Exploration on and of networks

Güldas, H.

CHAPTER 1
Introduction

The present thesis consists of two parts. Part I focusses on random walks on dynamic configuration models, Part II focusses on union complexity of random disk regions.

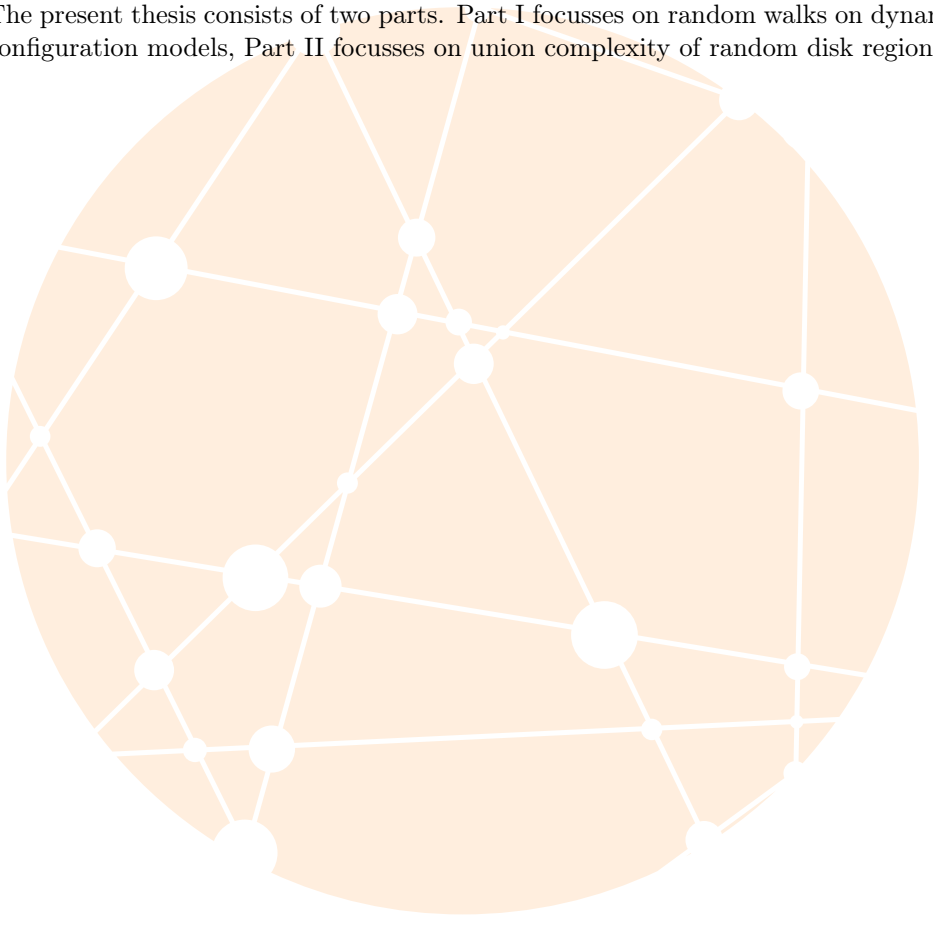

§1.1 Introduction to Part I

The theory of random graphs provides abstract mathematical tools for studying real-world problems related to complex networks. Questions concerning the formation, structure and functionality of real-world networks have been studied via random graph models. Many questions involve also a random process on the network. Mathematical notions for modelling random processes on networks include random walks, percolation and interacting particle systems. These notions have been studied extensively in the context of *static* random graphs. However, many real-world networks are *dynamic* in nature: their structure changes over time. Thus, it is natural to extend static settings to dynamic settings. Part I of the thesis contains our work in this direction: it studies mixing properties of random walks on dynamic random graphs.

In Section 1.1.1, we give an overview of random graph theory in the context of complex networks. We start by discussing the historical development of the study of complex networks. After that, we review two basic random graph models, the Erdős-Rényi Random Graph model and the configuration model, which are among the most widely studied models. We close Section 1.1.1 by discussing, very briefly, other random graph models that were invented in the context of complex networks. In Section 1.1.2, we give a review of the theory of random walks on static random graphs, with a focus on mixing times. We introduce the notion of mixing times in the more general framework of Markov chains and discuss several aspects. We close Section 1.1.2 by reviewing the literature on mixing times of random walks on random graphs. In Section 1.1.3, we review the modest literature on random walks on dynamic random graphs that has accumulated so far and list our contributions.

§1.1.1 Complex networks and random graphs

Since 1999, there has been a surge of interest in the study of real-world networks, and network science has emerged as an independent academic discipline. A number of journals are devoted to the field, and many universities offer courses and special programs on network science and *complex networks*. The term *'complex networks'* was used by Steven Strogatz in his 2001 paper [90] to refer to networks with non-trivial topological features, unlike regular networks such as chains, grids, lattices and fully-connected graphs. Many real-world networks are complex in this sense.

The research on complex networks was pioneered in a series of papers by Barabási and co-authors [14, 15, 3], Watts and Strogatz [98], and Newman and Watts [76, 77]. Their aim was to explain the mechanisms behind the formation of networks that give rise to certain statistical features that are observed in real-world networks, but cannot be explained by randomness alone. Empirical studies of real-world networks have shown that many real-world networks have certain key features in common, such as:

- **Sparsity:** the average number of links per node does not grow too large with the network size.
- **Power-law degrees:** the empirical distribution of the number of edges per ver-

tex follows a *power-law* distribution. In a power-law distribution, the frequency of vertices with k edges is roughly proportional to $k^{-\gamma}$ for some $\gamma > 2$ and for large k.

- **Small-world phenomenon:** the distance between typical vertex pairs does not grow too fast with the size of the network.
- **High clustering:** if two vertices are connected to a third vertex, then it is likely that these two vertices are connected to each other as well.

The above properties are observed in networks in many different settings, such as the World Wide Web, social networks and biological networks. For a review of early studies on real-world networks and their statistical properties, see [2, 78].

Why is it that different networks have so many features in common and how do these features affect the functionality of the network, for example, the vulnerability of a computer network or the spread of a disease in a population? The mathematical tools that can tackle these problems are provided by the mathematical theory of random graphs. Before going into details, we review two random graph models that are widely studied in the mathematics community.

Erdős-Rényi random graph model

The theory of random graphs was initiated by Paul Erdős and Alfred Rényi in their seminal papers dating back to 1959–1961 [42, 43, 44]. Although Erdős and Rényi studied random graphs for their own sake, their random graph model was motivated by the use of the *probabilistic method* for solving problems in extremal graph theory; see for example [41]. The probabilistic method can be summarised as the use of probabilistic arguments to prove deterministic statements, often in the form of existence results of combinatorial structures with certain properties. The basic principle is as follows: if a random combinatorial structure has a certain property with a positive probability, then there exists at least one deterministic structure with that property. The canonical reference for the probabilistic method is Alon and Spencer's textbook on the subject [8].

The remarkable discovery of Erdős and Rényi was the *threshold phenomenon* (also called *phase transition*) exhibited by random graphs. Their goal was to determine the number of edges a random graph must have in order to acquire a certain property, such as containing a set of given subgraphs, containing a *'giant'* component or being connected. They showed that these properties emerge rather suddenly and formalised this emergence in mathematical terms. Let us denote by $\mathrm{ER}_n(m)$ the random graph distributed uniformly over the set of graphs on n vertices with m edges. In the formal setting, one considers a family of random graphs whose number of vertices n tends to infinity and whose number of edges $m = m(n)$ depends on the number of vertices. In [43], Erdős and Rényi showed that $\mathrm{ER}_n(m)$ acquires certain properties, with high probability, only when the number of edges m grows faster than some threshold function $A(n)$, and they call this *threshold phenomenon*.

