Structural and Algorithmic Properties of Static and Mobile Random Geometric Graphs

Alexandre Stauffer

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2011-39
http://www.eecs.berkeley.edu/Pubs/TechRpts/2011/EECS-2011-39.html

May 5, 2011

Structural and Algorithmic Properties of Static and Mobile Random Geometric Graphs

by

Alexandre de Oliveira Stauffer

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Computer Science
and the Designated Emphasis

in

Communication, Computation and Statistics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Alistair Sinclair, Chair
Professor Richard M. Karp
Professor David Aldous

Spring 2011

Structural and Algorithmic Properties of Static and Mobile Random Geometric Graphs

Abstract

Structural and Algorithmic Properties of Static and Mobile Random Geometric Graphs

by

Alexandre de Oliveira Stauffer

Doctor of Philosophy in Computer Science

and the Designated Emphasis in Communication, Computation and Statistics

University of California, Berkeley

Professor Alistair Sinclair, Chair

We study fundamental problems for static and mobile networks. First, we consider the random geometric graph model, which is a well-known model for static wireless networks. In this model, n nodes are distributed independently and uniformly at random in the d-dimensional torus of volume n and edges are added between pairs of nodes whose Euclidean distance is at most some parameter r. We consider the case where r is a sufficiently large constant so that a so-called giant component (a connected component with $\Theta(n)$ nodes) exists with high probability. In this setting, we show that the graph distance between every pair of nodes whose Euclidean distance is sufficiently large is only a constant factor larger than their Euclidean distance. This result gives, as a corollary, that the diameter of the giant component is $\Theta(n^{1/d}/r)$. Then, we apply this result to analyze the performance of a broadcast algorithm known as the push algorithm. In this algorithm, at each discrete time step, each informed node chooses a neighbor independently and uniformly at random and informs it. We show that the push algorithm informs all nodes of the giant component of a random geometric graph within a number of steps that is only a constant factor larger then the diameter of the giant component.

In the second part of the thesis, we consider a model of mobile graphs that we call mobile geometric graphs, and which is an extension of the random geometric graph model to the setting where nodes are not static but are moving in space in continuous time. In this model, we start with a random geometric graph and let the nodes move as independent Brownian motions. Then, at any given time, there exists an edge between every pair of nodes whose Euclidean distance at that time is at most r. This model has been recently used as a model for mobile wireless networks. We study four fundamental problems in this model: detection (the time until a target point—fixed or moving—is within distance r of some node of the graph); coverage (the time until all points inside a finite box are detected by the graph); percolation (the time until a given node belongs to the giant component of the graph) and broadcast (the time until all nodes of the graph receive a piece of information that was

initially known by a single node). We obtain precise asymptotics for these quantities by combining ideas from stochastic geometry, coupling and multi-scale analysis.

Finally, in the last part of the thesis, we revisit the push algorithm described above and study its performance in general regular graphs. Our goal is to understand the relation between the performance of the push algorithm and the vertex expansion of the graph. We prove an upper bound for the runtime of this algorithm that depends on the vertex expansion of the graph and is tight up to polylogarithmic factors. Then, we show that there exists a substantial difference between the impact of vertex expansion and edge expansion on the performance of the push algorithm. In particular, we prove the existence of regular graphs (which are also vertex transitive) that have constant vertex expansion and for which the runtime of the push algorithm is a factor of $\Omega(\log n)$ slower than on any regular graph with constant edge expansion.

To my parents, Edna and Zediel,
to whom I owe everything I am.

Contents

Acknowledgments

First and foremost, I am very fortunate to have had Alistair Sinclair as my thesis advisor. I am thankful to Alistair for seeing mobile networks as a promising line of research and for guiding me during these years while, at the same time, giving me enough space to develop my own ideas. I am also thankful for all the time Alistair devoted to our discussions, his enthusiasm and care.

I cannot think of a better place to do theory than Berkeley. I am thankful to the members of the theory and probability groups, in particular Lorenzo Orecchia and Anindya De, who were both great colleagues and amazing friends. I am thankful to David Aldous, Dick Karp and Satish Rao for being in my qualifying exam and thesis committee, and to Elchanan Mossel for discussions during my first year. Also, I want to thank David Tse and Sylvia Ratnasamy for eye-opening discussions about practical aspects of mobile wireless networks.

I was extremely lucky to spend three fantastic months as a summer intern at Microsoft Research. I am grateful to Yuval Peres for accepting me for the internship, for his enthusiasm toward my research and for always being available to clarify my simplest doubts. I am also grateful to Itai Benjamini for lively and inspiring discussions and for inviting me to visit the Weizmann Institute of Science. At Microsoft, I had the opportunity to interact with many amazing people, and want to give special thanks to Perla Sousi for enjoyable moments and helpful discussions, and for her dedication to our joint work.

I was fortunate to collaborate with amazing researchers outside Berkeley. I am grateful to Jose Blanchet for nice visits to Harvard University and Columbia University, and for his interest and dedication to our paper. I am thankful to Thomas Sauerwald for our many inspiring discussions, and to Tobias Friedrich for attending my seminar at UC Berkeley, which became the first step of a fruitful collaboration. Also, I thank both Thomas and Tobias for having invited me to visit the Max-Planck Institut für Informatik.

My life in Berkeley would have been much harder without the support of many friends. I am thankful to Humberto Gonzalez for being an amazing housemate and great friend, to Stéphane Martinez for our amazing trips and chats, to Brenno Kenji for our weekly German meetings, to Decha Sermwittayawong for convincing me to play in his dodgeball team, to Paul Bogdan and Cristina Albu for our unforgettable road trips, to Pedro Cardoso for becoming my Google friend and to Vinicius Ramos for insisting in calling me as his advisor. I am also thankful to Eugenia Alexeev for the delightful Russian classes.

I want to thank Kate Leiva from IIE and La Shana from UC Berkeley for being highly efficient in responding my questions and handling my administrative problems. Also, I want to thank Valmir Barbosa, whose invaluable help allowed me to be admitted to Berkeley.

Above all, I have been extremely blessed to count on the unconditional love of my parents, Edna and Zediel, and my sister and brother-in-law, Corina and Wandemberg. Their unfailing support has been invaluable for my development, and their immaculate example has been serving as an inspiration to my life.

Chapter 1

Introduction

Wireless networks have been the source of much research within both the practical and theoretical communities in computer science, as well as other areas of research such as applied probability and electrical engineering. From a theoretical perspective, problems related to wireless networks have not only proved to be interesting in their own right, but have fostered the development of innovative proof techniques inspired by results on Random Graphs [71], Stochastic Geometry [83], and Percolation Theory [48].

In this thesis, we will be mostly concerned with the study of random graph models for static and mobile wireless networks. In 1959, Erdős and Rényi [37] and Gilbert [45] introduced two models for random graphs, which are now regarded as the foundations of random graph theory. In the model of Erdős and Rényi [37], for two fixed positive integers n and m, we construct a graph with n nodes and m edges by choosing the m edges uniformly at random from the set of all $\binom{n}{2}$ possible pairs of nodes. In the model of Gilbert [45], given a parameter $p \in (0, 1)$ and a positive integer n, we construct a random graph with n nodes by taking each of the $\binom{n}{2}$ pairs of nodes and adding an edge between them with probability p independently of all the other edges.

These models have been extensively studied and many of their properties are by now widely understood; see for example the books by Bollobás [13] and Janson, Łuczak and Ruciński [55]. However, as was pointed out by Gilbert as early as 1961 [46], these two models do not seem to capture very well the properties of some real-world networks, such as many types of wireless networks. For this reason, Gilbert [46] introduced a different model of random graphs, using as a motivating scenario the case where the nodes of the graph are stations with the capability of transmitting messages over short distances through the air. In this type of network, signals transmitted through the air have their amplitude attenuated as they move away from the transmitting node; thus, a pair of nodes can only exchange messages if they are sufficiently close to each other.

We now briefly introduce the model defined by Gilbert [46], which we refer to as the *random geometric graph model*, but which is also often called the *boolean model*. We let $\mathbb{S}$ be the torus of volume n in d-dimensional Euclidean space, and take the nodes of the graph to

be distributed as a Poisson point process with intensity λ over $\mathbb{S}$; i.e., the number of nodes in any bounded subregion $S \subseteq \mathbb{S}$ is a Poisson random variable with mean $\lambda \operatorname{vol}(S)$, where $\operatorname{vol}(S)$ is the volume of S. Note that the parameter $\lambda > 0$ governs the density of nodes per unit volume. Then, an edge is created between every pair of nodes for which their Euclidean distance is smaller than a certain parameter r, referred to as the *transmission range*. It follows by a standard scaling argument that the structure of the random geometric graph depends only on the product $\lambda \operatorname{vol}(B(0,r))$ (where $B(x,r)$ is the ball of radius r centered at $x \in \mathbb{R}^d$) [14], so we may fix r and parameterize the model on λ only. Asymptotic properties of random geometric graphs have been studied for n going to infinity while λ may or may not depend on n. For extensive background on this topic, we refer the reader to the books by Meester and Roy [65] and Penrose [71].

Instead of having a random number of nodes, we can consider the model where exactly λn nodes are placed independently and uniformly at random in $\mathbb{S}$. Not surprisingly, results that hold with high probability for random geometric graphs can be translated to the model with a fixed number of nodes using a technique called de-Poissonization [71]. We remark that Gilbert, in his original paper [46], considered only the case $\mathbb{S} = \mathbb{R}^d$, which is the infinite-volume limit as $n \to \infty$. Most of our analyses will consider the case $\mathbb{S} = \mathbb{R}^d$, which is usually mathematically cleaner. However, since many existing results on random geometric graphs depend on n and cannot be described in the case $\mathbb{S} = \mathbb{R}^d$, in this section we will consider the more general setting where $\mathbb{S}$ is the torus of volume n, and will explicitly use the term *infinite* random geometric graph to refer to the case $\mathbb{S} = \mathbb{R}^d$.

Clearly, increasing λ increases the average degree of the nodes. As is well known, there are two critical values of λ at which the connectivity properties of the random geometric graph undergo a significant change. First there is the *percolation threshold* $\lambda = \lambda_c$ (a constant that depends on the dimension d), so that if $\lambda > \lambda_c$ the network w.h.p.[1] has a unique "giant" component containing a constant fraction of the nodes, while if $\lambda < \lambda_c$ all components have size $O(\log n)$ w.h.p. [71]. Second, at the *connectivity threshold* $\lambda = \frac{\log n}{\operatorname{vol}(B(0,r))}$, the network becomes connected w.h.p. [50]; note that connectivity requires that λ, and hence the average degree, grows with n. The percolation threshold $\lambda = \lambda_c$ occurs also in infinite random geometric graphs, in which case the giant component is the unique infinite component (or "infinite cluster") with probability 1.

There are many theoretical results on routing and other algorithmic questions on random geometric graphs. Naturally most of these results are restricted to *connected* graphs, which requires the unrealistic assumption that the average degree is unbounded. We postpone a detailed discussion of these results to Section 1.2, and mention here only the seminal work of Gupta and Kumar [51, 52] (with refinements by Franceschetti et al. [41]), who examined the information-theoretic *capacity* (or throughput) of random geometric graphs above the connectivity threshold. The capacity is defined as the number of bits per unit time that

[1] We shall take the phrase "w.h.p." ("with high probability") to mean "with probability tending to 1 as $n \to \infty$."

each node u can transmit to some (randomly chosen) destination node t_u in steady state, assuming constant size buffers in the network. Gupta and Kumar showed that the capacity per unit node cannot be larger than $O(n^{-1/2})$ even when nodes are strategically deployed. Since the capacity tends to 0 as $n \to \infty$, this suggests a fundamental limitation on the scalability of such networks. Later, Franceschetti et al. [41] showed that this upper bound on the capacity is tight for random geometric graphs.

A central feature of many ad hoc networks is that the nodes are moving in space. One such case is *vehicular networks*, where wireless devices are installed in buses, taxis, private cars, and other means of transportation and are used to collect information regarding traffic accidents or vehicle breakdowns. This information is then routed among the vehicles in order to reach all the nearby cars or to reach a first response unit such as a police station [70]. Another application of mobile networks arises in the context of disaster recovery. When a major natural disaster, such as a hurricane or earthquake occurs, it may be unsafe to send humans into the field to gather information. One alternative is to launch moving sensors into the field from an aircraft to collect information and transmit it to a rescue unit. As shown in [65], static (non-mobile) nodes can only efficiently cover a region if either they are strategically deployed, which is not realistic when sensors are deployed from an aircraft, or their density or transmission range is unbounded, which is economically infeasible or unrealistic due to the power constraints of the sensors. In this situation, allowing the sensors to move inside the region may substantially improve the efficiency of a rescue operation.

Such mobile networks are frequently modeled using random geometric graphs, augmented by motion of the nodes. Once mobility is injected, the questions of interest naturally change from those in the static case. For example, connectivity no longer plays such a central role because mobility may allow nodes u, v to exchange messages even in the absence of a path between them at any given time: namely, u can route its message to v along a time-dependent path, opportunistically using other nodes to relay the message towards v. Networks of this kind are often termed "delay tolerant networks" [38]. This allows us to focus not on the rather artificial connectivity regime mentioned above (where the average degree grows with n), but instead on the case where both λ and r (and hence the average degree) are constants. This is obviously highly desirable as it makes the model much more realistic and scalable.

There are rather few rigorous results on wireless networks with mobile nodes, and those that do exist typically either make unrealistic assumptions about node mobility (such as unbounded range of motion [49, 28, 22]), or work in the connectivity regime which, as we have seen, requires unbounded density or transmission range [27, 44]. We discuss these and other existing results for mobile graphs in more detail in Section 1.2, and mention here only the result of Grossglauser and Tse [49] which contrasts with the result of Gupta and Kumar [51, 52] for static graphs mentioned above.

Grossglauser and Tse [49] (see also [28]) consider a very simple model of mobility where, at each discrete time step, each node chooses a position in $\mathbb{S}$ independently and uniformly at random and moves there instantaneously. They analyze the following two-hop routing

scheme: if a node u has a message to send to a node v, u first sends it to its current closest neighbor, and then this node waits until v is its closest neighbor before transmitting u's message to v. Grossglauser and Tse proved that this type of mobility model and this routing scheme achieve *constant* capacity, which dramatically improves upon the result mentioned above [51, 41] that the capacity of static networks goes to zero as $n \to \infty$. This result is among the first rigorous demonstrations that mobile networks may achieve substantially better performance than static networks; however, it requires the unrealistic assumption that nodes move a distance comparable to the diameter of the entire region $\mathbb{S}$ at each unit of time.

We conclude this section by discussing the broadcast problem, which is one of the most fundamental problems in networks. In this thesis, we not only analyze the broadcast problem for static and mobile wireless networks, but also take a more theoretical approach to the wider problem of broadcasting information on general graphs.

We consider a very simple algorithm for broadcasting called the *push algorithm*, which is also sometimes referred to as *randomized rumor spreading* or the *random phone call model*. We assume that initially there exists only one informed node, that is, there is a single node with the information to be broadcast. The algorithm runs in discrete time and, at each time step, each informed node chooses a neighbor independently and uniformly at random and informs it (which is referred to as a *push* operation) until all nodes of the network are informed.

The push algorithm has been applied successfully both in the context where a single piece of information has to be distributed from one node to all the others [53], as well as in the setting where information may be injected at various nodes at different times. The latter problem occurs, for example, when maintaining data integrity in distributed databases, such as name servers in large corporate networks [26, 57].

Despite looking naïvely simple, the push algorithm has been shown to be surprisingly efficient for many types of networks. In addition, it has also been shown to be robust against failures [56, 40]. From a more theoretical perspective, the runtime of the push algorithm has been observed to be related to other fundamental processes in computer science, such as the mixing time of Markov chains [15, 34, 77], the cover time of random walks [35] and the expansion of graphs [69, 77, 20, 19].

1.1 Summary of Results

Static Networks

We start by presenting our results for static networks, which will be more thoroughly discussed in Chapter 2. These results are based on a joint work with Milan Bradonjić, Robert Elsässer, Tobias Friedrich and Thomas Sauerwald [16]. We restrict our analysis here to finite random geometric graphs in two dimensions.

Let G be a random geometric graph in two dimensions, where the nodes are given by a Poisson point process with intensity λ over $\mathbb{S}$, the torus of volume n. We let the transmission range r be fixed, and add an edge between every pair of nodes whose Euclidean distance is at most r. We consider the case $\lambda > \lambda_c$, so that a giant component with $\Theta(n)$ nodes exists w.h.p., and denote by G' the graph induced by the giant component of G. In our discussion here, we assume that λ is a constant, but r may grow with n.

We take two nodes u, v of G' and study the graph distance between them; i.e., the length of the shortest path from u to v in G'. Since the transmission range is r, the graph distance between u and v, which we denote by $d_G(u, v)$, is at least $\|u - v\|_2/r$. Our first result (see Theorem 2.1) says that, w.h.p.,

$$d_G(u, v) = O(\|u - v\|_2/r) \quad \text{for all } u, v \text{ in } G' \text{ such that } \|u - v\|_2 = \Omega(\log^6 n/r^4). \tag{1.1}$$

A direct consequence of this result is that the diameter of the giant component of G is $\Theta\left(\sqrt{n}/r\right)$, a result that was previously known only above the connectivity threshold [33]. The proof of (1.1) uses a renormalization argument and ideas from continuum percolation.

We also study the problem of broadcasting information on the giant component of random geometric graphs with $\lambda > \lambda_c$, where initially there exists only one informed node in G'. We study the runtime of the push algorithm over G'. Recall that, in the push algorithm, at each time step, each informed node chooses a neighbor independently and uniformly at random and informs it. The runtime of the push algorithm is clearly lower bounded by the diameter of G', which we denote by $\mathrm{diam}(G')$. In our second main result, we establish a tight upper bound for the runtime of the push algorithm in G'; more formally, we show in Theorem 2.3 that, w.h.p., the push algorithm informs all nodes of G' within $\Theta(\mathrm{diam}(G))$ steps.

Mobile Networks

In our study of mobile networks, we employ the following model, which we refer to as the *mobile geometric graph model*. We begin at time 0 with an infinite random geometric graph G_0 over $\mathbb{S} = \mathbb{R}^d$. Nodes move independently in continuous time according to standard Brownian motion. It is not hard to verify that this produces a stationary sequence of graphs $(G_s)_{s\in\mathbb{R}_+}$; i.e., for any fixed s, the graph G_s is an instance of the random geometric graph model. However, it is crucial to notice that, though G_s and G_0 have the same distribution, they are *not* independent; it is this feature that makes mobility challenging to analyze.

This model is essentially equivalent to the "dynamic boolean model" introduced by van den Berg, Meester and White [12] in the context of dynamic continuum percolation. We choose to work principally in the infinite volume limit, $\mathbb{S} = \mathbb{R}^d$, which is matehmatically cleaner. However, where appropriate, results obtained in this setting can typically be translated to finite volume by working in a torus $\mathbb{S}$ of volume n (see Corollary 3.7 for an example).

The mobile geometric graph model inherits many properties from random geometric graphs. For example, it was proved in [12] that if $\lambda > \lambda_c$, where λ_c is the same critical value of λ discussed above for random geometric graphs, then G_s contains a giant component

at all times s almost surely. Similarly, if $\lambda < \lambda_c$ then G_s contains no giant component at all times almost surely. Using standard terminology, this means that the mobile geometric graph model has no *exceptional times*.

We address four fundamental problems in the mobile geometric graph model: detection, coverage, percolation and broadcast. The following results are presented in Chapter 3 and are based on joint work with Yuval Peres, Alistair Sinclair and Perla Sousi [80, 73].

Detection. Consider a target particle u that, at time 0, is located at the origin of $\mathbb{R}^d$. We say that u is *detected* by the graph at time s if there exists a node of G_s within distance r of u. We define the *detection time* T_{det} as the first time at which u is detected by the graph. Kesidis, Konstantopoulos and Phoha [58] (see also [60]) studied the case where u is non-mobile (only the nodes of the graph move) and, using ideas from stochastic geometry, they show that

$$\mathbf{P}\left(T_{\text{det}} > t\right) = \exp\left(-\lambda\, \mathbf{E}\left[\text{vol}\left(W_0(t)\right)\right](1 + o(1))\right), \tag{1.2}$$

where $W_0(t)$ is the so-called "Wiener sausage" up to time t (essentially the trajectory of a Brownian motion "fattened" by a disk of radius r). The expected volume of $W_0(t)$ is well known [82, 11] and, in two dimensions, it scales as $\Theta(t/\log t)$. We extend this result for the case where u also moves. We show in Theorem 3.1 that, in two dimensions, provided the motion of u is continuous, the best strategy for u to avoid detection is to stay fixed and not to move. The intuition behind this result is that, if u has remained at the origin and has not been detected from time 0 to some time s, then this implies that the density of nodes close to the origin is smaller than the typical density. Therefore, it is better for u to stay at its current position than to explore other regions of the space.

Coverage. Complementing the study of detection, we define the *coverage time* of a cube Q_R of side length R as the first time at which *all* points in Q_R are detected by the nodes of the mobile geometric graph. Let $T_{\text{cov}}(Q_R)$ be the coverage time of Q_R. A natural question proposed by Konstantopoulos [60] is to derive the asymptotics of $T_{\text{cov}}(Q_R)$ as $R \to \infty$. We establish that the expected coverage time of Q_R scales as

$$\mathbf{E}\left[T_{\text{cov}} Q_R\right] = \begin{cases} \Theta(\log^2 R), & \text{for } d = 1 \\ \Theta(\log R \log\log R), & \text{for } d = 2 \\ \Theta(\log R), & \text{for } d \geq 3. \end{cases} \tag{1.3}$$

We obtain the precise constants for the expressions above and also show that the expected coverage time of Q_R is concentrated around its expectation as $R \to \infty$. This result is given in Theorem 3.4.

Percolation. As discussed above, when $\lambda > \lambda_c$ the mobile geometric graph contains a giant component at all times. However, the giant component changes over time and, from the perspective of a given node u, it is important to known what is the first time at which u belongs to the giant component. We call this the *percolation time* and denote it by T_{perc}. First notice that the percolation time is at least as large as the dectection time. Therefore,

it follows from (1.2) and results for the volume of the Wiener sausage [82, 11] that

$$\mathbf{P}\left(T_{\text{perc}} > t\right) \geq \begin{cases} \exp(-ct/\log t), & \text{for } d = 2 \\ \exp(-ct), & \text{for } d \geq 3 \end{cases}$$

for some dimension-dependent constant c. (Note that a giant component does not exist when $d = 1$ for any value of λ.) We combine ideas from coupling and multi-scale analysis to derive an upper bound for $\mathbf{P}\left(T_{\text{perc}} > t\right)$ which is tight up to $\log t$ factors in the exponent. Our result for this problem, which is given in Theorem 3.6, establishes that

$$\mathbf{P}\left(T_{\text{perc}} > t\right) \leq \exp\left(-c\frac{t}{\log^{3+6/d} t}\right),$$

for some constant c. In proving this result, we develop a coupling argument that gives a general framework for handling dependencies over time. This technique has already been successfully employed in other contexts [9].

Broadcast. For the broadcast problem, we consider a finite mobile geometric graph over the torus $\mathbb{S}$ of volume n. We assume that an arbitrary node has a piece of information at time 0 and that a node gets informed whenever it is in the same connected component of an informed node. In mobile wireless networks, the speed of transmission is usually much higher than the speed of the motion of the nodes, so it is realistic to assume that messages can propagate instantaneously throughout an entire connected component. Our result for the percolation time can be used to show that this model of broadcast informs all the nodes of a mobile geometric graph with $\lambda > \lambda_c$ within $O(\log n(\log\log n)^{3+6/d})$ steps (see Corollary 3.7). This result is optimal up to $\log\log n$ factors.

Randomized Broadcast on General Graphs

In Chapter 4 we return to the push algorithm of Chapter 2 in a more general class of graphs, while Chapters 2 and 3 both focused on geometric graphs. Recall that, in this algorithm, initially there exists one informed node and, at each discrete time step, each informed node picks a neighbor independently and uniformly at random and informs it. It is easy to see that, for any graph, the push algorithm takes at least $\Omega(\text{diam}(G) + \log n)$ steps to inform all nodes, since the set of informed nodes can at most double in size at each step. A natural question to ask is for which classes of graphs is the push algorithm optimal up to constant factors. This is known for some types of graphs such as complete graphs [42, 76] and Erdős-Rényi random graphs [39, 40] (see the discussion in Section 1.2 for more results). Also, note that our results above establish that the push algorithm is optimal for (static) random geometric graphs.

We are mostly interested in studying how the notion of expansion on graphs affects the performance of the push algorithm. For a graph G with n nodes and node set V, we define the *conductance* of G by

$$\Phi = \Phi(G) = \min_{V' \subset V:\ \text{vol}(V') \leq \text{vol}(V)/2} \frac{E(V', V \setminus V')}{\text{vol}\,(V')}, \tag{1.4}$$

where $E(V', V \setminus V')$ is the number of edges with one endpoint in V' and another in $V \setminus V'$ and $\mathrm{vol}\,(V')$ is the sum of the degrees of the nodes in V'. A similar notion of expansion is the *vertex expansion*, which is defined as

$$\alpha = \alpha(G) = \min_{V' \subset V:\, |V'| \leq n/2} \frac{|\partial V'|}{|V'|}, \tag{1.5}$$

where $\partial V'$ is the set of nodes of $V \setminus V'$ that are neighbors of a node in V'.

It was shown by Mosk-Aoyama and Shah [69] that the runtime of the push algorithm is $O(\Phi^{-1} \log n)$ on any regular graph. This implies that, when $\Phi = \Theta(1)$, the push algorithm is optimal up to constant factors. Chierichetti, Lattanzi and Panconesi [20] posed the problem of whether the runtime of the push algorithm could be related to the vertex expansion of the graph. Our main result in this section says that, for any regular graph with vertex expansion α, the runtime of the push algorithm can be upper bounded by $O(\alpha^{-1} \log^5 n)$ (see Theorem 4.3). Since our result is not proved to be tight, one could ask whether it is possible to derive a result that is analogous to the result of Mosk-Aoyama and Shah [69] for the conductance; that is, whether the runtime of the push algorithm can be written as $O(\alpha^{-1} \log n)$. We show in Theorem 4.5 that this is false by constructing a regular graph with $\alpha = \Theta(1)$, and which is also vertex transitive, for which the runtime of the push algorithm is $\Omega(\log^2 n)$. Also, this graph has diameter $\Theta(\log n)$, which shows that even for regular graphs with constant vertex expansion the push algorithm may not be optimal. This establishes a substantial difference between these two notions of expansion in the context of broadcast on graphs. The results described above, which are discussed in Chapter 4, are based on a joint work with Thomas Sauerwald [78].

1.2 Related Work

Static Networks

Random geometric graphs have long been used as a model for *static* wireless networks, and by now their structural and algorithmic properties are rather well understood mathematically. We mention just a few results here that are directly related to the problems studied in this thesis, and refer the reader to the books by Meester and Roy [65] and Penrose [71] for extensive background on the subject.

We remark that most of the theoretical results on random geometric graphs consider networks above the connectivity threshold, which requires the unrealistic assumption that the network has an unbounded average degree. Our results for static networks discussed above consider the more realistic setting where λ and r are constants but $\lambda > \lambda_c$. Therefore, the graph has constant average degree and contains a giant component with $\Theta(n)$ nodes w.h.p. In this case, we can model the network by the giant component of the random geometric graph. We regard this as the minimum possible assumption to impose on the random geometric graph model so that a notion of network makes sense.

Even very simple fundamental properties of random geometric graphs are only known in the connected case. For example, Ellis, Martin and Yan [33] established in 2007 that the diameter of a two-dimensional random geometric graph above the connectivity threshold is $\Theta(\sqrt{n}/r)$. Our result in (1.1) about graph distances extends this result by giving as a corollary that, in two dimensions, the diameter of the giant component of a random geometric graph with $\lambda > \lambda_c$ is $\Theta(\sqrt{n}/r)$.

A process related to broadcast that has been studied for random geometric graphs is the cover time of random walks. It was shown by Avin and Ercal [5] that, for two-dimensional connected random geometric graphs, the cover time is $\Theta(n \log n)$ w.h.p., which is optimal up to constant factors. More recently, Cooper and Frieze [23] gave a more precise estimate of the cover time of random geometric graphs that extends also to higher dimensions.

In the context of wireless networks, the problem of determining the information-theoretic capacity of the network has received much attention. For static networks, as mentioned above, this problem was solved by Gupta and Kumar [51, 52] and Franceschetti et al. [41], who showed that the capacity of random geometric graphs goes to 0 as $n \to \infty$. In a slightly different context, Frieze et al. [43] studied the effect of physical obstacles that obstruct transmissions.

The *detection* problem that we study for mobile graphs has also been studied in the static case. It is well known that a target is detected w.h.p. only if λ grows with n. When λ is constant, Balister et al. [8] determine the maximum diameter of the uncovered regions, while Dousse, Tavoularis and Thiran [30] prove that, for any $\lambda > 0$, the detection time for a target moving in a fixed direction has an exponential tail. (Note that this is not a mobility result as the nodes are fixed.)

A different broadcast model known as *radio broadcasting* has also been studied on random geometric graphs [25, 62]. In this model, every transmission by a node is sent to all its neighbors simultaneously. However, if two (or more) transmissions are sent to the same node in one round, then this node cannot receive any of the transmissions. In order to derive an efficient algorithm for radio broadcasting on random geometric graphs, Lotker and Navarra [62] studied this problem first on a grid. Then, they emulated the corresponding grid protocol on random geometric graphs, and obtained an optimal algorithm (up to constant factors) for broadcasting above the connectivity threshold. However, the result of [62] only holds if each node is aware of its own position. Later, Czumaj and Wang [25] considered various scenarios with respect to the local knowledge of each node in the graph, and showed that, in many settings, radio broadcasting[2] can be solved in time $O(\mathrm{diam}(G))$.

Mobile Networks

The scope of mathematically rigorous work with mobile nodes is much more limited, and there is as yet no widespread agreement on an appropriate model for node mobility. The model we use in this paper was introduced by van den Berg, Meester and White [12] under

[2]In [25] the so-called gossiping problem was considered, where each node possesses a different message, and all these messages have to be disseminated efficiently to every node in the graph.

the name "dynamic boolean model." They proved that almost surely an infinite component exists at *all* times if $\lambda > \lambda_c$. We point out that in this model, in contrast to many others, node mobility is fixed and does not depend on n (the volume of the region $\mathbb{S}$ in a finite network).

Motivated by the fact mentioned above [51] that the capacity of static networks goes to zero as $n \to \infty$, Grossglauser and Tse [49] (see also [28]) showed how to exploit mobility to achieve constant capacity using a two-hop routing scheme. However, these results require the unrealistic assumption that nodes move a distance comparable to the diameter of the entire region $\mathbb{S}$ at each step. El Gamal et al. [44] study the tradeoff between capacity and delay in a realistic mobility model but above the connectivity threshold.

Clementi, Pasquale and Silvestri [22] show how to exploit mobility to enable broadcast in a random geometric graph sufficiently far above the percolation threshold. However, this result again assumes that the range of motion of the nodes grows with n. In a similar model, where the nodes are performing continuous-time random walks on $\mathbb{Z}^d$, Kesten and Sidoravicius [59] studied the case when λ is constant and showed that the rate at which a message can be disseminated through the graph is $\Theta(t)$. This means that, after t steps, the set of informed nodes contains, and is contained in, balls of radius $\Theta(t)$. Finally, Pettarin et al. [75] show that the broadcast time is optimal up to logarithmic factors on finite graphs with mobile nodes for the case where the intensity λ is also allowed to go to zero as $n \to \infty$.

The detection problem was addressed for mobile graphs by Liu et al. [61], assuming that each node moves continuously in a *fixed* randomly chosen direction. They show that the time it takes for the network to detect a target is exponentially distributed with expectation depending on the intensity λ. Also, for the special case of a stationary target, as observed by Kesidis, Konstantopoulos and Phoha in [58, 60], the detection time can be very precisely deduced from classical results on continuum percolation and stochastic geometry [83]: namely, it is shown in [58, 60] that $\mathbf{Pr}[T_{\text{det}} > t] = \exp(-\lambda \mathbf{E}[\text{vol}(W_0(t))])$, where $W_0(t)$ is the so-called "Wiener sausage" up to time t. This volume in turn is known quite precisely [82, 11]. Drewitz et al. [32], following initial analysis by Moreau et al. [67], established that the best strategy for a target to avoid detection is to stay put when particles move as continuous-time random walks in $\mathbb{Z}^d$. Very recently, Peres and Sousi [74] extended our results and showed that, for a mobile geometric graph in any dimension, and for any fixed time $t > 0$, the best strategy for the target to avoid detection up to time t is to stay put.

Recent work of Díaz et al. [27] in a similar model determines, for networks exactly at the connectivity threshold, the expected length of time for which the network stays connected (or disconnected) as the nodes move. However, this question makes sense only for very large values of λ (growing with n).

