Connectivity for line-of-sight networks in higher dimensions

Linda Farczadi

Department of Computer Science
McGill University
Montreal, Canada

September, 2010

A thesis submitted to the Faculty of Graduate Studies and Research
in partial fulfillment of the requirements of the degree of
Master of Science

Abstract

Line-of-sight networks were introduced by Frieze et al. [10] to model wireless communications in an urban setting. Their model is based on a two-dimensional grid of points in the shape of a torus. Two grid points are said to be mutually visible if they lie in the same row or column of the torus and if the distance between them is within a certain predetermined range. Distance is measured using the ℓ_1 norm. A random graph is obtained by placing a node at each grid point independently with the same placement probability and connecting all mutually visible pairs of nodes. Among the results proven by Frieze et al. is a threshold for the connectivity of the graph.

In this thesis we extend the model of Frieze et al. to higher dimensions and the general ℓ_p norm. Specifically we consider an underlying d-dimensional grid in the shape of a torus and we define two points to be mutually visible if they differ in at most k coordinates, for some $k < d$, and the distance between them is within a certain range. We then prove corresponding asymptotically tight connectivity thresholds for this generalized model.

Abrégé

Les réseaux de lignes de visée ont été introduits par Frieze et al. [10] pour modéliser les communications sans fil en milieu urbain. Leur modèle est basé sur une grille bidimensionnelle de points en forme de tore. Deux points de la grille sont dits mutuellement visibles si ils sont situés sur la même ligne ou la même colonne du tore et si la distance entre eux est dans un certain intervalle prédéterminé. La distance est mesurée en utilisant le norme ℓ_1. Un graphe aléatoire est obtenu en plaçant un noeud à chaque point de la grille indépendamment avec la même probabilité de placement et en reliant toutes les paires de noeuds qui sont mutuellement visible. En particulier, Frieze et. al déterminent le seuil de connectivité du graphe.

Dans cette thèse, nous généralisons le modèle de Frieze et al. à dimensions supérieures. Plus précisément, nous considérons comme un réseau sous-jacent une grille à d dimensions et nous disons que deux points sont mutuellement visible si ils ont un maximum de k coordonnées différentes, pour un certain $k < d$ et si la distance entre eux est dans un certain intervalle. Ensuite, nous prouvons un seuil pour la connectivité de ce modèle généralisé.

Acknowledgments

I would like to thank my advisor, Luc Devroye, for his guidance and support during my studies. It has been a highly challenging and enjoyable experience working with him and I have learned a great amount from being his student. I would also like to acknowledge all the staff and students at McGill who have welcomed me and made my stay here a very pleasant one. In particular, I appreciate Omar Fawzi's help with the French translation of the abstract of this thesis.

I would like to express my appreciation to my previous professors at Queen's University for their help and support during my undergraduate days and for encouraging me to pursue graduate studies. I am especially grateful for Prof. Navin Kashyap's patience and guidance during my first attempt at research.

Finally I would like to dedicate this thesis to my parents without whose encouragement and support none of this would be possible.

Contents

Abstract **2**

Abrégé **3**

Acknowledgments **4**

1 Introduction **8**

1.1 Motivation 8

1.2 Organization and contributions 15

2 Background **16**

2.1 Random graphs 16

2.2 Random geometric graphs 18

2.3 Line-of-sight model of Frieze et al. 20

3 Main Result **23**

3.1 Line-of-sight networks in higher dimensions 23

3.2 Connectivity theorem for line-of-sight networks in higher dimensions 24

3.3 Extension 1: no range constraints 25

3.4 Extension 2: no line-of-sight constraints 25

3.5 Correspondence between line-of-sight networks and random geometric graphs 26

4 Proof of main result **29**

4.1 Preliminaries and notation 29

4.2 The lower bound 30

4.3 The upper bound 33

4.4 The middle case 46

5 Conclusion **49**

References

Selected Bibliography Including Cited Works **50**

List of Figures

1.1 2D network with no obstructions. 9

1.2 2D network with obstructions. 10

1.3 2D line-of-sight network. 11

1.4 Connectivity range of a node when $k = 1$. 12

1.5 Example of a 3D line-of-sight network when $k = 1$. 13

1.6 Connectivity range of a node when $k = 2$. 14

1.7 Example of a 3D line-of-sight network when $k = 2$. 14

4.1 Square with side length $2\omega/\sqrt{2}$ inside ℓ_2 ball of radius ω. 35

4.2 Fixing of coordinates in the xy-plane. 41

Chapter 1

Introduction

1.1 Motivation

In recent years the area of complex networks has received extensive research from many different scientific fields. Due to the ever-increasing computational power provided by computers, researchers have been able to store and investigate data sets that are larger than ever before. This has allowed them to study such large networks that occur in the real world. Since these networks can occur almost anywhere, the approach has been highly interdisciplinary. Some common examples of such networks include wireless communications, electrical power grids, protein interactions, the World-Wide Web and social networks.

One of the main aims of the research in this field is to better understand the underlying structural properties of these networks since these properties govern their behaviour. For example the topology of social networks is the main governing force behind the spread of information, and the topology of forests can determine the spread of fires and disease [16]. A main requirement is to find models that accurately characterize these networks and capture their main properties. Since many of these networks can be described by their local behaviours, and since these behaviours are in fact probabilistic in nature [16], a natural choice for their study are random network models.

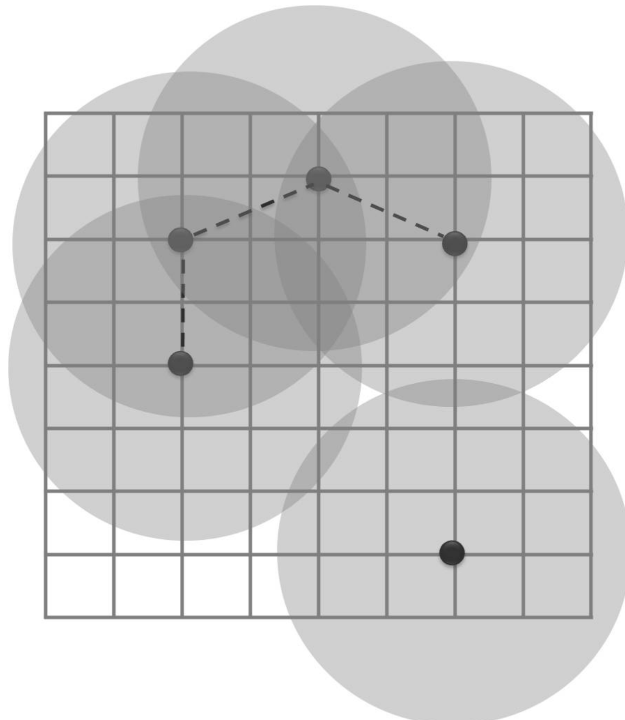

Figure 1.1: 2D network with no obstructions.

The subject of this thesis is line-of-sight networks, a random network model recently introduced by Frieze et al. [10]. Frieze et al. proposed this model in order to accurately capture the properties of wireless communications in a complex urban setting. Unlike in large open environments where communication between two nodes is only subject to range limitations, in an urban setting we have the added constraint of line-of-sight restrictions.

To better understand the motivation behind this model consider the example given in [10] where we have a two-dimensional grid representing the downtown of a major city. The rows of the grid represent the streets of the city while the columns represent the avenues. Figure 1.1 shows an example of such a grid. The black circles represent the nodes of the network. The shaded circles around each node denote the range of visibility of that node. For two nodes to communicate in this setting they have to be mutually visible. If there are no obstructions then two nodes can communicate as long as they lie inside each other's visibility range.

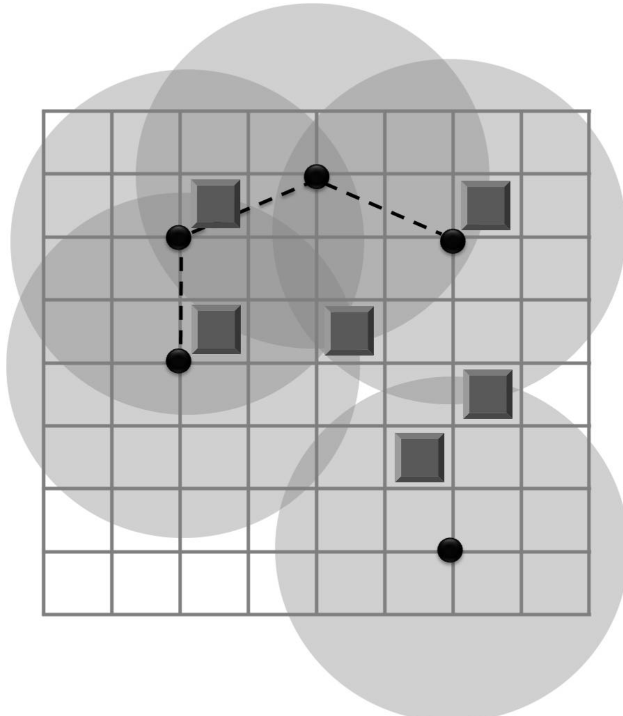

Figure 1.2: 2D network with obstructions.

However a city landscape is almost always filled with tall buildings and complex obstacles that act as obstructions to wireless communication. A more realistic depiction of this urban case is given in figure 1.2 where the dark squares denote possible obstructions. In this case visibility is no longer ensured by the nodes being in each other's visibility range. There is the additional requirement of the two nodes being on the same street or avenue.

Previous models such as random graphs [5] and random geometric graphs [15] are no longer well suited for this scenario since they are not designed to capture the line-of-sight constraints. Line-of-sight networks, on the other hand, provide a much better approximation to the true constraints of this model. Figure 1.3 shows the line-of-sight model where the shaded arms extending out of each node denote the locations on the torus that satisfy both line-of-sight and range constraints for the given nodes. Therefore a node can communicate with any other node that lies on one of its four arms denoted by the shaded regions.

