beta\sqrt{n}$ horizontal and $\beta\sqrt{n}$ vertical disjoint paths simultaneously with high probability. These paths form a grid, that we call the *highway system*.

Capacity of the percolation cluster

Along the paths of the highway system, we choose one node per edge, that relays the packets. This is possible as the paths are formed by open edges, which are associated to non-empty squares. The paths are thus made of a chain of nodes such that the distance between two consecutive nodes is at most $2\sqrt{2}c$.

To actually transport packets along the paths, we set up a TDMA scheme. When a node transmits, other nodes that are sufficiently far away can simultaneously transmit, without causing excessive interference. Theorem 12.3 makes this precise, ensuring that a constant rate R, independent of n, can be achieved on all the paths simultaneously as $n \to \infty$. Note that this theorem gives a more general result, that will be useful also in Section 12.5.3.

Theorem 12.3. *For any given integer $d > 0$ there exists a TDMA scheduling, such that one node per square can transmit to any destination located within a radius of d squares (in Manhattan distance) with fixed rate $R(d)$ independent of n.*

When d goes to infinity, the asymptotic behavior of the rate is given by

$$R(d) = \Omega\left(d^{-\alpha-2}e^{-\gamma cd}\right).$$

Proof: We take a coordinate system, and label each square with two integer coordinates (in our construction, the axis of the coordinate system are diagonal). Then we take an integer k, and consider the subset of squares whose two coordinates are a multiple of k (see left-hand side of Figure 12.6). By translation, we can construct k^2 disjoint equivalent subsets. This allows us to build the following TDMA scheme: we define k^2 time slots, during which only nodes from a particular subset are allowed to emit. We assume also that at most one node per square emits at the same time, and that they all emit with the same power P.

Let us consider one particular square. We suppose that the emitter in this square emits towards a destination located in a square at distance at most d. We compute the signal-to-interference ratio at the receiver. First, we choose the number of time slots k^2 as follows:

$$k = 2(d+1)$$

To find an upper bound to the interferences, we observe that with this choice, the emitters in the 8 first closest squares are located at a distance at least $d+2$ (in squares) from the receiver (see right-hand side of Figure 12.6). This means that the Euclidean distance between the receiver and the 8 closest interferers is at least $c(d+1)$. The 16 next closest squares are at distance at least $3d+4$ (in squares), and the Euclidean distance between the receiver and the 16 next interferers is therefore at least $c(3d+3)$, and so on. The sum of the interferences $I(d)$ can be bounded as

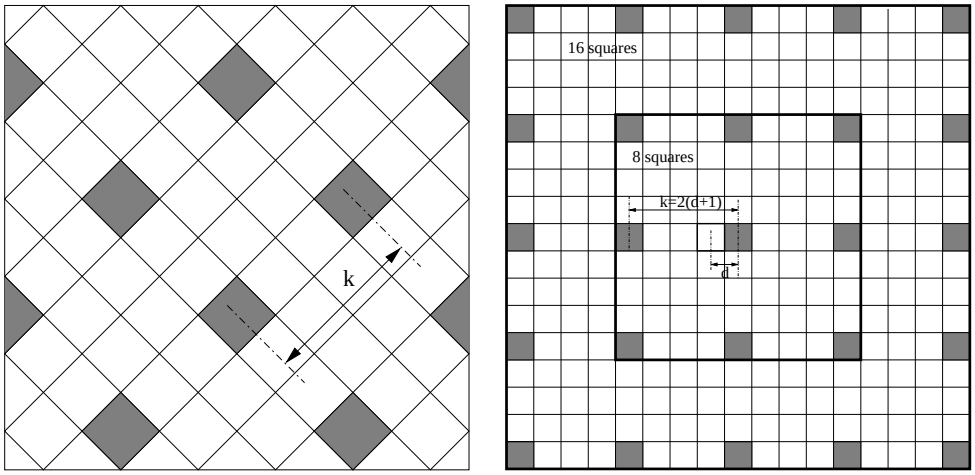

Figure 12.6: Left hand-side: A subset of squares. Right-hand side: Construction of the lower bound on the interference term.

follows:

$$\begin{aligned} I(d) &\leq \sum_{i=1}^{\infty} 8i\, Pl(c(2i-1)(d+1)) \\ &= \sum_{i=1}^{\infty} 8i\, P \min\{1, [c(2i-1)(d+1)]^{-\alpha}\} e^{-\gamma c(2i-1)(d+1)} \\ &\leq \sum_{i=1}^{\infty} 8i\, P[c(2i-1)(d+1)]^{-\alpha} e^{-\gamma c(2i-1)(d+1)} \\ &= P[c(d+1)]^{-\alpha} e^{-\gamma c(d+1)} \sum_{i=1}^{\infty} 8i\,(2i-1)^{-\alpha} e^{-\gamma c(d+1)(2i-2)}. \end{aligned}$$

The sum in $I(d)$ clearly converges if $\alpha > 2$ or $\gamma > 0$.

Now we want to bound from below the signal received from the emitter. We observe first that the distance between the emitter and the receiver is at most

$$\sqrt{(cd)^2 + c^2} \leq c(d+1).$$

The strength $S(d)$ of the signal at the receiver can be thus bounded by

$$\begin{aligned} S(d) &\geq Pl(c(d+1)) \\ &= P \min\{1, [c(d+1)]^{-\alpha}\} e^{-\gamma c(d+1)}. \end{aligned}$$

Finally, we obtain a bound on the signal-to-interference ratio

$$\begin{aligned} SNIR(d) &= \frac{S(d)}{N_0 + I(d)} \\ &\geq \frac{P \min\{1, [c(d+1)]^{-\alpha}\} e^{-\gamma c(d+1)}}{N_0 + P[c(d+1)]^{-\alpha} e^{-\gamma c(d+1)} \sum_{i=1}^{\infty} 8i\,(2i-1)^{-\alpha} e^{-\gamma c(d+1)(2i-2)}} \end{aligned}$$

As the above expression does not depend on n, the first part of the theorem is proven.

We now look at the asymptotic behavior of the SNIR for large d. If $c(d+1) \geq 1$, we can remove the minimum and write

$$\begin{aligned} SNIR(d) &\geq \frac{P[c(d+1)]^{-\alpha} e^{-\gamma c(d+1)}}{N_0 + P[c(d+1)]^{-\alpha} e^{-\gamma c(d+1)} \sum_{i=1}^{\infty} 8i\,(2i-1)^{-\alpha} e^{-\gamma c(d+1)(2i-2)}} \\ &= \frac{1}{N_0 [c(d+1)]^{\alpha} e^{\gamma c(d+1)}/P + \sum_{i=1}^{\infty} 8i\,(2i-1)^{-\alpha} e^{-\gamma c(d+1)(2i-2)}} \end{aligned}$$

The second term in the denominator clearly decreases when d goes to infinity. The first term grows like $d^{\alpha} e^{\gamma c d}$. The whole fraction therefore decreases like $1/d^{\alpha} e^{\gamma c d}$. The throughput on each link is given by $\log(1 + SNIR(d))$, and therefore also decreases like $1/d^{\alpha} e^{\gamma c d}$.

Now we have to divide this throughput by the number of time slots k^2 used in the TDMA scheme. As $k = 2(d+1)$, the number of time slots increases like d^2. So, finally, the actual throughput available in each square decreases like $d^{-\alpha-2} e^{-\gamma c d}$.

