

ASYMPTOTIC PROPERTIES OF WIRELESS MULTI-HOP NETWORK

THÈSE N° 3310 (2005)

PRÉSENTÉE À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS

Institut de systèmes de communication

SECTION DES SYSTÈMES DE COMMUNICATION

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Olivier DOUSSE

ingénieur physicien diplômé EPF
de nationalité suisse et originaire Du Mouret (FR)

acceptée sur proposition du jury:

Prof. P. Thiran, directeur de thèse
Prof. F. Baccelli, rapporteur
Prof. P. Kumar, rapporteur
Prof. M. Vetterli, rapporteur

Lausanne, EPFL
2005

Abstract

In this dissertation, we consider wireless multi-hop networks, where the nodes are randomly placed. We are particularly interested in their asymptotic properties when the number of nodes tends to infinity. We use percolation theory as our main tool of analysis.

As a first model, we assume that nodes have a fixed connectivity range, and can establish wireless links to all nodes within this range, but no other (Boolean model). We compute for one-dimensional networks the probability that two nodes are connected, given the distance between them. We show that this probability tends exponentially to zero when the distance increases, proving that pure multi-hopping does not work in large networks. In two dimensions however, an unbounded cluster of connected nodes forms if the node density is above a critical threshold (super-critical phase). This is known as the *percolation phenomenon*. This cluster contains a positive fraction of the nodes that depends on the node density, and remains constant as the network size increases. Furthermore, the fraction of connected nodes tends rapidly to one when the node density is above the threshold. We compare this *partial connectivity* to *full connectivity*, and show that the requirement for full connectivity leads to vanishing throughput when the network size increases. In contrast, partial connectivity is perfectly scalable, at the cost of a tiny fraction of the nodes being disconnected.

We consider two other connectivity models. The first one is a signal-to-interference-plus-noise-ratio based connectivity graph (STIRG). In this model, we assume deterministic attenuation of the signals as a function of distance. We prove that percolation occurs in this model in a similar way as in the previous model, and study in detail the domain of parameters where it occurs. We show in particular that the assumptions on the attenuation function dramatically impact the results: the commonly used power-law attenuation leads to particular symmetry properties. However, physics imposes that the received signal cannot be stronger than the emitted signal, implying a bounded attenuation function. We observe that percolation is harder to achieve in most cases with such an attenuation function.

The second model is an information theoretic view on connectivity, where two arbitrary nodes are considered connected if it is possible to transmit data from one to the other at a given rate. We show that in this model the same partial connectivity can be achieved in a scalable way as in the Boolean model. This result is however a pure connectivity result in the sense that there is no competition and interferences between data flows. We also look at the other extreme, the Gupta and Kumar scenario, where all nodes want to transmit data simultaneously. We show first that under point-to-point communication and bounded attenuation function the total transport capacity of a fixed area

network is bounded from above by a constant, whatever the number of nodes may be. However, if the network area increases linearly with the number of nodes (constant density), or if we assume power-law attenuation function, a throughput per node of order $1/\sqrt{n}$ can be achieved. This latter result improves the existing results about random networks by a factor $\sqrt{\log n}$.

In the last part of this dissertation, we address two problems related to latency. The first one is an intruder detection scenario, where a static sensor network has to detect an intruder that moves with constant speed along a straight line. We compute an upper bound to the time needed to detect the intruder, under the assumption that detection by disconnected sensors does not count. In the second scenario, sensors switch off their radio device for random periods, in order to save energy. This affects the delivery of alert messages, since they may have to wait for relays to turn on their radio to move further. We show that asymptotically, alert messages propagate with constant, deterministic speed in such networks.

Résumé

Ce travail de thèse a pour but d'étudier les réseaux sans fil, dans lesquels les nœuds agissent comme relais pour les autres nœuds (réseaux à relais multiples). Nous nous intéressons en particulier à leur comportement asymptotique lorsque le nombre nœuds tend vers l'infini. Nous utilisons principalement la théorie de la percolation pour parvenir à cette fin.

Comme premier modèle, nous faisons l'hypothèse que les nœuds sont capable d'établir une communication radio avec tous leurs voisins situés dans un certain périmètre autour d'eux, mais pas plus loin (modèle Booléen). Pour les réseaux unidimensionnels, nous calculons la probabilité que deux nœuds peuvent communiquer (avec l'aide d'autres nœuds agissant comme relais), en fonction de la distance qui les sépare. Nous montrons que cette probabilité tend exponentiellement vers zéro lorsque la distance augmente. Ainsi, aucune communication à longue distance n'est possible dans les réseaux unidimensionnels. Dans les réseaux bidimensionnels, si la densité de nœuds est suffisante (phase sur-critique), une clique géante de nœuds connectés se forme (*phénomène de percolation*). Cette clique contient une fraction non-nulle des nœuds, laquelle ne dépend que de la densité de nœuds et de leur portée. De plus, cette fraction tend rapidement vers 1 lorsque la densité augmente. Nous comparons cette connectivité *partielle* à la connectivité *complète* du réseau, et montrons que cette dernière est très coûteuse en termes de performance du réseau. Par contre, la connectivité partielle ne demande pas d'adaptation des paramètres du réseau pour être maintenue lorsque celui-ci s'agrandit.

Dans ce travail, nous étudions deux autres modèles de connectivité, plus sophistiqués. Dans le premier, on considère que deux nœuds peuvent établir une communication sans fil si le rapport signal-sur-interférence-plus-bruit est supérieur à un seuil prescrit. Pour calculer ce rapport, nous faisons l'hypothèse que les signaux s'affaiblissent selon une fonction déterministe de la distance parcourue (fonction d'atténuation). Nous montrons que dans ce modèle, le même phénomène de percolation a lieu, et nous décrivons les conditions sous lesquelles il se réalise. Notamment, nous montrons que les hypothèses faites sur la fonction d'atténuation ont un impact important sur les résultats. En particulier, l'hypothèse usuelle qui consiste à prendre l'atténuation égale à une puissance négative de la distance confère au modèle des propriétés de symétrie particulières, mais malheureusement irréalistes, vu que la fonction d'atténuation doit être bornée pour des raisons physiques évidentes.

Le second modèle consiste à déclarer deux nœuds connectés s'il est possible de transmettre de l'information de l'un à l'autre avec un certain débit prescrit, avec l'aide de tous les autres nœuds. Nous montrons que dans ce modèle, il est aussi possible de garder une fraction arbitraire de nœuds connectés entre eux,

sans changer le débit prescrit, lorsque la taille et le nombre de nœuds du réseau augmente.

Les résultats ci-dessus sont des résultat de connectivité pure, vu que nous ne regardons qu'une paire de nœuds à la fois, et qu'il n'y a donc pas d'interférence entre plusieurs flots de données. Dans une seconde étape, nous étudions le modèle initialement proposé par Gupta et Kumar, dans lequel chaque nœud veut transmettre des données vers une destination de son choix. Nous montrons que si le réseau occupe une surface bornée, et que la fonction d'atténuation est aussi bornée, le débit de donnée alloué à chaque nœud doit décroître de façon inversement proportionnelle au nombre de nœuds dans le réseau. Par contre, si la densité de nœuds est constante, et que la surface augmente linéairement avec le nombre de nœuds n , alors un débit égal et de l'ordre de $1/\sqrt{n}$ est réalisable pour chaque nœud. Ce résultat améliore la borne inférieure connue jusqu'ici d'un facteur $\sqrt{\log n}$.

La dernière partie de ce travail est consacrée à deux études portant sur le temps de latence de réseau de capteurs. Dans la première étude, le réseau doit détecter un intrus qui se déplace en ligne droite et à vitesse constante. Nous calculons une borne supérieure à la distribution du temps requis pour détecter l'intrus, sous l'hypothèse que l'intrus n'est détecté que quand il entre en contact avec la fraction de nœuds connectés du réseau. Dans la seconde étude, nous nous intéressons à un réseau de capteurs qui, pour économiser leur batterie, éteignent leur radio pour des périodes aléatoires, indépendamment les uns des autres. Dans ce réseau, un nœud émet un message d'alerte à tous les autres nœuds. Comme certains nœuds ont éteint leur radio, le message se propage jusqu'à ce qu'il ne puisse atteindre d'autres nœuds actifs; ensuite, il doit attendre qu'un nouveau nœud devienne actif pour continuer. Nous prouvons que dans un tel scénario, le rapport entre la distance parcourue par le message et le temps écoulé tend asymptotiquement vers une constante déterministe.

Acknowledgments

I want to thank first my advisor, Professor Patrick Thiran, for bringing me to this exciting research topic, and guiding me during my PhD. I'm deeply grateful to him for his availability and help, not only in research, but also in all other aspects of the PhD. I'm also very grateful to him for taking so much care of his students' personal career and promoting their results. Finally I want to thank him for giving me the great chance to teach a bit of his class, which I enjoyed much.

I also want to thank Martin Vetterli for giving great impulsion to the whole institute, and Martin Hasler for the very seminal meetings of the beginnings.

I would like to thank all the people in my lab for making it a friendly and lively working place, especially my close neighbors, Mathilde Durvy and Ruben Mertz for bearing my moods. I would like to thank specially Catherine Boutremans for the great support during the the first times. I'm also very thankful to the lab's staff, Angela Devenoge, Danielle Alvarez, Holy Cogliati, Jean-Pierre Dupertuis, Marc-André Lüthi and Philippe Chammartin for making everything work so smoothly.

A special thanks goes to my friends in Lausanne, *the Rock Family*, who made these years in Lausanne so exciting, and supported me in all critical moments.

Finally, my gratitude goes to my family for love, support and encouragement during all my studies.

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 1.1 | Motivations | 1 |
| 1.2 | Dissertation outline | 2 |
| 1.2.1 | Connectivity | 2 |
| 1.2.2 | Throughput | 3 |
| 1.2.3 | Latency | 4 |
| 1.3 | Related work | 5 |
| 1.4 | Contributions | 7 |
| | | |
| I | Connectivity | 9 |
| | | |
| 2 | The Boolean model | 11 |
| 2.1 | Model definition | 12 |
| 2.1.1 | Mathematical model | 12 |
| 2.1.2 | Specific model | 12 |
| 2.2 | Connectivity of a pure ad-hoc network | 13 |
| 2.2.1 | One-dimensional case: line | 13 |
| 2.2.2 | Two-dimensional case: plane | 15 |
| 2.2.3 | Two-dimensional case: strip | 16 |
| 2.3 | Number of paths and bottlenecks | 17 |
| 2.3.1 | Number of alternate paths between two nodes | 17 |
| 2.3.2 | Bottlenecks | 19 |
| 2.4 | Hybrid Networks | 22 |
| 2.4.1 | Nodes and base stations on a line | 24 |
| 2.4.2 | Nodes and base stations on a plane | 25 |
| 2.4.3 | Nodes and base stations on a strip | 26 |
| 2.5 | Non-homogeneous node distribution | 26 |
| 2.6 | Conclusion | 27 |
| | | |
| 3 | Impact of interferences | 31 |
| 3.1 | Model | 32 |
| 3.1.1 | A Bound on the Degree of the Nodes | 33 |
| 3.1.2 | Shot-Noise | 34 |
| 3.1.3 | Attenuation | 35 |
| 3.2 | Percolation | 35 |
| 3.2.1 | Existence of a percolation threshold for $\gamma = 0$ | 35 |
| 3.2.2 | Some observations on the graph with $\gamma > 0$ | 36 |

| | | |
|------------|---|------------|
| 3.2.3 | Percolation for non-zero values of γ | 37 |
| 3.3 | Asymptotic behavior | 47 |
| 3.3.1 | Scaling properties | 47 |
| 3.3.2 | Power law attenuation functions | 48 |
| 3.3.3 | Bounded attenuation function | 50 |
| 3.4 | A TDMA approach | 56 |
| 3.4.1 | Asymptotic behavior | 57 |
| 3.5 | Conclusion | 59 |
| II | Throughput | 61 |
| 4 | Connectivity vs Throughput | 63 |
| 4.1 | Partial vs full connectivity in the Boolean model | 63 |
| 4.2 | Taking interferences into account | 67 |
| 4.2.1 | Extended networks | 67 |
| 4.2.2 | Dense networks | 69 |
| 4.3 | Conclusion | 71 |
| 5 | Information theoretic connectivity | 73 |
| 5.1 | Network model | 74 |
| 5.2 | One dimensional network | 75 |
| 5.3 | Two dimensional network | 75 |
| 5.3.1 | Lower bound | 75 |
| 5.3.2 | Upper bound | 78 |
| 5.4 | Discussion | 79 |
| 5.5 | Conclusion | 80 |
| 6 | Throughput under uniform traffic matrix | 81 |
| 6.1 | Main results | 82 |
| 6.2 | Overview of the solution | 83 |
| 6.3 | Percolation results | 85 |
| 6.3.1 | Number of disjoint crossing paths | 87 |
| 6.4 | Capacity of the percolation cluster | 90 |
| 6.5 | Protocol | 93 |
| 6.6 | Dense Networks | 97 |
| 6.7 | Conclusion | 98 |
| III | Latency | 101 |
| 7 | Intruder detection | 103 |
| 7.1 | Model | 104 |
| 7.2 | Theorem formulation and insight of the proof | 104 |
| 7.3 | Construction and its independence properties | 105 |
| 7.3.1 | The size of finite clusters | 106 |
| 7.3.2 | Final computation | 108 |
| 7.4 | Extension to the continuous case | 109 |
| 7.4.1 | Mapping | 109 |
| 7.4.2 | Domination of the dependent model | 110 |

| | | |
|----------|---|------------|
| 7.5 | Conclusion | 111 |
| 8 | Uncoordinated power saving mechanism | 113 |
| 8.1 | Decentralized energy saving mechanism | 114 |
| 8.2 | Blinking Poisson Boolean model | 115 |
| 8.3 | Sensing coverage and transmission connectivity | 118 |
| 8.3.1 | Sensing area | 118 |
| 8.3.2 | Radio connectivity | 119 |
| 8.4 | Latency | 119 |
| 8.4.1 | Linear spreading of alarm messages | 119 |
| 8.4.2 | Duration of a transmission phase | 128 |
| 8.4.3 | Simulation studies | 129 |
| 8.5 | Conclusion | 134 |
| 9 | Conclusion | 135 |
| 9.1 | Future research | 136 |
| A | Appendices for Part I | 139 |
| A.1 | Derivation of the lower bound to the probability of connection in one-dimensional hybrid networks | 139 |
| A.2 | Bounded support attenuation function | 140 |
| A.3 | Invariance with respect to scaling | 143 |
| B | Appendices for Part II | 145 |
| B.1 | Shot-noise | 145 |
| C | Appendices for Part III | 147 |
| C.1 | Continuum growth model | 147 |
| | Notations | 149 |
| | Publications | 151 |
| | Curriculum Vitæ | 153 |
| | Bibliography | 154 |

Chapter 1

Introduction

1.1 Motivations

Multi-hop wireless networks are a promising technology when fixed, centralized infrastructure is not a good option. Currently, most wireless networks require base stations to operate; the most common examples are cellular telephony and IEEE 802.11 networks (in infrastructure mode). The idea of multi-hop networks is to exploit the possibility that the wireless devices (called hereafter *nodes*) can connect directly with one another, like walky-talkies. However, direct connection is only possible if the devices are close enough to each other. To achieve longer range communications, intermediate nodes can be used to *relay* the communications between the two end nodes. The possibility of using several consecutive relays in a single communication yields the name “multi-hop networks”.

Multi-hopping replaces therefore in theory the need for fixed infrastructure, since the network is only made of similar devices collaborating to provide connectivity, by acting as terminals and as relays simultaneously. Although it is quite easy to implement such networks on a small scale, things become much more complicated when the number of nodes is large. The first issue that comes to mind is connectivity: is it always possible to find appropriate relays between given source and destination in order to establish the communication? Or, are there disconnected nodes in the network? Furthermore, since nodes act as relays, they have to use their radio device not only for transmitting *their* data, but also the data from other nodes. The channel may thus become very busy, degrading the end-to-end throughput. The energy consumption of the nodes may increase dramatically because of relay traffic. The use of several relays implies also much greater end-to-end delays. If there are enough relays between two end nodes, one still needs a protocol that is able to *route* the data through the appropriate nodes to the destination. To do this, the routing protocol must be able to *identify* and *locate* the destination accurately. Finally, collaboration and fairness between nodes may also be an issue.

These different issues are more or less critical depending on the particular setting and the kind of application. For example, if the node density is very high, connectivity is less of a problem, whereas the throughput worsens as a result of the high number of interferers. In this work, we address the issues of

connectivity, throughput and in some particular scenarios latency. Our objective is to explain the trade-offs between these features, and the conditions required for good performance.

Since we address large scale networks, we mostly look at the network as a macroscopic object with properties that are not observable on a small scale. However, the network is still made of a multitude of small objects (the nodes) interacting locally (direct links between them). This setting is in fact very similar to the problems addressed in physics, where huge systems of particles are studied. Therefore, in this work, we use a tool that was originally developed by physicists, and that turns out to be very suitable for the study of multi-hop wireless networks: percolation theory.

1.2 Dissertation outline

1.2.1 Connectivity

In a first approach, we use the Poisson Boolean model to describe the network, as proposed initially in [GP60]. This model is very simple and implements the following two basic assumptions: the users are located randomly and independently of each other (Poisson point process) and they can connect to each other through a wireless link only if the distance between them is no more than their *connectivity range*.

This first model allows us to study fundamental connectivity properties of the network. In Chapter 2, we consider three geometries –one dimension, two dimensions and finite width infinite length strip– which model three of the most important settings for wireless networks. The second geometry is the most commonly used, and applies to most situations encountered in practice. The first geometry would apply for example to a network between cars on a highway. The third one would model users located in a strip-shaped area like a valley.

Under this model, the two-dimensional case exhibits a very interesting property: given a radio range and above a certain critical density of nodes, a giant cluster of (multi-hop) connected nodes appears. This phenomenon, known as *percolation*, allows connectivity along unbounded distances only via multiple hops. In one dimension and on a strip, this phenomenon does not occur, and multi-hopping is not an option for long range communication.

Multi-hopping can also be used in networks with infrastructure. In this case, nodes that are out of the coverage area of the base stations use the other nodes to relay their data towards the covered area. We look at the improvement provided by this technique, especially when pure multi-hopping is not efficient.

The major drawback of the Boolean model is that it does not allow interferences to be taken into account. When the number of nodes increases in the network, the wireless medium becomes more and more solicited. This competition for the channel may prevent successful transmissions, which would occur without the presence of interferences.

The physical justification for using the Boolean model and defining a “connectivity range” is that the signal strength decreases over distance, and that it must be still strong enough, compared to the background noise at the receiver. The required ratio between the signal and the background noise (SNR) thus defines the maximum possible distance between the transmitter and receiver. A

natural way to take interferences into account is to add the sum of the interfering signals coming from simultaneous transmissions to the background noise. We then adopt a different criterion for the existence of a link, which is that this new signal-to-interference-plus-noise-ratio (SINR) is above the threshold for successful decoding. Under the SINR model, we investigate in Chapter 3 the existence of a percolation phenomenon, and characterize the conditions under which it occurs. Since this model explicitly involves the signal power attenuation as a function of the distance, we also discuss how the choice of the model of power attenuation over the distance affects the connectivity results.

When the density of nodes in the network is high, interferences become stronger, since nodes are close to each other. We also investigate in the same chapter the asymptotic behavior of the SINR model when the node density increases to infinity. We find out that connectivity deteriorates, unless the interferences are moderated by an increasingly efficient CDMA system.

1.2.2 Throughput

Apart from connectivity, a fundamental feature of a wireless network is the rate at which it can transport data. When the bandwidth and emitting power of the nodes are limited, there is clearly a trade-off between connectivity and the throughput along links. For example, in 802.11, when two terminals are close enough, the data rate reaches its maximum. On the contrary, when the terminals are more distant, the protocol adjusts the rate to preserve connectivity. This kind of trade-off is intrinsic to wireless networks. In comparison to single-hop networks, multi-hop networks can also use the available bandwidth to relay packets. As a result, the end-to-end throughput may be limited not only by the length of the links (noise limited case) as in single hop networks, but also by the competition with other transmissions (interference limited case).

We study the noise-limited case in Chapter 5, where we compute bounds to the fraction of nodes that can be connected at a given rate provided that there is only one transmission at a time. This is a very natural extension of Chapter 2 that essentially treats noise-limited networks as well. The difference is that in Chapter 5, we adopt a more throughput-oriented connectivity model: we declare two nodes connected to each other if it is possible to relay data from one to the other at some minimum given rate. The interest of this model is that it requires no assumption on the transmission scheme, and is thus of a truly information theoretic nature. In the Boolean model, on the contrary, we implicitly assume point-to-point communication since we investigate the existence of paths between nodes which are essentially a sequence of point-to-point connections. The model in Chapter 5 is much more general since it allows for arbitrarily complicated relaying schemes as, for example, the use of several paths simultaneously.

To measure how many simultaneous data flows can be sustained by the network, the concept of *transport capacity* was introduced by Gupta and Kumar [GK00]. As a first result, they showed that when the node density increases, the total transport capacity of the network only increases like the square root of the number of nodes. This implies a decreasing available transport capacity per node. However, this result starts from the hypothesis that the strength of the signals decreases like the negative power of the distance (power law). We will see in Chapter 4 that power law attenuation functions gives very particular

symmetry properties with respect to scaling to the SINR model. Thanks to this symmetry, Gupta and Kumar's upper bound can be directly applied to the case where the network increases in size with a constant node density (*extended* networks, as opposed to *dense* networks considered originally).

However, from a physical point of view, the attenuation function cannot take values greater than one. Otherwise, the received power would be greater than the emitted power. In Chapter 4, we thus investigate the transport capacity of dense networks when the node density increases, assuming a bounded attenuation function (for example truncated power law). We show that the result differs greatly, and that the transport capacity is bounded from above by a constant under this new assumption.

The transport capacity is a suitable tool for computing upper bounds on the actual throughput experienced in different traffic scenarios. However, to obtain lower bounds, one needs to provide an explicit scheme for transporting packets from their source to their destination. In their paper, Gupta and Kumar presented a scheme that achieves a rate of order $1/\sqrt{n \log n}$ for each node in a random network of n nodes, when each node wants to send data to a randomly chosen destination (uniform traffic matrix). This result therefore achieves a total transport capacity of order $\sqrt{n}/\sqrt{\log n}$. There is thus a gap between the upper and the lower bound. In Chapter 6, we come up with a scheme that achieves a rate of order $1/\sqrt{n}$ per node under the same assumptions, thus closing the gap between previous bounds. The idea of this scheme is to adjust the length of the links in the network so that it becomes super-critical, but not fully connected. Inside the giant cluster, we show that one can find enough routes to carry all traffic with the desired rate. The fact that we do not require the network to be fully connected in a first step allows to remove the $\sqrt{\log n}$ factor from the Gupta and Kumar result, as already pointed out in [BBM03]. Next, in a separate time slot, we use a draining scheme to carry packets from the disconnected nodes to the giant cluster. In agreement with the results in Chapter 5, we show that the rate of this draining is higher than $1/\sqrt{n \log n}$, and thus *all* nodes can be served with the same throughput.

1.2.3 Latency

The last part of this dissertation is dedicated to a more specific kind of wireless network. We consider networks that are designed for sensing an event, and spreading a message describing this event. We focus on the the latency of such networks, both in detecting the event and transmitting the message.

In chapter 7, we consider a network of sensors across which an intruder moves with constant speed along a straight line. We use again the Boolean model, and assume that the nodes can detect intruders within a range r . The sensed area is thus a collection of disks, centered on the points of a Poisson process. Furthermore, we make the simplifying assumption that the connectivity range is equal to twice the sensing range so that two nodes are connected whenever their sensing disks overlap. In this chapter, we are interested in the time delay until the intruder gets detected. The distribution of the time before it enters the range of a sensor is actually well known (see e.g. [SKM95]). However, since we assumed that the sensors transmit the alarm in a multi-hop fashion, it is not enough to hit a random sensor to be successfully detected by the system. The sensor that detects the intruder must also be able to send the alarm message to

some remote sink. In other words, only sensors that belong to the giant cluster participate efficiently to the intruder detection. The appropriate question is to know how long the intruder will travel before it gets detected, and is therefore a new problem in percolation theory. We show that the probability that the intruder does not get detected within the first t seconds is upper bounded by an expression that decreases exponentially with t .

In Chapter 8, we look at a slightly different scenario. We consider again the Boolean model, and assume that the network is well connected. In addition, we add an energy saving mechanism in which nodes sleep for random periods in a totally desynchronized fashion. This mechanism of course harms the connectivity of the network, but is very simple to implement since it is purely local and random. In this scenario, a source node wants to broadcast a message to all other nodes (an alarm for instance). When the source starts to broadcast, the message is relayed to all currently awake and connected nodes. The spread of the message stops until some new nodes become active, allowing the message to progress further into the network.

This mechanism can lead to two distinct situations. First, if enough nodes are awake, the message can reach an unbounded number of nodes and therefore never stops propagating. In the second case, the network made of active nodes is disconnected, and the message propagates step by step when new nodes become active. We investigate the propagation speed in this latter scenario, and compute the asymptotic shape of the region of the network where the message has already arrived. We prove that the distance traveled by the message grows linearly with respect to the elapsed time. The propagation speed is thus asymptotically constant, and can be predicted from the model parameters.

1.3 Related work

Connectivity has received quite a lot of attention in the previous decade already in the context of packet radio networks, and has gained renewed interest recently in the context of ad-hoc and sensor networks. Many results apply to the full connectivity of a network made of a finite number of nodes. However, there are also numerous papers on the connectivity of infinite networks.

In extended networks, the connectivity problem is related to percolation theory, which is to find the probability that a node belongs to an infinite cluster of nodes. The most popular model is the Poisson Boolean model, where node locations follow a Poisson distribution, as in our work, but where only the distance between two nodes determines the existence of a link between them. This model is a particular case of the STIRG for $\gamma = 0$, and was introduced by Gilbert [Gil61], who started the field of continuum percolation. Despite the apparent simplicity of the Poisson Boolean model, the exact value of the critical density λ^* at which the transition occurs is still an open problem. Some bounds on λ^* have been obtained analytically in [Gil61], [MR96], [PPT89], and numerically by many others [QTZ00].

Percolation of a clustered wireless network in which the users (clients), who are distributed according to a Poisson process, are all covered by base stations that can connect to each other by a wireless link is studied in [BBFM03]. This model boils down to the Poisson Boolean Model if one base station is placed at each client.

The assumption that two nodes are connected to each other if and only if their distance is less than or equal to some predefined radius is certainly too simplistic. In reality, propagation in a real environment yields much more complex and irregular coverage shapes than circles. Interestingly, it appears that this irregularity makes the network percolate at lower values of the density than the Boolean model [BBCF03]. The latter appears thus as a conservative model for noisy channels. When interferences are taken into account however, this is no longer true: for the same patterns, the Boolean model may have an infinite component of the connectivity graph, whereas the STIRG could have no infinite component, or even no connectivity at all [DBT03]. This physical model was analyzed in [BB01] in the infinite plane case under Poisson assumptions within the context of CDMA networks. The level sets defining the contours around a node where the Signal to Interference has the same value can indeed have a very convoluted shape.

In dense networks, the full connectivity of the graph has been shown to hold with probability 1 when the distance r below which nodes can connect decreases at a rate slower than $\sqrt{\log n/n}$, with the number of nodes n tending to ∞ in [GK98].

A similar problem, where the power of a node is adjusted so that it can have a required number of neighbors, is studied in [XK04b]. When the node locations are restricted to be at the vertices of a grid instead of being scattered on the entire area as in the previous examples, results on the asymptotic connectivity of the network have been obtained in [SSS03].

The capacity of wireless ad hoc networks belongs to the very difficult field of multi-users information theory. In their seminal paper, Gupta and Kumar [GK00] proved that the aggregate transport capacity, defined as the sum of all the link capacities multiplied by the distances between the source and destination, is of the order of \sqrt{n} bit-meter per second for the case of dense networks with point-to-point codes, a uniform traffic matrix and a power law attenuation with a path loss exponent $\alpha > 2$. This means that the rate per user decreases as $1/\sqrt{n}$ and eventually reaches 0 for $n \rightarrow \infty$. These results rely on the assumption that the signal to noise and interference ratio exceeds some threshold β . This assumption, which may make sense in practice but is not truly of an information theoretic nature, is not needed, and Xie and Kumar [XK04a] have proved that for attenuation functions of the type e^{-ax}/x^α , with either $a > 0$ or $\alpha > 3$, the aggregate transport capacity is of the order of $n/\sqrt{\log n}$ bit-meter per second under the assumption that the minimum distance between any pair of nodes is bounded below by some constant. This restriction forces the domain size s to grow with n , and hence that capacity result applies to extended networks. These results were extended to power law attenuation functions with $\alpha > 2$ in [LT04] for extended networks without the minimum distance restriction. They encompass therefore the case where nodes are distributed as a Poisson process.

These general results are more optimistic if one considers some particular scenarios. For example, if there is only one active source/destination pair, while the $n - 2$ other nodes act as possible relays, the previous results would predict a constant capacity, independent of n , whereas allowing arbitrarily complex coding, this capacity is of the order of $\log n$ [GV02]. Another example is the increased capacity obtained when the nodes are mobile: with one hop relaying, the total transport capacity can be of the order of n [GT01]. Finally, the benefit of hybrid multi-hop cellular networks is investigated for the so-called protocol

model in [LLT03].

Since energy saving is central to sensor networks [Eph02], it has received a considerable amount of attention, driving routing algorithms (see e.g. [CT00, SWR98]), scheduling (see e.g. [PUBeG01, eGNP⁺02]), data collection and aggregation (see e.g. [KEW02]) and MAC (see e.g. [SWR98, YH03]). Scheduling strategies trading off energy saving and latency are devised in [YKP04], while a Markovian model exploring the performance of a wireless network with on/off periods is described in [CG04] under the assumption that there is a path from any sensor to the sink.

1.4 Contributions

The main contribution of this work is to introduce percolation theory as an efficient tool for studying large multi-hop wireless networks. The chapters of this thesis give several examples of application of percolation theory leading to novel results and new perspectives in the study of scaling properties of such networks.