In their paper [10] Frieze et al. prove several important results for two-dimensional line-of-

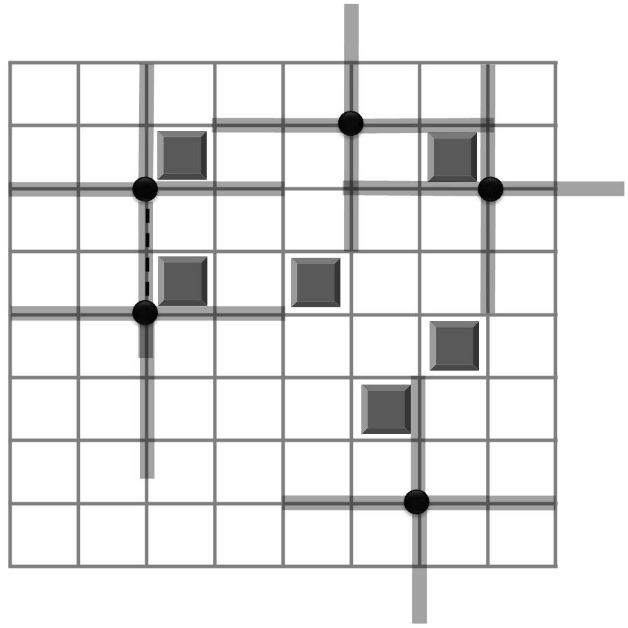

Figure 1.3: 2D line-of-sight network.

sight networks. In particular they give asymptotically tight thresholds for the k-connectivity of the graph. A graph is said to be k-connected if there does not exist a set of $k-1$ vertices whose removal, together with all incident edges, disconnects the graph. In other words, a graph is k-connected if there are at least k independent paths between any two vertices. For $k > 1$, this property becomes important for wireless networks since it makes them more reliable. If a certain link fails, there are still $k-1$ possible paths that can be used instead. In addition to k-connectivity Frieze et al. also study the emergence of a giant component. Furthermore they present an efficient algorithm for finding paths between nodes as well as an approximation algorithm for the relay placement problem.

In this thesis we chose to focus on the connectivity of these graphs. A graph is connected if there is a path between any two vertices of the graph. We chose to focus on connectivity since we believe that this is the most crucial property in the case of communication networks, especially mobile ad-hoc networks. These networks consist of a set of wireless devices that communicate with each other without any centralized control. These devices agree to route each other's data packets which means that they must forward all traffic that goes through

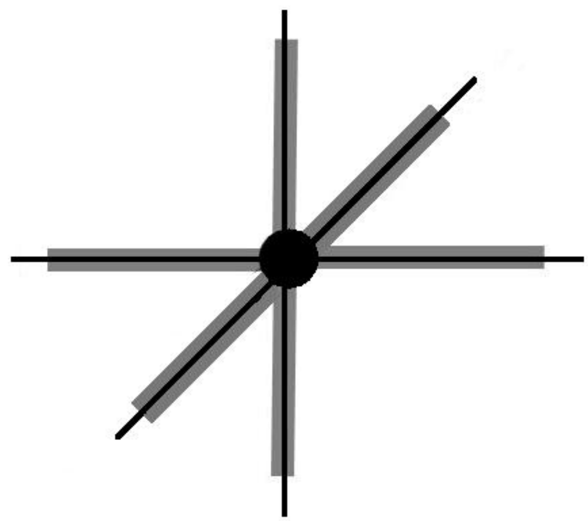

Figure 1.4: Connectivity range of a node when $k = 1$.

them even though it might be unrelated to their own use. Hence in order for this type of network to be functional it is absolutely critical that there exists a path between any pair of nodes in the network [12]. That is, the network must be connected.

Given a mathematical model for the network, it is of interest to determine the critical value of a certain parameter at which the graph becomes connected. In the case of line-of-sight networks the connectivity of the graph depends on a parameter called the placement probability. The higher the value of this parameter the denser the graph becomes. Therefore we want to determine the smallest value of this placement probability that will ensure that the graph is connected with high probability. We say that a graph is connected with high probability if the graph is asymptotically connected with probability one.

We are concerned with the connectivity of line-of-sight networks in higher dimensions. The initial motivation for higher dimensions comes from the three-dimensional case that occurs in scenarios where nodes can be placed both on the ground and also in space, for example on different floors of a building, in airplanes or in satellites [13]. The four-dimensional case can also be of interest in a situation where communication between nodes is dependent on time.

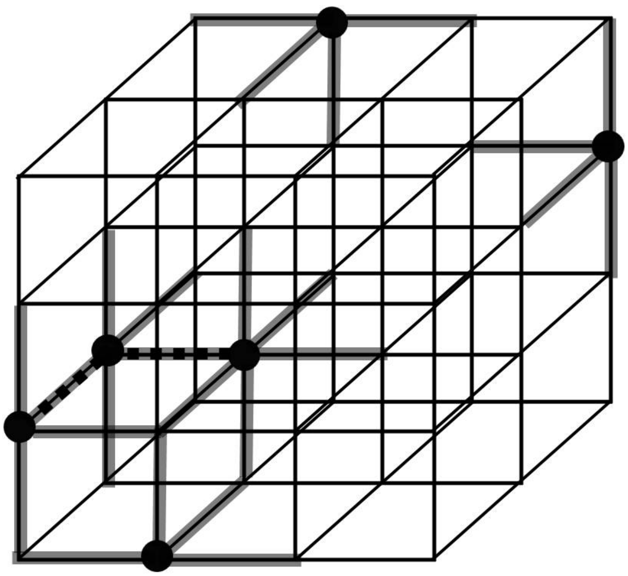

Figure 1.5: Example of a 3D line-of-sight network when $k = 1$.

We remark here that for dimensions greater than two there are multiple possibilities for defining the line-of-sight constraints. Take for example the three-dimensional case where we have an underlying grid of n^3 points forming a torus. We have two options here. First we can define two points to be mutually visible if they lie on the same one-dimensional line of the torus and are within a certain range. In this case each node is connected to any other nodes that lies on one of its one-dimensional arms denoted by the shaded regions in figure 1.4. Figure 1.5 shows an example of a line-of-sight network for this case.

Our second option is to define two points to be mutually visible if they lie on the same two-dimensional plane of the torus and are within a certain range. Figure 1.6 shows a node with its connectivity range denoted by the shaded regions. Figure 1.7 shows an example of a line-of-sight network for this case.

As the dimension increases the number of ways in which we can define the line-of-sight constraints also increases. Specifically for each integer k satisfying $1 \leq k < d$ we can say that two nodes satisfy the line-of-sight constraints if they agree in at least k coordinates.

Since the methods used for studying the connectivity of the three-dimensional model can be easily extended to any higher dimension we chose to present our results in terms of the

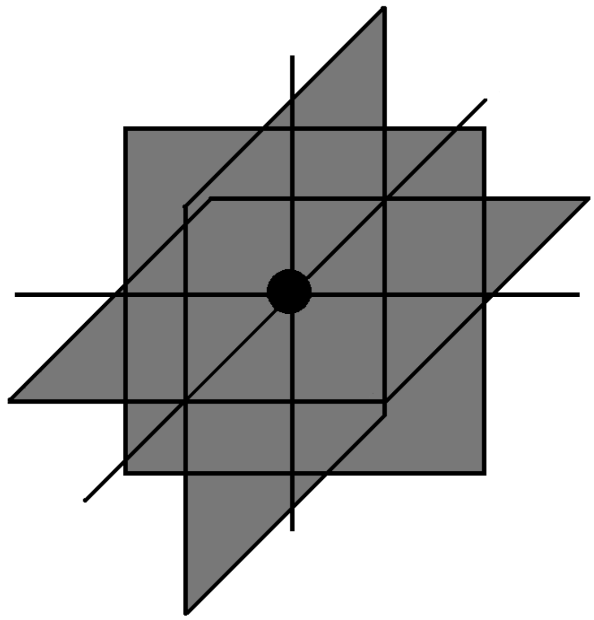

Figure 1.6: Connectivity range of a node when $k = 2$.

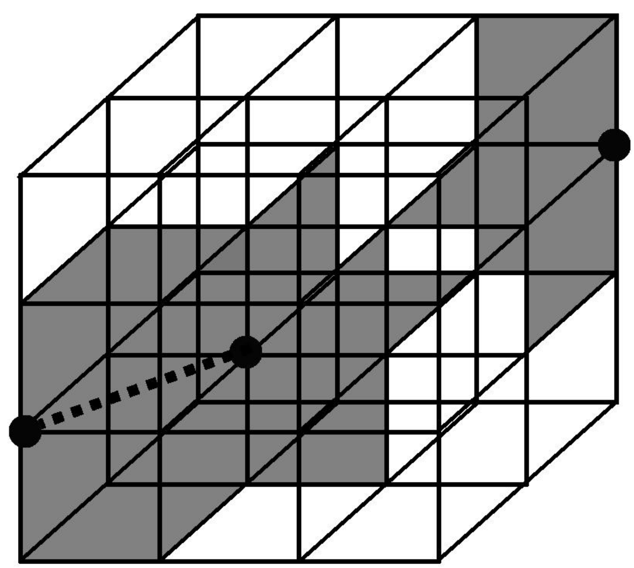

Figure 1.7: Example of a 3D line-of-sight network when $k = 2$.

general d-dimensional case.

1.2 Organization and contributions

This thesis is organized as follows. In chapter 2 we introduce two widely studied random graph models, specifically Erdős-Rényi random graphs and random geometric graphs. We also present well known connectivity results for these two models. We then formally introduce the line-of-sight model of Frieze et al. [10] and state the main connectivity result for this model. In chapter 3 we describe our generalized model of line-of-sight networks and we present our main connectivity result for this model. We also present connectivity thresholds for two extensions resulting from removing the line-of-sight and range constraints one at a time. In addition we demonstrate a correspondance between the model obtained by removing the line-of-sight constraints and random geometric graphs. In chapter 4 we prove our main theorem by considering each of the corresponding three cases: the lower bound, the upper bound and finally the middle case. In chapter 5 we conclude the thesis with a summary and suggestions for future work.

Chapter 2

Background

2.1 Random graphs

The study of random graphs was initiated by P. Erdős and A. Rényi in the late 1950's, see [8] and [9]. The initial motivation for the study of random graphs was to use them as a tool for proving the existence of graphs with certain properties. This method of proof is known as the probabilistic method and it is the subject of a book by Alon and Spencer [1]. The method can be easily described as follows: suppose that we want to prove the existence of a combinatorial structure that has a certain property; to do this we construct a probability space and prove that a randomly chosen object from this space has the desired property with a positive probability. Since its introduction by Erdős, the probabilistic method has found multiple uses in many branches of mathematics and theoretical computer science.

The Erdős-Rényi random graph model consists of a graph on n nodes, denoted by $G(n, p)$, where each potential edge is added independently with some probability p. In a famous paper [9], Erdős and Rényi looked at random graphs as living organisms and studied their evolution over time. This evolution can be described as follows. At the very beginning $p = 0$ and the graph consists of a set of n isolated nodes with no edges between them. As time evolves, p is increased continuously with time, and edges of the graph are born one by one. Finally at

$p = 1$ all edges are present resulting in the complete graph K_n.

An essential question regarding the evolution of random graphs is at what point does a specific property of the graph become likely to occur. In particular we are interested in monotone properties, that is properties whose likelihood of occurrence increases with p. Erdős and Rényi found that many important monotone properties of random graphs exhibit a threshold phenomena. Specifically there is a certain critical probability at which the desired property switches from being very unlikely to very likely.