■

Corollary 12.1. *For any given integer $d > 0$ there exists a TDMA scheduling, such that one node per square can receive from any emitter located within a radius of d squares with fixed rate $R(d)$ independent of n.*

When d goes to infinity, the asymptotic behavior of the rate is given by

$$R(d) = \Omega\left(d^{-\alpha-2} e^{-\gamma c d}\right).$$

Proof: This result is obtained by switching the role of emitters and receivers in the above proof. Distances remain the same, and all equations still hold.

■

12.5.3 Protocol

In this section, we describe the actual routing protocol and show that it achieves the desired $\Omega(\sqrt{n})$ throughput capacity.

The protocol uses 4 separate time slots: a first one for draining packets to the highway, a second one to transport packets on the "horizontal" highways connecting the left and right edges of the domain, a third one to transport packets on the "vertical" highways connecting the top and bottom edges of the domain, and a fourth one to deliver packets to the destination. The draining and delivery phases use direct transmission, while the highway phases use multiple hops. We show that the throughput bottleneck is in the highway phase that can sustain a rate per node proportional to $1/\sqrt{n}$ bit per second.

We start by proving two simple lemmas that will be useful for the capacity calculation. They are followed by three propositions that prove our main result.

Lemma 12.3. *Divide the network area in $\sqrt{n}/c$ boxes of side length c. The probability that there are less than $c^2 \log n$ nodes in each box tends to one when n goes to infinity.*

Proof: The number of nodes in each square is a Poisson random variable of parameter c^2. Let us denote one of these variables by X. Chernoff's inequality implies that

$$\mathbb{P}(X > c^2 \log n) \leq e^{-sc^2 \log n} \mathbb{E}[e^{sX}],$$

for any $s > 0$. We choose here $s = 2/c^2$ and obtain

$$\mathbb{P}(X \leq c^2 \log n) \geq 1 - n^{-2} e^{c^2(e^{2/c^2}-1)}.$$

As the numbers of nodes in all of the n/c^2 squares are i.i.d, we have

$$\mathbb{P}(X \leq c^2 \log n)^{\frac{n}{c^2}} \geq \left(1 - \frac{e^{c^2(e^{2/c^2}-1)}}{n^2}\right)^{\frac{n}{c^2}}.$$

The latter expression tends to one when n goes to infinity.

■

Lemma 12.4. *Divide the network into horizontal slabs of constant width $\sqrt{2}c/\beta$. The probability that each slab contains less than $2c\sqrt{2n}/\beta$ nodes tends to one when n goes to infinity.*

Proof: The number of nodes in the i-th slab is a Poisson random variable of parameter $c\sqrt{2n}/\beta$, that we denote by N_i here. We apply Chernoff's inequality:

$$\begin{aligned}
\mathbb{P}(N_i > 2\frac{c\sqrt{2n}}{\beta}) &\leq e^{-2\frac{c\sqrt{2n}}{\beta}s}\mathbb{E}[e^{sN_i}] \\
&= e^{-2\frac{c\sqrt{2n}}{\beta}s} e^{c\sqrt{2n}/\beta(e^s-1)} \\
&= e^{\frac{c\sqrt{2n}}{\beta}(e^s-2s-1)}.
\end{aligned}$$

Thus, if we take $s = 1$,

$$\mathbb{P}(N_i > 2\frac{c\sqrt{2n}}{\beta}) \leq e^{\frac{c\sqrt{2n}}{\beta}(e-3)}.$$

The probability that each of the $\beta\sqrt{n}/c\sqrt{2}$ slabs contains less than $2c\sqrt{2n}/\beta$ nodes is thus

$$\mathbb{P}(N_i \leq 2c\sqrt{2n}/\beta,\ \forall i) \geq \left(1 - e^{\frac{c\sqrt{2n}}{\beta}(e-3)}\right)^{\frac{\beta\sqrt{n}}{c\sqrt{2}}}.$$

The latter expression tends to one when n goes to infinity.

■

Proposition 12.1 (Draining Phase). *Each node can transmit packets to the highway system with rate of order*

$$\Omega\left((\log\sqrt{n})^{-\alpha-2}(\log n)^{-1} n^{-c\kappa\gamma/2}\right),$$

for any constant $\kappa > 0$ and if c is so large that (12.8) holds.

Proof: We consider carrying packets from sources to the highways. We start by dividing the network area into $\beta\sqrt{n}/c\sqrt{2}$ horizontal slabs. As there are exactly as many slabs as horizontal edge-disjoint paths, we can impose that nodes from the i-th slab send their packets to the i-th horizontal path. Note that each path may not be contained in its corresponding slab, but it may deviate from it (recall Fig. 12.1). However, Theorem 12.2 bounds the amount of deviation.

More precisely, to each source in the i-th slab, we assign an *entry point* on the i-th horizontal path. The entry point is defined as the node on the horizontal path closest to the vertical line drawn from the source point, see Fig. 12.7. The source then transmits its packet to the entry point in a single hop. Theorem 12.2 ensures that the distance between sources and entry points is never larger than $\kappa \log m = \kappa \log \frac{\sqrt{n}}{\sqrt{2}c}$ squares. This is because each rectangle $\overline{R}_n$ of width $\kappa \log m$ contains $\beta\kappa \log m$ paths, and therefore each source finds its highway within the same rectangle.

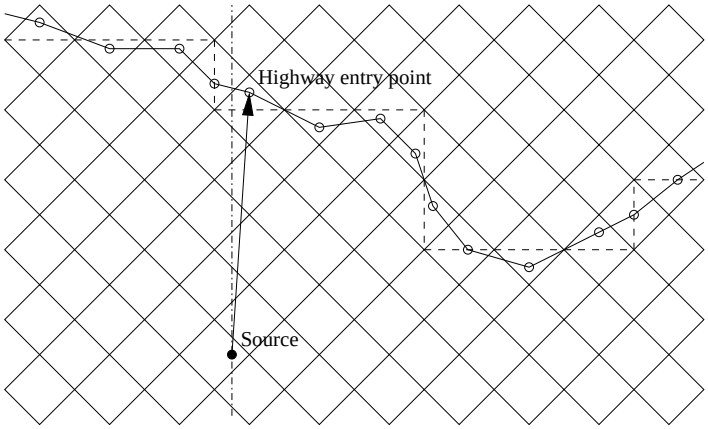

Figure 12.7: Draining of packets to the highways

To compute the rate at which nodes can send their packets to entry points, we let $d = \kappa \log \frac{\sqrt{n}}{\sqrt{2}c}$ in Theorem 12.3. We obtain that one node per square can send packets to its entry point with rate

$$\begin{aligned} R(\kappa \log \frac{\sqrt{n}}{\sqrt{2}c}) &= \Omega\left((\kappa \log \frac{\sqrt{n}}{\sqrt{2}c})^{-\alpha-2} e^{-\gamma c \kappa \log \frac{\sqrt{n}}{\sqrt{2}c}} \right) \\ &= \Omega\left((\log \sqrt{n})^{-\alpha-2} e^{-c\kappa\gamma \log \sqrt{n}} \right) \\ &= \Omega\left((\log \sqrt{n})^{-\alpha-2} n^{-c\kappa\gamma/2} \right). \end{aligned}$$

But as there are possibly many nodes in the squares, they have to share this bandwidth.

Using Lemma 12.3, we can conclude that the transmit rate of each node in the draining phase of our protocol is at least $R(d)/c^2 \log n$, which concludes the proof.