Randomized Broadcast on General Graphs

We review some of the existing results about the push algorithm. This algorithm was first analyzed by Frieze and Grimmett [42], who proved that, w.h.p., the runtime is $\log_2 n + \ln n + o(\log n)$ on complete graphs. This result was later improved by Pittel [76]. Feige et

al. [39] proved that, on any graph, the runtime is at most $O(n \log n)$ w.h.p., and that, for any bounded-degree graph, $O(\mathrm{diam}(G))$ rounds are sufficient. Furthermore, they established a runtime of $O(\log n)$ on hypercubes and sufficiently dense Erdős-Rényi random graphs. This result was extended by Elsässer and Sauerwald [34], who established an upper bound of $O(\mathrm{diam}(G) + \log n)$ on some classes of Cayley graphs.

To analyze the runtime of the push algorithm, it is very natural to look at the conductance of the set of informed nodes (see (1.4) for a definition), since this reflects the proportion of transmissions from the informed nodes that reach uninformed nodes. Many results bounding the runtime of the push algorithm rely heavily on lower bounds on the conductance [10, 29, 40, 77] (some of these also require lower bounds on other parameters). For regular graphs, it follows from a result of Mosk-Aoyama and Shah [69] and Sauerwald [77] that the push algorithm informs all nodes within $O(\Phi^{-1} \log n)$ steps w.h.p.

Very recently, Chierichetti, Lattanzi and Panconesi [20, 19] studied the so-called push-pull algorithm, which is a variant of the push algorithm in which, at each step, each node u (informed or uniformed) chooses a neighbor independently and uniformly at random and, if u is informed, it informs the chosen neighbor, while if u is not informed but the chosen neighbor is informed, then u becomes informed. This latter operation is referred to as a *pull* operation. Chierichetti et al. analyzed the relation between the conductance of a (possibly non-regular) graph and the performance of the push-pull algorithm. They show that this algorithm informs all nodes within $O\left(\frac{\log^2(\Phi^{-1})}{\Phi} \log n\right)$ steps. They also show in [18] that there are (non-regular) graphs with constant conductance for which the push algorithm requires a runtime of $\Omega(\mathrm{poly}(n))$ w.h.p.

We study in this thesis the relation between the *vertex expansion* (as defined in (1.5)) and the runtime of the push algorithm. This measure of expansion has been studied in diverse areas, such as random walks [6, 17], property testing [24], and graph theory [2]. In the context of broadcasting, our results answer a question of Chierichetti, Lattanzi and Panconesi [20], who stated that "an outstanding open problem in this area is whether [constant] vertex expansion implies that rumor spreading is fast."

Finally, we point out that the runtime of the push algorithm has also been related to other processes in computer science. For example, Boyd et al. [15], Elsässer and Sauerwald [34] and Sauerwald [77] studied the relation between the push algorithm and the mixing time of random walks on graphs. Also, the push algorithm has been related to the cover time of random walks by Elsässer and Sauerwald [35].

Chapter 2

Static Networks

In this chapter we consider a two-dimensional random geometric graph over the torus $\mathbb{S}$ of volume n. Recall that a random geometric graph G over $\mathbb{S}$ is obtained by taking the nodes to be distributed as a Poisson point process with intensity λ over $\mathbb{S}$, and creating an edge between every pair of nodes whose Euclidean distance is at most r. It is well known that the structure of G depends only on the product λr^2. Therefore, we take λ to be a fixed constant and let r be the only parameter of the model. We note that r may or may not depend on n. In contrast with the convention adopted in Chapter 1, we choose to parameterize the model on r here since this simplifies the discussion to follow. Using this convention, we have that there exists a critical value r_c such that, if $r > r_c$, then G contains a *giant component* with $\Theta(n)$ nodes w.h.p. (that is, with probability going to 1 as $n \to \infty$). Otherwise, if $r < r_c$, then all components of G have size $O(\log n)$.

We focus on the regime $r > r_c$ where the random geometric graph is likely to contain a giant component. We first show that, for any two nodes having sufficiently large Euclidean distance, their graph distance is just a constant factor larger than the optimum. In particular, this result shows that the diameter of the giant component is $\Theta(\sqrt{n}/r)$, which was only previously known for the case when G is connected w.h.p. [33]. We then use this result to show that, if a node of the giant component employs the push algorithm to broadcast a piece of information, then w.h.p. all nodes of the giant component get informed within $O(\sqrt{n}/r + \log n)$ rounds.

Now we state our results more formally. For any pair of nodes $v_1, v_2 \in G$, we say that v_1 and v_2 are *connected* if there exists a path in G from v_1 to v_2, and define $d_G(v_1, v_2)$ as the distance between v_1 and v_2 in G. Also, we denote the Euclidean distance between the locations of v_1 and v_2 by $\|v_1 - v_2\|_2$. Clearly, the smallest path between two nodes v_1 and v_2 in G must satisfy $d_G(v_1, v_2) \geq \|v_1 - v_2\|_2/r$. We then obtain the following theorem relating graph distances to Euclidean distances. We give the proof of this theorem in Section 2.1.

Theorem 2.1. *If $r > r_c$, for any two connected nodes v_1 and v_2 in G such that $\|v_1 - v_2\|_2 = \Omega(\log^6 n/r^4)$, we have $d_G(v_1, v_2) = \Theta(\|v_1 - v_2\|_2/r)$ w.h.p.*

The theorem above clearly yields the following bound on the diameter of the giant component.

Corollary 2.2. *If $r > r_c$, the diameter of the giant component of G is $\Theta(\sqrt{n}/r)$ w.h.p.*

For each node s of G, let $T_{\text{push}}(s, G)$ be the time it takes for the push algorithm to inform all nodes in the same connected component as s given that s is the only node that is initially informed. We show that, if $r > r_c$, for all nodes s in the giant component of G, w.h.p., $T_{\text{push}}(s, G)$ is within a constant factor of the minimum possible value based on the diameter. The theorem below is proved in Section 2.2.

Theorem 2.3. *For a random geometric graph G with $r > r_c$, $T_{\text{push}}(s, G) = O(\sqrt{n}/r + \log n)$ w.h.p. for all nodes s inside the giant component of G.*

2.1 Diameter and graph distance

We devote this section to proving Theorem 2.1. We consider G to be a random geometric graph with $r > r_c$ and assume that $r = O(\sqrt{\log n})$. (When $r = \omega(\sqrt{\log n})$, G is connected w.h.p. and Theorem 2.1 becomes a slightly different version of [33, Theorem 8].) We show that for any two connected nodes v_1 and v_2 of G such that $\|v_1 - v_2\|_2 = \Omega(\log^6 n/r^4)$, we have $d_G(v_1, v_2) = \Theta(\|v_1 - v_2\|_2/r)$ with probability $1 - O(n^{-1})$.

We first take two fixed nodes v_1 and v_2 satisfying the conditions above and show that $d_G(v_1, v_2) = O(\|v_1 - v_2\|_2/r)$ with probability $1 - O(n^{-3})$. Then, we would like to take the union bound over all pairs of nodes v_1 and v_2 to conclude the proof of Theorem 2.1; however, the number of nodes in G is a random variable and hence the union bound cannot be employed directly. We resort to the following lemma to extend the result to all pairs of nodes v_1 and v_2.

Lemma 2.4. *Let $E(w_1, w_2)$ be an event associated with a pair of nodes $w_1, w_2 \in G$. Assume that for all pairs of nodes, $\mathbf{P}(E(w_1, w_2)) \geq 1 - p$, with $p \in (0, 1)$. Then,*

$$\mathbf{P}\left(\bigcap_{w_1, w_2 \in G} E(w_1, w_2)\right) \geq 1 - 9\lambda^2 n^2 p - e^{-\Omega(n)}.$$

Proof. Let N be the number of nodes in G. We condition on $N \leq 3\lambda n$. Using a Chernoff bound for Poisson random variables, it follows easily that $\mathbf{P}(N > 3\lambda n) \leq e^{-\Omega(n)}$. Let $E^c(w_1, w_2)$ denote the complement of $E(w_1, w_2)$. Note that $\mathbf{P}(E^c(w_1, w_2) \mid N \leq 3\lambda n) \leq \frac{\mathbf{P}(E^c(w_1, w_2))}{\mathbf{P}(N \leq 3\lambda n)} \leq \frac{p}{1 - e^{-\Omega(n)}}$, for all $w_1, w_2 \in G$. Therefore, using the definition of conditional

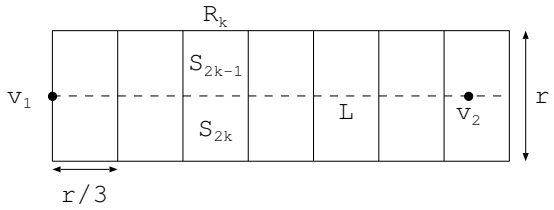

Figure 2.1: Illustration for the calculation of $d_G(v_1, v_2)$, with the large $r \times r/3$ rectangle R_k and the cells S_{2k-1} and S_{2k} contained in R_k.

probabilities and the union bound, we obtain

$$\begin{aligned}
&\mathbf{P}\left(\bigcup_{w_1,w_2\in G} E^{\mathrm{c}}(w_1,w_2)\right)\\
&\leq \mathbf{P}\left(\bigcup_{w_1,w_2\in G} E^{\mathrm{c}}(w_1,w_2)\;\middle|\; N\leq 3\lambda n\right)\mathbf{P}\left(N\leq 3\lambda n\right)+\mathbf{P}\left(N>3\lambda n\right)\\
&\leq 9\lambda^2 n^2 \max_{w_1,w_2\in G}\mathbf{P}\left(E^{\mathrm{c}}(w_1,w_2)\mid N\leq 3\lambda n\right)+e^{-\Omega(n)}\\
&\leq 9\lambda^2 n^2 p+e^{-\Omega(n)}.
\end{aligned}$$

□

We turn now to the main task of finding a short path between our two nodes v_1 and v_2 (see Figure 2.1). Take the line L that contains v_1 and v_2 and draw a sequence of adjacent rectangles starting from v_1 until we draw a rectangle that contains v_2. Each rectangle has two sides with length $r/3$ that are parallel to L and two other sides with length r that are perpendicular to L such that their middle point is contained in L. Let κ be the number of such rectangles and refer to them as $R_1, R_2, \ldots, R_\kappa$. For each $k \in [1, \kappa]$, L splits R_k into two identical, smaller rectangles which we denote by S_{2k-1} and S_{2k} and refer to as "cells."

Note that for any k and two points $x \in S_k$ and $x' \in S_{k+2}$, we obtain $\|x - x'\|_2 \leq \sqrt{(2r/3)^2 + (r/2)^2} \leq r$, that is, nodes in S_k and S_{k+2} are neighbors in G. For this reason, we say that the cell S_k is *adjacent* to the cells S_{k-2} and S_{k+2}. Note that v_1 belongs to both S_1 and S_2. We would like to find a path from v_1 to v_2 that starts at either S_1 or S_2 and moves along adjacent cells, but we have to handle the possibility that some S_k may contain no nodes.

Our choice for the length of the largest sides of the rectangle R_k is intended to achieve the following property. For any path in G that crosses the region $\bigcup_{i=1}^{\kappa} R_i$, in the sense that there exists an edge of the path that intersects $\bigcup_{i=1}^{\kappa} R_i$, it must be the case that the path contains a node inside $\bigcup_{i=1}^{\kappa} R_i$. This property is crucial in our analysis, since it guarantees that a path crossing two rectangles R_j and R_k provides a path from a node in R_j to a node in R_k in G and can be used to move around cells that contain no nodes.

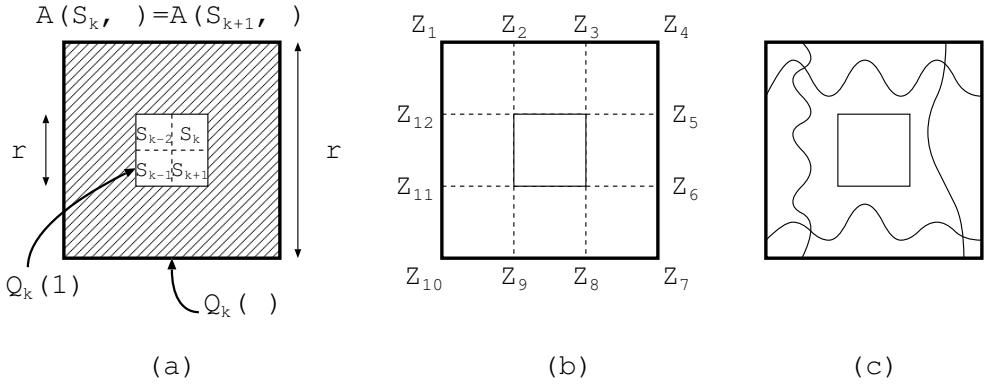

Figure 2.2: Illustration for the annulus $A(S_k, \gamma)$. Part (a) shows the annulus (highlighted region) and the cells S_{k-2}, S_{k-1}, S_k, and S_{k+1} in the middle. Part (b) shows the decomposition of $A(S_k, \gamma)$ into horizontal and vertical rectangles. And part (c) illustrates the event $F(A(S_k, \gamma))$ for the left-to-right and top-to-bottom crossings (depicted as curvy lines) of $A(S_k, \gamma)$.

We refer to a cell as *empty* if it contains no node. For any empty cell S_k with S_{k-2} being nonempty, we follow the shortest path from a node in S_{k-2} to some nonempty $S_{k'}$ for $k' \geq k+1$. Note that there is always such a k' since $R_\kappa = S_{2\kappa-1} \cup S_{2\kappa}$ contains v_2. Our aim is to give a bound for the length of the detour around empty cells. The path starts at $v_1 \in R_1$. For $3 \leq k \leq 2\kappa$, if S_k is empty and S_{k-2} is not empty, let D_k be the length of the shortest path from S_{k-2} to some $S_{k'}$ for $k' \geq k+1$. If S_k is not empty, we set $D_k = 0$. Also, if S_k and S_{k-2} are both empty, then we also set $D_k = 0$, since the detour around S_{k-2} will either go around S_k as well or lead to S_{k-1}, from which we can obtain an edge to S_{k+1} or a detour that goes around S_k. With these definitions we can write $d_G(v_1, v_2) \leq \kappa + \sum_{k=3}^{2\kappa} D_k$.

In order to calculate D_k, we exploit the idea of "crossings" from continuum percolation. For an odd number $k \geq 1$, we consider the cells $S_{k-2}, S_{k-1}, S_k, S_{k+1}$. Let $Q_k(1)$ be the rectangle containing all these cells, that is, $Q_k(1) = R_{(k-1)/2} \cup R_{(k+1)/2}$. Let $Q_k(\gamma)$ be a rectangle having the same center as $Q_k(1)$ and whose sides are parallel to those of $Q_k(1)$ and have length given by γ times the side lengths of $Q_k(1)$ (in other words, $Q_k(\gamma)$ is a stretched version of $Q_k(1)$). Then, for any odd number $k \geq 1$ and $\gamma > 1$, we define the annulus $A(S_k, \gamma) = A(S_{k+1}, \gamma) = Q_k(\gamma) \setminus Q_k(1)$ (see Figure 2.2(a)).

An annulus $A(S_k, \gamma)$ can be decomposed into two horizontal rectangles ($Z_1Z_4Z_5Z_{12}$ and $Z_{11}Z_6Z_7Z_{10}$ in Figure 2.2(b)) and two vertical rectangles ($Z_1Z_2Z_9Z_{10}$ and $Z_3Z_4Z_7Z_8$ in Figure 2.2(b)). For a horizontal rectangle, we define a horizontal crossing as a path in G completely contained in the rectangle and that connects the left to the right side of the rectangle, i.e., with the first node of the path being within distance r of the left side of the rectangle and the last node of the path being within distance r of the right side of the rectangle. Similarly, for a vertical rectangle, we define a vertical crossing as a path in G

that is completely contained in the rectangle and that connects the top to the bottom side of the rectangle. For an annulus $A(S_k, \gamma)$, we define $F(A(S_k, \gamma))$ as the event that both horizontal rectangles of $A(S_k, \gamma)$ have a horizontal crossing and that both vertical rectangles of $A(S_k, \gamma)$ have a vertical crossing. This event is illustrated in Figure 2.2(c). Note that when $F(A(S_k, \gamma))$ happens, the aforementioned crossings provide a cycle around S_k.

We now explain how to use the annuli to find detours around an empty cell S_k. Note that S_1 and S_2 contain v_1 and, consequently, are not empty. Now suppose that S_{k-2} is not empty and is connected to v_1, i.e., there is a path from v_1 to a node inside S_{k-2}. If S_k is also not empty, then the node inside S_k is a neighbor of the node in S_{k-2}, and we obtain a path from v_1 to S_k. Now, assume that S_k is empty. We want to use the path from v_1 to S_{k-2} to construct a path from v_1 to some $S_{k'}$ with $k' \geq k+1$. Clearly, for any $\gamma > 1$, the annulus $A(S_k, \gamma)$ intersects neither S_{k-2} nor S_k, but does intersect S_{k+2}. Take γ' such that $F(A(S_k, \gamma'))$ happens and let $H \subset Q_k(\gamma')$ be the largest region delimited by the cycle surrounding S_k that is induced by the crossings of $A(S_k, \gamma')$. If $v_1 \notin H$, then the path from v_1 to S_{k-2} provides a path from S_{k-2} to the crossings of $A(S_k, \gamma')$. If the crossings intersect some nonempty $S_{k'}$, $k' \geq k+1$, then there is a path entirely contained in H from S_{k-2} to a node inside $S_{k'}$. If such a $S_{k'}$ does not exist, it must be the case that $v_2 \in H$. Since v_1 and v_2 are connected, there is a path from v_2 to the crossing of $A(S_k, \gamma')$, and, consequently, there is a path from S_{k-2} to v_2 completely contained in H. Now, if $v_1 \in H$ and $S_{k'}$ as above exists, then the path from v_1 to v_2 intersects the crossings of $A(S_k, \gamma')$ and can be used to obtain a path completely contained in H from S_{k-2} to $S_{k'}$. Finally, if $v_1 \in H$ and $v_2 \in H$, then there is a path from v_1 to v_2 entirely contained in H.

This shows that whenever v_1 and v_2 are connected, we can use the annuli to move from S_k to $S_{k'}$, $k' \geq k+1$, or to move directly to v_2. Note that the constructions of S_k and $A(S_k, \gamma)$ apply to an arbitrary pair of nodes v_1 and v_2, independently of whether v_1 and v_2 are connected. This means that our calculations to follow are not conditioned on v_1 and v_2 being connected. However, when v_1 and v_2 turn out to be connected, this construction provides a path from v_1 to v_2.

Once we know that $F(A(S_k, \gamma))$ occurs for some γ, we can easily bound D_k by the following straightforward geometric lemma.

Lemma 2.5. *Let Q be a rectangle with side lengths s and αs. Let w_1 and w_2 be two nodes of G contained in Q. If there exists a path between w_1 and w_2 entirely contained in Q, then $d_G(w_1, w_2) \leq 11\alpha s^2/r^2$.*

Proof. The shortest path between w_1 and w_2 that is contained inside Q has the property that for any two non-consecutive nodes u and u' in the path, their distance is larger than r. Otherwise, we can take the edge (u, u') and make the path shorter. This means that if we draw a ball of radius $r/2$ around every other node of the path, then the balls will not overlap. Let m be the number of nodes in the path. There are $m/2$ non-overlapping balls of radius $r/2$. For each ball, at least $1/4$ of its area is contained inside Q. Therefore, it must

be the case that

$$m \leq 2\frac{\text{vol}\,(Q)}{\pi(r/2)^2/4} = \frac{32\alpha s^2}{\pi r^2}. \qquad \square$$

The lemma below gives an upper bound for the probability that $A(S_k, \gamma)$ does not have the required crossings.

Lemma 2.6. *There exist constants c and $\gamma_0 > 1$ such that, for all $\gamma > \gamma_0$ and $1 \leq k \leq 2\kappa$, we have that*

$$\mathbf{P}\left(F(A(S_k, \gamma))\right) \geq 1 - \exp\left(-c\gamma r\right).$$

Proof. We build upon ideas from the proof of [71, Lemma 10.5]. Recall the decomposition of $A(S_k, \gamma)$ into rectangles (refer to Figure 2.2(b)) and take the top rectangle $Z_1Z_4Z_5Z_{12}$. Its sides have lengths $(\gamma - 1)r/2$ and $2\gamma r/3$. Therefore, the aspect ratio of the rectangle is $3(\gamma - 1)/(4\gamma) \leq 3/4$, which increases with γ. We want to calculate the probability that such a rectangle has a horizontal crossing as γ increases. This is slightly different from the calculation in [71, Lemma 10.5], since there the aspect ratio is fixed and the side of the rectangle is allowed to vary. But clearly, for any rectangle with side lengths $(\gamma - 1)r/2$ and $2\gamma r/3$, we can stretch the largest sides (while keeping the smallest sides unchanged) to make the aspect ratio be $3(\gamma_0 - 1)/(4\gamma_0)$, which we can then fix. Also, if there is a horizontal crossing in the stretched rectangle, there must be a horizontal crossing in the original one. Following along the lines of the proof of [71, Lemma 10.5], we can then conclude that there are constants γ_0 and c such that, for all $\gamma \geq \gamma_0$, a rectangle of side lengths $(\gamma - 1)r/2$ and $2\gamma r/3$ has a horizontal crossing with probability larger than $1 - e^{-c\gamma r}/4$. Applying the union bound over the four rectangles composing $A(S_k, \gamma)$ concludes the proof. $\square$

Now we use this lemma to bound the length of a detour. For any k, let Γ_k be the smallest value of $\gamma > \gamma_0$ for which $F(A(S_k, \gamma))$ occurs. Suppose that S_k is empty and S_{k-2} is not empty. We want to obtain an upper bound for Γ_k. Note that once we know the value of Γ_k, we can apply Lemma 2.5 to conclude that $D_k \leq (22/3)\Gamma_k^2$. Since for each v_1 and v_2 there are at most $2\kappa = O(\sqrt{n})$ cells, Lemma 2.6 implies that, for all k, $D_k \leq c_1 \log^2 n/r^2$ with probability $1 - O(n^{-4})$ for some constant c_1. Let $E(v_1, v_2)$ be the event that $D_k \leq c_1 \log^2 n/r^2$ for a fixed pair of nodes v_1 and v_2, and all k. Thus, $\mathbf{P}\left(E(v_1, v_2)\right) \geq 1 - O(n^{-4})$.

We want to apply Azuma's inequality to $\sum_{k=3}^{2\kappa} D_k$ under the condition that $E(v_1, v_2)$ happens. Noting that $\mathbf{E}\left[D_k \mid E(v_1, v_2)\right] \leq \mathbf{E}\left[D_k\right] / \mathbf{P}\left(E(v_1, v_2)\right)$, we proceed to derive an upper bound for $\mathbf{E}\left[D_k\right]$. The probability that S_{k-2} is not empty and S_k is empty is $e^{-r^2/6}(1 - e^{-r^2/6})$. Recall that $D_k \leq (22/3)\Gamma_k^2$. Therefore,

$$\mathbf{P}\left(D_k \geq \ell\right) \leq 1 - \mathbf{P}\left(F(A(S_k, \sqrt{(3/22)\ell})\right) \leq \exp(-c\sqrt{(3/22)\ell}r).$$

We can then write

$$\mathbf{E}\left[D_k\right] = e^{-r^2/6}(1 - e^{-r^2/6}) \sum_{\ell=1}^{\infty} \mathbf{P}\left(D_k \geq \ell\right) \leq e^{-r^2/6} \int_0^{\infty} \mathbf{P}\left(D_k \geq \ell\right) d\ell,$$

where the last inequality follows from $\mathbf{P}(D_k \geq \ell)$ being a non-increasing function of ℓ. Since we have an exponential upper bound for $\mathbf{P}(D_k \geq \ell)$ with $\ell \geq (22/3)\gamma_0^2$, we obtain

$$\mathbf{E}[D_k] \leq e^{-r^2/6}(22/3)\gamma_0^2 + e^{-r^2/6}\int_{\ell=(22/3)\gamma_0^2}^{\infty} e^{-c\sqrt{(3/22)\ell}r}d\ell = O(1).$$

Using the linearity property of expectations and $\mathbf{P}(E(v_1,v_2)) \geq 1 - O(n^{-4})$, we obtain

$$\mathbf{E}[d_G(v_1,v_2) \mid E(v_1,v_2)] \leq \frac{\mathbf{E}[d_G(v_1,v_2)]}{\mathbf{P}(E(v_1,v_2))} = O(\kappa) = O(\|v_1 - v_2\|_2/r).$$

If the event $E(v_1,v_2)$ holds, we have $\mathbf{E}[D_k \mid E(v_1,v_2)] = O(1)$ and $\Gamma_k \leq c_1' \log n/r$ for all k and some constant c_1', which yields $D_k \leq c_1 \log^2 n/r^2$. Letting $\lambda = 4c_1' \log n/r$, this implies that for two cells S_k and $S_{k'}$ such that $|k - k'| \geq \lambda$, the annuli $A(S_k, \lambda/4)$ and $A(S_{k'}, \lambda/4)$ do not intersect, and consequently, the random variables D_k and $D_{k'}$ are independent. Now we split the random variables $D_1, D_2, \ldots, D_{2\kappa}$ into groups of independent random variables. Define the index set $I_j = \{k\colon 3 \leq k \leq 2\kappa, k \equiv j \pmod{\lambda}\}$. We can write

$$d_G(v_1,v_2) = \kappa + \sum_{j=0}^{\lambda-1}\sum_{k\in I_j} D_k,$$

where the second sum contains independent random variables. Note that

$$\begin{aligned}
&\mathbf{P}\left(\sum_{k\in I_j} D_k - \sum_{k\in I_j}\mathbf{E}[D_k] \geq c_2|I_j|\right) \\
&\leq \; 1 - \mathbf{P}(E(v_1,v_2)) + \mathbf{P}\left(\sum_{k\in I_j} D_k - \sum_{k\in I_j}\mathbf{E}[D_k] \geq c_2|I_j| \;\mid\; E(v_1,v_2)\right).
\end{aligned}$$

In order to apply Azuma's inequality to the last term, we need to write $\sum_{k\in I_j}\mathbf{E}[D_k]$ in terms of $\sum_{k\in I_j}\mathbf{E}[D_k \mid E(v_1,v_2)]$. Since

$$\mathbf{E}[D_k] \geq \mathbf{E}[D_k \mid E(v_1,v_2)]\,\mathbf{P}(E(v_1,v_2)) = \mathbf{E}[D_k \mid E(v_1,v_2)] - O(n^{-4}),$$

we derive that, for each j,

$$\mathbf{P}\left(\sum_{k\in I_j} D_k - \sum_{k\in I_j}\mathbf{E}[D_k] \geq c_2|I_j|\right) \leq 1 - \mathbf{P}(E(v_1,v_2)) + 2\exp\left(-\frac{(c_2 + O(n^{-4}))^2|I_j|^2 r^4}{2c_1^2|I_j|\log^4 n}\right).$$

Since $|I_j| \geq \kappa/\lambda = \Omega(\|v_1 - v_2\|_2/\log n)$, the probability above is smaller than $O(n^{-4}) + \exp\left(-\frac{c_3\|v_1-v_2\|_2 r^4}{\log^5 n}\right)$, for some constant $c_3 > 0$. We solve the first sum by the union bound, obtaining

$$\mathbf{P}\left(\sum_{k=1}^{2\kappa} D_k - \sum_{k=1}^{2\kappa} \mathbf{E}\left[D_k\right] \geq 2c_2\kappa\right) \leq O(\lambda n^{-4}) + \lambda \exp\left(-\frac{c_3\|v_1 - v_2\|_2 r^4}{\log^5 n}\right) = O(n^{-3}),$$

for any v_1 and v_2 such that $\|v_1 - v_2\|_2 \geq c_4 \log^6 n/r^4$, for some constant c_4. Hence, by setting the constant c_2 properly, for a fixed pair of nodes v_1, v_2 such that $\|v_1 - v_2\|_2 = \Omega(\log^6 n/r^4)$, $d_G(v_1, v_2) = O(\|v_1 - v_2\|_2/r)$ with probability $1 - O(n^{-3})$. Applying Lemma 2.4 concludes the proof of Theorem 2.1.

2.2 Broadcast time

In this section we prove Theorem 2.3. Given two nodes v_1 and v_2, let $T_{\text{push}}(v_1, v_2)$ be the time it takes for the push algorithm started at v_1 to inform v_2 for the first time. We assume in the sequel that v_1 and v_2 belong to the giant connected component of G and show that, provided $\|v_1 - v_2\|_2 = \Omega(\log^6 n/r^4)$, $T_{\text{push}}(v_1, v_2) = O(\|v_1 - v_2\|_2/r)$. Initially, we assume that $r = O(\sqrt{\log n})$. The case $r = \omega(\sqrt{\log n})$ is simpler[1], but since it uses different proof techniques, we deal with it at the end of this section.

We start our treatment for the case $r = O(\sqrt{\log n})$ with an easy lemma that shows that the time until a node informs a given neighbor is $O(\log^2 n)$ with high probability.

Lemma 2.7. *Let $r = O(\sqrt{\log n})$. There exists a constant c such that for all pairs of nodes w_1 and w_2 satisfying $\|w_1 - w_2\|_2 \leq r$, the following holds with probability $1 - O(n^{-1})$:*

$$T_{\text{push}}(w_1, w_2) \leq c \log^2 n.$$

Proof. Note that if the degree of w_1 in G is k, then the number of rounds until w_1 sends the information to w_2 is given by a geometric random variable with mean k. It is easy to check that there is a constant c_5 such that with probability $1 - O(n^{-3})$ all nodes of a random geometric graph have degree smaller than $c_5 \log n$ [71] provided $r = O(\sqrt{\log n})$. Therefore,

$$\mathbf{P}\left(T_{\text{push}}(w_1, w_2) \geq t\right) \leq \left(1 - \frac{1}{c_5 \log n}\right)^t \leq \exp\left(-\frac{t}{c_5 \log n}\right).$$

If we set $t = 3c_5 \log^2 n$, we obtain that $\mathbf{P}\left(T_{\text{push}}(w_1, w_2) \geq 3c_5 \log^2 n\right) \leq n^{-3}$ and, using Lemma 2.4 we conclude that $T_{\text{push}}(w_1, w_2) \leq 3c_5 \log^2 n$ for all w_1, w_2 with probability $1 - O(n^{-1})$. □

[1] In order to simplify the notation, we apply the convention that $f(n) = \omega(g(n))$ stands for $\limsup_{n\to\infty} f(n)/g(n) = \infty$.

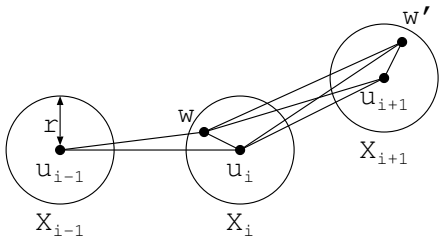

Figure 2.3: Illustration of the path considered to obtain $T_{\text{push}}(v_1, v_2)$. The picture shows three consecutive nodes u_{i-1}, u_i, and u_{i+1} of the path from v_1 to v_2 and the balls X_{i-1}, X_i, and X_{i+1} around them. Two other nodes $w \in X_i$ and $w' \in X_{i+1}$ are depicted to illustrate the edges that arise from the construction of the X_i's.

Before proceeding, note that the lemma above shows that $T_{\text{push}}(v_1, v_2)$ can be upper bounded by $O(d_G(v_1, v_2) \log^2 n)$. We derive a much better bound in the sequel. Let r' be defined such that $r_c < r' < r$. Note that such an r' exists since $r > r_c$. For convenience, write $r' = r(1 - 2\epsilon)$. Let G' be the subgraph of G obtained by using r' instead of r. Since $r' > r_c$, G' contains a connected component of size $\Omega(n)$ with probability $1 - e^{-\Omega(\sqrt{n})}$.

Our strategy to obtain an upper bound for $T_{\text{push}}(v_1, v_2)$ is the following. First, we assume that v_1 and v_2 belong to the giant component of G'. (We address the case where they do not belong to the giant component of G' later.) Then, we take a path in G' from v_1 to v_2. Instead of calculating the time it takes for the push algorithm to transmit the information along this path, which gives a rather pessimistic upper bound, we enlarge the path using the fact that G' is a subgraph of G and calculate the time it takes for the push algorithm to transmit the information along this enlarged path.

Let $u_1, u_2, \ldots, u_m$ be a path from v_1 to v_2 in G', where $u_1 = v_1$ and $u_m = v_2$. For each i, we define the region $X_i \subseteq \mathbb{S}$ in the following way. Set X_1 to be the point where u_1 is located and X_m to be the point where u_m is located; for $2 \leq i \leq m-1$, define X_i to be the ball with center at u_i and radius ϵr. Our goal is to get an upper bound for $T_{\text{push}}(v_1, v_2)$ by following the path $X_1, X_2, \ldots, X_m$ (refer to Figure 2.3).