For the case where the desired property of the graph is connectivity, a famous result of Erdős-Rényi [9] states that if $p = (\ln n + c_n)/n$ then

$$\lim_{n\to\infty} \Pr\{\text{graph is connected}\} \begin{cases} 0 & c_n \to -\infty \\ e^{-c} & c_n \to c \\ 1 & c_n \to \infty. \end{cases}$$

This means that if $p = c \ln n/n$ then for $c < 1$ the graph is non-connected with high probability while for $c > 1$ the graph is connected with high probability. We note here that this is also the threshold value for the existence of isolated nodes. To see this, let N be the number of isolated nodes and X_i the indicator variable that node i is isolated. Node i is isolated if there are no edges connecting it to any of the other $n - 1$ nodes in the graph. Therefore

$$\begin{aligned} \mathrm{E}\{N\} &= \sum_i \mathrm{E}\{X_i\} \\ &= \sum_i \Pr\{\text{node } i \text{ is isolated}\} \\ &= n(1-p)^{n-1} . \end{aligned}$$

If we let $p = (\ln n + c_n)/n$ then as $n \to \infty$

$$
\begin{aligned}
\mathrm{E}\{N\} &= n\left(1 - (\ln n + c_n)/n\right)^{n-1} \\
&\sim n e^{-\ln n - c_n} \\
&= e^{-c_n} \; .
\end{aligned}
$$

Therefore

$$
\lim_{n\to\infty} \mathrm{E}\{N\} = \begin{cases} \infty & c_n \to -\infty \\ e^{-c} & c_n \to c \\ 0 & c_n \to \infty. \end{cases}
$$

Hence the smallest p for which there are no longer any isolated nodes is also the smallest p at which the graph becomes connected with high probability. In the following section we show that this behaviour is also encountered in random geometric graphs.

2.2 Random geometric graphs

A more suitable model for the study of wireless ad-hoc networks are random geometric graphs, since they account for certain range limitations that might exist in wireless communications. A random geometric graph is formed by placing n nodes uniformly at random in a d-dimensional space and connecting all pairs of nodes that are within a certain range r. A detailed study of random geometric graphs can be found in a book by Penrose [15]. A major difference between random geometric graphs and the Erdős - Rényi model is that for random geometric graphs the events corresponding to the existence of different edges are no longer independent. This is because if a node x is close to a node y, and y is close to a node z then x will have to be fairly close to z as well. It is exactly this triangle property that makes random geometric graphs more suitable for modelling realistic scenarios such as wireless networks.

Similarly to the case of the Erdős-Rényi model, random geometric graphs also exhibit a threshold phenomena in the occurrence of certain monotone properties, see [11]. For the particular case of connectivity, Gupta and Kumar [12] showed that if the underlying space is the two-dimensional unit disk in $\mathbb{R}^2$ and distance is measured using the Euclidean norm then for $\pi r^2 = (\ln n + c_n)/n$ the following holds

$$\lim_{n\to\infty} \Pr\{\text{graph is connected}\} \begin{cases} 0 & c_n \to -\infty \\ e^{-c} & c_n \to c \\ 1 & c_n \to \infty. \end{cases}$$

The connectivity of random geometric graphs in higher dimensions using the ℓ_∞ norm has also been extensively studied by Appel and Russo, see [2], [3], [4]. Penrose [15] generalized connectivity results to all dimensions and any ℓ_p norm. For the case where the underlying space is the d-dimensional torus $[0,1]^d$ and θ is the volume of the unit ball in the norm of choice, a result of Penrose [14] states that if $\theta r^d = (\ln n + c_n)/n$ then

$$\lim_{n\to\infty} \Pr\{\text{graph is connected}\} \begin{cases} 0 & c_n \to -\infty \\ e^{-c} & c_n \to c \\ 1 & c_n \to \infty. \end{cases}$$

We note that similarly to the case of the Erdős-Rényi model, the threshold for connectivity is the same as the threshold for the existence of isolated nodes. Again let N be the number of isolated nodes, and let X_i be the indicator variable that node i is isolated. Node i is isolated if there are no other nodes in the ℓ_p ball of radius r centered at i. Since each node has a probability of θr^d of falling into a specific ℓ_p ball of radius r, we obtain

$$
\begin{aligned}
\mathrm{E}\{N\} &= \sum_i \mathrm{E}\{X_i\} \\
&= \sum_i \Pr\{\text{node } i \text{ is isolated}\} \\
&= n\left(1-\theta r^d\right)^{n-1} .
\end{aligned}
$$

If we let $\theta r^d = (\ln n + c_n)/n$ then as $n \to \infty$

$$
\begin{aligned}
\mathrm{E}\{N\} &= n\left(1-\frac{\ln n + c_n}{n}\right)^{n-1} \\
&\sim n e^{-\ln n - c_n} \\
&= e^{-c_n} .
\end{aligned}
$$

Therefore

$$
\lim_{n\to\infty} \mathrm{E}\{N\} = \begin{cases} \infty & c_n \to -\infty \\ e^{-c} & c_n \to c \\ 0 & c_n \to \infty . \end{cases}
$$

So the smallest value of the range parameter r for which there are no longer any isolated nodes is also the smallest value at which the graph becomes connected with high probability.

2.3 Line-of-sight model of Frieze et al.

Line-of-Sight networks were introduced by Frieze et al. [10] to model wireless communications in an urban setting. They have the added advantage over random graphs and random geometric graphs of incorporating both range and line-of-sight constraints. The model is based on an $n \times n$ two dimensional grid of points in the shape of a torus. The distance between two points is measured using the ℓ_1 norm. For a chosen range parameter ω, two points on the grid are said to be mutually visible if they agree in one coordinate and are a distance at most

ω apart. A random graph is obtained by placing a node at each grid point independently with some probability $p^* > 0$ and connecting all mutually visible pairs of nodes. Among the results proven by Frieze et al. [10] is a threshold for connectivity of the graph. Assuming that $\omega = n^\delta$, for some $0 < \delta < 1$ they show that if $p^* = \left(\left(1 - \frac{1}{2}\delta\right)\ln n + \frac{1}{2}\ln\ln n + c_n\right)/2\omega$. Then

$$\lim_{n\to\infty} \Pr\{\text{G is connected}\} = \begin{cases} 0 & c_n \to -\infty \\ e^{-\lambda} & c_n \to c \\ 1 & c_n \to \infty, \end{cases}$$

where $\lambda = \frac{1}{2}\left(1 - \frac{1}{2}\delta\right)e^{-2c}$ and $c_n = o(\ln\ln n)$.

We now show that similarly to the case of random graphs and random geometric graphs the threshold for connectivity is the same as the threshold for the existence of isolated nodes. A node in this two-dimensional line-of-sight network is isolated if there are no other nodes in its visibility range. For a given point i in the torus there are exactly 4ω points that are visible from i. Therefore we obtain

$$\begin{aligned} \mathrm{E}\{N\} &= \sum_i \mathrm{E}\{X_i\} \\ &= \sum_i \Pr\{\text{node } i \text{ is isolated}\} \\ &= n^2 p^* (1 - p^*)^{4\omega} \ . \end{aligned}$$

If we let $p^* = \left(\left(1 - \frac{1}{2}\delta\right)\ln n + \frac{1}{2}\ln\ln n + c_n\right)/2\omega$ then as $n \to \infty$

$$\begin{aligned} \mathrm{E}\{N\} &\sim n^2 \frac{\left(1 - \frac{1}{2}\delta\right)\ln n}{2\omega} e^{-\left(1-\frac{1}{2}\delta\right)\ln n - \frac{1}{2}\ln\ln n - c_n} \\ &= \frac{\left(1 - \frac{1}{2}\delta\right)e^{-c_n}}{2}. \end{aligned}$$

Therefore

$$\lim_{n\to\infty} \mathrm{E}\{N\} = \begin{cases} \infty & c_n \to -\infty \\ \frac{\left(1-\frac{1}{2}\delta\right)e^{-c}}{2} & c_n \to c \\ 0 & c_n \to \infty. \end{cases}$$

Hence the smallest value of the placement probability p^* for which there are no longer any isolated nodes is also the smallest value at which the graph becomes connected.

Chapter 3

Main Result

3.1 Line-of-sight networks in higher dimensions

Given positive integers d and k such that $d \geq 2$ and $1 \leq k < d$, we define an underlying $d-$dimensional grid in the shape of a torus

$$T = \{(x_1, \cdots, x_d) : x_i \in \{1, 2, ..., n\}, 1 \leq i \leq d\}.$$

We say that two points are mutually visible if they differ in at most k coordinates and are a distance at most ω apart. Distance is measured assuming that the points lie on a torus. That is given two points $x = (x_1, \cdots, x_d)$ and $y = (y_1, \cdots, y_d)$ on the torus we measure the distance between them as

$$d(x, y) = \| \min(|x_1 - y_1|, n - |x_1 - y_1|), \cdots, \min(|x_d - y_d|, n - |x_d - y_d|) \|_p .$$

The norm is the standard ℓ_p norm in $\mathbb{R}^d$ for $1 \leq p \leq \infty$.

We obtain a random graph G by placing a node at each grid point independently with some probability $p^* > 0$ and connecting all mutually visible pairs of nodes. We use the same assumption as in [10] that $\omega = n^{\delta}$, where $0 < \delta < \delta_0 < 1$ with

$$\delta_0 = \begin{cases} \frac{d}{d+k} & \text{if } d/2 < k < d \\ \frac{d}{d+k\left(\left\lceil \frac{d}{k} \right\rceil - 1\right)} & \text{if } k \leq d/2. \end{cases}$$

We define the constant a_p as $a_p = \binom{d}{k} \int_{B_p^k(0,1)} dx$ where $B_p^k(0,1) := \left\{x \in \mathbb{R}^k : \| x \|_p \leq 1\right\}$ and $\| x \|_p$ is the standard ℓ_p norm in $\mathbb{R}^d$ for $1 \leq p \leq \infty$. The exact expression for the integral can be found in [17] and is given by

$$\int_{B_p^k(0,1)} dx = \frac{2^k \Gamma\left(1 + \frac{1}{p}\right)^k}{\Gamma\left(1 + \frac{k}{p}\right)},$$

where Γ denotes the gamma function.

3.2 Connectivity theorem for line-of-sight networks in higher dimensions

We state here the main result of this thesis:

Theorem 1. *Let G be a line-of-sight network defined as in section 3.1 and let $p^* = \frac{(d-k\delta)\ln n + \ln\ln n + c_n}{a_p \omega^k}$, with $c_n = o(\ln\ln n)$. Then*

$$\lim_{n\to\infty} Pr\{G \text{ is connected}\} = \begin{cases} 0 & c_n \to -\infty \\ e^{-\lambda} & c_n \to c \\ 1 & c_n \to \infty, \end{cases}$$

where $\lambda = \frac{(d-k\delta)e^{-c}}{a_p}$.