■

Proposition 12.2 (Highway Phase). *Along the highway packets can be relayed at rate at least* $\beta/2k^2\sqrt{n}$ *bits/sec per-node.*

Proof: We now compute the rate that can be sustained during each highway phase of our protocol. Each node generates packets at constant rate W and we must carry these packets towards the point on the vertical highway appropriate for delivery.

We divide horizontal and vertical traffic, adopting the following simple routing policy: packets are carried along horizontal highways until they reach the crossing with their target vertical highway. Then, they are carried along vertical highways until they reach the appropriate point for delivery.

We start considering the horizontal traffic. We consider a node sitting on the i-th horizontal highway, and compute the traffic that goes through it. Actually, a packet will travel through this node if it was generated in the i-th slab, and has a destination on the other side of the node. So, at most, our node will relay all the traffic generated in the i-th slab.

According to Lemma 12.4, a node on a horizontal highway must therefore relay at most $2Wc\sqrt{2n}/\beta$ bits per second. As the maximal distance between hops is constant ($2\sqrt{2}c$), the throughput along highways is independent of n (see Section 12.5.2), one can set the rate per node to $W = \Omega(1/\sqrt{n})$ without overloading links, with high probability.

The problem for vertical traffic is the dual of the previous one. We can use the same arguments,

except that W now describes the *receiving rate* of the nodes. Since each node is the destination of exactly one source, the throughput per node becomes the same as above.

■

Proposition 12.3 (Delivery Phase). *Each destination node can receive packets from the highway at rate*

$$\Omega\left((\log\sqrt{n})^{-\alpha-2}(\log n)^{-1}n^{-c\kappa\gamma/2}\right),$$

for any constant $\kappa > 0$ and c so large that (12.8) holds.

Proof: The delivery phase consists in bringing the packets from the highway system to the actual destination. We proceed exactly in the same way as in Proposition 12.1, but in the other direction (horizontal delivery from the vertical highways).

We divide the network area into $\beta\sqrt{n}/c\sqrt{2}$ vertical slabs, and define a one-to-one mapping between slabs and vertical paths. We assume that packets have been transported by the highway system to their *exit point*, which is defined as the node of the vertical path closest to the horizontal line drawn from the destination. Again, the distance between exit points and destination is at most $\kappa\log\frac{\sqrt{n}}{\sqrt{2}c}$ squares. We can thus let $d = \kappa\log\frac{\sqrt{n}}{\sqrt{2}c}$ in Corollary 12.1, and conclude that each square can be served with rate $R(d) = \Omega\left((\log\sqrt{n})^{-\alpha-2}n^{-c\kappa\gamma/2}\right)$. As there are at most $c^2\log n$ in each square (Lemma 12.3), the throughput per node is at least equal to $R(d)/c^2\log n$.

■

We are now ready to complete the proof of Theorem 12.1.

Proof: *(of Theorem 12.1).* We observe in Propositions 12.1 and 12.3 that if $c\kappa\gamma/2 < 1/2$, the asymptotic throughput per node decreases slower than $1/\sqrt{n}$. In this case, the overall throughput of the protocol is limited by the highway phase only, and the first part of Theorem 12.1 immediately follows from Proposition 12.2. We thus have to make sure that we can choose values of c and κ such that Inequality (12.8) is verified, and such that

$$\frac{c\kappa\gamma}{2} < \frac{1}{2}.$$

A possible choice is $c > 2\gamma + \sqrt{4\gamma^2 + \log 6}$ and $\kappa = \frac{1}{2c\gamma}$. We remark that these choices depend on the physical attenuation factor γ.

■

12.6 Lower bound on transport capacity in dense networks: the percolation approach

For completeness, we consider the model where nodes are distributed according to a Poisson point process of intensity n over a square of unit area. Furthermore, we take an unbounded attenuation function l of the form

$$l(d) = d^{-\alpha}e^{-\gamma d}.$$

In this case, we divide that network into squares of size $c/\sqrt{n}$. We obtain thus the same number of little squares as in the previous model. The average number of nodes in each little square is also the same, namely c^2. Therefore, all the percolation results above still hold for this model, and we can find as many highways as above.

To derive the lower bound on the capacity, we have to compute the throughput along the highways, as well as the rate at which nodes can send data towards the highways. In fact, both of these throughputs were computed using Theorem 12.3, so it is enough here to give an adapted version of such theorem.

Theorem 12.4. *For any given integer $d > 0$ and when n is sufficiently large, there exists a TDMA scheduling, such that one node per square can transmit to any destination located within a radius of d squares (in Manhattan distance) with a fixed rate independent of n.*

When d goes to infinity, the asymptotic behavior of the rate is given by

$$R(d) = \Omega(d^{-2}).$$

Proof: We set up the same TDMA scheme as in Theorem 12.3, with k^2 time slots, where $k = 2(d+1)$. Similarly, the 8 closest interferers are located at least $d+2$ squares away from the receiver, the next 16 interferers at distance $3d+3$, and so on. The difference here is that squares have size $c/\sqrt{n}$ here. The sum of the interferences at the receiver can be bounded as follows:

$$\begin{aligned}
I(d,n) &\leq \sum_{i=1}^{\infty} 8i\, Pl\left(\frac{c[ki-(d+1)]}{\sqrt{n}}\right) \\
&= \sum_{i=1}^{\infty} 8i\, Pl\left(\frac{c(2i-1)(d+1)}{\sqrt{n}}\right) \\
&= \sum_{i=1}^{\infty} 8i\, P\left(\frac{c(2i-1)(d+1)}{\sqrt{n}}\right)^{-\alpha} e^{-\gamma c(2i-1)(d+1)/\sqrt{n}} \\
&\leq P\left(\frac{c(d+1)}{\sqrt{n}}\right)^{-\alpha} e^{-\gamma c(d+1)/\sqrt{n}} K'(d),
\end{aligned}$$

where

$$K'(d) = \sum_{i=1}^{\infty} 8i\,(2i-1)^{-\alpha} e^{-\gamma c(2i-2)(d+1)/\sqrt{n}}.$$

The latter sum clearly converges to a constant $\hat{K}$ if $\alpha > 2$ or $\gamma > 0$.

As the receiver is at most d squares away from the emitter, the Euclidean distance between them is less than $c(d+1)/\sqrt{n}$. The strength $S(d,n)$ of the signal at the receiver reads thus

$$S(d,n) = P\left(\frac{c(d+1)}{\sqrt{n}}\right)^{-\alpha} e^{-\gamma c(d+1)/\sqrt{n}}.$$

The SINR is thus

$$\begin{aligned}
SINR(d,n) &= \frac{S(d)}{N_0 + I(d)} \\
&\geq \frac{P\left(\frac{c(d+1)}{\sqrt{n}}\right)^{-\alpha} e^{-\gamma c(d+1)/\sqrt{n}}}{N_0 + P\left(\frac{c(d+1)}{\sqrt{n}}\right)^{-\alpha} e^{-\gamma c(d+1)/\sqrt{n}} K'(d)} \\
&= \frac{1}{N_0\left(\frac{c(d+1)}{\sqrt{n}}\right)^{\alpha} e^{\gamma c(d+1)/\sqrt{n}}/P + K'(d)}
\end{aligned}$$

When n goes to infinity, we observe that the first term of the denominator tends to zero, whereas the second term tends to $\hat{K}$. The whole fraction tends thus to $1/\hat{K}$, and thus

$$\lim_{n\to\infty} SINR(d,n) \geq \frac{1}{\hat{K}}.$$

It means that for any $\varepsilon > 0$, when n is large enough, the rate

$$R' = \frac{1}{2}\log(1 + \frac{1}{\hat{K}} - \varepsilon)$$

is achievable by this scheme during each time slot. This proves the first part of the theorem.