The detailed list of the contributions is as follows:

- In one-dimensional networks, we compute the probability that two nodes are connected through multiple hops under the Boolean model.
- In two-dimensions, we study the number of disjoint paths in the Boolean model. We derive a simplified version of the connectivity graph that contains fewer vertices, but exactly the same connectivity (in terms of disjoint paths) as the original graph.
- We study the connectivity under the Boolean model when we add fixed, wired base stations in the network. We come up with analytical bounds in the one-dimensional case.
- We study a third configuration, between the one-dimensional and two-dimensional case, where nodes are located on a strip of infinite length and finite width.
- We introduce a novel percolation model, based on the signal-to-interference-plus-noise-ratio (SINR), that is more realistic than the Boolean model. We show that the same percolation phenomenon occurs in this new model.
- We study in detail the domain of parameters of this new model for which percolation occurs. We characterize the asymptotic behavior of the percolation threshold when the node density tends to infinity.
- We show that, in this model, the percolation threshold can always be reached using a trivial time division multiple access (TDMA) scheme.
- We show that the assumption of power-law attenuation confers very specific scaling properties to the network. If the attenuation function does not diverge at the origin, the results for high density networks are dramatically different.

- We show that if we assume point-to-point communication and a bounded attenuation function, the total transport capacity does not increase with the node density (in contrast to the power-law attenuation case, where it increases like the square root of the number of nodes).
- We study an information theoretic model of connectivity, where nodes are considered connected if it is possible to transmit data from one to the other at a certain prescribed rate (possibly using all other nodes as relays). We show that under this model, a constant fraction of nodes are connected (independently of the network size), and we give bounds on this fraction.
- We present a new scheme that achieves a per node throughput of order $1/\sqrt{n}$ in the Gupta and Kumar setting with n randomly placed nodes.
- We show that in the Boolean model, the distance from an arbitrary point to the giant cluster has a distribution that is bounded from above by an exponential distribution. This result is applied to an intruder detection problem.
- Under a desynchronized energy saving mechanism, we show that a broadcast message propagates with asymptotically constant and deterministic speed.

Part I

Connectivity

Chapter 2

The Boolean model

In this chapter, we study the connectivity of wireless multi-hop networks in the case where there are no interferences. The idea behind this assumption is to identify all *possible* links in the network. Such an assumption applies to networks where there is very little traffic (for instance signalization messages), and where node *reachability* is more important than throughput between pairs of nodes. A more general discussion about the duality between connectivity and throughput will be made in Chapter 4.

As is it an essential feature, connectivity has received quite a lot of attention in the previous decades already, in the context of packet radio networks, and has gained renewed interest recently in the context of ad-hoc and sensor networks. Most results apply to the full connectivity of a network made of a finite number of nodes. A recursive formula giving the average number of hops between two connected nodes is found in [CR89], whereas the probability that a given number of nodes on a finite interval are all connected is computed in [SB02]. In the two-dimensional setting, relations between k -connectivity (the property that the graph has a minimal cutset equal to $k \geq 1$) and the node degree are studied in [Bet02], whereas this problem is addressed when the transmission powers of the nodes are different in [XK04b]. In this chapter, we assume that the number of nodes is not fixed nor on a bounded area, but that they are given as points of a Poisson process over \mathbb{R} , \mathbb{R}^2 , or a strip $\mathbb{R} \times [0, d]$. We do not make assumptions on its intensity, so that our results also apply to low density areas.

Since the number of nodes is not bounded, some of them will be disconnected. The problem is then related to percolation theory, which is to find the probability that a node belongs an infinite cluster of nodes. Since the pioneering work of Gilbert [Gil61], which started the field of continuum percolation, the exact value of this probability is still an open problem. Some bounds on the critical intensity λ^* below which it is zero have been obtained analytically in [Gil61], [MR96], [PPT89] for the Boolean Poisson Model, and numerically by many others [QTZ00]. Percolation of a clustered wireless network, in which the users (clients), who are distributed according to a Poisson process, are all covered by base stations that can connect to each other by a wireless link, is studied in [BBFM03]. This model reduces to the Poisson Boolean Model if one base station is placed at each client.

In the literature, connectivity of multi-hop wireless networks is often understood as the *full connectivity* of all its nodes [GK98, Bet02, SB03, XK04b]. We

will discuss in detail in Section 4.1 the differences between the approach in this chapter and asymptotic full connectivity.

2.1 Model definition

2.1.1 Mathematical model

We start with a stationary Poisson point process in \mathbb{R}^d of intensity λ . The points of this process are denoted by $\{X_i\}$, and called *germs*. On each germ, we center a random closed set, called *grain*, independent of the point process and the other grains (the grains are said to be i.i.d.). In this work, we only consider $d = 1, 2$ and disk shaped grains. Then the *Poisson Boolean model* $\mathcal{B}(\lambda, R)$ is just a union of randomly scattered disks (see Figure 2.1), i.e., the coverage process defines the occupied region

$$\mathcal{B}(\lambda, R) \doteq \bigcup_i B(X_i, R_i),$$

where $B(X_i, R_i)$ is the disk centered at X_i having radius R_i , and the radii R_i are i.i.d., independent of the point process $\{X_i\}$, and distributed as some reference random variable R .




Figure 2.1: Poisson Boolean model in \mathbb{R}^2 . Sub-critical intensity on the left and super-critical intensity on the right.

The occupied region can be divided into disjoint *clusters* which are formed by the overlapping disks. One way to measure the global connectivity is the size of the largest cluster. Let us denote by $W(A)$, $A \subset \mathbb{R}^d$, the union of all the occupied components that intersect A .

2.1.2 Specific model

In this chapter, we will assume that two nodes can communicate through a wireless link if the distance between them is less than the *range* r of their radio device. Furthermore, we will assume all nodes identical.

This situation can be modeled by a Poisson Boolean model, with fixed ball radius $R \equiv r_b$, denoted by $\mathcal{B}(\lambda, r_b)$, where $r_b = r/2$. This corresponds to assuming deterministic and isotropic wave propagation in space. r is thus the maximal range allowed by power constraints. Although these assumptions are

very simplistic and unrealistic, the results presented below are still qualitatively correct, since most of the percolation results we use still hold in the more general model, where the ball radii are random. Moreover, it is shown in [FBC⁺05] that randomness in the shape of the grains improves connectivity. It is even conjectured that in two dimensions, disks are the worst shape for percolation. Therefore, we consider that fixed range wireless links are a very conservative (and even pessimistic) assumption.

One can associate with the Boolean model a *random graph* as follows: we assign a vertex to each ball of $\mathcal{B}(\lambda, r_b)$, and put an edge between two vertices if the ball overlap (see Figure 2.2). We denote this graph by $\mathcal{G}(\lambda, r) = \mathcal{G}(\lambda, 2r_b)$.




Figure 2.2: The Boolean model (left) and the associated graph (right).

Scaling property

An interesting property of this model is that it can be rescaled, so that for any $a > 0$,

$$\mathcal{B}(\lambda, ar_b) = \mathcal{B}(a^d \lambda, r_b),$$

where d is the dimension. The same is true of course for the associated graph, so that $\mathcal{G}(\lambda, ar) = \mathcal{G}(a^d \lambda, r)$. This property implies that the model has actually only one degree of freedom. For convenience, one often uses the *average number of neighbors* (or *average node degree*, in the context of the associated graph) $N = \lambda \pi_d r^d = 2^d \lambda \pi_d r_b^d$ as a unique parameter, π_d being the volume of the unit ball in d dimensions.

2.2 Connectivity of a pure ad-hoc network

We begin with the connectivity of the Poisson Boolean model. More specifically, our goal is to evaluate the probability $P_c(x)$ that two arbitrary nodes i and j , whose Euclidean distance in \mathbb{R}^d is denoted by $d(X_i, X_j) = x$, are connected to each other, as a function of λ , r and x .

2.2.1 One-dimensional case: line

We begin with the one-dimensional case, where two nodes at a distance x of each other are connected if the entire interval between them is occupied. If there is a vacant region, the information cannot be relayed from a node to the

other. A hole appears if the interval between two consecutive nodes is longer than r . As $\lambda < \infty$, the probability for a vacant interval to occur between two consecutive nodes is strictly positive, whatever the value of λ and r . Since there is an infinite number of intervals, and since their lengths are independent, the probability of having no hole is zero. Hence $\lim_{x \rightarrow \infty} P_c(x) = 0$. The same result can be obtained by considering the network as an $M/D/\infty$ queue, in which the arrival process is the actual Poisson point process, and the "service" process has fixed duration r . The occurrence of a hole is equivalent to a visit to the empty-state of the queue. Since $\lambda < \infty$, this state will be visited almost surely after a finite time.

Related results include the computation that a given number of nodes on a finite interval are all connected to each other [SBV01]. A recursive formula giving the average number of hops between two connected nodes is given in [CR89].

Here we compute the probability distribution $P_c(x)$ for an infinite domain. Clearly, $P_c(x) = 1$ if $0 \leq x < r$. When $x \geq r$, $P_c(x)$ is the probability that there is another node between X_i and X_j , located at a distance ξ from B , with $0 \leq \xi < r$, and that this node is connected to X_i . We use this argument to condition $P_c(x)$ recursively on $P_c(x - \xi)$, with $0 \leq \xi < r$, and we establish the following result.

Theorem 2.1 *The probability $P_c(x)$ that two nodes distant of x space units are connected is*

$$P_c(x) = \begin{cases} 1 & \text{if } 0 \leq x < r \\ \sum_{i=0}^{\lfloor x/r \rfloor} \frac{(-\lambda e^{-\lambda r(x-ir)})^i}{i!} & \\ -e^{-\lambda r} \sum_{i=0}^{\lfloor x/r \rfloor - 1} \frac{(-\lambda e^{-\lambda r(x-(i+1)r)})^i}{i!} & \text{if } x \geq r \end{cases} \quad (2.1)$$

with $\lfloor x/r \rfloor$ denoting the largest integer smaller than or equal to x/r .

Proof: If $x \geq r$, $P_c(x)$ is the probability that (i) the next node towards Node i is located at a distance ξ from Node j , with $0 \leq \xi < r$, and that (ii) this new node is connected to i . This second event occurs with probability $P_c(x - \xi)$. The interval separating two consecutive nodes has an exponentially distributed length, whose pdf is thus $\lambda e^{-\lambda \xi}$ for $0 \leq \xi < x$. Therefore

$$\begin{aligned} P_c(x) &= \int_0^r P_c(x - \xi) \lambda e^{-\lambda \xi} d\xi \\ &= \lambda e^{-\lambda x} \int_{x-r}^x P_c(y) e^{\lambda y} dy. \end{aligned} \quad (2.2)$$

Taking the derivative of this expression with respect to x , we obtain

$$\frac{dP_c}{dx}(x) = -\lambda e^{-\lambda r} P_c(x - r) \quad (2.3)$$

for $x \geq r$. This is a first order linear delay differential equation, with initial condition defined on the interval $[0, r]$. We know that $P_c(x) = 1$ for $0 \leq x < r$. At $x = r$ however, $P_c(x)$ has a discontinuity: indeed, solving (2.2) for $x = r$, we

find $P_c(r) = 1 - e^{-\lambda r}$. Integrating (2.3) with this initial condition on $[0, r]$, we obtain, after some manipulations,

$$P_{c0}(x) = \sum_{i=0}^n (-1)^i \frac{(\lambda e^{-\lambda r} (x - ir))^i}{i!} - e^{-\lambda r} \sum_{i=0}^{n-1} \frac{(-\lambda e^{-\lambda r} (x - (i+1)r))^i}{i!}$$

for $nr \leq x < (n+1)r$ and $n \in \mathbb{N}$. \square

The following bounds on $P_c(x)$ will be needed in Section 2.4.

Theorem 2.2 *If $x \geq 2r$,*

$$(1 - e^{-\lambda r}) e^{-\lambda(x-2r)e^{-\lambda r}} - \lambda e^{-\lambda r} \leq P_c(x) \leq (1 - e^{-\lambda r}) e^{-\lambda(x-r)e^{-\lambda r}}. \quad (2.4)$$

Proof: Since $P_c(x)$ is a decreasing function of x , $P_c(x) \leq P_c(x-r)$ for any $x \geq r$. Inserting this inequality in (2.3), we get the following differential inequality

$$\frac{dP_c}{dx}(x) \leq -\lambda e^{-\lambda r} P_c(x) \quad (2.5)$$

for $x \geq r$. Since $P_c(r) = 1 - e^{-\lambda r}$, we have from the theory of Gronwall-type differential inequalities (see e.g [LLM89, Chapter 1]) that the solution of (2.3) is upper bounded by the solution of the linear ode (2.5), (with the inequality sign replaced by an equality sign), which is

$$P_c(x) \leq (1 - e^{-\lambda r}) e^{-\lambda(x-r)e^{-\lambda r}} \quad (2.6)$$

for all $x \geq r$. This is the upper bound in (2.4).

To obtain a lower bound, let us note that (2.6) provides an upper bound on $P_c(x-r)$, by replacing x by $(x-r)$ in this expression. Inserting it in (2.3), we get

$$\frac{dP_c}{dx}(x) \geq -\lambda e^{-\lambda r} (1 - e^{-\lambda r}) e^{-\lambda(x-2r)e^{-\lambda r}}$$

which we integrate for $x \geq 2r$ to obtain

$$P_c(x) \geq -\lambda r e^{-\lambda r} + (1 - e^{-\lambda r}) e^{-\lambda(x-2r)e^{-\lambda r}}$$

using the fact that $P_c(2r) = 1 - e^{-\lambda r} - \lambda r e^{-\lambda r}$ because of (2.1). This is the lower bound in (2.4). \square

We can thus conclude that in one dimension, the network is almost surely divided into an infinite number of bounded clusters, between which no communication is possible. A large-scale ad hoc network will therefore not work in one dimension with constant range.

2.2.2 Two-dimensional case: plane

The picture is very different in two dimensions. For example, the existence of an unbounded cluster does not imply full connectivity. The following result from percolation theory is our starting point.

Theorem 2.3 [MR96, pages 45–68] Consider a Poisson Boolean model $\mathcal{B}(\lambda, r_b)$ in \mathbb{R}^2 with fixed ball radius r_b . There exists a critical density $\lambda^* > 0$ such that

- in the sub-critical case, defined by $\lambda < \lambda^*$, all clusters are bounded almost surely (a.s.)
- in the super-critical case, defined by $\lambda > \lambda^*$, there exists a unique unbounded cluster a.s.

In terms of networking, this means that the sub-critical case is similar to the one-dimensional case, where the network is partitioned into an infinite number of bounded clusters. In the super-critical case however, the result is much more encouraging, because of the existence of an unbounded cluster. In the sequel, we will denote by C_∞ the unbounded cluster in $\mathcal{B}(\lambda, r_b)$.

In the super-critical phase, one can divide the nodes into two categories: those belonging to the unbounded cluster, and the others. We will denote by \mathcal{U} the set of the nodes in the first category. These nodes can communicate with nodes located arbitrarily far away, whereas the others are restricted to a finite area. Thus, the quality of the connectivity is related to the fraction θ of nodes belonging to the first category. Note that θ is most frequently defined as the probability of an arbitrary node to belong to the unbounded cluster, and is called *percolation probability*:

$$\theta \doteq \mathbb{P}(|W(\{0\})| = \infty).$$

Because of the spatial invariance of the Poisson process, the position of this node can be taken, without loss of generality, at the origin.

One can easily derive a lower bound to $P_c(x)$ using the percolation probability, by evaluating the probability that both end-nodes i and j belong to \mathcal{U} .

$$\begin{aligned} P_c(x) &\geq \mathbb{P}(i \in \mathcal{U} \text{ and } j \in \mathcal{U}) \\ &\geq \mathbb{P}(i \in \mathcal{U}) \mathbb{P}(j \in \mathcal{U}) \\ &= \theta^2, \end{aligned} \tag{2.7}$$

where Inequality (2.7) follows from the FKG inequality (see e.g. [MR96, Theorem 2.2]). This lower bound becomes tight when x tends to infinity.

In the sequel, we will often refer to θ as a function of the *average node degree* $\theta(\lambda\pi r^2)$. To date, there is unfortunately no explicit expression of $\theta(\lambda\pi r^2)$, nor of λ^* . Bounds on λ^* have been obtained in [Gil61, PPT89, MR96], whereas $\theta(\lambda\pi r^2)$ can be evaluated by simulation, as shown in Figure 2.3. It is however known that $\theta(\lambda\pi r^2)$ is continuous above the critical threshold, and tends to one when $\lambda\pi r^2$ tends to infinity. This latter property confirms the intuition that when nodes have a high degree in average, then most of the nodes are connected.

2.2.3 Two-dimensional case: strip

The last geometry we consider is a strip of infinite length and of fixed, finite width d . Let us assume that the strip is horizontal, and that the node density is λ . By projecting all nodes onto the x -axis, we obtain a one-dimensional Poisson point process of finite intensity λd . As the distances are smaller in the projected




Figure 2.3: Percolation probability θ in two dimensions as a function of the average node degree $\lambda\pi r^2$. The critical density point is reached when nodes have roughly 4.5 neighbors in average. The presence of a tail below the critical threshold is due to the finiteness of the domain, a simulation artifact.

space, for a constant ball radius, the existing connections are preserved, while others are created. Since percolation never occurs in one dimension, we can conclude that the strip never percolates.

The existence of an unbounded cluster requires thus the domain to be infinite in both dimensions. Therefore, $P_c(x)$ tends to zero when x tends to infinity. However, it is intuitively clear that the behavior of $P_c(x)$ for small distances (typically when $x \ll d$) must be similar to the two-dimensional case.

2.3 Number of paths and bottlenecks

In the previous section, we discussed the existence of at least one path between two nodes. In order to study the reliability of such a network, and to identify possible bottlenecks, we now examine the number of disjoint paths, which do not share any link with each other, in the super-critical phase. This amounts to compute the size of the minimal cutset in the associated graph, between both nodes.

2.3.1 Number of alternate paths between two nodes

Intuitively, the more distant two nodes, the lower the size of the minimal cutset in the associated graph $G(\lambda, r)$ between them. We consider first one of these two nodes, call it i . Suppose that $i \in \mathcal{U}$. We denote by $N_i(p)$ the size of the minimal cutset within a circle of radius p centered on X_i , that separates Node i from the nodes located outside of this circle.

More formally, denote by $\mathcal{C}_i(p)$ the set of the closed curves surrounding a node i , which are included in the disk of radius p centered on X_i . For each curve $\mathcal{C} \in \mathcal{C}_i(p)$, denote by $\mathcal{E}(\mathcal{C})$ the set of the edges of $\mathcal{G}(\lambda, r)$ intersected by \mathcal{C} . Function $N_i(p)$ associated with each node $i \in \mathcal{U}$ is therefore given by

$$N_i(p) = \min_{\mathcal{C} \in \mathcal{C}_i(p)} \text{card}(\mathcal{E}(\mathcal{C})) \quad (2.8)$$

This function is clearly decreasing in p because $\mathcal{C}_i(p_1) \subseteq \mathcal{C}_i(p_2)$ if $p_1 < p_2$. Furthermore, $N_i(p) > 0$, because $i \in \mathcal{U}$. Therefore, $N_i(p)$ admits a limit when $p \rightarrow \infty$, which we denote N_i^∞ :

$$\lim_{p \rightarrow \infty} N_i(p) = \inf_{p > 0} N_i(p) \doteq N_i^\infty. \quad (2.9)$$

This number can be seen as the size of the minimal cutset separating Node i from what we call the “remainder of the network”, i.e. nodes located far away from X_i .

Figure 2.4 shows typical values of $N_i(p)$ (scaled to the number of immediate neighbors of Node i , which is clearly equal to $\lim_{p \rightarrow 0} N_i(p)$).




Figure 2.4: Average ratio between $N_i(p)$ and the number of neighbors of Node i for sub- and super-critical densities. For sub-critical densities, $N_i(p)$ goes to zero, because X_i belongs a.s. to a bounded cluster. For super-critical densities, $N_i(p)$ tends to a limit N_i^∞ noticeably smaller than $\lim_{p \rightarrow 0} N_i(p)$.

Thanks to this number, we can now compute the size of the minimal cutset between two nodes $i, j \in \mathcal{U}$. But we need first the following result.

Lemma 2.1 *In the super-critical case, any unbounded¹ curve \mathcal{C} crosses an infinite number of edges of the associated graph.*

Proof: Consider a portion \mathcal{D} of the curve \mathcal{C} that does not cross any edge. Denote by $\mathcal{W} = \mathcal{B}(\lambda, r/2)$ the occupied region, and by $\mathcal{V} = \mathbb{R}^2 \setminus \mathcal{B}(\lambda, r/2)$ the vacant region of \mathbb{R}^2 . Take a point \mathbf{x} on \mathcal{D} . Either $\mathbf{x} \in \mathcal{V}$, or $\mathbf{x} \in \mathcal{W}$. If $\mathbf{x} \in \mathcal{W}$, it can be included in the ball of a node placed on either side of the curve, but not on both, as otherwise, there would be an edge that crosses the curve. Let d be the distance between this node and \mathbf{x} . We are then sure that there is a vacant region on the other side of the curve, at a distance $r/2 - d$ from \mathbf{x} . Therefore, there is always a vacant region at a distance smaller than $r/2$ from any point of the curve.

¹A curve $\mathcal{C} \subset \mathbb{R}^2$ is *unbounded* if for any two points \mathbf{x}, \mathbf{y} on the curve, $\sup_{\mathbf{x}, \mathbf{y} \in \mathcal{C}} d(\mathbf{x}, \mathbf{y}) = \infty$ where $d(\mathbf{x}, \mathbf{y})$ is the Euclidean distance in \mathbb{R}^2 between \mathbf{x} and \mathbf{y} .

One can thus draw a new curve $\tilde{\mathbf{D}}$ such that for any point \mathbf{y} of $\tilde{\mathbf{D}}$, $d(\mathbf{y}, \mathbf{D}) \leq r/2$ and \mathbf{y} is in a vacant region. $\tilde{\mathbf{D}}$ is thus included in a single vacant region. However, there is no unbounded vacant region in the super-critical case [MR96, pages 108–116]. Hence, $\tilde{\mathbf{D}}$ must be bounded. As every point of \mathbf{D} is at a finite distance of $\tilde{\mathbf{D}}$, \mathbf{D} must also be bounded.

One can then conclude that the portion of curve between two crossings must be bounded. As the curve \mathbf{C} is unbounded, there are infinitely many crossings. \square

We can now compute the following lower bound on the size of the minimal cutset, and thus the number of alternate paths, between two nodes i and j belonging to the unbounded cluster.

Theorem 2.4 *The number of alternate paths between two nodes i and j is lower-bounded by $\min\{N_i^\infty, N_j^\infty\}$.*

Proof: We can assume without loss of generality that $N_i^\infty \leq N_j^\infty$. Suppose that there are less than N_i^∞ paths from Node i to Node j . Then there exists a curve dividing \mathbb{R}^2 into two parts, one including X_i and the other including X_j , and that intersects less than N_i^∞ edges. This curve is either closed and surrounds X_i , either closed and surrounds X_j , or unbounded. The two first cases are excluded by assumption. The third possibility is excluded by Lemma 2.1. \square

The lower bound in Theorem 2.4 is tight. Indeed, since $N_i(p)$ is an integer-valued function, the infimum in (2.9) is reached for a finite p , which we denote by L_i . In other words,

$$L_i = \inf\{p > 0 \text{ such that } N_i(p) = N_i^\infty\}. \quad (2.10)$$

If the distance between X_i and X_j is larger than $\max\{L_i, L_j\}$, then the number of alternate paths is exactly equal to $\min\{N_i^\infty, N_j^\infty\}$. It follows that the number of paths between two nodes is limited by two local values (N_i^∞ and N_j^∞), which depend exclusively on their respective neighborhoods.

2.3.2 Bottlenecks

In this section, we analyze the structure of the graph $\mathcal{G}(\lambda, r)$. In particular, we would like to identify the possible bottlenecks that can occur in the network.

Parameter L_i is important for locating hot spots and bottlenecks. Indeed, it is the smallest distance from Node i at which the minimal cutset (which has size N_i^∞) with “the remainder of the network” is reached. A small value of L_i means that the number of alternate paths between Node i and a node located far away is approximately equal to the number of paths between Node i and its close neighbors. As a result, the number of paths is simply limited by the number of immediate neighbors of the two end nodes. On the contrary, if L_i is large, then long-distance communications are constrained by bottlenecks located far way from the end nodes.

As we can see on Figure 2.5, L_i is on average larger than zero. In other words, N_i^∞ is in average less than the number of immediate neighbors of Node i (what we had already observed on Figure 2.4). Nevertheless, L_i tends to zero when the density λ increases, for a fixed r . This means that for high densities, the number of paths is simply limited by the number of immediate neighbors of end nodes, and that “hot spots” do not occur, as already pointed out in [GK00].




Figure 2.5: $\mathbb{E}(L_i)$ as a function of λ . We see that L_i goes to zero when the density increases, meaning that N_i^∞ tends to the number of neighbors of Node i

The case where λ is just above the critical density is the situation where bottlenecks appear. The network appears then as a set of “islands”, inside which nodes are well connected. Islands are however connected by a few links only, those forming the minimal cutsets. Nodes belonging to different islands are thus connected by a number of alternate paths given by Theorem 2.4. These links are the bottlenecks of the network.

The remainder of this section is devoted to a formal definition of the concept of island. We begin first by that, more general, of a domain of a node.

Definition 2.1 For each node i , consider the closed curve $\mathcal{C}_i \in \mathcal{C}_i(L_i)$ with $\text{card}(\mathcal{E}(\mathcal{C}_i)) = N_i^\infty$ that surrounds the minimum number of nodes. We call domain of Node i the set \mathcal{S}_i of nodes surrounded by \mathcal{C}_i .

For all nodes j in the domain \mathcal{S}_i of Node i we have $N_j^\infty \leq N_i^\infty$, because \mathcal{C}_i surrounds also X_j . In other words, all the elements of \mathcal{S}_i share the same N_i^∞ paths to the “remainder of the network”.

Domain \mathcal{S}_i of Node i is called an island if no other domain contains \mathcal{S}_i . In other words:


Definition 2.2 Let $i \in \mathcal{U}$. Then \mathcal{S}_i is an island if there exists no node j such that $\mathcal{S}_j \supset \mathcal{S}_i$.

Islands are disjoint, and form a partition of the set \mathcal{U} . It is a consequence of the following lemma.

Lemma 2.2 Consider two nodes $i, j \in \mathcal{U}$. Then either $\mathcal{S}_i \cap \mathcal{S}_j = \emptyset$, either $\mathcal{S}_i \subseteq \mathcal{S}_j$, or $\mathcal{S}_j \subseteq \mathcal{S}_i$.

Proof: We discuss first the case where $j \in \mathcal{S}_i$. If $\mathcal{S}_j \not\subseteq \mathcal{S}_i$, as illustrated in Figure 2.6, \mathcal{C}_i and \mathcal{C}_j have two intersection points \mathbf{x} and \mathbf{y} .

If the number of edges cut by \mathcal{C}_i between \mathbf{x} and \mathbf{y} is less or equal to the number of edges cut by the portion of \mathcal{C}_j located between \mathbf{x} and \mathbf{y} and inside \mathcal{S}_i , then we can construct a new curve \mathcal{C}'_j surrounding X_j by replacing, in \mathcal{C}_j ,


Figure 2.6: Curves C_i and C_j when $j \in C_i$.

the latter portion between x and y by the former one. This new curve cuts fewer edges or surrounds fewer nodes, which is impossible by definition of C_j .

If the number of edges cut by C_i between x and y is strictly more than the number of edges cut by the corresponding portion of C_j , then one could construct a new curve C'_i surrounding X_i in the same way as above. This new curve cuts fewer edges than C_i , what is impossible by definition.

Therefore, the situation represented in Figure 2.6 is not possible, and Curve C_j is always surrounded by C_i . This implies $S_j \subseteq S_i$. Similarly, $i \in S_j$ implies $S_i \subseteq S_j$.

Let us discuss now the case where $i \notin S_j$ and $j \notin S_i$. Suppose that $S_i \cap S_j \neq \emptyset$. \mathcal{U} is therefore partitioned into four disjoint sets, $\mathcal{N} = S_i \cap S_j$, $\mathcal{A} = S_i \setminus S_j$, $\mathcal{B} = S_j \setminus S_i$ and $\tilde{\mathcal{U}} = \mathcal{U} \setminus (S_i \cup S_j)$. Denote by $n_{\mathcal{A}\mathcal{B}}$ the number of edges between elements of \mathcal{A} and elements of \mathcal{B} , by $n_{\mathcal{A}\mathcal{N}}$ the number of edges between elements of \mathcal{A} and elements of \mathcal{N} , and so on (see Figure 2.7).

Figure 2.7: Number of edges between sets \mathcal{A} , \mathcal{B} , \mathcal{N} and $\tilde{\mathcal{U}}$.

Remember that C_i surrounds the elements of $S_i = \mathcal{A} \cup \mathcal{N}$. Consider now a curve that surrounds the elements of \mathcal{A} only. By definition of C_i , this curve must intersect strictly more edges than C_i , because $\mathcal{N} \neq \emptyset$. In other words, if we count these edges, we must have

$$n_{\mathcal{A}\tilde{\mathcal{U}}} + n_{\mathcal{N}\tilde{\mathcal{U}}} + n_{\mathcal{B}\mathcal{N}} + n_{\mathcal{A}\mathcal{B}} < n_{\mathcal{A}\tilde{\mathcal{U}}} + n_{\mathcal{A}\mathcal{B}} + n_{\mathcal{A}\mathcal{N}}$$

and thus

$$n_{\mathcal{N}\tilde{\mathcal{U}}} + n_{\mathcal{B}\mathcal{N}} < n_{\mathcal{A}\mathcal{N}}. \quad (2.11)$$

Similarly, by considering $S_B = \mathcal{B} \cup \mathcal{N}$, we obtain

$$n_{\mathcal{N}\tilde{\mathcal{U}}} + n_{\mathcal{A}\mathcal{N}} < n_{\mathcal{B}\mathcal{N}} \quad (2.12)$$

Combining (2.11) and (2.12) together, we obtain a contradiction, which proves that $\mathcal{N} = \mathcal{S}_A \cap \mathcal{S}_B = \emptyset$. \square

A consequence of Lemma 2.2 is that for each node $j \in \mathcal{S}_i$, $\mathcal{S}_j \subseteq \mathcal{S}_i$. It follows then from Definition 2.2 and Lemma 2.2 that islands are disjoint. Moreover, every node of \mathcal{U} belongs to one (and only one) island:

Theorem 2.5 *For each node $i \in \mathcal{U}$, there exists a unique island \mathcal{I}_i containing i , and which is given by*

$$\mathcal{I}_i = \bigcup_{j: \mathcal{S}_j \ni i} \mathcal{S}_j. \quad (2.13)$$

Proof: (i) Existence: We have to prove that the set defined in (2.13) is actually an island. Consider $j, k \in \mathcal{U}$ with $\mathcal{S}_j \ni i$ and $\mathcal{S}_k \ni i$. As \mathcal{S}_j and \mathcal{S}_k are not disjoint, we know from Lemma 2.2 that either $\mathcal{S}_j \subseteq \mathcal{S}_k$ or $\mathcal{S}_k \subseteq \mathcal{S}_j$. Thus, the union in (2.13) is equal to \mathcal{S}_l for some node l .