We observe that p^* is the threshold value for the existence of isolated nodes. This is shown explicitly in section 4.1. We also note that the choice of $\omega = n^{\delta}$ is motivated by an observation of Frieze et al. [10] that if $\omega = o(\ln n)$ then the threshold value for connectivity is very close to one. Thus the requirement that $w \gg \ln n$ is needed for non-trivial results.

3.3 Extension 1: no range constraints

In these next two sections we look at what happens when either the range or line-of-sight constraints are removed one at a time. We first consider the case where there are no range constraints and connectivity between two points is subject only to line-of-sight constraints. This happens when $\omega = \infty$. To maintain the line-of-sight constraints we still require that $k < d$. Thus two points in the torus are mutually visible as long as they differ in at most k coordinates. Even though our bounds on δ exclude the case $\omega = \infty$ we state here the corresponding threshold for this case.

Theorem 2. *Assume $k < d$ and $\omega = \infty$. Let G be a line-of-sight network defined as in section 3.1. If $p^* = \frac{(d-k)\ln n + \ln\ln n + c_n}{\binom{d}{k} n^k}$, then*

$$\lim_{n\to\infty} Pr\{G \text{ is connected}\} = \begin{cases} 0 & c_n \to -\infty \\ \frac{(d-k)}{\binom{d}{k}} e^{-c} & c_n \to c \\ 1 & c_n \to \infty. \end{cases}$$

3.4 Extension 2: no line-of-sight constraints

We now consider the case where there are no line-of-sight constraints and connectivity between two points is subject only to range limitations. This happens when $k = d$. To maintain the range limitations we assume as in the original definition of line-of-sight networks that $\omega = n^\delta$ where δ satisfies the upper bound given in section 3.1. Then two nodes are mutually visible as long as they are within a distance of ω apart. Even though our bounds on k exclude the case $k = d$ we state here the corresponding threshold for this case. We note that the proof of the lower bound given in section 4.2 applies to this special case as well.

Theorem 3. *Assume $k = d$ and $\omega = n^\delta$ with δ satisfiying the upper bound in section* 3.1. *Let G be a line-of-sight network defined as in section* 3.1. *If $p^* = \frac{(d-d\delta)\ln n + \ln\ln n + c_n}{a_p \omega^d}$, then*

$$\lim_{n\to\infty} Pr\{G \text{ is connected}\} = \begin{cases} 0 & c_n \to -\infty \\ \frac{(d-d\delta)}{a_p} e^{-c} & c_n \to c \\ 1 & c_n \to \infty. \end{cases}$$

In the following section we describe how this case corresponds to that of random geometric graphs and we show that this correspondence also holds between the respective connectivity thresholds.

3.5 Correspondence between line-of-sight networks and random geometric graphs

In this section we describe a correspondence between line-of-sight networks where $k = d$ and a random geometric graph. For ease of notation we consider the two-dimensional case where distance is measured using the Euclidean norm.

Suppose we have an underlying two-dimensional torus $[0, 1]^2$. Consider a random geometric graph obtained by placing N points uniformly at random in $[0, 1]^2$ and connecting all pairs of nodes that are a distance at most r apart. Let $r = \sqrt{\frac{c \log N}{\pi N}}$ for some constant c. As described in section 2.2 the random geometric graph is connected with high probability if $c > 1$ and not connected with high probability if $c < 1$.

Now suppose that on the $[0, 1]^2$ torus we also have an n by n equally spaced grid of n^2 points. Let T denote the set of all grid points. We create a discretized version of the random geometric graph by mapping all its nodes to grid points. This mapping is done as follows: for each $1/n \times 1/n$ grid square we map all the nodes that fall into this square to the grid point given by the upper right corner of this square. That is a node (v_1, v_2) is mapped to the grid point $f((v_1, v_2))$ computed as follows:

$$f((v_1, v_2)) = \underset{(w_1,w_2)\in T: w_1 \geq v_1, w_2 \geq v_2}{\operatorname{argmin}} \sqrt{(w_1 - v_1)^2 + (w_2 - v_2)^2} \ .$$

Note that more than one node can be mapped to the same grid point. However if the grid size is small enough then the number of grid points that have more than one node mapped to them is negligible. The expected number of nodes that fall in a given square of area $1/n^2$ is N/n^2. Therefore if we require that $n >> \sqrt{N}$ then $N/n^2 << 1$ and the discretization of the random geometric graph described above is fine enough to ensure that the probability of two nodes being mapped to the same grid point is negligible.

After performing this mapping we connect a pair of mapped nodes if and only if they were connected in the non-discretized random geometric graph. Note that by requiring that $1/N << r$ this is roughly equivalent to connecting two nodes in the discretized version if and only if they are within a distance r of each other.

We now set $r = \omega/n$ and $p = N/n^2$. Consider the line-of-sight graph obtained by placing a node at each grid point independently with probability p and connecting all pairs of nodes that are a distance at most ω/n apart. We assume as before that the range parameter ω for the line-of-sight graph is of the form $\omega = n^\delta$ for some $0 < \delta < 1$. Then this line-of-sight graph and the discretized random geometric graph are equivalent models. We note that the number of nodes in the two graphs is not necessarily the same. The line-of-sight graph has an expected number of $n^2 p$ nodes while the discretized random geometric graph has a total of N nodes. However since we fixed the placement probability of the line-of-sight graph as $p = N/n^2$, the total number of nodes is roughly the same. Also since we fixed $r = \omega/n$, in both graphs two nodes are connected if and only if they are within a distance r of each other.

From $r^2 = \frac{c \log N}{\pi N}$ and $r = \omega/n$ we obtain

$$\frac{c \log N}{\pi N} = \frac{\omega^2}{n^2} = n^{2\delta-2}.$$

Thus using the equation $p = N/n^2$ and solving for p we obtain the following connectivity threshold for the line-of-sight graph

$$p = \frac{c\left((2-2\delta)\ln n + \ln\ln n\right)}{\pi n^{2\delta}}.$$

Therefore the line-of-sight graph is connected with high probability if $c > 1$ and not connected with high probability if $c < 1$. We note that this corresponds to the connectivity threshold for line-of-sight networks with $k = d$ given in section 3.4.

Chapter 4

Proof of main result

4.1 Preliminaries and notation

We say that a certain event or property holds with high probability if the probability that the event, respectively property, holds converges to one as $n \to \infty$. For an event E we let $\overline{E}$ denote its complement. We write $Bin(n,p)$ to denote a random variable with a binomial distribution with parameters n and p, and $Pois(\lambda)$ for a random variable with a Poisson distribution with parameter λ. For a subset X of $[0,n]^d$ we write $vol(X)$ to denote the volume of X. Given two distinct points $x = (x_1, \cdots, x_d)$ and $y = (y_1, \cdots, y_d)$ we say that x is lexicographically smaller than y if there exists $j \in 1, \cdots, d$ such that $x_j < y_j$ and $x_i = y_i$ for all $i < j$.

For a fixed point $i \in T$ let $V(i)$ denote the set of all points in T that are visible from i. Recall that these are all the points that differ in at most k coordinates from i and are within a distance ω from i. Let $S_1, \cdots, S_{\binom{d}{k}}$ denote all the distinct subsets formed by choosing k out of d coordinates. For each point $i \in T$ and each S_j, define the section $V_{S_j}(i)$ to be the set of all points that are within a distance ω from i and that can differ from i only in the k coordinates that are in S_j. Then $\cup_j V_{S_j}(i) = V(i)$. Recall that we defined the constant a_p as $a_p = \binom{d}{k} \int_{B_p^k(0,1)} dx$ where $B_p^k(0,1) := \left\{ x \in \mathbb{R}^k : \| x \|_p \leq 1 \right\}$ and $\| x \|_p$ is the standard ℓ_p

norm in $\mathbb{R}^d$ for $1 \leq p \leq \infty$. Using this definition we have $|V(i)| \sim a_p \omega^k$ as $\omega \to \infty$, $\omega < n/2$.

We let N denote the number of isolated nodes in G. With each point $i \in T$ we associate an indicator random variable X_i which is equal to 1 if point i contains an isolated node and is equal to 0 otherwise. Thus $N = \sum_{i \in T} X_i$. Letting $n \to \infty$ we obtain

$$\begin{aligned}
\mathrm{E}\{N\} &= \sum_i \mathrm{E}\{X_i\} \\
&= n^d p^* (1-p^*)^{|V(i)|} \\
&\sim \frac{n^d (d-k\delta) \ln n}{a_p \omega^k} e^{-p^* a_p \omega^k} \quad \text{as } n \to \infty \\
&= \frac{(d-k\delta) e^{-c_n}}{a_p}.
\end{aligned}$$

Therefore

$$\lim_{n\to\infty} \mathrm{E}\{N\} = \begin{cases} \infty & c_n \to -\infty, \\ \frac{(d-k\delta)e^{-c}}{a_p} = \lambda & c_n \to c \in \mathbb{R}, \\ 0 & c_n \to \infty. \end{cases}$$

4.2 The lower bound

Suppose $c_n \to -\infty$. We use the second moment method, [1], to show that with high probability there are isolated nodes, which implies that the random graph G is not connected with high probability. Using Chebyshev's inequality, we have

$$\Pr\{N=0\} \leq \frac{\mathrm{V}\{N\}}{\mathrm{E}^2\{N\}}.$$

Recall that $N = \sum_{i \in T} X_i$. For two points $i, j \in T$ we write $i \sim j$ when X_i and X_j are not independent. Then

$$\begin{aligned} \mathrm{V}\{N\} &= \sum_i \mathrm{E}\left\{(X_i - \mathrm{E}X_i)^2\right\} + \sum_{i \neq j} \mathrm{E}\{(X_i - \mathrm{E}X_i)\}\,\mathrm{E}\{(X_j - \mathrm{E}X_j)\} \\ &= \sum_i \mathrm{V}\{X_i\} + \sum_{i \sim j} (\mathrm{E}\{X_i X_j\} - \mathrm{E}\{X_i\}\,\mathrm{E}\{X_j\}) \\ &\leq \sum_i \mathrm{E}\{X_i\} + \sum_{i \sim j} \mathrm{E}\{X_i X_j\} \\ &= \mathrm{E}\{N\} + \sum_{i \sim j} \mathrm{E}\{X_i X_j\}. \end{aligned}$$

Thus

$$\Pr\{N = 0\} \leq \frac{1}{\mathrm{E}\{N\}} + \frac{\sum_{i \sim j} \mathrm{E}\{X_i X_j\}}{\mathrm{E}^2\{N\}} \to 0$$

if

$$\lim_{n \to \infty} \mathrm{E}\{N\} = \infty, \qquad \text{and} \qquad \lim_{n \to \infty} \frac{\sum_{i \sim j} \mathrm{E}\{X_i X_j\}}{\mathrm{E}^2\{N\}} = 0.$$

The first condition holds since $c_n \to -\infty$. To verify the second condition we first split the sum $\sum_{i \sim j} \mathrm{E}\{X_i X_j\}$ according to the number of coordinates that the points $i, j \in T$ differ in. For two distinct points $i, j \in T$ the number of coordinates that they differ in is at least one and can be at most d. Thus

$$\sum_{i \sim j} \mathrm{E}\{X_i X_j\} = \sum_{t=1}^{d} \sum_{i \sim j : t} \mathrm{E}\{X_i X_j\},$$

where the t under the sum denotes that i and j differ in exactly t coordinates.