The line of research initiated by Erdős and Rényi was later continued by others. The body of results with a similar flavor and of techniques used for proving these

results make up the core of the monograph entitled *Random Graphs* by Bollobás [25]. In fact, the majority of the works study a different but related model. For $p \in [0,1]$, let $\mathrm{ER}_n(p)$ be the random graph on n vertices where each of the $\binom{n}{2}$ edges is present with probability p, independently of each other. Although, in their original papers, Erdős and Rényi studied $\mathrm{ER}_n(m)$, the model with $\mathrm{ER}_n(p)$ is referred to as the *Erdős-Rényi Random Graph (ERRG)* model. The two models are practically interchangeable when m is close to $p\binom{n}{2}$ (see [25, Theorem 2.2]), and $\mathrm{ER}_n(p)$ also exhibits threshold phenomena for many properties. One particular feature of $\mathrm{ER}_n(p)$ in the sparse regime (where $p = \lambda/n$ for some $\lambda \in (0,\infty)$) is that the local neighborhoods of the vertices are tree-like in the following sense: the neighborhood of a uniformly chosen vertex in $\mathrm{ER}_n(p)$ can be coupled to the family tree of a branching process with a Poisson offspring distribution. This can be used to prove the phase transition for the emergence of a giant component in $\mathrm{ER}_n(p)$ (see [93, Chapter 4]).

While the ERRG model exhibits many mathematically interesting phenomena, it fails to capture many features exhibited by real-world networks. When $p = \lambda/n$ for some $\lambda \in (0,\infty)$, $\mathrm{ER}_n(p)$ is sparse, the average degree converges in probability to λ as n tends to ∞, and the distribution of the degree of a fixed vertex is close to a Poisson distribution with parameter λ for large n. Most real-world networks have power-law degree distributions, thus the ERRG model cannot capture the scale-free property of real-world networks. Another drawback of the model is that it is completely symmetric, in the sense that the vertices have the same degree distribution. To remedy these shortcomings, generalisations of ERRG model have been suggested, such as the generalised random graph model [29], the inhomogeneous random graph model [26], the Chung-Lu model [33] and the Norros-Reittu model [79]. For an extensive review, see [93, Chapter 6].

Configuration model

One of the problems with the ERRG model and its generalisations is that the resulting graph contains isolated vertices with positive probability in the sparse regime, which makes them impractical in many cases for the study of real-world networks. One possible solution is to fix the degrees of the graph beforehand and generate a random graph with the prescribed degree sequence. The *configuration model* is one such model. Let $\mathbf{d} = (d_1, \dots, d_n)$ be the given degree sequence on n vertices. The configuration model is constructed by attaching d_i half-edges to vertex v_i for each $i \in [n] \coloneqq \{1, \dots, n\}$ and then pairing these half-edges uniformly at random. A pairing of the half-edges is called a *configuration* and a uniformly distributed configuration is denoted by $\mathrm{Conf}_n(\mathbf{d})$. The resulting graph is denoted by $\mathrm{CM}_n(\mathbf{d})$, which need not be *simple*. A simple graph is a graph that does not contain any loop (an edge with the same vertex at both ends) or multiple edges between any pair of vertices. For graphs that are not necessarily simple, the term *multi*-graph is used. In the rest of this section, we use the term graph to refer to multi-graphs.

In the configuration model, each configuration is identified with a graph, and a graph can be obtained via several distinct configurations. Letting $\ell_n = \sum_{i\in[n]} d_i$, we see that there are $(\ell_n - 1)!! \coloneqq (\ell-1) \times (\ell-3) \times \cdots \times 3 \times 1$ distinct configurations in

total, so each configuration has probability $\frac{1}{(\ell_n-1)!!}$. We can identify a graph G with a matrix $(X_{ij})_{i,j\in[n]}$, where X_{ij} denotes the number of edges between vertices v_i and v_j for $i \neq j$ and X_{ii} the number of self-loops at v_i. This gives

$$d_i = X_{ii} + \sum_{j\in[n]} X_{ij}, \quad i \in [n].$$

The distribution of $\mathrm{CM}_n(\mathbf{d})$ is given by the following proposition:

Proposition ([93, Proposition 7.7]). *Let $G = (x_{ij})_{i,j\in[n]}$ be a graph on n vertices such that*

$$d_i = x_{ii} + \sum_{j\in[n]} x_{ij} \text{ for } i \in [n].$$

Then

$$\mathbb{P}(\mathrm{CM}_n(\mathbf{d}) = G) = \frac{1}{(\ell_n-1)!!} \frac{\prod_{i\in[n]} d_i!}{\prod_{i\in[n]} 2^{x_{ii}} \prod_{1\le i\le j\le n} x_{ij}!}.$$

From this, we see that $\mathrm{CM}_n(\mathbf{d})$ is not distributed uniformly over the set of graphs on n vertices with degree sequence $\mathbf{d}$. However, when G is a simple graph, i.e. when $x_{ii} = 0$ for all $i \in [n]$ and $x_{ij} \in \{0, 1\}$ for all $i, j \in [n]$, we have

$$\mathbb{P}(\mathrm{CM}_n(\mathbf{d}) = G) = \frac{\prod_{i\in[n]} d_i!}{(\ell_n-1)!!}. \tag{1.1}$$

This shows that the distribution of $\mathrm{CM}_n(\mathbf{d})$ conditionally on being simple is uniform on the set of simple graphs on n vertices with degree sequence $\mathbf{d}$.

The configuration model was introduced by Bollobás [24] to count labelled simple regular graphs. This work was inspired by ideas introduced by Bender and Canfield [17]. To count the labelled simple graphs, we can use the configuration model as follows. We know that, conditional on being simple, $\mathrm{CM}_n(\mathbf{d})$ is distributed uniformly on the set of simple graphs with degree sequence $\mathbf{d}$. Letting $N_n(\mathbf{d})$ denote the number of simple graphs with degree sequence $\mathbf{d}$, we see that for any simple graph G with degree sequence $\mathbf{d}$,

$$\mathbb{P}(\mathrm{CM}_n(\mathbf{d}) = G \mid \mathrm{CM}_n(\mathbf{d}) \text{ is simple}) = \frac{1}{N_n(\mathbf{d})}.$$

Combining this with (1.1), we see that

$$N_n(\mathbf{d}) = \frac{(\ell_n-1)!!}{\prod_{i\in[n]} d_i!} \mathbb{P}(\mathrm{CM}_n(\mathbf{d}) \text{ is simple}).$$

So, if we can estimate the probability that the resulting graph is simple, then we can estimate the number of simple graphs with the given degree sequence. This is exactly how Bollobás obtained his result on the number of simple regular graphs.

The configuration model and uniform simple graphs with given degree sequences were later studied by others. Łuczak showed that when the minimum degree is three, the resulting graph is connected with a probability tending to 1 as n tends to infinity

[68]. Molloy and Reed studied the emergence of a giant component and its size for sparse random graphs with nonregular degree sequences [72, 73]. They established the phase transition for the giant component, similar to that of the ERRG model. In fact, some results for $\mathrm{ER}_n(p)$ and $\mathrm{ER}_n(m)$ can be obtained through the results for uniform random graph with a given degree sequence, by noting that the latter has the same distribution as the former conditioned on the degree sequence. Fernholz and Ramachandran [48] established that the diameter of the uniform random graph with a given degree sequence on n vertices is $c \log n + o(\log n)$, for some $c > 0$ depending on the degree sequence, under certain sparsity conditions. Using this result and the above observation, they refined earlier results on the diameter of the $\mathrm{ER}_n(p)$ [32, 26]. The diameter and typical distances for the configuration model, especially for the scale-free degree sequences, was also studied by van der Hofstad and co-authors (see for example [31, 95, 96]).

When studying the configuration model, one typically considers fixed degree sequences indexed by the number of vertices n, i.e. $(\mathbf{d}_n = (d_i^n)_{i\in[n]})_{n\in\mathbb{N}}$ with certain statistical properties, and studies properties of $\mathrm{CM}_n(\mathbf{d}_n)$ as n tends to infinity. A typical condition on the degrees is the convergence of the empirical degree distribution to a deterministic limit distribution. Let D_n be the random variable whose distribution function is given by

$$F_n(x) = \sum_{i=1}^{n} \mathbf{1}_{\{d_i^n \le x\}},$$

i.e., D_n is the degree of a uniformly chosen vertex on the graph with n vertices whose degree sequence is given by $\mathbf{d}_n = (d_i^n)_{i\in[n]}$ and F_n is the empirical distibution of the degrees. The following regularity conditions on the degrees are common:

Condition 1.1.1 (Regularity of the degrees)**.** *There exists an integer-valued random variable D with $\mathbb{P}(D > 0) = 1$ such that*

(a) $D_n \xrightarrow{d} D$ *as* $n \to \infty$*, where* $\xrightarrow{d}$ *denotes convergence in distribution,*

(b) $\lim_{n\to\infty} \mathbb{E}[D_n] = \mathbb{E}[D] < \infty$*,*

(c) $\lim_{n\to\infty} \mathbb{E}[D_n^2] = \mathbb{E}[D^2] < \infty$*.*

In the context of complex networks, the limiting degree distribution D is usually assumed to follow a power-law distribution. Conditions 1.1.1(a)-(b) ensure that the degree sequences are sparse. Condition 1.1.1(c) controls the variance of the degrees and the maximum degree. Under Conditions 1.1.1(a)-(c), the asymptotics of the probability that $\mathrm{CM}_n(\mathbf{d_n})$ is a simple graph is given by the following theorem:

Theorem 1.1.2 ([93, Theorem 7.12])**.** *Assume that the sequence of degree sequences $(\mathbf{d}_n)_{n\in\mathbb{N}}$ satisfies Conditions 1.1.1(a)-(c). Then the probability that $\mathrm{CM}_n(\mathbf{d}_n)$ is simple converges to*

$$\mathrm{e}^{-\nu/2-\nu^2/4},$$

where $\nu = \mathbb{E}[D(D-1)]/\mathbb{E}[D]$.