On the other hand, the condition in Definition 2.2 is verified by construction. \mathcal{I}_i is therefore an island.

(ii) Uniqueness. As islands are disjoint, two islands cannot contain Node i simultaneously. \square

Since every node of \mathcal{U} belongs to an island, the set \mathcal{U} can be partitioned into islands. We can thus define a new simplified graph with islands as nodes, and connections between islands as edges. This new graph, with a coarser granularity than $\mathcal{G}(\lambda, r)$, describes the “backbone” structure of the unbounded cluster. An example of a connectivity graph and its simplified version is given in Figure 2.8. The interesting property of the simplified graph is that the minimal cutset between any two nodes (provided that they are sufficiently distant) has the same size in the original graph and in the simplified graph.

This division of the network into island suggests a very natural way of grouping nodes together (*clustering*). Such clustering is often required by protocols, in order to deal with the large number of nodes. For example, many existing routing protocols use a clustering algorithm. The interest of the islands is that they are connected by bottleneck edges exclusively. Therefore, any algorithm running on the simplified graph would address the limited resource directly, as better connectivity is available inside the islands (by construction).

The development of algorithms for identifying the islands efficiently is however outside the scope of this work.

2.4 Hybrid Networks

The introduction of base stations allows distant nodes to communicate through a fixed, wired infrastructure. It represents a trade-off between today’s cellular networks and large-scale multi-hop networks. In the first case, every node connects to the nearest base station. To prevent nodes from being isolated from the network, base stations must therefore cover the whole space. In the second case, as we have seen in the previous sections, the density and radius must be large enough to keep the proportion of nodes which do not belong to \mathcal{U} small enough.

In this section we want to evaluate whether the introduction of base stations is helpful in decreasing the probability that an arbitrary node remains unconnected to the network, for the three geometries of a line, a plane and a strip.




Figure 2.8: A connectivity graph $\mathcal{G}(\lambda, r)$ (upper graph) and the corresponding simplified graph (lower graph). In the upper graph, edges connecting two islands together are drawn in red. These red edges are also the edges of the lower graph, where all nodes of each island are grouped into a single node (placed on the center of gravity of the island, in this illustration). Note that multiple edges can connect the same two vertices (not visible on the figure), when several edges connect the same two islands in the original graph. To approximate an infinite domain, we used periodic boundary conditions in this simulation.

We assume here that the radius of connectivity of the base stations is the same as the one of the wireless nodes.

2.4.1 Nodes and base stations on a line

Base stations are placed every L distance units, say at every nL for $n \in \mathbb{Z}$, and therefore define intervals $[nL, (n+1)L]$ of length L on the line \mathbb{R} .

With no loss of generality, we only consider here nodes that are located in the first interval $[0, L]$. Such a node is connected to the left (respectively, right) base station if and only if all intervals separating any pair of consecutive nodes located in $[0, x]$ (resp., $[x, L]$) have a length less than r . Because of the Poisson assumption, these intervals are independent, so that the probability that the node located at position x be connected to both base stations is the product of $P_c(x)$ (the probability that the node is connected to the left base station) and $P_c(L-x)$ (the probability that the node is connected to the right base station). Therefore the probability that an arbitrary node located at distance $0 \leq x \leq L$ from the origin is connected to a base station is

$$P_{cb}(x) = P_c(x) + P_c(L-x) - P_c(x)P_c(L-x)$$

where $P_c(x)$ is given by (2.1).

Of course, the closer we pick a node from a base station, the more likely it will be connected. The average of the connection probability of a node to a base station, over all positions that this node can have (which are uniformly distributed over a given interval, because of the Poisson assumption) is $\int_0^L P_{cb}(x)dx/L$. A more interesting value is a lower bound on P_{cb} , which is valid for the worst possible location of a node, which is at mid distance between the base stations:

$$\begin{aligned} P_{cb}^+ &= \inf_{0 \leq x \leq L} \{P_{cb}(x)\} = P_{cb}(L/2) \\ &= 2P_c(L/2) - P_c^2(L/2). \end{aligned} \quad (2.14)$$

Using the bounds (2.4), we obtain the following lower bound on P_{cb}^+ .

Theorem 2.6 P_{cb}^+ is lower bounded by:

$$P_{cb}^+ \geq 1 - (1 + \lambda L/2)^2 e^{-2\lambda r}. \quad (2.15)$$

Proof: If $0 \leq L \leq 2r$, then $P_{cb}^+ = 1$, obviously: in this case, the base stations cover the entire line. If $2r \leq L \leq 4r$, then $P_c(L/2) = 1 - e^{-\lambda r} - \lambda L e^{-\lambda r}/2$ because of (2.1). Inserting this value in (2.14) yields $P_{cb}^+ = 1 - (1 + \lambda L/2)^2 e^{-2\lambda r}$.

Finally, if $L \geq 4r$, we insert the bounds (2.4) in (2.14), which becomes

$$\begin{aligned} P_{cb}^+ &\geq 2(1 - e^{-\lambda r}) e^{-\lambda(L-4r)e^{-\lambda r}/2} \\ &\quad - 2\lambda r e^{-\lambda r} - (1 - e^{-\lambda r})^2 e^{-\lambda(L-2r)e^{-\lambda r}} \end{aligned}$$

The expansion of the right hand side of this inequality in Taylor series yields (2.15) after lengthy manipulations detailed in Appendix A.1. \square

One deduces from (2.15) the maximal spacing between consecutive base stations to keep the probability of finding a node not connected to a base station, in an arbitrary interval, below a given value $P_{uc}^+ = 1 - P_{cb}^+$. This spacing is

at least equal to $2r$, in which case the base stations do cover the entire line ($P_{uc}^+ = 0$). Combining this with (2.15), we see that the maximal distance between base stations is lower bounded by the following expression

$$\left(\frac{L}{r}\right)_{\max} \geq \max \left\{ 2, \frac{2}{\lambda r} \left(e^{\lambda r} \sqrt{P_{uc}^+} - 1 \right) \right\}. \quad (2.16)$$

The distance between base stations can grow exponentially with λr , provided it is large enough to bring the right hand side of this inequality above 2.

2.4.2 Nodes and base stations on a plane

We now place the base stations on the nodes of a square lattice, with edges of length equal to L . In the two-dimensional case, P_{cb} cannot be computed analytically. We compute it therefore numerically. Figure 2.9 shows P_{cb} as a function of L in both the super- and sub-critical cases.

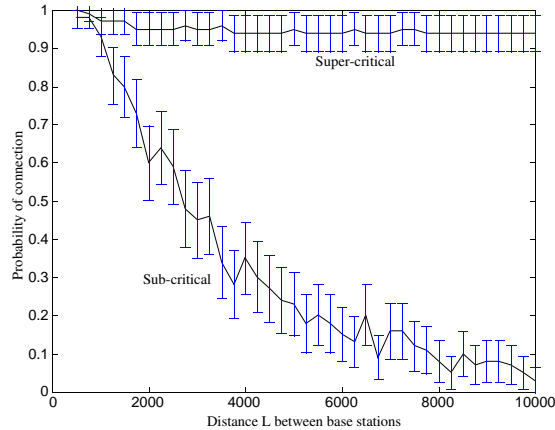


Figure 2.9: Probability of connection P_{cb} in 2 dim. with base stations for a sub-critical density $\lambda = 2.1 \cdot 10^{-5}$ and a super-critical density $\lambda = 2.56 \cdot 10^{-5}$.

For sub-critical densities, the probability of connection is now non-zero, but still far away from one, except in the limit case where base stations cover almost the whole space. Connectivity is thus not significantly better than with a standard cellular network, where only one hop is allowed between nodes and base stations.

For super-critical densities, we can see that the probability of connection is slightly better than without base stations. Furthermore, the probabilities take almost identical values in both cases for high densities. Let us explain why.

If $\lambda > \lambda^*$, there exists a unique unbounded cluster C_∞ . This cluster is in contact with a base station almost surely (the probability for each base station to be connected to the unbounded cluster is positive, and there are an infinite number of them). All nodes in \mathcal{U} are thus connected to the base stations. The situation for this subset of nodes is actually the same as that of a purely ad hoc network. The difference comes from nodes that are connected to a base station through a bounded cluster.

The quantity in which we are interested is then the fraction of bounded clusters connected to a base station. We know from [MR96] that the average number of nodes contained in a bounded cluster tends to one when $\lambda \rightarrow \infty$. Therefore the probability that a bounded cluster is connected to a base station tends to the fraction of space covered by the base stations, which is $\pi r^2/L^2$. The total probability that an arbitrary node is connected is thus

$$P_{cb} \stackrel{\lambda \rightarrow \infty}{\simeq} \theta(\lambda\pi r^2) + (1 - \theta(\lambda\pi r^2)) \frac{\pi r^2}{L^2}.$$

This shows that the gain of connectivity for high densities grows linearly with the fraction of space covered by the base stations.

We saw in Section 2.3 that one can divide the unbounded cluster into *islands*, and that the bottlenecks appear right between them. Inside an island, the connectivity is better, meaning that if a base station is present inside an island, the number of paths between the nodes of this island and the remainder of the network will be increased. As the size of the islands is of order of L_i , we can see on the graph of Figure 2.5 that this size decreases quickly with an increasing density, meaning the probability of an island to contain a base station becomes small. In other words, for high densities, the benefit (in terms of connectivity and number of paths) of inserting a fixed infrastructure of base stations is small. For intermediate densities, the benefit needs further study. One needs to investigate how the lattice of base stations covers the islands, following probably the approach of [BB01, BBT02].

2.4.3 Nodes and base stations on a strip

The behavior of the network is very different in one and in two dimensions. It is therefore interesting to explore the transition between these two behaviors, thanks to the infinite strip of width d of Section 2.2.3. Remember that percolation never occurs for d finite: one expects therefore to benefit from base stations for connectivity. This will certainly be the case for very small values of d . On the other hand, as $d \rightarrow \infty$, one gets closer a network on \mathbb{R}^2 , and the need for base stations to ensure connectivity should vanish.

In Figure 2.10, we set the node density to $\lambda = \frac{\lambda_1}{d} + \lambda_2$, so that the model tends to a one-dimensional Poisson point process of intensity λ_1 when $d \rightarrow 0$, and to a two-dimensional Poisson point process of intensity λ_2 when $d \rightarrow \infty$. We computed by simulation the maximum distance L between base stations, as a function of the strip width d , in order to achieve a probability of connection larger than 90%. We chose $\lambda_2 > \lambda^*$, so that the 90% connectivity is achieved without base station for $d = \infty$. The result in Figure 2.10 confirms our expectation. One can see that $L(d)$ grows regularly, meaning that there is continuous transition between the one-dimensional and two-dimensional cases.

2.5 Simulation study of a non-homogeneous node distribution

In this section, instead of generating the positions of the nodes with a Poisson point process, we use node distributions that are based on real population data².

²Source : Population census 1990, OFS GEOSTAT




Figure 2.10: Maximal distance L between the base stations as a function of the width d of the strip. $r = 300$ pixels, $\lambda_1 = 8 \cdot 10^{-3}$ and $\lambda_2 = 2.3 \cdot 10^{-5}$

We consider two regions of Switzerland: the Zurich area, and the Alpine Valley of Surselva (Fig 2.11). The first region present a quite uniform population two-dimensional density, whereas the second one has an almost one-dimensional shape. We assume that each resident has a probability 0.01 to own a (turned-on) device, and we simulate the network for different values of r .

Let us discuss the region of Zurich first. In the pure ad-hoc case, when r increases, we do observe the phase transition predicted by the model (see Figure 2.12). Furthermore, when r is sufficiently large, nodes are connected with high probability (more than 99.9% for $r = 2500\text{m}$). From a connectivity point of view, a pure ad-hoc network would therefore be feasible in this region.

If we insert base stations every 5km, the gain of connectivity is not significant above $r = 1500$ m. We actually obtain full connectivity only for $r = 2500$ m, as with a pure ad-hoc network.

In the Surselva Valley, the picture is very different. In fact, without base stations, we never observe clusters that spread from one end of the valley to the other. It means that long range communication is impossible. If we add base stations every 5km (i.e. 11 units), the probability of connection then becomes much larger, already for transmission ranges of a few hundreds of meters. However, a 99 % connection probability is only reached when base stations cover almost entirely the area. To allow more distance between base stations, one needs a larger fraction of active nodes than 0.01 in this little populated region.

We can conclude from these simulations that the qualitative behavior of the network remains the same with a realistic node distribution, as with a Poisson distribution.

2.6 Conclusion

We reviewed in this chapter the connectivity properties of the Boolean model. Many of the results presented above will be used in the next chapters, since




Figure 2.11: Swiss population density in 1990. In the frames, the two considered regions of Zurich (zoom at bottom left) and Surselva Valley (zoom at bottom right).




Figure 2.12: Probability of connection in the region of Zurich without base stations and with base stations every 5km as a function of r .




Figure 2.13: Probability of connection in Surselva Valley with base stations every 5km as a function of r .

continuum percolation is a central tool in this dissertation.

We found out that in one dimension, pure multi-hop relaying never allows long range communication. We computed the probability that two nodes are connected as a function of the distance between them, and showed that it is an exponentially decreasing function. However, in one dimension, the use of a sparse network of wired base stations improves dramatically the connectivity. We also computed bounds on the probability that a node is connected in this case.

The picture is very different in two dimensions. The percolation phenomenon allows unlimited range communications for a certain fraction of the nodes. Furthermore, this fraction converges rapidly to one when the node density increases. The number of paths between two arbitrary nodes remains however limited by some bottlenecks. Although it does not help much for connectivity in two dimensions, the presence of base stations can be a good answer to the problem of low capacity at large node density [LLT03]. However, we will see in Chapter 6 that the cluster of connected nodes is sufficiently meshed, so that it can sustain a throughput of the same order as on a regular lattice.

Chapter 3

Impact of interferences

We have seen in Chapter 2 that the Boolean model is suitable for studying connectivity, but requires the assumption that interferences are inexistent or at least negligible. In the case where interferences cannot be neglected, the Boolean model cannot be used anymore. In order to take them into account, a model based on the signal to interference ratio was used in [GK00]. In this last paper, which departs from a deterministic and finite population setting, all stations are assumed to have the same power, and some attenuation function is given. Station A can receive a signal from station B if the ratio of the power it receives from B to the total power received from all other stations is above a threshold.

The same physical model was analyzed in [BB01] in the infinite plane case under Poisson assumptions within the context of CDMA networks. The corresponding coverage process has connection with Poisson shot noise processes.

In this chapter, we study the connectivity of infinite wireless multihop networks under the physical model alluded to above. The parametric setting will be that of an homogeneous Poisson point process. Our main goal within this context is to learn whether the percolation phenomenon that was established in Chapter 2 for the case without interference still holds within this more realistic context.

By analogy with CDMA networks, we will introduce some orthogonality factor γ , which can vary from 0 to 1, and which stems from the imperfect orthogonality of the codes used in CDMA. The case with $\gamma = 0$ (perfect orthogonality) boils down to the case considered in Chapter 2.

As we will see, there are essential differences between the case $\gamma = 0$ and $\gamma > 0$. In some natural cases, for the same patterns, the first case could have an infinite component of the connectivity graph, whereas the second one could have no infinite component, or even no connectivity at all.

We show in Section 3.2 that percolation holds under conditions similar to those of the Boolean model provided the orthogonality factor γ is small enough. In this sense, connectivity of wireless multihop networks scales well with the size of the network even in the case of models that take interferences into account.

In Section 3.3, we study the asymptotic behavior of the model for large densities. We show that the shape of the attenuation function $l(\cdot)$ has a crucial impact on this behavior. In particular, we show that the most popular attenuation function, namely the power law function $l(x) = x^{-\alpha}$, where α is the

path loss exponent enjoys particular scaling properties and a singularity at the origin that lead to an increasingly good connectivity when the node density increases. However, in the most realistic case where $l(x)$ is bounded from above, the network becomes disconnected when the node density is too high.

The vast majority of experimental power-law attenuation models are valid at large scales, relatively far away from the emitting antenna. Experimental models valid at a smaller scale (from a few meters to a few kilometers) are usually piecewise power laws, with different path loss exponents, increasing with the distance from the antenna [Has93]. For example, [BH96] found by regression a piecewise power law with two segments, the first one having a path loss exponent between 1.4 and 2 for a range of distances between 1 and 500 meters, and the second one having a path loss exponent between 4 and 10 for a range of distances between 500 and 1000 meters. The attenuation function decreases slower close to the origin than a strict power law, and is bounded from above for physical reasons. For models neglecting interferences, and for extended networks, these changes from a strict power law attenuation function may be only second order effects on the performance of the network. For models of dense networks taking interferences into account, such as the STIRG, we will show that it is no longer so.

Since percolation may hold for very small values of $\gamma > 0$, narrow band communications may not be possible if we let all nodes emit simultaneously. A remedy is to use TDMA, so that each node is allowed to emit every t -th time slot. We show in Section 3.4 that a very simple TDMA scheme achieves connectivity similar to the previous one, with γ/t . We also look at the asymptotic behavior of the number of time slots required for percolation, when the node density increases, and show that it grows linearly with λ .

The type of random graphs that are introduced in this chapter are of independent interest. In particular, this class of random graphs which are built on the points of a Poisson point process, may simultaneously have infinite components, bounded range (each edge is of bounded length), and bounded degree (each vertex is of bounded degree).

Note that the first proof of the existence of a percolation threshold in this model was given in [DBT03]. However, this proof only holds under the assumption of an attenuation function with a bounded support, and for sufficiently high density. We give in this chapter a more general proof, that applies to a large class of attenuation functions. We also show that the minimum node density for percolation is exactly the same as the critical density of the Boolean model.

3.1 Model

We consider a multiple-hop ad-hoc network where nodes are distributed according to a Poisson point process of constant spatial intensity λ . Depending on its location, number of neighbors, and battery level, each node i will adjust its emitting power P_i within a given range $[0, P]$, where P is the maximal power of a node, which is finite. The power of the signal emitted by Node i and received by Node j is $P_i L(X_i - X_j)$, where X_i and X_j are the positions of Node i and j in the plane, respectively, and $L(\cdot)$ is the attenuation function in the wireless medium.

To describe the available links between these nodes, we use the *physical*

model as described in [GK00]. We define the signal to noise and interference ratio (SINR) of the signal received by Node j from Node i as

$$\beta_{ij} = \frac{P_i L(X_i - X_j)}{N_0 + \gamma \sum_{k \neq i, j} P_k L(X_k - X_j)}, \quad (3.1)$$

where N_0 is the power of the thermal background noise. The coefficient γ is the inverse of the processing gain of the system, it weights the effect of interferences, depending on the orthogonality between codes used during simultaneous transmissions. It is equal to 1 in a narrow band system, and is smaller than 1 in a broadband system that uses CDMA. The physical model of Gupta and Kumar [GK00] assumes $\gamma = 1$; other models [GT01] allow γ to be smaller than 1.

Next, we make the assumption that two nodes are able to communicate directly if the SINR of the signal they receive from each other is above a certain threshold β :

Definition 3.1 *There exists a link between Node i and Node j if*

$$\beta_{ij} > \beta \quad \text{and} \quad \beta_{ji} > \beta \quad (3.2)$$

With this definition, we obtain a non-oriented graph that describes the connectivity of the network. We call this graph STIRG (Signal to interference ratio graph), and study its properties throughout this chapter. In this work, we choose to neglect unidirectional links, which are difficult to exploit in wireless networks [RCM02].

As our model has many more parameters than degrees of freedom, we will focus on the node density λ and the orthogonality factor γ . The other parameter are supposed constant in the sequel. We will thus denote by $\mathcal{H}(\gamma, \lambda)$ the connectivity graph.

3.1.1 A Bound on the Degree of the Nodes

In the following theorem, we will prove that if $\gamma > 0$, the number of neighbors of each node is bounded from above (note that this is not the case in the Boolean Model with $\gamma = 0$).

Theorem 3.1 *Each node can have at most $1 + 1/\gamma\beta$ neighbors.*

Proof: Pick any node (called hereafter Node 0), and let N be the number of its neighbors (i.e. the number of nodes to which Node 0 is connected). If $N \leq 1$, the claim is trivially proved. Suppose next that $N > 1$, and denote by 1 the node whose signal power received by Node 0 is the smallest but is non zero, namely is such that

$$P_1 L(X_1 - X_0) \leq P_i L(X_i - X_0), \quad i = 2 \dots N. \quad (3.3)$$

Since it is connected to Node 0, (3.2) imposes that

$$\frac{P_1 L(X_1 - X_0)}{N_0 + \gamma \sum_{i=2}^{\infty} P_i L(X_i - X_0)} \geq \beta. \quad (3.4)$$

Taking (3.3) into account, (3.4) implies that

$$\begin{aligned}
P_1 L(X_1 - X_0) &\geq \beta N_0 + \beta \gamma \sum_{i=2}^{\infty} P_i L(X_i - X_0) \\
&\geq \beta N_0 + \beta \gamma (N - 1) P_1 L(X_1 - X_0) \\
&\quad + \beta \gamma \sum_{i=N+1}^{\infty} P_i L(X_i - X_0) \\
&\geq \beta \gamma (N - 1) P_1 L(X_1 - X_0),
\end{aligned}$$

from which we deduce that

$$N \leq 1 + \frac{1}{\beta \gamma}.$$

□

In CDMA cellular networks, this kind of bound is known under the name of pole capacity (see e.g. [VS99], [BB01]).

As a consequence of Theorem 3.1, we see that if $\gamma > 1/\beta$, each node has at most one neighbor. This is a very general and restrictive condition, that imposes the network to use efficient spread-spectrum encoding in order to keep γ small, or to introduce a scheduling between nodes to avoid having them emitting all the same time – we will investigate such a scheme in Section 3.4.

3.1.2 Shot-Noise

The sum in the denominator of (3.1) is a random variable that depends on the position of almost all nodes in the network. We can write it as $N_0 + \gamma J(X_j) - \gamma P_i L(X_i - X_j)$ where

$$J(\mathbf{x}) = \sum_{i, X_i \neq \mathbf{x}} P_i L(X_i - \mathbf{x}) \quad (3.5)$$

is the *interference contribution*. This kind of variable is called a Poisson *shot-noise*. As it is an infinite sum, it may diverge to infinity, making connections impossible.

If we assume that the sequence $\{P_i\}$ is uniformly bounded from below by a strictly positive constant, and that the attenuation function $L(\cdot)$ has the form $L(\mathbf{x}) = l(\|\mathbf{x}\|)$ (isotropic) where $l(t)$ is a non increasing function of t , the necessary and sufficient condition for the sum

$$\sum_i P_i L(X_i - \mathbf{x})$$

to be a.s. finite is given in [Dal71]:

$$\int_y^{\infty} l(t) t dt < \infty, \text{ for a sufficiently large } y. \quad (3.6)$$

This condition remains valid if $\inf_i \{P_i\} = 0$ but the sequence $\{\mathbf{1}_{\{P_i < \varepsilon\}}\}$ is i.i.d. and independent from the point process, for some $\varepsilon > 0$ (like in Section 3.4).

We notice that for free space propagation (i.e. $l(t) = 1/t^2$), the integral in (3.6) is divergent and thus no connection is possible in this case whenever $\gamma > 0$.

This situation is similar to the setting of Olber's paradox, according to which the sky should be infinitely bright (see e.g. [Har87]).

By letting $y = 0$ in (3.6), we obtain the condition for *integrability*, which is stronger. This last property holds for all stationary point processes with finite intensity see e.g. [SKM95], and in particular in the homogeneous Poisson case.

3.1.3 Attenuation

For the attenuation, the most common function is

$$l(t) = \frac{1}{t^\alpha}, \quad (3.7)$$

with α ranging from 3 to 6. It makes sense to assume the attenuation to be a bounded function in the vicinity of the antenna. The following two functions:

- $l(t) = A[\max(t, r_0)]^{-\alpha}$,
- $l(t) = (1 + At)^{-\alpha}$,

with $A > 0$, are bounded modifications of the latter considered in [BB01].

3.2 Percolation

As our model is ergodic (it is a deterministic construction on a Poisson point process), the probability that there exists a cluster of infinite size¹ is either 0 or 1, depending on the parameters λ and γ . In the first case, as there are a.s. only finite clusters, the network is said *sub-critical*, whereas in the second case, it is said *super-critical*.

In the sub-critical phase, long range connections in multiple hops are not possible, contrary to the super-critical phase. It is thus a crucial property to establish in a network.

We begin this central section by noticing that the much simpler Boolean model is a particular case for our model when $\gamma = 0$. We then make some preliminary observations on simulations to show the difference between the graphs obtained when the interferences are neglected (which amounts to set $\gamma = 0$) or not (when $\gamma > 0$). In a third step, we prove that percolation occurs (i.e. an infinite cluster exists) for small, but nonzero values of γ . We finally give some asymptotic results for large node densities.

3.2.1 Existence of a percolation threshold for $\gamma = 0$

Let us first note that if we let $\gamma = 0$, the model described in Section 3.1 becomes equivalent to a generalized Boolean model, where two nodes are connected if and only if they are in a ball of radius r , independently from all the other nodes. Assuming all nodes emit at the maximum power P , this radius r is then constant and found from (3.1) and (3.2) to be

$$r = \sup \left\{ \rho \text{ such that } l(\rho) \geq \frac{\beta N_0}{P} \right\}.$$

¹We conjecture moreover that whenever it exists, the infinite cluster is also unique. The proof of this conjecture is out of the scope of this work.




Figure 3.1: **An example of graph $\mathcal{H}(0, \lambda)$ with no interference (Boolean Model).** As the node density is super-critical ($\lambda > \lambda^*$), most of the nodes belong to the same connected component. [This simulation was run in a square of $65,536 \times 65,536$ pixels with parameters $\lambda = 9.3110^{-4}$, $\beta = 1$, $\gamma = 0$, $N_0 = 1$, $P_i = 100,000 \forall i$.]

For example, for the attenuation function (3.7) we obtain $r = (P/(\beta N_0))^{1/\alpha}$. This is the model we have studied in the previous chapter, and for which many results from continuous percolation theory apply [MR96]. The most important one is mentioned above, namely that there is a critical density λ^* , above which the graph contains an infinite connected component.

3.2.2 Some observations on the graph with $\gamma > 0$

If $\gamma > 0$, it is clear that for the same realization of the spatial point process giving the position of the nodes, the graph obtained with $\gamma > 0$ misses some edges in the graph obtained with $\gamma' = 0$. In other words, $\mathcal{H}(\gamma, \lambda) \subseteq \mathcal{H}(0, \lambda)$. As a result, it is not sure that percolation still occurs for nonzero values of γ . At least, for $\lambda < \lambda^*$, we are sure that $\mathcal{H}(\gamma, \lambda)$ is always sub-critical. However, for $\lambda > \lambda^*$, we know that

1. for $\gamma = 0$, the network is super-critical
2. for $\gamma > 1/\beta$, the network is sub-critical.

Therefore, there exists a critical value $0 \leq \gamma^*(\lambda) \leq 1/\beta$ at which a phase transition occurs. The next subsection will prove that $\gamma^*(\lambda)$ is strictly positive for sufficiently large values of λ .

We have computed by simulation the value of the percolation threshold $\gamma^*(\lambda)$, with $L(\mathbf{x}) = \max(1, \|\mathbf{x}\|)^{-3}$. The simulation results are shown in Figures 3.1, 3.2, 3.3 and 3.4. In the simulations, all nodes emit with the same power P . We observe in Figure 3.1 and 3.2 that $\mathcal{H}(0.02, \lambda) \subseteq \mathcal{H}(0, \lambda)$. We observe in Figure 3.3 that $\gamma^*(\lambda)$ exhibits a maximum at a certain density $\tilde{\lambda}$. Below $\tilde{\lambda}$, increasing the node density helps for connectivity, whereas after the maximum,




Figure 3.2: **An example of graph $\mathcal{H}(\gamma, \lambda)$ with interferences ($\gamma = 0.02$).** This simulation was run with the same parameters as in Figure 3.1, except γ that is now nonzero. Due to the interferences, the graph is split into many small components.

the impact of interferences becomes preponderant, and $\gamma^*(\lambda)$ becomes decreasing. Figure 3.4 illustrates that a value of γ slightly smaller than $\gamma^*(\lambda)$ makes the graph percolate.

3.2.3 Percolation for non-zero values of γ

We have shown above that if γ exceeds some finite, positive critical value, percolation does not occur. We want now to show that percolation can occur for nonzero values of γ . We make the simplifying assumption that every node emits at maximal power P : $P_i = P$ for all i . This corresponds to the worst power assignment for the interfering communications.

We make the assumptions that $L(\mathbf{x}) = l(\|\mathbf{x}\|)$, and that $l(\cdot)$ fulfills the following conditions:

1. $l(x) < cx^{-\alpha}$ for some constants $c < \infty$ and $\alpha > 2$.
2. $l(0) > \frac{\beta N_0}{P}$.
3. $l(x) \leq 1$.
4. $l(x)$ is continuous and strictly decreasing.

The first assumption ensures that the attenuation function verifies Inequality (3.6). The second assumption is necessary for the existence of links. Otherwise, even when nodes are very close to each other, the signal-to-noise ratio at the receiver ($\simeq Pl(0)/N_0$) would be already too weak. These two assumptions are necessary conditions for percolation. The two last assumptions have been introduced for mathematical convenience, but clearly make sense in the real world. In particular, the third assumption imposes that the received power cannot be higher than the emitter power, as discussed in Section 3.1.3.

We can now state the main theorem of this chapter:




Figure 3.3: **Critical value of γ as a function of the node density λ .** The curve shows the critical value of γ below which the network percolates. [The parameters of this simulation are $\beta = 1$, $N_0 = 10^4$ and $P_i = 10^5 \forall i$.]




Figure 3.4: **A barely supercritical graph with interferences.** This simulation was run with the same parameters as in Figure 3.2, except that the node density is higher ($\lambda = 2.79 \cdot 10^{-3}$). The graph percolates despite the interferences because here $\gamma < \gamma^*(\lambda)$. One can observe that fewer edges are needed to achieve percolation than in Figure 3.1.