Now for two point $i, j \in T$ we have $i \sim j$ if and only if $V(i) \cap V(j) \neq \emptyset$, and this happens if and only if i and j are a distance at most 2ω apart and they differ in at most $2k$ coordinates. Since $\mathrm{E}\{X_i X_j\} = 0$ if i and j are neighbours, we are only interested in summing over pairs

i and j that are not neighbours. For a fixed $i \in T$, as $n \to \infty$ there are $\sim \omega^t$ points $j \in T$ such that $j \sim i$ and j differs from i in exactly t coordinates. Thus

$$\begin{aligned}
\sum_{i \sim j : t} \mathrm{E}\{X_i X_j\} &\leq \text{const.} \cdot n^d \omega^t \max_{i,j \in T} \Pr\{ \ i, j \text{ isolated } \} \\
&= \text{const.} \cdot n^{d+t\delta} p^{*2} (1-p^*)^{2|V(i)| - \max_{i,j\in T}|V(i) \cap V(j)|}.
\end{aligned}$$

If $t \leq k$, then there must be a coordinate in which j differs by more than ω from i, since otherwise i and j would be neighbours. Thus $|V(i) \cap V(j)| \leq \frac{1}{2}|V(i)|$ for all such pairs $i, j \in T$. Hence $\max_{i,j\in T} |V(i) \cap V(j)| \leq \frac{1}{2}|V(i)|$ and:

$$\begin{aligned}
\sum_{t=1}^{k} \sum_{i \sim j : t} \mathrm{E}\{X_i X_j\} &\leq \sum_{t=1}^{k} \text{const.} \cdot n^{d+t\delta} p^{*2} (1-p^*)^{\frac{3}{2}|V(i)|} \\
&\sim \sum_{t=1}^{k} n^{d+(t-2k)\delta} e^{-\frac{3}{2}(d-k\delta)\ln n} \text{ as } n \to \infty \\
&= \sum_{t=1}^{k} n^{-\frac{d}{2} + \left(t - \frac{k}{2}\right)\delta} \to 0 \qquad \text{for } \delta < 1.
\end{aligned}$$

Next if $k < t \leq \min\{2k, d\}$ then $\max_{i,j\in T} |V(i) \cap V(j)| \leq \text{const.} \cdot \omega^{2k-t}$. Hence we obtain

$$\begin{aligned}
\sum_{t=k+1}^{\min\{2k,d\}} \sum_{i \sim j : t} \mathrm{E}\{X_i X_j\} &\leq \sum_{t=1}^{\min\{2k,d\}} n^{d+t\delta} p^{*2} (1-p^*)^{2|V(i)| - \text{const.} \cdot \omega^{2k-t}} \\
&\leq \sum_{t=1}^{\min\{2k,d\}} n^{d+t\delta} p^{*2} e^{p^*\left(2|V(i)| - \text{const.} \cdot \omega^{2k-t}\right)}
\end{aligned}$$

Using the fact that $2k - t < k$ we note that

$$e^{p^*\left(\text{const.} \cdot \omega^{2k-t}\right)} \leq e^{p^*\left(\text{const.} \cdot \omega^{k-1}\right)} \sim e^{\frac{\ln n}{n^\delta}} \to 1$$

Therefore

$$\begin{aligned}
\sum_{t=k+1}^{\min\{2k,d\}} \sum_{i\sim j:t} \mathrm{E}\{X_i X_j\} &\sim \sum_{t=1}^{\min\{2k,d\}} n^{d+t\delta} {p^*}^2 e^{p^* 2|V(i)|} \\
&\sim \sum_{t=1}^{\min\{2k,d\}} n^{d+(t-2k)\delta} e^{-2(d-k\delta)\ln n} \text{ as } n \to \infty \\
&= \sum_{t=1}^{\min\{2k,d\}} n^{-d+t\delta} \to 0 \qquad \text{for } \delta < 1.
\end{aligned}$$

Finally if $2k < d$ and $t > 2k$ then $|V(i) \cap V(j)| = \emptyset$. Therefore we have shown that $\sum_{i\sim j} \mathrm{E}\{X_i X_j\} \to 0$ and thus

$$\lim_{n\to\infty} \frac{\sum_{i\sim j} \mathrm{E}\{X_i X_j\}}{\mathrm{E}^2\{N\}} = 0.$$

4.3 The upper bound

Consider now the case where $c_n \to \infty$. We note that our proof for the upper bound follows the general structure of the one in [10].

In section 4.1 we have shown that if $c_n \to \infty$ then $\lim_{n\to\infty} \mathrm{E}\{N\} = 0$. This implies that $\Pr\{N > 0\} \to 0$ so with high probability we do not have any isolated nodes. Using the same construction as in Frieze et al. [10], we add nodes according to a two stage process. In the first stage we place a node at each point in T independently with probability

$$p_1 := p^* - \frac{1}{a_p \omega^k \ln n}.$$

In the second stage we place a node at each point in T independently with probability p_2. We pick p_2 so that the two stage process is equivalent to the original process with probability p^*. That is

$$(1 - p_1)(1 - p_2) = 1 - p^*,$$

thus

$$p_2 \sim \frac{1}{a_p \omega^k \ln n} \text{as } n \to \infty.$$

We refer to the nodes placed in the first stage as *red* nodes, and nodes placed in the second stage as *blue* nodes. We let H denote the subgraph of G consisting of only red nodes.

Recall that for a point $i \in T$ and a set of k coordinates S_j, the section $V_{S_j}(i)$ is the set of all points that are within a distance ω from i and that can differ from i only in the k coordinates that are in S_j. Hence each $V_{S_j}(i)$ section is in the form of a k-dimensional ℓ_p ball of radius ω centered at i. For the purpose of our proof we want to work with sections that are in the form of a cube. In the case of the ℓ_∞ norm each $V_{S_j}(i)$ section is already in the shape of a k-dimensional cube centered at i with side length 2ω. For $1 \leq p < \infty$ we note that each $V_{S_j}(i)$ section contains a subsection in the shape of a k-dimensional cube with side length $2\omega/k^{\frac{1}{p}}$.

We define b_p as follows:

$$b_p = \begin{cases} \lfloor 2\omega \rfloor & \text{if } p = \infty \\ \left\lfloor 2\omega/k^{\frac{1}{p}} \right\rfloor & \text{if } 1 \leq p < \infty. \end{cases}$$

Then each section $V_{S_j}(i)$ contains a subsection that is a k-dimensional cube of side length b_p centered at i. We take the floor in order to ensure that b_p is an integer. Figure 4.1 shows an example for the ℓ_2 norm and $k = 2$. The outer circle is an ℓ_2 ball of radius ω centered at the origin, the center point is the orgin and the four other points denote the corners of the inner square.

We now let $m \geq 2$ be a positive even integer. Without loss of generality we assume that b_p is a multiple of m. For each point $i \in T$ and each section $V_{S_j}(i)$, we partition the inner cubical

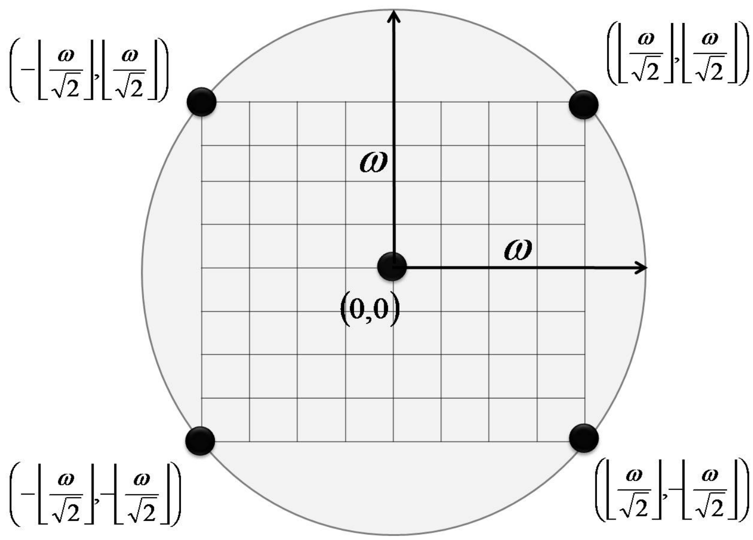

Figure 4.1: Square with side length $2\omega/\sqrt{2}$ inside ℓ_2 ball of radius ω.

section of $V_{S_j}(i)$ into m^k equal size subsections. That is we tile the inner cube contained in each section into m^k smaller cubes. Each of these subsections is a k-dimensional cube with side length b_p/m. We let $n' = (b_p/m)^k$. Then the number of red nodes in a subsection is $Bin(n', p_1)$ with mean

$$n'p_1 \quad = \quad \left(\frac{b_p}{m}\right)^k p_1.$$

As $n \to \infty$ we have $n'p_1 \sim \gamma \ln n$ where

$$\gamma := \begin{cases} \left(\frac{2}{m}\right)^k \frac{(d-k\delta)}{a_p} & \text{if } p = \infty \\ \left(\frac{2}{k^{\frac{1}{p}} m}\right)^k \frac{(d-k\delta)}{a_p} & \text{if } 1 \le p < \infty. \end{cases}$$

Since $a_p = \binom{d}{k} \int_{B_p^k(0,1)} dx \ge 1$, $m \ge 2$ and $k^{\frac{k}{p}} \ge 1$ we have:

$$\begin{aligned} \gamma &\le d - k\delta \\ &< d \qquad \text{since } k \le d \text{ and } \delta < 1. \end{aligned}$$

Still following the same argument and terminology of Frize et al. [10] we say that a section $V_{S_j}(i)$ is *mighty* if all the subsections in its inner cube have at least $\gamma \ln n/10$ red nodes.

Let $\beta := \lfloor d/\gamma \rfloor$. We note that since $\gamma < d$ we have $\beta > 0$. Given two nodes $v, u \in H$ and two sets S_j and S_ℓ we say that the section $V_{S_\ell}(u)$ is *orthogonal* to the section $V_{S_j}(v)$ if $S_j \cap S_\ell = \emptyset$. Still following the proof method in [10], we now define several events and show that they hold with high probability.