Using this result and the above arguments about the number of simple graphs with a given degree sequence, we see that, under Conditions 1.1.1(a)-(c), $N_n(\mathbf{d_n})$ grows

asymptotically as

$$\frac{(\ell_n - 1)!!}{\prod_{i\in[n]} d_i^n!} \mathrm{e}^{-\nu/2-\nu^2/4}.$$

As in the ERRG model, one of the characteristics of the configuration model is its tree-like structure. The neighborhood of a uniformly chosen vertex can be coupled to the family tree of a branching process whose offspring distribution is related to the degree sequence. Again, this can be used to prove results about the phase transition for the giant component, typical distances and the diameter (see [94, Chapters 3 and 6]). The tree-like structure is best formalized in the framework of *local weak convergence*, which was developed in [20, 6, 5]. Under certain regularity conditions, local weak limits of configuration models and Erdős-Rényi Random Graph models are unimodular Galton-Watson trees [94]. We emphasize the tree-like structure, since it is also important for the study of random walks on these graphs. As we will see later, it plays a crucial role in the analysis of the mixing time of random walks.

Apart from the counting perspective, early studies on the configuration model also focused on the problem of uniformly sampling a graph with a given degree sequence. The simplest approach is to repeatedly sample from the configuration model until the resulting graph is simple. In [100], the author adopts this approach to generate a uniform simple regular graph. However, when the degrees are large this method is prohibitively inefficient, as is evident from Theorem 1.1.2, because ν is large. A sampling procedure is considered to be efficient when the expected running time of the underlying algorithm is polynomial in the size of the graph. Later, McKay and Wormald [71] relaxed the degree constraints by incorporating a *switching* mechanism into this procedure, and Gao and Wormald [52] relaxed the degree constraints further. If one is content with *approximately* uniform samples, then Markov Chain Monte Carlo methods are a fallback option. We will discuss this approach later in more detail, in the context of dynamic random graphs.

While the configuration model provides great flexibility in terms of degree sequences, particularly in the study of scale-free complex networks, it exhibits a low clustering compared to many real-world networks. This is mainly due to its tree-like structure. To overcome this shortcoming, extensions of the configuration model were proposed in which clustering information can be incorporated. Examples of these models include the configuration model with household structure [92], and the hierarchical configuration model [97]. For a more thorough discussion of these extensions, see [93, Chapter 7].

Complex networks and random graphs

A big leap in the study of random graphs came with the works on complex networks appearing in the physics literature in late 1990's. The small-world model of Watts and Strogatz [98] and the preferential attachment model of Barabási and Albert [14] led to an explosion of research in the field of random graphs. The study of complex networks within the physics community was largely based on the empirical study of real-world networks and on simulation of random graph models. These studues usually lack mathematical rigour. However, mathematicians were quick to fill in

the gaps and push the theory forward. For example, Bollobás, Riordan, Spencer and Tusnády [27] showed that the preferential attachment mechanism suggested by Barabási and Albert [14] indeed gives rise to a power-law degree sequence, as they predicted using simulations and a heuristic argument. There have been many works on this theme since 2000, on many different aspects of random graph models related to the phenomena occuring in real-world networks. For a detailed account of these, we refer to the book by Durrett [38] on random graph dynamics and the two books by van der Hofstad [93, 94] on random graphs and complex networks.

§1.1.2 Random walks on static random graphs

In this section, we give an overview of the study of random walks on static random graphs. The classical theory of random walks focuses on random walks on infinite graphs with a simple structure, such as $\mathbb{Z}^d$. The canonical example is the nearest-neighbour simple random walk on $\mathbb{Z}^d$, in which jumps to neighbouring vertices occur with equal probability at each time unit. Results on this type of random walks go back almost a hundred years. In a classical result for the nearest-neighbour simple random walk on $\mathbb{Z}^d$ [84], Pólya established that when $d \leq 2$ the random walk comes back to its initial position infinitely many times with probability one, i.e., the random walk is *recurrent*, while when $d \geq 3$ the random walk visits every vertex at most finitely many times with probability one, i.e., the random walk is *transient*. Also random walks in random environments have been considered. In this setting, the jump probabilities are themselves random. This line of research mostly focuses on qualitative behaviour of the random walk, such as recurrence versus transience, or laws of large numbers, central limit theorems and large deviation principles.

The study of random walks on more general but finite graphs has attracted much attention more recently. In contrast to the classical theory of random walks, this line of research focuses on the asymptotics of the finite-time behaviour of the random walk when the size of the graph tends to infinity. Within this framework, the main objects of study are hitting times, mixing times and cover times. The simplest example is again the simple random walk, in which jumps to neighbouring vertices occur with equal probabilities, as in the case of the simple random walk on $\mathbb{Z}^d$. However, since the degrees need not be the same for all vertices, the jump probabilities are inhomogeneous. These inhomogeneities make it harder to study random walks on general graphs in detail. In a more complicated setting, the underlying graph is also random. While random walks on random graphs can be viewed as an example of random walks in random environments, the research focuses on different settings and different questions, and the mathematical techniques used in the two areas differ greatly. One of the main reasons for this is that the natural questions in finite and infinite graph settings are vastly different.

In this thesis, we concentrate on the mixing time of random walks on random graphs. We first review the mixing time of general Markov chains.

Mixing time of general Markov chains

Let $(X_t)_{t\in\mathbb{N}_0}$ be a time-homogeneous, irreducible and aperiodic Markov chain on a finite state space $\mathcal{S}$ of size n, whose transition probabilities are given by a $n\times n$ matrix P whose rows and columns are indexed by the elements of $\mathcal{S}$, i.e.,

$$\mathbb{P}(X_t = y \mid X_{t-1} = x) = P(x,y) \text{ for all } t \in \mathbb{N} \text{ and } x, y \in \mathcal{S}.$$

Then there exists a unique probability measure π on $\mathcal{S}$, called the *stationary distribution*, which solves the equation

$$\pi P = \pi,$$

where π is viewed as a row vector of length n. Moreover, t-step transition probabilities converge to π as $t \to \infty$, in the following sense:

$$\lim_{t\to\infty} P^t(x,y) = \pi(y) \text{ for any } x, y \in \mathcal{S}.$$

In practice, one is interested in the speed of convergence towards π. The usual way of measuring how far the distribution of the Markov chain is from the stationary distribution is the *total variation* distance. The total variation distance between two probability measures μ and ν on a countable state space $\mathcal{S}$ is defined as

$$\|\mu - \nu\|_{\text{TV}} \coloneqq \frac{1}{2}\sum_{x\in\mathcal{S}} |\mu(x) - \nu(x)|.$$

In the literature, convergence to stationarity is usually studied in the *worst-case* setting as we explain next. Let

$$\mathcal{D}_x(t) \coloneqq \|P^t(x,\cdot) - \pi(\cdot)\|_{\text{TV}} \text{ and } \mathcal{D}(t) = \max_{x\in\mathcal{S}} \mathcal{D}_x(t).$$

It is easy to see that both $\mathcal{D}_x(t)$, for all $x \in \mathcal{S}$, and $\mathcal{D}(t)$, are non-decreasing in t. A classical result states that the convergence to the stationary distribution measured in total variation distance happens exponentially fast:

Theorem (Theorem 4.9 in [65]). *Suppose that P is the matrix of transition probabilities of an irreducible and aperiodic Markov chain on a finite state space $\mathcal{S}$ with stationary distribution π. Then there exist two constants $\alpha \in (0,1)$ and $C > 0$ such that*

$$\mathcal{D}(t) \le C\alpha^t.$$

Using this theorem, we see that the asymptotic rate of exponential convergence can be bounded from above

$$\limsup_{t\to\infty} \mathcal{D}(t)^{1/t} \le \alpha.$$

When the Markov chain is also time-reversible, i.e., if the detailed balance condition

$$\pi(x)P(x,y) = \pi(y)P(y,x), \quad x, y \in \mathcal{S},$$

holds, then the asymptotic rate of exponential convergence is given by

$$\lim_{t\to\infty} \mathcal{D}(t)^{1/t} = \lambda_*,$$

where

$$\lambda_* = \max\{|\lambda| : \lambda \text{ is an eigenvalue of } P \text{ and } \lambda \neq 1\}$$

For a proof of this, see for example [65, Corollary 12.7].