Now if d increases with n, we notice that the above limit still holds whenever $(d+1)/\sqrt{n}$ tends to zero. Therefore, if $d = O(\sqrt{n})$, the rate R' is achievable for each active transmission when n is large enough. However, as there are $k^2 = 4(d+1)^2$ time slots in our TDMA scheme, the actual throughput available for each square is $R(d) = R'/k^2$, and thus

$$R(d) = \Omega(d^{-2}).$$

■

13

Scale-Free Graphs

In Chapter 10, we have studied the small-world property, which is observed in many real networks in a broad array of contexts. In this chapter, we will study another such robust empirical property, that of degree distributions that follow a power law. We will also sketch some probabilistic models that have been proposed to explain this effect.

Networks with a power law degree distribution are referred to as *scale-free* because a power law does not possess a characteristic scale. This is in contrast to, for example, the exponential distribution, which drops off very quickly beyond the mean. A power law has a heavy tail, which makes values far beyond the mean much more likely than for light-tailed distributions.

It is important to note that power laws have a very long and controversial history. Power laws have been observed in many contexts, such as income distributions, file sizes on computers, object sizes on the web, lengths of phone calls, the sizes of cities, word frequencies in prose, etc. Whenever a phenomenon occurs in such a wide variety of seemingly unrelated scenarios, it is tempting to suspect that some very simple underlying mechanism can be discovered that explains this phenomenon. As an example, another very prominent probability distribution, the normal law, occurs very naturally whenever independent random variables are added, which is captured in the central limit theorem (CLT).

For power laws, the quest for such a hidden mechanism is not quite as settled as for the normal law, but we have two interesting leads. The first class of mechanisms could be termed "the rich get richer"; its study has been pioneered by Yule in the 1920s, and by the economist Herbert Simon in the 1950s. In the context of network degree distributions, this mechanism has been rediscovered under the name of *preferential attachment* by Albert and Barabasi.

The second class of mechanisms, which we will only touch upon briefly, postulates that power laws occur as the outcomes of optimization problems. Again, such models go back to at least the 1950s, and remain an active area of inquiry as well.

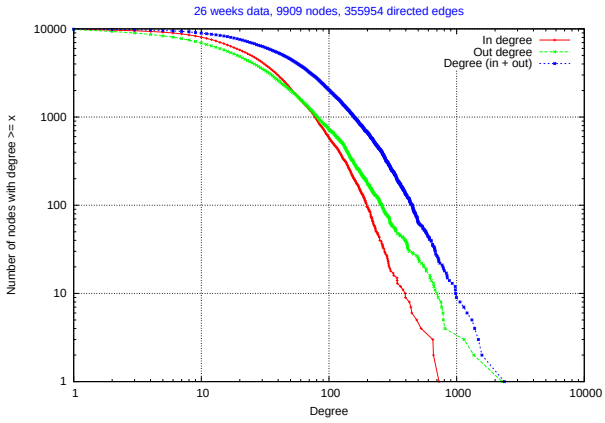

Figure 13.1: The ccdf of the degree distribution of the EPFL email graph.

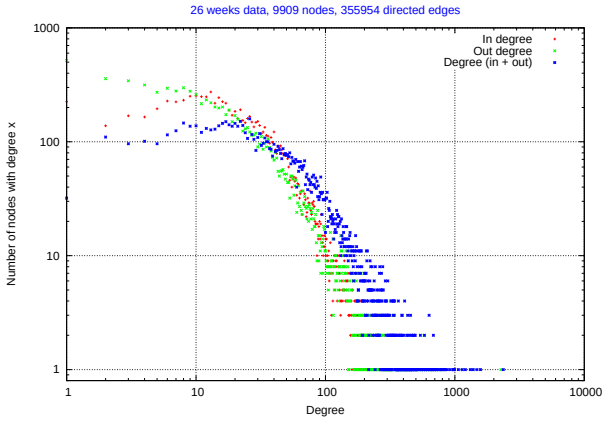

Figure 13.2: The degree density of the EPFL email graph.

13.1 Introduction

The most prominent and simple example of a power law distribution is the Pareto law. The Pareto distribution has the following ccdf.

$$F(x) = \left(\frac{x}{\beta}\right)^{-\gamma}, \tag{13.1}$$

which requires $X \geq \beta$. The exponent γ is sometimes called the Pareto index.

The moments of the Pareto distribution only exist up to $k < \gamma$; higher moments are infinite.

$$\mathbb{E}\left[X^k\right] = \left(\frac{\beta^k \gamma}{\gamma - k}\right), k < \gamma, \tag{13.2}$$

and ∞ otherwise. So if $\gamma \leq 1$, X has infinite mean, and for $1 < \gamma \leq 2$, then X has infinite variance.

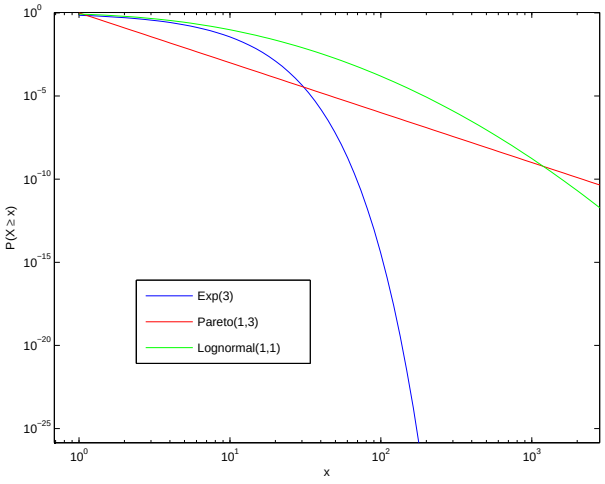

Figure 13.3: The ccdf of the exponential, Pareto, and lognormal distributions.

13.2 The Simon model: the rich get richer

Herbert Simon, a famous economist, introduced a model in a 1955 paper to explain the heavy-tailed distribution in incomes first observed by Vilfredo Pareto[1]. His model could be termed "the rich get richer"; the intuition is that someone who is rich has more opportunities to develop additional wealth, while a poor person has trouble getting out of poverty.

He explains his model in terms of word frequencies in prose; these frequencies have been shown to satisfy power laws. More specifically, assume that $X_i(t)$ is the number of different words that have occurred i times among the first t words in the text. Then, he assumes that with some probability α, the next word typed is a new word, i.e., a word that has not occurred among the first t words; and with probability $1-\alpha$, the next word is identical to some previously observed word, and the probability that that word has previously occurred i times is proportional to $iX_i(t)$. In other words, there is a bias for the next word to be chosen among those that have already occurred quite often.

Another intuitive example arises when we consider the sizes of cities. Suppose you move to a new country and try to decide what city or town to settle in. It is unlikely that you will sample from the set of cities and towns uniformly to decide where to move; rather, your behavior is probably representative of the behavior of people before you, i.e., you are much more likely to settle in a large city than in a small town. In other words, you sample from the *population* rather from the set of cities and towns. This may go same way towards explaining the power law in city sizes. What is remarkable that even if the sizes of cities are initially equal, the cumulative effect of random decisions by new arrivals would give rise to the power law over time.

13.3 Preferential attachment: the BA model

The Simon model is an example of a growth model, where objects arrive sequentially in a system, i.e., the system is never in equilibrium. The Barabási-Albert (BA) model [2] proceeds along exactly the same lines: at each time step, a node node arrives to the system, and establishes one or several edges to existing nodes. We now reproduce an argument that shows why this gives to power laws for the

[1] Who taught at the University of Lausanne.

degree distribution, based on a slight departure from the BA model for ease of exposition.