Define the random variable $T(X_i, X_{i+1})$, $1 \leq i \leq m-1$, as the time the push algorithm takes to first inform a node in X_{i+1} given that it started at a node chosen uniformly at random from X_i. Note that, for any two nodes $w \in X_i$ and $w' \in X_{i+1}$, the triangle inequality and the definition of X_i give $\|w - w'\|_2 \leq 2\epsilon r + \|u_i - u_{i+1}\|_2 \leq r$. Therefore, w and w' are neighbors in G. Moreover, for any i, once the push algorithm informs a node inside X_i, then the node that receives the information is a uniformly random node from X_i. Thus, we can clearly obtain the following upper bound:

$$T_{\text{push}}(v_1, v_2) \leq \sum_{i=1}^{m-1} T(X_i, X_{i+1}).$$

Note that Lemma 2.7 implies $T(X_{m-1}, X_m) = O(\log^2 n)$ with probability $1 - O(n^{-1})$, for each choice of v_1 and v_2. The next lemma gives the expectation of $T(X_i, X_{i+1})$ for each $1 \leq i \leq m-2$.

Lemma 2.8. *For any $1 \leq i \leq m-2$, we have that $\mathbf{E}[T(X_i, X_{i+1})] \leq 1/\epsilon^2$.*

Proof. Let w be a node chosen uniformly at random from X_i. Assume $w \notin X_{i+1}$ (otherwise, the broadcast time from w to X_{i+1} is zero). Let Y be the number of neighbors of w and let Y' be the number of nodes in X_{i+1}. Therefore, $\mathbf{E}[T(X_i, X_{i+1})] = \mathbf{E}[Y/Y']$. We know that $Y \geq 1$ and $Y' \geq 1$, so $Y - 1$ and $Y' - 1$ are Poisson random variables with means πr^2 and $\pi\epsilon^2 r^2$, respectively. Conditional on $Y - 1 = k$, the value of $Y' - 1$ is given by a Binomial distribution with mean $k\frac{\pi\epsilon^2 r^2}{\pi r^2} = k\epsilon^2$. We obtain

$$\begin{aligned}
\mathbf{E}[T(X_i, X_{i+1})] &= \sum_{k=0}^{\infty}\sum_{i=0}^{k} \frac{k+1}{i+1} e^{-\pi r^2}\frac{(\pi r^2)^k}{k!}\binom{k}{i}(\epsilon^2)^i(1-\epsilon^2)^{k-i} \\
&= \frac{1}{\epsilon^2}\sum_{k=0}^{\infty} e^{-\pi r^2}\frac{(\pi r^2)^k}{k!}\sum_{i=0}^{k}\binom{k+1}{i+1}(\epsilon^2)^{i+1}(1-\epsilon^2)^{k-i} \\
&\leq \frac{1}{\epsilon^2}\sum_{k=0}^{\infty} e^{-\pi r^2}\frac{(\pi r^2)^k}{k!} = \frac{1}{\epsilon^2}.
\end{aligned}$$

□

For any two connected nodes v_1 and v_2 such that $\|v_1 - v_2\|_2 = \Omega(\log^6 n/r^4)$ (we come back to the case $\|v_1 - v_2\|_2 = o(\log^6 n/r^4)$ at the end of this section), we know that there is a path like the ones derived in Section 2.1 for the proof of Theorem 2.1. In particular, we know that there is a path $v_1 = u_1, u_2, \ldots, u_{m-1}, u_m = v_2$ such that $m = O(\sqrt{n}/r')$ and, provided $E(v_1, v_2)$ holds, the annuli $A(S_k, \Gamma_k)$ and $A(S_{k'}, \Gamma_{k'})$ are disjoint if $|k - k'| \geq \lambda$. Recall that the cells $S_1, S_2, \ldots, S_{2\kappa}$ have side lengths $r/2$ and $r/3$; therefore, we need to take six adjacent cells together to obtain a rectangle with largest side length $2r$. Recall also that only every other cell is adjacent. Then, for k and k' such that $|k - k'| \geq \lambda + 12$, the distance between any point in $A(S_k, \Gamma_k)$ and any point in $A(S_{k'}, \Gamma_{k'})$ is at least $2r$. Each annulus has at most $c_1 \log^2 n/r^2$ nodes in the path, so letting $\lambda' = (c_1 \log^2 n/r^2)(\lambda + 12) = O(\log^3 n/r^3)$, we obtain that for any two nodes u_i and u_j in the path such that $|i - j| \geq \lambda'$, $\|u_i - u_j\|_2 \geq 2r$ and, consequently, $T(X_i, X_{i+1})$ and $T(X_j, X_{j+1})$ are independent.

It is important to remark that the path has length $m = O(\sqrt{n}/r')$, for all v_1 and v_2. Conditional on the existence of this particular path, the Poisson point process over $\mathbb{S} \setminus \bigcup_{i=1}^{m}\{u_i\}$, where the union is over the points where the nodes of the path are located, remains unchanged since $\bigcup_{i=1}^{m}\{u_i\}$ spans a set of measure 0 in $\mathbb{S}$.

Let the index set $J_j = \{1 \leq i \leq m \colon i \equiv j \pmod{\lambda'}\}$. We can write

$$T_{\text{push}}(v_1, v_2) = O(\log^2 n) + \sum_{j=0}^{\lambda'-1}\sum_{i \in J_j} T(X_i, X_{i+1}),$$

where the first term comes from the time it takes for the push algorithm to inform v_2 once any neighbor of v_2 is informed. For all j, the term $\sum_{i\in J_j} T(X_i, X_{i+1})$ is given by the sum of independent geometric random variables. Using a standard Chernoff bound for geometric random variables (cf. Lemma A.4), we obtain that for each j,

$$\mathbf{P}\left(\sum_{i\in J_j} T(X_i, X_{i+1}) \geq (1+\epsilon)\sum_{i\in J_j} \mathbf{E}\left[T(X_i, X_{i+1})\right]\right) \leq \exp\left(-\epsilon^2 \frac{|J_j|}{2(1+\epsilon)}\right).$$

Note that $|J_j| = \Omega(d_G(v_1, v_2)r^3/\log^3 n) = \Omega(\log n)$, since $d_G(v_1, v_2) \geq \|v_1 - v_2\|_2/r = \Omega(\log^6 n/r^5)$. Using the fact that $\mathbf{E}\left[T(X_i, X_{i+1})\right] = O(1)$ for all i, and taking the union bound over all j, allows us to conclude that for all pairs of connected nodes v_1 and v_2 such that $\|v_1 - v_2\|_2 = \Omega(\log^6 n/r^4)$, there is a constant c_6 for which

$$\mathbf{P}\left(\sum_{j=0}^{\lambda'-1}\sum_{i\in J_j} T(X_i, X_{i+1}) \geq c_6(m-2)\right) \leq n^{-3}.$$

Applying Lemma 2.4, we can conclude that for any two nodes v_1 and v_2 in the giant component of G for which $\|v_1 - v_2\|_2 = \Omega(\log^6 n/r^4)$, we obtain $T_{\text{push}}(v_1, v_2) = \Theta(\|v_1 - v_2\|_2/r)$. Note that there exist $v_1, v_2 \in G$ for which $\|v_1 - v_2\|_2 = \Theta(\sqrt{n})$ and, consequently, $T_{\text{push}}(v_1, v_2) = \Theta(\sqrt{n}/r)$.

Now we treat two remaining cases. First, since G' is a subgraph of G, there may exist some nodes in the giant component of G that do not belong to the giant component of G'. Nevertheless, it is a known fact from random geometric graphs [71, Theorem 10.18] that the second largest component of G' contains $O(\log^2 n)$ nodes with probability $1 - O(n^{-1})$. Therefore, since $T_{\text{push}}(w_1, w_2) = O(\log^2 n)$ for every pair of neighbors w_1 and w_2, we conclude that the time it takes to inform all the remaining nodes is $O(\log^4 n)$, which is negligible in comparison to $\Theta(\sqrt{n}/r)$. The second case corresponds to the nodes that are within distance $o(\log^6 n/r^4)$ of the initially informed node, which we take to be v_1. Take Q to be a square centered at v_1 with side length $c_7 \log^6/r^4$, for some constant c_7 (the orientation of Q does not matter). Note that Q contains all nodes within distance $o(\log^6 n/r^4)$ of v_1. Now, take Q' to be a square centered at v_1, with the same orientation as Q but with sides having twice the length of the sides of Q. Clearly, $Q' \setminus Q$ is an annulus centered at v_1 and Lemma 2.6 can be used to show that $F(Q' \setminus Q)$ holds with probability $1 - e^{-\Omega(\log^6 n/r^4)}$. Thus, all nodes within distance $o(\log^6 n/r^4)$ of v_1 are contained inside the crossings of $Q' \setminus Q$ and their distance to v_1 in G must be smaller than $44c_7^2 \log^{12} n/r^{10}$ by Lemma 2.5. So using Lemma 2.7 we conclude that all nodes within distance $o(\log^6 n/r^4)$ of v_1 are informed after $O(\log^{14} n/r^{10})$ rounds, which is also negligible in comparison to $\Theta(\sqrt{n}/r)$.

Case $r = \omega(\sqrt{\log n})$.

In this section we prove the following lemma, which deals with Theorem 2.3 in the remaining case $r = \omega(\sqrt{\log n})$.

Lemma 2.9. *If $r = \omega(\sqrt{\log n})$, then for all node $s \in G$ we obtain $T_{\text{push}}(s, G) = O(\sqrt{n}/r + \log n)$ with probability $1 - O(n^{-1})$.*

Remark 2.10. We point out that Lemma 2.9 can be generalized to random geometric graphs in higher dimensions. For dimension $d \geq 2$, the lemma holds with $\mathbb{S} = [0, n^{1/d}]^d$ and $r = \omega(\log^{1/d} n)$ as long as d is a constant independent of n.

In order to prove Lemma 2.9, we consider a tessellation of $\mathbb{S}$ into squares of side-length $\min\{r/3, \sqrt{n}/2\}$, which we refer to as "cells." (If $\sqrt{n}$ is not a multiple of $r/3$, then we make the cells in the last row or column of the tessellation be smaller than the others.) It is very easy to verify that nodes in the same cell are neighbors in G and that a node in a given cell can have neighbors in at most 49 different cells. Let $a_{\min}$ be the number of nodes inside the cell that contains the smallest number of nodes, and let $a_{\max}$ be the number of nodes inside the cell that contains the largest number of nodes. Since $r = \omega(\sqrt{\log n})$, a standard Chernoff bound for Poisson random variables (cf. Lemma A.3) can be used to show that there are constants $c_1 < c_2$ such that a fixed cell contains at least $c_1 r^2$ nodes and at most $c_2 r^2$ nodes with probability larger than $1 - n^{-2}$. Using the union bound over the cells of the tessellation, we obtain that $a_{\min}$ and $a_{\max}$ are $\Theta(r^2)$ with probability $1 - O(n^{-1})$.

We are now in position to start our proof of Lemma 2.9. We index the cells by $i \in \mathbb{Z}^2$ and let Z_i be the event that the cell i contains at least one informed node. We say that cells i and j are *adjacent* if and only if they share an edge. Therefore, each cell has exactly four adjacent cells and this adjacency relation induces a 4-regular graph C over the cells.

Given two adjacent cells i and j, at any round of the push algorithm, an informed node in cell i chooses a node from cell j with probability larger than $a_{\min}/(49 a_{\max}) = \Theta(1)$. We want to derive the time until $Z_i = 1$ for all i. Given a path between two cells $j_1, j_2 \in C$, the number of rounds the information takes to be transmitted along this path can be upper bounded by the sum of independent geometric random variables with mean $\Theta(1)$. Applying Lemma A.4, we infer that the number of rounds required to transmit the information from j_1 to j_2 is smaller than $O(\text{diam}(C) + \log n)$ with probability $1 - e^{-\Omega(\text{diam}(C) + \log n)}$. Since there are $O(n/r^2)$ cells and $\text{diam}(C) = O(\sqrt{n}/r)$, we obtain that with probability $1 - O(n^{-1})$, $Z_i = 1$ for all i after $O(\sqrt{n}/r + \log n)$ rounds.

Now, we consider a faulty version of the push algorithm, in which each transmission may fail with probability $p \in [0, 1)$ independently of all other transmissions. Then, a node that is not informed at the beginning of the algorithm can only become informed if it receives the information from a transmission that did not fail. We denote by $T^p_{\text{push}}(s, G)$ the runtime of the faulty version of the push algorithm initiated at node $s \in G$. We use the following relation between $T_{\text{push}}(s, G)$ and $T^p_{\text{push}}(s, G)$.

Lemma 2.11 ([36, Theorem 6])**.** *For any graph G, any node $s \in G$, and any $p \in [0, 1)$, there exists a coupling between $T^p_{\text{push}}(s, G)$ and $T_{\text{push}}(s, G)$ such that*

$$T^p_{\text{push}}(s, G) = O\left(\frac{T_{\text{push}}(s, G)}{1 - p}\right).$$

Assume that each cell contains at least one informed node. We want to bound the number of additional rounds until all nodes in G become informed. Note that each cell constitutes a clique with $\Theta(r^2)$ nodes. According to the push algorithm, at any round, a node chooses a neighbor inside its own cell with probability larger than $a_{\min}/(49a_{\max}) = \Theta(1)$. Therefore, a standard coupling argument can be used to show that the time until all nodes from a given cell get informed can be upper bounded by the time the faulty version of the push algorithm with failure probability $\Theta(1)$ takes to inform all nodes of a complete graph with $\Theta(r^2)$ nodes. Thus, using [42, Theorem 5.2], which establishes the runtime of the push algorithm in complete graphs, and Lemma 2.11 above, we obtain that all nodes of a given cell get informed within $O(\log r^2)$ steps with probability $1 - O(r^{-2})$.

Now we need to extend this result to all cells. For each cell i, let W_i be an independent geometric random variable with parameter ρ (and thus mean $1/\rho$), where we assume $\rho = 1 - O(r^{-2})$. Therefore, once $Z_i = 1$ for all cell i, then the time it takes until all nodes get informed can be upper bounded by $O(\log r^2) \max_i W_i$, where the maximum is taken over all cells. Since we have $\Theta(n/r^2)$ cells, we obtain that all W_i's are smaller than $c \log(n/r^2)$ for some constant c with probability $(1 - (1 - \rho)^{c \log(n/r^2)})^{\Theta(n/r^2)} \geq 1 - O(n^{-1})$ for a proper choice of c. Therefore, we obtain that $T_{\text{push}}(s, G) \leq O(\text{diam}(C) + \log n + \log(r^2) \log(n/r^2)) = O(\sqrt{n}/r + \log n)$, which concludes the proof of Lemma 2.9.

Chapter 3

Mobile Networks

We consider the following random graph model, which was introduced by van den Berg, Meester and White [12] and which we refer to as the *mobile geometric graph* model. Let $\Pi_0 = \{X_i\}_i$ be a Poisson point process on $\mathbb{R}^d$ of intensity λ. To avoid ambiguity, we refer to the points of a point process as *nodes*. We let each node X_i move according to a standard Brownian motion $(\zeta_i(s))_{s\geq 0}$ independently of the other nodes, and set $\Pi_s = \{X_i + \zeta_i(s)\}_i$ to be the point process obtained after the nodes of Π_0 have moved for time s. By standard arguments [12] it follows that Π_s is again a Poisson point process of the same intensity λ.

At any given time s we construct a graph G_s by putting an edge between any two nodes of Π_s that are at distance at most r. In what follows we take r to be an arbitrary but fixed constant. There exists a critical intensity $\lambda_c = \lambda_c(d)$ such that if $\lambda > \lambda_c$, then a.s. there exists a unique infinite connected component in G_s, which we denote by $C_\infty(s)$, while if $\lambda < \lambda_c$ then all connected components are finite a.s. [65, 71].

In this chapter, unless otherwise stated, we take the infinite mobile geometric graph defined in the whole of $\mathbb{R}^d$. We study four problems for this model: detection, coverage, percolation and broadcast. We define these problems and state our results below.

Detection. Consider a "target" particle u which is initially placed at the origin, and whose position at time s is given by $g(s)$, which is a continuous process in $\mathbb{R}^d$. We are interested in the time it takes for u to be *detected* by the mobile geometric graph, namely how long it takes until a node of Π_0 has come within distance at most r from u. More formally we define

$$T_{\text{det}} = \inf\Big\{t \geq 0 : g(t) \in \bigcup_i B(X_i + \zeta_i(t), r)\Big\},$$

where the union is taken over all the nodes $\{X_i\}$ of Π_0 and $B(x, r)$ denotes the ball of radius r centered at x. In Section 3.1 we prove the following theorem, which extends previous classical results on detection for non-mobile particles u [83, 58].

Theorem 3.1. *In two dimensions, for any fixed λ and any g independent of the motions of the nodes of Π_0, we have*

$$\mathbf{P}\left(T_{\text{det}} > t\right) \leq \exp\left(-2\pi\lambda \frac{t}{\log t}(1+o(1))\right).$$

In addition, if g is an independent Brownian motion then the above bound is tight, i.e.,

$$\mathbf{P}\left(T_{\text{det}} > t\right) = \exp\left(-2\pi\lambda \frac{t}{\log t}(1+o(1))\right). \tag{3.1}$$

Remark 3.2. Theorem 3.1 is stated for $d = 2$. In Section 3.1 it is extended to all dimensions $d \geq 1$ (see Theorem 3.12), where the tail of T_{det} is shown to be $\exp(-\Theta(\sqrt{t}))$ for $d = 1$ and $\exp(-\Theta(t))$ for $d \geq 3$. All these results exploit a connection between the detection time and the volume of the Wiener sausage, as explained in Section 3.1.

Remark 3.3. From [83, 58], we have that for $d = 2$ the result in (3.1) also holds when u does *not* move (i.e., $g \equiv 0$). Thus Theorem 3.1 establishes that, asymptotically, the best strategy for a particle u (that is not informed of the motion of the nodes of Π_0) to avoid detection is to stay put. In Section 3.1 we show that this is true for $d = 1$ for any fixed t (not only asymptotically).

Coverage. Let A be a subset of $\mathbb{R}^d$. We are interested in the time it takes for *all* the points of A to be detected. We thus define

$$T_{\text{cov}}(A) = \inf\left\{t \geq 0 : A \subset \bigcup_{s \leq t} \bigcup_i B(X_i + \zeta_i(s), r)\right\}.$$

For $R \in \mathbb{R}_+$, let Q_R be the cube in $\mathbb{R}^d$ of side length R. A natural question proposed by Konstantopoulos [60] is to determine the asymptotics of $\mathbf{E}\left[T_{\text{cov}}(Q_R)\right]$ as $R \to \infty$.

In Section 3.2 we prove Theorem 3.4 below, which gives the asymptotics for the expected time to cover the set Q_R as $R \to \infty$ and shows that $T_{\text{cov}}(Q_R)$ is concentrated around its expectation. We write $f \sim g$ as $x \to \infty$ to mean that $\frac{f(x)}{g(x)} \to 1$ as $x \to \infty$.

Theorem 3.4. *We have that, as $R \to \infty$,*

$$ET_{\text{cov}}(Q_R) \sim \begin{cases} \frac{\pi}{8\lambda^2}(\log R)^2 & \text{for } d = 1 \\ \frac{1}{\lambda\pi} \log R \log\log R & \text{for } d = 2 \\ \frac{d \log R}{\lambda c(d) r^{d-2}} & \text{for } d \geq 3 \end{cases} \quad \text{and} \quad \frac{T_{\text{cov}}(Q_R)}{\mathbf{E}\left[T_{\text{cov}}(Q_R)\right]} \to 1 \text{ in probability,}$$

where $c(d) = \frac{\Gamma\left(\frac{d}{2}-1\right)}{2\pi^{\frac{d}{2}}}$ and Γ stands for the Gamma function.

Remark 3.5. Instead of covering a whole cube, we could ask for the coverage time of other sets. We prove Theorem 3.4 in this general setting in Section 3.2; for instance, we show that $\mathbf{E}[T_{\text{cov}}]$ for a line segment of length R is smaller than $\mathbf{E}[T_{\text{cov}}(Q_R)]$ by a factor of $\frac{1+o(1)}{d}$ and also obtain asymptotics for fractal sets (see Theorem 3.15).

Percolation. Let u be an extra node which is initially at the origin and which moves independently of the nodes of Π_0 according to some function g. We investigate the time it takes until u belongs to the infinite connected component. We denote this time by T_{perc}, which can be more formally written as

$$T_{\text{perc}} = \inf\{t \geq 0 : \exists y \in C_\infty(t) \text{ s.t. } \|g(t) - y\|_2 \leq r\}.$$

The detection time clearly provides a lower bound on the percolation time, so we may deduce from Theorem 3.1 and Remark 3.2 above that, as $t \to \infty$, $\mathbf{P}(T_{\text{perc}} > t)$ is at least $\exp(-O(t/\log t))$ for $d = 2$ and at least $\exp(-O(t))$ for $d \geq 3$, when u is non-mobile or moves according to an independent Brownian motion. We will prove the following stretched exponential upper bound in all dimensions $d \geq 2$ in Section 3.3:

Theorem 3.6. *For all dimensions $d \geq 2$, if $\lambda > \lambda_c(d)$ then there exist constants c and t_0, depending only on d, such that*

$$\mathbf{P}(T_{\text{perc}} > t) \leq \exp\left(-c\frac{\lambda t}{\log^{3+6/d} t}\right), \text{ for all } t \geq t_0.$$

This holds when u is non-mobile or moves according to an independent Brownian motion.

We briefly mention some of the ideas used in the proof, which we believe are of independent interest. The key technical challenge is the dependency of the G_s's over time. To overcome this, we partition $\mathbb{R}^d$ into subregions of suitable size and show via a multi-scale argument that all such subregions contain sufficiently many nodes for a large fraction of the time steps. This is the content of Proposition 3.19. This result allows us to *couple* the evolution of the nodes in each subregion with those of a fresh Poisson point process of slightly smaller intensity $\lambda' < \lambda$ which is still larger than the critical value λ_c. After a number of steps Δ that depends on the size of the subregion, we are able to guarantee that the coupled processes match up almost completely. As a result, we can conclude that there are $\Theta(t/\Delta)$ time steps for which the mobile geometric graph contains an *independent* Poisson point process with intensity $\lambda'' > \lambda_c$. This fact, which we believe is of wider applicability, is formally stated in Proposition 3.18. This independence is sufficient to complete the proof.

Broadcast. As a sample application of Theorem 3.6, we consider the time taken to broadcast a message in finite mobile geometric graph. Consider a mobile geometric graph in a torus $\mathbb{S}$ of volume n. Suppose a message originates at an arbitrary node at time 0, and at each integer time step s each node that has already received the message broadcasts it to all

nodes in the same connected component. (Here we are making the reasonable assumption that the speed of transmission is much faster than the motion of the nodes, so that messages can travel throughout a connected component before it is altered by the motion.) Let $T_{\rm bc}$ denote the time until all nodes have received the message. We prove the following result in Section 3.4.

Corollary 3.7. *In a mobile geometric graph on the torus of volume n with any fixed $\lambda > \lambda_{\rm c}$, the broadcast time $T_{\rm bc}$ is $O(\log n(\log\log n)^{3+6/d})$ w.h.p. in any dimension $d \geq 2$.*

3.1 Detection time

In this section we give the proof of Theorem 3.1. We first state a generalization of a well-known result [83], which we will use in several proofs; we include its proof here for the sake of completeness.

Lemma 3.8. *Suppose that u starts from the origin at time 0 and its position at time s is given by a deterministic function $g(s)$. Let $W_g(t) = \cup_{s\leq t} B(g(s) - \zeta(s), r)$ be the so-called "Wiener sausage with drift" up to time t. Then, for any dimension $d \geq 1$, the detection probability satisfies*

$$\mathbf{P}\left(T_{\rm det} > t\right) = \exp(-\lambda\,\mathbf{E}\left[\mathrm{vol}(W_g(t))\right]),$$

where $\mathrm{vol}\,(A)$ stands for the Lebesgue measure of the set A in $\mathbb{R}^d$.

Proof. Let Φ be the set of points of Π_0 that have detected u by time t, that is

$$\Phi = \{X_i \in \Pi_0 : \exists s \leq t \text{ s.t. } g(s) \in B(X_i + \zeta_i(s), r)\}.$$

Since the ζ_i's are independent we have that Φ is a thinned Poisson point process with intensity given by

$$\Lambda(x) = \lambda\,\mathbf{P}\left(x \in \cup_{s\leq t} B(g(s) - \zeta(s), r)\right),$$

where ζ is a standard Brownian motion.

So for the probability that the detection time is greater than t we have that

$$\mathbf{P}\left(T_{\rm det} > t\right) = \exp(-\lambda \int_{\mathbb{R}^d} \mathbf{P}\left(x \in \cup_{s\leq t} B(g(s) - \zeta(s), r)\right) dx)$$
$$= \exp(-\lambda\,\mathbf{E}\left[\mathrm{vol}(\cup_{s\leq t} B(g(s) - \zeta(s), r))\right]) = \exp(-\lambda\,\mathbf{E}\left[\mathrm{vol}\left(W_g(t)\right)\right]).$$

□

Remark 3.9. The preceding lemma implies that when the motion g of u is random and independent of the motions of the nodes of the Poisson point process Π_0 then

$$\mathbf{P}\left(T_{\rm det} > t\right) = \mathbf{E}\left[\exp(-\lambda\,\mathbf{E}\left[\mathrm{vol}(W_g(t)) \mid (g(s))_{s\leq t}\right])\right].$$

Remark 3.10. We note that the above proof can be easily generalized to show that the time $T_{\text{det}}(K)$ until we detect some point in a compact set $K \subset \mathbb{R}^d$ satisfies

$$\mathbf{P}\left(T_{\text{det}}(K) > t\right) = \exp(-\lambda\, \mathbf{E}\left[\text{vol}\left(\cup_{s\le t}(K^r - \zeta(s))\right)\right]), \tag{3.2}$$

where K^r stands for the r-enlargement of K, i.e. $K^r = \cup_{x\in K} B(x,r)$.

From Lemma 3.8 we see that estimating $\mathbf{P}\left(T_{\text{det}} > t\right)$ translates to deriving estimates for $\mathbf{E}\left[\text{vol}\left(W_g(t)\right)\right]$. When u does not move (i.e., $g \equiv 0$), it is well known that in two dimensions [82, 11]

$$\mathbf{E}\left[\text{vol}\left(W_0(t)\right)\right] = \frac{2\pi t}{\log t}(1 + o(1)).$$

The following Lemma implies that $\mathbf{E}\left[\text{vol}\left(W_g(t)\right)\right] \ge \mathbf{E}\left[\text{vol}\left(W_0(t)\right)\right](1-o(1))$ for any function g that is deterministic and continuous.

Lemma 3.11. *Let ζ be a standard Brownian motion in two dimensions and let g be a deterministic continuous function, $g : \mathbb{R}_+ \to \mathbb{R}^2$. Let $W_g(t) = \cup_{s\le t} B(g(s) - \zeta(s), r)$ be a Wiener sausage with drift g up to time t. We then have that as $t \to \infty$*

$$\mathbf{E}\left[\text{vol}\left(W_g(t)\right)\right] \ge \frac{2\pi t}{\log t}(1 - o(1)).$$

Proof. We may write

$$\mathbf{E}\left[\text{vol}(W_g(t))\right] = \int_{\mathbb{R}^2} \mathbf{P}\left(y \in \cup_{s\le t} B(g(s) - \zeta(s), r)\right) dy = \int_{\mathbb{R}^2} \mathbf{P}\left(\tau_{B(y,r)} \le t\right) dy,$$

where τ_A is the first hitting time of the set A by $g - \zeta$. Define

$$Z_y = \int_0^t \mathbf{1}(g(s) - \zeta(s) \in B(y,r))\, ds,$$

i.e., the time that the process $g - \zeta$ spends in the ball $B(y,r)$ before time t. It is clear by the continuity of $g - \zeta$ that $\{Z_y > 0\} = \{\tau_{B(y,r)} \le t\}$. Clearly $\mathbf{P}\left(Z_y > 0\right) = \frac{\mathbf{E}[Z_y]}{\mathbf{E}[Z_y \mid Z_y > 0]}$ and for the first moment we have

$$\begin{aligned}\mathbf{E}\left[Z_y\right] &= \int_0^t \mathbf{P}\left(g(s) - \zeta(s) \in B(y,r)\right)\, ds \\ &= \int_0^t \int_{B(y,r)} \frac{1}{2\pi s} e^{-\frac{\|z - g(s)\|_2^2}{2s}}\, dz\, ds = \int_0^t \int_{B(0,r)} \frac{1}{2\pi s} e^{-\frac{\|z + y - g(s)\|_2^2}{2s}}\, dz\, ds.\end{aligned}$$

For the conditional expectation $\mathbf{E}[Z_y \mid Z_y > 0]$, if we write τ for the first time before time t that $g - \zeta$ hits the boundary of the ball $B(y,r)$, denoted by $\partial B(y,r)$, then we get

$$\begin{aligned}
\mathbf{E}[Z_y \mid Z_y > 0] &= \mathbf{E}\left[\int_0^{t-\tau} \mathbf{1}(g(s+\tau) - \zeta(s+\tau) \in B(y,r))\, ds\right] \\
&\leq 1 + \mathbf{E}\left[\int_1^{(t-\tau)\vee 1} \mathbf{1}(g(s+\tau) - \zeta(s+\tau) \in B(y,r))\, ds\right] \\
&\leq 1 + \max_{x \in \partial B(y,r)} \int_1^t \int_{B(y,r)} \mathbf{E}\left[\frac{1}{2\pi s} e^{-\frac{\|z - g(s+\tau) - x + g(\tau)\|_2^2}{2s}}\right] dz\, ds \\
&\leq 1 + \int_1^t \int_{B(y,r)} \frac{1}{2\pi s}\, dz\, ds \leq 1 + r^2 \frac{\log t}{2}.
\end{aligned}$$

So, putting everything together we obtain that

$$\begin{aligned}
\mathbf{E}[\mathrm{vol}(W_g(t))] &= \int_{\mathbb{R}^2} \frac{\mathbf{E}[Z_y]}{\mathbf{E}[Z_y \mid Z_y > 0]}\, dy \\
&\geq \frac{\int_0^t \int_{B(0,r)} \left(\int_{\mathbb{R}^2} \frac{1}{2\pi s} e^{-\frac{\|z+y-g(s)\|_2^2}{2s}}\, dy\right) dz\, ds}{1 + r^2 \frac{\log t}{2}} = \frac{2\pi t r^2}{2 + r^2 \log t}
\end{aligned}$$

and hence as $t \to \infty$

$$\mathbf{E}[\mathrm{vol}(W_g(t))] \geq \frac{2\pi t}{\log t}(1 - o(1)).$$

□

We are now ready to prove Theorem 3.1.

Proof of Theorem 3.1. From Remark 3.9 we have

$$\mathbf{P}(T_{\mathrm{det}} > t) = \mathbf{E}[\exp(-\lambda \mathbf{E}[\mathrm{vol}(W_g(t)) \mid (g(s))_{s \leq t}])],$$

where g is independent of ζ. By Lemma 3.11, we have the upper bound

$$\mathbf{P}(T_{\mathrm{det}} > t) \leq \exp\left(-2\pi\lambda \frac{t}{\log t}(1 - o(1))\right), \text{ as } t \to \infty.$$

So it remains to show the lower bound on this probability for the case when g is a standard Brownian motion independent of the motions of the nodes of Π_0. Letting $R = \log t$, it is clear that

$$\mathbf{P}(T_{\mathrm{det}} > t) \geq \mathbf{P}\left(u \text{ stays in } B(0,R) \text{ for all } s \leq t, T_{B(0,R)} > t\right), \tag{3.3}$$

where $T_{B(0,R)}$ is the detection time of the ball $B(0,R)$, i.e.,

$$T_{B(0,R)} = \inf\{s > 0 : \exists i \text{ s.t. } B(X_i + \zeta_i(s), r) \cap B(0,R) \neq \emptyset\}.$$

Since the motions of u and the nodes of Π_0 are independent, we get that

$$\begin{aligned} &\mathbf{P}\left(u \text{ stays in } B(0,R) \text{ for all } s \leq t, T_{B(0,R)} > t\right) \\ &\quad = \mathbf{P}\left(u \text{ stays in } B(0,R) \text{ for all } s \leq t\right) \mathbf{P}\left(T_{B(0,R)} > t\right). \end{aligned} \tag{3.4}$$

From Remark 3.10 with $K = B(0,R)$ we get

$$\mathbf{P}\left(T_{B(0,R)} > t\right) = \exp(-\lambda\, \mathbf{E}\left[\text{vol}\left(\cup_{s\leq t} B(0, R+r) - \zeta(s)\right)\right])$$

and writing $\cup_{s\leq t}(B(0,R+r) - \zeta(s)) = R \cup_{s\leq t} B\left(-\frac{\zeta(s)}{R}, 1 + \frac{r}{R}\right)$ and for R large enough we get that, for all s,

$$B\left(-\frac{\zeta(s)}{R}, 1\right) \subset B\left(-\frac{\zeta(s)}{R}, 1 + \frac{r}{R}\right) \subset B\left(-\frac{\zeta(s)}{R}, 2\right).$$

For any $x > 0$ we have by Brownian scaling that

$$\mathbf{E}\left[\text{vol}\left(\bigcup_{s\leq t} B\left(-\frac{\zeta(s)}{R}, x\right)\right)\right] = \mathbf{E}\left[\text{vol}\left(\bigcup_{s' \leq \frac{t}{R^2}} B\left(\tilde{\zeta}(s'), x\right)\right)\right],$$

where $\tilde{\zeta}$ is a standard Brownian motion. So finally, using the asymptotic expression for the expected volume of the Wiener sausage in two dimensions [11, 82], i.e.,

$$\mathbf{E}\left[\text{vol}\left(\cup_{s\leq t} B(\zeta(s), x)\right)\right] \sim 2\pi \frac{t}{\log t} \quad \text{as } t \to \infty,$$

for any x independent of t, we get that

$$\mathbf{E}\left[\text{vol}\left(\cup_{s\leq t} B(0, R+r) - \zeta(s)\right)\right] \sim R^2 2\pi \frac{t/R^2}{\log t - \log R^2} \sim 2\pi \frac{t}{\log t}.$$

Hence,

$$\mathbf{P}\left(T_{B(0,R)} > t\right) = \exp\left(-2\pi\lambda \frac{t}{\log t}(1 + o(1))\right), \text{ as } t \to \infty.$$

Thus we only need to lower bound the probability that u stays in the ball $B(0,R)$ for all times $s \leq t$.