Theorem 3.2 *Let λ^* be the critical node density when $\gamma = 0$ (as defined in Section 3.2.1) and assume that the attenuation function has the form $L(\mathbf{x}) = l(\|\mathbf{x}\|)$, with $l(\cdot)$ fulfilling assumptions [1–4]. Then for any node density $\lambda > \lambda^*$, there exists $\gamma^*(\lambda) > 0$ such that for $\gamma \leq \gamma^*(\lambda)$, the connectivity graph $\mathcal{H}(\gamma, \lambda)$ percolates.*

The idea of the proof is the following: we fix the node density $\lambda > \lambda^*$, so that the graph $\mathcal{H}(0, \lambda)$ is super-critical. As the difference $\lambda - \lambda^*$ is strictly positive, one can increase a bit the value of N_0 to $N'_0 = N_0 + \varepsilon$ while keeping the graph super-critical. Then we remove all points of the graph where the interference term $\gamma \sum_{k \neq i, j} Pl(\|X_k - X_j\|)$ is larger than ε . By setting γ sufficiently small, we remove an arbitrary small number of points. We finally show that the graph without these points still percolates, i.e. that the origin still belongs to an infinite cluster with positive probability.

To deal with the dependence between the links, we start the proof by mapping the problem onto a lattice. We end up with a dependent bond percolation model, where the probability that an edge is open can be made arbitrarily large by tuning the scale of the mapping (renormalization argument). This mapping is chosen in a way that the existence of an infinite open component in the bond percolation model implies the existence of an infinite connected component in the original graph.

In a second step, we show that although the edges of the bond percolation model are dependent, the probability that a collection of n edges are all closed is less than q^n , for a certain constant $q < 2/9$ independent of the choice of the edges. This latter property implies the existence of an infinite open component. The next two sections describe in detail these two steps.

Mapping on a lattice

As we just saw in Section 3.2.1, if we set $\gamma = 0$, we obtain a fixed radius Poisson Boolean model, and thus

$$\mathcal{H}(0, \lambda) = \mathcal{G}(\lambda, 2r_b),$$

with

$$2r_b = l^{-1} \left(\frac{\beta N_0}{P} \right).$$

Since l is continuous and larger than $\beta N_0/P$ at the origin, we are sure that $l^{-1}(\beta N_0/P)$ exists.

We consider next the Boolean model $\mathcal{B}(\lambda, r_b)$ where the node density λ is slightly higher than the critical value λ^* , and the nodes have range r_b . This graph is clearly super-critical. Moreover, by rescaling the model (see Section 2.1.2), we can establish that the critical radius for a node density λ is

$$r^*(\lambda) = \sqrt{\frac{\lambda^*}{\lambda}} r_b < r_b.$$

Thus, the Boolean model $\mathcal{B}(\lambda, r)$ with $r^*(\lambda) < r < r_b$ is also super-critical.

We map this latter model on a discrete percolation model as follows. We denote by $\mathcal{L} = d \cdot \mathbb{L}^2$ the two-dimensional square lattice whose vertices are located at all points of the form (dx, dy) with $(x, y) \in \mathbb{Z}^2$. For each horizontal edge $a \in \mathcal{L}$, we denote by $\mathbf{z}_a = (x_a, y_a)$ the point in the middle of the edge, and




Figure 3.5: A horizontal edge a that fulfills the two conditions for having $A_a = 1$.

introduce the random variable A_a that takes value 1 if the following two events (illustrated in Figure 3.5) happen, and 0 otherwise:

1. the rectangle $[x_a - 3d/4, x_a + 3d/4] \times [y_a - d/4, y_a + d/4]$ is crossed from left to right by an occupied component in $\mathcal{B}(\lambda, r)$, and
2. both squares $[x_a - 3d/4, x_a - d/4] \times [y_a - d/4, y_a + d/4]$ and $[x_a + d/4, x_a + 3d/4] \times [y_a - d/4, y_a + d/4]$ are crossed from top to bottom by an occupied component in $\mathcal{B}(\lambda, r)$.

We define A_a similarly for vertical edges, by rotating the above conditions by 90° .

According to Corollary 4.1 in [MR96], the probability that $A_a = 1$ can be made as large as we like by choosing d large. The variables A_a are not independent in general. However, if a and b are not adjacent, then A_a and A_b are independent: adjacent edges depend on the realization of the Poisson point process on a common surface $R_a \cap R_b$. These variables thus define a 1-dependent edge percolation process.

We define now a second random field, B_a , $a \in \mathcal{L}$, as follows: we define first a shifted function \tilde{l} of the attenuation function l as:

$$\tilde{l}(x) = \begin{cases} l(0) & x \leq \frac{\sqrt{10}d}{4} \\ l(x - \frac{\sqrt{10}d}{4}) & x > \frac{\sqrt{10}d}{4} \end{cases}$$

We define then the *shot-noise* processes I and \tilde{I} as follows:

$$I(\mathbf{z}) = \sum_k l(\|\mathbf{z} - X_k\|)$$

and

$$\tilde{I}(\mathbf{z}) = \sum_k \tilde{l}(\|\mathbf{z} - X_k\|)$$

where $\mathbf{z} \in \mathbb{R}^2$ is an arbitrary point.

We define now the second random field B_a that takes value 1 if the value of the shot-noise process never exceeds a certain threshold M in the rectangle $R(\mathbf{z}_a) = [x_a - 3d/4, x_a + 3d/4] \times [y_a - d/4, y_a + d/4]$. As the distance between any point of the rectangle and its center \mathbf{z}_a is at most $\sqrt{10}d/4$, the triangle inequality yields that $\|\mathbf{z}_a - X_k\| \leq \sqrt{10}d/4 + \|\mathbf{z} - X_k\|$, and thus that $I(\mathbf{z}) \leq \tilde{I}(\mathbf{z}_a)$ for all $\mathbf{z} \in R(\mathbf{z}_a)$. Therefore, $\tilde{I}(\mathbf{z}_a) < M$ is a sufficient condition for having $B_a = 1$.

Percolation in the lattice

Now we have to prove that the percolation process defined by the product $C_a \doteq A_a B_a$ is super-critical for appropriate values of M and d . To do this, we prove that the probability that the origin belongs to an infinite open cluster is positive. This happens if in the dual lattice, the probability to find a closed circuit that surrounds the origin is less than one. We will compute an upper bound to the latter probability.

First, we consider a circuit of length n , and denote by $\{a_i\}_{i=1}^n$ the edges of this circuit (clearly, n has to be even). For the sake of simplicity, we will abusively write $A_{a_i} = A_i$, $B_{a_i} = B_i$ and so on.

$$\begin{aligned} p(n) &\doteq \mathbb{P}(C_1 = 0 \cap C_2 = 0 \cap \dots \cap C_{n-1} = 0 \cap C_n = 0) \\ &\leq \mathbb{P}(C_1 = 0 \cap C_3 = 0 \cap \dots \cap C_{n-1} = 0) \end{aligned}$$

There are $n/2$ events in the above probability. We have furthermore

$$\begin{aligned} p(n) &\doteq \mathbb{P}(C_1 = 0 \cap C_3 = 0 \cap \dots \cap C_{n-1} = 0) \\ &= \mathbb{E}((1 - C_1)(1 - C_3) \dots (1 - C_{n-1})) \\ &= \mathbb{E}((1 - A_1 B_1) \dots (1 - A_{n-1} B_{n-1})). \end{aligned}$$

It is easy to see that

$$1 - A_i B_i \leq (1 - A_i) + (1 - B_i) \doteq \bar{A}_i + \bar{B}_i.$$

Let $(k_i)_{i=1}^{n/2}$ be a binary sequence of length $n/2$, with $k_i = 0$ or 1 for all $1 \leq i \leq n/2$. We denote by \mathcal{K} the set of the $2^{n/2}$ such sequences. Then we can write

$$\begin{aligned} p(n) &\leq \mathbb{E}((\bar{A}_1 + \bar{B}_1)(\bar{A}_3 + \bar{B}_3) \dots (\bar{A}_{n-1} + \bar{B}_{n-1})) \\ &\leq \mathbb{E}\left(\sum_{(k_i) \in \mathcal{K}} \prod_{i:k_i=0} \bar{A}_{2i-1} \prod_{i:k_i=1} \bar{B}_{2i-1}\right) \\ &= \sum_{(k_i) \in \mathcal{K}} \mathbb{E}\left(\prod_{i:k_i=0} \bar{A}_{2i-1} \prod_{i:k_i=1} \bar{B}_{2i-1}\right) \\ &\leq \sum_{(k_i) \in \mathcal{K}} \sqrt{\mathbb{E}\left(\prod_{i:k_i=0} \bar{A}_{2i-1}^2\right) \mathbb{E}\left(\prod_{i:k_i=1} \bar{B}_{2i-1}^2\right)} \\ &= \sum_{(k_i) \in \mathcal{K}} \mathbb{E}\left(\prod_{i:k_i=0} \bar{A}_{2i-1}\right)^{1/2} \mathbb{E}\left(\prod_{i:k_i=1} \bar{B}_{2i-1}\right)^{1/2} \end{aligned} \quad (3.8)$$

where the last inequality follows from Schwartz's inequality, and the last equality from the observation that $\bar{A}_i^2 = \bar{A}_i$ and $\bar{B}_i^2 = \bar{B}_i$.

As the edges x_1, x_3, \dots, x_{n-1} are not adjacent, the events A_1, A_3, \dots, A_{n-1} are independent. We have thus

$$\mathbb{E}\left(\prod_{i:k_i=0} \bar{A}_{2i+1}\right) = \prod_{i:k_i=0} \mathbb{E}(\bar{A}_{2i+1}) = \prod_{i:k_i=0} (1 - p_A) \quad (3.9)$$

Let us look at the second term. The events B_i are not independent, so we need some more complicated argument.

Lemma 3.1 For a sequence $\{a_i\}_{i=1}^L$ of edges, $L \in \mathbb{N}$, and the respective events $\{\bar{B}_i\} = \{\bar{B}_{a_i}\}$, we have

$$\mathbb{E}(\bar{B}_1 \bar{B}_2 \dots \bar{B}_L) \leq (1 - p_B)^L$$

for some $p_B > 0$.

Proof:

To simplify notations, we denote by \mathbf{z}_i the center \mathbf{z}_{a_i} of Edge a_i . Using generalized Chernov's bound, we have for any $s \geq 0$

$$\begin{aligned} \mathbb{E}(\bar{B}_1 \bar{B}_2 \dots \bar{B}_L) &= \mathbb{P}(B_1 = 0, B_2 = 0, \dots, B_L = 0) \\ &\leq \mathbb{P}\left(\tilde{I}(\mathbf{z}_1) > M, \tilde{I}(\mathbf{z}_2) > M, \dots, \tilde{I}(\mathbf{z}_L) > M\right) \\ &\leq e^{-sLM} \mathbb{E}\left(e^{s \sum_{i=1}^L \tilde{I}(\mathbf{z}_i)}\right) \end{aligned}$$

Using Campbell's theorem (see e.g. [Kin93, pages 28–31]), we obtain

$$\mathbb{E}\left(e^{s \sum_{i=1}^L \tilde{I}(\mathbf{z}_i)}\right) = \exp\left(\lambda \int_{\mathbb{R}^2} (e^{s \sum_{i=1}^L \tilde{l}(\|\mathbf{x} - \mathbf{z}_i\|)} - 1) d\mathbf{x}\right). \quad (3.10)$$

We need now to bound the exponent $s \sum_{i=1}^L \tilde{l}(\|\mathbf{x} - \mathbf{z}_i\|)$. As $\{\mathbf{z}_i\}$ are centers of edges, they are located on a square lattice with edge length $d/\sqrt{2}$. So, if we consider the square in which \mathbf{x} is located, the contribution to $\sum_{i=1}^L \tilde{l}(\|\mathbf{x} - \mathbf{z}_i\|)$ coming from the four corners of this square is at most equal to 4, since $\tilde{l}(x) \leq 1$. Around this square, there are 12 nodes, each located at distance at least $d/\sqrt{2}$ from \mathbf{x} . Further away, there are 20 other nodes at distance at least $2d/\sqrt{2}$, and so on.

On the other hand, $l(\mathbf{x}) \leq c\|\mathbf{x}\|^{-\alpha}$ because of Assumption [1] on the attenuation function, so that

$$\tilde{l}(\mathbf{x}) \leq \left(\frac{\sqrt{10}d}{4} + c^{1/\alpha}\right)^\alpha \|\mathbf{x}\|^{-\alpha} \doteq \tilde{c}\|\mathbf{x}\|^{-\alpha}$$

with $\alpha > 2$. Consequently,

$$\begin{aligned} \sum_{i=1}^L \tilde{l}(\|\mathbf{x} - \mathbf{z}_i\|) &\leq \sum_{i=1}^{\infty} \tilde{l}(\|\mathbf{x} - \mathbf{z}_i\|) \\ &\leq 4 + \sum_{k=1}^{\infty} (4 + 8k) \tilde{c} \left(\frac{\sqrt{2}}{kd}\right)^\alpha \\ &= 4 + \sum_{k=1}^{\infty} \left[4\tilde{c} \left(\frac{\sqrt{2}}{kd}\right)^\alpha + 8k\tilde{c} \left(\frac{\sqrt{2}}{kd}\right)^\alpha\right] \\ &\leq 4 + \sum_{k=1}^{\infty} \left[4\tilde{c} \left(\frac{\sqrt{2}}{kd}\right)^\alpha + \frac{8\sqrt{2}\tilde{c}}{d} \left(\frac{\sqrt{2}}{kd}\right)^{\alpha-1}\right] \doteq K. \end{aligned}$$

Both sums converge as $\alpha > 2$, and thus $K < \infty$.

The computation made above holds for any $s \geq 0$. We now take $s = 1/K$. In this case, we have $s \sum_{i=1}^L l(\|\mathbf{x} - \mathbf{z}_i\|) \leq 1$, for all \mathbf{x} . Furthermore, since $e^x - 1 < 2x$ for all $x \leq 1$, we have

$$e^{s \sum_{i=1}^L \tilde{l}(\|\mathbf{x} - \mathbf{z}_i\|)} - 1 < 2s \sum_{i=1}^L \tilde{l}(\|\mathbf{x} - \mathbf{z}_i\|) = \frac{2}{K} \sum_{i=1}^L \tilde{l}(\|\mathbf{x} - \mathbf{z}_i\|).$$

Replacing this in (3.10), we obtain

$$\begin{aligned} \mathbb{E}\left(e^{\sum_{i=1}^L \tilde{I}(\mathbf{z}_i)/K}\right) &< \exp\left(\lambda \int_{\mathbb{R}^2} \frac{2}{K} \sum_{i=1}^L \tilde{l}(\|\mathbf{x} - \mathbf{z}_i\|) d\mathbf{x}\right) \\ &= \exp\left(\frac{2}{K} \sum_{i=1}^L \lambda \int_{\mathbb{R}^2} \tilde{l}(\|\mathbf{x} - \mathbf{z}_i\|) d\mathbf{x}\right) \\ &= \exp\left(\frac{2L\lambda}{K} \int_{\mathbb{R}^2} \tilde{l}(\|\mathbf{x}\|) d\mathbf{x}\right) \\ &= \left[\exp\left(\frac{2\lambda}{K} \int_{\mathbb{R}^2} \tilde{l}(\|\mathbf{x}\|) d\mathbf{x}\right)\right]^L. \end{aligned}$$

Finally we have that

$$\begin{aligned} &\mathbb{P}\left(\tilde{I}(\mathbf{z}_1) > M, \tilde{I}(\mathbf{z}_2) > M, \dots, \tilde{I}(\mathbf{z}_L) > M\right) \\ &\leq e^{-sLM} \mathbb{E}\left(e^{s \sum_{i=1}^L \tilde{I}(\mathbf{z}_i)}\right) \\ &< e^{-LM/K} \left[\exp\left(\frac{2\lambda}{K} \int_{\mathbb{R}^2} \tilde{l}(\|\mathbf{x}\|) d\mathbf{x}\right)\right]^L \\ &= \left[\exp\left(\frac{2\lambda}{K} \int_{\mathbb{R}^2} \tilde{l}(\|\mathbf{x}\|) d\mathbf{x} - \frac{M}{K}\right)\right]^L. \end{aligned}$$

Defining

$$p_B \doteq 1 - \left[\exp\left(\frac{2\lambda}{K} \int_{\mathbb{R}^2} \tilde{l}(\|\mathbf{x}\|) d\mathbf{x} - \frac{M}{K}\right)\right], \quad (3.11)$$

the lemma is proved. \square

Lemma 3.1 allows use to bound from above the second term of (3.8).

$$\mathbb{E}\left(\prod_{k(i)=1} \bar{B}_{2i+1}\right) \leq \prod_{k(i)=1} (1 - p_B).$$

Combining this bound with (3.9), we obtain

$$p(n) \leq \sum_{(k_i) \in \mathcal{K}} \prod_{i:k_i=0} (1 - p_A)^{1/2} \prod_{i:k_i=1} (1 - p_B)^{1/2}.$$

We observe that this last expression is the $n/2$ -th power of a binomial. Indeed, by grouping together the sequences (k_i) that contain j zeros, we obtain

$$\begin{aligned} p(n) &\leq \sum_{j=0}^{n/2} \binom{n/2}{j} (1 - p_A)^{j/2} (1 - p_B)^{(n/2-j)/2} \\ &= (\sqrt{1 - p_A} + \sqrt{1 - p_B})^{n/2} \doteq q^n \end{aligned}$$

where we defined $q = (\sqrt{1-p_A} + \sqrt{1-p_B})^{1/2}$.

We consider now the *dual lattice* \mathcal{L}' of \mathcal{L} , obtained by placing a vertex at the center of each face of \mathcal{L} , and an edge between two vertices whenever the two corresponding faces of \mathcal{L} have a common edge. The main property of dual lattices is the one-to-one relation between edges of the original lattice and the edges of its dual. Furthermore, in our particular case, the dual lattice is also a square lattice (see Figure 3.6), and we can set $\mathcal{L}' = \mathcal{L} + (d/2, d/2) = \{\langle \mathbf{v}, \mathbf{w} \rangle \in \mathbb{R}^4 : \langle \mathbf{v} - (d/2, d/2), \mathbf{w} - (d/2, d/2) \rangle \in \mathcal{L}\}$.




Figure 3.6: Lattice \mathcal{L} (plain) and its dual \mathcal{L}' (dashed)

We look at the number of possible circuits that surround the origin in the dual lattice \mathcal{L}' . Clearly, a circuit that surrounds the origin has to cross the right half of the horizontal axis $\{(x, 0) | x > 0\}$. In the dual lattice, there is an infinite number of edges that cross this axis (see Figure 3.7), but if the circuit has length n , it has to contain one of the $n/2$ first such edges. Then, we can construct the circuit, starting from this edge, and choosing one of the three possible directions at each step. The number of circuits of length n is thus at most $\frac{n}{2} 3^{n-1}$.

We can now bound the probability to find a closed circuit surrounding the




Figure 3.7: Circuits in the dual lattice that surround the origin. All such circuits must contain at least one of the bold edges.

origin:

$$\begin{aligned}
 \mathbb{P}(\exists \text{ closed circ.}) &\leq \sum_{n=4}^{\infty} \mathbb{P}(\exists \text{ closed circ. of length } n) \\
 &\leq \sum_{n=4}^{\infty} \frac{n}{2} 3^{n-1} p(n) \\
 &\leq \sum_{n=4}^{\infty} \frac{n}{2} 3^{n-1} q^n \\
 &\leq \sum_{n=4}^{\infty} \frac{nq}{2} (3q)^{n-1} \\
 &\leq \frac{q}{2} \sum_{n=1}^{\infty} n (3q)^{n-1} \\
 &\leq \frac{q}{2} \left(\frac{1}{1-3q} \right)^2
 \end{aligned}$$

The latter expression is smaller than one if $q < 2/9$, i.e if $\sqrt{1-p_A} + \sqrt{1-p_B} < 4/81$. It is thus sufficient to have $p_A > 1 - 4/9^4$ and $p_B > 1 - 4/9^4$. We must now choose d and M such that these conditions are verified. First, we know that we can choose d sufficiently large so that $p_A > 1 - 4/9^4$. Second, it is easy to see in Equation (3.11) that one can choose M sufficiently large so that $p_B > 1 - 4/9^4$. It is thus possible to find appropriate parameters M and d to have percolation in the bond percolation model.

Percolation in the STIRG model

To conclude the proof of Theorem 3.2, we must show that percolation of C_a implies percolation in the SINR model, with appropriate γ . Actually, if $B_a = 1$, the interference level in the rectangle $R(z_a)$ is at most equal to M . Therefore,




Figure 3.8: Two adjacent edges a (plain) and b (dashed) with $A_a = 1$ and $A_b = 1$. The crossings overlap, and form a unique connected component.

for two nodes X_i and X_j such that $\|X_i - X_j\| \leq r$, we have

$$\begin{aligned} \frac{Pl(\|X_i - X_j\|)}{N_0 + \gamma \sum_{k \neq i, j} Pl(\|X_k - X_j\|)} &\geq \frac{Pl(\|X_i - X_j\|)}{N_0 + \gamma PM} \\ &\geq \frac{Pl(2r)}{N_0 + \gamma PM}. \end{aligned}$$

As $r < r_b$ and as $l(\cdot)$ is strictly decreasing (Assumption [4] on the attenuation function), we pick

$$\gamma = \frac{N_0}{PM} \left(\frac{l(2r)}{l(2r_b)} - 1 \right) > 0 \quad (3.12)$$

yielding

$$\frac{Pl(2r)}{N_0 + \gamma PM} = \frac{Pl(2r_b)}{N_0} = \beta.$$

Therefore, there exists a positive value of γ such that any two nodes separated by a distance less than r are connected in the SINR model. This means that in the rectangle $R(z_a)$ all connections of $\mathcal{B}(\lambda, r)$ also exist in the SINR model.

Furthermore, if $A_a = 1$, there exist crossings along edge a , as shown in Figure 3.5. These crossings are designed such that if for two adjacent edges a and b , $A_a = 1$ and $A_b = 1$, the crossings overlap, and they all belong to the same connected component (see Figure 3.8). Thus, an infinite cluster of such sites implies an infinite cluster in $\mathcal{B}(\lambda, r)$.

Here we have an infinite cluster of edges such that $A_a = 1$ and $B_a = 1$ simultaneously for each edge a . This means that the crossings exist in the SINR model, and thus form an infinite connected component. \square

3.3 Asymptotic behavior of the critical threshold for large node densities

3.3.1 Scaling properties

In this section, we study the symmetry properties of the STIRG model with respect to scaling. We derive a number of useful tools that we will need later to study its properties. Throughout this section, we keep the assumption that all nodes emit with the same power P .

The Poisson point process

We start by defining the dilation operator h_a of \mathbb{R}^2 , with $a \in \mathbb{R}_+$ as follows:

$$\begin{aligned} h_a : \mathbb{R}^2 &\longrightarrow \mathbb{R}^2 \\ \mathbf{x} &\longmapsto a\mathbf{x} \end{aligned}$$

Let us denote by Φ_λ the Poisson process over \mathbb{R}^2 of intensity λ (i.e. the random set containing the points). If we apply h_a to the Poisson point process, we obtain another Poisson point process, but with density λ/a^2 :

$$h_a\Phi_\lambda = \Phi_{\lambda/a^2}.$$

Denoting by Ω the probability space, and by ω one of its elements, one can define Φ_λ and Φ_{λ/a^2} such that if we apply h_a to a particular realization $\Phi_\lambda(\omega)$ of Φ_λ , we obtain the same set as with the corresponding realization of Φ_{λ/a^2} :

$$h_a\Phi_\lambda(\omega) = \Phi_{\lambda/a^2}(\omega) \quad \forall \omega \in \Omega.$$

Thus if x_1 and x_2 are two points of $\Phi_{\lambda,\omega}$, ax_1 and ax_2 belong to $\Phi_{\lambda/a^2,\omega}$. In particular, the distance between these points is multiplied by a when we apply h_a .

This technique, called *coupling*, allows us to compare Poisson processes of different density. Briefly said, decreasing the node density by a factor a^2 is equivalent to multiply all the distances by a (see [MR96, pages 28–31] for more details on scaling and coupling).

The physical model

If we introduce now on top of the Poisson process our physical model, we can derive its scaling properties. Let us construct the STIRG for a given node density λ and attenuation function $L(\mathbf{x}) = l(\|\mathbf{x}\|)$, and pick two nodes X_1, X_2 that are directly connected. In other words, (3.2) holds for $i = 1$ and $j = 2$. We apply then h_a to the entire graph. We want to show that the dilated (if $a > 1$) or contracted (if $a < 1$) graph is another STIRG. We want thus to derive a condition on aX_1 and aX_2 of the same form as (3.2). Let us try with a new attenuation function $l'(\cdot)$. We want thus

$$\frac{P_1 l'(\|aX_1 - aX_2\|)}{N_0 + \gamma \sum_{k \neq 1,2} P_k l'(\|aX_k - aX_2\|)} > \beta.$$

This condition is always true if we choose l' such that

$$l'(ax) = l(x).$$

Consequently, the STIRG $\mathcal{H}(\lambda, \gamma)$ with attenuation function $l(x)$ is equivalent by dilation to $\mathcal{H}'(\lambda/a^2, \gamma)$ with $l'(x) = l(x/a)$.

Scale-free attenuation functions

We saw that rescaling the STIRG model amounts to modify the attenuation function. We now look for a specific category of these functions, such that if there was no background noise ($N_0 = 0$), the STIRG model would be invariant with respect to h_a . In other words, if Node 1 sends to Node 2, and Node 3 sends to Node 4, we would like to have for all $a, x, y > 0$

$$\frac{l(\|aX_1 - aX_2\|)}{l(\|X_1 - X_2\|)} = \frac{l(\|aX_3 - aX_4\|)}{l(\|X_3 - X_4\|)},$$

or equivalently

$$\frac{l(ax)}{l(x)} = \frac{l(ay)}{l(y)}.$$

We show in Appendix A.3 that the only class of functions with this property happens to be the commonly used power law function $l(x) = cx^{-\alpha}$. This implies that the model acquires very specific symmetry properties with respect to scaling when this class of attenuation functions is considered. In the scaled graph, we have

$$\frac{P_1 a^{-\alpha} \|X_1 - X_2\|^{-\alpha}}{N_0 + \gamma \sum_{k \neq 1, 2} P_k a^{-\alpha} \|X_k - X_2\|^{-\alpha}} > \beta,$$

and thus

$$\frac{P_1 \|X_1 - X_2\|^{-\alpha}}{a^\alpha N_0 + \gamma \sum_{k \neq 1, 2} P_k \|X_k - X_2\|^{-\alpha}} > \beta. \quad (3.13)$$

Rescaling the network in this case is therefore equivalent to multiply the background noise N_0 by a factor a^α . The attenuation function remains unchanged.

3.3.2 Power law attenuation functions

As we saw in Section 3.3.1, when $l(x) = x^{-\alpha}$, the model acquires new scaling properties. In particular, increasing the node density by a factor a^2 is equivalent to multiply the background noise by a factor $a^{-\alpha}$. We show in this section that the critical value $\gamma^*(\lambda)$ is an increasing function of the node density λ in this particular case.

Theorem 3.3 *For an attenuation function $l(x) = x^{-\alpha}$, $\alpha > 2$, the critical threshold $\gamma^*(\lambda)$ for percolation of the STIRG model, if it exists, is an increasing function of the node density.*

Proof: Let λ_1 and λ_2 be two node densities such that $\lambda_1 \leq \lambda_2$. We consider a realization of $\mathcal{H}(\lambda_1, \gamma)$, and pick an edge of the graph, say between Node i and

Node j . We define also $a = \sqrt{\lambda_1/\lambda_2} \leq 1$. The SINR at Node j is

$$\begin{aligned} \beta_{ij} &= \frac{Pl(\|X_i - X_j\|)}{N_0 + \gamma \sum_{k \neq i,j} Pl(\|X_k - X_j\|)} \\ &= \frac{Pa^\alpha l(\|X_i - X_j\|)}{a^\alpha N_0 + \gamma \sum_{k \neq i,j} Pa^\alpha l(\|X_k - X_j\|)} \\ &\leq \frac{Pl(a\|X_i - X_j\|)}{N_0 + \gamma \sum_{k \neq i,j} Pl(a\|X_k - X_j\|)}. \end{aligned} \quad (3.14)$$

On the other hand, if we apply the contraction h_a to the graph, we obtain a node density $a^2\lambda_1 = \lambda_2$. The contracted graph is thus the corresponding realization of $\mathcal{H}(\lambda_2, \gamma)$. The condition for the existence of a link between Node i and Node j in $\mathcal{H}(\lambda_2, \gamma)$ is

$$\beta'_{ij} = \frac{Pl(\|X'_i - X'_j\|)}{N_0 + \gamma \sum_{k \neq i,j} Pl(\|X'_k - X'_j\|)} > \beta,$$

where the coordinates X'_k are measured in the contracted space. We have thus for any node k

$$X'_k = aX_k.$$

Therefore, because of (3.14),

$$\beta'_{ij} = \frac{Pl(a\|X_i - X_j\|)}{N_0 + \gamma \sum_{k \neq i,j} Pl(a\|X_k - X_j\|)} > \beta_{ij} > \beta.$$

This proves that any link in $\mathcal{H}(\lambda_1, \gamma)$ also exists in $\mathcal{H}(\lambda_2, \gamma)$. Super-criticality of the first implies thus super-criticality of the second. However, we know that $\mathcal{H}(\lambda_1, \gamma)$ is super-critical if $\gamma < \gamma^*(\lambda_1)$. Therefore

$$\gamma < \gamma^*(\lambda_1) \Rightarrow \gamma < \gamma^*(\lambda_2).$$

We thus proved that $\gamma^*(\lambda_1) \leq \gamma^*(\lambda_2)$. \square

Theorem 3.3 is quite a surprising result: even though all nodes are emitting at the same time and thus interfere more and more, increasing the node density helps for connectivity. This result is counter-intuitive, and is due to the very particular shape of the attenuation function. In fact, as $l(x)$ tends to infinity when x becomes small, the signal level remains comparable to the interference level when the node density increases.