Every new node connects to one of the existing nodes uniformly at random with probability α. With probability $1-\alpha$, it connects to each node v with a probability proportional to $d_{in}(v)$.

Let $X_j(t)$ denote the number of nodes v with $d_{in}(v) = j$ at time t (to simplify notation, we write X_j for $X_j(t)$). Therefore, the total number of nodes and the total number of edges at time t is always t.

For $j \geq 1$,

$$\mathbb{P}\{X_j(t+1) = X_j(t) + 1\} = \alpha\frac{X_{j-1}}{t} + (1-\alpha)(j-1)\frac{X_{j-1}}{t}. \tag{13.3}$$

$$\mathbb{P}\{X_j(t+1) = X_j(t) - 1\} = \alpha\frac{X_j}{t} + (1-\alpha)j\frac{X_j}{t}. \tag{13.4}$$

$$\frac{dX_j}{dt} = \frac{\alpha(X_{j-1} - X_j) + (1-\alpha)[(j-1)X_{j-1} - jX_j]}{t} \tag{13.5}$$

For $j = 0$, we get

$$\frac{dX_0}{dt} = 1 - \frac{\alpha X_0}{t} \tag{13.6}$$

Let us assume that the fraction of nodes of degree j converges, i.e., $X_j/t \to c_j$. We then solve for c_j:

$$c_0 = \frac{1}{1+\alpha} \tag{13.7}$$

$$\frac{c_j}{c_{j-1}} = \frac{\alpha + (j-1)(1-\alpha)}{1+\alpha+j(1-\alpha)} = 1 - \frac{2-\alpha}{1+\alpha+j(1-\alpha)} \sim 1 - \frac{2-\alpha}{1-\alpha}j^{-1}. \tag{13.8}$$

Note that

$$\left(\frac{j}{j-1}\right)^{-\beta} = (1-\frac{1}{j})^{\beta} \sim 1 - \frac{\beta}{j}. \tag{13.9}$$

Therefore, the tail behavior of c_j is

$$c_j \sim c \cdot j^{-\frac{2-\alpha}{1-\alpha}} \tag{13.10}$$

Therefore, the exponent of the ccdf is $\gamma = 1/(1-\alpha)$.

Obviously, the above argument is not rigorous; in (13.5), we have replaced a discrete random process with a continuous deterministic function. Also, we have assumed that X_j/t converges, which requires a proof, of course. This type of approximation is often referred to as "continuum theory" [2][2]. A mathematically rigorous argument proceeds by first showing that $X_j(t)$, appropriately rescaled, is a martingale, and then applying a concentration result known as Azuma's inequality (which can be roughly viewed as equivalent to the CLT for martingales). This allows us to show that the process $X_j(t)$ is unlikely to depart too much from its expectation. An example of a rigorous treatment of the subject is given in [27].

We now turn to the BA model. Start with m_0 nodes. At each time step t, add $m \leq m_0$ edges from the new node to the set of existing nodes, such that

$$\mathbb{P}\{\text{connect to } v\} = \frac{d_v}{\sum_{u=1}^{t+m_0} d_u} \tag{13.11}$$

[2]although "continuum approximation" would be a more appropriate label.

After t steps, there are $m_0 + t$ nodes and $mt + e_0$ edges in the graph, where e_0 is the number of initial edges.

The analysis of the BA model proceeds in analogy to the one for the version with directed edges above, and with $\alpha = 0$. However, in their model, $\gamma = 3$.

13.4 Optimization: the FKP model

We briefly discuss the FKP model [13] to give a flavor of a growth model with optimization to explain power laws.

Start with a unit square. Nodes arrive sequentially, and are placed i.i.d. uniformly in the unit square. We generate a sequence G_i of growing trees as follows.

The ith node attaches itself to one node j that minimizes

$$\alpha d_{ij} + h_j. \tag{13.12}$$

where d_{ij} is the Euclidean distance between i and j, and h_j is a measure of "centrality" of node j. For example, h_j can be taken to be the average or maximum number of hops in the graph G_i from j to all other nodes, or the number of hops from a fixed center of the tree.

If $\alpha < 2^{-1/2}$, then G_i is a star centered at node 0. For $\alpha = \Omega(\sqrt{n})$, the degree distribution has exponential tail. For $\alpha \geq 4$, but $\alpha = o(\sqrt{n})$, the degree distribution follows a power law.

13.5 Power law vs. lognormal law

There is another probability law called the *lognormal* law, which is easily confused with a power law. A random variable Y has lognormal distribution if it can be written as $X = \log Y$ and X has normal distribution.

The ccdf of a lognormal random variable with parameters μ and σ^2 is given by

$$F(x) = \int_{z=x}^{\infty} \frac{dz}{\sqrt{2\pi}\sigma z} e^{-(\log z - \mu)^2/2\sigma^2}. \tag{13.13}$$

To see why the lognormal distribution is hard to distinguish from a power law, write its pdf as

$$\begin{aligned} \log p(x) &= -\log x - \frac{1}{2}\log(2\pi\sigma^2) - \frac{(\log x - \mu)^2}{2\sigma^2} \\ &= -\frac{(\log x)^2}{2\sigma^2} + \left(\mu/\sigma^2 - 1\right)\log x - \left(\frac{1}{2}\log(2\pi\sigma^2) + \frac{\mu^2}{2\sigma^2}\right) \end{aligned} \tag{13.14}$$

Note however that the tail of the lognormal distribution is light enough for all moments to exist. This is because

$$\exp(-(\log x)^2) = \frac{1}{x^{\log x}}, \tag{13.15}$$

which decreases faster than any polynomial.

14

Navigability

We had seen in Chapter 10 that a surprising feature of many real networks is the existence of short paths betweeen pairs of nodes. Milgram's experiments had established the existence of such paths in social networks, by forwarding letters along social links between a sender and a target who did not know each other personally.

While the lengths of the paths in Milgram's experiment is certainly surprising, we have been at least partially able to convince ourselves that the randomness in social ties could explain this feature. However, a major puzzle remains, which in fact may be much more challenging to explain convincingly: *people were actually able to find these paths.* This fact is remarkable because every person typically has only full knowledge of their own social ties. While we sometimes know a few of our friends' friends, it is unlikely that we have complete knowledge of our two-hop neighborhood - and certainly our knowledge decreases further the more removed our indirect acquaintances.

In this chapter, we discuss a class of models that has been advanced to explain why some networks - and in particular, social networks - appear to be navigable, i.e., that it is possible to compute efficient route from sources to destinations where each intermediate node only has information about its neighbors.

The first two models we discuss use an underlying geometry to provide a network-wide reference system. This reference system provides a geometric addressing scheme for the destination; also, a node knows its own location in the geometry and the locations of its neighbors, and it uses this information to forward a message towards a destination. The second models we discuss use a hierarchy, or multiple hierarchies, to organize the nodes, rather than a geometry.

These models bring out an interesting puzzle: navigability is not a robust property; rather, it requires a particular way of forming shortcuts between nodes. Why actual social networks exhibit navigability remains an open question, and suggests that we have only started to uncover the principles underlying the structure of such networks, and the mechanisms that give rise to this structure.