Lemma 4. *If $k \leq \frac{d}{2}$, let ϵ_1 denote the event that there does not exist a red node v that has a section $V_{S_j}(v)$ on which we can find β red nodes each having a non-mighty section orthogonal to $V_{S_j}(v)$.*

If $\frac{d}{2} < k < d$, let ϵ_1 denote the event that there does not exist a red node v that has a section $V_{S_j}(v)$ on which we can find β red nodes $\{u_1, \cdots, u_\beta\}$ such that for some set S_ℓ the sections $V_{S_\ell}(u_i)$ are all non-mighty and pairwise non-intersecting.

Then ϵ_1 holds with high probability.

Proof. For a fixed point i and section $V_{S_j}(i)$ we have

$$
\begin{aligned}
\Pr\left\{\ V_{S_j}(i) \text{ not mighty }\right\} \ &\leq\ m^k \Pr\left\{Bin\,(n', p_1) \leq (\gamma \ln n)/10\right\} \\
&\sim\ m^k e^{\frac{\gamma}{10}\ln n - \gamma \ln n - \frac{\gamma}{10}\ln n \ln(0.1)} \text{ as } n \to \infty \\
&\leq\ m^k e^{-\gamma \ln n},
\end{aligned}
$$

where we used the following Chernoff tail bound $\Pr\left\{Bin\,(n,p) \leq t\right\} \leq e^{t - np - t\ln(t/np)}$ found in [6] and [1].

Note that γ is a strictly positive constant. For both $k \leq \frac{d}{2}$ and $k > \frac{d}{2}$ we can use the following upper bound:

$$\begin{aligned}
\Pr\left\{\overline{\epsilon_1}\right\} &\leq n^d p_1 \binom{d}{k}^2 \binom{\left|V_{S_j}(i)\right|}{\beta} p_1^\beta \left(m^k e^{-\gamma \ln n}\right)^\beta \\
&\sim \text{const.} \cdot n^{d-k\delta} n^{-\gamma\beta} (\ln n)^\beta \qquad \text{as } n \to \infty. \\
&\to 0 \qquad \text{as } n \to \infty.
\end{aligned}$$

$\square$

Lemma 5. *Let ϵ_2 denote the event that there does not exists a red node v with deg$(v) < \ln \ln n$ that has a red neighbour w such that w has a non-mighty section. Then ϵ_2 holds with high probability.*

Proof. We obtain:

$$\begin{aligned}
\Pr\left\{\overline{\epsilon_2}\right\} &\leq n^d p_1 \sum_{t=1}^{\ln \ln n} \binom{|V(i)|}{t} p_1^t (1-p_1)^{|V(i)|-t} \binom{d}{k} m^k e^{-\gamma \ln n} \\
&\sim \text{const.} \cdot n^{-\gamma} (\ln n)^{\ln \ln n} \ln \ln n \qquad \text{as } n \to \infty. \\
&\to 0 \qquad \text{as } n \to \infty.
\end{aligned}$$

$\square$

Lemma 6. *Let ϵ_3 denote the event that every red vertex has at least one red neighbour. Then ϵ_3 holds with high probability.*

Proof. We obtain:

$$\begin{aligned}
\Pr\left\{\overline{\epsilon_3}\right\} &\leq n^d p_1 (1-p_1)^{|V(i)|} \\
&\sim n^d p^* (1-p^*)^{|V(i)|} \qquad \text{as } n \to \infty. \\
&\to 0 \qquad \text{as } n \to \infty.
\end{aligned}$$

$\square$

Lemma 7. *Let ϵ_4 denote the event that every blue vertex has at least one red neighbour. Then ϵ_4 holds with high probability.*

Proof. We obtain:

$$\begin{aligned}\Pr\{\overline{\epsilon_4}\} &\leq n^d p_2 (1-p_1)^{|V(i)|} \\ &\leq n^d p_1 (1-p_1)^{|V(i)|} \quad \text{for } n \text{ large enough} \\ &\to 0 \quad \text{as } n \to \infty.\end{aligned}$$

$\square$

Let $s := b_p/m$. We define the box $B(s) := \{y \in T : \| y_i \| \leq s \ , \forall 1 \leq i \leq d\}$. For a point $x \in T$ we let $B(x, s) := x + B(s)$, where we use addition within the torus. We now state a lemma that is the counterpart of lemma 2.5 in [10].

Lemma 8. *Assume that the high probability events $\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4$ all hold and that G does not have any isolated nodes. Then for each point $x \in T$ and each node $v \in H$ there exists a node $\phi(x, v)$ that is in the same component of H as v and lies in the box $B(x, s)$.*

Proof. All the nodes mentioned in this proof are assumed to be red nodes. We note here that the proof of this lemma differs in style from the proof of lemma 2.5 in [10]. Fix a point $x \in T$ and a node $v \in H$. We let $\phi(x, v)$ be the node returned by Algorithm 1.

We define some terms used in the algorithm. For a given point $u \in T$, u_i denotes its i^{th} coordinate. The distance between a subsection and a point $x \in T$ is given by

$$\max_{u:\ u \in \text{subsection}} \| u - x \|_{\ell_p} ,$$

where the maximum is taken over all points u in the subsection and the distance is measured assuming that the points lie on a torus. By the closest subsection to the origin, we mean the subsection that minimizes this distance.

The algorithm starts at node v and moves from neighbouring node to neighbouring node until it reaches a node inside $B(x, s)$. At the start of the algorithm all coordinates are declared unfixed. Each time the algorithm moves to a new node it calls Algorithm 2 which checks whether each of the coordinates of the current node are within a distance s of the corresponding coordinate of the point x and declares them as fixed if they are. Once all coordinates become fixed the algorithm has reached a node inside $B(x, s)$.

We now prove that Algorithm 1 works correctly. The algorithm first checks if v lies in $B(x, s)$ in which case it just returns the node v. Otherwise, it moves to a neighbour of v. If $\deg(v) < \ln \ln n$, then it selects any neighbour w of v and picks a set of k coordinates such that the number of unfixed coordinates in this set is maximized. We know that with high probability there is at least one neighbour w from ϵ_3. Furthermore ϵ_2 guarantees us that all sections of w are mighty. In particular the section corresponding to the chosen set is mighty.

If instead $\deg(v) \geq \ln \ln n$, the way the algorithm picks a neighbour w of v depends on whether $k \leq d/2$ or $k > d/2$. In the first case it picks a w such that the sections of w which are orthogonal to the section of v containing w are mighty. There is always a choice of at least one orthogonal section since $k \leq d/2$. It follows from ϵ_1 that there always is a choice of $\ln \ln n - \binom{d}{k}\beta$ neighbours of v with mighty orthogonal sections. The algorithm then selects a set of k coordinates such that the number of unfixed coordinates in this set is maximized and the section of w corresponding to this set is mighty. In the second case the algorithm first selects a set that has the maximum number of unfixed coordinates among all sets S_ℓ for which there exists a section of v that contains more than β nodes with pairwise non-intersecting V_{S_ℓ} sections. To see that such a choice of set always exists note that since $\deg(v) \geq \ln \ln n$ and the maximum number of nodes in a section must be greater than the average, there must be a section of v that has at least $\ln \ln n / \binom{d}{k}$ nodes. Let this section be $V_{S_j}(v)$. Then there must be a coordinate i such that there are at least $(\ln \ln n / \binom{d}{k})^{1/k}$ nodes

Algorithm 1 Compute $\phi(x, v)$

Input: node $v \in H$, point $x \in T$

Output: node $\phi(x, v) \in B(x, s)$ that is in same component of H as v

for $i = 1$ to d **do**
  coordinate$_i$ $\leftarrow$ unfixed
end for
update(v)
if $v \in B(x, s)$ **then**
  return v
end if
if $deg(v) < \ln \ln n$ **then**
  currNode $\leftarrow$ any neighbour of v
  currSet $\leftarrow$ set S_ℓ with max number of unfixed coordinates
else
  if $k \leq d/2$ **then**
    currNode $\leftarrow$ neighbour of v with mighty orthogonal sections
    currSet $\leftarrow$ set S_ℓ with max number of unfixed coordinates such that V_{S_ℓ}(currNode) is mighty
  else
    currSet $\leftarrow$ set S_ℓ with max number of unfixed coordinates such that there exists a section of currNode that contains more than β nodes with pairwise non-intersecting V_{S_ℓ} sections
    currNode $\leftarrow$ neighbour w such that the section $V_{\text{currSet}}(w)$ is mighty
  end if
end if
update(currNode)
while $\exists\ i$ such that coordinate$_i$ $=$ unfixed **do**
  currSection $\leftarrow V_{\text{currSet}}$(currNode)
  currSubsection $\leftarrow$ subsection in currSection that is closest to x
  if $k \leq d/2$ **then**
    currNode $\leftarrow$ node in currSubsection with mighty orthogonal sections
    currSet $\leftarrow$ set S_ℓ with max number of unfixed coordinates such that V_{S_ℓ}(currNode) is mighty
  else
    currSet $\leftarrow$ set S_ℓ with max number of unfixed coordinates such that currSubsection contains more than β nodes with pairwise non-intersecting V_{S_ℓ} sections
    currNode $\leftarrow$ neighbour w such that the section $V_{\text{currSet}}(w)$ is mighty
  end if
  update(currNode)
end while
return currNode

Algorithm 2 Update(w)

for $i = 1$ to d **do**
  if $|w_i - x_i| \leq s$ **then**
    coordinate$_i \leftarrow$ fixed
  end if
end for