The particularity of the power law attenuation function also affects the capacity of the network. The case of a finite area network with increasing node density has been studied in [GK00], with this kind of attenuation function (in the physical model). The result is that the total transport capacity of the network grows at most like \sqrt{n} . This is a negative result, but we will see in Section 4.2.2 that the transport capacity remains constant when the node density increases if the attenuation function is bounded.

We conclude that power law attenuation functions lead to scalable connectivity and “half-scalable” transport capacity in the case of dense ad hoc networks. This is due to their invariance property with respect to dilation.

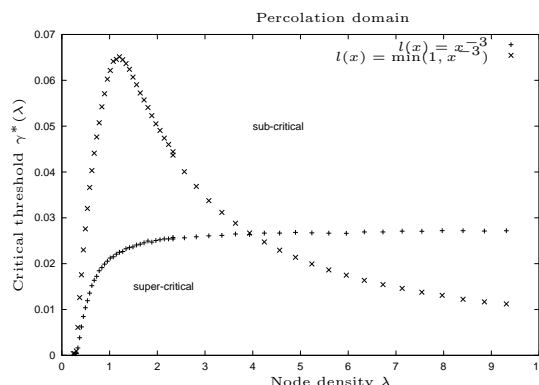


Figure 3.9: Percolation domain computed by simulation, with $l(x) = x^{-3}$ and $l(x) = \min(1, x^{-3})$. The parameters are $P = 100,000$, $\beta = 1$, and $N_0 = 10,000$

3.3.3 Bounded attenuation function

In this section, we address the case where the attenuation function is uniformly bounded, namely:

$$l(x) \leq 1, \quad \forall x \in \mathbb{R}^+.$$

When the node density becomes high, the average level of interference in the network becomes also high (its average grows linearly with λ). If we assume that the attenuation function is bounded, we see intuitively that communication becomes impossible. One can observe that directly in the expression of the signal-to-noise ratio, whose numerator is always bounded, while the denominator diverges:

$$\frac{\overbrace{Pl(\|X_i - X_j\|)}^{\text{bounded}}}{N_0 + \gamma \underbrace{\sum_{k \neq i, j} Pl(\|X_k - X_j\|)}_{\text{unbounded}}}.$$

We consequently guess that the critical value $\gamma^*(\lambda)$ has to decrease to zero when the node density λ increases. Simulation results confirm this intuition, as shown in Figure 3.9. It is interesting to notice that for moderate densities, the bounded attenuation function leads to a much higher critical threshold than the power law. This is due to the fact that the boundedness at the origin also reduces the interference term.

To see more precisely what happens, let us look at a simple example (depicted in Figure 3.10), with attenuation function $l(x) = \min(1, x^{-\alpha})$: consider a node (say Node 0) and assume that it is connected to its two first neighbors (Nodes 1 and 2). We look at the SINR available from the third neighbor (Node 3). When Node 0 listens to Node 3, most of the interference comes from Nodes 1 and 2. If Node 3 is located at a distance greater than 1 from Node 0, the strength of its signal at Node 0 is the same as with $l(x) = x^{-\alpha}$. However, if at the same time Nodes 1 and 2 are located at a distance smaller than one (this situation happens with a reasonable probability when λ has a suitable value),




Figure 3.10: A situation where the truncated attenuation performs better than the unbounded one.

their signals (taken as interference here) have a magnitude bounded by 1. Thus, Node 3 has more chance to be connected, than with an unbounded attenuation function.

We will now study rigorously the asymptotic behavior of $\gamma^*(\lambda)$ when λ goes to infinity. We start by finding an upper bound and show that it tends to zero when λ increases.

Upper bound

Theorem 3.4 *If the attenuation function fulfills the same assumptions as in Theorem 3.2, there exists a constant $c_1 < \infty$ such that*

$$\gamma^*(\lambda) < \frac{c_1}{\lambda}. \quad (3.15)$$

To prove this theorem, we will need the following simple lemma:

Lemma 3.2 *Let X be a Poisson random variable of parameter μ , and $0 < \varepsilon < 1$ a positive constant. Then*

$$\lim_{\mu \rightarrow \infty} \mathbb{P}[X \leq (1 - \varepsilon)\mu] = 0.$$

Proof: Using Chebycheff inequality:

$$\mathbb{P}[|X - \mu| \geq \varepsilon\mu] \leq \frac{\text{Var}(X)}{\varepsilon^2\mu^2} = \frac{1}{\varepsilon^2\mu}.$$

Thus

$$\lim_{\mu \rightarrow \infty} \mathbb{P}[|X - \mu| \geq \varepsilon\mu] = 0,$$

which implies

$$\lim_{\mu \rightarrow \infty} \mathbb{P}[X \leq (1 - \varepsilon)\mu] = 0. \quad \square$$

We also need to make the following observation: since $l(0) > \beta N_0/P$ by assumption, and since l is continuous, there exists a real number $\delta > 0$ such that

$$l(x) \geq \frac{\beta N_0}{P} \quad \forall x \leq \delta. \quad (3.16)$$




Figure 3.11: Mapping of the continuous STIRG model onto a site percolation model

We then map our network onto a lattice. We proceed as follows: we divide \mathbb{R}^2 into squares of edge length $\delta/2$, as depicted in Figure 3.11. The centers of these squares form a lattice \mathcal{L}'' . We assume furthermore, without loss of generality, that the origin of \mathbb{R}^2 coincides with the origin of the lattice.

Because of the boundedness of the attenuation function, if there are n nodes in a square, the sum of the interferences anywhere in the square is greater than nP . Lemma 3.3 then follows:

Lemma 3.3 *If there are more than*

$$N' = \frac{(1 + 2\beta\gamma)P}{\beta^2\gamma N_0} \quad (3.17)$$

nodes inside a square, all nodes in this square are isolated.

Proof: Pick any node i inside the square, and another node j (inside or outside the square). As $L(\cdot)$ is bounded from above by 1, we have

$$PL(X_j - X_i) \leq P.$$

Because of (3.16), we also have

$$\begin{aligned} \sum_{k \neq i, j} PL(X_k - X_i) &\geq \sum_{k \text{ in the sq., } k \neq i, j} PL(X_k - X_i) \\ &\geq \sum_{k \text{ in the sq.}} PL(X_k - X_i) - 2P \\ &\geq N'P \frac{\beta N_0}{P} - 2P \\ &= \beta N' N_0 - 2P \end{aligned}$$

Therefore we have:

$$\begin{aligned} \frac{PL(X_j - X_i)}{N_0 + \gamma \sum PL(X_k - X_i)} &\leq \frac{P}{N_0 + \gamma(\beta N' N_0 - 2P)} \\ &\leq \frac{P}{\gamma(\beta N' N_0 - 2P)} \end{aligned}$$

The above expression is clearly smaller than β when $N' > (1 + 2\beta\gamma)P/\beta^2\gamma N_0$, which implies that Node i is isolated. \square

We can now define a site percolation model by declaring a square *open* if it contains at most $2N' = 2(1 + 2\beta\gamma)P/\beta^2\gamma N_0$ nodes. It is declared *closed* otherwise. It is clear that each square is open or closed independently from the others. Therefore, the origin is a.s. surrounded by a closed circuit (i.e. a circuit formed by closed squares) if:

$$\mathbb{P}(\text{a square is closed}) > p_{\text{site}}, \quad (3.18)$$

where p_{site} is the critical site percolation threshold, whose value is around 0.59 [Gri99, page 56]. The number of nodes inside a square is a Poisson random variable of parameter $\lambda\delta^2/4$. Lemma 3.2 implies that if

$$2N' \leq (1 - \varepsilon)\lambda\delta^2/4, \quad (3.19)$$

we have

$$\lim_{\lambda \rightarrow \infty} \mathbb{P}(\text{a square is closed}) = 1,$$

which means that above a certain value of λ , Inequality (3.18) holds.

Inequality (3.19) is verified if:

$$\frac{2(1 + 2\beta\gamma)P}{\beta^2\gamma N_0} \leq \frac{(1 - \varepsilon)\lambda\delta^2}{4}, \quad (3.20)$$

which can be recast as

$$\gamma \geq \frac{8P}{\beta[(1 - \varepsilon)\beta\lambda\delta^2 N_0 - 16P]}.$$

When $\lambda \geq 16P/\beta\varepsilon\delta^2 N_0$, a sufficient condition is

$$\gamma \geq \frac{8P}{(1 - 2\varepsilon)\beta^2\lambda\delta^2 N_0} \doteq \frac{c_2}{\lambda}. \quad (3.21)$$

We thus proved that for sufficiently high densities, if $\gamma \geq c_1/\lambda$, the origin is a.s. surrounded by a closed circuit in the discrete model. We now have to prove that in this case, the origin belongs to a finite cluster in the continuous model.

Because of Lemma 3.3, when a site is closed, the square centered on this site contains only isolated nodes. Therefore, in the continuous model, when $\gamma \geq c_1/\lambda$, the origin is surrounded by a chain of closed squares with no link inside. To make sure that the origin belongs to a finite cluster, we have to prove that no link can cross this chain.

Let us consider two nodes i and j , such that Node i is located inside an open square surrounded by the chain, and Node j is also located inside an open square, but on the other side of the chain. As these nodes are separated by the chain of closed squares, the distance between them $q \doteq \|X_i - X_j\|$ is larger than $\delta/2$.

We consider two cases. First, we assume that $\delta/2 < q < \delta$. In this case we construct the disk D_1 of radius δ centered on X_i and the disk D_2 of radius δ centered on X_j , as depicted in Figure 3.12. As the chain of closed squares




Figure 3.12: The chain of closed squares separating the two nodes.

separates X_i and X_j , there exists at least one closed square Q that has a non-empty intersection with the segment $[X_i, X_j]$. Moreover, the shortest distance between $[X_i, X_j]$ and $\mathbb{R}^2 \setminus (D_1 \cup D_2)$ is

$$\sqrt{\delta^2 - \frac{q^2}{4}} \geq \frac{\sqrt{3}}{2} \delta.$$

As the diagonal of the square Q has length $\delta/\sqrt{2}$, Q cannot have a non-empty intersection with $[X_i, X_j]$ and with $\mathbb{R}^2 \setminus (D_1 \cup D_2)$ at the same time. Therefore $Q \subset D_1 \cup D_2$.

Furthermore, we count the number of nodes inside three different subsets of Q :

$$\begin{aligned} N_1 &= \Phi_\lambda(Q \cap (D_1 \setminus D_2)) \\ N_2 &= \Phi_\lambda(Q \cap (D_2 \setminus D_1)) \\ N_3 &= \Phi_\lambda(Q \cap D_1 \cap D_2). \end{aligned}$$

As Q is a closed square, we have by assumption $N_1 + N_2 + N_3 \geq 2N'$. This implies that either

$$N_1 + N_3 \geq N'$$

or

$$N_2 + N_3 \geq N'.$$

Let us assume without loss of generality that the first inequality holds. There are thus at least N' nodes located inside D_1 . As D_1 has radius δ and because of (3.16), the signal received by Node i from each of these nodes is at least $P\beta N_0/P = \beta N_0$. The SINR at Node i received from Node j is thus upper-bounded by

$$\beta_{ji} \leq \frac{P}{N_0 + \gamma N' \beta N_0}.$$

Plugging the value of N' into this expression, we verify that

$$\beta_{ji} \leq \beta,$$

which means that no link between Node i and Node j exists. The same is true if $N_2 + N_3 \geq N'$.

Let us now address the case where $q > \delta$ (the case $q = \delta$ appears with probability zero). In this case, we draw the same disks D_1 and D_2 , but with radius q instead of δ . There exists at least one square Q of the chain such that $Q \subset D_1 \cup D_2$. We define N_1 , N_2 and N_3 in the same way as above. Thus, either $N_1 + N_3 \geq N'$ or $N_2 + N_3 \geq N'$.

Let us assume without loss of generality that $N_1 + N_3 \geq N'$. This implies that there are at least N' nodes inside D_1 . Node j is, by construction, on the border of D_1 . Therefore, all these nodes are closer to Node i than Node j . As we assumed that $l(x)$ is decreasing, the SINR at i from Node j is bounded above by

$$\beta_{ji} \leq \frac{Pl(q)}{N_0 + \gamma PN'l(q)} \leq \frac{1}{\gamma N'}.$$

From (3.16) and (3.17), we verify that

$$N' > \frac{1}{\beta\gamma},$$

and therefore

$$\beta_{ji} \leq \beta,$$

meaning that the link cannot exist.

Consequently, we have proved that if the origin is surrounded by a chain of closed squares in the discrete model, then the continuous model is sub-critical. We conclude that when (3.21) holds, the network is sub-critical. We have thus obtained an upper bound on the critical value $\gamma^*(\lambda)$ by proving that if $\gamma \geq c_1/\lambda$, the origin belongs a.s. to a finite cluster.

Lower bound

To obtain a lower bound to the percolation threshold, we review the proof in Section 3.2.3. First we note that in Equation (3.12), γ is defined as inverse proportional to the parameter M . Second, M has been chosen so that p_B is greater than $1 - 4/9^4$. We observe in Equation (3.11) that it is enough to let M grow linearly with λ to keep p_B constant. Therefore, the value of γ used in the proof can be expressed with the form c_2/λ . As the model percolates for this value of γ , we have a valid lower bound to the critical $\gamma^*(\lambda)$. In conclusion, there exists $c_2 > 0$ such that

$$\gamma^*(\lambda) \geq \frac{c_2}{\lambda}.$$

Summary

As both upper and lower bounds on the critical threshold have this form, we have obtained the asymptotic behavior of the threshold $\gamma^*(\lambda)$ for $\lambda \rightarrow \infty$.

Theorem 3.5 *For λ tending to infinity, the critical value of γ has the following asymptotic behavior:*

$$\gamma^*(\lambda) = \Theta\left(\frac{1}{\lambda}\right).$$

Figure 3.13 illustrates the typical shape of the function $\gamma^*(\lambda)$, and summarizes the known bounds on this function.




Figure 3.13: Illustration of the bounds on the super-critical domain

3.4 A TDMA approach

We can conclude, from the previous sections, that unless γ can be made sufficiently small, long-range communications are impossible if we allow all nodes to emit simultaneously, because the graph $\mathcal{H}(\gamma, \lambda)$ may remain in a sub-critical phase for all λ . Having a small γ requires nodes to use CDMA for transmission, which can be complex to implement in an ad-hoc network (node synchronization may be difficult in the presence of mobility). An alternative is to avoid having all nodes emitting at the same time, and thus to use a TDMA scheme. We assume that each time interval is divided into t time slots. An optimal TDMA scheme poses also a quite complex challenge to assign the slots to each node, which is clearly beyond the scope of this work. In this section, we keep the strategy sub-optimal but very simple and totally decentralized: each node picks randomly a number i between 1 and t , and only emits during the i -th time slot. All nodes are listening at all times. We also assume, for the sake of simplicity, that all nodes emit with the same power P . We denote by $\mathcal{H}^t(\lambda, \gamma)$ the graph obtained by superposing the t graphs derived for each slot.

In Appendix A.2, we show that for bounded support attenuation functions and fixed γ , one can reach the supercritical phase by choosing t large enough (Corollary A.2). We conjecture that the same is true for attenuation functions that satisfy assumptions [1–4] of Section 3.2.3. To prove this conjecture, we need to modify the proof of Theorem 3.2 to take the time slots into account. We leave this proof for future work.

By applying this TDMA strategy, we actually reduce the number of interfering nodes by a factor t . It is therefore interesting to compare the connectivity of the graph obtained by superposing the t graphs derived for each slot, to that of the original graph $\mathcal{H}(\gamma, \lambda)$ obtained when all nodes emit at the same time. Let us introduce the following notations for the interference contribution, which is another shot noise, at each time slot:

$$J_i(\mathbf{x}) = \sum_{k \in S_i} PL(X_k - \mathbf{x})$$

where S_i , $i = 1, \dots, t$ is the set of the indices of the nodes that emit during the

i -th time slot. It follows immediately that

$$\sum_{i=1}^t J_i(\mathbf{x}) = J(\mathbf{x}).$$

The expected value of the interference term of (3.2) in the TDMA scheme is t times lower than in the regular scheme

$$\mathbb{E}[\gamma J_i(\mathbf{x})] = \frac{1}{t} \mathbb{E}[\gamma J(\mathbf{x})]$$

We computed by simulation the critical threshold $\gamma^*(\lambda)$ for this scheme. Figure 3.14 presents the results, compared to those of the regular scheme. As expected, we observe that the threshold in the TDMA scheme is about t times higher. This means that introducing an t -slots TDMA system is somehow equivalent to dividing γ by t .

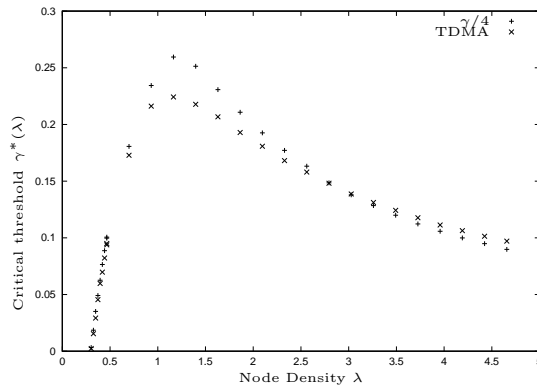


Figure 3.14: Comparison between the critical threshold in the TDMA case ($t = 4$ time slots) and in the original model with all nodes allowed to emit simultaneously. To make comparison easier, the critical value in the second case has been multiplied by 4. [Simulation parameters are the same as in Figure 3.3]

3.4.1 Asymptotic behavior

In this section, we compute a lower bound on the necessary number of time slots $t^*(\lambda)$ to restore connectivity when $\gamma = 1$ (no CDMA).

Theorem 3.6 *If the attenuation function verifies the same assumptions as in Theorem 3.2, there exists a constant $c_3 > 0$ such that*

$$t^*(\lambda) > \lambda c_3.$$

The derivation of this lower bound is very similar to that of the Section 3.3.3. We start thus with the same assumptions on $l(\cdot)$. As each node chooses its category randomly between 1 and t , the process formed by the nodes of a particular category is still a Poisson point process, but with intensity λ/t .

We divide the plane into squares in the same way as above, i.e. as depicted in Figure 3.11. We then construct a site percolation model on this grid by

declaring a site of \mathbb{Z}^2 *closed* if the square centered on this site contain at least $2N' = 2(1+2\beta)M/\beta m$ nodes of each category. Thus, if a site is declared closed, during each time slot, according to Lemma 3.3 (with $\gamma = 1$), the square is full of isolated nodes.

We compute now the probability that a site is closed. The number of nodes of each category in a square is a Poisson random variable of parameter $\delta^2\lambda/2t$. Let us call N_i the variable counting the number of nodes of category i in the square. We have

$$\begin{aligned} \mathbb{P}(\text{a site is closed}) &= \mathbb{P}\left(\bigcap_{i=1}^t \{N_i > 2N'\}\right) \\ &= \prod_{i=1}^t \mathbb{P}(N_i > 2N') \\ &= \mathbb{P}^t(N_i > 2N'). \end{aligned}$$

Therefore, if

$$\mathbb{P}(N_i > 2N') > \sqrt[t]{1 - p_{site}}, \quad (3.22)$$

where $p_{site} \cong 0.59$ is the site percolation threshold on \mathbb{Z}^2 , we obtain

$$\mathbb{P}(\text{a site is open}) < p_{site}.$$

For having (3.22), the intensity of the Poisson random variable N_i has to be large enough. Let μ be the minimal intensity of a Poisson random variable X_μ such that $\mathbb{P}(X_\mu > 2N') \geq \sqrt[t]{1 - p_{site}}$. Then a sufficient condition for having (3.22) is

$$\frac{\delta^2\lambda}{4t} > \mu$$

or equivalently

$$t < \frac{\delta^2\lambda\mu}{4}.$$

We obtained thus the bound on $t^*(\lambda)$ by proving that if

$$t < \lambda c_3$$

for some fixed constant $c_3 = \delta^2\mu/4$, the network is almost surely disconnected. This proves Theorem 3.6.

As each category has a node density λ/t , and as $t > \lambda c_3$, we conclude that the node density of the nodes allowed to emit at each time slot is uniformly bounded above by c_3 , independently from the number of nodes. This means that good connectivity can be achieved, but that nodes can emit less and less frequently. This result suggests that the rate throughput per node decreases like $1/\lambda$.

However, since our TDMA scheme is very simple, our negative result might be possibly due to its sub-optimality. In Section 4.2.2, we study the transport capacity of dense networks with bounded attenuation function in a much more general context. We will see that the available rate per node has to decrease like $1/\lambda$, no matter what scheduling or power allocation is used.

3.5 Conclusion

The STIRG $\mathcal{H}(\gamma, \lambda)$ is radically different from the graph obtained in the Boolean Model, where $\gamma = 0$: the node degree is now bounded (Theorem 3.1), and the existence of an edge between two nodes depends not only on the location of these two nodes, but on the location of all others. We showed that if γ is too large, all clusters are almost surely finite. Our main result is that percolation, and thus long range communications, are however still possible if γ is small enough, but non zero, as soon as the node density is greater than the critical threshold of the Boolean model (Theorem 3.2). If this had not been the case, it would have been a serious impediment for multiple hops large scale ad hoc networks.

We also proved that when the node density λ tends to infinity, the critical value $\gamma^*(\lambda)$ decreases as $1/\lambda$ provided the attenuation function is bounded from above and from below in a small neighborhood of the origin. The main result of this chapter is a first picture of the shape of the region in the (λ, γ) plane where percolation occurs.

This asymptotic behavior of the critical value $\gamma^*(\lambda)$ has a very important meaning: if the bandwidth is limited, decreasing γ means decreasing the throughput on the links. Therefore, when the node density increases, one has to decrease the throughput in order to maintain connectivity. Good connectivity and high throughput are conflicting goals in wireless networks. We will see in the next two chapters that this opposition is a fundamental property of wireless networks, and study the involved trade-offs.

As a small value of γ requires very efficient and thus complex CDMA codes, an alternative is to use a TDMA system, where each node emits during 1 slot every t time slots. We showed that such a system led to a connectivity similar to the original scheme with an orthogonality factor γ/t .

Part II

Throughput

Chapter 4

Connectivity vs Throughput

In this chapter, we look at scaling laws for large finite networks, from the point of view of connectivity and throughput. We show that under many circumstances, these two properties have contradictory requirements.

Two factors essentially limit the performance of wireless networks: the ambient noise, which limits the capacity of the links, and interferences, which are due to the nodes sharing the same wireless medium. In the two next sections, we study these two limitations separately.

In practice, a wireless network is made of a finite number of nodes scattered on a finite area. To study how their throughput and connectivity scale, we need to make the number of nodes grow to infinity. There are two ways of doing this. We can first keep the network area constant, and increase the node density. We speak then of a *dense network*. Such a network typically suffers more interferences than thermal noise, since nodes get asymptotically very close to each other. Alternatively, we can keep the node density constant, and increase the network area, to create an *extended network*, where the average number of nodes grows proportionally to the area. This kind of network is more difficult to study, as it can be noise or interference limited, depending on the traffic load. On the other hand, the constant node density makes it quite insensitive to the behavior of the attenuation function $L(\cdot)$ near the origin (near field effects), contrary to a dense network.

4.1 Partial vs full connectivity in the Boolean model

In this section, we assume that there is no competition between transmissions in the network, i.e. only one link is active at a time in the whole network. In this case, there is already an obvious trade-off between connectivity and throughput: assuming a bounded emitting power, one has to make a compromise between the length and the capacity of the wireless links. A very general information theoretic model for this problem is presented in Chapter 5. In this section, as an illustration, we discuss briefly the Boolean case, which corresponds to the

case where a minimum SNR is required for connectivity, and that links with very low throughput are ignored.

As we have seen in Section 3.3.1, the Boolean model can always be rescaled. Contrary to the STIRG model of Chapter 3, there is therefore no difference, from the connectivity point of view, between dense and extended networks. We consider thus, without loss of generality, only extended networks in the sequel of this section.

We consider a finite area network, and model the node distribution using a 2-dimensional Poisson point process of intensity λ (λ is thus the average number of nodes per square meter). *Full connectivity* occurs when one can find a path joining any two nodes.

As a first observation, it is clear that the probability that the network is fully connected is always strictly less than one, whenever the diameter of the network area is larger than r . Therefore, full connectivity can only be an asymptotic property, in the sense that this probability can only *tend* to one. Moreover, if one considers the (unrealistic) case where the network area is infinite, then the probability that the network is fully connected is always exactly zero.

However, in the case of an ad hoc network, one can say that a network is still *well connected*, if disconnected nodes may exist but always represent a small fraction of the total number of nodes. We say that a node is *disconnected* if it is not connected to the *majority* of the other nodes. In fact, in the context of wireless multi-hop networks, we would like most nodes to belong to the same huge connected component (which forms the network itself). A disconnected node that is not connected to any other node is called an *isolated node*.

As already mentioned in Section 2.2.2, percolation theory addresses the case where the network area is infinite, and the fundamental result is that if the node density λ and the range r are such that $\pi\lambda r^2 > N^*$, for a special constant $N^* \simeq 4.5$, then the network is indeed formed by a huge connected component (the network), plus a multitude of finite components (disconnected nodes). Moreover, the fraction of connected nodes is a deterministic function θ of the average node degree $\pi\lambda r^2$.

Therefore, this infinite network model is a good approximation for large networks. However, networks are never infinite, and one needs more specific results for the finite case. Penrose and Pisztor [PP96] showed that for a large but finite area, the fraction of connected nodes is always close to the deterministic function $\theta(\pi\lambda r^2)$. We call this *partial connectivity*.

Theorem 4.1 (Penrose and Pisztor 1996) *Let $B(m)$ denote the square $[0, m]^2$ and set $\Phi_{\lambda, m} \doteq \Phi_\lambda \cap B(m)$, a Poisson point process of intensity λ on $B(m)$. Suppose that $\lambda r^2 > N^*$. Let $0 < \varepsilon < \frac{1}{2}$, and let $E(m)$ be the event that (i) there is a unique cluster $C_b(B(m))$ on $\Phi_{\lambda, m}$ containing more than $\varepsilon\lambda\theta(\pi\lambda r^2)m^2$ points of $\Phi_{\lambda, m}$, and (ii)*

$$(1 - \varepsilon)\lambda m^2\theta(\pi\lambda r^2) \leq \text{card}(C_b(B(m)) \cap \Phi_{\lambda, m}) \leq (1 + \varepsilon)\lambda m^2\theta(\pi\lambda r^2).$$

Then there exist constants $c_1 > 0$ and $m_0 > 0$, such that

$$\mathbb{P}[E(m)] \geq 1 - \exp(-c_1 m), \quad m \geq m_0.$$

We can observe in the above theorem that if we let the area of the network m^2 tend to infinity (and thus also the number of nodes), then the fraction of connected nodes tends to the constant $\theta(\pi\lambda r^2)$. This result matches the first percolation result for infinite networks. However, there is a slight difference between the two: in the first one, we consider only an *infinite* network area, whereas in the second one, we consider a sequence of *finite* networks, and derive an *asymptotic* property when the number of nodes tend to infinity.

The same approach can be applied for full connectivity. As we only consider (larger and larger but still) finite networks, full connectivity can happen with positive probability. This approach is very frequent in the literature, and the following result has been proved first by Penrose in 1997 [Pen97].

Theorem 4.2 (Penrose 1997) *Let M_n denote the length of the longest edge of the minimal spanning tree connecting all nodes in the network. Then*

$$\lim_{n \rightarrow \infty} \mathbb{P}[\pi M_n^2 - \log n \leq \alpha] = \exp(e^{-\alpha}), \quad \forall \alpha \in \mathbb{R}.$$

It is now interesting to observe how these two results, namely the one about partial connectivity (Theorem 4.1) and the one about full connectivity (Theorem 4.2), are related. We will start with an intuitive – but not rigorous – reasoning, that allows to link the probability that the network is fully connected with the function $\theta(\pi\lambda r^2)$. A rigorous and detailed version of this reasoning can be found in [GK98].

We assume without loss of generality that nodes are distributed according to a Poisson point process of unit density ($\lambda = 1$), and that the network area increases and is equal to n . The total number of nodes in the network is therefore approximately equal to n . We want to compute the critical range $r(n)$ to keep the network fully connected.

We start with the simple observation that the network is fully connected if and only if no node is disconnected. So we obtain immediately a lower bound on the probability that the network is connected:

$$\mathbb{P}(\text{the network is connected}) \geq 1 - n\mathbb{P}(\text{node } i \text{ is disconnected}).$$

Furthermore, we know from Theorem 4.1 that the probability that a given node is isolated tends to $1 - \theta(\pi\lambda r^2)$ when the network area tends to infinity. Therefore

$$\mathbb{P}[\text{the network is connected}] \geq 1 - n(1 - \theta(\pi\lambda r^2)). \quad (4.1)$$

In order to have the above probability tend to one when n tends to infinity, we must have $\theta(\pi\lambda r^2) \rightarrow 1$. This is only possible if $\pi\lambda r^2 = \pi\lambda r^2(n)$ grows with n .

To compute how fast $r(n)$ should grow, we derive an approximation of the function $\theta(\pi\lambda r^2)$ when $\pi\lambda r^2$ is large. According to Propositions 6.4-6.6 in [MR96], when the average node degree is large, the ratio between the probability that a node is disconnected, and the probability that a node is isolated (i.e. has degree zero) tends to one. Therefore, asymptotically, all disconnected nodes are isolated nodes. The probability that a node is isolated is easy to compute:

$$\mathbb{P}[\text{a given node is isolated}] = \exp(-\pi\lambda r^2).$$

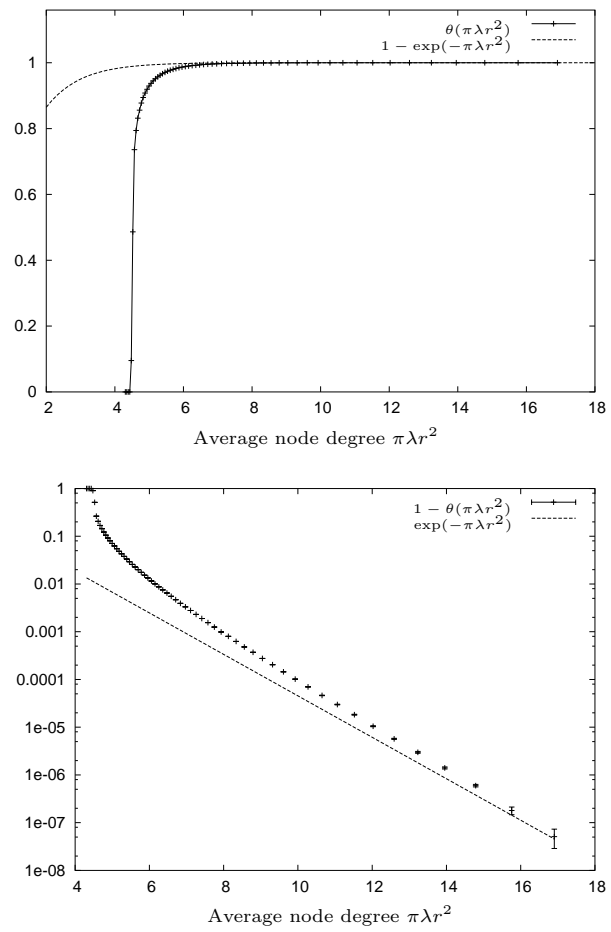


Figure 4.1: The fraction of disconnected nodes in the Boolean model.