14.1 Lattice-induced networks

Consider the following model [24]. We start with the lattice $[n]^2$ with n^2 vertices, and define distance $d(u, v)$ between two vertices u and v as Manhattan distance. The graph topology then has two components, local edges and remote edges or *shortcuts*. Every vertex v has local edges to its neighbors. Every vertex v also has $r \geq 0$ remote edges, whose endpoints are selected randomly i.i.d. over the set of other vertices, with

$$\mathbb{P}\{(u, v) \text{ is a remote edge}\} \propto (d(u, v))^{-\gamma}, \tag{14.1}$$

with $\gamma \geq 0$ a constant.

Note that $\gamma = 0$ corresponds roughly to the Watts-Strogatz model, in that remote edges are uniform over the set of vertices. Larger γ concentrates the remote edges closer to v. Therefore, γ controls how far-reaching the shortcuts through remote edges are.

A *decentralized routing algorithm* is allowed to make local forwarding decisions for a message based only on knowledge of the lattice, as well as knowledge of the remote edges of the current vertex making the decision, as well as all of the previous vertices encountered by the message.

We are interested in the expected number of iterations for any decentralized routing algorithm, for the source s and destination t chosen uniformly at random over the lattice.

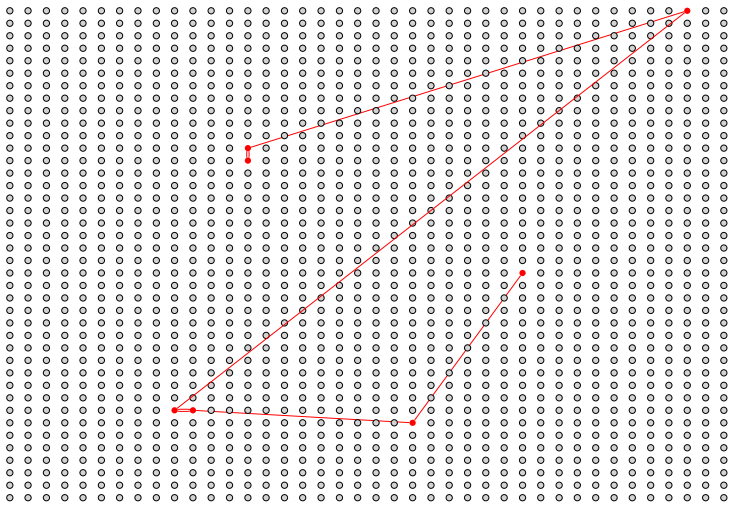

Figure 14.1: A sample shortest path in a lattice-induced topology with $\gamma = 1$. The path is short, but does not follow the geometry of the underlying lattice.

Theorem 14.1 (Efficient for $\gamma = 2$). *There exists a decentralized algorithm whose expected delivery time is $O((\log n)^2)$ for n large enough.*

Proof: We show the result for $r = 1$ only; it is obvious that larger r can only help.

The proof is constructive: we define an algorithm to forward a message from the source s to the target t, and show that this algorithm has expected cost of $O((\log n)^2)$.

In this algorithm, the current message holder selects the neighbor (on the lattice or through a shortcut) that is closest to the target. Note that it is always possible to make progress towards the target through the lattice; therefore, the algorithm always terminates correctly.

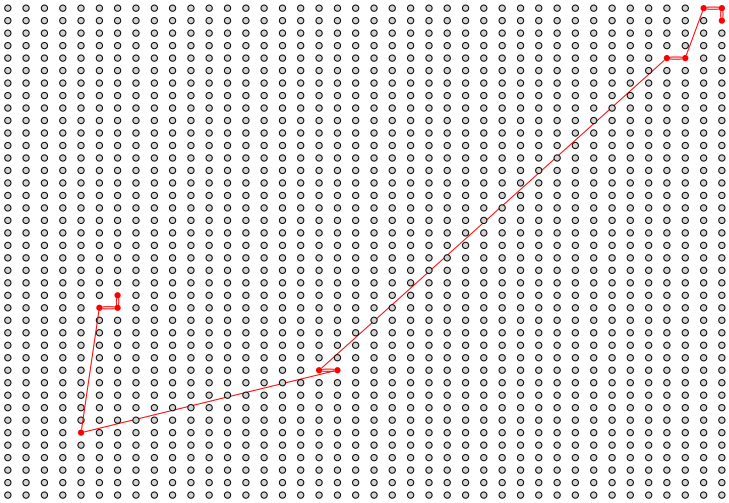

Figure 14.2: A sample shortest path in a lattice-induced topology with $\gamma = 2$, the critical exponent. The path is short, but most hops are still in the direction of the destination.

The algorithm is loop-free because by definition the message always makes progress. Therefore, we can assume that the shortcuts are generated at a node u when that node receives the message.

Define the annuli U_j as the set of points at lattice distance in $[2^j + 1, 2^{j+1}]$. Also, the box B_j is the set of points at lattice distance at most 2^j. The algorithm is in phase j while the message is in U_j. Thus, the initial value of j is at most $\log_2 n$.

Consider phase j, $\log\log n \leq j < \log n$. We bound the probability that the algorithm finishes phase j in the next step, and moves into phase $j + 1$. This requires that the message is passed to a node in B_j in the next step. The size of the box B_j is at least

$$|B_j| \geq 2^{2j-1}, \tag{14.2}$$

and the maximum lattice distance between u and a node in B_j is $2^{j+1} + 2^j < 2^{j+2}$. The probability of hitting a node in B_j is at least

$$q_j \geq \frac{|B_j|}{(2^{j+2})^\gamma 4\ln(6n)}. \tag{14.3}$$

Therefore, the message enters B_j with probability at least

$$q_j \geq \frac{2^{2j-1}}{4\ln(6n)2^{2j+4}} \geq \frac{1}{128\ln(6n)} \tag{14.4}$$

We next bound the expected number of steps in phase j. Let X_j denote this number. Note that X_j is stochastically upper-bounded by a geometric random variable with mean $1/q_j$, and therefore

$$\mathbb{E}[X_j] \leq 128\ln(6n) \tag{14.5}$$

Now the total cost of the algorithm is at most $\sum_{j=0}^{\log n} X_j = O(\log^2 n)$.

■

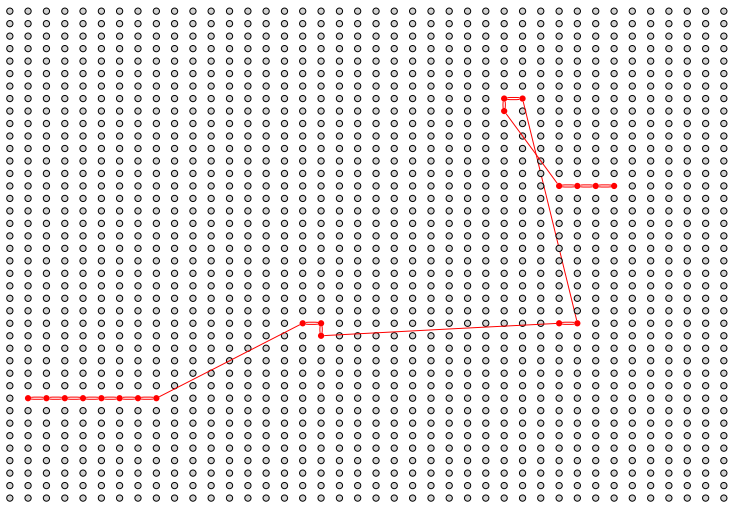

Figure 14.3: A sample shortest path in a lattice-induced topology with $\gamma = 3$. Note that the shortest path is quite long, but that most hops on the shortest path are in the direction of the destination.