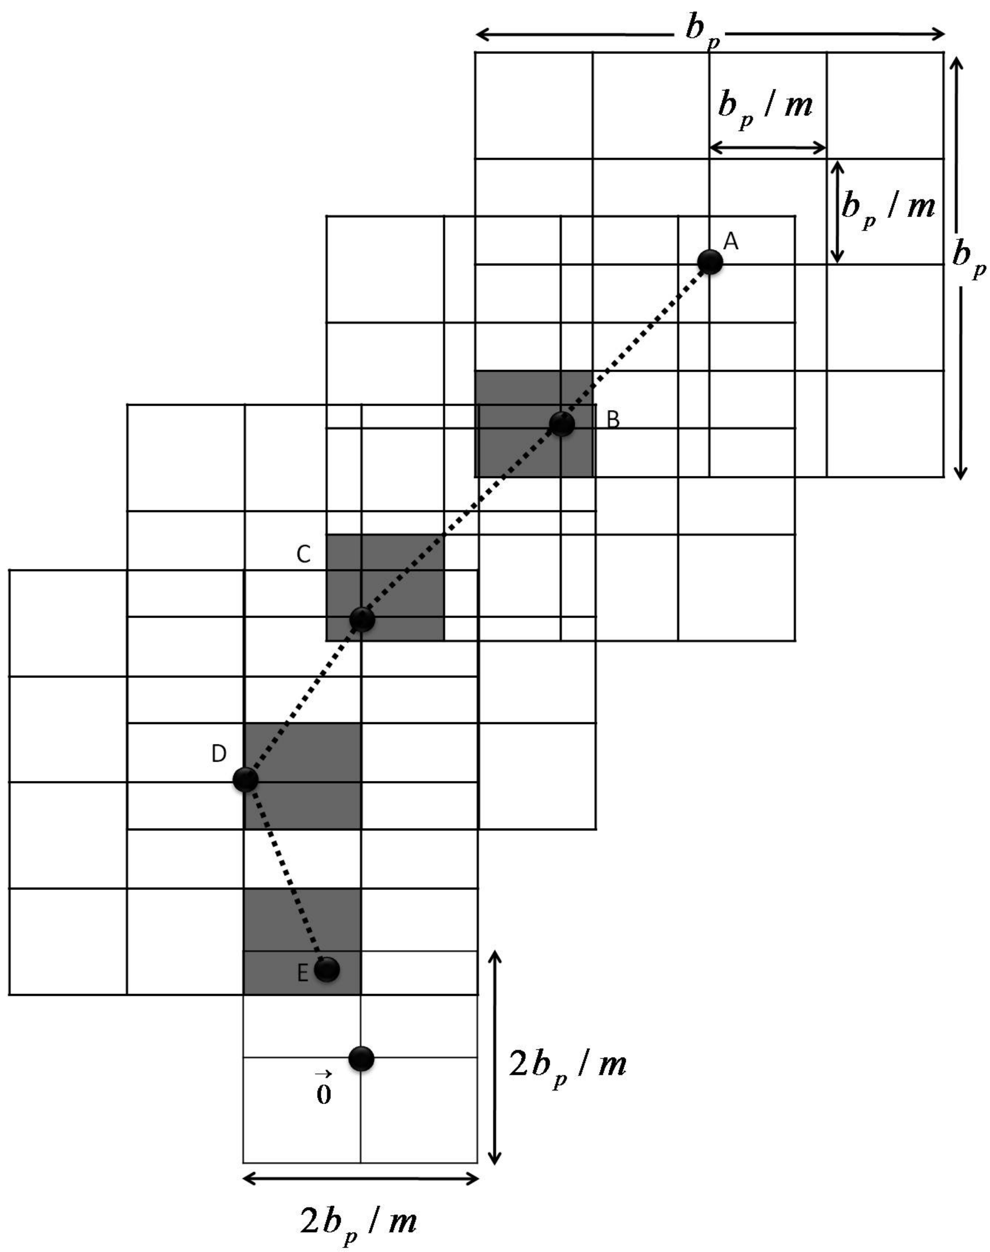

Figure 4.2: Fixing of coordinates in the xy-plane.

in $V_{S_j}(v)$ with distinct i^{th} coordinates. Choose S_ℓ such that $\ell \neq j$ and $i \notin S_\ell$. The number of choices for such an S_ℓ is least one since $k \leq d-1$. Then any two nodes that differ in the i^{th} coordinate have non-intersecting V_{S_ℓ} sections. The algorithm then picks w to be a neighbour of v with a mighty V_{S_ℓ} section. Note that it follows from ϵ_1 that there is always a node $w \in V_{S_j}(v)$ with a mighty $V_{S_\ell}(w)$ section.

At the start of every while loop the variable currNode holds the current position of the algorithm and currSet refers to a chosen set of k coordinates. Each iteration of the loop consists of three steps. First the algorithm makes the assignment currSection $\leftarrow V_{\text{currSet}}(\text{currNode})$. Note that currSection will always be mighty. Second, it chooses the subsection in currSection that is closest to the point x. Note that this subsection will always have at least $\gamma \ln n/10$ nodes. And third, it moves to a node in this subsection and chooses a new set of k coordinates. This is done again according to whether $k \leq d/2$ or $k > d/2$. In the first case a node is picked such that the sections of this node which are orthogonal to currSection are mighty. Again from ϵ_1 we know there is always a choice of $\gamma \ln n/10 - \beta$ such nodes. The algorithm then selects a set of k coordinates such that the number of unfixed coordinates in this set is maximized and the section of the new node corresponding to this set is mighty. In the second case where $k > d/2$ the algorithm first selects a set that has the maximum number of unfixed coordinates among all sets S_ℓ for which there are more than β nodes in the chosen subsection with pairwise non-intersecting V_{S_ℓ} sections. Note that there exists a coordinate i such that there are at least $(\gamma \ln n/10)^{1/k}$ nodes in the current subsection with distinct i^{th} coordinates so choosing S_ℓ such that $i \notin S_\ell$ gives us such a set. Finally it picks a node in the chosen subsection that has a mighty section corresponding to this choice of set. It follows from ϵ_1 that there always at least one such node.

It remains to show that each time the algorithm moves to a new neighbour in a while loop iteration the distance to x is decreased for all the unfixed coordinates that are in the current set, and stays within a range of s for all the fixed coordinates in the current set. Suppose that the current set contains the first k coordinates and let u denote the current node. Assume without loss of generality that $m = 4$, $u_i > x_i$ for all i, and $u_i - x_i > s$ for all $1 \leq i \leq j$, while $u_i - x_i < s$ for all $j+1 \leq i \leq k$. That is the first j coordinates are not within a range s of x while the last $k - j$ are. Then the algorithm selects the subsection $[u_1 - s, u_1 - 2s] \times \cdots \times [u_j - s, u_j - 2s] \times [u_{j+1}, u_{j+1} - s] \times \cdots \times [u_k, u_k - s]$. Thus for $1 \leq i \leq j$ the distance between the i^{th} coordinate of any point in this subsection and x_i is smaller than $u_i - x_i$, and for $j + 1 \leq i \leq k$ the distance between the i^{th} coordinate of any point in this subsection and x_i is at most s.

This implies that after a sufficient number of iterations the algorithm fixes all the coordinates, and once a coordinate is fixed it remains within a range s of the corresponding x coordinate. Therefore when the algorithm terminates it returns the node $\phi(x, v)$ that lies in $B(x, s)$. Since we always move from neighbour to neighbour, this node is in the same component of H as the starting node v.

Figure 4.2 shows an example of how the algorithm works. In this example $k = 2$, $d \geq 4$ and $m = 4$. The plane depicted in the figure is the xy plane and we are interested in seeing how the x and y coordinates become fixed. We note that the algorithm is not able to move in the same plane in two consecutive steps. For the purposes of this example we assume that in between any two moves in the xy plane the algorithm moves in an orthogonal plane to the xy plane so that the x and y coordinates remain unchanged.

Suppose now that the algorithm starts at node A and wants to reach a node that is inside the box of side length $2s$ centered at the origin. Since the x and y coordinates of node A are not within a distance s of zero they both start off as unfixed. The algorithm selects the subsection of node A that is closest to the origin. This is given by the shaded subsection in the bottom left corner. The algorithm then selects a node in the chosen subsection whose orthogonal planes are mighty. Note that the algorithm has no control over where in the chosen subsection this node is located, it just knows that such a node exists. In this example we assume that the selected node is node B. Since the coordinates of node B are not within a distance s of zero they remain unfixed. Again the subsection of node B that is closest to the origin is in the bottom left corner. The algorithm moves to node C in this subsection.

Now the x coordinate of node C is exactly zero therefore it becomes fixed. Since the y coordinate is not within a distance s of zero it remains unfixed. There are now two subsections of node C that achieve the minimum distance to the origin. These are the shaded subsection containing node D and the subsection immediately to its right. The algorithm can chose either one of these two subsections. Suppose that it chooses the shaded one and moves to node D. Node D happens to be at the left most edge of the chosen subsection. Hence by moving from node C to node D the algorithm actually increases the distance to the origin in the x coordinate. However the x coordinate of node D cannot be further than a distance s from the zero, hence it remains fixed. Since the y coordinate is still unfixed the algorithm performs another iteration and moves to node E. Now the y coordinate becomes fixed and node E is inside the box of side length $2s$ centered at the origin. If the algorithm were to perform additional iterations in the xy-plane after both the x and y coordinates have been fixed the x and y coordinates of the new nodes will still each remain within a distance s of zero. $\square$

We now show how placing the blue nodes in the second stage can guarantee that the final graph is connected with high probability, following the proof method in [10]. Let L be the set of points in T whose coordinates are multiples of 3ω. For each component K of H let v_K be the lexicographically smallest node in K. Given two distinct points $x = (x_1, \cdots, x_d)$ and $y = (y_1, \cdots, y_d)$ we say that x is lexicographically smaller than y if there exists $j \in \{1, \cdots, d\}$ such that $x_j < y_j$ and $x_i = y_i$ for all $i < j$. Now suppose we have two distinct components K and J of H. For each $x \in L$ let $\phi(x, v_K)$ and $\phi(x, v_J)$ be the unique nodes returned by algorithm 1 or 2 depending on each case. Note that by the assumption that K and J are distinct components it follows that the nodes $\phi(x, v_K)$ and $\phi(x, v_J)$ are also distinct for all $x \in L$.

For $d/2 < k < d$ let $z(J, K, x)$ be the unique point in $B(x, s)$ that agrees with $\phi(x, v_K)$ in the first $d - k$ coordinates and with $\phi(x, v_J)$ in the last k coordinates. Then $z(J, K, x)$

is visible from both $\phi(x, v_K)$ and $\phi(x, v_J)$. Therefore if we were to place a blue node at $z(J, K, x)$ then $\phi(x, v_K)$ and $\phi(x, v_J)$ would become linked up.

For $k \leq d/2$ let $p(J, K, x)$ be the unique path whose points all lie in $B(x, s)$ that is picked as follows. We let the first point of the path agree with $\phi(x, v_K)$ in the first $d - k$ coordinates and with $\phi(x, v_J)$ in the last k coordinates. Hence the first point of the path is mutually visible with $\phi(x, v_K)$. We let the second point agree with $\phi(x, v_K)$ in the first $d - 2k$ coordinates and with $\phi(x, v_J)$ in the last $2k$ coordinates. This ensures that the second point of the path is mutualy visible with the first point of the path. We continue in this way until the last point is mutually visible with $\phi(x, v_J)$. Then this path has at most $\lceil d/k \rceil - 1$ points. The first point of the path is visible from $\phi(x, v_K)$ and the last point of the path is visible from $\phi(x, v_J)$. Therefore if we were to place a blue node at each point of the path then $\phi(x, v_K)$ and $\phi(x, v_J)$ would become linked up.