Therefore, when $\pi\lambda r^2$ is large, $\theta(\pi\lambda r^2) \simeq 1 - \exp(-\pi\lambda r^2)$. Figure 4.1 shows a simulation based evaluation of $\theta(\pi\lambda r^2)$ for large values of $\pi\lambda r^2$.

Using this approximation in (4.1), we obtain

$$\mathbb{P}[\text{the network is connected}] \geq 1 - n \exp(-\pi\lambda r^2(n))$$

The latter expression tends to one provided $n \exp(-\pi\lambda r^2(n))$ tends to zero. Taking the logarithm of this expression leads to the conclusion that the network is asymptotically connected if

$$\pi\lambda r^2(n) - \log n \rightarrow \infty.$$

In fact, the lower bound (4.1) is asymptotically tight, as when the number of nodes becomes large, the events that different nodes are connected become almost independent. Therefore we have

$$\mathbb{P}[\text{the network is connected}] \simeq \theta^n(\pi\lambda r^2),$$

making the above condition necessary and sufficient.

This means that the average node degree must grow approximately like $\log n$, when the number of nodes in the network increases. However, it appears in this intuitive derivation of the critical range for full connectivity that the *most isolated node* is determining the result. In fact, the network becomes connected when the last node joins the network. This means that full connectivity is not really a global property of the network; it just answers the question “how isolated is the most isolated node”. As we just saw above, the distance to the first neighbor of the most isolated node increases with the number of nodes. But this is a pure statistical effect: we are taking a set of randomly distributed distances, and pick the largest one. As we increase the sample set, the largest element becomes longer and longer. This explains why the range of the nodes has to increase, even though the node density remains constant.

This behavior is not only observed in the Boolean model, but also in large noise-limited networks. We will see in Chapter 5 that the same results hold for a much more general model.

4.2 Taking interferences into account

In this section, we discuss the case where several links compete for the shared medium. The limitations of the network due to this factor are very difficult to evaluate in general. However, more can be said if we consider specific communication models, rather than the general problem. Here, we discuss this issue in the context of the STIRG model.

4.2.1 Extended networks

We consider a finite portion of the STIRG model. Although it was not formally proved yet, we can conjecture that a similar theorem as Theorem 4.1 exists for this model, and that for a finite network, the fraction of connected nodes tend to the percolation probability when the network area grows. However, so far, we did not discuss this probability in the context of the STIRG model, but only the




Figure 4.2: Effective throughput across a piece of network, computed as the number of vertex disjoint paths times γ . The maximum is reached for $\gamma \simeq 0.25$. The simulation was run on a 65536×65536 domain, with 1,000,000 points. We took a power law attenuation function with exponent 3, and the following parameters: $P_i = 1$, $N_0 = 1.8 \cdot 10^{-4}$, $\beta = 2.5 \cdot 10^{-2}$

percolation threshold. Let us start with some observations on the percolation probability for this model, which we will denote here by $\theta_{\text{STIRG}}(\gamma, \lambda)$.

Since we assume the node density fixed (extended network), the only parameter that can be tuned is γ . Clearly, $\theta_{\text{STIRG}}(\gamma, \lambda)$ is a decreasing function of γ , and we have

$$\theta_{\text{STIRG}}(0, \lambda) = \theta(\lambda\pi r_b^2),$$

where $\theta(\cdot)$ denotes the percolation probability of the corresponding Boolean model, and $2r_b = l^{-1}(\beta N_0/P)$ (see Section 3.2.1). It appears thus clearly that the only way to improve the fraction of connected nodes is to reduce γ , until the network becomes noise-limited, and nothing more can be done with interferences. The price to pay for reducing γ is throughput. Indeed, one needs longer codewords to make them more orthogonal in a CDMA system. In fact, the throughput of the system decreases linearly with γ .

We end up thus with the following situation: by reducing γ , we reduce the throughput on the links, but make the network more connected, i.e. add more links. As we have seen in Section 3.2, if γ is too large, the network is sub-critical, and is therefore unable to transport packet over long distances (zero throughput). On the other hand, if γ is very small, the rate on the links is also very small, but the number of available links is bounded from above by the Boolean connectivity graph that corresponds to the case $\gamma = 0$. Therefore, the effective throughput also goes to zero. From these two extremes, we conclude that there must be an optimal value for γ , that optimizes the density and rate of the links.

A simulation of this trade-off is presented in Figure 4.2. In this simulation, we counted the number of vertex disjoint paths from the left border to the right border of a square simulation domain, and multiplied it by γ , to obtain a value proportional to the effective throughput across the domain.

The *transport capacity* of a network is defined as the sum over all links of the

product of their length and their rate. If we compute it in our STIRG model, we end up with a constant spacial density of transport capacity, since the model is stationary and ergodic. Therefore, the total transport capacity simply grows linearly with the network area (and thus with the average number of nodes).

This transport capacity of order $\Theta(n)$ (where n denotes the number of nodes) for extended networks does only give an upper bound to the amount of information that can be carried in this point-to-point relaying fashion. However, given a certain traffic demand, the available links may not be positioned at the right place, and paths with sufficient capacity may not be found between sources and destinations. The construction of a lower bound requires thus an explicit routing scheme. This issue is done in Chapter 6 in the case of uniform traffic matrix.

4.2.2 Dense networks

The picture is very different for dense networks. As seen in Section 3.3.1, if we take a power law attenuation function, the network behaves similarly to extended networks. But if we assume the attenuation function to be bounded, which makes sense from a physical point of view, we obtain a very surprising result: the total transport capacity of dense ad hoc networks is bounded above by a constant, independently from the power allocation strategy and number of nodes. This means that the actual throughput in the whole network is of the same order as if we let each node transmit one after the other.

To prove this result, we consider a compact region S of \mathbb{R}^2 , in which the network is confined. We introduce furthermore the following technical assumption on the attenuation function:

$$l(\|x - y\|) > 0 \quad \forall x, y \in S. \quad (4.2)$$

This condition is always verified if the attenuation is decreasing and has unbounded support.

Theorem 4.3 *If the attenuation function is bounded and verifies (4.2), the total transport capacity of the network is bounded above by a constant, independently of the number of nodes.*

Proof: As S is a compact set, we have

$$m \doteq \inf_{x, y \in S} l(\|x - y\|) > 0.$$

Hereafter, we will denote by P_i the actual power emitted by Node i , and by P_{tot} the total power emitted in the network. Hence

$$P_{tot} = \sum_i P_i.$$

Let us pick an arbitrary real number $0 < \varepsilon < 1/2$, and call *dominant* the nodes i such that $P_i/P_{tot} > \varepsilon$. Clearly, there are no more than $1/\varepsilon$ dominant nodes in the network.

Let us now look at the transport capacity. As above, we assume that Node i is transmitting towards a destination $r(i)$ located at a distance δ_i . The available

SINR at Node $r(i)$ is

$$\beta_{i,r(i)} = \frac{P_i l(\delta_i)}{N_0 + \gamma \sum_{k \neq i, r(i)} P_k l(\|X_k - X_{r(i)}\|)}.$$

As $l(\|\mathbf{x} - \mathbf{y}\|) > m$ inside S by assumption, one can bound the sum of the interferences from below

$$\sum_{k \neq i, r(i)} P_k l(\|X_k - X_{r(i)}\|) \geq m(P_{tot} - P_i - P_{r(i)}),$$

and obtain a bound on $\beta_{i,r(i)}$

$$\beta_{i,r(i)} \leq \frac{P_i l(\delta_i)}{N_0 + \gamma m(P_{tot} - P_i - P_{r(i)})}.$$

Next, let us consider two distinct classes of nodes. The first class \mathcal{C}_1 contains all nodes i that are *not* dominant and emit towards a destination $r(i)$ that is *not* dominant. For these nodes, we use the following looser bound

$$\beta_{i,r(i)} \leq \frac{P_i l(\delta_i)}{\gamma m P_{tot} (1 - 2\varepsilon)}.$$

The second class \mathcal{C}_2 contains the nodes that do not belong to \mathcal{C}_1 (i.e. they are either dominant, or emit towards dominant node, or both). For these nodes, we use the following bound:

$$\beta_{i,r(i)} \leq \frac{P_{max} l(\delta_i)}{N_0}.$$

We notice that the second class contains a *finite* number of nodes (this number is smaller than $2/\varepsilon$).

We can now sum the individual transport capacities, and obtain the total capacity of the network. In a first step, we will add separately the transport capacities of the nodes of the first class:

$$\begin{aligned} \sum_{i \in \mathcal{C}_1} \delta_i C(\beta_{i,r(i)}) &\leq \sum_{i \in \mathcal{C}_1} \frac{1}{2} \delta_i \beta_{i,r(i)} \\ &\leq \sum_{i \in \mathcal{C}_1} \delta_i \frac{P_i l(\delta_i)}{2\gamma m P_{tot} (1 - 2\varepsilon)} \\ &\leq \max_i [\delta_i l(\delta_i)] \sum_{i \in \mathcal{C}_1} \frac{P_i}{2\gamma m P_{tot} (1 - 2\varepsilon)} \\ &\leq \max_i [\delta_i l(\delta_i)] \frac{P_{tot}}{2\gamma m P_{tot} (1 - 2\varepsilon)} \\ &= \frac{1}{2\gamma m (1 - 2\varepsilon)} \max_i [\delta_i l(\delta_i)]. \end{aligned}$$

We next add up the transport capacities of the nodes of the second class:

$$\begin{aligned} \sum_{i \in \mathcal{C}_2} \delta_i C(\beta_{i,r(i)}) &\leq \sum_{i \in \mathcal{C}_2} \frac{1}{2} \delta_i \beta_{i,r(i)} \\ &\leq \sum_{i \in \mathcal{C}_2} \frac{P_{max}}{2N_0} \delta_i l(\delta_i) \\ &\leq \frac{P_{max}}{\varepsilon N_0} \max_i [\delta_i l(\delta_i)]. \end{aligned}$$

The total transport capacity is thus bounded above by

$$C_{tot} \leq \left(\frac{P_{max}}{\varepsilon N_0} + \frac{1}{2\gamma m(1-2\varepsilon)} \right) \max_i [\delta_i l(\delta_i)].$$

As $l(\cdot)$ is uniformly bounded and decreases at least as fast as x^{-2} , the function $xl(x)$ is also bounded. This ensures that the above maximum exists and is uniformly bounded, even if the number of nodes goes to infinity. \square

We have therefore proved that the total transport capacity of the network is uniformly bounded. Assuming a uniform traffic matrix for example, this implies that the available bandwidth per node decreases like $1/n$.

4.3 Conclusion

We have seen that in the Boolean model, full connectivity does not scale when the network size increases. On the contrary, if we allow for a (possibly very small) fraction of disconnected nodes, then the range does not need to be adjusted when the number of nodes tend to infinity.

When several flows have to share the available bandwidth and when interferences are critical, full connectivity turns out to be very costly in terms of throughput. In fact, keeping the most isolated nodes connected consumes a lot of resources, and affects greatly the overall performance of the network. This situation leads to a trade-off between capacity and connectivity. Under several models, keeping the connectivity graph quite sparse leads to an optimal throughput.

This trade-off has been observed in Chapter 3 already, where we saw that in dense networks, assuming a bounded attenuation function, one has to decrease the throughput in order to keep the network connected. In this chapter, we proved that this behavior holds with arbitrary power assignment, making this opposition more general. Indeed, we showed that in such networks, under the STIRG model, the total transport capacity is bounded from above by a constant. This result differs greatly from the well known \sqrt{n} transport capacity bound, and is the result of the boundedness of the attenuation function.

In the next chapter, we will show that full connectivity does not scale, even if we release any assumption on the relaying scheme.

Chapter 5

Information theoretic definition of connectivity

We have seen in Section 4.1 that in the context of the Boolean model, requiring partial connectivity instead of full connectivity allows to preserve the scalability of the network. In this chapter, we consider a more information theoretic model where there is only one active source-destination pair at any given time, while the remaining nodes act as possible relays. It has been shown in [GV02] that the transmission rate in dense networks can grow at most as¹ $O(\log n)$, under the assumption that around each of the two active nodes there is a dead zone of finite nonzero radius without any node. For extended networks (and without dead zone assumption) it has been shown that the rate decreases as $O((\log n)^{-d\alpha})$, where $\alpha > 2$ is the exponent of the power attenuation function and $d \in \{1, 2\}$ denotes the dimension of the network [LS04].

The conclusion drawn from all these works is rather pessimistic, since the rate offered to each node always tends asymptotically to zero as the number of nodes grows —except for relay transmission in dense networks [GV02]. On the other hand, the common requirement for all the works mentioned above, is that *every* pair of nodes can be connected at that rate.

In this chapter we show that the price to pay to operate the network at a given rate is precisely its full connectivity. We prove that if we allow an *arbitrary small* fraction of the nodes to be disconnected, then a non-vanishing rate can be achieved in two-dimensional extended networks and in the relay scenario of [LS04]. On the contrary, it turns out that in the one-dimensional case, a non-zero rate is impossible even if we allow an *arbitrary large* fraction of nodes to be disconnected.

We want to spend here a few words on the intuition behind these results. The original result of [GV02] for dense relay networks can be easily seen as an application of the capacity formula for multi-antenna channels: the addition of more nodes in a finite area, each of which is capable of working as a relay transmitting at constant power, improves the transmission rate by a logarithmic

¹We use the following notation throughout the next two chapters: $f = O(g)$ if $\limsup_{n \rightarrow \infty} \frac{f(n)}{g(n)} < +\infty$; $f = \Omega(g)$ if $g = O(f)$; $f = \Theta(g)$ if $f = O(g)$ and $g = O(f)$. Thus all $O(\cdot)$ results are upper bounds, $\Omega(\cdot)$ results are lower bounds and $\Theta(\cdot)$ results are sharp scaling estimates.

factor of the total power. Our result for extended relay networks of constant density can be seen as a consequence of percolation theory in the Boolean model: by choosing the constant density of the nodes λ larger than a critical value λ^* , a giant connected component forms. Inside this component every pair of nodes can communicate at a constant bit rate. As seen in Section 2.2.2, percolation theory tells us even more: this giant component contains a constant fraction of the nodes that can be made arbitrarily close to one by an appropriate choice of λ . This good news is counter-balanced by a corresponding pessimistic result that immediately leads to a corresponding upper bound: in an extended network, no matter how small the rate of transmissions, there will always be a non zero fraction of the nodes that will not be able to communicate to the rest of the network at that rate, even if we allow arbitrary cooperation between the nodes. This is proved by recasting the constraint on a minimal rate from the source to all other nodes, as a constraint on the value of a shot noise at the source location, and by showing that the fraction of nodes verifying this constraint is strictly less than one.

We point out that our bounds tend to zero when the fraction of the nodes required to sustain the given rate tends to one, in agreement with [LS04]; and that they diverge as the density increases, in agreement with [GV02].

5.1 Network model

We define connectivity as follows: we consider two nodes *connected* if one can send data to the other at *rate* $R > 0$ and vice versa (possibly with the help of all other nodes, as we assume that at each instant there is only one node transmitting and one node receiving). We also say in this case that the end to end *throughput* between the two nodes is at least equal to R . The two terms “rate” and “throughput” are thus used interchangeably.

We assume that each node has a maximum emitting power P , and that the power attenuation from Point x to Point y is deterministic, and given by some function $l(\|y - x\|)$, where $\|\cdot\|$ denotes the euclidean norm. We denote by N_0 the power of the background white Gaussian noise added to each received signal (AWGN). Furthermore, for technical reasons (see Section 3.1.2), we assume that l is decreasing and such that

$$\int_y^\infty xl(x)dx < \infty \text{ for } y > 0.$$

We will make no further assumptions on the communication protocol and look at information theoretical bounds on the fraction of nodes that can communicate at rate R with each other.

We consider one-dimensional and two-dimensional random networks. In one dimension, we assume that nodes are scattered according to a Poisson point process of unit intensity over the interval $[0, n]$. The average number of nodes in the network is thus equal to n . Similarly, in two dimensions, we consider a Poisson point process of unit intensity over the square $[0, \sqrt{n}] \times [0, \sqrt{n}]$, so that the average number of nodes is also equal to n .

We look at asymptotic connectivity results when n tends to infinity, while the node density remains constant (extended network). In the following, we

will use the expression *with high probability* (w.h.p.) to qualify an event whose probability tends to one when n tends to infinity.

5.2 One dimensional network

For one-dimensional networks, the results are similar for Boolean and information theoretic connectivity. Namely, for any fixed rate $R > 0$, the fraction of connected nodes tends to zero when n tends to infinity.

If a node is disconnected from the origin, then all nodes further away from the origin are also disconnected. Thus, as the network size increases (and as the node density remains constant), there are w.h.p arbitrarily large gaps in the network, and thus we expect it to be eventually disconnected, for any fixed rate $R > 0$. This intuition matches the result in [LS04], which shows that the rate has to decrease to zero as a function of n .

Now if we require that only a (positive) fraction of the nodes has to be connected, we can obtain the same negative result. Assume that for a given rate R , nodes are connected until node x . Then all nodes in the interval $[x, n]$ are disconnected. In fact, the fraction of connected nodes is thus $x/(n - x)$, which tends to zero when n goes to infinity. Therefore, w.h.p, the fraction of connected nodes is below any positive number.

The idea of requiring only partial connectivity does thus not help, and the same asymptotic result holds. In fact, in one dimension, *partial connectivity* and *full connectivity* are asymptotically equivalent. The picture is definitely not the same in two dimensions, as we will show in the next section.

5.3 Two dimensional network

We look at the bounds on the rate at which a given fraction of the nodes can exchange data with each other. In other words, if we discard a given fraction of the nodes (the worst positioned), what are the bounds on the rate? We will see that discarding the worst nodes (up to a given percentage, that can be arbitrarily small) dramatically changes the asymptotic behavior of the rate, as it then stays constant when n tends to infinity, in sharp contrast to the vanishing rate obtained in one-dimensional networks.

5.3.1 Lower bound

We construct an explicit scheme that achieves a constant rate, for an arbitrary (but smaller than one) fraction of the nodes. The following Theorem gives the rigorous formulation of our result.

Theorem 5.1 *For any $0 < \check{\theta} < 1$, there exists a rate $R > 0$ independent of n , such that there exists a subset of nodes of size $n\check{\theta}$ in which each node can send data to any other node at rate R w.h.p.*

To prove Theorem 5.1, we need the two following lemmas (5.1 and 5.2). In the first one, we consider the Boolean model $\mathcal{G}(1, r)$, whose node density is equal to one, and where nodes have a connectivity range equal to r . We compute an achievable throughput along the shortest path between two nodes that belong




Figure 5.1: A shortest path from X to Y in a Boolean model. Filled balls do not overlap, and neither do the empty balls.

to the same cluster. We show that this throughput only depend on r , and not on the number of hops. Thus, the same throughput is achievable between arbitrary distant nodes, if the Boolean model is super-critical.

Lemma 5.1 *In a Boolean model $\mathcal{G}(1, r)$, the following throughput is achievable along the shortest path between any two nodes of the same cluster:*

$$R = \frac{1}{6} \log \left(1 + \frac{Pl(r)}{N_0 + P \sum_{k=1}^{\infty} 6(k+1)l(kr)} \right).$$

Proof: We look first at the properties of shortest paths in Boolean models. By construction, the distance between any two consecutive nodes on the path is smaller than r . We observe furthermore that if we consider every second ball along the path, these balls do not overlap. Otherwise, if they had overlapped, they would have made it possible to take a shortcut between them, which would have avoided at least one other ball and thus give a shorter path (see Figure 5.1), which is impossible. More generally, the same reason implies that any ball can overlap only with its predecessor and its successor along the shortest path.

We color now the with three colors –say blue red and green– so that the path follows the sequence blue red green blue red etc. Because of the above observation, two balls of the same color never overlap. Using this coloring, we can set up a three time slots TDMA scheme, where nodes of each color emit successively with full power P .

Let us assume that blue nodes are currently emitting, and that red nodes are listening to their blue predecessor along the path. We consider a particular (red) receiver, and compute a lower bound on the signal-to-interference-plus-noise ratio (SINR) at this node. First, the emitter is at distance at most r , since the receiver listen to its predecessor. The power of the received signal is thus at least $Pl(r)$, since $l(\cdot)$ is decreasing.

We look now at the interferences. We notice that all interfering (blue) nodes are located at distance at least r from the receiver. Furthermore, the interferers are at distance at least r from each other (since blue balls do not overlap). Thus, we can derive an upper bound to the power of the interferences by packing an

infinite number of blue balls around the receiver's ball. The optimal packing is obtained by placing the nodes on the vertices of a honey comb lattice.

Let us compute the sum of the sum of the powers of interfering signals: in a honey comb lattice, there are 6 neighbors at distance r , 6 at distance $\sqrt{3}r$, etc. We obtain:

$$\begin{aligned} \sum_{i=1}^{\infty} Pl(\|X_{2i} - X_0\|) &\leq P[6l(r) + 6l(\sqrt{3}r) + 6l(2r) + 12l((2 + \frac{\sqrt{3}}{2})r) + \dots] \\ &\leq P \sum_{k=1}^{\infty} 6(k+1)l(kr) \doteq PK(r). \end{aligned}$$

The sum $K(r)$ converges because $\int xl(x)dx < \infty$. Therefore, the SINR at the receiver is at least:

$$\text{SINR} \geq \frac{Pl(r)}{N_0 + PK(r)}$$

where N_0 is the power of the background noise. The following throughput is thus achievable between the emitter and the receiver

$$\frac{1}{2} \log \left(1 + \frac{Pl(r)}{N_0 + PK(r)} \right). \quad (5.1)$$

The same bound applies to each red receiver, and by rotating the colors, to all other emitter-receiver pairs in the two other time slots. Therefore, the overall throughput of the TDMA scheme is equal to one third of the throughput in (5.1), since we used three time slots. \square

The second lemma follows from Theorem 4.1.

Lemma 5.2 *In the Boolean model $\mathcal{G}(\lambda, r)$, for any $0 < \check{\theta} < 1$, there exists an $r < \infty$ such that there is a connected cluster that contains a fraction of nodes larger than $\check{\theta}$.*

Proof: $\theta(\lambda\pi r^2)$ is continuous for $r > r^*$ and tends to 1 when r tends to infinity, see [MR96]. Thus, given $0 < \check{\theta} < 1$, one can choose r large enough so that $\theta(\lambda\pi r^2) = \check{\theta}/(1 - \varepsilon)$ for some $\varepsilon > 0$. Theorem 4.1 ensures that the number of nodes inside the unbounded cluster \mathcal{U} is at least $(1 - \varepsilon)\theta(\lambda\pi r^2)n$ w.h.p. Thus the fraction of nodes inside this cluster is larger than $\check{\theta}$. \square

We can now prove Theorem 5.1.

Proof of Theorem 5.1: To prove this theorem, we use Lemma 5.2. Given $\check{\theta}$, this ensures that there exists r such that under the Boolean model, there exists a connected cluster of size greater than $n\check{\theta}$ w.h.p.

Then, we use Lemma 5.1 to show that along the shortest path in the Boolean model, the rate

$$R = \frac{1}{6} \log \left(1 + \frac{Pl(r)}{N_0 + P \sum_{k=1}^{\infty} 6(k+1)l(kr)} \right)$$

is achievable between any two connected nodes. Therefore, a fraction $\check{\theta}$ of the nodes can exchange data at rate at least R , independently of n . \square

The simple TDMA construction of Lemma 5.1, along with the percolation theory result of Lemma 5.2 have been enough to prove our theorem. A better

bound on the throughput can be obtained using more complex schemes than our simple TDMA strategy. One could use, for example, the multiple relay scheme described in [XK04a] and [KGG03], that leads in our case to an achievable throughput

$$R = \frac{1}{2} \log \left(1 + \frac{Pl(r)}{N_0} \right). \quad (5.2)$$

However, the asymptotic behavior remains the same as the one of our very simple scheme and only the pre-constant is improved.

It is also important to notice that the proof does not work for $\tilde{\theta} = 1$, as the fraction of connected nodes in a Boolean model is never equal to one. This is consistent with the fact that the rate must decrease to zero if we want to keep all nodes connected. It is therefore impossible to find a fixed rate $R > 0$ such that the fraction of connected nodes is equal to one.

5.3.2 Upper bound

We now derive an information theoretic bound on the rate at which a given fraction of the nodes can send data to any destination. This result does not depend on the adopted strategy for transmitting information to the destination.

Theorem 5.2 *For any rate $R > 0$, the fraction of nodes that can send data to any destination at that rate is at most $\hat{\theta}$ w.h.p., where*

$$\hat{\theta} = \mathbb{P} \left(I(0) \geq \frac{N_0}{P} (e^{2R} - 1) \right),$$

where $I(0)$ is the value of the shot-noise at the origin, defined by

$$I(0) = \sum_k l(\|X_k\|) \quad (5.3)$$

and $\{X_k\}$ are the points of a Poisson process of unit density over \mathbb{R}^2 .

Proof: To prove Theorem 5.2, we proceed in two steps. First, we show that the rate at which a node can send data to any destination is bounded above by a function of a shot-noise at its location. Then we show that the fraction of nodes such that this shot-noise at their location is lower than a certain threshold is $\hat{\theta}$.

In the first step, we use the max-flow min-cut theorem 14.10.1 from [CT91]. Accordingly, we divide the network in the following way: on one side the sender X_0 , and on the other side all the other nodes $X_k, k \neq 0$. The max-flow min-cut theorem ensures that the maximum throughput between these two sides can be upper-bounded by the multiple receiver Gaussian channel formed by the emitting node and all others listening to it. The rate of this channel is computed in [Tel99] and corresponds to the rate at which the emitter X_0 can send data to the rest of the nodes X_k :

$$R = \frac{1}{2} \log \left(1 + \frac{P \sum_{k=1}^{\infty} l(\|X_k - X_0\|)}{N_0} \right).$$

The sum in the above expression is a shot-noise process evaluated in X_0 , that we will denote by S later on:

$$S(X_0) \doteq \sum_{k=1}^{\infty} l(\|X_k - X_0\|).$$

From this result, we conclude that a necessary condition for achieving rate R from node X_0 is

$$S(X_0) \geq \frac{N_0}{P} (e^{2R} - 1) \doteq M. \quad (5.4)$$

We can then use Theorem B.1 in Appendix B.1 to compute the fraction of nodes that fulfill the above condition. As the sequence of squares $[0, \sqrt{n}] \times [0, \sqrt{n}]$ is a convex averaging sequence (see Definition B.1 in Appendix B.1), we conclude from this theorem that the fraction of nodes that do not fulfill the condition is equal to

$$\mathbb{P}(I(0) < M) \doteq 1 - \hat{\theta}$$

where $I(0)$ is given by (5.3). As they do not fulfill the necessary condition (5.4), a fraction $1 - \hat{\theta}$ of nodes are unable to send data at rate R to any destination, and are thus isolated. The fraction of connected nodes is therefore at most $\hat{\theta}$. \square

The bound given in Theorem 5.2 is not explicit, as it involves the cumulative distribution function of a shot-noise. To obtain an explicit bound, we can use Chernoff's inequality. We know from Campbell's theorem that [Kin93, page 28]

$$\mathbb{E}[e^{sI(0)}] = \exp \left\{ \lambda \int_{\mathbb{R}^2} (e^{sI(\|\mathbf{x}\|)} - 1) d\mathbf{x} \right\}.$$

Therefore, Chernoff's inequality yields

$$\begin{aligned} \mathbb{P}[I(0) \geq M] &\leq \inf_{s \in \mathbb{R}} e^{-sM} \mathbb{E}[e^{sI(0)}] \\ &= \inf_{s \in \mathbb{R}} \exp \left\{ \lambda \int_{\mathbb{R}^2} (e^{sI(\|\mathbf{x}\|)} - 1) d\mathbf{x} - sM \right\}. \end{aligned}$$

This bound gives a good approximation when R is large, but becomes loose when R is smaller.

5.4 Discussion

We plot in Figure 5.2 the upper and lower bounds on the fraction of the nodes that can connect to each other at a given rate R .

The lower bound indicates that when R is close to zero, the fraction of nodes that can achieve this rate tends to one. This case corresponds to the results in [LS04].

At the other extreme, if R is too large, the lower bound becomes zero. In practice, our lower bound represents a percolation curve that marks a transition at a critical rate value *below which* a non-zero fraction of the nodes in the network can sustain a constant rate.

The upper bound computed by Chernoff's inequality is not informative for small values of R , but decreases to zero for large values of R . Contrary to the lower bound, the curve has a tail for large rates. Although upper and lower bounds are not tight, we believe that there is also a critical rate, *above which* each node can only connect to a finite number of other nodes. This means that under our definition of connectivity, the network would not percolate in this case.




Figure 5.2: Upper and lower bounds on the fraction of nodes that can achieve a given rate R . The lower bound has been improved using the multiple relay channel result given in (5.2), and the upper bound has been computed using the bound presented in Section 5.3.2.

5.5 Conclusion

We saw in Chapter 4 that in the Boolean model, the requirement of full connectivity results in bad scaling properties of the network. We showed in this chapter that even in the most optimistic information-theoretic setting (arbitrary complexity of the network encoding, all nodes acting as relay for one pair source-destination arbitrarily picked in the network), if we want to impose a given rate $R > 0$ to any possible transmission in the network, a fraction $1 - \hat{\theta}(R)$ of nodes will automatically be disconnected. This result is obtained using tools from shot noise processes and information theory.

On the other hand, if we allow some non zero, but arbitrary small fraction $\check{\theta}$ of nodes to be disconnected, then Theorem 5.1 shows that it is possible to find a rate $R(\check{\theta})$ that any other pair of nodes can enjoy in an arbitrarily large network. The theorem is proved by continuum percolation techniques, and therefore holds for two-dimensional networks, but not for one-dimensional networks, because percolation does not occur in dimension 1.