Theorem 14.2 (Inefficient for $0 \leq \gamma < 2$: no geometric clues). *The expected delivery time of any decentralized algorithm for $0 \leq \gamma < 2$ is $\Omega(n^\delta)$, with $\delta = (2-\gamma)/3 < 1$.*

Proof: For the converse, we need to assume that every node has r shortcuts (as opposed the the achievability proof, where we could constrain to $r = 1$).

The probability that the current message holder u chooses node v as one of its r shortcuts is $d(u,v)^{-\gamma} / \sum_{v \neq u} d(u,v)^{-\gamma}$. We bound the normalization constant as

$$\begin{aligned}
\sum_{v \neq u} d(u,v)^{-\gamma} &\geq \sum_{j=1}^{n/2} j \cdot j^{-\gamma} \\
&= \sum_{j=1}^{n/2} j^{1-\gamma} \\
&\geq \int_1^{n/2} x^{1-\gamma} dx \\
&\geq \frac{(n/2)^{2-\gamma} - 1}{2-\gamma} \\
&\geq \frac{n^{2-\gamma}}{(2-\gamma)2^{3-\gamma}} \qquad \text{(assuming } n \text{ large enough s.t. } n^{2-\gamma} \geq 2^{3-\gamma}\text{).} \qquad (14.6)
\end{aligned}$$

Let the box B denote the set of nodes within lattice distance n^δ of the target t. Note that there are at most $1 + \sum_{j=1}^{n^\delta} 4j \leq 4n^{2\delta}$ nodes in B. Let $\lambda = (2^{8-\gamma} r^2)^{-1}$.

B is difficult to hit. We now bound the probability that a set of nodes of a given size has at least one shortcut into B. Consider an arbitrary node u. We make no assumptions about where u is located w.r.t to t (the preceding argument even allows it to be within B). Let E'_i be the event that at the ith step, the message holder has a shortcut into B. Let E' be the event that any message holder in the first λn^δ steps has a shortcut into B,

$$E' = \cup_{i \leq \lambda n^\delta} E'_i, \tag{14.7}$$

where $\lambda = 2^{\gamma-8}/r$ is a constant.

$$\begin{aligned} \mathbb{P}\{E'_i\} &\leq r|B|\frac{(2-\gamma)2^{3-\gamma}}{n^{2-\gamma}} \quad \text{(because } d(u,v)^{-\gamma} \leq 1\text{, and using (14.6))} \\ &\leq \frac{(2-\gamma)2^{3-\gamma}r4n^{2\delta}}{n^{2-\gamma}} \\ &= \frac{(2-\gamma)2^{5-\gamma}rn^{2\delta}}{n^{2-\gamma}}, \end{aligned} \tag{14.8}$$

and by the union bound

$$\begin{aligned} \mathbb{P}\{E'\} &\leq \sum_{i \leq \lambda n^\delta} \mathbb{P}\{E'_i\} \\ &\leq \frac{(2-\gamma)2^{3-\gamma}r4n^{2\delta}}{n^{2-\gamma}} \\ &= (2-\gamma)2^{5-\gamma}\lambda r \leq 1/4. \end{aligned} \tag{14.9}$$

Cost is high if no shortcut into B. If s and t are chosen uniformly at random, then $d(s,t) > \lambda n^\delta$ *a.a.s.*. Also, from (14.9), with non-vanishing probability, there is no shortcut into B in the first λn^δ steps.

■

Essentially, what the previous theorem says is that shortcuts are too random. This makes it difficult to find a shortcut that takes the message close to the target. Although a short path exists, the step-by-step discovery of the decentralized algorithm is unable to find this path, given that it cannot zero in to the target in an exponentially decreasing sequence of boxes around t, as had been the case for $\gamma = 2$.

In the next theorem, we consider the opposite case, where shortcuts are very local. Here, the problem is not that geometry is not a useful guide towards the target, but that shortcuts are simply too short to get to the target efficiently.

Theorem 14.3 (Inefficient for $\gamma > 2$: shortcuts are too local.). *The expected delivery time of any decentralized algorithm is $\Omega(n^\beta)$, where $\beta = (\gamma-2)/(\gamma-1)$.*

Proof: Write $\epsilon = \gamma - 2$, and $\lambda = \min(\epsilon, 1)/8r$.

$$\begin{aligned} \mathbb{P}\{d(u,v) > d\} &\leq \sum_{j=d+1}^{2n-2} 4j \cdot j^{-\gamma} \\ &= 4\sum_{j=d+1}^{2n-2} j^{1-\gamma} \\ &\leq \int_d^\infty x^{1-\gamma} dx \\ &\leq \frac{d^{2-\gamma}}{\gamma-2} = d^{-\epsilon}/\epsilon. \end{aligned} \tag{14.10}$$

Let E_i' denote the event that the message reaches a node u that has a shortcut of length at least $n^{1-\beta}$. Let $E' = \cup_{i \leq \lambda n^\beta} E_i'$ be the event that any of the first λn^β nodes has such a shortcut.

By the union bound,

$$\begin{aligned} \mathbb{P}\{E'\} &\leq \sum_{i \leq \lambda n^\beta} \mathbb{P}\{E_i\} \\ &\leq \lambda n^\beta n^{-\epsilon(1-\beta)}/\epsilon \quad \text{(combining (14.10) and def. of } E_i') \\ &= \lambda/\epsilon \leq 1/4. \quad (\epsilon(1-\beta) = \beta) \end{aligned} \tag{14.11}$$

As in the previous proof, we choose s and t uniformly at random, so that they are at typical distance $\Theta(n)$. We want to show that the algorithm requires more than $O(n^\beta)$ steps to complete on average. It follows from (14.11) that with non-vanishing probability, the message can only find shortcuts of distance less than $n^{1-\beta}$ in the first λn^β steps, for a total progress of order λn. In that case, the message cannot reach the target within the required number of steps.

■

Kleinberg studies several generalizations. For example, he allows for more than 4 local connections, which, unsurprisingly, does not affect the scaling results. The model can also be generalized to a d-dimensional lattice. The critical exponent is then $\gamma = d$.

14.2 Tree-induced networks

The network model underlying the results on decentralized routing in the previous section is arguably quite peculiar. In particular, there is a priori no reason why a network should "live" in a Euclidean space, which is crucial to allow decentralized routing at $\gamma = 2$. It would therefore be desirable to study a more general class of networks to shed insight In a followup paper [25], Kleinberg introduces two additional models. The first model uses a tree instead of a lattice as the underlying generator. The second model does not explicitly derive from an underlying generator. Rather, he postulates a set of conditions for a general underlying group structure, and again generates random networks based on this group structure.

