1

Introduction

Random networks arise when nodes are randomly deployed on the plane and randomly connected to each other. Depending on the specific rules used to construct them, they create structures that can resemble what is observed in real natural, as well as in artificial, complex systems. Thus, they provide simple models that allow us to use probability theory as a tool to explain the observable behaviour of real systems and to formally study and predict phenomena that are not amenable to analysis with a deterministic approach. This often leads to useful design guidelines for the development and optimal operation of real systems.

Historically, random networks has been a field of study in mathematics and statistical physics, although many models were inspired by practical questions of engineering interest. One of the early mathematical models appeared in a series of papers starting in 1959 by the two Hungarian mathematicians Paul Erdös and Alfréd Rényi. They investigated what a 'typical' graph of n vertices and m edges looks like, by connecting nodes at random. They showed that many properties of these graphs are almost always predictable, as they suddenly arise with very high probability when the model parameters are chosen appropriately. This peculiar property generated much interest among mathematicians, and their papers marked the starting point of the field of random graph theory. The graphs they considered, however, were abstract mathematical objects and there was no notion of geometric position of vertices and edges.

Mathematical models inspired by more practical questions appeared around the same time and relied on some notion of geometric locality of the random network connections. In 1957, British engineer Simon Broadbent and mathematician John Hammersley published a paper introducing a simple discrete mathematical model of a random grid in which vertices are arranged on a square lattice, and edges between neighbouring vertices are added at random, by flipping a coin to decide on the presence of each edge. This simple model revealed extreme mathematical depth, and became one of the most studied mathematical objects in statistical physics.

Broadbent and Hammersley were inspired by the work they had done during World War II and their paper's motivation was the optimal design of filters in gas masks. The gas masks of the time used granules of activated charcoal, and the authors realised that proper functioning of the mask required careful operation between two extremes. At one extreme, the charcoal was highly permeable, air flowed easily through the cannister, but the wearer of the mask breathed insufficiently filtered air. At the other extreme,

the charcoal pack was nearly impermeable, and while no poisonous gases got through, neither did sufficient air. The optimum was to have high charcoal surface area and tortuous paths for air flow, ensuring sufficient time and contact to absorb the toxin. They realised that this condition would be met in a critical operating regime, which would occur with very high probability just like Erdös and Rényi showed later for random graph properties, and they named the mathematical framework that they developed *percolation theory*, because the meandering paths reminded them of water trickling through a coffee percolator.

A few years later, in 1961, American communication engineer Edgar Gilbert, working at Bell Laboratories, generalised Broadbent and Hammersley's theory introducing a model of random planar networks in continuum space. He considered nodes randomly located in the plane and formed a random network by connecting pairs of nodes that are sufficiently close to each other. He was inspired by the possibility of providing long-range radio connection using a large number of short-range radio transmitters, and marked the birth of continuum percolation theory. Using this model, he formally proved the existence of a critical transmission range for the nodes, beyond which an infinite chain of connected transmitters forms and so long-distance communication is possible by successive relaying of messages along the chain. By contrast, below critical transmission range, any connected component of transmitters is bounded and it is impossible to communicate over large distances. Gilbert's ingenious proof, as we shall see, was based on the work of Broadbent and Hammersley, and on the theory of branching processes, which dated back to the nineteenth-century work of Sir Francis Galton and Reverend Henry William Watson on the survival of surnames in the British peerage.

Additional pioneering work on random networks appears to be the product of communication engineers. In 1956, American computer scientist Edward Moore and information theory's father Claude Shannon wrote two papers concerned with random electrical networks, which became classics in reliability theory and established some key inequalities, presented later in this book, which are important steps towards the celebrated threshold behaviours arising in percolation theory and random graphs.

As these early visionary works have been generalised by mathematicians, and statistical physicists have used these simple models to explain the behaviour of more complex natural systems, the field of random networks has flourished; its application to communication, however, has lagged behind. Today, however, there is great renewed interest in random networks for communication. Technological advances have made it plausible to envisage the development of massively large communication systems composed of small and relatively simple devices that can be randomly deployed and 'ad hoc' organise into a complex communication network using radio links. These networks can be used for human communication, as well as for sensing the environment and collecting and exchanging data for a variety of applications, such as environmental and habitat monitoring, industrial process control, security and surveillance, and structural health monitoring. The behaviour of these systems resembles that of disordered particle systems studied in statistical physics, and their large scale deployment allows us to appreciate in a real setting the phenomena predicted by the random models.