These two results shed some new light on the throughput scaling laws of random wireless networks. Relaxing the full connectivity requirement and allowing a small fraction of the nodes to be disconnected is shown to be both necessary (Theorem 5.2) and sufficient (Theorem 5.1) to have nodes communicating at a non zero, positive rate, with other nodes acting as potential relays, on an arbitrary large network. Contrary to the full connectivity case, the dimensionality of the network is now a crucial factor.

This connectivity model relies on the assumption that only one data flow is active at a time. In the next chapter, we address the same question of the throughput scaling, but in a case where all nodes are simultaneously acting as a source.

Chapter 6

The throughput under uniform traffic matrix

In their seminal paper [GK00], Gupta and Kumar showed that in a network of n arbitrarily located nodes the throughput must go to zero, as n tends to infinity, at least at rate $1/\sqrt{n}$. They also showed that in a network of randomly located nodes, a throughput of order $1/\sqrt{n \log n}$ can be achieved. The purpose of this chapter is to show that a $1/\sqrt{n}$ throughput is also achievable in networks of randomly located nodes. Our result follows from a natural application of percolation theory techniques [Gri99, MR96].

Following the work in [GK00], throughput scaling in wireless networks has received considerable attention [eGMPS04, JVK04, KV02, LT04, TG04, XK04a]. As we have seen in Chapter 4, there are two ways of letting the number of nodes n tend to infinity. One can either keep the area on which the network is deployed constant, and make the node density λ tend to infinity (*dense* networks); or one can keep the node density λ constant, and increase the area to infinity (*extended* networks). In both of these settings, network theoretic lower bounds on achievable transmission rates can be obtained constructively, for given communication strategies and power attenuation laws; while information theoretic upper bounds must be obtained allowing arbitrary communication strategies and assuming only the power decay law in the propagation medium.

The work in [GK00] considered the dense network case, and a traffic scenario where each node generates packets for a destination non-vanishingly far away. Using a network theoretic approach based on multi-hop communication, it showed a lower bound on the per-node rate of $\Omega(1/\sqrt{n})$ bit/sec, if nodes are arbitrarily located; and a lower bound of $\Omega(1/\sqrt{n \log n})$ bit/sec if nodes are randomly located. These results rely on point to point connections delivering higher power as nodes tend to be closer to each other. In practice, this holds only as long as near field effects can be neglected. When the physical constraint of bounded power in the near field is enforced, computed bounds reduce to $1/n$, as seen in Section 4.2.2.

For extended networks, near field effects do not play a fundamental role and the power received by every node can be bounded without affecting the final result. In this case, the work in [XK05] presents an information theoretic bound of $\Theta(1/\sqrt{n})$ bit/sec per node, for arbitrarily located nodes satisfying a minimum

distance constraint, random phase fading and a power attenuation function that exhibits a power law behavior with exponent $\alpha > 4$, or an exponential attenuation. When nodes are randomly located, the work in [LT04] shows an upper bound of $O(1/n^{1/2-1/\alpha})$ that holds for $\alpha > 2$.

Constructive strategies proposed for networks of randomly located nodes [eGMPS04, GK00, KV02, TG04] achieve only $\Omega(1/\sqrt{n \log n})$ per-node bit rate, somehow suggesting that at least a $\sqrt{\log n}$ factor is the price to pay for randomness. Instead, we show the contrary, namely that it is possible to achieve a per-node throughput capacity of $\Omega(1/\sqrt{n})$ in random networks. Hence, nodes in a random network can transmit at the same rate than nodes in an arbitrary network and there is no price to pay—at least asymptotically—for the additional randomness present in the system. Our result holds in the same setting of [GK00], i.e., for dense networks, as long as near field effects are negligible. It also holds in the more general case of extended networks with bounded transmitted and received power, assuming only a power decay law whose tail exhibits a power law behavior with exponent $\alpha > 2$, or an exponential attenuation (which is typical if there is absorption in the medium [FBS04]). The proposed routing also achieves the optimal average delay required for a packet to reach its destination, as defined in [eGMPS04].

Finally, we want to spend few words on the connection we establish with the field of percolation theory. In this chapter, we construct a mapping such that the open grid edges of a percolation model correspond to the presence of wireless transmitters in certain locations of the plane, and the open percolating paths represent a wireless backbone that is used to multi-hop packets across the network. Accordingly, to maximize throughput, we want to operate at $p > 1/2$, above the percolation threshold, so that we can guarantee the existence of many paths that traverse the network, but also have $p < 1$, so to avoid overcrowding and excessive interference. We show that controlling the parameter p corresponds in the wireless network to controlling the transmission bit-rate of the nodes as $n \rightarrow \infty$, and we find that the optimal scaling law corresponds to some $1/2 < p < 1$.

6.1 Main results

We construct a random extended network by placing nodes according to a Poisson point process of unit intensity on the plane and focus our attention to the square $[0, \sqrt{n}] \times [0, \sqrt{n}]$. Similarly, we construct a dense network by placing nodes according to a Poisson point process of intensity n over a square of unit area. We are mainly concerned with events that occur inside these squares with high probability (w.h.p.), that is, with probability tending to one as $n \rightarrow \infty$. We denote the Euclidean distance between two nodes i and j by d_{ij} . We pick uniformly at random a matching of source-destination pairs, so that each node is the destination of exactly one source. We assume all nodes transmit at constant power P , and that node j receives the transmitted signal from node i with power $Pl(i, j)$, where $l(i, j)$ indicates the path loss between i and j . In this chapter we are concerned with lower bounds on achievable rates, hence, we assume a model of multi-hop communication such that two nodes can establish

a direct wireless link of capacity

$$C(i \rightarrow j) = \frac{1}{2} \log \left(1 + \frac{Pl(i, j)}{N_0 + \sum_{k \neq i} Pl(k, j)} \right) \text{ bps/Hz},$$

where N_0 is the ambient noise power at the receiver.

The per-node throughput capacity $T(n)$ of the network is defined as the number of bits per second that every node can transmit w.h.p. to its destination.

The per-packet delay $D(n)$ is the sum of the times the packet spends at each relay node. When we measure this, we scale the packet size by $T(n)$. In this way the transmission delay at each node is constant and the per-packet delay corresponds to the number of hops needed to reach its destination. In practice, thanks to scaling, $D(n)$ captures the dynamics of network and not of the transmission delay, see also [eGMPS04].

Our results are the following.

Theorem 6.1 *Assuming a power attenuation function of the type $l(i, j) = \min\{1, e^{-\gamma d_{ij}}/d_{ij}^\alpha\}$ with $\alpha > 0, \gamma > 0$ or $\alpha > 2, \gamma = 0$, a per-node throughput capacity of*

$$T(n) = \Omega(1/\sqrt{n}) \text{ bit/sec}$$

is achievable in a random extended network, with a corresponding average packet delay at most

$$\mathbb{E}(D(n)) = O(\sqrt{n}).$$

Theorem 6.2 *Assuming a power attenuation function of the type $l(i, j) = e^{-\gamma d_{ij}}/d_{ij}^\alpha$ with $\alpha > 0, \gamma > 0$, or $\alpha > 2, \gamma = 0$, a per-node throughput capacity of*

$$T(n) = \Omega(1/\sqrt{n}) \text{ bit/sec}$$

is achievable in a random dense network, with a corresponding average packet delay at most

$$\mathbb{E}(D(n)) = O(\sqrt{n}).$$

Note that the result of Theorem 6.2 relies on an ideal power attenuation function $l(\cdot)$ that is singular at the origin. In practice, this result holds as long as near field effects are negligible and it is presented for completeness, as it matches the dense network model of [GK00]. The per-node throughput in dense networks and in presence of bounded attenuation function cannot scale better than $1/n$, as we saw in Section 4.2.2. Most of the chapter is devoted to the proof of Theorem 6.1, while Theorem 6.2 is proved in Section 6.6, by adapting a few steps of the main proof.

6.2 Overview of the solution (extended networks)

The main idea of our proposed solution is to have a wireless backbone of nodes that carry packets across the network at constant rate, using short hops, and to drain the rest of the traffic to the wireless backbone using single hops of longer length. See Figure 6.1 for a schematic representation.




Figure 6.1: Nodes inside a slab of constant width access a path of the wireless backbone in single hops of length at most proportional to $\log \sqrt{n}$. Packets are carried across a distance \sqrt{n} along the wireless backbone in constant-length hops.

Since all nodes transmit at the same power, we expect the longer hops that access the backbone to have a lower bit-rate, due to higher power loss across longer distances; however, one needs to take into account other components that influence the bit-rate: namely, interference, and cost of relay packets from other nodes. It turns out that when all these components are accounted for, the bottleneck is due to the traffic on the backbone.

The wireless backbone consists of paths of constant length hops. We will show that the interference caused by nodes sufficiently far away along the path is small enough to allow a constant transmission rate along the path. Each path, however, needs to relay packets coming from other nodes that access the backbone in single hops. This traffic is at most proportional to \sqrt{n} , if we associate to each path only nodes that are within a slab of constant width that crosses the network area, see Figure 6.1. Hence, the per-node throughput on the backbone can be only of order $1/\sqrt{n}$.

Now, let us look at the throughput of the nodes that access the wireless backbone in single hops. We will show that these single hops are of length at most proportional to $\log \sqrt{n}$ and can sustain a rate higher than $1/\sqrt{n}$, even in presence of exponential power attenuation. Furthermore, there is no relay burden for nodes accessing the backbone in single hops. It follows that the bottleneck is represented by the nodes on the backbone that transmit at a rate of $1/\sqrt{n}$.

There are three key points in our reasonings: i) there exist paths of constant hop length that cross the entire network forming the wireless backbone, ii) these paths can be put into a one to one correspondence with \sqrt{n} slabs of constant width, each containing at most a constant times \sqrt{n} number of nodes, and iii) these paths are somehow regularly spaced so that there is always one within a $\log \sqrt{n}$ distance factor from any node in the network.

In the following, Theorem 6.3 ensures the existence of many paths using percolation theory arguments. Theorem 6.4 shows that each path in the wireless backbone can transport packets at a constant rate, and that packets can be drained to the backbone at a rate higher than \sqrt{n} . Finally, Lemmas 6.3 and 6.4 are needed to bound the number of nodes that access any given path.




Figure 6.2: Construction of the bond percolation model. We declare each square on the left-hand side of the picture open, if there is at least a Poisson point inside it, closed otherwise. This corresponds to associate an edge to each square, traversing it diagonally, as depicted on the right-hand side of the figure, and declare the edge either open or closed according to the state of the corresponding square.

6.3 Percolation results

In this section we establish the percolation results that are needed to show existence of a cluster of nodes forming the wireless backbone. The objective is to formally construct a mesh of paths that can simultaneously carry information across the network at a constant rate, independent of the number of nodes n . We call this mesh the *highway system*, and will use it to carry packets over most of the distance.

To begin our construction, we divide the area into squares of constant side length c , as depicted in the left-hand of Figure 6.2. By adjusting c , we can adjust the probability that a square contains at least one point:

$$\mathbb{P}(\text{a square contains at least one point}) = 1 - e^{-c^2} \doteq p. \quad (6.1)$$

We say that a square is *open* if it contains at least one point, and *closed* otherwise; note that the status of the squares is i.i.d.

We now map our construction to a bond percolation model. We draw an horizontal edge across half of the squares, and a vertical edge across the others, as shown on the right-hand side of Figure 6.2. In this way we obtain a lattice of horizontal and vertical edges, each edge being open, independently of all other edges, with probability p . We call a path *open* (resp. *closed*) if it contains only open (resp. closed) edges. Note that, for c large enough, our construction produces winding open paths that cross the network area, see Figure 6.3. Next, we turn to the question of how many of these paths there are.

Throughout the proof, we will introduce several subdivisions of the networks and other objects of different sizes. We summarize them in Table 6.1 to ease the reading of the proof.




Figure 6.3: Horizontal paths in a 40×40 bond percolation model obtained by computer simulation. Each square is traversed by an open edge with probability p ($p = 0.7$ here). Closed edges are not depicted. We find seven disjoint open path crossing the area from left to right.

| Object | Size | Lattice size | Nb. of objects |
|-----------------------|---|----------------------------|------------------------------------|
| Whole network area | $\sqrt{n} \times \sqrt{n}$ | $m \times m$ | 1×1 |
| Rectangle \bar{R}_n | $\sqrt{n} \times \sqrt{2c\kappa} \log \frac{\sqrt{n}}{\sqrt{2c}}$ | $m \times \kappa \log m$ | $1 \times \frac{m}{\kappa \log m}$ |
| Slab (Lemma 6.4) | $\sqrt{n} \times \frac{\sqrt{2c}}{\beta}$ | $m \times \frac{1}{\beta}$ | $1 \times \beta m$ |

Table 6.1: Divisions of the network. Recall that $m = \sqrt{n}/\sqrt{2c}$.




Figure 6.4: The network area is divided into $m/\kappa \log m$ horizontal rectangles of lattice size $m \times \kappa \log m$. A left to right crossing of rectangle \bar{R}_n is shown.

6.3.1 Number of disjoint crossing paths

Let us divide the network area into horizontal rectangles \bar{R}_n , of size $\sqrt{n} \times \sqrt{2c\kappa \log \frac{\sqrt{n}}{\sqrt{2c}}}$, for some constant $\kappa > 0$, see Figure 6.4. Each of these rectangles has thus lattice size $m \times \kappa \log m$ in the bond percolation model, with $m = \frac{\sqrt{n}}{\sqrt{2c}}$ (as the edges have length $\sqrt{2c}$). We want to show that there exist many disjoint open paths from left to right inside such rectangles.

Theorem 6.3 *For any constant $\kappa > 0$ and if c is so large that*

$$c^2 > \log 6 + \frac{2}{\kappa}, \quad (6.2)$$

then there exists a strictly positive constant $\beta = \beta(c, \kappa)$ such that w.h.p. there exist $\beta\kappa \log m = \beta\kappa \log \frac{\sqrt{n}}{\sqrt{2c}}$ disjoint open paths inside each rectangle \bar{R}_n , that cross it from left to right.

In order to prove this theorem we need some preliminary results expressed by the lemmas below. We denote by P_p the product measure with open edge density p .

Lemma 6.1 *Let S_n be a square lattice of size $n \times n$ and $0 < p \leq 1$. The probability that there exists an open path from the center 0 of S_n to its boundary ∂S_n is upper bounded by*

$$P_p(0 \leftrightarrow \partial S_n) \leq \frac{4}{3}(3p)^n.$$

Proof: If a path starts at 0 and touches ∂S_n , its length is at least n . Thus, denoting by $N(n)$ the number of open paths of length n starting at the origin:

$$P_p(0 \leftrightarrow \partial S_n) \leq P_p(N(n) \geq 1).$$

As paths of length n are open with probability p^n , we can bound this probability by

$$P_p(N(n) \geq 1) \leq p^n \sigma(n),$$

where $\sigma(n)$ denotes the number of paths of length n starting at the origin. This number is obviously not larger than

$$\sigma(n) \leq 4 \cdot 3^{n-1}.$$

Combining these three inequalities, we obtain

$$P_p(0 \leftrightarrow \partial S_n) \leq 4 \cdot 3^{n-1} p^n = \frac{4}{3} (3p)^n.$$

□

The next well known results follows directly from Theorem 2.45 in [Gri99] and the remarks thereafter.

Lemma 6.2 *Let R_n be a rectangle embedded in the square lattice. Let A_n be the event that there exists an open path between the left and right sides of R_n and $I_r(A_n)$ the event that there exist r edge-disjoint such crossings. We have*

$$1 - P_p(I_r(A_n)) \leq \left(\frac{p}{p-p'} \right)^r [1 - P_{p'}(A_n)]$$

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

Proof of Theorem 6.3: We consider bond percolation in the rectangle \bar{R}_n , with each edge having probability p' to be open, independently of all other edges. Let the *dual graph* of \bar{R}_n be obtained by placing a vertex in each square of the percolation lattice, and joining two such vertices by an edge whenever the corresponding squares share a side, see Figure 6.5. An edge of the dual is open if it crosses an open edge of the original lattice, it is closed otherwise. Let A_n be the event of having at least one open path inside \bar{R}_n that crosses it from left to right, and let B_n be the event that a closed path crosses the rectangle vertically in the dual lattice. We have $A_n \cap B_n = \emptyset$, because if both A_n and B_n occur, then there must be an intersection between an open edge of \bar{R}_n and a closed edge of its dual, which is impossible. Moreover, whenever A_n does not occur then B_n occurs (one can be convinced of this again by looking at Figure 6.5). It follows that A_n and B_n are disjoint events which partition the sample space, and hence $P_{p'}(A_n) + P_{p'}(B_n) = 1$.

We index the nodes at the base of the dual graph by index i , and denote by $i \leftrightarrow \square \bar{R}_n$ the existence of a closed path in the dual from node i to the opposite side $\square \bar{R}_n$ at the top of the dual graph. As \bar{R}_n has width $\kappa \log m$, we have for any $0 \leq q \leq 1$:

$$P_q(i \leftrightarrow \square \bar{R}_n) \leq P_q(0 \leftrightarrow \partial S_{\kappa \log m}).$$

Therefore, as edges are closed with probability $1 - p'$,

$$\begin{aligned} P_{p'}(B_n) &\leq \sum_{i=1}^m P_{1-p'}(i \leftrightarrow \square \bar{R}_n) \\ &\leq \sum_{i=1}^m P_{1-p'}(0 \leftrightarrow \partial S_{\kappa \log m}) \\ &\leq \frac{4m}{3} (3(1-p'))^{\kappa \log m} \end{aligned}$$




Figure 6.5: A picture of \overline{R}_n (solid line) and its dual graph (dotted line). Note that if there is no open path traversing the rectangle from left to right, then there is a closed path in the dual graph traversing it from top to bottom.

where the first inequality is a union bound, and the third inequality follows from Lemma 6.1.

Now we look at the event that $\beta\kappa \log m$ disjoint paths exist. We apply Lemma 6.2 to rectangle \overline{R}_n with $r = \beta\kappa \log m$, obtaining

$$1 - P_p(I_{\beta\kappa \log m}(A_n)) \leq \left(\frac{p}{p-p'}\right)^{\beta\kappa \log m} P_{p'}(B_n), \quad (6.3)$$

for any $p' < p$. Let us choose $p' = 2p - 1$. We have thus

$$\frac{p}{p-p'} = e^{c^2} - 1 < e^{c^2},$$

and

$$1 - p' = 2(1 - p) = 2e^{-c^2}.$$

Hence, Equation (6.3) becomes

$$\begin{aligned} 1 - P_p(I_{\beta\kappa \log m}(A_n)) &\leq \left(e^{c^2}\right)^{\beta\kappa \log m} P_{p'}(B_n) \\ &\leq m^{\beta\kappa c^2} \frac{4m}{3} (6e^{-c^2})^{\kappa \log m} \\ &= m^{\beta\kappa c^2} \frac{4m}{3} m^{-\kappa c^2 + \kappa \log 6} \\ &= \frac{4}{3} m^{(\beta-1)\kappa c^2 + \kappa \log 6 + 1}. \end{aligned}$$

The probability to find at least $\beta\kappa \log m$ paths in \overline{R}_n is thus

$$P_p(I_{\beta\kappa \log m}(A_n)) \geq 1 - \frac{4}{3} m^{(\beta-1)\kappa c^2 + \kappa \log 6 + 1}.$$

As this happens independently in each of the $\frac{m}{\kappa \log m}$ rectangles, the probability of having $\beta\kappa \log m$ disjoint paths in each rectangle is

$$P_p(I_{\beta\kappa \log m}(A_n))^{\frac{m}{\kappa \log m}} \geq \left(1 - \frac{4}{3} m^{(\beta-1)\kappa c^2 + \kappa \log 6 + 1}\right)^{\frac{m}{\kappa \log m}}.$$

Finally note that if $(\beta - 1)\kappa c^2 + \kappa \log 6 + 1 \leq -1$, the above expression tends to 1 when m goes to infinity. Thus, if $c^2 > \log 6 + 2/\kappa$, one can choose

$$\beta(c, \kappa) = 1 - \frac{\kappa \log 6 + 2}{\kappa c^2} > 0$$

such that the above condition is fulfilled. \square

Joining all the rectangles together, we obtain $\beta\sqrt{n}$ paths in the whole network. The same is true of course if we divide the area into vertical rectangles and look for paths crossing the area from bottom to top. Using a simple union bound argument, we conclude that there exist $\beta\sqrt{n}$ horizontal and $\beta\sqrt{n}$ vertical disjoint paths simultaneously with high probability. These paths form a grid, that we call the *highway system*.

6.4 Capacity of the percolation cluster

Along the paths of the highway system, we choose one node per edge, that relays the packets. This is possible as the paths are formed by open edges, which are associated to non-empty squares. The paths are thus made of a chain of nodes such that the distance between two consecutive nodes is at most $2\sqrt{2}c$.

To actually transport packets along the paths, we set up a TDMA scheme. When a node transmits, other nodes that are sufficiently far away can simultaneously transmit, without causing excessive interference. Theorem 6.4 makes this precise, ensuring that a constant rate R , independent of n , can be achieved on all the paths simultaneously as $n \rightarrow \infty$. Note that this theorem gives a more general result, that will be useful also in Section 6.5.

Theorem 6.4 *For any given integer $d > 0$ there exists a TDMA scheduling, such that one node per square can transmit to any destination located within a radius of d squares (in Manhattan distance) with fixed rate $R(d)$ independent of n .*

When d goes to infinity, the asymptotic behavior of the rate is given by

$$R(d) = \Omega(d^{-\alpha-2}e^{-\gamma cd}).$$

Proof: We take a coordinate system, and label each square with two integer coordinates (in our construction, the axis of the coordinate system are diagonal). Then we take an integer k , and consider the subset of squares whose two coordinates are a multiple of k (see Figure 6.6). By translation, we can construct k^2 disjoint equivalent subsets. This allows us to build the following TDMA scheme: we define k^2 time slots, during which only nodes from a particular subset are allowed to emit. We assume also that at most one node per square emits at the same time, and that they all emit with the same power P .

Let us consider one particular square. We suppose that the emitter in this square emits towards a destination located in a square at distance at most d . We compute the signal-to-interference ratio at the receiver. First, we choose the number of time slots k^2 as follows:

$$k = 2(d + 1)$$

To find an upper bound to the interferences, we observe that with this choice, the emitters in the 8 first closest squares are located at a distance at least $d + 2$ (in squares) from the receiver (see Figure 6.7). This means that the Euclidean distance between the receiver and the 8 closest interferers is at least $c(d + 1)$. The 16 next closest squares are at distance at least $3d + 4$ (in squares), and the




Figure 6.6: A set of squares assigned to a given time slot.

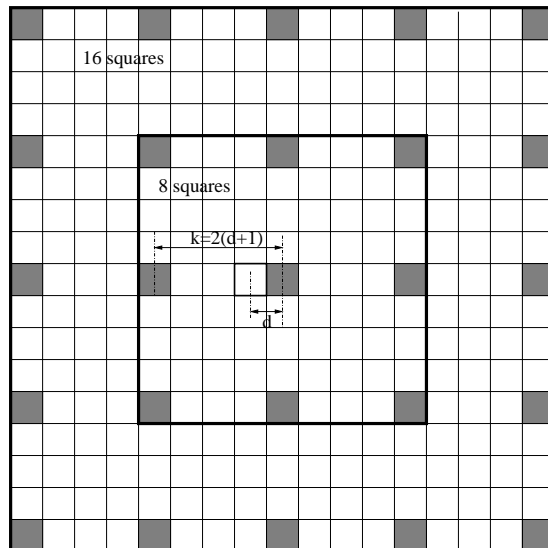


Figure 6.7: Construction of the lower bound on the interference term.

Euclidean distance between the receiver and the 16 next interferers is therefore at least $c(3d+3)$, and so on. The sum of the interferences $I(d)$ can be bounded as follows:

$$\begin{aligned} I(d) &\leq \sum_{i=1}^{\infty} 8i Pl(c(2i-1)(d+1)) \\ &= \sum_{i=1}^{\infty} 8i P \min\{1, [c(2i-1)(d+1)]^{-\alpha}\} e^{-\gamma c(2i-1)(d+1)} \\ &\leq \sum_{i=1}^{\infty} 8i P [c(2i-1)(d+1)]^{-\alpha} e^{-\gamma c(2i-1)(d+1)} \\ &= P[c(d+1)]^{-\alpha} e^{-\gamma c(d+1)} K(d). \end{aligned}$$

where

$$K(d) = \sum_{i=1}^{\infty} 8i (2i-1)^{-\alpha} e^{-\gamma c(d+1)(2i-2)}.$$

The latter sum clearly converges if $\alpha > 2$ or $\gamma > 0$.

Now we want to bound from below the signal received from the emitter. We observe first that the distance between the emitter and the receiver is at most

$$\sqrt{(cd)^2 + c^2} \leq c(d+1).$$

The strength $S(d)$ of the signal at the receiver can be thus bounded by

$$\begin{aligned} S(d) &\geq Pl(c(d+1)) \\ &= P \min\{1, [c(d+1)]^{-\alpha}\} e^{-\gamma c(d+1)}. \end{aligned}$$

Finally, we obtain a bound on the signal-to-interference ratio

$$\begin{aligned} SINR(d) &= \frac{S(d)}{N_0 + I(d)} \\ &\geq \frac{P \min\{1, [c(d+1)]^{-\alpha}\} e^{-\gamma c(d+1)}}{N_0 + P[c(d+1)]^{-\alpha} e^{-\gamma c(d+1)} K(d)}. \end{aligned}$$

As the above expression does not depend on n , the first part of the theorem is proven.

We now look at the asymptotic behavior of the SINR for large d . If $c(d+1) \geq 1$, we can remove the minimum and write

$$\begin{aligned} SINR(d) &\geq \frac{P[c(d+1)]^{-\alpha} e^{-\gamma c(d+1)}}{N_0 + P[c(d+1)]^{-\alpha} e^{-\gamma c(d+1)} K(d)} \\ &= \frac{1}{N_0 [c(d+1)]^{\alpha} e^{\gamma c(d+1)} / P + K(d)}. \end{aligned}$$

The first term grows like $d^{\alpha} e^{\gamma cd}$ when d goes to infinity, whereas $K(d)$ decreases to zero. The whole fraction therefore decreases like $1/d^{\alpha} e^{\gamma cd}$. The throughput on each link is given by $\log(1 + SINR(d))$, and therefore also decreases like $1/d^{\alpha} e^{\gamma cd}$.

Now we have to divide this throughput by the number of time slots k^2 used in the TDMA scheme. As $k = 2(d+1)$, the number of time slots increases like d^2 . So, finally, the actual throughput available in each square decreases like $d^{-\alpha-2} e^{-\gamma cd}$. \square

Corollary 6.5 *For any given integer $d > 0$ there exists a TDMA scheduling, such that one node per square can receive from any emitter located within a radius of d squares with fixed rate $R(d)$ independent of n .*

When d goes to infinity, the asymptotic behavior of the rate is given by

$$R(d) = \Omega(d^{-\alpha-2}e^{-\gamma cd}).$$

Proof: This result is obtained by switching the role of emitters and receivers in the above proof. Distances remain the same, and all equations still hold. \square

6.5 Protocol

In this section, we describe the actual routing protocol and show that it achieves the desired $\Omega(\sqrt{n})$ throughput capacity.

The protocol uses 4 separate time slots: a first one for draining packets to the highway, a second one to transport packets on the “horizontal” highways connecting the left and right edges of the domain, a third one to transport packets on the “vertical” highways connecting the top and bottom edges of the domain, and a fourth one to deliver packets to the destination. The draining and delivery phases use direct transmission, while the highway phases use multiple hops. We show that the throughput bottleneck is in the highway phase that can sustain a rate per node proportional to $1/\sqrt{n}$ bit per second.

We start by proving two simple lemmas that will be useful for the capacity calculation. They are followed by three propositions that prove our main result.

Lemma 6.3 *Divide the network area in n/c^2 boxes of side length c . The probability that there are less than $c^2 \log n$ nodes in each box tends to one when n goes to infinity.*

Proof: The number of nodes in each square is a Poisson random variable of parameter c^2 . Let us denote one of these variables by X . Chernoff’s inequality implies that

$$\mathbb{P}(X > c^2 \log n) \leq e^{-s c^2 \log n} \mathbb{E}(e^{sX}),$$

for any $s > 0$. We choose here $s = 2/c^2$ and obtain

$$\mathbb{P}(X \leq c^2 \log n) \geq 1 - n^{-2} e^{c^2(e^{2/c^2} - 1)}.$$

As the numbers of nodes in all of the n/c^2 squares are i.i.d, we have

$$\mathbb{P}(X \leq c^2 \log n)^{\frac{n}{c^2}} \geq \left(1 - \frac{e^{c^2(e^{2/c^2} - 1)}}{n^2}\right)^{\frac{n}{c^2}}.$$

The latter expression tends to one when n goes to infinity. \square

Lemma 6.4 *Divide the network into horizontal slabs of constant width $\sqrt{2}c/\beta$. The probability that each slab contains less than $2c\sqrt{2n}/\beta$ nodes tends to one when n goes to infinity.*

Proof: The number of nodes in the i -th slab is a Poisson random variable of parameter $c\sqrt{2n}/\beta$, that we denote by N_i here. We apply Chernoff's inequality:

$$\begin{aligned} \mathbb{P}\left(N_i > 2\frac{c\sqrt{2n}}{\beta}\right) &\leq e^{-2\frac{c\sqrt{2n}}{\beta}s} \mathbb{E}(e^{sN_i}) \\ &= e^{-2\frac{c\sqrt{2n}}{\beta}s} e^{c\sqrt{2n}/\beta(e^s-1)} \\ &= e^{\frac{c\sqrt{2n}}{\beta}(e^s-2s-1)}. \end{aligned}$$

Thus, if we take $s = 1$,

$$\mathbb{P}\left(N_i > 2\frac{c\sqrt{2n}}{\beta}\right) \leq e^{\frac{c\sqrt{2n}}{\beta}(e-3)}.$$

The probability that each of the $\beta\sqrt{n}/\sqrt{2}c$ slabs contains less than $2\sqrt{2nc}/\beta$ nodes is thus

$$\mathbb{P}\left(N_i \leq 2c\sqrt{2n}/\beta, \forall i\right) \geq \left(1 - e^{\frac{c\sqrt{2n}}{\beta}(e-3)}\right)^{\frac{\beta\sqrt{n}}{c\sqrt{2}}}.$$

The latter expression tends to one when n goes to infinity. \square

Proposition 6.1 (Draining Phase) *Each node can transmit packets to the highway system with rate of order*

$$\Omega\left((\log \sqrt{n})^{-\alpha-2}(\log n)^{-1}n^{-c\kappa\gamma/2}\right),$$

for any constant $\kappa > 0$ and if c is so large that (6.2) holds.