For any $t \geq R^2$ and any dimension $d \geq 1$, we have by [21] that

$$\mathbf{P}\left(u \text{ stays in } B(0,R) \text{ for all } s \leq t\right) \geq \exp(-ct/R^2), \tag{3.5}$$

for a positive constant c and hence, since $R = \log t$, we get that

$$\begin{aligned}\mathbf{P}\left(T_{\text{det}} > t\right) &\geq \exp\left(-2\pi\lambda\frac{t}{\log t}(1+o(1))\right)\exp\left(-c\frac{t}{(\log t)^2}\right)\\ &= \exp\left(-2\pi\lambda\frac{t}{\log t}(1+o(1))\right).\end{aligned}$$

□

General dimensions and the Wiener sausage

For $d = 1$, the volume of $W_g(t)$ can be computed from the maximum and minimum values of $g - \zeta$ via the formula

$$\text{vol}\left(W_g(t)\right) = 2r + \max_{s\leq t}(g(s) - \zeta(s)) - \min_{s\leq t}(g(s) - \zeta(s)).$$

Let $t^\star$ and $t_\star$ be the random times in the interval $[0, t]$ at which $-\zeta$ achieves its maximum and minimum values respectively. Then we have

$$\begin{aligned}\mathbf{E}\left[\text{vol}\left(W_g(t)\right)\right] &\geq 2r + \mathbf{E}\left[(g(t^\star) - \zeta(t^\star)) - (g(t_\star) - \zeta(t_\star))\right]\\ &= 2r + \mathbf{E}\left[-\zeta(t^\star) + \zeta(t_\star)\right] = \mathbf{E}\left[\text{vol}\left(W_0(t)\right)\right],\end{aligned} \tag{3.6}$$

where $\mathbf{E}\left[g(t^\star)\right] = \mathbf{E}\left[g(t_\star)\right]$ holds since $t^\star$ and $t_\star$ have the same distribution. Thus, for $d = 1$ the inequality $\mathbf{E}\left[\text{vol}\left(W_g(t)\right)\right] \geq \mathbf{E}\left[\text{vol}\left(W_0(t)\right)\right]$ holds for all fixed t; for $d = 2$ Lemma 3.11 gives this inequality only asymptotically as $t \to \infty$.

For dimensions $d \geq 3$, the proof of Lemma 3.11 can be used to obtain the following weaker result: there exists a positive constant c such that

$$\mathbf{E}\left[\text{vol}\left(W_g(t)\right)\right] \geq c\,\mathbf{E}\left[\text{vol}\left(W_0(t)\right)\right]. \tag{3.7}$$

The expected volume of the Wiener sausage with $g \equiv 0$ is known to satisfy [82, 11]

$$V_0(t) = \mathbf{E}\left[\text{vol}\left(W_0(t)\right)\right] = \begin{cases}\sqrt{\frac{8t}{\pi}} + 2r & \text{for } d = 1\\ \frac{2\pi t}{\log t}(1+o(1)) & \text{for } d = 2\\ c(d)r^{d-2}t(1+o(1)) & \text{for } d \geq 3,\end{cases} \tag{3.8}$$

where $c(d) = \frac{\Gamma\left(\frac{d}{2}-1\right)}{2\pi^{d/2}}$. (For $d = 1$ the quantity above follows from well-known results for Brownian motion [68]; namely $\mathbf{E}\left[\max_{s\leq t}\zeta(s)\right] = -\,\mathbf{E}\left[\min_{s\leq t}\zeta(s)\right] = \sqrt{\frac{2t}{\pi}}$.) Hence plugging (3.6–3.8) into Lemma 3.8 gives an upper bound for $\mathbf{P}\left(T_{\text{det}} > t\right)$ when $d = 1$ and $d \geq 3$. Regarding the lower bound, when g is an independent Brownian motion, the same strategy as in the proof of Theorem 3.1 works for $d = 1$ and $d \geq 3$ provided we set R properly. For

$d = 1$ it suffices to take $R = t^{1/3}$; for $d \geq 3$ we can set $R = r$. Then we obtain a positive constant c_1 such that, as $t \to \infty$,

$$\mathbf{P}\left(T_{B(0,R)} > t\right) \geq \begin{cases} \exp\left(-\sqrt{\frac{8t}{\pi}}(1+o(1))\right) & \text{for } d = 1 \\ \exp\left(-c_1 r^{d-2} t(1+o(1))\right) & \text{for } d \geq 3. \end{cases}$$

This together with (3.4) and (3.5) give us the following theorem, which holds in all dimensions $d \geq 1$.

Theorem 3.12. *Let g be a continuous process in $\mathbb{R}^d$, for $d \geq 1$. If g is a deterministic continuous function, then we have*

$$-\frac{1}{\lambda} \log \mathbf{P}\left(T_{\text{det}} > t\right) = \mathbf{E}\left[\text{vol}\left(W_g(t)\right)\right].$$

If g is random but independent of the motion of the nodes of Π_0, we obtain a positive constant β_1 such that as $t \to \infty$

$$-\frac{1}{\lambda} \log \mathbf{P}\left(T_{\text{det}} > t\right) \geq \begin{cases} V_0(t) & \textit{for } d = 1 \\ (1-o(1))V_0(t) & \textit{for } d = 2 \\ (1-o(1))\beta_1 V_0(t) & \textit{for } d \geq 3, \end{cases}$$

where $V_0(t)$ is defined in (3.8).

If g is a standard Brownian motion, then we obtain a positive constant β_2 such that as $t \to \infty$

$$-\frac{1}{\lambda} \log \mathbf{P}\left(T_{\text{det}} > t\right) \leq \begin{cases} (1+o(1))V_0(t) & \textit{for } d = 1, 2 \\ (1+o(1))\beta_2 V_0(t) & \textit{for } d \geq 3. \end{cases}$$

Remark 3.13. Very recently, Peres and Sousi [74] established that

$$\mathbf{E}\left[\text{vol}\left(W_g(t)\right)\right] \geq \mathbf{E}\left[\text{vol}\left(W_0(t)\right)\right]$$

for all dimensions $d \geq 1$ and any fixed $t > 0$. This result was previously known only for particles performing continuous-time random walks on $\mathbb{Z}^d$ [67, 32].

3.2 Coverage time

In this section we will prove a more general version of Theorem 3.4. Let A be a subset of $\mathbb{R}^d$. For $R \in \mathbb{R}_+$ we define the set $RA = \{Ra : a \in A\}$. We recall the definition of Minkowski dimension, which can be found, e.g., in [64].

Definition 3.14. *Let A be a non-empty bounded subset of $\mathbb{R}^d$. For $\epsilon > 0$ let $M(A, \epsilon)$ be the smallest number of balls of radius ϵ needed to cover A:*

$$M(A, \epsilon) = \min \left\{ k : A \subset \bigcup_{i=1}^{k} B(x_i, \epsilon) \text{ for some } x_i \in \mathbb{R}^d \right\}.$$

The Minkowski dimension *of A is defined as*

$$\dim_M(A) = \lim_{\epsilon \to 0} \frac{\log M(A, \epsilon)}{\log \epsilon^{-1}},$$

whenever this limit exists.

We now proceed to state the more general version of Theorem 3.4.

Theorem 3.15. *Let A be a bounded subset of $\mathbb{R}^d$ of Minkowski dimension α. We have that as $R \to \infty$*

$$ET_{\text{cov}}(RA) \sim \begin{cases} \frac{\alpha^2 \pi}{8\lambda^2} (\log R)^2 & \text{for } d = 1 \\ \frac{\alpha}{2\pi\lambda} \log R \log\log R & \text{for } d = 2 \\ \frac{\alpha \log R}{\lambda c(d) r^{d-2}} & \text{for } d \geq 3 \end{cases} \quad \text{and} \quad \frac{T_{\text{cov}}(RA)}{\mathbf{E}\left[T_{\text{cov}}(RA)\right]} \to 1 \text{ in probability,}$$

where $c(d) = \frac{\Gamma\left(\frac{d}{2}-1\right)}{2\pi^{\frac{d}{2}}}$ and Γ stands for the Gamma function.

Proof. In the proof we will drop the dependence on RA from $T_{\text{cov}}(RA)$ and $\mathbf{E}\left[T_{\text{cov}}(RA)\right]$ to simplify the notation. We will carry the proof for the case $d = 2$ only and discuss how to adapt the proof for other dimensions at the end.

Let $M(A, \epsilon) = \min\{k \geq 1 : \exists \quad B_1, \ldots, B_k$ balls of radius ϵ covering $A\}$; then it is easy to see that $M(RA, \epsilon) = M(A, \frac{\epsilon}{R})$. By the assumption that A has Minkowski dimension α, for any $\delta > 0$ we can find ϵ_0 small enough such that

$$\epsilon^{-\alpha+\delta} \leq M(A, \epsilon) \leq \epsilon^{-\alpha-\delta}, \text{ for any } \epsilon < \epsilon_0. \tag{3.9}$$

We will first show that

$$\limsup_{R \to \infty} \frac{\mathbf{E}\left[T_{\text{cov}}\right]}{\frac{\alpha}{2\pi\lambda} \log R \log\log R} \leq 1.$$

To do so, we are going to cover the set RA by $M = M(RA, \epsilon)$ balls of radius $0 < \epsilon < r$.

From (3.9) we get that

$$\left(\frac{\epsilon}{R}\right)^{-\alpha+\delta} \leq M \leq \left(\frac{\epsilon}{R}\right)^{-\alpha-\delta} \quad \text{for } R \text{ sufficiently large.} \tag{3.10}$$

Let Z_t be the number of balls not covered by the nodes at time t. It is clear that $\{T_{\text{cov}} > t\} \subset \{Z_t \geq 1\}$. For the first moment of Z_t we have

$$\mathbf{E}[Z_t] \leq M\,\mathbf{P}\left(\text{a given ball } B(x,\epsilon) \text{ is not covered by time } t\right).$$

The probability that a ball $B(x,\epsilon)$ is covered by time t is lower bounded by the probability that a node of the Poisson point process Π_0 has entered the ball $B(x, r-\epsilon)$ before time t. Hence, $\mathbf{P}\left(B(x,\epsilon) \text{ is not covered by time } t\right)$ is at most the probability that x has not been detected by time t by a mobile geometric graph with radius $r-\epsilon$. From Lemma 3.8 we obtain

$$\mathbf{P}\left(B(x,\epsilon) \text{ is not covered by time } t\right) \leq e^{-\lambda\,\mathbf{E}[\text{vol}(W_{0,r-\epsilon}(t))]},$$

where $W_{z,\rho}(t) = \cup_{s\leq t} B(z+\zeta(s),\rho)$.

We are now prove the **upper bound**. Let $\delta' > 0$ be small. For t large enough we have that (see (3.8))

$$(1-\delta')2\pi\frac{t}{\log t} \leq \mathbf{E}[\text{vol}(W_{0,r-\epsilon}(t))] \leq (1+\delta')2\pi\frac{t}{\log t} \tag{3.11}$$

and hence,

$$\mathbf{E}[Z_t] \leq Me^{-2\pi\lambda(1-\delta')\frac{t}{\log t}}. \tag{3.12}$$

By Markov's inequality we have that $\mathbf{P}(T_{\text{cov}} > t) \leq \mathbf{E}[Z_t]$. Also,

$$\mathbf{E}[T_{\text{cov}}] = \int_0^\infty \mathbf{P}(T_{\text{cov}} > t)\,dt \leq t_*(R) + \int_{t_*(R)}^\infty Me^{-2\pi\lambda(1-\delta')\frac{t}{\log t}}\,dt, \tag{3.13}$$

where $t_*(R)$ satisfies

$$M\exp\left(-2\pi\lambda(1-\delta')\frac{t_*(R)}{\log t_*(R)}\right) = 1 \text{ and } t_*(R) > e, \text{ for sufficiently large } R. \tag{3.14}$$

We claim that the last integral appearing in (3.13) is $o(t_*(R))$. To see this set $c = 2\pi\lambda(1-\delta')$ and use a change of variable, $x = \frac{t}{\log t}$, which gives $\frac{dx}{dt} \geq \frac{1}{2\log t} \geq \frac{1}{4\log x}$ for t large enough. Hence setting $x_* = \frac{t_*(R)}{\log t_*(R)}$ the integral is upper bounded by

$$\int_{x_*}^\infty 4M(\log x)e^{-cx}\,dx \leq c'(1/x_* + \log x_*) = o(t_*(R)).$$

So we finally obtain that

$$\mathbf{E}[T_{\text{cov}}] \leq t_*(R)(1+o(1)), \text{ as } R\to\infty.$$

From (3.14) and (3.10) we get

$$(\alpha-\delta)\log\frac{R}{\epsilon}\leq c\frac{t_*(R)}{\log t_*(R)}\leq(\alpha+\delta)\log\frac{R}{\epsilon}\quad\text{ for } R \text{ large enough,}$$

and thus we conclude that

$$\limsup_{R\to\infty}\frac{\mathbf{E}\left[T_{\text{cov}}\right]}{\frac{\alpha}{2\pi\lambda}\log R\log\log R}\leq 1,$$

which follows by letting δ and δ' go to 0.

We now proceed to show the **lower bound**

$$\liminf_{R\to\infty}\frac{\mathbf{E}\left[T_{\text{cov}}\right]}{\frac{\alpha}{2\pi\lambda}\log R\log\log R}\geq 1.$$

To do so, we are going to use the equivalent definition of Minkowski dimension involving packings [64, Chapter 5]. Letting

$$K(A,\epsilon)=\max\{k\geq 1:\exists\quad B_1,\cdots,B_k \text{ disjoint balls of radius } \epsilon \text{ centered in } A\},$$

it is clear that $K(RA,\epsilon)=K(A,\frac{\epsilon}{R})$. For $\delta>0$ there exist $K=K(RA,1)$ disjoint balls with centers in RA and radius 1 satisfying

$$R^{\alpha-\delta}\leq K\leq R^{\alpha+\delta},\text{ for } R \text{ large enough.}$$

So, for R large enough, we can pack the set RA with points $x_1,\cdots,x_K$ (the centers of the balls) that are at distance at least 2 from each other. Let U_t denote the number of centers $x_1,\cdots,x_K$ that have not been detected by time t. Obviously we have that $\{T_{\text{cov}}>t\}\supset\{U_t\geq 1\}$.

Recall that the Wiener sausage $W_{z,r}(t)=\cup_{s\leq t}B(z+\zeta(s),r)$ in two dimensions satisfies, for $\delta'>0$,

$$(1-\delta')2\pi\frac{t}{\log t}\leq\mathbf{E}\left[\text{vol}\left(W_{z,r}(t)\right)\right]\leq(1+\delta')2\pi\frac{t}{\log t},\text{ for } t \text{ large enough.}\tag{3.15}$$

Let $\epsilon>0$ be small and let $t^*=t^*(R)>e$ satisfy the equation

$$\frac{t^*}{\log t^*}=\frac{\alpha-\epsilon-\delta}{2\pi\lambda(1+\delta')}\log R.\tag{3.16}$$

Applying the second moment method to the random variable U_{t^*} we obtain

$$\mathbf{P}\left(T_{\text{cov}}>t^*\right)\geq\frac{(\mathbf{E}\left[U_{t^*}\right])^2}{\mathbf{E}\left[U_{t^*}^2\right]},$$

so in order to obtain a lower bound for $\mathbf{P}(T_{\text{cov}} > t^*)$ it suffices to lower bound the first moment of U_{t^*} and upper bound its second moment. We will show that $\mathbf{P}(T_{\text{cov}} > t^*) \geq \frac{1}{1+o(1)}$, hence we will get that $\mathbf{E}[T_{\text{cov}}] \geq t^* \frac{1}{1+o(1)}$.

We have that $\mathbf{E}[U_{t^*}] = \sum_{i=1}^{K} \mathbf{P}(x_i \text{ not detected by time } t^*)$, and using Lemma 3.8 we obtain that

$$\mathbf{P}(x_i \text{ not detected by time } t^*) = \exp(-\lambda \mathbf{E}[\text{vol}(W_{x_i})]),$$

where $W_x = W_{x,r}(t^*) = \bigcup_{s \leq t^*} B(x + \zeta(s), r)$.

Obviously $\mathbf{E}[\text{vol}(W_x)]$ is independent of x, and hence we get that

$$\mathbf{E}[U_{t^*}] = K \exp(-\lambda \mathbf{E}[\text{vol}(W_0)]). \tag{3.17}$$

Now, for the second moment of U_{t^*} we have

$$\mathbf{E}\left[U_{t^*}^2\right] = \sum_{i=1}^{K} \sum_{j \neq i} \mathbf{P}(x_i, x_j \text{ not detected by time } t^*) + \mathbf{E}[U_{t^*}] \tag{3.18}$$

and using Remark 3.10 we get that

$$\mathbf{P}(x_i, x_j \text{ not detected by time } t^*) = \exp(-\lambda \mathbf{E}\left[\text{vol}(W_{x_i} \cup W_{x_j})\right]).$$

(Note that the two Wiener sausages W_{x_i} and W_{x_j} use the same driving Brownian motion.)

Writing

$$\text{vol}(W_{x_i} \cup W_{x_j}) = \text{vol}(W_{x_i}) + \text{vol}(W_{x_j}) - \text{vol}(W_{x_i} \cap W_{x_j}),$$

equation (3.18) becomes

$$\mathbf{E}\left[U_{t^*}^2\right] \leq \exp(-2\lambda \mathbf{E}[\text{vol}(W_0)]) \sum_{i=1}^{N} \sum_{j \neq i} \exp(\lambda \mathbf{E}\left[\text{vol}(W_{x_i} \cap W_{x_j})\right]) + \mathbf{E}[U_{t^*}]. \tag{3.19}$$

Thus it remains to upper bound $\mathbf{E}\left[\text{vol}(W_{x_i} \cap W_{x_j})\right]$ for all i and j. If $\|x_i - x_j\|_2 \leq (\log R)^2$, then we may use the bound $\text{vol}(W_{x_i} \cap W_{x_j}) \leq \text{vol}(W_{x_i})$.

Recall from (3.16) that $t^*(R) = \Theta(\log R \log\log R)$. The idea is that if x_i and x_j are at distance greater than $(\log R)^2$ apart, then it is very unlikely that the two sets W_{x_i} and W_{x_j} will intersect. Specifically, when $\|x_i - x_j\|_2 \geq (\log R)^2$ it is easy to see that the probability that the two sausages, W_{x_i} and W_{x_j}, intersect is smaller than the probability that a 2-dimensional Brownian motion has traveled distance greater than $\frac{1}{2}(\log R)^2$ in t^* time steps, and this last probability is bounded above by $ce^{-c(\log R)^2}$ by the standard bound for the tail of a Gaussian.

When $\|x_i - x_j\|_2 \geq (\log R)^2$, writing $S_1 = B(x_i, R)$ for the ball of radius R centered at x_i and defining inductively $S_k = B(x_i, 2^{k-1}R) \setminus B(x_i, 2^{k-2}R)$ for all $k \geq 2$, we can split the

volume of $W_{x_i} \cap W_{x_j}$ as follows:

$$\begin{aligned}\mathbf{E}\left[\mathrm{vol}(W_{x_i} \cap W_{x_j})\right] &= \sum_{n=1}^{\infty} \mathbf{E}\left[\mathrm{vol}(W_{x_i} \cap W_{x_j} \cap S_n)\right] \\ &\leq c\pi R^2 e^{-c(\log R)^2} + \sum_{n=2}^{\infty} c' 2^{2n} R^2 e^{-c' 2^{2n} R} \leq c'' R^{-M},\end{aligned}$$

where c, c', c'' and M are all positive constants. The first part of the first inequality follows from the discussion above, namely that if the intersection is nonempty, then the Brownian motion must have traveled distance greater than $\frac{1}{2}(\log R)^2$ in less than t^* steps. If this has happened, then we simply bound the intersection of the two Wiener sausages in the ball $B(x, R)$ by the volume of the ball. The second part of the first inequality follows by the same type of argument, since now in order to have a nonempty intersection in the set S_n, the Brownian motion must have traveled distance at least $2^{n-2}R$ in less than t^* steps, which again is exponentially small.

Finally, the sum appearing in (3.19) is bounded above by

$$\sum_{i=1}^{K} \sum_{j \neq i} \mathbf{1}(\|x_i - x_j\|_2 \leq (\log R)^2) e^{\lambda \mathbf{E}[\mathrm{vol}(W_0)]} + \sum_{i=1}^{K} \sum_{j \neq i} \mathbf{1}(\|x_i - x_j\|_2 > (\log R)^2) e^{c'' \lambda R^{-M}}. \tag{3.20}$$

By (3.15) and the definition of t^* given in (3.16) we get that (3.20) is bounded from above by

$$c_1 K (\log R)^4 R^{\alpha - \delta - \epsilon} + K^2 e^{c'' \lambda R^{-M}}$$

and hence

$$\mathbf{E}\left[U_{t^*}^2\right] \leq \exp(-2\lambda \mathbf{E}\left[\mathrm{vol}(W_0)\right]) \left(c_1 K (\log R)^4 R^{\alpha - \delta - \epsilon} + K^2 e^{c'' \lambda R^{-M}}\right) + \mathbf{E}\left[U_{t^*}\right]. \tag{3.21}$$

Therefore, putting all the estimates together we get that

$$\mathbf{P}\left(U_{t^*} > 0\right) \geq \frac{(\mathbf{E}\left[U_{t^*}\right])^2}{\mathbf{E}\left[U_{t^*}^2\right]} \geq \frac{1}{e^{c'' \lambda R^{-M}} + \frac{1}{K}((\log R)^4 R^{\alpha - \delta - \epsilon} + e^{\lambda \mathbf{E}[\mathrm{vol}(W_0)]})}.$$

Using the lower bound $K \geq R^{\alpha - \delta}$, the upper bound for the expected volume from (3.15) and the definition (3.16) of t^*, we deduce that

$$\mathbf{P}\left(U_{t^*} > 0\right) \geq \frac{1}{1 + o(1)},$$

and thus $\mathbf{E}\left[T_{\mathrm{cov}}\right] \geq \frac{1}{1+o(1)} t^*$. Since t^* satisfies (3.16), we deduce that

$$\liminf_{R \to \infty} \frac{\mathbf{E}\left[T_{\mathrm{cov}}\right]}{\frac{\alpha}{2\pi\lambda} \log R \log\log R} \geq \frac{1 - \frac{\epsilon}{\alpha} - \frac{\delta}{\alpha}}{1 + \delta'}$$

and hence, letting ϵ, δ and δ' go to 0, we get that

$$\liminf_{R\to\infty} \frac{\mathbf{E}\left[T_{\text{cov}}\right]}{\frac{\alpha}{2\pi\lambda}\log R \log\log R} \geq 1.$$

So, we have shown that

$$\mathbf{E}\left[T_{\text{cov}}\right] \sim \frac{\alpha}{2\pi\lambda}\log R \log\log R, \text{ as } R\to\infty. \tag{3.22}$$

Now, for $d = 2$ it only remains to show the last part of the theorem, namely that $\frac{T_{\text{cov}}}{\mathbf{E}[T_{\text{cov}}]}$ converges to 1 in probability as $R \to \infty$. For any $\gamma > 0$ we have that

$$\begin{aligned}\mathbf{P}\left(\left|\frac{T_{\text{cov}}}{\mathbf{E}\left[T_{\text{cov}}\right]} - 1\right| > \gamma\right) &= \mathbf{P}\left(T_{\text{cov}} > (1+\gamma)\,\mathbf{E}\left[T_{\text{cov}}\right]\right) + \mathbf{P}\left(T_{\text{cov}} < (1-\gamma)\,\mathbf{E}\left[T_{\text{cov}}\right]\right) \\ &\leq \mathbf{E}\left[Z_{(1+\gamma)\,\mathbf{E}[T_{\text{cov}}]}\right] + \mathbf{P}\left(T_{\text{cov}} < (1-\gamma)\,\mathbf{E}\left[T_{\text{cov}}\right]\right).\end{aligned}$$

From (3.12) and the definition of M we have that

$$\mathbf{E}\left[Z_t\right] \leq R^{\alpha+\delta}\exp\left(-2\pi\lambda(1-\delta')\frac{t}{\log t}\right).$$

Plugging in $t = (1+\gamma)\,\mathbf{E}\left[T_{\text{cov}}\right]$, using (3.22) and taking δ' sufficiently small gives that

$$\mathbf{E}\left[Z_{(1+\gamma)\,\mathbf{E}[T_{\text{cov}}]}\right] \to 0, \text{ as } R\to\infty.$$

For $\epsilon, \delta, \delta'$ small enough we get that $(1-\gamma)\,\mathbf{E}\left[T_{\text{cov}}\right] < t^*$, so

$$\mathbf{P}\left(T_{\text{cov}} < (1-\gamma)\,\mathbf{E}\left[T_{\text{cov}}\right]\right) \leq \mathbf{P}\left(T_{\text{cov}} \leq t^*\right) \leq \mathbf{P}\left(U_{t^*} = 0\right) \leq 1 - \frac{\left(\mathbf{E}\left[U_{t^*}\right]\right)^2}{\mathbf{E}\left[U_{t^*}^2\right]} = o(1).$$

Hence we get the desired result that

$$\mathbf{P}\left(\left|\frac{T_{\text{cov}}}{\mathbf{E}\left[T_{\text{cov}}\right]} - 1\right| > \gamma\right) \to 0, \text{ as } R\to\infty.$$

For dimensions $d \neq 2$, the same arguments carry through by employing the proper expression for the expected volume of the Wiener sausage given in (3.8). Then, we need to set $t_\star(R)$ and $t^\star(R)$ correspondingly. From (3.14) and (3.16), it suffices to set $t_\star$ to satisfy

$$\begin{aligned}\exp\left(\lambda(1-\delta')\sqrt{\frac{8t_\star(R)}{\pi}}\right) &= M \quad \text{for } d = 1 \\ \exp\left(\lambda(1-\delta')c(d)r^{d-2}t_\star(R)\right) &= M \quad \text{for } d \geq 3,\end{aligned}$$

and $t^\star$ to satisfy

$$\lambda(1+\delta')\sqrt{\frac{8t^\star(R)}{\pi}} = (\alpha-\epsilon-\delta)\log R \quad \text{for } d=1$$
$$\lambda(1+\delta')c(d)r^{d-2}t^\star(R) = (\alpha-\epsilon-\delta)\log R \quad \text{for } d\geq 3.$$

□

Remark 3.16. While the limit defining Minkowski dimension in Definition 3.14 may not exist, the corresponding lim sup is denoted by $\overline{\dim}_M(A)$ and always exists. The proof of Theorem 3.15 also shows that for $d=2$

$$\limsup_{R\to\infty} \frac{\mathbf{E}\left[T_{\text{cov}}(RA)\right]}{\frac{1}{2\pi\lambda}\log R\log\log R} = \overline{\dim}_M(A)$$

and similarly for lim inf and other dimensions.

Remark 3.17. The estimates in the proof of Theorem 3.15 actually imply that, as $R\to\infty$, we have $\frac{T_{\text{cov}}(RA)}{\mathbf{E}[T_{\text{cov}}(RA)]}\to 1$.

3.3 Percolation time

In this section we give the proof of Theorem 3.6. We will observe the process $(G_i)_{i\geq 0}$ in discrete time steps $i=0,1,\ldots$ in order to be able to apply a multi-scale argument. For a nonnegative integer i we define the event J_i that u does not belong to the infinite component at time i; more formally,

$$J_i = \{u \notin \cup_{y\in C_\infty(i)} B(y,r)\}.$$

Then it is easy to see that, for all t, we have

$$\mathbf{P}\left(T_{\text{perc}} > t\right) \leq \mathbf{P}\left(\cap_{i=0}^{\lfloor t\rfloor} J_i\right).$$

We define Q_L to be the cube with side length L centered at the origin and with sides parallel to the axes of $\mathbb{R}^d$. We tessellate Q_L into subcubes of side length $\ell < L$, which we call *cells*. We now state two key propositions that lie at the heart of our argument.

The first proposition says that, provided every cell of the tessellation contains sufficiently many nodes, then we can couple the positions of these nodes after sufficiently many steps with the nodes of an independent Poisson point process of only slightly smaller intensity on a smaller cube. We prove this proposition in Section 3.3.1.

Proposition 3.18. *Fix $K > \ell > 0$ and consider the cube Q_K tessellated into cells of side length ℓ. Let Φ_0 be an arbitrary point process at time 0 that contains at least $\beta\ell^d$ nodes at each cell of the tessellation for some $\beta > 0$. Let Φ_Δ be the point process obtained at time Δ*

from Φ_0 after the nodes have moved according to standard Brownian motion for time Δ. Fix $\epsilon \in (0,1)$ and let Ξ be an independent Poisson point process with intensity $(1-\epsilon)\beta$. Then there exists a coupling of Ξ and Φ_Δ and constants c_1, c_2, c_3 depending only on d such that, if $\Delta \geq \frac{c_1 \ell^2}{\epsilon^2}$ and $K' \leq K - c_2\sqrt{\Delta \log \epsilon^{-1}} > 0$, then the nodes of Ξ are a subset of the nodes of Φ_Δ inside the cube $Q_{K'}$ with probability at least

$$1 - \frac{K^d}{\ell^d} \exp(-c_3 \epsilon^2 \beta \ell^d).$$

The second proposition, which we prove in Section 3.3.2, says that the above condition that each cell contains sufficiently many nodes is satisfied at an arbitrary constant fraction of time steps with high probability.

Proposition 3.19. *Let $t > 0$ be a sufficiently large integer and $\xi, \epsilon \in (0,1)$ be two constants. Suppose that the cube Q_L, for $L = t$, is tessellated into cells of side length ℓ, where $\ell^d \geq C \log^3 t$ for some sufficiently large constant C. For $i = 0, 1, \ldots$ let*

$$A_i = \{\text{at time } i \text{ all cells contain } \geq (1-\xi)\lambda \ell^d \text{ nodes of } \Pi_i\}.$$

Then there exists a positive constant c such that

$$\mathbf{P}\left(\sum_{i=0}^{t-1} \mathbf{1}(A_i) \geq (1-\epsilon)t\right) \geq 1 - \exp\left(-c \frac{\lambda t}{\log^{3+6/d} t}\right). \tag{3.23}$$

Proof of Theorem 3.6. Let u be a node that is at the origin at time 0 independent of the nodes of Π_0. We assume that u is non-mobile; the proof can easily be extended to mobile u using translated cubes that track the motion of u as in [80, Section 4].