Various questions are of interest in this renewed context. The first and most basic one deals with connectivity, which expresses a global property of the system as a whole: can information be transferred through the network? In other words, does the network allow at least a large fraction of the nodes to be connected by paths of adjacent edges, or is it composed of a multitude of disconnected clusters? The second question naturally follows the first one: what is the network capacity in terms of sustainable information flow under different connectivity regimes? Finally, there are questions of more algorithmic flavour, asking about the form of the paths followed by the information flow and how these can be traversed in an efficient way. All of these issues are strongly related to each other and to the original 'classic' results on random networks, and we attempt here to give a unifying view.

We now want to spend a few words on the organisation of the book. It starts by introducing random network models on the infinite plane. This is useful to reveal *phase transitions* that can be best observed over an infinite domain. A phase transition occurs when a small variation of the local parameters of the model triggers a macroscopic change that is observed over large scales. Obviously, one also expects the behaviour that can be observed at the infinite scale to be a good indication of what happens when we consider finite models that grow larger and larger in size, and we shall see that this is indeed the case when considering scaling properties of finite networks. Hence, after discussing in Chapter 2 phase transitions in infinite networks, we spend some words in Chapter 3 on connectivity of finite networks, treating full connectivity and almost connectivity in various models. In order to deal with the information capacity questions in Chapter 5, we need more background on random networks on the infinite plane, and Chapter 4 provides all the necessary ingredients for this. Finally, Chapter 5 is devoted to studying the information capacity of a random network, applying the scaling limit approach of statistical physics in an information-theoretic setting, and Chapter 6 presents certain algorithmic aspects that arise in trying to find the best way to navigate through a random network.

The remainder of this chapter introduces different models of random networks and briefly discusses their applications. In the course of the book, results for more complex models often rely on similar ones that hold for simpler models, so the theory is built incrementally from the bottom up.

1.1 Discrete network models

1.1.1 The random tree

We start with the simplest structure. Let us consider a *tree* T composed of an infinite number of vertices, where each vertex has exactly $k > 0$ children, and draw each edge of the tree with probability $p > 0$, or delete it otherwise, independently of all other edges. We are then left with a random infinite subgraph of T, a finite realisation of which is depicted in Figure 1.1. If we fix a vertex $x_0 \in T$, we can ask how long is the line of descent rooted at x_0 in the resulting random network. Of course, we expect this to be on average longer as p approaches one. This question can also be phrased in more general

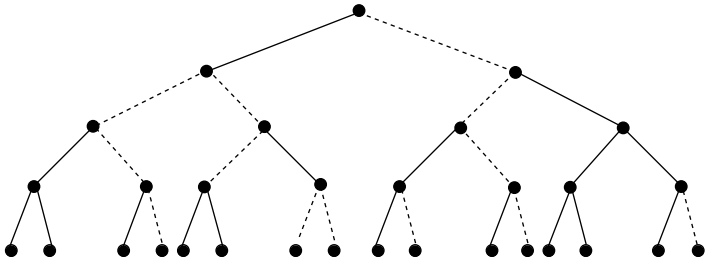

Fig. 1.1 A random tree $T(k, p)$, with $k = 2$, $p = 1/2$; deleted edges are represented by dashed lines.

terms. The distribution of the number of children at each node of the tree is called the *offspring distribution*, and in our example it has a Bernoulli distribution with parameters k and p. A natural way to obtain a random tree with arbitrary offspring distribution is by a so-called *branching process*. This has often been used to model the evolution of a population from generation to generation and it is described as follows.

Let Z_n be the number of members of the nth generation. Each member i of the nth generation gives birth to a random number of children, X_i, which are the members of the $(n+1)$th generation. Assuming $Z_0 = 1$, the evolution of the Z_i can be represented by a random tree structure rooted at Z_0 and where

$$Z_{n+1} = X_1 + X_2 + \cdots + X_{Z_n}, \tag{1.1}$$

see Figure 1.2. Note that the X_i are random variables and we make the following assumptions,

(i) the X_i are independent of each other,
(ii) the X_i all have the same offspring distribution.