Theorem 14.4 (Efficient for $\gamma = 1$). *There exists a tree-induced model with exponent $\gamma = 1$ and outdegree of $\Theta(\log n)$ such that a decentralized routing algorithm can achieve search time $O(\log n)$.*

Theorem 14.5 (Inefficient for $0 \leq \gamma < 1$: no hierarchical clues). *There does not exist a tree-induced model with exponent $\gamma < 1$ and outdegree $O(\log n)$ such that a decentralized routing algorithm can achieve polylogarithmic search time.*

Theorem 14.6 (Inefficient for $\gamma > 1$: shortcuts too local). *There does not exist a tree-induced model with exponent $\gamma > 1$ and outdegree $O(\log n)$ such that a decentralized routing algorithm can achieve polylogarithmic search time.*

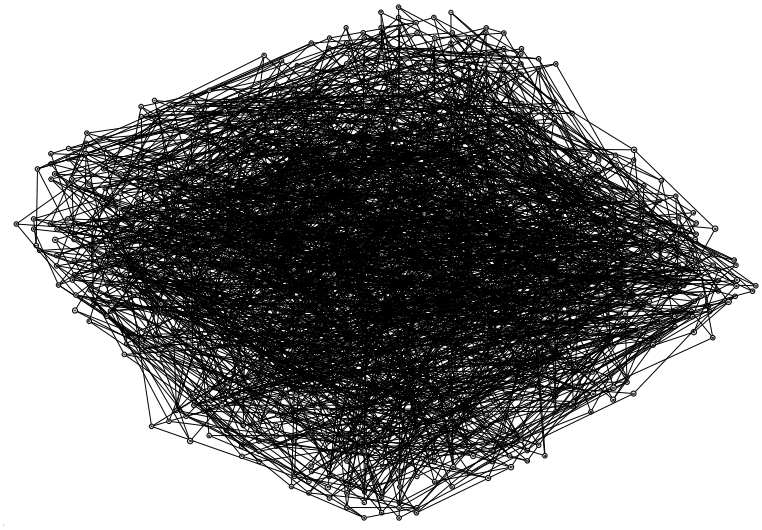

Figure 14.4: A tree-induced network with $n = 1024$, $b = 2$, and $\alpha = 0.5$. This network has a lot of long-range links.

Figure 14.5: A tree-induced network with $n = 1024$, $b = 2$, and $\alpha = 2$. This network has mostly local links.

Bibliography

[1] F. Baccelli and B. Blaszczyszyn. On a coverage process ranging from the boolean model to the poisson voronoi tessellation, with applications to wireless communications. *Adv. Appl. Prob.*, 33(2):293–323, 2001.

[2] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, vol. 286:pp. 509–512, 1999.

[3] A. D. Barbour, L. Holst, and S. Janson. *Poisson Approximation*. Oxford Science Publications, 1992.

[4] B. Bollobas and F. R. K. Chung. The Diameter of a Cycle plus a Random Matching. *SIAM J. Disc. Math.*, 1(3), August 1988.

[5] B. Bollobas and F. de la Vega. The Diameter of Random Regular Graphs. *Combinatorica*, 2:125–134, 1982.

[6] S. R. Broadbent and J. M. Hammersley. Percolation processes i. crystals and mazes. In *Proceedings of the Cambridge Philosophical Society*, pages 629–641, 1957.

[7] F. Chung and L. Lu. The Diameter of Sparse Random Graphs. *Advances in Applied Mathematics*, 26:257–279, 2001.

[8] D. J. Daley. The definition of a multi-dimensional generalization of shot noise. *J. Appl. Prob.*, 8:128–135, March 1971.

[9] O. Dousse, F. Baccelli, and P. Thiran. Impact of interferences on connectivity of ad hoc networks. *ACM/IEEE Transactions on Networking*, 13(2):425–436, April 2005.

[10] O. Dousse, M. Franceschetti, N. Macris, R. Meester, and P. Thiran. Percolation in the signal to interference ratio graph. *Journal of Applied Probability*, 43(2), June 2006.

[11] O. Dousse, M. Franceschetti, and P. Thiran. The costly path from percolation to full connectivity. In *Proc. 42nd Annual Allerton Conference on Communication, Control and Computing*, Monticello, September 2004.

[12] O. Dousse and P. Thiran. Connectivity vs capacity in ad hoc networks. In *Proc. IEEE Infocom*, Hong Kong, March 2004.

[13] A. Fabrikant, E. Koutsoupias, and C. H. Papadimitriou. Heuristically Optmized Tradeoffs: A new paradigm for power laws in the Internet. 2002.

[14] M. Franceschetti, J. Bruck, and L. Schulman. A random walk model of wave propagation. *IEEE Transactions on Antennas and Propagation*, May 2004. to appear.

[15] M. Franceschetti, O. Dousse, D. Tse, and P. Thiran. Closing the gap in the capacity of random wireless networks. In *Proc. of Information Theory Symposium (ISIT)*, Chicago, Illinois, July 2004.

[16] M. Franceschetti and R. Meester. Critical node lifetimes in random networks via the chen-stein method. In *Proc. Spaswin'05*, Riva del Garda, Italy, April 2005.

[17] A. El Gamal, J. Mammen, B. Prabhakar, and D. Shah. Throughput delay trade-off in wireless networks. In *Proc. IEEE Infocom*, Hong Kong, March 2004. to appear.

[18] G. R. Grimmett. *Percolation*. Springer, 1999.

[19] G. R. Grimmett and D. R. Stirzaker. *Probability and Random Processes*. Oxford Science Publications, 1992.

[20] M. Grossglauser and D. Tse. Mobility increases the capacity of ad-hoc wireless networks. In *Proc. IEEE Infocom*, Anchorage, Alaska, April 2001.

[21] P. Gupta and P. R. Kumar. Critical power for asymptotic connectivity in wireless networks. *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming*, 1998. Edited by W.M. McEneany, G. Yin, and Q. Zhang, (Eds.) Birkhuser.

[22] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, March 2000.

[23] J. F. C. Kingman. *Poisson Processes*. Clarendon Press, Oxford, 1993.

[24] J. Kleinberg. The small-world phenomenon: An algorithmic perspective. 2000.

[25] J. Kleinberg. Small-World Phenomena and the Dynamics of Information. *Advances in Neural Information Processing Systems (NIPS)*, 14, 2001.

[26] S. R. Kulkarni and P. Viswanath. A deterministic approach to throughput scaling in wireless networks. In *Proc. IEEE International Symposium on Information Theory*, November 2002.

[27] R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, and E. Upfal. Stochastic models for the Web graph. pages 57–65, 2000.

[28] O. Lvque, E. Telatar, and D. Tse. Upper bounds on the capacity of ad-hoc wireless networks. In *Proc. Winter School on Coding and Information Theory*, Monte Verit, Switzerland, February 2003.

[29] R. Meester and R. Roy. *Continuum percolation*. Cambridge University Press, 1996.

[30] M. Molloy and B. Reed. A Critical Point for Random Graphs with a Given Degree Sequence. *Random Structures and Algorithms*, 6(2-3):161–179, 1995.

[31] P. Erdös and A. Rényi. On random graphs I. *Publ. Math. Debrecen*, 6:290–297, 1959.

[32] M. Penrose. The longest edge of the random minimal spanning tree. *Ann. Appl. Probability*, 7:340–361, 1997.

[33] M. Penrose and A. Pisztora. Large deviations for discrete and continuous percolation. *Adv. Appl. Prob.*, 28:29–52, 1996.

[34] V. Ramasubramanian, R. Chandra, and D. Mosse. Providing a bidirectional abstraction for unidirectional adhoc networks. In *Proc. IEEE Infocom*, Ney-York, June 2002.

[35] S. M. Ross. *Stochastic Processes*. Wiley, 1996.

[36] S. Toumpis and A.J. Goldsmith. Large wireless networks under fading, mobility, and delay constraints. In *Proc. IEEE Infocom*, Hong Kong, 2004. to appear.

[37] V. Veeravalli and A. Sendonaris. The coverage-capacity tradeoff in cellular cdma systems. *IEEE Transactions on Vehicular Technology*, 48:1443–1451, 1999.

[38] Liang-Liang Xie and P. R. Kumar. A network information theory for wireless communication: Scaling laws and optimal operation. *IEEE Transactions on Information Theory*, 50(5):748–767, 2004.