Let t be an integer sufficiently large. We consider the cube Q_L, for $L = t$. Set H_t to be the event that u has *never* been in the infinite component from time 0 to $t-1$. More formally, we define

$$H_t = \cap_{i=0}^{t-1} J_i = \cap_{i=0}^{t-1} \{u \notin \cup_{y \in C_\infty(i)} B(y, r)\}.$$

We say that a cube Q_L has a *crossing component* at a given time i if among the nodes in Q_L there exists a connected component that has a path connecting each pair of opposite faces of Q_L. (A path *connects* two faces of Q_L if for each face there is at least one node of the path within distance r of the face.) We then define $\tilde{H}_t$ to be the event that u has never been within distance r of a crossing component of Q_L from time 0 to $t-1$. Let K_t be the event that, in each step from 0 to $t-1$, there exists a unique crossing component of Q_L and it intersects the infinite component. Therefore, if K_t holds and u belongs to a crossing component of Q_L at some time step from 0 to $t-1$, then at the same time step u will also belong to the infinite component. We can then conclude that $K_t \cap \tilde{H}_t^c \subseteq H_t^c$, which gives

$$\mathbf{P}(H_t) \leq \mathbf{P}(K_t^c \cup \tilde{H}_t) \leq \mathbf{P}(K_t^c) + \mathbf{P}(\tilde{H}_t).$$

By [72, Theorems 1 and 2] and by taking the union bound over all time steps, we have

$$\mathbf{P}\left(K_t^{\mathrm{c}}\right) \leq t \exp(-c_1 L). \tag{3.24}$$

We will now derive an upper bound for $\mathbf{P}(\tilde{H}_t)$. Let $\xi > 0$ be a sufficiently small constant such that $(1-\xi)\lambda > \lambda_{\mathrm{c}}$. Take the cube Q_{2L} and tessellate it into cells of side length ℓ, where $\ell = C_1 \log^{3/d} t$, for C_1 a sufficiently large constant in order to satisfy the assumptions of Proposition 3.19. Call a cell *dense* if it contains more than $(1-\xi)\lambda\ell^d$ nodes. For $\delta > 0$, let D be the event that all cells inside Q_{2L} are dense for at least $(1-\delta)t$ time steps. Applying Proposition 3.19 we obtain a constant c_2 such that

$$\mathbf{P}\left(D\right) \geq 1 - \exp\left(-c_2 \frac{\lambda t}{\log^{3+6/d} t}\right).$$

We use the event D to obtain an upper bound for $\mathbf{P}(\tilde{H}_t)$ via

$$\mathbf{P}(\tilde{H}_t) \leq \mathbf{P}(\tilde{H}_t \cap D) + \mathbf{P}(D^{\mathrm{c}}) \leq \mathbf{P}(\tilde{H}_t \cap D) + \exp\left(-c_2 \frac{\lambda t}{\log^{3+6/d} t}\right). \tag{3.25}$$

On the event D, by definition, we can find a collection S of $(1-\delta)t$ time steps for which all cells of side length ℓ are dense inside the cube Q_{2L}. We set $\Delta = C_2\ell^2$ for some sufficiently large constant C_2. We define τ_1 as the first time step for which all cells of Q_{2L} are dense. We now define τ_{i+1} recursively as the first time step after $\tau_i + \Delta$ for which all cells are dense. Obviously, $\tau_1 < \tau_2 < \cdots$ and if we take $k = c_3 t/\Delta = c_3' t/\log^{6/d} t$ for some constant c_3, then we can ensure that on D we have $\tau_k \leq t-1$.

For each i, let A_i be the event that u does *not* belong to a crossing component of Q_L at time $\tau_i + \Delta$. Since when D holds we have $\tau_k \leq t-1$, we can write

$$\mathbf{P}(\tilde{H}_t \cap D) \leq \mathbf{P}\left(\bigcap_{i=1}^{k} A_i \cap D\right). \tag{3.26}$$

For each i, let $\mathcal{F}_i$ be the σ-field induced by the locations of the nodes of Π_0 from time 0 to τ_i. We now claim that for t sufficiently large there exists a positive constant c_4 such that

$$\mathbf{P}\left(A_i \mid \mathcal{F}_i\right) < e^{-c_4}. \tag{3.27}$$

We will define two events E_1, E_2 such that for any $F \in \mathcal{F}_i$ we have

$$\mathbf{P}\left(A_i \mid F\right) \leq \mathbf{P}\left(E_1^{\mathrm{c}} \mid F\right) + \mathbf{P}\left(E_2^{\mathrm{c}} \mid F\right). \tag{3.28}$$

Take $\epsilon > 0$ sufficiently small so that $(1-\epsilon)(1-\xi)\lambda > \lambda_{\mathrm{c}}$, and let Ξ be an independent Poisson point process of intensity $(1-\epsilon)(1-\xi)\lambda$. We define the events

$$E_1 = \{u \text{ belongs to a crossing component of } \Xi \text{ in } Q_L\} \text{ and}$$
$$E_2 = \{\exists \text{ a coupling of } \Xi \text{ and } \Pi_{\tau_i+\Delta} \text{ so that } \Xi \subset \Pi_{\tau_i+\Delta} \text{ in } Q_L\},$$

where "$\Xi \subset \Pi_{\tau_i+\Delta}$ in Q_L" means that the nodes of Ξ that lie inside the cube Q_L are a subset of the nodes of $\Pi_{\tau_i+\Delta}$.

Note that when E_1 and E_2 both hold, then u belongs to a crossing component of Q_L at time $\tau_i + \Delta$, which implies that A_i does not hold. Since the intensity of Ξ is strictly larger than λ_c and E_1 is independent of F by construction, we obtain $\mathbf{P}(E_1 \mid F) \geq c_5$ for some constant $c_5 \in (0,1)$ by [72, Theorem 1].

All cells are dense at time τ_i, by the definition of τ_i. Taking K and K' appearing in Proposition 3.18 to be $K = 2L$ and $K' = L$, we see by the choice of Δ that for large enough t the condition for K' in Proposition 3.18 is satisfied and thus we obtain, uniformly over all $F \in \mathcal{F}_i$, that, for a positive constant c_6,

$$\mathbf{P}(E_2^c \mid F) \leq \exp\left(-c_6 \lambda \log^3 t\right).$$

Plugging everything into (3.28) we get

$$\mathbf{P}(A_i \mid F) \leq 1 - c_5 + \exp\left(-c_6 \lambda \log^3 t\right),$$

which can be made strictly smaller than 1 by taking t sufficiently large. This establishes (3.27).

Note that by definition we have $\tau_i + \Delta < \tau_{i+1}$ for all i, which gives $A_i \in \mathcal{F}_{i+1}$. We can write (3.26) as

$$\mathbf{P}(\tilde{H}_t \cap D) \leq \mathbf{P}\left(\cap_{i=1}^k A_i\right) = \prod_{i=2}^k \mathbf{P}\left(A_i \mid \cap_{j=1}^{i-1} A_j\right) \mathbf{P}(A_1),$$

which by (3.27) translates to

$$\mathbf{P}(\tilde{H}_t \cap D) \leq \exp(-c_4 k) \leq \exp\left(-c_7 \frac{t}{\log^{6/d} t}\right),$$

for a positive constant c_7. Plugging this into (3.25) concludes the proof of Theorem 3.6. □

3.3.1 Coupling

In this section we give the proof of Proposition 3.18. We begin by stating and proving a small technical lemma that will be used in the proof.

Lemma 3.20. *Assume $\epsilon \in (0,1)$ and $\rho > 0$. Let $\Delta \geq 16 d^2 \rho^2 / \epsilon^2$ and $R \geq 2\sqrt{d\Delta \log(8d\epsilon^{-1})}$. Define*

$$g(z) = \frac{1}{(2\pi\Delta)^{d/2}} \exp\left(-\frac{(\|z\|_2 + \rho)^2}{2\Delta}\right)$$

on $\mathbb{R}^d$. Then we have

$$\int_{B(0,R)} g(z)\, dz \geq 1 - \epsilon/2.$$

Proof. Let $\psi(x) = \frac{1}{(2\pi\Delta)^{1/2}} \exp\left(-\frac{(|x|+\rho)^2}{2\Delta}\right)$, for $x \in \mathbb{R}$.
Note that $\sum_{i=1}^{d}(|z_i|+\rho)^2 = \|z\|_2^2 + 2\rho\|z\|_1 + \rho^2 d \geq (\|z\|_2 + \rho)^2$, so

$$\prod_{i=1}^{d} \psi(z_i) \leq g(z), \text{ for } z = (z_1, \ldots, z_d) \in \mathbb{R}^d. \tag{3.29}$$

Next observe that

$$\int_{-\infty}^{\infty} \psi(x)\, dx = 1 - \int_{-\rho}^{\rho} \frac{1}{(2\pi\Delta)^{1/2}} \exp\left(-\frac{y^2}{2\Delta}\right) dy \geq 1 - \frac{2\rho}{\sqrt{2\pi\Delta}} \geq 1 - \frac{\rho}{\sqrt{\Delta}} \geq 1 - \frac{\epsilon}{4d}.$$

By the Gaussian tail bound we have that

$$\int_{R/\sqrt{d}}^{\infty} \psi(x)\, dx \leq \exp\left(-\frac{R^2}{2d\Delta}\right) \leq \frac{\epsilon^2}{64d^2} \leq \frac{\epsilon}{8d},$$

for any $\epsilon \in (0,1)$. Thus $\int_{-R/\sqrt{d}}^{R/\sqrt{d}} \psi(x)\, dx \geq 1 - \frac{\epsilon}{2d}$. Since $[-R/\sqrt{d}, R/\sqrt{d}]^d \subset B(0,R)$, we deduce from (3.29) that

$$\int_{B(0,R)} g(z)\, dz \geq \int_{[-R/\sqrt{d}, R/\sqrt{d}]^d} \prod_{i=1}^{d} \psi(z_i)\, dz \geq \left(1 - \frac{\epsilon}{2d}\right)^d \geq 1 - \epsilon/2.$$

□

We now proceed to the proof of Proposition 3.18.

Proof of Proposition 3.18. We will construct Ξ via three Poisson point processes. We start by defining Ξ_0 as a Poisson point process over Q_K with intensity $(1-\epsilon/2)\beta$. Recall that Φ_0 has at least $\beta\ell^d$ nodes in each cell of Q_K. Then, in any fixed cell, Ξ_0 has fewer nodes than Φ_0 if Ξ_0 has less than $\beta\ell^d$ nodes in that cell, which by a standard Chernoff bound (cf. Lemma A.3) occurs with probability larger than $1 - \exp\left(-\frac{\epsilon'^2(1-\epsilon/2)\beta\ell^d}{2}(1-\epsilon'/3)\right)$ for ϵ' such that $(1+\epsilon')(1-\epsilon/2) = 1$. Since $\epsilon \in (0,1)$ we have $\epsilon' \in (\epsilon/2, 1)$, and the probability above can be bounded below by $1 - \exp\left(-c\epsilon^2\beta\ell^d\right)$ for some constant $c = c(d)$. Let $\{\Xi_0 \preceq \Phi_0\}$ be the event that Ξ_0 has fewer nodes than Φ_0 in every cell of Q_K. Using the union bound over cells we obtain

$$\mathbf{P}\left(\Xi_0 \preceq \Phi_0\right) \geq 1 - \frac{K^d}{\ell^d} \exp(-c\epsilon^2\beta\ell^d). \tag{3.30}$$

If $\{\Xi_0 \preceq \Phi_0\}$ holds, then we can map each node of Ξ_0 to a unique node of Φ_0 in the same cell. We will now show that we can couple the motion of the nodes in Ξ_0 with the motion of their respective pairs in Φ_0 so that the probability that an arbitrary pair is at the same location at time Δ is sufficiently large.

To describe the coupling, let v' be a node from Ξ_0 located at $y' \in Q_K$, and let v be the pair of v' in Φ_0. Let y be the location of v in Q_K, and note that since v and v' belong to the same cell we have $\|y - y'\|_2 \leq \sqrt{d}\ell$. We will construct a function $g(z)$ that is smaller than the densities for the motions of v and v' to the location $y' + z$, uniformly for $z \in \mathbb{R}^d$. That is,

$$g(z) \leq \frac{1}{(2\pi\Delta)^{d/2}} \exp\left(-\frac{\max\{\|z\|_2^2, \|y' + z - y\|_2^2\}}{2\Delta}\right) \tag{3.31}$$

for all $z \in \mathbb{R}^d$.

To this end we set

$$g(z) = \frac{1}{(2\pi\Delta)^{d/2}} \exp\left(-\frac{(\|z\|_2 + \sqrt{d}\ell)^2}{2\Delta}\right). \tag{3.32}$$

Note that this definition satisfies (3.31) since by the triangle inequality $\|y' + z - y\|_2 \leq \|y' - y\|_2 + \|z\|_2$ and $\|y' - y\|_2 \leq \sqrt{d}\ell$. Define $\psi = 1 - \int_{\mathbb{R}^d} g(z)\, dz$. Then, with probability $1 - \psi$ we can use the density function $\frac{g(z)}{1-\psi}$ to sample a single location for the position of both v and v' at time Δ, and then set Ξ'_0 to be the Poisson point process with intensity $(1-\psi)(1-\epsilon/2)\beta$ obtained by *thinning* Ξ_0 (i.e., deleting each node of Ξ_0 with probability ψ). At this step we have crucially used the fact that the function $g(z)$ in (3.32) is oblivious of the location of v and, consequently, is independent of the point process Φ_0. (If one were to use the maximal coupling suggested by (3.31), then the thinning probability would depend on Φ_0, and Ξ'_0 would not be a Poisson point process.)

Let Ξ'_Δ be obtained from Ξ'_0 after the nodes have moved according to the density function $\frac{g(z)}{1-\psi}$. Thus we are assured that the nodes of the Poisson point process Ξ'_Δ are a subset of the nodes of Φ_Δ and are independent of the nodes of Φ_0, where Φ_Δ is obtained by letting the nodes of Φ_0 move from time 0 to time Δ.

By Lemma 3.20 we get that if Δ and $K - K'$ are large enough, then the integral of $g(z)$ inside the ball $B = B(0, (K - K')/2)$ is larger than $1 - \epsilon/2$. (We are interested in the ball B since for all $z \in Q_{K'}$ we have $z + B \subset Q_K$.)

When $\{\Xi_0 \preceq \Phi_0\}$ holds, Ξ'_Δ consists of a subset of the nodes of Φ_Δ. Note that Ξ'_Δ is a *non-homogeneous* Poisson point process over Q_K. It remains to show that the intensity of Ξ'_Δ is strictly larger than $(1 - \epsilon)\beta$ in $Q_{K'}$ so that Ξ can be obtained from Ξ'_Δ via thinning; since Ξ'_Δ is independent of Φ_0, so is Ξ.

For $z \in \mathbb{R}^d$, let $\mu(z)$ be the intensity of Ξ'_Δ. Since Ξ'_0 has no node outside Q_K, we obtain for any $z \in Q_{K'}$,

$$\mu(z) \geq (1 - \psi)(1 - \epsilon/2)\beta \int_{z+B} \frac{g(z - x)}{1 - \psi}\, dx = (1 - \epsilon/2)\beta \int_B g(x)\, dx,$$

where the inequality follows since $z + B \subset Q_K$ for all $z \in Q_{K'}$. From Lemma 3.20, choosing the constants c_1 and c_2 sufficiently large we have $\int_B g(x)\, dx \geq 1 - \epsilon/2$. We then obtain

$\mu(z) \geq (1-\epsilon/2)^2\beta \geq (1-\epsilon)\beta$, which is the intensity of Ξ. Therefore, when $\{\Xi_0 \preceq \Phi_0\}$ holds, which occurs with probability given by (3.30), the nodes of Ξ are a subset of the nodes of Φ_Δ, which completes the proof of Proposition 3.18. □

3.3.2 Density

In this section we prove Proposition 3.19 using a multi-scale argument. Since the argument is rather involved, we begin with a high-level overview.

Proof overview

Our goal is to show that if we tessellate the cube Q_L, with $L = t$, into cells of volume of order $(\log t)^c$, then the probability that all cells contain sufficiently many nodes for a fraction $1-\epsilon$ of the time steps is at least the expression given in Proposition 3.19.

We start at scale 1 with the cube Q_{L_1} where $L_1 > L$. We tessellate Q_{L_1} into cells that are so large that we can easily show that with very high probability during *all* time steps all these cells contain sufficiently many nodes. We refer to this as the event that "the density condition is satisfied at all steps for scale 1." Then, when going from scale $j-1$ to scale j, we take a smaller cube Q_{L_j} with $L_j < L_{j-1}$, and tessellate it into cells that are smaller than the cells at the previous scale (see Figure 3.1). We define the density condition for scale j at a given time step as the event that all the cells at scale j contain a number of nodes that is sufficiently large but strictly smaller than the one used for the density condition for scale $j-1$. Since this density requirement becomes less strict when going from scale $j-1$ to scale j, we will be able to show that the density condition for scale j is satisfied for a large fraction of the time steps at which the density condition is satisfied for scale $j-1$. We repeat this procedure until we obtain, at the last scale, the cube Q_L and cells of side length ℓ.

The importance of the multi-scale approach is that it allows us to recover quickly from instances of low density, i.e., if the density condition holds in scale $j-1$ but fails (at some time) in scale j, there are enough nodes nearby to recover density shortly thereafter.

We now proceed to the detailed argument.

Full proof

Let κ be the number of scales; we will see in a moment that $\kappa = O(\log t)$ will suffice. Let $L_1 > L_2 > \cdots > L_\kappa = L$ such that $L_1 = t^2$ and $L_\kappa = t$.

Let $\ell_1 > \ell_2 > \cdots > \ell_\kappa = \ell$. At scale j, we consider the cube Q_{L_j} and tessellate it into cells of side length ℓ_j (see Figure 3.1(b–c)). We say that a cell is *dense* at a given time step for scale j if it contains more than $(1-\xi_j)\lambda\ell_j^d$ nodes at that step, where the ξ_j satisfy

$$\frac{\xi}{2} = \xi_1 < \xi_2 < \cdots < \xi_\kappa = \xi \quad \text{and} \quad \xi_j - \xi_{j-1} = \frac{\xi}{2(\kappa-1)}, \text{ for all } j.$$

We start by analyzing the event that all cells are dense for scale 1 during all time steps, which we denote by D_1. The next lemma shows that D_1 occurs with very high probability.

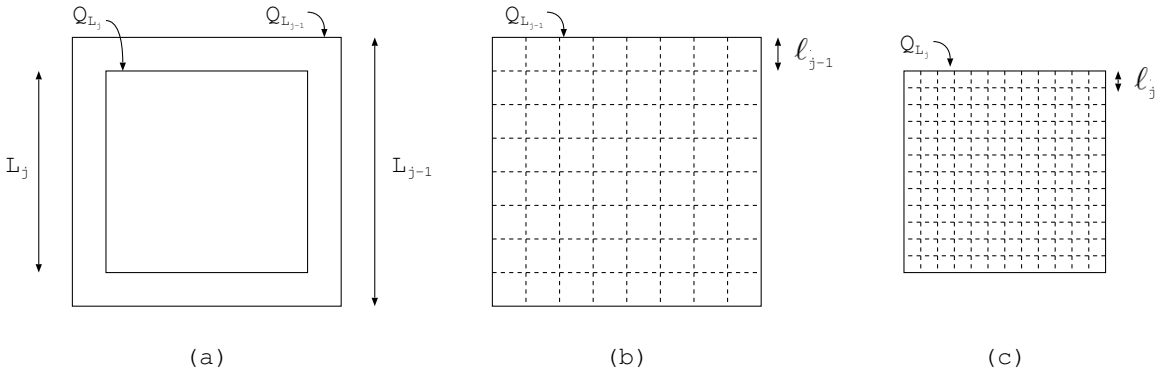

Figure 3.1: (a) The cube $Q_{L_{j-1}}$ and the smaller cube Q_{L_j} obtained when going from scale $j-1$ to scale j. (b) The tessellation of $Q_{L_{j-1}}$ into cells of side length ℓ_{j-1}. (c) The finer tessellation of Q_{L_j} into cells of side length $\ell_j < \ell_{j-1}$.

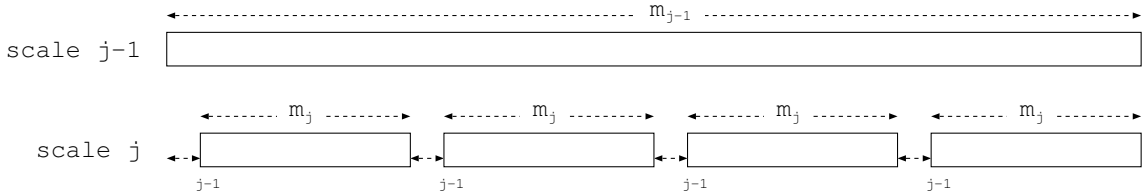

Figure 3.2: Illustration for how 1 time interval of scale $j-1$ gives 4 subintervals of scale j.

Lemma 3.21. *If $\ell_1^d > C \log t$ for some large enough constant C, then there exists a constant c such that*

$$\mathbf{P}\left(D_1\right) \geq 1 - \exp\left(-c\lambda\ell_1^d\right).$$

Proof. For any fixed time i and cell k, the number of nodes in k at time i is given by a Poisson random variable with mean $\lambda\ell_1^d$. Then, using a Chernoff bound (cf. Lemma A.3), we obtain that there are more than $(1-\xi_1)\lambda\ell_1^d$ nodes in that cell at that time step with probability larger than $1 - \exp(-\xi_1^2\lambda\ell_1^d/2)$. The number of cells inside Q_{L_1} is $O(t^{2d})$ by our choice of L_1 and ℓ_1. The proof is completed by taking the union bound over all cells and time steps, and using the assumption on ℓ_1. □

We will need to disregard some time steps when going from one scale to the next. During this discussion it will be useful to refer to Figure 3.2. Let s_j be the number of time steps considered for scale j. We start with $s_1 = t$ so that at scale 1 all time steps are considered; we will have $s_1 > s_2 > \cdots > s_\kappa$. For each scale j, we will split time into intervals of m_j

consecutive time steps. We start with $m_1 = t$, so that at scale 1 we have only one time interval of length t.

In each interval $[a, a + m_{j-1})$ at scale $j - 1$ we consider the following four separated subintervals of length m_j (see Figure 3.2):

$$[a + k\Delta_{j-1} + (k-1)m_j, a + k\Delta_{j-1} + km_j), \text{ for } k = 1, 2, 3, 4, \tag{3.33}$$

where

$$m_j = \frac{m_{j-1} - 4\Delta_{j-1}}{4}. \tag{3.34}$$

We will set the Δ_j in a moment. We skip Δ_{j-1} steps in order to allow the nodes to move far enough and enable the application of the coupling from Proposition 3.18. Note that this gives $s_j = s_{j-1}\left(1 - \frac{4\Delta_{j-1}}{m_{j-1}}\right)$.

For a given scale j, we say that a time interval is *dense* if all cells are dense during all the time steps contained in this time interval, i.e., each cell contains more than $(1 - \xi_j)\lambda\ell_j^d$ nodes at all time steps.

Let $0 = \epsilon_1 < \epsilon_2 < \cdots < \epsilon_\kappa = \epsilon$ satisfy $\epsilon_j - \epsilon_{j-1} = \frac{\epsilon}{\kappa - 1}$. For each scale $j \geq 1$, we define the event

$$D_j = \{\text{a fraction of at least} \geq \left(1 - \frac{\epsilon_j}{2}\right) \text{ time intervals of scale } j \text{ are dense}\}. \tag{3.35}$$

If D_κ holds, the number of time steps for which all cells are dense for the last scale κ is at least

$$\left(1 - \frac{\epsilon_\kappa}{2}\right) s_\kappa = \left(1 - \frac{\epsilon_\kappa}{2}\right) s_1 \prod_{j=1}^{\kappa-1} \left(1 - \frac{4\Delta_j}{m_j}\right) \geq \left(1 - \frac{\epsilon}{2}\right) t \left(1 - \sum_{j=1}^{\kappa-1} \frac{4\Delta_j}{m_j}\right). \tag{3.36}$$

Since we are aiming to obtain $(1 - \epsilon)t$ time steps for which the density condition is satisfied for the last scale, we set Δ_j to satisfy

$$\frac{\Delta_j}{m_j} = \frac{\epsilon}{8\kappa} \tag{3.37}$$

for all j. The value of Δ_j must be sufficiently large to allow nodes to move over a distance ℓ_j. We then define ℓ_j by

$$\Delta_j = C'\ell_j^2\kappa^2, \tag{3.38}$$

where C' is a sufficiently large constant.

From (3.34), (3.37) and (3.38), we obtain

$$\frac{\ell_j^2}{m_j} = \frac{\epsilon}{8C'\kappa^3} \quad \text{and} \quad \ell_{j+1} = \ell_j\sqrt{\frac{1}{4} - \frac{\epsilon}{8\kappa}}. \tag{3.39}$$

Since $m_1 = t$, we get that $\ell_1^2 = \frac{\epsilon}{8C'\kappa^3}t \leq \frac{\epsilon}{8C'}t$ and since we want to get $\ell_\kappa^d = \ell^d \geq C(\log t)^3$, it is easy to see that $\kappa = O(\log t)$ is sufficient.

For any time step i, let $\mathcal{F}_i$ be the σ-field induced by the locations of the nodes of Π_0 from time 0 up to time i.

Lemma 3.22. *Let $A = [a, a + m_j)$ be a time interval considered in scale j. We write $b = a - \Delta_{j-1}$ and $E = \{$at time b all cells are dense for scale $j - 1\}$. Let $\ell^d \geq C(\log t)^3$ for some sufficiently large constant $C > 0$. Then there exists a constant c such that*

$$\mathbf{P}\left(A \text{ not dense}, E \mid \mathcal{F}_b\right) \leq \exp\left(-c\lambda\ell_j^d/\kappa^2\right).$$

Proof. For any $F \in \mathcal{F}_b$ such that $F \cap E = \emptyset$ the lemma clearly holds. We then take $F \cap E \neq \emptyset$ and give an upper bound for $\mathbf{P}\left(A \text{ not dense} \mid E, F\right)$. Let Φ_b be the point process obtained at time b after conditioning on $F \cap E$. We first fix a time $w \in A$ and derive an upper bound for

$$\mathbf{P}\left(\text{at time } w \text{ not all cells are dense for scale } j \mid E, F\right).$$

Since we condition on E, all cells are dense for scale $j-1$ at time b. We now set δ such that $(1-\delta)^2(1-\xi_{j-1}) = 1-\xi_j$, which implies $\delta = \Theta(\xi_j - \xi_{j-1})$. We also choose a constant c and the constant C' appearing in the definition of Δ_j in (3.38) so that, setting

$$L_j \leq L_{j-1} - c\sqrt{\Delta_{j-1}\log\frac{1}{\delta}}, \tag{3.40}$$

allows us to apply Proposition 3.18 with $K = L_{j-1}$ and $K' = L_j$. Thus we obtain a fresh Poisson point process Ξ with intensity $(1-\delta)(1-\xi_{j-1})\lambda$ that can be coupled with Φ_w (which is the point process obtained at time w after the points of Φ_b have moved for time $w-b$) in such a way that Ξ is stochastically dominated by Φ_w inside Q_{L_j} with probability at least

$$1 - \exp\left(-c_1\delta^2(1-\xi_{j-1})\lambda\ell_{j-1}^d\right), \tag{3.41}$$

for some positive constant c_1. We note that the choice of $L_1 = t^2$ and the fact that $\kappa = O(\log t)$ together with equation (3.38) gives that it is always possible to choose the L_j's satisfying (3.40) and such that $L_\kappa = t$.

A given cell is dense for scale j at time w if Ξ contains at least $(1-\xi_j)\lambda\ell_j^d$ nodes in that cell, which by the choice of δ happens with probability at least $1-\exp\left(-c_2\delta^2(1-\delta)(1-\xi_{j-1})\lambda\ell_j^d\right)$ for some constant c_2 (cf. Lemma A.3). The proof is completed by taking the union bound over all cells and over all time steps in A and using the condition for ℓ. □

We now use Lemma 3.22 to give an upper bound for $\mathbf{P}\left(D_j^c \cap D_{j-1}\right)$ that holds for all j, where D_j was defined in (3.35).

Lemma 3.23. *If $\ell^d \geq C(\log t)^3$ for some large enough C, then there exists a constant c such that for any $j \geq 2$ we have*

$$\mathbf{P}\left(D_j^c \cap D_{j-1}\right) \leq \exp\left(-\frac{c\lambda t(\log t)^{3-6/d}}{\kappa^6}\right).$$

Proof. If D_{j-1} happens, then there are at least $\left(1-\frac{\epsilon_{j-1}}{2}\right)\frac{s_{j-1}}{m_{j-1}}$ dense time intervals for scale $j-1$. When we go to scale j, these intervals will give us

$$4\left(1-\frac{\epsilon_{j-1}}{2}\right)\frac{s_{j-1}}{m_{j-1}} \tag{3.42}$$

time intervals that we will consider for scale j. On the other hand, if the event D_j^c holds, then there are less than

$$\left(1-\frac{\epsilon_j}{2}\right)\frac{s_j}{m_j} \tag{3.43}$$

dense intervals for scale j. Let w be obtained by subtracting (3.43) from (3.42), that is,

$$w = \frac{s_j}{m_j}\left(\frac{\epsilon_j - \epsilon_{j-1}}{2}\right).$$

Let Z be the number of subintervals $[a, a+m_j)$ of scale j that are not dense for scale j, but are such that the time step $a - \Delta_{j-1}$ is dense for scale $j-1$. (We call a time step dense if all cells are dense at that time.) It is easy to see that if both D_{j-1} and D_j^c happen, then $Z \geq w$.

We can write Z as a sum of s_j/m_j indicator random variables I_k, one for each time interval of scale j. Although the I_k's depend on one another, Lemma 3.22 gives that the probability that $I_k = 1$ given an arbitrary realization of the previous $k-1$ indicators is smaller than $\rho_j = \exp\left(-c_1\xi^2\lambda\ell_j^d/\kappa^2\right)$ for some constant c_1. Therefore, Z is stochastically dominated by a random variable Z' obtained as a sum of s_j/m_j i.i.d. Bernoulli random variables with mean ρ_j. Using a Chernoff bound (cf. Lemma A.2), we obtain

$$\begin{aligned}\mathbf{P}\left(Z' \geq w\right) &= \mathbf{P}\left(Z' - \mathbf{E}\left[Z'\right] \geq \frac{s_j}{m_j}\left(\frac{\epsilon_j - \epsilon_{j-1}}{2} - \rho_j\right)\right) \\ &\leq \exp\left(-\frac{s_j}{m_j}\left(\frac{\epsilon_j - \epsilon_{j-1}}{2}\right)\left(\log\left(\frac{\epsilon_j - \epsilon_{j-1}}{2\rho_j}\right) - 1\right)\right). \end{aligned} \tag{3.44}$$

Note that $\epsilon_j - \epsilon_{j-1} = \frac{\epsilon}{\kappa - 1}$ and $\log(\rho_j^{-1}) = \Theta(\xi^2\lambda\ell_j^d/\kappa^2)$. Also $\ell_j^d \geq \ell^d \geq C(\log t)^3$ and $\kappa = O(\log t)$, so we obtain a constant c_2 such that

$$\mathbf{P}\left(Z' \geq w\right) \leq \exp\left(-c_2\xi^2\lambda s_j\frac{\ell_j^d}{m_j}\frac{\epsilon}{\kappa^3}\right).$$

Recall from (3.39) that $\frac{\ell_j^2}{m_j} = \frac{\epsilon}{8C'\kappa^3}$. By (3.36) and (3.37) we have that $s_{j-1} = \Theta(t)$ for all j, so we finally obtain

$$\mathbf{P}\left(Z' \geq w\right) \leq \exp\left(-c_3 \frac{\epsilon^2 \xi^2}{\kappa^6} \lambda \ell_j^{d-2} t\right)$$

for some constant $c_3 > 0$. Using $\ell_j \geq \ell$ and the assumption on ℓ in the statement of the lemma completes the proof. □

We are now in a position to prove Proposition 3.19.

Proof of Proposition 3.19. To prove Proposition 3.19, we need to derive an upper bound for $\mathbf{P}\left(D_\kappa^{\mathrm{c}}\right)$. Note that $\mathbf{P}\left(D_\kappa^{\mathrm{c}}\right) \leq \mathbf{P}\left(D_\kappa^{\mathrm{c}} \cap D_{\kappa-1}\right) + \mathbf{P}\left(D_{\kappa-1}^{\mathrm{c}}\right)$. Applying this inequality recursively for the term $\mathbf{P}\left(D_{\kappa-1}^{\mathrm{c}}\right)$ we obtain

$$\mathbf{P}\left(D_\kappa^{\mathrm{c}}\right) \leq \sum_{j=2}^{\kappa} \mathbf{P}\left(D_j^{\mathrm{c}} \cap D_{j-1}\right) + \mathbf{P}\left(D_1^{\mathrm{c}}\right).$$

Each term in the sum can be bounded using Lemma 3.23 and the last term can be bounded using Lemma 3.21. The proof is completed since $\kappa = O(\log t)$ and the initial value

$$\ell_1 = \frac{\epsilon}{8C'\kappa^3} t \geq c_1 \frac{t}{(\log t)^3}.$$

□

3.4 Broadcast time

In this section we use Theorem 3.6 to prove Corollary 3.7 for a finite mobile network of volume n.

We may relate the mobile geometric graph model on the torus to a model on $\mathbb{R}^d$ as follows. Let S_n denote the cube $Q_{(n)^{1/d}}$. The initial distribution of the nodes is a Poisson point process over $\mathbb{R}^d$ with intensity λ on S_n and zero elsewhere. We allow the nodes to move according to Brownian motion over $\mathbb{R}^d$ as usual, and at each time step we project the location of each node onto S_n so that nodes "wrap around" S_n when they reach the boundary.

Proof of Corollary 3.7. Let $t = C \log n (\log\log n)^{3+6/d}$ for some sufficiently large constant $C = C(d)$. We define a *giant component* as a connected component that contains at least two nodes at distance larger than $\frac{(n)^{1/d}}{4}$. It follows from [72, Theorem 2] and the union bound over time steps that, with probability $1 - e^{-\Theta(n^{1/d})}$, G_i contains a *unique* giant component for all integer $i \in [0, 2t-1]$.

The proof proceeds in two stages. First, we show that for any fixed $i \in [0, 2t-1]$, w.h.p. the giant component of G_i has at least one node in common with the giant component of G_{i+1}. This means that, once the message has reached the giant component, it will reach any node v as soon as v itself belongs to the giant component. Then we show that, after t steps, all nodes have belonged to the giant component w.h.p. This implies that broadcast is achieved after $2t$ steps w.h.p.