The process described above could in principle evolve forever, generating an infinite tree. One expects that if the offspring distribution guarantees that individuals have a sufficiently large number of children, then the population will grow indefinitely, with positive probability at least. We shall see that there is a critical value for the expected

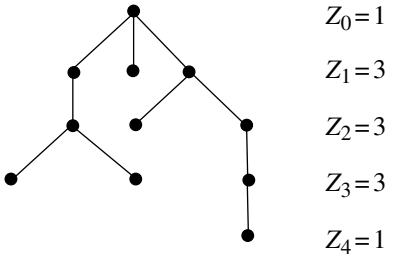

Fig. 1.2 A random tree obtained by a branching process.

offspring that makes this possible and make a precise statement of this in the next chapter. Finally, note that the branching process reduces to our original example if we take the offspring distribution to be Bernoulli of parameters k and p.

1.1.2 The random grid

Another basic structure is the random *grid*. This is typically used in physics to model flows in porous media (referred to as percolation processes). Consider an infinite square lattice $\mathbb{Z}^2$ and draw each edge between nearest neigbours with probability p, or delete it otherwise, independently of all other edges. We are then left with a random infinite subgraph of $\mathbb{Z}^2$, see Figure 1.3 for a realisation of this on a finite domain. It is reasonable to expect that larger values of p will lead to the existence of larger connected components in such subgraphs, in some well-defined sense. There could in principle even be one or more infinite connected subgraphs when p is large enough, and we note that this is trivially the case when $p = 1$.

What we have described is usually referred to as a *bond percolation model* on the square lattice. Another similar random grid model is obtained by considering a *site percolation model*. In this case each box of the square lattice is occupied with probability p and empty otherwise, independently of all other boxes. The resulting random structure, depicted in Figure 1.4, also induces a random subgraph of $\mathbb{Z}^2$. This is obtained by calling boxes that share a side neighbours, and considering connected neighbouring boxes that are occupied. It is also interesting to note that if we take a tree instead of a grid as the underlying structure, then bond and site percolation can be viewed as the same process, since each bond can be uniquely identified with a site and vice versa.

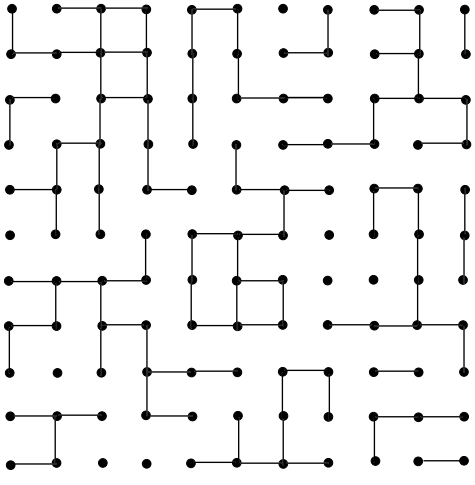

Fig. 1.3 The grid (bond percolation).

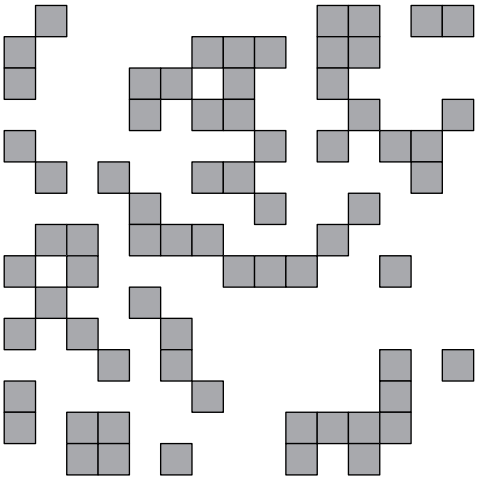

Fig. 1.4 The grid (site percolation).

1.2 Continuum network models

1.2.1 Poisson processes

Although stochastic, the models described above are developed from a predefined deterministic structure (tree and grid respectively). In continuum models this is no longer the case as the positions of the nodes of the network themselves are random and are formed by the realisation of a *point process* on the plane. This allows us to consider more complex random structures that often more closely resemble real systems.