Lemma 9. *For distinct points x and y in L the points $z(J, K, x)$ and $z(J, K, y)$ are distinct and the paths $p(J, K, x)$ and $p(J, K, y)$ do not have any points in common.*

Proof. This follows since $z(J, K, x)$ and $p(J, K, x)$ lie in $B(x, s)$, while $z(J, K, y)$ and $p(J, K, y)$ lies in $B(y, s)$. These boxes are disjoint since they each have side length $2s <\leq 2\omega$ and x and y are at least 3ω apart. $\square$

Therefore for a fixed pair of components J and K there are $(n/3\omega)^d$ such points (or paths) that could link them up. The probability of not placing a blue node at a fixed point is $1 - p_2$, while the probability of not placing a blue node at every point in a fixed path of length ℓ is $1 - p_2^\ell$. The total number of components of H is upper bounded by $(2s)^d \leq (2\omega)^d$ since for any fixed point $x \in L$ each component has a point in the box $B(x, s)$ and this box has a total of $(2s)^d$ points.

For $d/2 < k < d$ we have

$$\begin{aligned}
\Pr\{H \text{ is not connected}\} &\leq \Pr\{\exists \text{ components } J \text{ and } K \text{ that are not linked up}\} \\
&= (2\omega)^d (1-p_2)^{n^d/(3\omega)^d} \\
&\leq (2\omega)^d e^{-\frac{n^{d-(d+k)\delta}}{a_p 3^d \ln n}} \\
&\to 0 \qquad \text{if } \delta < \frac{d}{d+k}.
\end{aligned}$$

And similarly for $k \leq d/2$

$$\begin{aligned}
\Pr\{H \text{ is not connected}\} &\leq \Pr\{\exists \text{ components } J \text{ and } K \text{ that are not linked up}\} \\
&\leq (2\omega)^d (1-p_2^{\lceil d/k \rceil - 1})^{n^d/(3\omega)^d} \\
&\leq (2\omega)^d e^{-\frac{n^{d-\left(d+k\left(\lceil \frac{d}{k} \rceil - 1\right)\right)\delta}}{3^d (a_p \ln n)^{\lceil d/k \rceil - 1}}} \\
&\to 0 \qquad \text{if } \delta < \frac{d}{d+k\left(\lceil \frac{d}{k} \rceil - 1\right)}.
\end{aligned}$$

Thus assuming the high probability events $\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4$ hold and G has no isolated nodes, H is connected with high probability. From ϵ_4 we know that every blue node has at least one red neighbour, and therefore it follows that G must be connected with high probability.

4.4 The middle case

We finally consider the case $c_n \to c$. The proof in this section also follows the general structure of the one in [10]. In Section 4.1 we have shown that if $c_n \to c$ then $\lim_{n\to\infty} \mathrm{E}\{N\} = \lambda$, where N is the number of isolated nodes. Let S_0 be the set of isolated nodes. Let N'' denote the number of pairs of nodes $i, j \in S_0$ such that $i \sim j$, that is X_i and X_j are not independent. Recall that X_i and X_j are not independent if and only if $V(i) \cap V(j) \neq \emptyset$. Let

N' denote the number of nodes $i \in S_0$ such that X_i and X_j are independent for all $j \in S_0$. Then

$$N' \leq N \leq N' + N''.$$

We have

$$\mathrm{E}\left\{N''\right\} = \sum_{i \sim j} \mathrm{E}\left\{X_i X_j\right\} \to 0,$$

as shown in Section 4.2.

Therefore $N = N'$ with high probability and

$$\lim_{n\to\infty} \mathrm{E}\left\{N'\right\} = \lim_{n\to\infty} \mathrm{E}\left\{N\right\} = \lambda. \tag{4.1}$$

Let t be a positive integer. For a random variable X we define $(X)_t = t!\binom{X}{t}$. We obtain the following upper and lower bounds on $\mathrm{E}\left\{(N')_t\right\}$

$$\begin{aligned}
\mathrm{E}\left\{(N')_t\right\} &\leq t!\binom{n^d}{t}\left(p^*(1-p^*)^{|V(i)|}\right)^t \\
&\sim \left(n^d p^*(1-p^*)^{|V(i)|}\right)^t \\
&= \left(\mathrm{E}\left\{N\right\}\right)^t \to \lambda^t \qquad \text{as } n \to \infty \\
\mathrm{E}\left\{(N')_t\right\} &\geq \left(\left(n^d - t\left(|V(i)|\right)^2\right) p^*(1-p^*)^{|V(i)|}\right)^t \\
&\sim \left(n^d p^*(1-p^*)^{|V(i)|}\right)^t \\
&= \left(\mathrm{E}\left\{N\right\}\right)^t \to \lambda^t \qquad \text{as } n \to \infty.
\end{aligned}$$

Therefore

$$\lim_{n\to\infty} \mathrm{E}\left\{(N')_t\right\} = \lambda^t. \tag{4.2}$$

We state here a lemma which is Theorem 8.3.1 in the book by Alon and Spencer [1]. Let $X = X_1 + \cdots + X_m$ where $X_1, \cdots, X_m$ are indicator variables for random events.

Lemma 10. ***(Alon and Spencer)*** *Suppose that $E\{X\} \to \mu$ for some constant μ and for every fixed r we have $E\left\{\frac{(X)_t}{t!}\right\} \to \frac{\mu^t}{t!}$. Then*

$$Pr\{X = 0\} \to e^{-\mu},$$

and

$$Pr\{X = t\} \to \frac{\mu^t}{t!}e^{-\mu}.$$

Using equations 4.1 and 4.2 and applying Lemma 10 we obtain that N' is asymptotically Poisson with mean λ. This implies that

$$\lim_{n\to\infty} \Pr\{G \text{ has an isolated vertex}\} \quad = \quad 1 - \Pr\{Pois(\lambda){=}0\} = 1 - e^{-\lambda}.$$

Finally, conditioning on the event that G has no isolated nodes we can use the proof in Section 4.3 to show that G is connected with high probability. Therefore

$$\lim_{n\to\infty} \Pr\{G \text{ is connected}\} \quad = \quad e^{-\lambda}.$$

Chapter 5

Conclusion

In this thesis we introduced a generalized model of line-of-sight networks and we proved connectivity results for this model. We showed that the threshold value for connectivity is given by $p^* = \left((d - k\delta)\ln n + \ln\ln n + c_n\right)/a_p\omega^k$, where n is the side length of the underlying grid, d is the dimension of the grid, k is the maximum number of coordinates that two grid points can differ in in order to be mutually visible, $\omega = n^\delta$ is the range parameter, $c_n = o(\ln\ln n)$ and a_p is a constant dependent on the choice of norm.

We proved this result in three stages. In section 4.2 we first proved that if $c_n \to -\infty$ then the graph is non-connected with high probability. We used a proof method known as the second moment method, a detailed description of which can be found in [1].

In section 4.3 we proved that if $c_n \to \infty$ then the graph is connected with high probability. To show this we first showed that with high probability the graph does not have any isolated nodes. We then extended a proof found in the original paper by Frieze et al. [10]. This method consists in first creating an equivalent line-of-sight network using a two stage process for placing the nodes. We then showed that if there are any two disconnected components after the first stage then they will be linked up with high probability by the nodes placed in the second stage.

In section 4.4 we showed that if $c_n \to c$ then the probability that the graph is connected is asymptotically $e^{-\lambda}$. We proved this by showing that the number of isolated nodes is asymptotically Poisson with mean λ.

Our results for connectivity of line-of-sight networks in higher dimensions parallel those of the two-dimensional model introduced by Frieze et al. as well as random graphs and random geometric graphs. Specifically if we start with single nodes and increase the placement probability $p*$ (respectively edge probability p for random graphs and range parameter r for random geometric graphs) continuously, the resulting graph becomes connected with high probability at the exact instant when there are no more isolated nodes.

Opportunities for future work include the study of other structural properties of the generalized model such as the existence of a giant component as well as algorithms for passing messages between nodes. It could also be of interest to explore modifications of this model to better capture the specific properties of certain wireless networks such as networks where nodes or links can fail, networks were nodes become active at certain time intervals or networks with different geometric configurations.

In all these cases the connectivity of the network is of fundamental importance but there are also other interesting properties to be studied such as the capacity of the network, efficient algorithms for information propagation as well as routing and broadcasting protocols.

Selected Bibliography Including Cited Works

[1] N. Alon and J.H. Spencer. *The Probabilistic Method*. Wiley-Interscience, 2000.

[2] M. Appel and R. Russo. The Maximum Vertex Degree of a Graph on Uniform Points in $[0,1]^d$. *Advances in Applied Probability*, 29:567–581, 1997.

[3] M.J.B. Appel and R.P. Russo. The minimum vertex degree of a graph on uniform points in $[0,1]^d$. *Advances in Applied Probability*, 29(3):582–594, 1997.

[4] M.J.B. Appel and R.P. Russo. The connectivity of a graph on uniform points on $[0,1]^d$. *Statistics & Probability Letters*, 60(4):351–357, 2002.

[5] B. Bollobás. *Random Graphs*. Cambridge University Press, 2001.

[6] H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, 23(4):493–507, 1952.

[7] A. Czumaj and X. Wang. Communication problems in random line-of-sight ad-hoc radio networks. *Stochastic Algorithms: Foundations and Applications*, pages 70–81, 2007.

[8] P. Erdős and A. Rényi. On random graphs. *Publ. Math. Debrecen*, 6(290-297):156, 1959.

[9] P. Erdős and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(1):17–60, 1960.

[10] A. Frieze, J. Kleinberg, R. Ravi, and W. Debany. Line-of-sight networks. *Combinatorics, Probability and Computing*, 18(1-2):145–163, 2008.

[11] A. Goel, S. Rai, and B. Krishnamachari. Sharp thresholds for monotone properties in random geometric graphs. In *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*, page 586. ACM, 2004.

[12] P. Gupta and P.R. Kumar. Critical power for asymptotic connectivity in wireless networks. *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of WH Fleming*, 3(20):547–566, 1998.

[13] P. Gupta and P.R. Kumar. Internets in the sky: The capacity of three dimensional wireless networks. *Communications in Information and Systems*, 1(1):33–49, 2001.

[14] M. Penrose. On k-connectivity for a geometric random graph. *Random Structures and Algorithms*, 15(2):145–164, 1999.

[15] M. Penrose. *Random Geometric Graphs*. Oxford University Press, USA, 2003.

[16] R. Van Der Hofstad. Random graphs and complex networks. *preparation, see http://www. win. tue. nl/ ~ rhofstad/NotesRGCN. pdf*, 2009.

[17] X. Wang. Volumes of generalized unit balls. *Mathematics Magazine*, 78(5):390–395, 2005.