To establish the first stage, let $\epsilon > 0$ be sufficiently small so that $(1-\epsilon)\lambda > \lambda_c$. We use the thinning property to split Π_i into two Poisson point processes, Π'_i and Π''_i, with intensities $(1-\epsilon)\lambda$ and $\epsilon\lambda$ respectively. Let G'_i and G'_{i+1} be the graphs induced by Π'_i and Π'_{i+1} respectively. Then with probability $1 - e^{-\Theta(n^{1/d})}$ both G'_i and G'_{i+1} contain a unique giant component [72, Theorem 2]. We show that at least one node from Π''_i belongs to both giant components. For any node v of Π''_i, the probability that v belongs to the giant component of G'_i is larger than some constant $c = c(d)$. Moreover, using the FKG inequality we can show that v belongs to the giant components of both G'_i and G'_{i+1} with probability larger than c^2. Therefore, using the thinning property again, we can show that the nodes from Π''_i that belong to the giant components of both G'_i and G'_{i+1} form a Poisson point process with intensity $\epsilon\lambda c^2$, since c does not depend on Π''_i. Hence, there will be at least one such node inside S_n with probability $1 - e^{-\epsilon c^2 n}$, and this stage is concluded by taking the union bound over time steps i.

We now proceed to the second stage of the proof. We first need to show that the tail bound on T_{perc} from Theorem 3.6 also holds when applied to the finite region S_n defined above. Note that all the derivations in the proof of Theorem 3.6 were restricted to the cube Q_{L_1}, where $L_1 = t^2$ was defined in Section 3.3.2. We have that Q_{L_1} is contained inside S_n for all sufficiently large n since $L_1 = t^2 = O(\log^2 n (\log\log n)^{6+12/d})$ while S_n has side length $n^{1/d}$. In order to check that the toroidal boundary conditions do not affect the result, it suffices to observe that, during the time interval $[0, 2t]$, no node moved distance larger than $\frac{n^{1/d}}{2}$ w.h.p.

Now note that, by a Chernoff bound, G has at most $(1+\delta)\lambda n$ nodes with probability larger than $1 - e^{-\Omega(n)}$ for any fixed $\delta > 0$. These nodes are indistinguishable, so letting ρ be the probability that an arbitrary node has percolation time at least t, we can use the union bound to deduce that this applies to *at least one* node in G with probability at most $(1+\delta)\lambda n\rho$. Let v be an arbitrary node. In order to relate ρ to the result of Theorem 3.6, we can use translation invariance and assume that v is at the origin. Then, by the "Palm theory" of Poisson point processes [83], ρ is equivalent to the tail of the percolation time for a node added at the origin, which is precisely $\mathbf{P}(T_{\text{perc}} > t)$. Thus finally, using Theorem 3.6 we get $\rho \leq \exp(-c\frac{t}{(\log t)^{3+6/d}})$, which can be made $o(1/n)$ by setting C sufficiently large in the definition of t.

We then obtain that with probability $1 - o(1/n)$ all nodes of G have been in the giant component during the time interval $[0, t-1]$, which implies that at time step $t-1$, the nodes of the giant component contain the message being broadcast. By stationarity, with

probability $1 - o(1/n)$ all nodes have been in the giant component during the time interval $[t, 2t - 1]$, and thus have received the message by time $2t$. This completes the proof of Corollary 3.7. □

Remark 3.24. It is easy to see that the above result also holds in the case where the graph has *exactly* λn nodes. The proof above shows that, by setting C large enough, we can ensure $\mathbf{P}(T_{\text{bc}} > 2t) = o(1/n)$ for the given value of t. Also, it is well known that a Poisson random variable with mean λn takes the value λn with probability $p = \Theta(1/\sqrt{n})$. Therefore, for a graph with exactly λn nodes, we have $\Pr[T_{\text{bc}} < t] = \frac{p - o(1/n)}{p} = 1 - o(1/\sqrt{n})$.

Chapter 4

Randomized Broadcast on General Networks

In this chapter, we consider the push algorithm, which was discussed in Chapter 2 for random geometric graphs, and study its performance in *general* networks. Recall that, in this algorithm, at the beginning, one node has a piece of information and, at each discrete step, each informed node chooses a neighbor independently and uniformly at random and informs it. We denote by I_t the set of nodes that have the information by the end of step t.

Let $G = (V, E)$ be an undirected, simple, connected graph with n nodes. For a node $u \in V$, $\deg(u)$ denotes the degree of u and $\Gamma(u)$ denotes the set of neighbors of u in G; therefore, $\deg(u) = |\Gamma(u)|$. A graph is called *vertex transitive* if, for each pair of nodes $u, v \in V$, there exists a permutation $\pi \colon V \to V$ such that $\pi(u) = v$ and, for each pair of nodes $r, s \in V$, we have that $\{r, s\} \in E$ if and only if $\{\pi(r), \pi(s)\} \in E$.

Definition 4.1 (Vertex Expansion)**.** *The vertex expansion of a graph $G = (V, E)$ is*

$$\alpha = \min_{S \subseteq V, |S| \leq n/2} \frac{|\partial(S)|}{|S|},$$

where $\partial(S)$ is the boundary of S, i.e., the nodes in $V \setminus S$ having an edge to a node in S.

Clearly $\alpha \leq 1$ for any graph G, and if G is connected then $\alpha > 0$. When $\alpha = \Theta(1)$, we say that G is a *vertex expander*. For any two sets of nodes $S, T \subseteq V$ we define $E(S, T)$ as the set of edges between a node in S and a node in T. We now define conductance.

Definition 4.2 (Conductance)**.** *The conductance of a graph $G = (V, E)$ is*

$$\Phi = \min_{S \subseteq V, 0 < \mathrm{vol}(S) \leq \mathrm{vol}(V)/2} \frac{|E(S, V \setminus S)|}{\mathrm{vol}\,(S)},$$

where $\mathrm{vol}\,(S) = \sum_{v \in S} \deg(v)$.

It is easy to see that for any d-regular graph G, $\alpha/d \leq \Phi \leq \alpha$. Hence any d-regular graph with constant conductance also has constant vertex expansion, but the converse is not true. An elementary example is obtained by connecting two cliques with $n/2$ nodes by a matching, which yields a graph with constant vertex expansion, but with conductance of order $\Theta(1/n)$.

Another important class of graphs for which the vertex expansion can be potentially much larger than the conductance is vertex transitive graphs. While vertex transitive graphs always have a vertex expansion of at least $1/(4D)$ [6, Theorem 3.2], where D is the diameter of the graph, the conductance can be much smaller (even exponentially smaller) than $1/D$. Even for the hypercube with n nodes, there is already a substantial difference between vertex expansion $\alpha = \Theta(1/\sqrt{\log n})$ and conductance $\Phi = \Theta(1/\log n)$. The difference between these two expansion measures is closely linked to the fact that vertex expansion is monotone under adding edges, while conductance is not.

Our first result, which we prove in Section 4.1, relates vertex expansion to the runtime of the push algorithm.

Theorem 4.3. *For any d-regular graph with vertex expansion α, all nodes are informed within $O\left((1/\alpha)\log^5 n\right)$ steps by the push algorithm, with high probability.*

We remark that this result does not extend to non-regular graphs. To see this, consider two cliques of size $\sqrt{n}$ and $n - \sqrt{n}$, and for each node of the smaller clique add an edge between it and one node of the larger clique in such a way that we obtain a matching of size $\sqrt{n}$ between the two cliques. It is easy to verify that this graph has constant vertex expansion. Now if the information is initially placed in the large clique, then we need $\Omega(\sqrt{n})$ steps in expectation until one node of the large clique informs one in the smaller clique, since each of the $\sqrt{n}$ nodes of the larger clique having an edge to the smaller clique will transmit the information through that edge with probability $\Theta(1/n)$.

Another interesting application of Theorem 4.3 is to vertex transitive graphs. The runtimes of the push algorithm on certain vertex transitive graphs (e.g., the hypercube) have been analyzed in [39, 34, 77]. Moreover, much research has been devoted to the construction of efficient determininistic algorithms on vertex transitive graphs, e.g., [47, 54, 79]. Our result provides a nearly optimal runtime that applies to *all* vertex transitive graphs.

Corollary 4.4. *Let G be any vertex transitive graph with diameter D. Then all nodes are informed in $O(D\log^5 n)$ steps, with high probability.*

The corollary above follows immediately from Theorem 4.3 and a result of Babai [6, Theorem 3.2] saying that every vertex transitive graph with diameter D has vertex expansion at least $1/(4D)$.

Since D is a trivial lower bound, Corollary 4.4 implies that the runtime on vertex transitive graphs is "fast," i.e., of order $O(\text{polylog}(n))$, if and only if $D = O(\text{polylog}(n))$. From another perspective, as $\alpha \geq 1/(4D)$ and we need at least D steps to inform all nodes,

the runtime for any vertex transitive graph is between $\Omega(1/\alpha)$ and $O((1/\alpha)\log^5 n)$. This shows that the dependency on $1/\alpha$ in Theorem 4.3 is the best possible (neglecting logarithmic factors).

Another natural question is whether there exists a substantial difference between the performance of the push algorithm on graphs with constant conductance and on graphs with constant vertex expansion. We show that this is indeed the case by proving that there are graphs with constant vertex expansion on which the push algorithm takes $\omega(\log n)$ steps, while on any graph with constant conductance the push algorithm is known to take at most $O(\log n)$ steps (cf. [69, 20, 77]). We give the proof of the theorem below in Section 4.2.

Theorem 4.5. *There is a regular graph with constant vertex expansion (which is also vertex transitive) for which the push algorithm requires $\Omega(\log^2 n)$ steps with high probability.*

Finally, we conclude this chapter by giving in Section 4.3 applications of the push algorithm to other processes such as the cover time of random walks, and by discussing extensions of our results to non-regular graphs in Section 4.4.

4.1 Upper bound

This section is devoted to the proof of Theorem 4.3. In the proof we consider a different version of the push algorithm, which we call the *sequential model*. In the sequential model, each time step is divided into n substeps. At each substep $1/n, 2/n, 3/n, \ldots$, we choose a node $u \in V$ uniformly at random and a neighbor $v \in \Gamma(u)$ also uniformly at random. Then, if u is informed, it transmits the information to v. We remark that for regular graphs, this procedure translates to choosing at each step an undirected edge $\{u, v\} \in E$ uniformly at random along with a uniformly random orientation of its endpoints.

Before proceeding, we emphasize that we refer to a substep of the sequential model as a time interval of length $1/n$. If t *substeps* of the sequential model are performed, then we say that t/n *time steps* occurred, where t is an integer but t/n may be non-integer. As shown in [77, Theorem 3], for any graph G and worst-case initial node, the sequential model and the (standard) push algorithm take the same number of time steps (up to constant factors) with high probability.

In what follows we only consider the sequential model. Our main motivation in resorting to this model is that it has a useful symmetry property. In order to explain the property, consider two nodes $u, v \in V$ and a time step t, and let $p^t_{u,v} = \mathbf{P}(v \in I_t \mid I_0 = \{u\})$, where I_t for the sequential model is defined as the set of informed nodes at time step t. In other words, $p^t_{u,v}$ is the probability that v is informed within t time steps given that initially u was the only informed node. We extend this definition to sets of nodes by considering two sets $X, X' \subseteq V$ and defining $p^t_{X,X'}$ as the probability that at least one node of X' is informed given that initially only the nodes of X were informed. More formally, $p^t_{X,X'} = \mathbf{P}(I_t \cap X' \neq \emptyset \mid I_0 = X)$. The symmetry property of the sequential model is expressed in the following lemma.

Lemma 4.6. *Let $G = (V, E)$ be any regular graph. Let $u, v \in V$ be two nodes and $t \in \{0, 1/n, 2/n, \ldots\}$ be arbitrary but fixed. Then for the sequential model we have $p^t_{u,v} = p^t_{v,u}$. More generally, for any two subsets $X, X' \subseteq V$ we have $p^t_{X,X'} = p^t_{X',X}$.*

Proof. Let $\mathcal{L}$ be an instance of t time steps of the sequential model. We can see $\mathcal{L}$ as a list of tn oriented pairs of nodes $\mathcal{L} = \{(u_1, v_1), (u_2, v_2), \ldots, (u_{tn}, v_{tn})\}$, where (u_i, v_i) is an edge of G for all $i \in [1, tn]$. The existence of the pair (u_i, v_i) in $\mathcal{L}$ means that at substep i the sequential model selected the node u_i and its neighbor v_i and u_i informed v_i if u_i was informed at that moment. From $\mathcal{L}$ we can construct another instance $\mathcal{L}' = \{(u'_1, v'_1), (u'_2, v'_2), \ldots, (u'_{tn}, v'_{tn})\}$ of the sequential model by setting $(u'_i, v'_i) = (v_{tn-i+1}, u_{tn-i+1})$ for all $i \in [1, tn]$. In other words, $\mathcal{L}'$ is obtained by reversing the order in which the transmissions of $\mathcal{L}$ happen and also inverting the direction according to which each transmission occurs in $\mathcal{L}$. Therefore, if there exists a sequence of nodes $w_1, w_2, \ldots, w_k$ such that w_i informed w_{i+1} according to $\mathcal{L}$ for all $1 \leq i \leq k-1$, then w_{i+1} informs w_i in $\mathcal{L}'$ for all $1 \leq i \leq k-1$. The lemma then follows since $\mathcal{L}$ and $\mathcal{L}'$ occur with the same probability for any regular graph. □

In the sequel, whenever we say that u *informs* a set of nodes X within t time steps, we mean that if the sequential model is initiated at u, then all nodes in X will be informed within t time steps. Thus, u informing X does not necessarily mean that u transmits the information directly to each node in X. We also note that this notion is different from $p^t_{u,X}$ which was defined as the probability that u informs at least *one* element in X within t time steps.

We now start the proof of Theorem 4.3 by showing that in the first steps of the algorithm the number of informed nodes grows exponentially fast.

Lemma 4.7. *Let $G = (V, E)$ be any d-regular graph and fix an arbitrary node $u \in V$. Let $t = \log d + 10$. Then with probability at least $1/2$, u informs at least $d/2$ nodes within t time steps. Moreover,*

$$\mathbf{E}\left[|I_t|\right] = \sum_{v \in V} p^t_{u,v} \leq e^{10} d.$$

Proof. We start with the upper bound. Assume that whenever an informed node v is chosen at a certain substep, then the set of informed nodes increases by 1. Hence, we obtain the inequality

$$\mathbf{E}\left[|I_t|\right] = \mathbf{E}\left[\mathbf{E}\left[|I_t| \;\middle|\; |I_{t-1/n}|\right]\right] \leq \mathbf{E}\left[|I_{t-1/n}|\right] + \frac{\mathbf{E}\left[|I_{t-1/n}|\right]}{n} = \left(1 + \frac{1}{n}\right) \mathbf{E}\left[|I_{t-1/n}|\right],$$

and by induction

$$\mathbf{E}\left[|I_t|\right] \leq \left(1 + \frac{1}{n}\right)^{nt} \mathbf{E}\left[|I_0|\right] \leq \exp(t), \tag{4.1}$$

and the upper bound on $\mathbf{E}[|I_t|]$ follows for $t = \log d + 10$.

We proceed to the lower bound. For any values $1 \le i < j \le d/2$, let $\Delta_{i,j}$ be the number of time steps required to increase the number of informed nodes from i to j; therefore, $\Delta_{i,j} \in \{0, 1/n, 2/n, 3/n, \ldots\}$. Let $t = \Delta_{1,i}$ be the first time step with $|I_t| = i$. Then, $\Delta_{i,i+1}$ is $1/n$ times a geometric random variable with parameter $|E(I_t, I_t^c)|/(2|E|)$. Assuming I_t to constitute a clique, we obtain the lower bound

$$\frac{|E(I_t, I_t^c)|}{2|E|} \ge \frac{id - 2\frac{i(i-1)}{2}}{nd} \ge \frac{i(d-i+1)}{nd} \ge \frac{i(1-i/d)}{n}.$$

Hence, the expected number of time steps until u informs at least $d/2$ nodes is upper bounded by

$$\mathbf{E}\left[\Delta_{1,1+d/2}\right] = \sum_{i=1}^{d/2} \mathbf{E}\left[\Delta_{i,i+1}\right] \le \frac{1}{n} \sum_{i=1}^{d/2} \frac{n}{i(1-i/d)} = \sum_{i=1}^{d/2} \left(\frac{1}{i} + \frac{1}{d-i}\right) \le \log d + 2.$$

Our next goal is to prove that with constant probability $\Delta_{1,1+d/2}$ is smaller than $\log d + 10$. Using Markov's inequality and (4.1) we obtain that for any $x \ge 0$

$$\mathbf{P}\left(\Delta_{1,1+d/2} \le \log d - x\right) = \mathbf{P}\left(|I_{\log d - x}| \ge 1 + d/2\right) \le \frac{\mathbf{E}\left[|I_{\log d - x}|\right]}{1 + d/2} \le \frac{e^{\log d - x}}{1 + d/2} \le 2e^{-x}. \quad (4.2)$$

Now, if we write $\mathbf{E}\left[\Delta_{1,1+d/2}\right] = \sum_{x=0}^{\infty} x\, \mathbf{P}\left(\Delta_{1,1+d/2} = x\right)$, where the sum is over the values of x in $\{0, 1/n, 2/n, 3/n, \ldots\}$, we obtain

$$\begin{aligned}
\log d + 2 &\ge \mathbf{E}\left[\Delta_{1,1+d/2}\right] \\
&= \sum_{x=-\log d}^{\infty} (\log d + x)\, \mathbf{P}\left(\Delta_{1,1+d/2} = \log d + x\right) \\
&\ge \log d + \sum_{x=2+1/n}^{\log d} (-x)\, \mathbf{P}\left(\Delta_{1,1+d/2} = \log d - x\right) \\
&\quad + (-2)\, \mathbf{P}\left(\log d - 2 \le \Delta_{1,1+d/2} \le \log d + 9 + (n-1)/n\right) \\
&\quad + 10\, \mathbf{P}\left(\Delta_{1,1+d/2} \ge \log d + 10\right). \qquad (4.3)
\end{aligned}$$

Note that changing the index of the sum to $y = x - 2$ we can write

$$\begin{aligned}
&\sum_{x=2+1/n}^{\log d} x\,\mathbf{P}\left(\Delta_{1,1+d/2} = \log d - x\right) \\
&= \sum_{y=1/n}^{\log d-2} (2+y)\,\mathbf{P}\left(\Delta_{1,1+d/2} = \log d - 2 - y\right) \\
&= 2\,\mathbf{P}\left(\Delta_{1,1+d/2} \le \log d - 2 - 1/n\right) + \sum_{y=1/n}^{\log d-2} y\,\mathbf{P}\left(\Delta_{1,1+d/2} = \log d - 2 - y\right) \\
&= 2\,\mathbf{P}\left(\Delta_{1,1+d/2} \le \log d - 2 - 1/n\right) + \frac{1}{n}\sum_{y=1/n}^{\log d-2} \mathbf{P}\left(\Delta_{1,1+d/2} \le \log d - 2 - y\right) \\
&\le (2+1/n)\,\mathbf{P}\left(\Delta_{1,1+d/2} \le \log d - 2 - 1/n\right) + \int_{1/n}^{\log d-2} \mathbf{P}\left(\Delta_{1,1+d/2} \le \log d - 2 - y\right) dy,
\end{aligned}$$

where the last inequality comes from the fact that $\mathbf{P}\left(\Delta_{1,1+d/2} \le \log d - 2 - y\right)$ decreases as y increases, and therefore,

$$\begin{aligned}
&\frac{1}{n}\sum_{y=1/n}^{\log d-2} \mathbf{P}\left(\Delta_{1,1+d/2} \le \log d - 2 - y\right) \\
&\le \frac{1}{n}\,\mathbf{P}\left(\Delta_{1,1+d/2} \le \log d - 2 - 1/n\right) + \int_{1/n}^{\log d-2} \mathbf{P}\left(\Delta_{1,1+d/2} \le \log d - 2 - y\right) dy.
\end{aligned}$$

Now we use (4.2) to obtain

$$\sum_{x=2+1/n}^{\log d} x\,\mathbf{P}\left(\Delta_{1,1+d/2} = \log d - x\right) \le 2(2+1/n)e^{-2-1/n} + 2e^{-2-1/n} - 2e^{-\log d} \le 1,$$

for all $n \ge 2$. Plugging the result above and (4.2) into (4.3), and subtracting $\log d$ from both sides, we obtain

$$-1 - 2\,\mathbf{P}\left(\log d - 2 \le \Delta_{1,1+d/2} \le \log d + 9 + (n-1)/n\right) + 10\,\mathbf{P}\left(\Delta_{1,1+d/2} \ge \log d + 10\right) \le 2.$$

Using $\mathbf{P}\left(\log d - 2 \le \Delta_{1,1+d/2} \le \log d + 9 + (n-1)/n\right) \le 1$ and rearranging the terms yield

$$\mathbf{P}\left(\Delta_{1,1+d/2} \ge \log(d/2) + 10\right) \le \frac{1}{10}\,(1+2+2) = \frac{1}{2}.$$

This implies

$$\mathbf{P}\left(\Delta_{1,1+d/2} \le \log d + 10\right) = \mathbf{P}\left(|I_{\log d+10}| \ge d/2\right) \ge \frac{1}{2}$$

and therefore

$$\mathbf{E}\left[|I_{\log d+10}|\right] \geq \frac{1}{2}\frac{d}{2} = \frac{d}{4},$$

which completes the proof. □

Remark 4.8. The fact that the upper and lower bounds in Lemma 4.7 differ only by a constant factor will turn out to be crucial in our proof of Theorem 4.3.

For the analysis to follow it will be important to distinguish the nodes that are likely to be informed by u.

Definition 4.9. *Let $X \subseteq V$ be any subset and u be an arbitrary node. We call u a* friend of X *if the probability that u informs $d/4$ or more nodes in X within $\log d + 10$ time steps is at least $1/2$.*

The following lemma is a direct consequence of Lemma 4.7.

Lemma 4.10. *Let $X \subseteq V$ be any subset and u be an arbitrary node. Then u is a friend of X or of $V \setminus X$.*

Note that u may be a friend of both X and $V \setminus X$. In order to explain the motivation behind Lemma 4.10, let X be the set of informed nodes and $u \notin X$. Then, if u is a friend of X, u will become informed within $\log d + 10$ time steps with probability at least $1/2$. This follows by the symmetry property of Lemma 4.6. On the other hand, consider the case when u is a friend of $V \setminus X$. Then, if u is only connected by one edge to X, it may take up to d steps in expectation until u becomes informed. However, after that, since u is a friend of $V \setminus X$, u will inform $d/4$ nodes in $V \setminus X$ within $\log d + 10$ time steps with probability at least $1/2$.

We are now ready to proceed with the proof of Theorem 4.3. We split the proof into three phases according to the size of I_t. Let τ_1, τ_2 and τ_3 be the number of time steps spent in each phase, respectively.

- **Phase 1:** $1 \leq |I_t| \leq d/2$. Thus τ_1 is the smallest integer such that $|I_{\tau_1}| \geq d/2$. The analysis of this phase follows almost directly from Lemma 4.7 and we have that $\tau_1 = O(\log d \log n)$.

- **Phase 2:** $d/2 \leq |I_t| \leq d(8/\alpha)$. Thus τ_2 is the smallest integer such that $|I_{\tau_1+\tau_2}| \geq 8d/\alpha$. We show that $\tau_2 = O((1/\alpha)\log^3 n)$ with high probability. The analysis of this phase relies on Lemma 4.10.

- **Phase 3:** $d(8/\alpha) \leq |I_t| \leq n/2 + 1$. Thus τ_3 is the smallest integer such that $|I_{\tau_1+\tau_2+\tau_3}| \geq n/2 + 1$. In this phase, for any t satisfying $\tau_1 + \tau_2 \leq t \leq \tau_1 + \tau_2 + \tau_3$, $|I_t|$ is large enough such that after one time step, some nodes in $|\Gamma(I_t) \setminus I_t|$ become informed. After that, each of these nodes may inform up to $\Omega(d)$ other nodes within $O(\log d)$ time steps, but dependencies may arise. This is the most challenging phase and we prove that $\tau_3 = O((1/\alpha)\log^5 n)$ with high probability.

Once we have verified the claimed time bounds for the three phases, the theorem follows immediately. After $\tau_1 + \tau_2 + \tau_3 = O((1/\alpha)\log^5 n)$ time steps, $n/2+1$ nodes have been informed with high probability. From Lemma 4.6 we know that in order to keep a node v uninformed at step $2(\tau_1+\tau_2+\tau_3)$, with high probability at least $n/2+1$ nodes must be uninformed in time step $\tau_1+\tau_2+\tau_3$. Since this does not happen, we obtain that with high probability v is informed at time step $2(\tau_1+\tau_2+\tau_3)$ for all v.

We now present three lemmas that deal with each of the phases described above. These will complete our proof.

Lemma 4.11 (Phase 1)**.** *Assume that $|I_0| = 1$. For $\tau_1 = O(\log d \log n)$, we have $|I_{\tau_1}| \geq d/2$ with high probability.*

Proof. Let u be the initially informed node. Then by Lemma 4.7 we know that the expected number of time steps until u informs at least $d/2$ nodes is $O(\log d)$. Hence the high probability result is obtained by repeating this procedure $O(\log n)$ times. □

Lemma 4.12 (Phase 2)**.** *If $|I_t| \geq d/2$ for some t, then for $\tau_2 = O((1/\alpha)\log^3 n)$ we obtain $|I_{t+\tau_2}| \geq d(8/\alpha)$ with high probability.*

Proof. Let t be any time step such that $d/2 \leq |I_t| \leq d(8/\alpha)$. From now on we condition on I_t and consider a case analysis concerning the nodes in $\Gamma(I_t) \setminus I_t$. The first case is that at least half of the nodes in $\Gamma(I_t) \setminus I_t$ are friends of I_t. Using Lemma 4.6 we conclude that each node in $\Gamma(I_t) \setminus I_t$ which is a friend of I_t becomes informed within $\log d + 10$ time steps with probability at least $1/2$. Hence we can upper bound the expected number of newly informed nodes after $k = \log d + 10$ additional time steps as follows

$$\begin{aligned}\mathbf{E}\left[|I_{t+k} \setminus I_t|\right] &\geq \mathbf{E}\left[|I_{t+k} \cap (\Gamma(I_t) \setminus I_t)|\right] \\ &= \sum_{u \in \Gamma(I_t)\setminus I_t} \mathbf{P}\left(u \in I_{t+k}\right) \geq \left(\frac{1}{2}\left|\Gamma(I_t) \setminus I_t\right|\right)\frac{1}{2} = \frac{1}{4}|\Gamma(I_t) \setminus I_t|.\end{aligned}$$

Let $Z = |I_{t+k} \cap (\Gamma(I_t) \setminus I_t)|$ and $\mu = \mathbf{E}[Z] \geq \frac{1}{4}|\Gamma(I_t) \setminus I_t|$. Applying Markov's inequality we obtain that

$$\mathbf{P}\left(Z \leq \frac{1}{2}\mu\right) = \mathbf{P}\left(|\Gamma(I_t) \setminus I_t| - Z \geq |\Gamma(I_t) \setminus I_t| - \frac{1}{2}\mu\right) \leq \frac{|\Gamma(I_t) \setminus I_t| - \mu}{|\Gamma(I_t) \setminus I_t| - \frac{1}{2}\mu} \leq \frac{6}{7}.$$

Using this inequality and $|\Gamma(I_t) \setminus I_t| \geq \alpha|I_t|$ we get that

$$\mathbf{P}\left(|I_{t+k}| \geq \left(1+\frac{\alpha}{8}\right)|I_t|\right) \geq \mathbf{P}\left(|I_{t+k} \setminus I_t| \geq \frac{1}{8}|\Gamma(I_t) \setminus I_t|\right) \geq \mathbf{P}\left(Z \geq \frac{1}{2}\mu\right) \geq \frac{1}{7}.$$

Let Y_1 be the number of additional time steps spent in the first case. Note that whenever we happen to be in the first case, $|I_t|$ is increased by a factor of $(1+\alpha/8)$ after $k =$

$\log d + 10$ time steps with probability at least $\frac{1}{7}$. Hence the expected time to achieve such an increase is smaller than $7k$. Now observe that, in order to increase I_t from $d/2$ to $d(8/\alpha)$ it suffices to repeat this procedure $\left(\frac{8}{\alpha} + \frac{1}{2}\right)\log(16/\alpha)$ times since $(d/2)(1+\alpha/8)^{\left(\frac{8}{\alpha}+\frac{1}{2}\right)\log(16/\alpha)} \geq (d/2)e^{\log(16/\alpha)} = d(8/\alpha)$, which follows from the inequality $(1+t/n)^{n+t/2} \geq e^t$ for all $t, n > 0$. We then obtain

$$\mathbf{E}[Y_1] \leq \left(\frac{8}{\alpha} + \frac{1}{2}\right)\log(16/\alpha)7k = O\left(\frac{1}{\alpha}\log n \log d\right).$$

We now consider the second case where at least half of the nodes in $\Gamma(I_t) \setminus I_t$ are friends of $V \setminus I_t$. We can upper bound the expected time until the first node in $\Gamma(I_t) \setminus I_t$ which is a friend of $V \setminus I_t$ becomes informed by $k_1 = \frac{1}{n}\frac{2|E|}{|\Gamma(I_t)\setminus I_t|/2} \leq \frac{2d}{\alpha|I_t|}$, where the factor $1/n$ comes from the fact that one time step comprises n substeps. By Markov's inequality, a node in $\Gamma(I_t) \setminus I_t$ which is a friend of $V \setminus I_t$ becomes informed within $2k_1$ time steps with probability at least $1/2$. Once such a node becomes informed, it will inform $d/4$ nodes in $V \setminus I_t$ within an additional $k_2 = \log d + 10$ time steps with probability at least $1/2$. Combining these insights we have that $\mathbf{P}\left(|I_{t+2k_1+k_2} \setminus I_t| \geq \frac{d}{4}\right) \geq \frac{1}{2}\frac{1}{2} = \frac{1}{4}$.

Thus, whenever we are in the second case, we increase $|I_t|$ by $d/4$ after $2k_1 + k_2$ time steps with probability at least $\frac{1}{4}$. Let Y_2 be the number of steps spent in the second case. We have

$$\begin{aligned}\mathbf{E}[Y_2] &\leq \sum_{i=1}^{(d\frac{8}{\alpha})/(\frac{d}{4})} 4(2k_1 + k_2) \leq \sum_{i=1}^{\lceil\frac{32}{\alpha}\rceil}\left(8\frac{2d}{\alpha(id/4)} + 4(\log d + 10)\right) \\ &= O\left(\frac{1}{\alpha}\log(1/\alpha) + \frac{1}{\alpha}\log d\right).\end{aligned}$$

Clearly, the number of time steps spent in the second phase satisfies $\tau_2 = Y_1 + Y_2$, giving $\mathbf{E}[\tau_2] = O\left((1/\alpha)\log^2 n\right)$. By Markov's inequality, $\mathbf{P}(\tau_2 \geq 2\,\mathbf{E}[\tau_2]) \leq 1/2$. By repeating this $\Theta(\log n)$ times we obtain $\tau_2 = O((1/\alpha)\log^3 n)$ with high probability. □

Lemma 4.13 (Phase 3)**.** *Assume that $|I_t| \geq d(8/\alpha)$ for some t. Then after the push algorithm runs for $\tau_3 = O((1/\alpha)\log^5 n)$ additional time steps, we have $|I_{t+\tau_3}| \geq n/2 + 1$ with high probability.*

Proof. Let t be any time step where $d(8/\alpha) \leq |I_t| \leq n/2$. As in the proof of Lemma 4.11 above, we split the analysis into two cases. The first case is when at least half of the nodes in $\Gamma(I_t) \setminus I_t$ are friends of I_t. Let Y_1 be the number of time steps spent in the first case. Then, as in the proof of Lemma 4.12, it follows that after $\log d + 10$ time steps, we increase $|I_t|$ by a factor of $(1 + \alpha/8)$ with probability at least $1/7$, and $\mathbf{E}[Y_1] \leq O\left(\frac{1}{\alpha}\log n \log d\right)$.

Consider now the more difficult case when at least half of the nodes in $\Gamma(I_t) \setminus I_t$ are friends of $V \setminus I_t$. Let $S \subseteq \Gamma(I_t) \setminus I_t$ be the set of nodes in $\Gamma(I_t) \setminus I_t$ which are friends of $V \setminus I_t$. Assume that each node in S has only one edge to I_t; if a node $u \in S$ has more

than one edge to I_t, then we select one edge $\{u, v\}$ with $v \in I_t$ and ignore transmissions along the other edges. At first, we only consider transmissions from I_t to S. Let T be the number of additional time steps until $\frac{8}{d}|S|$ transmissions have been sent to S. Since the expected time until a transmission from I_t to S occurs is $\frac{1}{n}\frac{nd}{|S|}$, we have $\mathbf{E}[T] \leq \frac{d}{|S|}\frac{8}{d}|S| = 8$. Hence by Markov's inequality, $\mathbf{P}(T \leq 16) \geq \frac{1}{2}$. From now on, from time step t to $t+T$ we only consider transmissions from a node in I_t to a node in S and disregard all the other transmissions. This assumption clearly yields an upper bound for τ_3.