For our purposes, we can think of a point process as a random set of points on the plane. Of course, one could think of a more formal mathematical definition, and we refer to the book by Daley and Vere-Jones (1988) for this. We make use of two kinds of point processes. The first one describes occurrences of unpredictable events, like the placement of a node in the random network at a given point in space, which exhibit a certain amount of statistical regularity. The second one accounts for more irregular network deployments, while maintaining some of the most natural properties.

We start by motivating our first definition listing the following desirable features of a somehow regular, random network deployment.

(i) *Stationarity*. We would like the distribution of the nodes in a given region of the plane to be invariant under any translation of the region to another location of the plane.
(ii) *Independence*. We would like the number of nodes deployed in disjoint regions of the plane to be independent.
(iii) *Absence of accumulation.* We would like only finitely many nodes in every bounded region of the plane and this number to be on average proportional to the area of that region.

We now describe a way to construct a process that has all the features listed above and later give its formal definition. Consider first a square of side length one. Imagine we partition this square into n^2 identical subsquares of side length $1/n$ and assume that the probability p that a subsquare contains exactly one point is proportional to the area of the subsquare, so that for some $\lambda > 0$,

$$p = \frac{\lambda}{n^2}. \tag{1.2}$$

We assume that having two or more points in a subsquare is impossible. We also assume that points are placed independently of each other. Let us look at the probability that the (random) number of points N in the whole unit square is k. This number of points is given by the sum of n^2 independent random variables, each of which has a small probability λ/n^2 of being equal to one, and which are equal to zero otherwise. It is well known and not difficult to see that, as $n \to \infty$, this sum converges to the Poisson distribution of parameter λ, which is sometimes referred to as the *law of rare events*. Indeed,

$$\begin{aligned}
\lim_{n\to\infty} P(N=k) &= \lim_{n\to\infty} \binom{n^2}{k} \left(\frac{\lambda}{n^2}\right)^k \left(1-\frac{\lambda}{n^2}\right)^{n^2-k} \\
&= \lim_{n\to\infty} \frac{n^2!}{k!(n^2-k)!} \left(\frac{\lambda}{n^2}\right)^k \left(1-\frac{\lambda}{n^2}\right)^{n^2} \left(1-\frac{\lambda}{n^2}\right)^{-k} \\
&= \lim_{n\to\infty} \frac{\lambda^k}{k!} \left(1-\frac{\lambda}{n^2}\right)^{n^2} \frac{n^2!}{n^{2k}(n^2-k)!} \left(1-\frac{\lambda}{n^2}\right)^{-k} \\
&= \frac{\lambda^k}{k!} e^{-\lambda}.
\end{aligned} \tag{1.3}$$

The construction in the unit square clearly satisfies the three desired properties, and we now want to extend it to the whole plane. Consider two disjoint unit squares and look for the distribution of the number of points inside them. This is the sum of two independent Poisson random variables, and a simple exercise in basic probability shows that it is a Poisson random variable of parameter 2λ. This leads to the idea that in our point process on the plane, the number of points in any given region A should have a Poisson distribution of parameter $\lambda|A|$, where $|\cdot|$ denotes area. This intuition leads to the following definition.

Definition 1.2.1 (Poisson process) *A random set of points $X \subset \mathbb{R}^2$ is said to be a Poisson process of density $\lambda > 0$ on the plane if it satisfies the conditions*

(i) *For mutually disjoint domains of $\mathbb{R}^2$ $D_1, \ldots, D_k$, the random variables $X(D_1), \ldots, X(D_k)$ are mutually independent, where $X(D)$ denotes the random number of points of X inside domain D.*
(ii) *For any bounded domain $D \subset \mathbb{R}^2$ we have that for every $k \geq 0$*

$$P(X(D)=k) = e^{-\lambda|D|} \frac{(\lambda|D|)^k}{k!}. \tag{1.4}$$

Note that we have $E(X([0,1]^2) = \lambda$ and the density of the process corresponds to the expected number of points of the process in the unit area. We also note that the definition does not say explicitly how to construct a Poisson process, because it does not say how the points are distributed on the plane, but only what the distribution of their number looks like.