While the above setting describes the asymptotic convergence properties of a fixed Markov chain P as $t \to \infty$, research in the past few decades has focused on the finite-time asymptotics of a family of Markov chains indexed by the size of the state space in the limit as the size tends to infinity. In this setting, a precision level ε is fixed and one usually looks at the first time when the distance to the stationary distribution falls below ε. The *ε-mixing time* is defined as

$$t_{\text{mix}}(\varepsilon) \coloneqq \min\{t \in \mathbb{N} : \mathcal{D}(t) \leq \varepsilon\}.$$

In fact, mixing times at different precision levels can be related to each other by using the sub-multiplicativity of $\mathcal{D}(t)$. Indeed, for all $s, t \in \mathbb{N}$, $\mathcal{D}(s+t) \leq \mathcal{D}(s)\mathcal{D}(t)$, and therefore it is possible to show that (see [65, Section 4.5])

$$t_{\text{mix}}(\varepsilon) \leq \lceil \log_2 \varepsilon^{-1} \rceil t_{\text{mix}}(1/4).$$

The techniques used to study mixing times vary greatly, and in specific examples one has to come up with ad hoc methods. However, there are a few techniques that can be used in more general settings. One such method is *coupling*. A coupling of two probability distributions μ and ν on state spaces $\mathcal{S}_1$ and $\mathcal{S}_2$, respectively, is a pair of random variables $(X, Y) \in \mathcal{S}_1 \times \mathcal{S}_2$ such that the marginal distribution of X is μ and the marginal distribution of Y is ν, i.e., $\mathbb{P}(X \in \cdot) = \mu(\cdot)$ and $\mathbb{P}(Y \in \cdot) = \nu(\cdot)$. The notion of coupling is related to total variation distance via the following relation (see [91, Chapter 1]):

$$\|\mu - \nu\|_{\text{TV}} = \inf\{\mathbb{P}(X \neq Y) : (X, Y) \text{ is a coupling of } \mu \text{ and } \nu\}. \tag{1.2}$$

To see how this can be used in the context of mixing times of Markov chains, we need to define another notion of distance to stationarity. Let $\overline{\mathcal{D}}(t) = \max_{x,y \in \mathcal{S}} \|P^t(x, \cdot) - P^t(y, \cdot)\|_{\text{TV}}$. Although the definition of this quantity does not involve the stationary distribution directly, it satisfies the inequalities (see [65, Lemma 4.10])

$$\mathcal{D}(t) \leq \overline{\mathcal{D}}(t) \leq 2\mathcal{D}(t). \tag{1.3}$$

Given a Markov chain on state space $\mathcal{S}$ with transition matrix P, a *Markovian coupling* of two copies of the same Markov chain is a Markov chain $(X_t, Y_t)_{t \in \mathbb{N}_0}$ on $\mathcal{S} \times \mathcal{S}$ such that, for all $x, y, x', y' \in \mathcal{S}$ and $t \in \mathbb{N}$,

$$\begin{aligned} \mathbb{P}(X_t = x' \mid X_{t-1} = x, Y_{t-1} = y) &= P(x, x'), \\ \mathbb{P}(Y_t = y' \mid X_{t-1} = x, Y_{t-1} = y) &= P(y, y'). \end{aligned}$$

Consider a Markovian coupling of two copies of the Markov chain given by P, with the additional property

$$\text{if } X_t = Y_t, \text{ then } X_s = Y_s \text{ for all } s \geq t,$$

so that whenever the two components meet they stick together. Let τ_{couple} be the first time the two components meet. Then $\{X_t \neq Y_t\} \subset \{\tau_{\text{couple}} > t\}$ and hence $\mathbb{P}(X_t \neq Y_t) \leq \mathbb{P}(\tau_{\text{couple}} > t)$. Using (1.2) and the first inequality of (1.3), we see that

$$\mathcal{D}(t) \leq \mathbb{P}(\tau_{\text{couple}} > t).$$

For examples of how to obtain upper bounds on mixing times via coupling, see [65, Chapter 5] and references therein.

Another useful tool to obtain upper bounds on mixing times is the notion of *strong stationary times*. A strong stationary time for a Markov chain on state space $\mathcal{S}$ starting from state $x \in \mathcal{S}$ with stationary distribution π is a randomised stopping time τ such that, for all $t \in \mathbb{N}_0$ and $y \in \mathcal{S}$,

$$\mathbb{P}_x(X_t = y, \tau = t) = \pi(y)\mathbb{P}_x(\tau = t),$$

or equivalently,

$$\mathbb{P}_x(X_t = y, \tau \leq t) = \pi(y)\mathbb{P}_x(\tau \leq t),$$

where $\mathbb{P}_x$ denotes the law of the Markov chain starting from $x \in \mathcal{S}$. In words, X_τ has distribution π and is independent of τ. The total variation distance to the stationary distribution can be bounded by the tail probability of the strong stationary time (see [65, Proposition 6.11]):

$$\mathcal{D}_x(t) \leq \mathbb{P}_x(\tau > t).$$

In fact, strong stationary times may be optimal in some sense. To explain this, we need another notion of distance to stationarity, called the *separation distance*, defined by

$$s_x(t) \coloneqq \max_{y \in \mathcal{S}} \left[1 - \frac{P^t(x,y)}{\pi(y)}\right].$$

Let $s(t) = \max_{x \in \mathcal{S}} s_x(t)$. The total variation distance to the stationary distribution can be bounded by the separation distance (see [65, Lemma 6.16])

$$\mathcal{D}_x(t) \leq s_x(t).$$

If the Markov chain is reversible, then (see [65, Lemma 6.17])

$$s(2t) \leq 1 - (1 - \bar{\mathcal{D}}(t))^2 \leq 2\bar{\mathcal{D}}(t) \leq 4\mathcal{D}(t).$$

There are strong stationary times that are optimal in the following sense:

Proposition ([65, Proposition 6.21])**.** *For every starting state $x \in \mathcal{S}$, there exists a strong stationary time τ such that, for all $t \in \mathbb{N}_0$*

$$s_x(t) = \mathbb{P}_x(\tau > t).$$

For examples of how to obtain upper bounds on mixing times via strong stationary times, see [65, Chapter 6] and references therein. In our proofs in Chapters 2,3 and 4, we use a randomised stopping time that is not a strong stationary time, but very close to one.