We now concentrate on the first $\frac{8}{d}|S|$ transmissions to S. Let $S' \subseteq S$ be the set of nodes in S that receive a transmission. Given our assumption that each node in S has only one edge to I_t, a transmission from I_t to S is equally likely to hit each node in S. Moreover, since from time step t to $t+T$ we only consider direct transmissions from I_t to S, we obtain

$$\begin{aligned}\mathbf{E}[|S'|] &= |S|\left(1 - \left(1 - \frac{1}{|S|}\right)^{\frac{8}{d}|S|}\right) \geq |S|\left(1 - \exp\left(-\frac{8}{d}\right)\right) \\ &\geq |S|\left(1 - \frac{1}{1+\frac{8}{d}}\right) = |S|\frac{\frac{8}{d}}{1+\frac{8}{d}} \geq |S|\frac{2}{d},\end{aligned}$$

where the last inequality holds for all $d \geq 3$. (If $d = 2$, then G is a cycle, and we never enter Phase 3 since the condition on $|I_t|$ in the statement of the lemma is never satisfied.) By applying Markov's inequality we obtain that

$$\mathbf{P}\left(|S'| \leq \frac{1}{2}\mathbf{E}[|S'|]\right) = \mathbf{P}\left(\frac{8}{d}|S| - S' \geq \frac{8}{d}|S| - \frac{1}{2}\mathbf{E}[|S'|]\right) \leq \frac{\frac{8}{d}|S| - \mathbf{E}[|S'|]}{\frac{8}{d}|S| - \frac{1}{2}\mathbf{E}[|S'|]} \leq \frac{6}{7}.$$

Hence we have shown that

$$\mathbf{P}\left(|S'| \geq \frac{1}{d}|S|\right) \geq \frac{1}{7}. \tag{4.4}$$

(By assumption we have $|I_t| \geq d(8/\alpha)$ and $|S| \geq (1/2)|\Gamma(I_t) \setminus I_t| \geq (1/2)\alpha|I_t|$, which implies $\frac{1}{d}|S| \geq 4$.) From now on we condition our analysis on the event $|S'| \geq \frac{1}{d}|S|$ and we address the conditioning at the end.

We now briefly explain the remainder of the proof. We run for an additional number $k_1 = \log d + 10$ of time steps, after which every node in S' has informed $d/4$ nodes in $V \setminus I_t$ with probability at least $1/2$. It will be helpful to imagine the set of nodes that are informed by a node $s \in S'$ as a directed graph, which we refer to as the *broadcast graph* B_s with root s. The existence of a directed edge $(u \to v)$ in B_s means that u transmitted the information to v. Note that v may receive the information from more than one node, so the in-degree of v may be larger than one. We assign a label to each edge in such a way that $(u \to v)$ has label k if u first transmitted the information to v at time step k. We would like to argue that the informed nodes of S inform in total $\Theta(d|S|)$ nodes in $V \setminus I_t$. The problem is that the

broadcast graphs rooted at different nodes may collide, giving rise to dependencies. To cope with this, we first analyze the *independent case*, where all broadcast graphs are constructed independently from one another. After that we design a coupling between the independent case and the case where dependencies occur, which we refer to as the *parallel case*. We prove that each path of transmissions appearing in the independent case can be translated to a path in the parallel case. Our goal is to show that the paths translated to the parallel case occur in a time interval that is not much larger than the path in the independent case. We now proceed to the formal proof.

We first study the independent case. Consider the broadcast graph B_s rooted at some node $s \in S'$. For each $s \in S'$, let Z_s be the indicator random variable which is 1 if B_s contains at least $d/4$ nodes from $V \setminus I_t$. Recall that $\mathbf{P}(Z_s = 1) \geq 1/2$ for all $s \in S'$. Let $Z = \sum_{s \in S'} Z_s$. Then $\mathbf{E}[Z] \geq (1/2)|S'|$, and conditioning on $|S'| \geq \frac{1}{d}|S|$ we obtain $\mathbf{E}[Z] \geq \frac{1}{2d}|S| \geq 2$. Then, by a Chernoff bound (cf. Lemma A.1), we get

$$\mathbf{P}(Z \leq (1/4)|S'|) \leq \mathbf{P}(Z \leq (1/2)\,\mathbf{E}[Z]) \leq \exp\left(-\frac{(1/2)^2\,\mathbf{E}[Z]}{2}\right) \leq \exp\left(-\frac{1}{4}\right).$$

Now we bound the number of times a node $v \in V \setminus I_t$ occurs in a broadcast graph B_s, which we denote by the random variable Z'_v. Recall that from time step t to $t+T$ we disregard all transmissions to nodes out of S. Recall also that we consider that each node of S has only one edge to I_t. Under these assumptions, for any $u \in S$ we have that $\mathbf{P}(u \in I_{t+T}) = 1 - (1 - 1/|S|)^{\frac{8}{d}|S|} \leq 8/d$, and by linearity of expectations, we obtain

$$\begin{aligned}\mathbf{E}[Z'_v] &= \sum_{u \in S'} p_{uv}^{k_1} = \sum_{u \in S} \mathbf{P}(u \in I_{t+T})\, p_{uv}^{k_1} \leq \frac{8}{d} \sum_{u \in S} p_{uv}^{k_1} \leq \frac{8}{d} \sum_{u \in V} p_{uv}^{k_1} \\ &= \frac{8}{d} \sum_{u \in V} p_{vu}^{k_1} \leq 8e^{10} = c_1,\end{aligned}$$

where the second inequality follows from Lemma 4.6 and the third one follows from Lemma 4.7. Note that $Z'_v = \sum_{u \in S'} \mathbf{1}_{v \in B_u}$ is the sum of independent Bernoulli random variables, so applying a Chernoff bound (cf. Lemma A.1)

$$\mathbf{P}(Z'_v \leq 6c_1 \log n) \geq 1 - n^{-3}.$$

Taking the union bound, we obtain

$$\mathbf{P}\left(\max_{v \in V} Z'_v \leq 6c_1 \log n\right) \geq 1 - n^{-2}. \tag{4.5}$$

Hence, by taking the union bound one more time, we can lower bound the total number of nodes in the union of all broadcast graphs as follows:

$$\begin{aligned}\mathbf{P}\left(|\cup_{s\in S'} B_s| \geq \frac{(1/4)|S'|d/4}{\min\{6c_1\log n, |S'|\}}\right) &\geq \mathbf{P}\left((Z \geq (1/4)|S'|) \cap \left(\max_{v\in V} Y_v \leq 6c_1\log n\right)\right)\\ &\geq \mathbf{P}\left(Z \geq (1/4)|S'|\right) - \mathbf{P}\left(\max_{v\in V} Y_v \geq 6c_1\log n\right)\\ &\geq 1 - \exp(-1/4) - n^{-2}. \qquad (4.6)\end{aligned}$$

We now define a coupling between the independent and the parallel cases. Let $\mathcal{L}_s(u)$ be the ordered list of nodes to which u transmits the information in the broadcast graph B_s; note that the list may be empty. Let $\mathcal{L}(u)$ be an ordered list of nodes defined as the concatenation of $\mathcal{L}_s(u)$ for all $s \in V$, where the order in which the lists are concatenated is arbitrary. Note that the elements of $\mathcal{L}(u)$ are independently and uniformly random neighbors of u. For each node $u \in V$, we use the list $\mathcal{L}(u)$ to define the nodes to which u transmits in the parallel case. In other words, when u becomes informed in the parallel case, the neighbors that u chooses to subsequently transmit the information to are obtained following the list $\mathcal{L}(u)$. When the end of the list is reached, u will choose the neighbors independently from the independent case.

Now we look at paths in a broadcast graph B_s from the root s to a node v. Take the path $P = (v_0 = s, v_1, \ldots, v_\ell = v)$ such that for all i, v_{i+1} received the information for the first time from v_i. Note that the labels of the edges of P in the broadcast graph B_s are monotonically increasing with the path.

From [77, Lemma 8], it holds that with probability at least $1 - n^{-2}$, $\ell \leq 2000\log n$ for all paths in a fixed broadcast graph B_s. Taking the union bound over all broadcast graphs, this holds for all broadcast graphs simultaneously with probability at least $1 - n^{-1}$. We proceed to upper bound $|\mathcal{L}_s(u)|$. The broadcast graph B_s is restricted to the time interval $[t+T, t+T+k_1]$ and hence the number of transmissions of u is the sum of nk_1 independent bernoulli random variables with sucess probability $1/n$ each. Using a standard Chernoff bound, we obtain a constant $c_2 > 0$ such that $\mathbf{P}\left(|\mathcal{L}_s(u)| \leq c_2\log n\right) \geq 1 - n^{-4}$, and hence by taking the union bound over all nodes $s \in S'$ and $u \in V$,

$$\mathbf{P}\left(\max_{s\in S', u\in V} |\mathcal{L}_s(u)| \leq c_2\log n\right) \geq 1 - n^{-2}.$$

Together with the upper bound for the number of broadcast graphs in which a node u occurs (4.5), we obtain the following upper bound for $\max_{u\in V}|\mathcal{L}(u)|$:

$$\mathbf{P}\left(\max_{u\in V} |\mathcal{L}(u)| \leq 6c_1\log n c_2\log n\right) \leq 1 - 2n^{-2}.$$

Consider a path $P = (v_0 = s, v_1, \ldots, v_\ell = v)$ from the root $s \in S'$ to v in the broadcast graph B_s. We want to bound the time until v becomes informed in the parallel case using

our coupling. Consider the substep at which $v_i, 0 \leq i \leq \ell$, becomes informed for the first time. At that moment, either v_{i+1} is already informed, in which case we consider the step at which v_{i+1} becomes informed for the first time, or v_{i+1} is not already informed. In the latter, we know from the coupling that v_i will transmit the information to v_{i+1} after at most $|\mathcal{L}(v_i)| \leq 6c_1 \log nc_2 \log n$ transmissions. Therefore, v_ℓ becomes informed after at most

$$(6c_1 \log nc_2 \log n)\ell \leq (6c_1 \log nc_2 \log n)2000 \log n = c_3 \log^3 n$$

transmissions. Regardless of where we are on the path P, a transmission from the current node on P occurs in a given substep with probability $1/n$. Hence the number of time steps until the path P is completed can be upper bounded by $1/n$ times the sum of $c_3 \log^3 n$ independent geometric random variables each of which has parameter $1/n$. Hence by a Chernoff bound (cf. Lemma A.4), it follows that the path P is completed within $c_4 \log^3 n$ steps with probability at least $1 - n^{-3}$ for some constant $c_4 > 0$. Now for every node v for which there is a node s with $v \in B_s$, we can take the path from s to v in B_s. Taking the union bound over all possible nodes v, we obtain that the probability that the parallel case informs all the nodes informed by the independent case within $c_4 \log^3 n$ time steps is at least $1 - n^{-2}$. Combining this with (4.6), we get

$$\begin{aligned}
&\mathbf{P}\left(|I_{t+16+c_4\log^3 n} \setminus I_t| \geq \frac{(1/4)|S'|d/4}{6c_1 \log n}\right) \\
&\geq \mathbf{P}\left(\left(|I_{t+16+c_4\log^3 n} \setminus I_t| \geq \frac{(1/4)|S'|d/4}{6c_1 \log n}\right) \cap (T \leq 16)\right) \\
&\geq \mathbf{P}\left(\left(|I_{t+T+c_4\log^3 n} \setminus I_t| \geq \frac{(1/4)|S'|d/4}{6c_1 \log n}\right) \cap (T \leq 16)\right).
\end{aligned}$$

We know that $\mathbf{P}(T \leq 16) \geq \frac{1}{2}$ and that the event $\left\{|I_{t+T+c_4\log^3 n} \setminus I_t| \geq \frac{(1/4)|S'|d/4}{6c_1 \log n}\right\}$ is independent of T if we fix $|S'| = \lceil \frac{1}{d}|S| \rceil$ (we deal with the complementary case in a moment using (4.4)). We obtain

$$\begin{aligned}
&\mathbf{P}\left(\left(|I_{t+T+c_4\log^3 n} \setminus I_t| \geq \frac{(1/4)|S'|d/4}{6c_1 \log n}\right) \cap (T \leq 16)\right) \\
&= \mathbf{P}(T \leq 16)\,\mathbf{P}\left(\left|I_{t+T+c_4\log^3 n} \setminus I_t\right| \geq \frac{(1/4)|S'|d/4}{6c_1 \log n} \;\middle|\; T \leq 16\right) \\
&\geq \frac{1}{2}\,\mathbf{P}\left(\left|I_{t+T+c_4\log^3 n} \setminus I_t\right| \geq \frac{(1/4)|S'|d/4}{6c_1 \log n}\right) \\
&\geq \frac{1}{2}\left(\mathbf{P}\left(|\cup_{s\in S'} B_s| \geq \frac{(1/4)|S'|d/4}{\min\{6c_1 \log n, |S'|\}}\right) - \mathbf{P}\left(I_{t+T+c_4\log^3 n} \not\supseteq \cup_{s\in S'} B_s\right)\right) \\
&\geq \frac{1}{2}\left(1 - \exp(-1/4) - 2n^{-2}\right).
\end{aligned}$$

Recall that $\mathbf{P}\left(S' \geq \frac{1}{d}|S|\right) \geq \frac{1}{7}$ (cf. (4.4)); therefore,

$$\mathbf{P}\left(|I_{t+16+c_4\log^3 n} \setminus I_t| \geq \frac{(1/16)|S|}{6c_1 \log n}\right) \geq \frac{1}{14}\left(1 - \exp(-(1/4)) - 2n^{-2}\right).$$

Using the vertex expansion of G, we obtain $|S| \geq \frac{1}{2}|\Gamma(I_t) \setminus I_t| \geq \frac{\alpha}{2}|I_t|$, which yields

$$\mathbf{P}\left(|I_{t+16+c_4\log^3 n} \setminus I_t| \geq \frac{(1/32)\alpha|I_t|}{6c_1 \log n}\right) \geq \frac{1}{14}\left(1 - \exp(-(1/4)) - 2n^{-2}\right).$$

Let W be the indicator random variable which is 1 if $|I_{t+16+c_4\log^3 n} \setminus I_t| \geq \frac{(1/32)\alpha|I_t|}{6c_1 \log n}$. We are going to iterate this procedure many times and will denote by W_i the event $\{W = 1\}$ for the ith iteration. From the calculations above, we obtain that $\mathbf{P}(W_i = 1) \geq c_5$ for some constant $c_5 > 0$ uniformly over all realizations of the random variables W_j, $j \neq i$. We assume that if $W_i = 0$ no node becomes informed during the ith iteration and that $W_i = 1$ implies that the set of informed nodes increases by a factor of $\Omega(\alpha/\log n)$ during the ith iteration. Therefore, in order to increase the set of informed nodes from $d(8/\alpha)$ to $n/2+1$ it suffices to obtain $O(\log^2 n/\alpha)$ values of i for which $W_i = 1$. Since the lower bound $\mathbf{P}(W_i = 1) \geq c_5$ holds for all i independently of the other W_j's, we can apply a Chernoff bound and find a number $k_2 = O(\log^2 n/\alpha)$ such that

$$\mathbf{P}\left(|I_{t+k_2(16+c_4\log^3 n)}| \geq n/2+1\right) \geq 1 - n^{-1}.$$

Therefore with high probability $\tau_3 \leq Y_1 + k_2(16+c_4\log^3 n) = O((1/\alpha)\log^5 n)$. This completes the proof of Lemma 4.13. □

4.2 A vertex expander with large broadcast time

In this section we give the proof of Theorem 4.5 and show the existence of a regular graph with constant vertex expander for which the push algorithm takes $\Omega(\log^2 n)$ steps to inform all nodes of the graph.

First, we recall the following result of Lubotzky, Phillips and Sarnak [63] about the explicit construction of Ramanujan graphs. (Recall that the *girth* of a graph is the length of its shortest cycle.)

Theorem 4.14 (Lubotzky, Phillips and Sarnak [63])**.** *For any pair of unequal prime numbers p, q, both congruent to 1 mod 4, there exists a $(p+1)$-regular vertex transitive graph $G = G(p,q)$ with $\Theta(q(q^2-1))$ nodes such that*

- *the girth of G is $2\log_p q - \log_p 4$,*
- *the second largest eigenvalue of the adjacency matrix in absolute value λ_2, satisfies $|\lambda_2| \leq 2\sqrt{p}$.*

We use this theorem to show the existence of a vertex expander with large girth, which we use in our proof of Theorem 4.5. The result below is well known, but we give a proof for the sake of completeness.

Lemma 4.15. *There is an infinite, increasing sequence of values $(n_k)_{k=1}^{\infty}$ such that for each k, there exists a n_k-node graph G such that: G has girth $\Omega(\log n_k)$, the degree is constant, and the vertex expansion α is constant.*

Proof. We will use Theorem 4.14. Choosing $p = O(1)$ and $q = \Theta(n)$, it follows that $G = G(p, q)$ has girth $\Omega(\log n)$ and bounded degree. It only remains to verify that the upper bound on λ_2 implies that G has vertex expansion > 0. To this end, we use the following bound from Tanner [84]. This bound states that for any subset $S \subseteq V$ of a d-regular graph, we have

$$|\Gamma(S)| \geq \frac{d^2 |S|}{\lambda_2^2 + (d^2 - \lambda_2^2)|S|/n}.$$

Therefore if $\lambda_2 \leq C\sqrt{d}$ for some constant $0 < C < \sqrt{d}$,

$$|\Gamma(S) \setminus S| \geq \frac{d^2 |S|}{\lambda_2^2(1 - |S|/n) + d^2|S|/n} - |S| = |S| \left(\frac{1}{C^2/d(1 - |S|/n) + |S|/n} - 1 \right),$$

where the right factor is larger than zero, since $C^2/d < 1$ by assumption on C. This ends the proof of Lemma 4.15. □

Now we recall the definition of the Cartesian product of two graphs.

Definition 4.16. *Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be two graphs. The Cartesian product $G = G_1 \times G_2$ is the graph with node set $V = V_1 \times V_2$ and and edges between each pair of nodes (u_1, u_2) and (v_1, v_2) for which either $\{u_1, v_1\} \in E_1$ and $u_2 = v_2$ or $\{u_2, v_2\} \in E_2$ and $u_1 = v_1$.*

We use the Cartesian product of the graph from Lemma 4.15 with a complete graph. The lemma below shows that the resulting graph is also a vertex expander.

Lemma 4.17. *Let n_1 and n_2 be arbitrary integers. Let G_{n_1} be a graph with n_1 nodes and vertex expansion α. Then $G_{n_1} \times K_{n_2}$ is a graph with vertex expansion at least $\frac{\alpha^2}{64}$, where K_{n_2} is the complete graph on n_2 nodes.*

Proof. We first make the following simple claim about the vertex expansion of large subsets.

Claim. *Let G be any graph with n nodes and vertex expansion α. Then, for any $m \leq (1 + \frac{\alpha}{4})\frac{n}{2}$, $\partial_m \geq \frac{3}{5}\alpha m$, where $\partial_m = \min_{S \subseteq V : |S| = m} \partial(S)$.*

Let us first prove this claim. We know that for any $X \subseteq V$ and $|X| = \frac{n}{2}$, it holds that $|\Gamma(X)\backslash X| \geq \alpha\frac{n}{2}$. Clearly, $|\Gamma(X)\backslash X|$ decreases by at most one if the set X is enlarged by one node. More generally, for any $\frac{n}{2} \leq m < n$, it holds that $\partial_m \geq \alpha\frac{n}{2} - (m - \frac{n}{2})$. Hence, for $\frac{n}{2} \leq m \leq (1 + \frac{\alpha}{4})\frac{n}{2}$ we have

$$\partial_m \geq \alpha\frac{n}{2} - \left(m - \frac{n}{2}\right) \geq \frac{3}{4}\alpha\frac{n}{2} \geq \frac{3}{4}\alpha\frac{m}{1+\frac{\alpha}{4}} \geq \frac{3}{5}\alpha m,$$

and thus for any $m \leq (1 + \frac{\alpha}{4})\frac{n}{2}$, $\delta_m \geq \min\{\frac{3}{5}\alpha, \alpha\}m \geq \frac{3}{5}\alpha m$, which yields the claim.

Now let $X \subseteq V$ be any subset with $|X| \leq \frac{n_1 n_2}{2}$. We will frequently use the fact that $G = G_{n_1} \times K_{n_2}$ consists of n_1 node-disjoint K_{n_2} subgraphs which are connected by edges induced by G_{n_1}. We distinguish between two cases:

1. A portion of at most $\frac{\alpha}{8}$ of nodes in X is located in a complete graph K_{n_2} which contains less than $(1 - \frac{\alpha}{8})n_2$ nodes of X. Let X' be the set of these nodes. Note that every such node in X' has at least $\frac{\alpha}{8}$ neighbors, considering only edges induced by K_{n_2}. Denote by $\Gamma_{K_{n_2}}(Y)$ the set of neighbors of some set $Y \subseteq V$ which are adjacent to some node in Y by an edge induced by K_{n_2}. We have

   $$|\Gamma(X)\backslash X| \geq |\Gamma_{K_{n_2}}(X)\backslash X| \geq |\Gamma_{K_{n_2}}(X')\backslash X| = |\Gamma_{K_{n_2}}(X')\backslash X'| \geq |X'|\frac{\alpha}{8} \geq \frac{\alpha^2}{64}|X|,$$

   and the first case is finished.

2. A portion of at least $1 - \frac{\alpha}{8}$ of nodes in X is located in a complete graph K_{n_2} which contains more than $(1 - \frac{\alpha}{8})n_2$ nodes of X. In other words, $|X'| \geq (1 - \frac{\alpha}{8})|X|$. Let y be the number of K_{n_2}'s containing more than $(1 - \frac{\alpha}{8})n_2$ nodes. We claim that $y \leq \frac{1+\frac{\alpha}{4}}{2}n_1$. Assuming the converse, we obtain

   $$|X| \geq y\left(1 - \frac{\alpha}{8}\right)n_2 > \frac{1+\frac{\alpha}{4}}{2}n_1\left(1 - \frac{\alpha}{8}\right)n_2 = \frac{1 + \frac{\alpha}{4} - \frac{\alpha}{8} - \frac{\alpha^2}{32}}{2}n_1 n_2 \geq \frac{n_1 n_2}{2},$$

   which contradicts the assumption $|X| \leq \frac{n_1 n_2}{2}$. We can also lower bound y by $y \geq \frac{|X'|}{n_2} \geq \frac{(1-\frac{\alpha}{8})|X|}{n_2} \geq \frac{7}{8}\frac{|X|}{n_2}$. Now, as $y \leq \frac{1+\frac{\alpha}{4}}{2}n_1$, the Claim above implies that the number of K_{n_2}, which are neighbors of a K_{n_2} containing more than $(1 - \frac{\alpha}{8})n_2$ nodes of X, is at least $\frac{3}{5}\alpha y$. Hence,

   $$|\Gamma(X')\backslash X'| \geq \frac{3}{5}\alpha y\left(1 - \frac{\alpha}{8}\right)n_2 \geq \frac{3}{5}\alpha\frac{7}{8}\frac{|X|}{n_2}\frac{7}{8}n_2 \geq \frac{3}{8}\alpha|X|,$$

   and finally

   $$\begin{aligned}|\Gamma(X)\backslash X| \geq |\Gamma(X')\backslash X| &\geq |\Gamma(X')\backslash X'| - |X\backslash X'| \\ &\geq \frac{3}{8}\alpha|X| - \frac{\alpha}{8}|X| = \frac{1}{4}\alpha|X| \geq \frac{1}{64}\alpha^2|X|,\end{aligned}$$

and the lemma follows. □

Remark 4.18. We suspect that the dependency on α in Lemma 4.17 could be made linear. But, as α is constant in our application, the quadratic dependency suffices.

The proof of Theorem 4.5 is concluded with the next lemma, which shows that the push algorithm takes $\Omega(\log^2 n)$ steps in the graph obtained from the Cartesian product.

Lemma 4.19. *Let $G_{\sqrt{n}}$ be an arbitrary d-regular graph with $\sqrt{n}$ nodes, degree $d = O(1)$ and girth $c \log n$ for some constant $c > 0$. Then for the graph $G = K_{\sqrt{n}} \times G_{\sqrt{n}}$, the push algorithm takes $\Omega(\log^2 n)$ steps with high probability.*

Proof. Let $s = (s_1, s_2)$ with s_1 in $K_{\sqrt{n}}$, and s_2 in $G_{\sqrt{n}}$ be the initially informed node. Note that the subgraph in $G_{\sqrt{n}}$ formed by s_2 and all nodes within distance at most $\min\{\frac{c}{2}\log n - 1, \frac{1}{16}\log_\Delta n\} - 1 = c' \log n$ forms a tree with root s_2. For $0 \le i \le c' \log n$, let $\mathcal{B}_i$ be the set of nodes in $G_{\sqrt{n}}$ with distance i to s; so $\mathcal{B}_0 = \{s_2\}$. We will lower bound the time X until the first node in $K_{\sqrt{n}} \times \mathcal{B}_{c' \log n}$ becomes informed. Let $X_i, 1 \le i \le \lceil (c'/2) \log n \rceil$ be the number of steps until the first node in $K_{\sqrt{n}} \times \mathcal{B}_{2i}$ becomes informed after $K_{\sqrt{n}} \times \mathcal{B}_{2i-2}$ contains one informed node. By definition, $X = \sum_{i=1}^{\lceil (c'/2) \log n \rceil} X_i$.

Consider an arbitrary but fixed X_i. In order to lower bound X_i, we may assume that $K_{\sqrt{n}} \times \mathcal{B}_{2i-2}$ is completely informed. Fix an arbitrary $u_{2i-1} \in \mathcal{B}_{2i-1}$ and let $Y(t)$ be the set of informed nodes in $K_{\sqrt{n}} \times u_{2i-1}$ at step t. By a Chernoff bound, we obtain with probability $1 - n^{-4}$ that $K_{\sqrt{n}} \times \mathcal{B}_{2i-2}$ informs at most $24 \log n$ nodes in $K_{\sqrt{n}} \times (\Gamma(u_{2i-1}) \cap \mathcal{B}_{2i-1})$ in one step. Moreover, a set of informed nodes can only inform a set of the same size in one step. Therefore, $|Y(t+1)| \le 2|Y(t)| + 24 \log n$ with probability $1 - n^{-3}$, provided that no node in $K_{\sqrt{n}} \times (\Gamma(\Gamma(u_{2i-1})) \cap \mathcal{B}_{2i})$ is informed. Assuming that no node in $K_{\sqrt{n}} \times (\Gamma(\Gamma(u_{2i-1})) \cap \mathcal{B}_{2i})$ becomes informed and applying the union bound, we conclude that $|Y(\frac{1}{96}\log n)| \le n^{\frac{1}{16}}$ with probability $1 - n^{-2}$. As long as $|Y(t)| \le n^{\frac{1}{16}}$ in a step t, the probability that there is a node in $Y(t)$ which informs a node in $K_{\sqrt{n}} \times (\Gamma(\Gamma(u_{2i-1})) \cap \mathcal{B}_{2i})$ in one step is bounded by $n^{\frac{1}{16}} \frac{\Delta}{\binom{\sqrt{n}}{2} + \Delta} \le n^{-\frac{3}{8}}$. By the Union bound over $\frac{1}{96}\log n$ steps, no node in $K_{\sqrt{n}} \times (\Gamma(\Gamma(u_{2i-1})) \cap \mathcal{B}_{2i})$ becomes informed during these $\frac{1}{96}\log n$ steps with probability at least $1 - n^{-\frac{1}{4}}$. Using a Union bound over all nodes in $\mathcal{B}_{2i}$ ($|\mathcal{B}_{2i}| \le \Delta^{2i} \le n^{\frac{1}{8}}$), this establishes $X_i \ge \frac{1}{96}\log n$ with probability $1 - |\mathcal{B}_{2i}| n^{-\frac{1}{4}} \ge 1 - n^{-\frac{1}{8}}$. Finally, by applying a Chernoff bound on $\sum_{i=1}^{\lceil (c'/2) \log n \rceil} X_i \le X$, we conclude that $X = \Omega(\log^2 n)$ with high probability, and the claim follows. □

4.3 Applications to other models

In this section, we give some applications of our results to other processes such as the cover time of random walks and random subgraph generation.

4.3.1 Cover time

Let G be a graph and $\mathsf{C}(G)$ be the expected time for a random walk to visit all nodes of G (we assume that the random walk starts at the node maximizing this expectation). We first state the known upper bound on the cover time in terms of the vertex expansion, which is a slight reformulation of [17, Theorem 12].

Theorem 4.20 (cf. [17, Theorem 12])**.** *For any graph $G = (V, E)$ with minimum degree δ and vertex expansion α, $\mathsf{C}(G) = O\left(\frac{1}{\alpha^2}\frac{|E|}{\delta}\log n\right)$.*

Note that for regular graphs, the upper bound becomes $O\left(\frac{1}{\alpha^2} n \log n\right)$. In [17], the authors remark on Theorem 4.20 that "it is unknown whether the quadratic dependence on $1/\alpha$ is necessary." To address this question, we make use of the following relation between the runtime of the push algorithm and the cover time, shown in [35]. Let τ be the expected runtime of the push algorithm on G considering the worst-case initial node. Then, as in [35, Theorem 3.8], we have for any graph G that $\mathsf{C}(G) = O\left(\frac{|E|}{\delta}\log n\tau\right)$.

Combining this with Theorem 4.3 we immediately get the following result.

Corollary 4.21. *For any regular graph $G = (V, E)$ with vertex expansion α, $\mathsf{C}(G) = O\left(\frac{1}{\alpha} n \log^6 n\right)$.*

We believe that this result represents some progress towards the question raised in [17] concerning the dependency on the vertex expansion. In addition, we now prove that Corollary 4.21 is tight up to logarithmic factors in the following sense. For any integer k, let C_k be a cycle with k nodes. For any x that divides n, consider the graph $G = K_x \times C_{n/x}$ (cf. Definition 4.16). Note that $D = \Theta(n/x)$, and since G is a vertex transitive graph, [6, Theorem 3.2] implies that G has vertex expansion $\alpha = \Omega(1/D) = \Omega(x/n)$. Plugging this into Corollary 4.21, we obtain an upper bound of $O(n^2/x \log^6 n)$.

We argue that this bound is tight up to logarithmic factors. The random walk has to spend $\Theta(x)$ steps in expectation in each clique before leaving it. Clearly, the random walk has to traverse at least $D = \Theta(n/x)$ different cliques to cover all nodes. Since the cliques are arranged in a cycle, we obtain that the cover time is $\Omega(xD^2) = \Omega(n^2/x)$, which proves the tightness up to logarithmic factors.

4.3.2 Random subgraphs

Given a parameter $p \in [0, 1]$, we construct a *random subgraph* $\mathcal{G}_p$ of a graph G by adding all nodes of G to $\mathcal{G}_p$, and adding each edge of G to $\mathcal{G}_p$ independently with probability p. This is related to the so-called bond percolation process on graphs, but in this section we will be mostly concerned with whether $\mathcal{G}_p$ is a *connected* subgraph of G.

Bond percolation has been studied on many lattices [48], and the main focus of existing work on bond percolation regards the emergence of the so-called giant component. An

example of a random subgraph as described above is the so-called Erdős-Rényi random graphs [13], which is a random subgraph of the complete graph. Also, the expansion properties of random subgraphs have been studied recently [3, 7]. The corollary below shows how our results can be interpreted in terms of random subgraphs.

Corollary 4.22. *Let G be any d-regular graph with n nodes and vertex expansion α, and $\mathcal{G}_p$ be the subgraph obtained via the process described above with $p = \min\{1, 4C\frac{\log^5 n}{\alpha d}\}$, where C is the constant from Theorem 4.3. Then with high probability, $\mathcal{G}_p$ is connected and has diameter $O((1/\alpha)\log^5 n)$.*

Proof. We can assume that $4C\frac{\log^5 n}{\alpha d} \leq 1$, since otherwise the claim is trivial. The proof uses a standard technique known also as "Poissonization". Let POIS be a continuous-time version of the push algorithm defined as follows. Let $\mathcal{P}_n$ be a Poisson process with rate n. Whenever the clock ticks, we choose a random node $v \in V$ and let this node forward the information to a random neighbor if v was informed earlier. Using a Chernoff bound for Poisson random variables (cf. Lemma A.3), it follows that the runtime of this process is the same as the runtime of the sequential algorithm used in the proof of Theorem 4.3 up to factor of 2 with high probability. Hence we can use Theorem 4.3 to conclude that this continuous-time version of the push algorithm has finished before time

$$2C\frac{\log^5 n}{\alpha d}$$

with high probability, where $C > 0$ is the constant from Theorem 4.3.