However, a constructive procedure is suggested by the following observations. Let $B \subset A$ be bounded sets. By conditioning on the number of points inside A, and applying Definition 1.2.1, we have

$$P(X(B) = m \,|\, X(A) = m+k) = \frac{P(X(B) = m,\ X(A) = m+k)}{P(X(A) = m+k)}$$
$$= \frac{P(X(A \setminus B) = k,\ X(B) = m)}{P(X(A) = m+k)} = \frac{P(X(A \setminus B) = k)\, P(X(B) = m)}{P(X(A) = m+k)}$$
$$= \binom{m+k}{m} \left(\frac{|A| - |B|}{|A|}\right)^k \left(\frac{|B|}{|A|}\right)^m. \tag{1.5}$$

We recognise this expression as a binomial distribution with parameters $m+k$ and $|B|/|A|$. Hence, if we condition on the number of points in a region A to be $m+k$, then we can interpret the number of points that end up in $B \subset A$ as the number of successes in $m+k$ experiments with success probability $|B|/|A|$. This means that each of the $m+k$ points is randomly and uniformly distributed on A, and the positions of the different points are independent of each other.

Hence, to construct a Poisson point process in any bounded region A of the plane we should do the following: first draw a random number N of points from a Poisson distribution of parameter $\lambda|A|$, and then distribute these uniformly and independently over A.

Is it obvious now that this procedure indeed leads to a Poisson process, that is, a process that satisfies Definition 1.2.1? Strictly speaking, the answer is no: we described a necessary property of a Poisson process, but if we then find a process with this property, it is not yet clear that this property satisfies all the requirements of a Poisson process. More formally: the property is necessary but perhaps not sufficient. However, it turns out that it is in fact sufficient, and this can be seen by using a converse to (1.5) which we discuss next.

Suppose we have random variables N and $M_1, \ldots, M_r$ with the following properties: (i) N has a Poisson distribution with paramater μ, say; (ii) The conditional distribution of the vector $(M_1, \ldots, M_r)$ given $N = s$ is multinomial with parameters s and $p_1, \ldots, p_r$. We claim that under these conditions, $M_1, \ldots, M_r$ are mutually independent Poisson distributed random variables with parameters $\mu p_1, \ldots, \mu p_r$ respectively. To see this, we perform a short computation, where $m_1 + \cdots + m_r = s$,

$$P(M_1 = m_1, \ldots, M_r = m_r) = P(M_1 = m_1, \ldots, M_r = m_r | N = s) P(N = s)$$
$$= \frac{s!}{m_1! \cdots m_r!} p_1^{m_1} \cdots p_r^{m_r} e^{-\mu} \frac{\mu^s}{s!}$$
$$= \prod_{i=1}^{r} \frac{p_i^{m_i}}{m_i!} e^{-\mu p_i}, \tag{1.6}$$

proving our claim. The relevance of this is as follows: N represents the number of points in A in the construction above, and $M_1, \ldots, M_r$ represent the number of points ending up in regions $B_1, \ldots, B_r$ into which we have subdivided A. Since the properties of the construction are now translated into properties (i) and (ii) above, the conclusion is that the number of points in disjoint regions are mutually independent with the correct Poisson distribution. Hence we really have constructed a Poisson process on A.

Finally, we note that the independence property of the process also implies that if we condition on the event that there is a point at $x_0 \in \mathbb{R}^2$, apart from that point, the rest of the process is not affected by the conditioning event. This simple fact can be stated with arbitrarily high level of formality using *Palm calculus* and we again refer to the book of Daley and Vere-Jones (1988) for the technical details.

The definition of a Poisson point process can be generalised to the case when the density is not constant over the plane, but it is a function of the position over $\mathbb{R}^2$. This gives a non-stationary point process that is useful to describe non-homogeneous node deployments. We first describe a way to construct such a process from a standard Poisson point process and then give a formal definition. Let X be a Poisson point process with density λ on the plane, and let $g : \mathbb{R}^2 \to [0, 1]$. Consider a realisation of X and delete each point x with probability $1 - g(x)$, and leave it where it is with probability $g(x)$, independently of all other points of X. This procedure is called *thinning* and generates an inhomogeneous Poisson point process of density function $\lambda g(x)$. The formal definition follows.