The usual setting in recent works on mixing times of Markov chains involves a family of Markov chains $P^{(n)}$, $n \in \mathbb{N}$, where $P^{(n)}$ is a Markov chain on a state space $\mathcal{S}^{(n)}$ and the size of the state space tends to infinity as n tends to infinity. Letting $t^{(n)}_{\text{mix}}(\varepsilon)$ denote the ε-mixing time of the Markov chain $P^{(n)}$, one is usually interested in finding an expression for $t^{(n)}_{\text{mix}}(\varepsilon)$ as a function of n. Within this setting, a remarkable phenomenon observed for many different families of Markov chains is the so-called *cutoff* phenomenon, in which $t^{(n)}_{\text{mix}}(\varepsilon)$ virtually does not depend on ε. More precisely, a family of Markov chains is said to exhibit a cutoff phenomenon if, for all $\varepsilon > 0$,

$$\lim_{n\to\infty} \frac{t^{(n)}_{\text{mix}}(\varepsilon)}{t^{(n)}_{\text{mix}}(1-\varepsilon)} = 1.$$

In this case, the distance to the stationary distribution stays near 1 for awhile and subsequently drops to near 0 in a relatively short time window. The notion of *cutoff window* is formalised in the following way. A family of Markov chains is said to have cutoff with a window of size $O(w_n)$ if $w_n = o(t^{(n)}_{\text{mix}}(\frac{1}{4}))$ and

$$\lim_{\alpha\to-\infty} \liminf_{n\to\infty} \mathcal{D}^{(n)}(t^{(n)}_{\text{mix}}(\tfrac{1}{4}) + \alpha w_n) = 1,$$
$$\lim_{\alpha\to\infty} \limsup_{n\to\infty} \mathcal{D}^{(n)}(t^{(n)}_{\text{mix}}(\tfrac{1}{4}) + \alpha w_n) = 0.$$

In early works, Diaconis and Shahshahani [36] established cutoff for uniform transposition random walks on permutation groups, and Aldous [7] for random walks on hypercubes. Later, cutoff was observed for many other Markov chains, including random walks on groups, random walks on random graphs, and interacting particle systems. For a more detailed discussion of the cutoff phenomenon, we refer to [65, Chapter 18].

Mixing times of random walks on static random graphs

The study of mixing times of random walks on random graphs has received considerable attention in last few decades. While these studies are interesting in their own right, they are motivated by problems in combinatorics and computer science [35]. A Markov chain on a state space of size n is said to be *rapidly mixing* if the mixing time is polynomial in $\log n$. Rapid mixing is especially useful in algorithmic applications, where one usually deals with very large state spaces, typically exponential in the system size. In such cases, rapid mixing is used to obtain polynomial time algorithms. Many examples of random walks on random graphs have been shown to be rapidly mixing.

In an early work on mixing times of random walks on random graphs [56], Hildebrand studied simple random walks on random regular graphs with degrees $d = (\log n)^a$ for some $a \geq 2$ and on $\text{ER}_n(p)$ with $p = d/n$, and showed that the mixing time is $\log n / \log d$ in both cases. Later, Benjamini and Mossel [19] studied simple random walk on the part of the infinite cluster of supercritical percolation on $\mathbb{Z}^d$ within the box $[-n, n]^d$. Using the average conductance method developed by Lovász and Kannan [66], they showed that the mixing time is of order n^2. In [51], Fountoulakis

and Reed improved the conductance method of Lovász and Kannan. Using improved estimates, they showed that the mixing time of simple random walk on the giant component of the supercritical Erdős-Rényi random graph is of order $\log^2 n$ [50].

Benjamini, Kozma and Wormald [18] obtained the same result on the mixing time of simple random walk on the giant component by using different techniques, namely, via expansion properties of the giant component. Nachmias and Peres [75] showed that the mixing time of simple random walk on the largest component of $\mathrm{ER}_n(p)$ is of order n within the critical window, i.e., when $p = (1 + \lambda n^{-1/3})/n$ with $\lambda \in \mathbb{R}$. Ding, Lubetzky and Peres [37] studied the near-critical case, where $p = (1+\varepsilon)/n$ with $n^{-1/3} \ll \varepsilon \ll 1$, and showed that the mixing time is of order $\varepsilon^{-3} \log^2(\varepsilon^3 n)$, which interpolates between the critical and the supercritical case.

In 2010, Lubetzky and Sly [67] established cutoff for simple random walk and nonbacktracking random walk on d-regular graphs with $d \geq 3$, and derived precise asymptotics of the total variation distance, via path counting arguments. Nonbacktracking random walk is the same as simple random walk, except that it cannot traverse the same edge twice in a row.

Theorem (Lubetzky and Sly [67])**.** *Let G be a uniformly random d-regular graph on n vertices with $d \geq 3$. Then, for all $\varepsilon \in (0,1)$, with high probability:*

(a) *the mixing time of nonbacktracking random walk on G satisfies*

$$t_{\mathrm{mix}}^{(n)}(\varepsilon) = \log_{d-1}(dn) + O_{\mathbb{P}}(1).$$

(b) *the mixing time of simple random walk on G satisfies*

$$t_{\mathrm{mix}}^{(n)}(\varepsilon) = \frac{d}{d-2} \log_{d-1} n + (\Lambda + o_{\mathbb{P}}(1))\Phi^{-1}(\varepsilon)\sqrt{\log_{d-1} n},$$

where $\Lambda = \frac{2\sqrt{d(d-1)}}{(d-2)^{3/2}}$ *and* $\Phi(x) = \frac{1}{2\pi}\int_x^\infty \exp(-u^2/2)du,\ x \in \mathbb{R}$.

Here, *with high probability* refers to the randomness of the underlying graph and means that the asymptotics holds with a probability that tends to 1 as n tends to infinity. A crucial ingredient in the proof in [67] is the locally tree-like structure of random regular graphs.

Extensions to the configuration model with non-regular degrees were later obtained by Berestycki, Lubetzky, Peres and Sly [21] (for simple and nonbacktracking random walks), and by Ben-Hamou and Salez [16] (for nonbacktracking random walks). In [21], the authors study the mixing time of random walks starting from a *typical* position, instead of a worst-case position. As in [67], they use the locally tree-like structure of the configuration model, but the presence of inhomogeneities in the degrees of vertices requires a more sophisticated approach. In [16], the authors obtain finer asymptotics for the mixing time of nonbacktracking random walk compared to [21], under less restrictive conditions, including the exact order of the cutoff window. The results of [16] are of special interest to us, since we use them in our works on random walks on dynamic random graphs that are presented in Chapters 3 and 4.

§1.1.3 Random walks on dynamic random graphs

The study of random walks on dynamic random graphs involves random walks on graphs that change over time in a random fashion. While this line of research can be considered to be part of the more general framework of random walks in dynamic random environments, the problems and techniques considered are very different. The typical setting in the study of random walks in dynamic random environments involves random walks on $\mathbb{Z}^d$ with the jump probabilities of the random walk changing over time in a random fashion. For a review, see [11, 23] and references therein. Another line of research focuses on random walks on general graphs in which edge or vertex weights change over time, rather than the graph structure itself, such as reinforced random walks (see [63, 81]). Our interest in this thesis lies mainly in models of random walks on graphs whose structure changes over time. For examples of such models see [9, 30, 49, 57, 86].

Mixing times of random walks on dynamic random graphs have been studied only very recently. In [83], Peres, Stauffer and Steif studied the mixing time of random walk on dynamical percolation on $\mathbb{Z}^d$ in the subcritical regime. In the dynamical percolation model, which was introduced by Häggström, Peres and Steif [80], each edge of $\mathbb{Z}^d$ gets refreshed independently at random times, given by exponential clocks, and when refreshed an edge becomes open with probability p and closed with probability $1-p$ independently of the state it had before. Peres, Stauffer and Steif considered the case $p < p_c(\mathbb{Z}^d)$, where $p_c(\mathbb{Z}^d)$ is the critical probability for bond percolation on $\mathbb{Z}^d$, and looked at the problem on the discrete torus with side length n and on the lattice $\mathbb{Z}^d$. By using a regeneration time argument, they estimated the mixing time for the joint Markov chain of the random walk and dynamical percolation.

More recently, Peres, Sousi and Steif [82] considered the same model in the supercritical regime $p > p_c(\mathbb{Z}^d)$. Using the evolving sets method developed by Morris and Peres [74], they obtain an upper bound for the mixing time of the random walk on a quenched realisation of dynamical percolation with the additional constraint that $\theta(p) > 1/2$, where $\theta(p)$ is the probability that the origin belongs to the infinite cluster in bond percolation on $\mathbb{Z}^d$ with probability p. In [89], Sousi and Thomas considered random walk on a dynamical Erdős-Rényi graph in the supercritical regime $p = \lambda/n$ with $\lambda > 1$. The dynamical Erdős-Rényi graph is the same as dynamical percolation, except that the underlying graph is the complete graph on n vertices K_n, instead of $\mathbb{Z}^d$. They showed that the joint Markov chain exhibits cutoff, by showing that the random walk component mixes much faster than the dynamic graph component. The proof was based on a regeneration time argument. More sophisticated variations of the dynamical Erdős-Rényi graph, without the random walk, was also studied very recently in different contexts. In [101], the authors study a variation of the dynamical Erdős-Rényi graph, in which the edge refresh rates also randomly change over time, in the context of complex networks. In [70], the authors study a similar but more complicated model. In particular, they consider the evolution of the number of edges, with explicit results for the corresponding moments, functional central limit theorems and large deviations asymptotics.