By standard properties of Poisson processes, POIS can be also regarded as the union of $2|E|$ independent Poisson processes for every $(u, v) \in V \times V$ with $\{u, v\} \in E$, where each such Poisson process has rate 1. Now fix an (undirected) edge $e = \{u, v\} \in E$. The probability that (u, v) or (v, u) is used in POIS is at most

$$2\,\mathbf{P}\left(\mathsf{Pois}\left(2C\frac{\log^5 n}{\alpha d}\right) > 0\right) \leq 2\left(1 - \exp\left(-2C\frac{\log^5 n}{\alpha d}\right)\right) \leq 4C\frac{\log^5 n}{\alpha d},$$

where we have used the fact that $1 - \exp(x) \leq x$ for any $x \in \mathbb{R}$. Note that the events $\{$ POIS uses $\{u, v\}$ $\}$ with $\{u, v\} \in E$ are independent. Hence there is a coupling between an execution of POIS and a randomly chosen subgraph of G such that if a transmissions occurs along the edge $\{u, v\}$ in either direction within the first $2C\frac{\log^5 n}{\alpha d}$ steps, then this edge also appears in the random subgraph. In particular, if POIS informs all nodes within $2C\frac{\log^5 n}{\alpha d}$ steps, then by the coupling above, the random subgraph is also connected. This finishes the proof. □

This corollary essentially says that sparse random subgraphs of vertex expanders are connected and have small diameter. However, the random subgraph may have small vertex expansion. Consider for example the graph G which consists of two cliques of size $n/2$

connected by a matching. This yields a $(n/2)$-regular graph with n nodes and constant vertex expansion. However, for $p = O(\frac{\text{polylog}(n)}{n})$, only $O(\text{polylog}(n))$ edges of the matching are present in the random subgraph $\mathcal{G}_p$. Hence, the vertex expansion is reduced vastly from a constant in G to $O(\frac{\text{polylog}(n)}{n})$ in the random subgraph $\mathcal{G}_p$.

4.4 Extension to non-regular graphs

We now extend some of our results to non-regular graphs. We first consider the push-pull algorithm, and then derive some stronger results for random subgraphs.

Push-pull algorithm

We first recall the push-pull algorithm. In this algorithm, at each step, not only does each informed node choose a neighbor uniformly at random and send the information to that neighbor, but also, each uninformed node chooses a neighbor uniformly at random and becomes informed if the neighbor is informed. We have already observed that for non-regular graphs it is essential to consider both push and pull transmissions. We now prove a sublinear runtime bound for the push-pull algorithm on arbitrary expanders. We first prove the following lemma that corresponds to Lemma 4.7.

Lemma 4.23. *Let G be an arbitrary, possibly non-regular graph with vertex expansion α. Suppose that initially only one node v with degree $\deg(v)$ is informed. Then after $2\deg(v)^{1/8}$ steps, at least $(1/4)\deg(v)^{1/4}$ nodes become informed with probability at least $1/2$.*

Proof. If half of the neighbors of v have degree at most $\deg(v)^{7/8}$, then each of those nodes becomes informed via pull within the first $\deg(v)^{1/8}$ steps with probability at least $(1/2)\deg(v)^{-6/8}$. Hence the expected number of informed nodes is at least

$$\deg(v)/2(1/2)(1/2)\deg(v)^{-6/8} = (1/4)\deg(v)^{1/4}.$$

Since the number of informed nodes is binomially distributed, it follows that with probability at least $1/2$, at least $(1/4)\deg(v)^{1/4}$ nodes become informed via pull.

The other case to consider is when half of the neighbors of v have degree at least $\deg(v)^{7/8}$. Then after $\deg(v)^{1/8}$ steps, a set $\mathcal{I}$ which consists of at least $\frac{1}{2}\deg(v)^{1/8}$ of such neighbors become informed via push with probability at least $3/4$. The probability that during the next $\deg(v)^{1/8}$ steps, two nodes in $\mathcal{I}$ inform the same node via push can be bounded from above by

$$\binom{\deg(v)^{1/8}}{2}\binom{\deg(v)^{1/8}}{2}\frac{1}{\deg(v)^{7/8}} \leq \frac{1}{4\deg(v)^{3/8}}.$$

If no two nodes in $\mathcal{I}$ inform the same node within the next $\deg(v)^{1/8}$, then all the nodes in $\mathcal{I}$ together inform at least $\frac{1}{2}\deg(v)^{1/8}\deg(v)^{1/8} = \frac{1}{2}\deg(v)^{1/4}$ with probability at least $1 - 2 \cdot 1/4 = 1/2$. □

As in Definition 4.9, we say that a node v is a *friend* of a set V if the push-pull algorithm initialized with v informs at least $(1/8)\deg(v)^{1/4}$ nodes in V. Using Lemma 4.23, we conclude as in the regular case that for any set $X \subseteq V$, a node v is a friend of X or a friend of $V \setminus X$.

Theorem 4.24. *Let G be an arbitrary, possibly non-regular graph with vertex expansion α. Then the push-pull algorithm informs all nodes in time $O((1/\alpha)n^{1-\epsilon})$.*

Proof. We divide the time into two phases. The first phase covers the case where $1 \leq |I_t| \leq n^{7/8}$ and the second phase covers $n^{7/8} \leq |I_t| \leq n/2$.

For the first phase, consider any step t with $1 \leq |I_t| \leq n^{7/8}$. Let $|I_t| = x$. By definition of vertex expansion, $|\Gamma(I_t) \setminus I_t| \geq \alpha x$. Our goal is now to show that with probability $\alpha/2$, at least one node in $V \setminus I_t$ becomes informed by a push transmission. This is clearly true, if there is a node $u \in I_t$ with $\deg(u) \geq 2x$, since then half of the neighbors of u are in $V \setminus I_t$. Hence suppose that for all nodes $u \in I_t$, $\deg(u) \leq 2x$. This means every node in $u \in \Gamma(I_t) \setminus I_t$ becomes informed with probability at least $\frac{1}{2x}$. Let $Z_u = 1$ if $u \in I_{t+1}$ and $Z_u = 0$ otherwise. Clearly, for any subset $S \subseteq \Gamma(I_t) \setminus I_t$ we have

$$\mathbf{P}\left(Z_u = 0 \,\middle|\, \bigcap_{v \in S} (Z_v = 0)\right) \leq \mathbf{P}\left(Z_u = 0\right),$$

and hence,

$$\mathbf{P}\left(\bigcap_{v \in \Gamma(I_t) \setminus I_t} (Z_v = 0)\right) \leq \prod_{v \in \Gamma(I_t) \setminus I_t} \mathbf{P}\left(Z_v = 0\right) \leq \left(1 - \frac{1}{2x}\right)^{\alpha x} \leq e^{-\alpha/2} \leq 1 - \frac{\alpha}{8},$$

as $\alpha \leq 1$. Hence the expected time to finish the first phase is $8\, n^{7/8}/\alpha$.

Let us now analyze the second phase. First note that we may assume that half of the nodes in $\Gamma(I_t) \setminus I_t$ have degree at least $n^{6/8}$, as otherwise all of them become informed by a pull transmission within $O(n^{6/8} \log n)$ steps.

Consider again first the case where there is a node in $u \in I_t$ with $\deg(u) \geq 4x$. Then at least $3x$ of its neighbors are outside I_t. Moreover, we may assume that a quarter of those neighbors are not friends of I_t (otherwise a quarter of those neighbors become informed within $O(n^{1/8})$ steps by Lemma 4.23). Then with high probability, at least one such node v becomes informed within $O(\log n)$ steps. After an additional $\deg(v)^{1/8}$ rounds, this neighbor informs by itself $\deg(v)^{1/4}$ nodes by Lemma 4.23, as it is a friend of $V \setminus I_t$.

Hence we may now assume that all nodes in I_t have degree at most $4x$. However, then we can conclude with the same calculations as above that, with probability $\alpha/2$, at least

one node $u \in \Gamma(I_t) \setminus I_t$ which has degree at least $n^{6/8}$ and is a friend of $V \setminus I_t$ becomes informed by push. Hence after an additional $O(\log n)$ rounds, this node informs by itself $(1/4)(n^{6/8})^{1/4} = (1/4)n^{6/32}$ nodes within $2(n^{6/8})^{1/8} = 2n^{3/32}$ steps due to Lemma 4.23.

Note that in all these cases, there is a number y such that within y steps we inform at least $y^{1+1/8}$ nodes. Moreover, note that $y \geq n^{3/32}$. Using a similar analysis as in Lemma 4.12, it follows that after $O(n^{1-\epsilon})$ steps ($\epsilon = 3/32 \cdot 1/8 = 3/256$), the number of informed nodes is at least $n/2 + 1$ with high probability. As in the proof of Theorem 4.3, this shows that after $2O(n^{1-\epsilon})$ steps all nodes are informed with high probability. □

Random subgraphs

We now return to the problem of constructing a random subgraph of G, but without assuming that G is regular.

We use the concept of resistance in electrical networks. We view a graph G as an electrical network with each edge corresponding to a unit resistor. The resistance $\mathsf{R}(u, v)$ between two nodes u and v is defined as the voltage difference induced between u and v by passing a current flow of one between them [31]. Note that $\mathsf{R}(u, v) = \mathsf{R}(v, u)$.

We first note the following slight extension of a result of [17, Theorem 5.2].

Corollary 4.25 (cf. [17, Theorem 5.2])**.** *For any graph G with vertex expansion α and minimum degree d,*

$$\mathsf{R}(s, t) \leq \frac{24}{\alpha^2(\widetilde{d}(s, t) + 1)},$$

where $\widetilde{d}(s, t) = \min\{\deg(s), \deg(t)\}$.

The proof of Corollary 4.25 is exactly the same as the proof of [17, Theorem 5.2], except that the minimum degree of G is replaced by the "local" minimum degree of s and t, $\widetilde{d}(s, t)$.

We also observe the following basic fact about resistances.

Lemma 4.26. *For any graph G and two nodes s and t, we have $\mathsf{R}(s, t) \geq \frac{1}{\widetilde{d}(s,t)}$. Moreover, $\sum_{\{i,j\} \in E} \mathsf{R}(i, j) \geq \frac{n}{2}$.*

Proof. It is known [1] that $\mathsf{R}(s, t)$ is the minimum cost of a unit flow $f_{s,t}$ from s to t with cost function $\sum_{e \in E} (f_{s,t}(e))^2$, where $f_{s,t}(e)$ is the amount of flow routed over edge e by $f_{s,t}$. Assume without loss of generality that $\deg(s) \leq \deg(t)$. Then

$$\sum_{e \in E} (f_{s,t}(e))^2 \geq \sum_{j \in N(s)} (f_{s,t}(s, j))^2 .$$

As $f_{s,t}$ is a unit flow from s to t, $\sum_{j \in N(s)} f_{s,t}(s, j) = 1$ and hence

$$\sum_{j \in N(s)} (f_{s,t}(s, j))^2 \geq \sum_{j \in N(s)} \left(\frac{1}{\deg(s)}\right)^2 \geq \frac{1}{\deg(s)},$$

which proves the first statement. For the second statement, note that

$$\sum_{\{i,j\}\in E} \mathsf{R}(i,j) \geq \sum_{\{i,j\}\in E} \frac{1}{\min\{\deg(i),\deg(j)\}}.$$

Since $\frac{1}{\min\{a,b\}} \geq \frac{1}{2}\left(\frac{1}{a}+\frac{1}{b}\right)$, we can lower bound this term by

$$\sum_{\{i,j\}\in E} \mathsf{R}(i,j) \geq \sum_{\{i,j\}\in E} \frac{1}{2}\left(\frac{1}{\deg(i)}+\frac{1}{\deg(j)}\right) = \frac{1}{2}n.$$

□

We use the following theorem, which is a direct consequence of a sparsification result of Spielman and Srivastava [81].

Theorem 4.27 ([81, Theorem 1])**.** *Let $C \geq 1$ be a sufficiently large constant. Consider a random subgraph H which is obtained from G by picking $Cn\log n$ edges $e=\{u,v\}$, each of which is chosen with probability p_e proportional to $\mathsf{R}(u,v)$ (with replacement). Then with probability at least $1/2$, H is connected.*

We point out that [81] proved a stronger result which relates the expansion of a (weighted version) of H to the expansion of G. We now use Theorem 4.27 and Lemma 4.26 to prove the following result.

Corollary 4.28. *Let F be the subgraph obtained when every node of a graph G with vertex expansion α chooses $4C(\log n/\alpha^2)$ incident edges uniformly at random and with replacement ($C>0$ is the constant from Theorem 4.27). Then H is connected with probability $1/2-o(1)$.*

Proof. In this corollary, we relate several randomized procedures that all construct a subgraph of G and are defined as follows:

1. G_1. Every edge $e=\{u,v\}$ is associated to an independent Poisson process with rate $\frac{\mathsf{R}(u,v)}{\sum_{\{i,j\}\in E}\mathsf{R}(i,j)}$. The subgraph G_1 is obtained by taking every edge whose Poisson process ticks until time $2Cn\log n$.

2. G_2. Every edge $e=\{u,v\}$ is associated to an independent Poisson process with rate $\frac{48}{\alpha^2\min\{\deg(u),\deg(v)\}}$. The subgraph G_2 is obtained by taking every edge whose Poisson process ticks until time $2Cn\log n$.

3. G_3. Every node is associated to an independent Poisson process with rate n. Whenever, this process ticks at a certain node u, then u chooses a neighbor uniformly at random. The subgraph G_3 is obtained by taking every edge until time $2C\frac{48}{\alpha^2}n\log n$.

The strategy of the proof is to construct a coupling of H, G_1, G_2 and G_3 such that with probability $1 - o(1)$ the following inclusions hold:

$$H \overset{(1)}{\subseteq} G_1 \overset{(2)}{\subseteq} G_2 \overset{(3)}{\subseteq} G_3 \overset{(4)}{\subseteq} F. \tag{4.7}$$

As H is connected with probability at least $1/2$ by Theorem 4.27, this implies that F is also connected with probability $1/2 - o(1)$.

In the following, we shall make use of some standard properties of Poisson processes (cf. [66, Chapter 8]). One of these properties is that the union of two independent Poisson processes with respective rates λ_1 and λ_2 forms a Poisson process with rate $\lambda_1 + \lambda_2$. The following converse also holds. Suppose that there is a Poisson process with rate λ and everytime the Poisson process ticks, we classify this step as type A with probability p and otherwise as type B. Then the original Poisson process restricted to those steps which are classified as A forms a Poisson process with rate λp, and the original Poisson process restricted to those steps which are classified as B forms a Poisson process with rate $\lambda(1-p)$. In addition, these two Poisson processes are independent.

We now proceed to prove each of the inclusions in (4.7) above.

Proof of (1). Clearly, the Poisson process involved in the construction of G_1 is equivalent to a Poisson process with rate 1 such that whenever this process ticks, we choose an edge $e = \{u, v\}$ with probability $p_e = \frac{\mathsf{R}(u,v)}{\sum_{\{i,j\}\in E} \mathsf{R}(i,j)}$. Using a Chernoff bound, we obtain that with probability $1 - n^{-1}$, the Poisson process ticks at least $Cn \log n$ times until time $2Cn \log n$. Coupling the choices of edges in the construction of H and in the construction of G_1, we obtain

$$\mathbf{P}\left(H \subseteq G_1\right) \geq 1 - n^{-1}.$$

Proof of (2). Combining Corollary 4.25 and Lemma 4.26, we obtain that

$$\frac{\mathsf{R}(u,v)}{\sum_{\{i,j\}\in E} \mathsf{R}(i,j)} \leq \frac{\frac{24}{\alpha^2 \min\{\deg(u),\deg(v)\}}}{\frac{n}{2}} \leq \frac{48}{\alpha^2 \min\{\deg(u), \deg(v)\}}.$$

Hence we can couple all Poisson processes involved in G_1 and G_2 such that whenever a process in G_1 ticks at some step $0 \leq t \leq 2C \log n$, then so does the corresponding process in G_2. Hence,

$$\mathbf{P}\left(G_1 \subseteq G_2\right) = 1.$$

Proof of (3). The model for constructing G_3 is equivalent to a model, where every edge $e = \{u, v\}$ is associated to an independent Poisson process with rate $\frac{1}{\deg(u)} + \frac{1}{\deg(v)} \geq \frac{1}{\min\{\deg(u),\deg(v)\}}$. Since a Poisson process with rate λ_1 until time t_1 is the same as a Poisson process with rate λ_2 until time t_2 if $\lambda_1 t_1 = \lambda_2 t_2$, we conclude that

$$\mathbf{P}\left(G_2 \subseteq G_3\right) = 1.$$

Proof of (4). Using a Chernoff bound for Poisson random variables (cf. Lemma A.3), it follows that every node u ticks not more than $4C\frac{48}{\alpha^2}\log n$ times with probability $1 - n^{-1}$. Since this holds for all nodes, we can couple the choices made by each node $u \in V$ in the construction of G_3 with the choices made by node u in the construction of F. Therefore, with probability at least $1 - n^{-1}$,

$$\mathbf{P}\left(G_3 \subseteq F\right) \geq 1 - n^{-1}.$$

□

We can also derive a similar result for random subgraphs for which the edges are picked in a centralized fashion.

Corollary 4.29. *Let F be the subgraph obtained when every edge $e = \{u, v\}$ of the graph G is included with probability $p_e = \min\{\frac{48}{\alpha^2}\frac{C\log n}{\min\{\deg(u),\deg(v)\}}, 1\}$. Then F is connected with probability $1/2 - o(1)$.*

Note that for non-regular graphs, it is essential to pick edges with different probabilities. Recall the graph which consists of two cliques of size $\sqrt{n}$ and $n - \sqrt{n}$ which are connected by a matching. This graph has constant vertex expansion, but if we pick $\Theta(n \log n)$ edges chosen uniformly at random, then with high probability no edge of the matching is picked and the random subgraph is disconnected.

Bibliography

[1] D. Aldous and J. Fill. Reversible Markov chains and random walks on graphs, in progress. The manuscript is available at `http://www.stat.berkeley.edu/~aldous/RWG/book.html`.

[2] N. Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986.

[3] N. Alon, I. Benjamini, and A. Stacey. Percolation on finite graphs and isoperimetric inequalities. *The Annals of Probability*, 32(3):1727–1745, 2004.

[4] N. Alon and J.H. Spencer. *The probabilistic method*. John Wiley & Sons, 3rd edition, 2008.

[5] C. Avin and G. Ercal. On the cover time and mixing time of random geometric graphs. *Theoretical Computer Science*, 380:2–22, 2007.

[6] L. Babai. Local expansion of vertex-transitive graphs and random generation in finite groups. In *23rd Annual ACM Symposium on Theory of Computing (STOC'91)*, pages 164–174, 1991.

[7] A. Bagchi, A. Bhargava, A. Chaudhary, D. Eppstein, and C. Scheideler. The effect of faults on network expansion. *Theory of Computing Systems*, 39(6):903–928, 2006.

[8] P. Balister, Z. Zheng, S. Kumar, and P. Sinha. Trap coverage: Allowing coverage holes of bounded diameter in wireless sensor networks. In *Proceedings of the 28th IEEE Conference on Computer Communications*, pages 19–25, 2009.

[9] I. Benjamini and A. Stauffer. Perturbing the hexagonal circle packing: a percolation perspective, 2011. Preprint at arXiv:1104.0762v1.

[10] P. Berenbrink, R. Elsässer, and T. Friedetzky. Efficient randomized broadcasting in random regular networks with applications in peer-to-peer systems. In *27th Annual ACM Principles of Distributed Computing (PODC'08)*, pages 155–164, 2008.

[11] A.M. Berezhovskii, Yu.A. Makhovskii, and R.A. Suris. Wiener sausage volume moments. *Journal of Statistical Physics*, 57:333–346, 1989.

[12] J. van den Berg, R. Meester, and D.G. White. Dynamic boolean model. *Stochastic Processes and their Applications*, 69:247–257, 1997.

[13] B. Bollobás. *Random Graphs*. Cambridge University Press, Cambridge, UK, second edition, 2001.

[14] B. Bollobás and O. Riordan. *Percolation*. Cambridge University Press, 2006.

[15] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. randomized gossip algorithms. *IEEE Transactions on Information Theory and IEEE/ACM Transactions on Networking*, 52(6):2508–2530, 2006.

[16] M. Bradonjić, R. Elsässer, T. Friedrich, T. Sauerwald, and A. Stauffer. Efficient broadcast on random geometric graphs. In *Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1412–1421, 2010.

[17] A.K. Chandra, P. Raghavan, W.L. Ruzzo, R. Smolensky, and P. Tiwari. The electrical resistance of a graph captures its commute and cover times. *Computational Complexity*, 6(4):312–340, 1997.

[18] F. Chierichetti, S. Lattanzi, and A. Panconesi. Rumor spreading in social networks. In *36th International Colloquium on Automata, Languages, and Programming (ICALP'09)*, pages 375–386, 2009.

[19] F. Chierichetti, S. Lattanzi, and A. Panconesi. Almost tight bounds on rumour spreading by conductance. In *41st Annual ACM Symposium on Theory of Computing (STOC'10)*, pages 399–408, 2010.

[20] F. Chierichetti, S. Lattanzi, and A. Panconesi. Rumour spreading and graph conductance. In *21th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'10)*, pages 1657–1663, 2010.

[21] Z. Ciesielski and S.J. Taylor. First passage times and sojourn times for Brownian motion in space and the exact Hausdorff measure of the sample path. *Transactions of the American Mathematical Society*, 103:434–450, 1962.

[22] A. Clementi, F. Pasquale, and R. Silvestri. MANETS: high mobility can make up for low transmission power. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP)*, 2009.

[23] C. Cooper and A. Frieze. The cover time of random geometric graphs. In *Proceedings of the 20th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2009.

[24] A. Czumaj and C. Sohler. Testing expansion in bounded-degree graphs. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 570–578, 2007.

[25] A. Czumaj and X. Wang. Fast message dissemination in random geometric ad-hoc radio networks. In *18th International Symposium on Algorithms and Computation (ISAAC'07)*, pages 220–231, 2007.

[26] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *6th Annual ACM Principles of Distributed Computing (PODC'87)*, pages 1–12, 1987.

[27] J. Díaz, D. Mitsche, and X. Pérez-Giménez. On the connectivity of dynamic random geometric graphs. In *Proceedings of the 19th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 601–610, 2008.

[28] S. Diggavi, M. Grossglauser, and D. Tse. Even one-dimensional mobility increases adhoc wireless capacity. *IEEE Transactions on Information Theory*, 51:3947–3954, 2005.

[29] B. Doerr, T. Friedrich, and T. Sauerwald. Quasirandom rumor spreading: Expanders, push vs. pull, and robustness. In *36th International Colloquium on Automata, Languages, and Programming (ICALP'09)*, pages 366–377, 2009.

[30] O. Dousse, C. Tavoularis, and P. Thiran. Delay of intrusion detection in wireless sensor networks. In *Proceedings of the 7th ACM International Conference on Mobile Computing and Networking (MobiCom)*, pages 155–165, 2006.

[31] P.G. Doyle and J.L. Snell. *Random walks and electric networks*. The Mathematical Association of America, 1984.

[32] A. Drewitz, J. Gärtner, A.F. Ramírez, and R. Sun. Survival probability of a random walk among a poisson system of moving traps, 2010. Preprint at arXiv:1010.3958v1.

[33] R.B. Ellis, J.L. Martin, and C. Yan. Random geometric graph diameter in the unit ball. *Algorithmica*, 47(4):421–438, 2007.

[34] R. Elsässer and T. Sauerwald. Broadcasting vs. mixing and information dissemination on Cayley graphs. In *24th International Symposium on Theoretical Aspects of Computer Science (STACS'07)*, pages 163–174, 2007.

[35] R. Elsässer and T. Sauerwald. Cover time and broadcast time. In *26th International Symposium on Theoretical Aspects of Computer Science (STACS'09)*, pages 373–384, 2009.

[36] R. Elsässer and T. Sauerwald. On the runtime and robustness of randomized broadcasting. *Theoretical Computer Science*, 410(36):3414–3427, 2009.

[37] P. Erdős and A. Rényi. On random graphs i. *Publicationes Mathematicae*, 6:290–297, 1959.

[38] K. Fall. A delay-tolerant network architecture for challenged internets. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures and Protocols for Computer Communications*, pages 27–34, 2003.

[39] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Random Structures and Algorithms*, 1(4):447–460, 1990.

[40] N. Fountoulakis, A. Huber, and K. Panagiotou. Reliable broadcasting in random networks and the effect of density. In *29th IEEE Conference on Computer Communications (INFOCOM'10)*, 2010.

[41] M. Franceschetti, O. Dousse, D. Tse, and P. Thiran. Closing the gap in the capacity of random wireless networks via percolation theory. *IEEE Transactions on Information Theory*, 53:1009–1018, 2007.

[42] A. Frieze and G. Grimmett. The shortest-path problem for graphs with random-arc-lengths. *Discrete Applied Mathematics*, 10:57–77, 1985.

[43] A. Frieze, J. Kleinberg, R. Ravi, and W. Debany. Line-of-sight networks. *Combinatorics, Probability and Computing*, 18:145–163, 2009.

[44] A. El Gamal, J. Mammen, B. Prabhakar, and D. Shah. Throughput-delay trade-off in wireless networks. In *Proceedings of the 23rd IEEE Conference on Computer Communications*, pages 464–475, 2004.

[45] E.N. Gilbert. Random graphs. *Annals of mathematical statistics*, 4:1141–1144, 1959.

[46] E.N. Gilbert. Random plane networks. *Journal of the society for industrial and applied mathematics*, 9:533–543, 1961.

[47] C. GowriShankaran. Broadcasting on recursively decomposable Cayley graphs. *Discrete Applied Mathematics*, 53(1–3):171–182, 1994.

[48] G. Grimmett. *Percolation*. Springer-Verlag, second edition, 1999.

[49] M. Grossglauser and D. Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE Transactions on Networking*, 10:477–486, 2002.

[50] P. Gupta and P.R. Kumar. Critical power for asymptotic connectivity in wireless networks. In W.M. McEneany, G. Yin, and Q. Zhang, editors, *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming*, pages 547–566. Birkhäuser, Boston, 1998.

[51] P. Gupta and P.R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46:388–404, 2000. Correction in *IEEE Transactions on Information Theory* **49** (2000), p. 3117.

[52] P. Gupta and P.R. Kumar. Internets in the sky: The capacity of three dimensional wireless networks. *Communications in Information and Systems*, 1:33–49, 2001.

[53] S.M. Hedetniemi, S.T. Hedetniemi, and A.L. Liestman. A survey of gossiping and broadcasting in communication networks. *Networks*, 18(4):319–349, 1988.

[54] J. Hromkovič, R. Klasing, A. Pelc, P. Ružička, and W. Unger. *Dissemination of Information in Communication Networks*. Springer, 2004.

[55] S. Janson, T. Łuczak, and A. Ruciński. *Random Graphs*. John Wiley & Sons, 2000.

[56] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *41st Annual IEEE Symposium on Foundations of Computer Science (FOCS'00)*, pages 565–574, 2000.

[57] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *44rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'03)*, pages 482–491, 2003.

[58] G. Kesidis, T. Konstantopoulos, and S. Phoha. Surveillance coverage of sensor networks under a random mobility strategy. In *Proceedings of the 2nd IEEE International Conference on Sensors*, 2003.

[59] H. Kesten and V. Sidoravicius. The spread of a rumor or infection in a moving population. *The annals of probability*, 33:2402–2462, 2005.

[60] T. Konstantopoulos. Response to Prof. Baccelli's lecture on modelling of wireless communication networks by stochastic geometry. *Computer Journal Advance Access*, 2009.

[61] B. Liu, P. Brass, O. Dousse, P. Nain, and D. Towsley. Mobility improves coverage of sensor networks. In *Proceedings of the 6th ACM International Conference on Mobile Computing and Networking (MobiCom)*, 2005.

[62] Z. Lotker and A. Navarra. Managing random sensor networks by means of grid emulation. In *Networking*, pages 856–867, 2006.

[63] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.

[64] P. Mattila. *Geometry of Sets and Measures in Euclidean Spaces*. Cambridge University Press, 1995.

[65] R. Meester and R. Roy. *Continuum Percolation*. Cambridge University Press, 1996.

[66] M. Mitzenmacher and E. Upfal. *Probability and Computing: randomized algorithms and probabilistic analysis*. Cambridge University Press, 2005.

[67] M. Moreau, G. Oshanin, O. Bénichou, and M. Coppey. Lattice theory of trapping reactions with mobile species. *Phys. Rev.*, E 69, 2004.

[68] P. Mörters and Y. Peres. *Brownian Motion*. Cambridge University Press, 2010.

[69] D. Mosk-Aoyama and D. Shah. Fast distributed algorithms for computing separable functions. *IEEE Transactions on Information Theory*, 54(7):2997–3007, 2008.

[70] T. Nadeem, S. Dashtinezhad, C. Liao, and L. Iftode. TrafficView: traffic and data dissemination using car-to-car communication. *Mobile Computing and Communications Review*, 8(3):6–19, 2004.

[71] M. Penrose. *Random Geometric Graphs*. Oxford University Press, 2003.

[72] M. Penrose and A. Pisztora. Large deviations for discrete and continuous percolation. *Advances in Applied Probability*, 28:29–52, 1996.

[73] Y. Peres, A. Sinclair, P. Sousi, and A. Stauffer. Mobile geometric graphs: detection, coverage and percolation. In *Proceedings of the 22st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 412–428, 2011.

[74] Y. Peres and P. Sousi. An isoperimetric inequality for the Wiener sausage, 2011. Preprint at arXiv:1103.6059v1.

[75] A. Pettarin, A. Pietracaprina, G. Pucci, and E. Upfal. Tight bounds on information dissemination in sparse mobile networks. In *30th Annual ACM Principles of Distributed Computing (PODC'11)*, 2011.

[76] B. Pittel. On spreading rumor. *SIAM Journal on Applied Mathematics*, 47(1):213–223, 1987.

[77] T. Sauerwald. On mixing and edge expansion properties in randomized broadcasting. *Algorithmica*, 56(1):51–88, 2010.

[78] T. Sauerwald and A. Stauffer. Rumor spreading and vertex expansion on regular graphs. In *Proceedings of the 22st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 462–475, 2011.

[79] J. Sheu, C. Wu, and T. Chen. An optimal broadcasting algorithm without message redundancy in star graphs. *IEEE Transactions on Parallel and Distributed Systems*, 6(6):653–658, 1996.

[80] A. Sinclair and A. Stauffer. Mobile geometric graphs, and detection and communication problems in mobile wireless networks, 2010. Preprint at arXiv:1005.1117v1.

[81] D. Spielman and N. Srivastava. Graph sparsification by effective resistances. In *40th Annual ACM Symposium on Theory of Computing (STOC'08)*, pages 563–568, 2008.

[82] F. Spitzer. Electrostatic capacity, heat flow, and Brownian motion. *Z. Wahrscheinlichkeitstheorie verw. Geb.*, 3:110–121, 1964.

[83] D. Stoyan, W.S. Kendall, and J. Mecke. *Stochastic Geometry and its Applications*. John Wiley & Sons, 2nd edition, 1995.

[84] R.M. Tanner. Explicit concentrators from generalized N-gons. *SIAM Journal Algebraic Discrete Methods*, 5(3):287–293, 1984.

Appendix A

Large Deviation Results

We use the following standard Chernoff bounds.

Lemma A.1 (Chernoff bound for binomial, see [4, Lemma A.1.4]). *Let $X_1, X_2, \ldots, X_n$ be independent Bernoulli random variable such that $\mathbf{E}[X_i] = p_i$. Let $X = \sum_{i=1}^{n} X_i$. Then, for any $\epsilon > 0$,*

$$\mathbf{P}\left(X \geq (1+\epsilon)\,\mathbf{E}[X]\right) \leq \exp\left(-\frac{2\epsilon^2(\mathbf{E}[X])^2}{n}\right),$$

and

$$\mathbf{P}\left(P \leq (1-\epsilon)\lambda\right) \leq \exp\left(-\frac{\lambda\epsilon^2}{2}\right).$$

Lemma A.2 (Another Chernoff bound for binomial [4, Corollary A.1.10]). *Let X be the sum of n i.i.d. Bernoulli random variables with mean p. Then,*

$$\mathbf{P}\left(X \geq np + a\right) \leq \exp\left(a - (pn+a)\log\left(1 + \frac{a}{pn}\right)\right).$$

Lemma A.3 (Chernoff bound for Poisson). *Let P be a Poisson random variable with mean λ. Then, for any $0 < \epsilon < 1$,*

$$\mathbf{P}\left(P \geq (1+\epsilon)\lambda\right) \leq \exp\left(-\frac{\lambda\epsilon^2}{2}(1-\epsilon/3)\right),$$

and

$$\mathbf{P}\left(P \leq (1-\epsilon)\lambda\right) \leq \exp\left(-\frac{\lambda\epsilon^2}{2}\right).$$

Lemma A.4 (Chernoff bound for geometric). *Let $X_1, \ldots, X_n$ be independent geometric random variables, each having parameter p, and let $X = \sum_{i=1}^{n} X_i$. Then, for any $\epsilon > 0$,*

$$\mathbf{P}\left(X \geq (1+\epsilon)\,\frac{n}{p}\right) \leq \exp\left(-\frac{\epsilon^2}{2(1+\epsilon)}\,n\right).$$