Definition 1.2.2 (Inhomogeneous Poisson process) *A countable set of points $X \subset \mathbb{R}^2$ is said to be an inhomogeneous Poisson process on the plane with density function $\Lambda : \mathbb{R}^2 \to [0, \infty)$, if it satisfies the conditions*

(i) *For mutually disjoint domains of $\mathbb{R}^2$ $D_1, \ldots, D_k$, the random variables $X(D_1), \ldots, X(D_k)$ are mutually independent, where $X(D)$ denotes the random number of points inside domain D.*
(ii) *For any bounded domain $D \subset \mathbb{R}^2$ we have that for every $k \geq 0$*

$$P(X(D) = k) = e^{-\int_D \Lambda(x)dx} \frac{[\int_D \Lambda(x)dx]^k}{k!}. \tag{1.7}$$

In the case $\int_D \Lambda(x)dx = \infty$, this expression is interpreted as being equal to zero.

We claim that the thinning procedure that we decribed above leads to an inhomogeneous Poisson process with density function $\lambda g(x)$. To see this, we argue as follows.

We denote by $\tilde{X}$ the point process after the thinning procedure. The independence property is immediate from the construction, and the distribution of $\tilde{X}$ can be computed as follows:

$$P(\tilde{X}(A) = k) = \sum_{i=k}^{\infty} P(X(A) = i) P(\tilde{X}(A) = k | \, X(A) = i). \tag{1.8}$$

We have from (1.5) that given the event $\{X(A) = i\}$, the i points of X in A are uniformly distributed over A. Thus the conditional distribution of $\tilde{X}$ given $X(A) = k$ is just

$$P(\tilde{X}(A) = 1 | X(A) = 1) = |A|^{-1} \int_A g(x)dx, \tag{1.9}$$

and more generally,

$$P(\tilde{X}(A) = k | X(A) = i) = \binom{i}{k} \left(|A|^{-1} \int_A g(x)dx \right)^k \times \left(1 - |A|^{-1} \int_A g(x)dx \right)^{i-k}. \tag{1.10}$$

Hence,

$$\begin{aligned} P(\tilde{X}(A) = k) &= e^{-\lambda|A|} \frac{(\lambda \int_A g(x)dx)^k}{k!} \\ &\quad \times \sum_{i=k}^{\infty} \frac{(\lambda|A|[1 - |A|^{-1} \int_A g(x)dx])^{i-k}}{(i-k)!} \\ &= e^{-\lambda|A|} \frac{(\lambda \int_A g(x)dx)^k}{k!} e^{\lambda|A|(1-|A|^{-1} \int_A g(x)dx)} \\ &= \frac{(\lambda \int_A g(x)dx)^k}{k!} e^{-\lambda \int_A g(x)dx}. \end{aligned} \tag{1.11}$$

A few final remarks are appropriate. The definition of an inhomogeneous Poisson point process is more general than the described thinning procedure: since g is defined from $\mathbb{R}^2$ into $[0, \infty)$, it also allows accumulation points. Note also that $E(X([0,1]^2) = \int_{[0,1]^2} \Lambda(x)dx$, which is the expected number of points of the process in the unit square. Finally, note that one can obtain a Poisson process from its inhomogeneous version by taking $\Lambda(x) \equiv \lambda$.

1.2.2 Nearest neighbour networks

We can now start to consider networks of more complex random structure. *Nearest neighbour networks* represent a natural mathematical construction that has been used, for example, to model multi-hop radio transmission, when a message is relayed between two points along a chain of successive transmissions between nearest neighbour stations.

Let X be a Poisson point process of unit density on the plane. We place edges between each point of X and its k nearest neighbours in Euclidean distance, where k is some chosen positive integer. The result is a random network, whose structure depends on the random location of the points and on the choice of k.

Note that the density of the Poisson process is just a scaling factor that does not play a role in the geometric properties of the graph. To see this, take a realisation of the Poisson process, and imagine scaling all lengths by a certain factor, say $1/2$. As shown