The main subject of this thesis is random walks on dynamic random graphs in

which the degrees do not change over time, but the edges are *rewired* randomly. Such dynamic random graph models, without the random walk, were studied earlier in the context of approximate uniform sampling of graphs with given degree sequences. As we have pointed out earlier, one method for sampling graphs is the Markov Chain Monte Carlo (MCMC) method. In MCMC, an ergodic Markov chain, whose stationary distribution is the target sampling distribution, is constructed and run for a sufficiently long time, i.e., longer than the mixing time, so that the resulting sample will be approximately distributed as the target distribution. In order to decide how long MCMC the must run, we need good estimates of mixing times.

The study of MCMC methods and their mixing times in the context of sampling graphs with given degree sequences goes back to the 1990's. In [60], Jerrum and Sinclair gave an algorithm for sampling regular graphs that is based on a Markov chain whose mixing time is polynomial in the number of vertices. Their algorithm works for a large class of degree sequences, in particular, for regular graphs. In [61], Kannan, Tetali and Vempala studied a Markov chain for sampling bipartite graphs and tournaments. Their Markov chain involved a *switch* mechanism, in which endpoints of two randomly chosen edges are switched, provided that the switch results in a simple graph. Using the canonical paths method of Jerrum and Sinclair [88], they showed that the mixing time is polynomial in the number of vertices.

In [34], Cooper, Dyer and Greenhill adapted the switch chain of Kannan et al. to the case of random regular graphs and gave an upper bound for the mixing time that is polynomial in the number of vertices. In [54], Greenhill studied the switch chain for non-regular degree sequences and, in [53], Greenhill and Sfragara studied the switch chain for non-regular simple graphs and directed graphs. They obtained upper bounds for the mixing time via the multicommodity flow argument of Sinclair [87]. In [69], Mahlmann and Schindelhauer introduced a variant of the switch chain, which they called *flip* chain, for sampling regular graphs. Their algorithm is more *local*, in the sense that the random choice of switching edges is made locally, namely, they have to be incident to the same edge. In [46], Feder, Guetz, Mihail and Saberi obtained an upper bound for the mixing time of the flip chain on regular graphs by using a Markov chain comparison argument.

Our contribution

The three chapters of Part I of this thesis are based on our three papers on the mixing time of random walks on dynamic random graphs. In Chapter 2, we present our work on the mixing time of random walks on dynamic configuration models. The dynamic configuration model is a dynamic version the configuration model, in which dynamic random graphs with a fixed degree sequences are generated. A fraction α_n of the edges is rewired at each unit of time, where n is the number of vertices. This model was introduced in [12], which forms the basis of Chapter 2. In Chapter 2, we study the mixing time of the random walk without backtracking on the dynamic configuration model in the *supercritical* regime. By supercritical we mean $\lim_{n\to\infty} \alpha_n(\log n)^2 = \infty$. We show that, under some regularity conditions on the degree sequence, the mixing time is of order $\alpha_n^{-1/2}$, which is of order $o(\log n)$ and hence is much smaller than

the mixing time of the random walk on the static configuration model. To prove our main result, we use a randomised stopping time argument. We consider the first time τ at which the random walk crosses an edge that was rewired before, and show that τ behaves like a strong stationary time. The key part of the proof involves an exploitation of the tree-like structure of the configuration model.

In Chapter 3 of the thesis, we present our work on the extension of the results of [12] to the *critical* and the *subcritical* regimes, which correspond to $\lim_{n\to\infty} \alpha_n(\log n)^2 \in (0,\infty)$ and $\lim_{n\to\infty} \alpha_n(\log n)^2 = 0$, respectively. The mixing time was analysed in [13], which forms the basis of Chapter 3. Together with the main result of [12], we see that there is *trichotomy* for the mixing time of random walks on dynamic configuration models:

- $\lim_{n\to\infty} \alpha_n(\log n)^2 = \infty$: the mixing time is of order $\alpha_n^{-1/2}$ and there is *no cutoff*,
- $\lim_{n\to\infty} \alpha_n(\log n)^2 \in (0,\infty)$: the mixing time is of order $\log n$ and there is *one-sided cutoff*,
- $\lim_{n\to\infty} \alpha_n(\log n)^2 = 0$: the mixing time is of order $\log n$ and there is *two-sided cutoff*.

The latter regime includes the random walk on the static configuration model.

We again use the randomised stopping time argument from [12]. However, because the distances over which the random walk can travel are now of order $\log n$, we cannot rely on the tree-like structure of the configuration model only. Instead, we show that the random walk path is with high probability self-avoiding, which in fact is a consequence of the locally tree-like structure of the configuration model, and this in turn ensures that the randomised stopping time τ behaves like a strong stationary time.

In Chapter 4 of the thesis, we present our work on the mixing time of random walks on more general dynamically rewired random graphs. We consider a more general setting where the vertices and vertex degrees are fixed as before, but the edges are rewired according to a more general prescribed rule. This setting includes the dynamic configuration model as a special case. We show that, under some conditions on the graph dynamics, the total variation distance for the random walk on the dynamically rewired random graph can be linked to the total variation distance for the random walk on the static configuration model. We also introduce a specific model that we call *the random walk with local rewirings*. In this model, the rewiring occurs only *along* the random walk path, with probability α_n for each time unit. Using our result for the general framework, we show that this random walk exhibits the same trichotomy as the random walk on the dynamic configuration model, but on a different time scale. More precisely, we identify three regimes:

- $\lim_{n\to\infty} \alpha_n \log n = \infty$: the mixing time is of order α_n^{-1} and there is *no cutoff*,
- $\lim_{n\to\infty} \alpha_n \log n \in (0,\infty)$: the mixing time is of order $\log n$ and there is *one-sided cutoff*,

- $\lim_{n\to\infty} \alpha_n \log n = 0$: the mixing time is of order $\log n$ and there is two-sided cutoff.

Again, the latter regime includes the random walk on the static configuration model.

To prove the above result, we show that the random walk on the dynamically rewired random graph can be coupled to a *modified* random walk on the static configuration model. While this coupling was implicit in the proof used in [13], we show that the same argument can be used in a more general setting. We show that the total variation distance for the modified random walk can be expressed in terms of the tail probability of the randomised stopping time τ (the first time the random walk crosses an edge that was rewired before) and the total variation distance for the random walk on the static configuration model.

§1.2 Introduction to Part II

In Part II of the thesis, we study the problem of *union complexity* of random disk regions. We have a collection of disks, whose centers are distributed randomly within a region of the plane, and we are interested in the expected number of boundary arcs. Formally, let $\mathcal{D} = \{D_1, \ldots, D_n\}$ be a set of n disks in $\mathbb{R}^2$. The set of boundary disks of $\mathcal{D}$, denoted by $\mathsf{BD}(\mathcal{D})$, is the set of disks in $\mathcal{D}$ whose boundaries are not completely covered by other disks, i.e.,

$$\mathsf{BD}(\mathcal{D}) = \{D \in \mathcal{D} : \partial D \setminus \cup_{D' \in \mathcal{D} \setminus \{D\}} D' \neq \emptyset\},$$

where ∂D denotes the boundary of D. A boundary arc of a boundary disk D is a connected component of the set $\partial D \setminus \cup_{D' \in \mathcal{D} \setminus \{D\}} D'$ and union complexity of $\mathcal{D}$ is the total number of boundary arcs of the boundary disks of $\mathcal{D}$. It was shown that, when $\mathcal{D}$ consists of disks, the total number of boundary arcs of $\mathcal{D}$ is bounded from above by $6|\mathsf{BD}(\mathcal{D})| - 12$ [62]. Hence, we study the union complexity by focusing on the number of boundary disks and we state our results in terms of the number of boundary disks.

The union complexity of geometric regions is important in several combinatorial and algorithmic problems in different fields ranging from linear programming and robotics to molecular modeling and geographic information systems [1]. A special case of the union-complexity problem is motivated by an algorithm to compute a *conflict-free colouring* for unit disks. Besides, the union-complexity problem is similar to the problem of the complexity of the convex hull, when the number of vertices of the convex hull of a random point set is studied. In Section 1.2.1, we motivate the union complexity problem in the context of conflict-free colouring of unit disk regions. In Section 1.2.2, we introduce the convex hull problem, review the relevant literature, and show its relation to the union complexity problem. In Section 1.2.3, we briefly report our contribution.

§1.2.1 Conflict-free colouring

The problem of conflict-free colouring was introduced by Even et al. [45] for simple geometric regions . In the general setting, we are given a set $\mathcal{D}$ of regions in the plane. The *coverage area* of $\mathcal{D}$ is defined as $\mathsf{Cov}(\mathcal{D}) := \cup_{D \in \mathcal{D}} D$. For a point $p \in \mathsf{Cov}(\mathcal{D})$, we let $\mathcal{D}(p) := \{D \in \mathcal{D} : p \in D\}$, i.e., the set of regions that contain p. A colouring of $\mathcal{D}$ is a function $\chi : \mathcal{D} \to \mathbb{N}$. A conflict-free (CF) colouring is defined as follows:

Definition 1.2.1. *A conflict-free colouring of $\mathcal{D}$ is a colouring χ of $\mathcal{D}$ such that for every $p \in \mathsf{Cov}(\mathcal{D})$, there exists a $D \in \mathcal{D}(p)$ with $\chi(D) \neq \chi(D')$ for any other $D' \in \mathcal{D}(p)$, i.e., for each point in the coverage area there is a region with a* unique *colour among the regions containing that point. A minimum CF colouring of $\mathcal{D}$ is a CF colouring of $\mathcal{D}$ that uses a minimum number of colours among all possible CF colourings of $\mathcal{D}$ (see Figure 1.1).*

Even et al. [45] show that when $\mathcal{D}$ consists of n unit disks (disks with unit radius), it is always possible to achieve a conflict-free colouring with $O(\log n)$ colours by giving

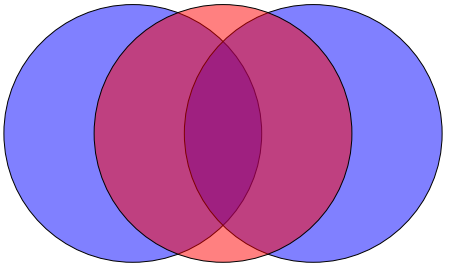

Figure 1.1: A minimum conflict-free colouring of three disks. Two outer disks are coloured blue and the middle disk is coloured red.

an algorithm that uses $O(\log n)$ colours for any set of unit disks in the plane. They also show that when the centres of the disks lie on a straight line and every pair of disks intersects, the minimum number of colours required is $\Omega(\log n)$. So, in the worst-case scenario, minimum CF colourings of unit disks in the plane use $\Theta(\log n)$ colours. In Chapter 5, we consider the average-case scenario for disks in the plane. We take the average over all possible configurations of disks whose centres are uniformly distributed in a convex compact region.

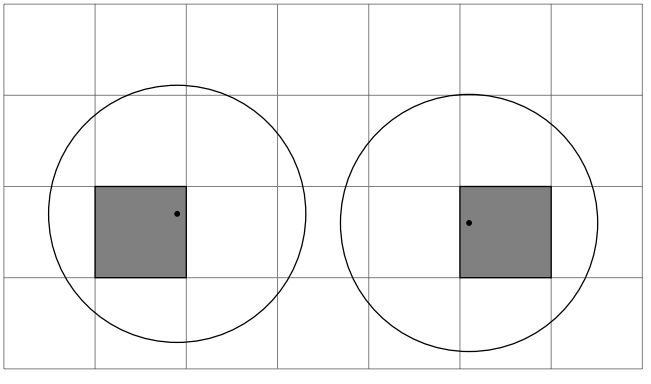

Figure 1.2: Disks whose centres lie at least 3 squares apart do not intersect.

Lev-Tov and Peleg [64] give a constant-factor approximation algorithm for the minimum CF colouring of a set of unit disks in the plane, i.e., an algorithm that achieves a solution that uses, for any given input, at most a constant times more colours than the minimum number of colours needed for that input. The algorithm proceeds by dividing the plane into squares and treating the disks whose centers lie in each square separately. The diagonals of the squares are of length 1, so if a square contains a disk center, then it is completely covered by that disk. Another useful property of the algorithm is that two disks whose centres lie in squares that are 3 squares apart do not intersect (see Figure 1.2). Hence the disks of every fourth square in a row and in a column can be coloured by using the same colour set.

The algorithm given by Lev-Tov and Peleg [64] relies on CF colourings of *chains* of regions. A chain is a set of regions that can be ordered in a natural way. The formal definition is as follows:

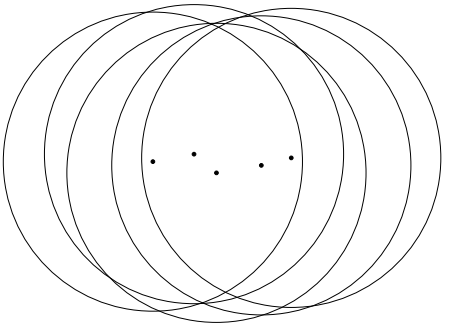

Figure 1.3: A chain of five disks whose centers are marked by crosses.

Definition 1.2.2. *A set $\mathcal{D}$ of n regions is called a* chain *if there exists an indexing of the regions $\mathcal{D} = \{D_1, \ldots, D_n\}$ such that, for every pair $i, j \in [n]$ with $i < j$, the set $\cap_{k\in[i,j]} D_k \setminus \cup_{k\notin[i,j]} D_k$ is non-empty and for every $p \in \mathsf{Cov}(\mathcal{D})$ there exist $i, j \in [n]$ with $i < j$ such that $p \in \cap_{k\in[i,j]} D_k \setminus \cup_{k\notin[i,j]} D_k$ (see Figure 1.3).*

A collection of disks whose centres lie on a line and for which every pair of disks intersect is an example of a chain. It is proven by Even et al. [45] that, for a chain of length n, minimum CF colourings always use $\Theta(\log n)$ colours.

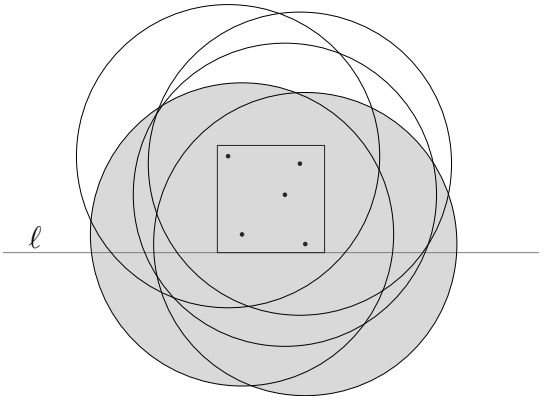

Figure 1.4: Lower boundary diks of five disks highlighted in grap.

In the algorithm given by Lev-Tov and Peleg [64], CF colourings of disks whose centers lie in a square involves CF colourings of chains. For a square Q, let $\mathcal{D}^Q$ denote the set of disks whose centres lie in Q. Let ℓ be the line passing through the lower side of Q. Let B be the intersection of the half plane below ℓ and $\mathsf{Cov}(\mathcal{D}^Q)$, and let $\mathcal{D}^Q_{\text{lower}}$ be the set of disk whose bounding arcs appear on the boundary of B. Disks in $\mathcal{D}^Q_{\text{lower}}$ are called *lower boundary disks* (see Figure 1.4). Upper, left and right boundary segments are defined similarly, and they are denoted by $\mathcal{D}^Q_{\text{upper}}$, $\mathcal{D}^Q_{\text{left}}$, $\mathcal{D}^Q_{\text{right}}$, respectively. A crucial observation is that each of these sets forms a chain (see [64, Lemma 2.5]) and their union covers $\mathsf{Cov}(\mathcal{D}^Q)$ entirely (since the diagonal of the square is 1). So the problem of minimum CF colourings of $\mathcal{D}^Q$ reduces to the

problem of minimum CF colourings of $\mathcal{D}^Q_{\text{bound}} = \mathcal{D}^Q_{\text{lower}} \cap \mathcal{D}^Q_{\text{upper}} \cap \mathcal{D}^Q_{\text{left}} \cap \mathcal{D}^Q_{\text{right}}$, the set of boundary disks. Thus, it is important to estimate the average number of boundary disks for the average-case analysis of the algorithm.

§1.2.2 Complexity of the convex hull of random point sets

Let $X = \{X_1, \ldots, X_n\}$ be a set of n independently sampled random points in the plane. The convex hull of X, denoted as $\mathsf{CH}(X)$, is the smallest convex set that contains X, which is also the set of all possible convex combinations of points of X, i.e.,

$$\mathsf{CH}(X) = \{x \in \mathbb{R}^2 : x = \sum_{i=1}^{n} \alpha_i X_i \text{ with } \sum_{i=1}^{n} \alpha_i = 1 \text{ and } \alpha_i \geq 0 \text{ for all } i = 1, \ldots, n\}.$$

Let $V(X)$ be the set of vertices of $\mathsf{CH}(X)$, i.e., the set of points of X that are at the boundary of its convex hull (note that, contrary to standard terminology, we also call a point from X a vertex when it lies in the interior of an edge of $\mathsf{CH}(X)$). The study of convex hulls of random point sets goes back to 1960's. Rényi and Sulanke [85] showed that, as n tends to infinity, the expectation of $|V(X)|$, where X is a set of n randomly sampled points, scales like $\sqrt{\log n}$ if the points are sampled according to the normal distribution, like $\log n$ if the points are sampled uniformly in a polygon, and like $n^{1/3}$ if the points are sampled uniformly in a convex compact region with a smooth boundary. Later studies focused on exact asymptotics for the expectation of $|V(X)|$ and its variance, on limit theorems and on generalization to higher dimensions. For a survey of results on convex hull problems, see [99].

The usual proof strategy involves computing the probability of two fixed points forming an edge of the convex hull and multiplying this probability by $\binom{n}{2}$. This gives the expected number of edges of the convex hull, which is the same as the number of vertices. Two points of X form an edge if all the other points fall on the same side of the line passing through these two points. Efron [39] devised a different approach that exploits the relation between the area of the convex hull and the number of vertices. Using this approach, Har-Peled [55] presented a simple proof for the asymptotic scaling of $|V(X)|$. His arguments can be viewed as a discretization of the classical arguments of Rényi and Sulanke in [85].

The convex-hull problem is directly related to the union-complexity problem. Let $\mathcal{D}(X, r) = \{D_1, \ldots, D_n\}$, where D_i is the disk with radius r centered at X_i. Let us use the shorthand notation $\mathsf{BD}(X, r) = \mathsf{BD}(\mathcal{D}(X, r))$ for the set of boundary disks of $\mathcal{D}(X, r)$. Although we do not need the following proposition in the proofs of Theorem 5.1.1 below, we state the connection between the convex hull and the boundary disks.

Proposition 1.2.3. *Suppose that $X = \{X_1, \ldots, X_n\}$ is a set of n points in $\mathbb{R}^2$. Then, for $i = 1, \ldots, n$, the point X_i is a vertex of the convex hull of X, i.e. $X_i \in V(X)$, if and only if for any $r > 0$ the disk with radius r centered at X_i is a boundary disk of $\mathcal{D}(X, r)$, i.e. $D_i \in \mathsf{BD}(X, r)$.*

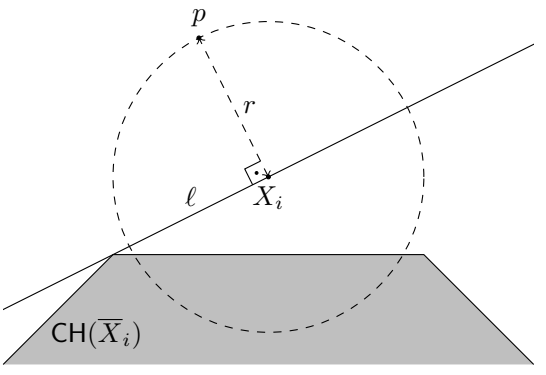

Figure 1.5: Illustration of the first part of the proof of Proposition 1.2.3.

Proof. Fix $i \in [n]$. First assume that $X_i \in V(X)$ and fix $r > 0$. Let ℓ be a line through X_i that is tangent to $\mathsf{CH}(\overline{X}_i)$ where $\overline{X}_i \coloneqq X \setminus \{X_i\}$. Let h^+ denote the closed half-plane bounded by ℓ such that $\mathsf{CH}(X) \subset h^+(\ell)$, and let h^- be the opposite half-plane. Finally, let $p \in h^-$ be the point at distance r from X_i such that the segment pX_i is perpendicular to ℓ (see Figure 1.5). Clearly, $p \in \partial D_i$ and all other points in X have distance greater than r to p. Hence, D_i is a boundary disk.

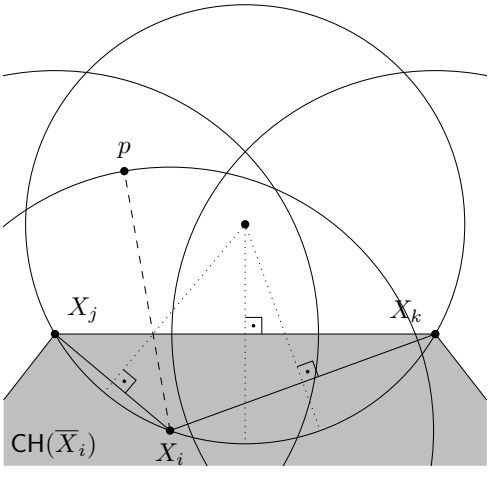

Figure 1.6: Illustration of the second part of the proof of Proposition 1.2.3. Dotted lines are perpendicular bisectors of the triangle $X_iX_jX_k$.

Next assume that $X_i \notin V(X)$. Then X_i is in the interior of $\mathsf{CH}(\overline{X}_i)$. For distinct $j, k \in [n]$ with $j, k \neq i$, let $r_{j,k}$ denote the radius of the circumscribed circle of the points X_i, X_j and X_k. Set $r = \max_{j,k\in[n]} r_{j,k} + \mathsf{diam}(X)$, where $\mathsf{diam}(X)$ is the diameter of X. Then the boundary of the disk D_i with radius r centered at X_i is not a boundary disk of $\mathcal{D}(X, r)$. Indeed, fix any point on the boundary of D_i, say p. If $p \in \mathsf{CH}(X)$, then p is covered by all the other disks, since $r > \mathsf{diam}(X)$. If $p \notin \mathsf{CH}(X)$, then consider the line segment between the points p and X_i. Suppose that this line

segment passes through the edge X_jX_k of the convex hull for some $j, k \in [n]$. Consider the perpendicular bisector ℓ_1 of the line segment X_iX_j and the perpendicular bisector ℓ_2 of the line segment X_iX_k. Assume without loss of generality that X_i lies to the right of ℓ_1 and to the left of ℓ_2. Since $r > r_{j,k}$, p lies to the left of ℓ_1 or to the right of ℓ_2. Then, $\min\{d(p, X_j), d(p, X_k)\} < d(p, X_i) = r$ and at least one of the disks D_j and D_k contains p (see Figure 1.6). So any point on the boundary of D_i is covered by some other disk in $\mathcal{D}(X, r)$, and hence D_i is not a boundary disk of $\mathcal{D}(X, r)$.

□

§1.2.3 Our contribution

In Chapter 5, we study the average-case union complexity for disks whose centers are sampled uniformly and independently at random in a region S in $\mathbb{R}^2$ and whose radii are greater than the diameter of S. We obtain union-complexity results by bounding the number of boundary disks. We consider two cases: the case where S is a square and the case where S is a disk. We show that, in both cases, the expected number of boundary disks scales like $n^{1/3}$ as n, the number of random disks, tends to infinity. Our proof is an adaptation of the proof of Har-Peled [55] for the convex hull problem.

Our result for the square case is important in the context of the average-case analysis of Lev-Tov and Peleg's algorithm [64] for the conflict-free colouring of disk regions [64], since their colouring scheme relies on the colouring of the boundary disks. Considering the connection between the union complexity and the convex-hull complexity, our results are surprising. For the convex-hull problem, the order of the expected number of vertices of the convex hull is different for the two cases: it is $\log n$ for the square case and $n^{1/3}$ for the disk case. However, for the union-complexity problem, the order of the number of boundary disks is $n^{1/3}$ in both cases.