Proof: We consider carrying packets from sources to the highways. We start by dividing the network area into $\beta\sqrt{n}/\sqrt{2}c$ horizontal slabs. As there are exactly as many slabs as horizontal edge-disjoint paths, we can impose that nodes from the i -th slab send their packets to the i -th horizontal path. Note that each path may not be contained in its corresponding slab, but it may deviate from it (recall Figure 6.1). However, Theorem 6.3 bounds the amount of deviation.

More precisely, to each source in the i -th slab, we assign an *entry point* on the i -th horizontal path. The entry point is defined as the node on the horizontal path closest to the vertical line drawn from the source point, see Figure 6.8. The source then transmits its packet to the entry point in a single hop. Theorem 6.3 ensures that the distance between sources and entry points is never larger than $\kappa \log m = \kappa \log \frac{\sqrt{n}}{\sqrt{2}c}$ squares. This is because each rectangle \bar{R}_n of width $\kappa \log m$ contains $\beta\kappa \log m$ paths, and therefore each source finds its highway within the same rectangle.

To compute the rate at which nodes can send their packets to entry points, we let $d = \kappa \log \frac{\sqrt{n}}{\sqrt{2}c}$ in Theorem 6.4. We obtain that one node per square can send packets to its entry point with rate

$$\begin{aligned} R\left(\kappa \log \frac{\sqrt{n}}{\sqrt{2}c}\right) &= \Omega\left(\left(\kappa \log \frac{\sqrt{n}}{\sqrt{2}c}\right)^{-\alpha-2} e^{-\gamma c\kappa \log \frac{\sqrt{n}}{\sqrt{2}c}}\right) \\ &= \Omega\left((\log \sqrt{n})^{-\alpha-2} e^{-c\kappa\gamma \log \sqrt{n}}\right) \\ &= \Omega\left((\log \sqrt{n})^{-\alpha-2} n^{-c\kappa\gamma/2}\right). \end{aligned}$$




Figure 6.8: Draining of packets to the highways

But as there are possibly many nodes in the squares, they have to share this bandwidth.

Using Lemma 6.3, we can conclude that the transmit rate of each node in the draining phase of our protocol is at least $R(d)/c^2 \log n$, which concludes the proof. \square

Proposition 6.2 (Highway Phase) *Along the highway packets can be relayed at rate at least $\beta/2k^2\sqrt{n}$ bits/sec per-node.*

Proof: We now compute the rate that can be sustained during each highway phase of our protocol. Each node generates packets at constant rate W and we must carry these packets towards the point on the vertical highway appropriate for delivery.

We divide horizontal and vertical traffic, adopting the following simple routing policy: packets are carried along horizontal highways until they reach the crossing with their target vertical highway. Then, they are carried along vertical highways until they reach the appropriate point for delivery.

We start considering the horizontal traffic. We consider a node sitting on the i -th horizontal highway, and compute the traffic that goes through it. Actually, a packet will travel through this node if it was generated in the i -th slab, and has a destination on the other side of the node. So, at most, our node will relay all the traffic generated in the i -th slab.

According to Lemma 6.4, a node on a horizontal highway must therefore relay at most $2Wc\sqrt{2n}/\beta$ bits per second. As the maximal distance between hops is constant ($2\sqrt{2}c$), the throughput along highways is independent of n (see Section 6.4), one can set the rate per node to $W = \Omega(1/\sqrt{n})$ without overloading links, with high probability.

The problem for vertical traffic is the dual of the previous one. We can use the same arguments, except that W now describes the *receiving rate* of the nodes. Since each node is the destination of exactly one source, the throughput per node becomes the same as above. \square

Proposition 6.3 (Delivery Phase) *Each destination node can receive packets from the highway at rate*

$$\Omega\left((\log \sqrt{n})^{-\alpha-2}(\log n)^{-1}n^{-c\kappa\gamma/2}\right),$$

for any constant $\kappa > 0$ and c so large that (6.2) holds.

Proof: The delivery phase consists in bringing the packets from the highway system to the actual destination. We proceed exactly in the same way as in Proposition 6.1, but in the other direction (horizontal delivery from the vertical highways).

We divide the network area into $\beta\sqrt{n}/c\sqrt{2}$ vertical slabs, and define a one-to-one mapping between slabs and vertical paths. We assume that packets have been transported by the highway system to their *exit point*, which is defined as the node of the vertical path closest to the horizontal line drawn from the destination. Again, the distance between exit points and destination is at most $\kappa \log \frac{\sqrt{n}}{\sqrt{2c}}$ squares. We can thus let $d = \kappa \log \frac{\sqrt{n}}{\sqrt{2c}}$ in Corollary 6.5, and conclude that each square can be served with rate $R(d) = \Omega((\log \sqrt{n})^{-\alpha-2} n^{-c\kappa\gamma/2})$. As there are at most $c^2 \log n$ in each square (Lemma 6.3), the throughput per node is at least equal to $R(d)/c^2 \log n$. \square

We are now ready to complete the proof of Theorem 6.1.

Proof of Theorem 6.1: We observe in Propositions 6.1 and 6.3 that if $c\kappa\gamma/2 < 1/2$, the asymptotic throughput per node decreases slower than $1/\sqrt{n}$. In this case, the overall throughput of the protocol is limited by the highway phase only, and the first part of Theorem 6.1 immediately follows from Proposition 6.2. We thus have to make sure that we can choose values of c and κ such that Inequality (6.2) is verified, and such that

$$\frac{c\kappa\gamma}{2} < \frac{1}{2}.$$

A possible choice is $c > 2\gamma + \sqrt{4\gamma^2 + \log 6}$ and $\kappa = \frac{1}{2c\gamma}$. We remark that these choices depend on the physical attenuation factor γ .

As for the second part of the theorem, we note that since we route packets along wandering paths of the highways, it may seem possible to have a delay higher than the optimal \sqrt{n} , which is achieved by straight line routes [eGMPS04]. However, a simple counting argument can show the opposite.

First, we can bound the average number of hops of a packet in the network. This is given by the number of hops on the highways, which is at most twice the length of a crossing path in the rectangle \bar{R}_n . As there are $\beta\kappa \log m$ disjoint crossing paths in \bar{R}_n , we have

$$\begin{aligned} & \mathbb{E}(\text{number highway hops}) \\ & \leq \frac{2}{\beta\kappa \log m} \mathbb{E} \left(\sum_{i=1}^{\beta\kappa \log m} \text{path length}_i \right) \\ & \leq \frac{2\kappa m \log m}{\beta\kappa \log m} \\ & = \frac{2}{\beta} m, \end{aligned}$$

where the last inequality holds because the total number of little squares in rectangle \bar{R}_n is $\kappa m \log m$ and the crossing paths are disjoint. Since the packet size scales with the throughput, each packet is forwarded in a constant time at each relay, and the average delay is at most proportional to \sqrt{n} , as $m = \sqrt{2n}/c$. \square

6.6 Dense Networks

In this section, we consider the model where nodes are distributed according to a Poisson point process of intensity n over a square of unit area. Furthermore, we take an attenuation function l of the form

$$l(d) = d^{-\alpha} e^{-\gamma d}.$$

In this case, we divide that network into squares of size c/\sqrt{n} . We obtain thus the same number of little squares as in the previous model. The average number of nodes in each little square is also the same, namely c^2 . Therefore, all the percolation results above still hold for this model, and we can find as many highways as above.

To derive the lower bound on the capacity, we have to compute the throughput along the highways, as well as the rate at which nodes can send data towards the highways. In fact, both of these throughputs were computed using Theorem 6.4, so it is enough here to give an adapted version of such theorem.

Theorem 6.6 *For any given integer $d > 0$ and when n is sufficiently large, there exists a TDMA scheduling, such that one node per square can transmit to any destination located within a radius of d squares (in Manhattan distance) with a fixed rate independent of n .*

When d goes to infinity, the asymptotic behavior of the rate is given by

$$R(d) = \Omega(d^{-2}).$$

Proof: We set up the same TDMA scheme as in Theorem 6.4, with k^2 time slots, where $k = 2(d+1)$. Similarly, the 8 closest interferers are located at least $d+2$ squares away from the receiver, the next 16 interferers at distance $3d+3$, and so on. The difference here is that squares have size c/\sqrt{n} here. The sum of the interferences at the receiver can be bounded as follows:

$$\begin{aligned} I(d, n) &\leq \sum_{i=1}^{\infty} 8i Pl \left(\frac{c[k_i - (d+1)]}{\sqrt{n}} \right) \\ &= \sum_{i=1}^{\infty} 8i Pl \left(\frac{c(2i-1)(d+1)}{\sqrt{n}} \right) \\ &= \sum_{i=1}^{\infty} 8i P \left(\frac{c(2i-1)(d+1)}{\sqrt{n}} \right)^{-\alpha} e^{-\gamma c(2i-1)(d+1)/\sqrt{n}} \\ &\leq P \left(\frac{c(d+1)}{\sqrt{n}} \right)^{-\alpha} e^{-\gamma c(d+1)/\sqrt{n}} K'(d), \end{aligned}$$

where

$$K'(d) = \sum_{i=1}^{\infty} 8i (2i-1)^{-\alpha} e^{-\gamma c(2i-2)(d+1)/\sqrt{n}}.$$

The latter sum clearly converges to a constant \hat{K} if $\alpha > 2$ or $\gamma > 0$.

As the receiver is at most d squares away from the emitter, the Euclidean distance between them is less than $c(d+1)/\sqrt{n}$. The strength $S(d, n)$ of the signal at the receiver reads thus

$$S(d, n) = P \left(\frac{c(d+1)}{\sqrt{n}} \right)^{-\alpha} e^{-\gamma c(d+1)/\sqrt{n}}.$$

The SINR is thus

$$\begin{aligned} \text{SINR}(d, n) &= \frac{S(d)}{N_0 + I(d)} \\ &\geq \frac{P \left(\frac{c(d+1)}{\sqrt{n}} \right)^{-\alpha} e^{-\gamma c(d+1)/\sqrt{n}}}{N_0 + P \left(\frac{c(d+1)}{\sqrt{n}} \right)^{-\alpha} e^{-\gamma c(d+1)/\sqrt{n}} K'(d)} \\ &= \frac{1}{N_0 \left(\frac{c(d+1)}{\sqrt{n}} \right)^\alpha e^{\gamma c(d+1)/\sqrt{n}} / P + K'(d)} \end{aligned}$$

When n goes to infinity, we observe that the first term of the denominator tends to zero, whereas the second term tends to \hat{K} . The whole fraction tends thus to $1/\hat{K}$, and thus

$$\lim_{n \rightarrow \infty} \text{SINR}(d, n) \geq \frac{1}{\hat{K}}.$$

It means that for any $\varepsilon > 0$, when n is large enough, the rate

$$R' = \frac{1}{2} \log\left(1 + \frac{1}{\hat{K}} - \varepsilon\right)$$

is achievable by this scheme during each time slot. This proves the first part of the theorem.

Now if d increases with n , we notice that the above limit still holds whenever $(d+1)/\sqrt{n}$ tends to zero. Therefore, if $d = O(\sqrt{n})$, the rate R' is achievable for each active transmission when n is large enough. However, as there are $k^2 = 4(d+1)^2$ time slots in our TDMA scheme, the actual throughput available for each square is $R(d) = R'/k^2$, and thus

$$R(d) = \Omega(d^{-2}).$$

□

6.7 Conclusion

We have found that the capacity of wireless networks of randomly located nodes has the same asymptotic behavior of the capacity of arbitrary networks: nodes in a random network can transmit at the same rate than nodes in an arbitrary network and there is no price to pay—at least asymptotically—for the additional randomness present in the system. This result closes a previous gap between upper and lower bounds on the optimal per-node transmission rate that consistently appeared in different proofs proposed in the literature.

It is interesting to link this result with the one in the previous chapter. Indeed, we obtained a throughput of order $\Omega(1/\sqrt{n})$ for *all* the nodes in the

network. This implies that they are all connected. The surprise is that full connectivity comes for free in this scenario. We saw that the throughput per node has to decrease like $1/\sqrt{n}$ as the number of nodes n increases because of the relay traffic. Therefore, according to the results seen in Chapter 5, the fraction of connected nodes increases to one, making the network asymptotically fully connected. In the scheme described in this chapter, full connectivity is realized by the draining phase of the protocol.

Part III

Latency

Chapter 7

Intruder detection

In this chapter, we consider a wireless multi-hop sensor network whose task is to detect the intrusion of a moving object in the area in which the sensors are deployed. The first question that comes to mind is to know how far the object can move in the monitored area before being detected by a sensor. When the object moves along a straight line, stochastic geometry provides the answer under the name of *linear contact (or hit) distribution function* (see e.g. [SKM95, page 80]). For example, when the spatial distribution of the sensors is Poisson, this function decreases exponentially with the distance traveled by the object.

The sensor that has detected the intruder needs however to convey this information to a monitoring station that collects the sensed data, the *sink*. Nodes between the detecting sensor and the sink act as relays for the message. However, various sources of noise, battery failures in the nodes or simply their random locations, will inevitably result in having some sensors disconnected from the network. If the sensor that detected the intruder is not connected to the sink, the intruder can continue to progress in the monitored area without actually being detected by the central sink. Consequently, a second, more difficult question is to know the distance that can be traveled by the intruder until the sink (and not only any sensor) can be notified. This chapter addresses precisely this question.

Tracking and detecting a moving object is an important application of a sensor network, and has thus received some attention. Most of the work is so far devoted to the problem of computing the linear contact (or hit) distribution function, i.e. the distance traveled by the intruder until detection by a sensor without checking that the sensor that spawns the alarm is actually connected to the network. The simplest case is to compute this distance when the network is modeled by a Boolean model; it is known to be exponentially distributed (see e.g. [SKM95], p.80). More results from this approach are derived in [LBD⁺05]. When the motion of the object does not follow a straight line, but a Brownian motion, explicit formulas and/or bounds are obtained in [KKP03]. To save energy, nodes periodically switch off their batteries, which implies trade-offs between the speed of detection of an intruder and the energy savings incurred by rendering many nodes into sleep mode. This trade-off is examined in [GM04].




Figure 7.1: The intruder moving along a straight line across the sensor network. Since it is moving with constant speed, the time it takes before the intruder gets caught is proportional to the distance from its starting point to the giant cluster. In this figure, disconnected clusters are represented in gray.

7.1 Model

We assume that the sensor network can be described by a 2-dimensional Poisson Boolean model $\mathcal{B}(\lambda, R)$: nodes (sensors) are distributed according to a 2-dimensional homogeneous Poisson process of intensity λ , and their connectivity range $2R_i$ are randomly and identically distributed, independently from all other nodes and from the Poisson process. We also assume that $\lambda > \lambda^*$, so that there exists a giant cluster of connected sensors (see Chapter 2). Furthermore, we assume that the sink belongs to the giant cluster, so that it can be reached by messages from most of the sensors.

Each sensor i can sense the intruder within a certain *sensing range*, which is assumed here to be exactly equal to R_i , half of its connectivity range. In this way, we can use the same Boolean model $\mathcal{B}(\lambda, R)$ to describe connectivity and detection capacity. Indeed, two nodes are connected when their sensing range overlap.

We suppose that the intruder starts at the origin and moves along a straight line with constant speed in any arbitrary direction, as shown in Figure 7.1. The time needed by the network to detect the intruder is therefore proportional to the distance from the origin to the giant cluster along the direction chosen by the intruder.

Without loss of generality, in this chapter, we will assume that the intruder moves along the positive part of the x -axis $\{(x, 0) : x > 0\}$.

7.2 Theorem formulation and insight of the proof

We consider the set of vertices \mathbb{Z}^2 , and denote by \mathbb{L}^2 the set of the edges joining all adjacent vertices (i.e. $\mathbb{L}^2 = \{(x, y) \in \mathbb{Z}^4 : |x_1 - y_1| + |x_2 - y_2| = 1\}$). We construct the standard independent bond percolation model by declaring each

edge *open* with probability p , independently of all the other edges, and *closed* with probability $1 - p$. We say that two vertices are *connected* if there exists a sequence of contiguous open edges that joins them. This definition yields the partitioning of \mathbb{Z}^2 into *connected components* (or *clusters*).

We know from percolation theory that if p is greater than the *percolation threshold* $p^* = 1/2$, then the probability $\theta(p)$ that the origin belongs to an infinite cluster is strictly positive. Furthermore, in this case, the infinite cluster (also called *giant cluster*) is unique a.s..

Our main result states that the distance between the origin and the infinite cluster is upper bounded by a geometric random variable. We denote by N the coordinate of the first vertex on the right part of the horizontal axis that belongs to the giant cluster. Thus, if the origin belongs to the giant cluster, we have $N = 0$.

Theorem 7.1 *When $p > p^*$, there exist constants c_1 and $c_2 \in \mathbb{R}^+$ such that*

$$P_p(N > n) \leq c_1 \exp(-c_2 n).$$

7.3 Construction and its independence properties

To prove Theorem 7.1, we will perform the following construction: we start at the origin, and look at the set of points of the half-plane $\mathbb{H} = \mathbb{Z}^+ \times \mathbb{Z} = \{x = (x_1, x_2) \in \mathbb{Z}^2 \mid x_1 \geq 0\}$ (all the points located on the right hand side of the vertical axis) that are connected to it. If this set is infinite, then $N = 0$. If not, we look at the size of this connected component. More precisely, if we denote by $\partial B^+(n)$ the set of vertices $x = (x_1, x_2)$ of \mathbb{H} such that $x_1 = n$ or $|x_2| = n$, and by $\{0 \leftrightarrow S\}$ the event that the origin is connected to some element of the set S , we look at the events $\{0 \leftrightarrow \partial B^+(n)\}$, $n \in \mathbb{N}$, and find the first integer n such that this event is not true. Note that this can be done iteratively, starting with $n = 1$ and increasing its value one by one. Therefore, if we define the variable $M = \min\{n > 0 : 0 \not\leftrightarrow \partial B^+(n)\}$, the event $\{M = n\}$ only depends on the state of the edges inside the box $B^+(n) = [0, n] \times [-n, n]$.

Once we have found the value of M , we iterate the same procedure, starting at the point $(M, 0)$, and looking at the connected vertices in the half-plane $\mathbb{H} + (M, 0) = \{x = (x_1, x_2) \in \mathbb{Z}^2 \mid x_1 \geq M\}$. Clearly, the outcome of this second step is independent of what happened in the first step, because we only look at the half-plane located to the right of the box $B^+(n)$.

Then we iterate this algorithm until we hit an infinite connected component in the half-plane. The coordinate of the stopping point of the algorithm (denoted hereafter by N') gives an upper bound of N .

However, we have to check that this algorithm stops after a finite number of steps a.s. whenever an infinite cluster exists in the whole plane. In fact, it is enough to show that the probability $\theta^+(p)$ that the origin belongs to an infinite cluster in \mathbb{H} is strictly positive (in this case, the number of steps before stopping is a geometric random variable). This result follows directly from Theorem 7.2 in [Gri99]. Let $B(k) = [-k, k]^d$ be the box with side length $2k$ centered at the origin.




Figure 7.2: The construction of the proof. Starting from the origin, the first cluster we meet is finite. We have that $0 \leftrightarrow \partial B^+(4)$. Thus, $M = 4$ in this example. Then we move to $(4, 0)$ and look only at the edges on the right, that are independent of M . The cluster at $(4, 0)$ is also finite, so we move to $(7, 0)$. The cluster at $(7, 0)$ is infinite, and therefore $(7, 0)$ belongs to the giant component. We conclude that $N' = 7$, and thus that $N \leq 7$.

Theorem 7.2 [Gri99, page 148] For $d \geq 2$, let F be an infinite connected subset of \mathbb{Z}^d with percolation threshold $p^*(F) < 1$. For each $\eta > 0$ there exists an integer k such that $p^*(2kF + B(k)) \leq p^* + \eta$.

By choosing $F = \mathbb{H}$ we have $p^*(2kF + B(k)) \rightarrow p^*$ as $k \rightarrow \infty$ and the theorem implies that this is true when $p^*(F) = p^*$ (this result is also given in [Gri99, page 162]). In other words, the percolation threshold is the same in the right half plane \mathbb{H} as in the entire plane \mathbb{Z}^2 .

7.3.1 The size of finite clusters

According to our algorithm, the random variable N' is the sum of a random number (the number of steps) of i.i.d variables (identical to the variable M above). In this section, we derive a bound on the probability that M is larger than a given number n . Note that this amounts to bounding the radius of a finite cluster in the half-plane.

Let C^+ be the open cluster at the origin on the half-plane \mathbb{H} and consider the event of an open path from the origin to the surface $\partial B^+(n)$ when C^+ is finite. We claim that there exist finite constants $k_1, k_2 > 0$ for $p^* < p < 1$ such that

$$P_p(0 \leftrightarrow \partial B^+(n), |C^+| < \infty) \leq k_1 n^2 e^{-k_2 n} \text{ for all } n.$$

This result is a fairly straightforward extension of Theorems 8.18 and 8.21 in [Gri99]. For the open cluster C at the origin of \mathbb{Z}^2 and the box $B(n) = [-n, n]^d$ with surface $\partial B(n)$, the theorems state that there exist finite constants $A(p, d), \sigma(p) > 0$ for $d \geq 2$ and $p^* < p < 1$ such that

$$P_p(0 \leftrightarrow \partial B(n), |C| < \infty) \leq A(p, d) n^2 e^{-\sigma(p)n} \text{ for all } n.$$

In our case, the cluster C^+ in \mathbb{H} results from the restriction of the cluster C in \mathbb{Z}^2 to its vertices and edges lying in \mathbb{H} .

We proceed by modifying the construction proofs for Theorem 8.18 and Lemma 8.27 of [Gri99], so that the result applies to the cluster C^+ when $|C^+| < \infty$. Rather than repeating the entire proof here, we refer the reader to [Gri99] and highlight the necessary modifications. Firstly, we replace C in the proof with the cluster C^+ where $x_1 \geq 0$ for every $(x_1, x_2) \in C^+$. The cluster has minimum and maximum extremities in the horizontal coordinate direction $L_1 = \min\{x_1 : (x_1, x_2) \in C^+\}$ and $R_1 = \max\{x_1 : (x_1, x_2) \in C^+\}$ (similarly L_2 and R_2 for the vertical coordinate) and diameter $\text{diam}(C^+) = \max\{R_i - L_i : 1 \leq i \leq 2\}$. We follow the same construction of two finite clusters whose widths in the horizontal coordinate direction are m and n such that the clusters only lie in the half-plane, one to the right of the other. Following the proofs of Lemma 8.27 and Theorem 8.18 of [Gri99], we obtain the same upper and lower bounds for $\text{diam}(C^+)$ as for $\text{diam}(C)$. In particular, we find the bound

$$P_p(\text{diam}(C^+) = k) \leq \frac{d^2}{p^2(1-p)^{2d-2}} (2k+1)^d e^{-(k+2)\sigma(p)}.$$

Since the diameter of C^+ is at least as large as its radius, for the half box $B^+(n) = [0, n] \times [-n, n]$ we have

$$\begin{aligned} & P_p(0 \leftrightarrow \partial B^+(n), |C^+| < \infty) \\ & \leq P_p(n \leq \text{diam}(C^+) < \infty) \\ & \leq \frac{d^2}{p^2(1-p)^{2d-2}} \sum_{m=n}^{\infty} (2m+1)^d e^{-(m+2)\sigma(p)}. \end{aligned}$$

This equation gives us the constants k_1 and $k_2 = \sigma(p)$.

To show that $k_2 > 0$ when $p > p^*$, we must consider Theorem 8.21 in [Gri99] (which, for the case $d = 2$, is implied by Theorem 11.24), and show that it is also true for the finite cluster on the half-plane. We refer the reader to the proof of Theorem 11.24 which shows that $0 < \xi(p) = \xi(1-p) < \infty$ when $p > p^*$, where $\xi(p) = \sigma^{-1}(p)$ is known as the correlation length. We adapt the proof by considering the whole plane \mathbb{Z}^2 but treating edges on the left half of the plane as closed with probability 1 to imply results on \mathbb{H} . Consider the vertex $e_n = (n, 0)$ and let $\tau_p^{f+}(0, e_n)$ denote the probability that the origin lies in a finite open cluster containing the vertex e_n in \mathbb{H} . We need to find an upper bound to this probability. We slightly modify the definition of the event A_n [Gri99, page 298]: denoting by $Y^+ = \{(-\frac{1}{2}, k + \frac{1}{2}) : k \geq 0\}$ and by $Y^- = \{(-\frac{1}{2}, -k - \frac{1}{2}) : k \geq 0\}$ the two vertical half-axes, we define the events A_n (respectively B_n) that some vertex in Y^+ (respectively Y^-) is joined by a closed path in the dual lattice to some vertex in X_d^+ . Using the BK inequality [Gri99, Theorem 2.15], we have that

$$\tau_p^{f+}(0, e_n) \leq P_p(A_n)P_p(B_n) = P_p(A_n)^2,$$

because $P_p(A_n) = P_p(B_n)$. Now,

$$\begin{aligned} P_p(A_n) & \leq \sum_{k=0}^{\infty} \sum_{l=n}^{\infty} P_p\left(\left(-\frac{1}{2}, k + \frac{1}{2}\right) \text{ joined to } \left(l + \frac{1}{2}, \frac{1}{2}\right) \right. \\ & \quad \left. \text{by a closed dual path}\right) \\ & \leq \sum_{k=0}^{\infty} \sum_{l=n}^{\infty} P_{1-p}((0, k) \leftrightarrow (l, 0)). \end{aligned}$$

However, $1 - p < p^*$ since $p > p^*$, and hence

$$P_{1-p}((0, k) \leftrightarrow (l, 0)) \approx e^{-(l+k)/\xi(1-p)}$$

as $k + l \rightarrow \infty$, from which we deduce, as in [Gri99],

$$\sigma(p) = \liminf_{n \rightarrow \infty} \left\{ -\frac{1}{n} \log \tau_p^{f^+}(0, e_n) \right\} \geq \frac{2}{\xi(1-p)}.$$

7.3.2 Final computation

In the two previous subsections, we showed that there exist $k_1, k_2 > 0$ such that

$$P_p(M > n) \leq k_1 n^2 e^{-k_2 n} \text{ for all } n. \quad (7.1)$$

As $k_2 > 0$, we can choose a new constant $e^{-k_2} < p_1 < 1$, and find a constant $k_3 \in \mathbb{N}$ such that

$$k_1 n^2 e^{-k_2 n} \leq (1 - p_1) p_1^{n - k_3}. \quad (7.2)$$

Let us define a sequence of i.i.d geometric random variables $\{X_i\}_{i \in \mathbb{N}}$ of parameter p_1 , whose probability mass function is $\mathbb{P}(X_i = n) = (1 - p_1) p_1^n$ for all $n \in \mathbb{N}$. Since $\mathbb{P}(X_i > n - k_3) = p_1^{n - k_3}$, combining (7.1) and (7.2), we have that

$$\mathbb{P}(M > n) \leq \mathbb{P}(X_i + k_3 > n).$$

Moreover, the variable N' (describing the stopping point of the algorithm of the beginning of Section 7.3) is the sum of K independent variables M_i , identically distributed as M (note that $M_1 = M$), where K is the number of steps in the algorithm:

$$N' = \sum_{i=1}^K M_i.$$

As the algorithm stops when we hit an infinite cluster in \mathbb{H} , and this happens with independent probability at each step (denoted by $\theta^+(p)$), we have $\mathbb{P}(K = n) = \theta^+(p)(1 - \theta^+(p))^n$. Thus, K is also a geometric random variable with parameter $p_2 \doteq 1 - \theta^+(p)$. We can write

$$\begin{aligned} \mathbb{P}(N' > n) &= \mathbb{P}\left(\sum_{i=1}^K M_i > n\right) \\ &\leq \mathbb{P}\left(\sum_{i=1}^K (X_i + k_3) > n\right) \\ &= \mathbb{P}\left(\sum_{i=1}^K X_i + k_3 K > n\right). \end{aligned} \quad (7.3)$$

The rest of this section is devoted to the computation of the above probability, and is purely technical.

We define an auxiliary random variable $N'' = \sum_{i=1}^K X_i$. Its generating function is $G_{N''}(z) = G_K(G_X(z))$, where G_K and G_X denote the generating functions of K and X_i , respectively. We have thus

$$\begin{aligned} G_{N''}(z) &= \frac{p_2}{1 - (1 - p_2) \frac{p_1}{1 - (1 - p_1)z}} \\ &= p_2 + (1 - p_2) \frac{p_3}{1 - (1 - p_3)z}, \end{aligned}$$

where $p_3 = p_1 p_2 / (1 - p_1 + p_1 p_2)$. We see that N'' is a geometric variable of parameter p_3 , with some extra weight on $\{N'' = 0\}$. To obtain the probability in (7.3), we must consider the sum of N'' and $k_3 K$. The tail of the distribution of this sum decreases exponentially, as fast as the slowest term. Here, depending on the value of p_1 , p_2 and k_3 , either N'' or $k_3 K$ dominates the sum in (7.3). In all cases, there exists $k_4 \in \mathbb{N}$ such that

$$\mathbb{P}(N'' + k_3 K > n) \leq p_4^{n-k_4},$$

with $p_4 = \max\{p_3, p_2^{1/k_3}\} < 1$. Finally as

$$\mathbb{P}(N > n) \leq \mathbb{P}(N' > n) \leq \mathbb{P}(N'' + k_3 K > n),$$

we have

$$\mathbb{P}(N > n) \leq p_4^{n-k_4} = c_1 \exp(-c_2 n),$$

for $c_1 = p_4^{-k_4}$ and $c_2 = -\log(p_4)$.

7.4 Extension to the continuous case

In this section, we address a different percolation model –the Poisson Boolean model– defined as follows: we consider a homogeneous Poisson point process of intensity λ over the plane \mathbb{R}^2 . On each point of the process we center a ball, whose radius R is also random, independent of and identically distributed as the other balls' radii. This divides the plane into an *occupied region*, and a *vacant region*, which is not covered by any ball.

Given the distribution of R such that $\mathbb{E}(R^2) < \infty$, there exists a critical density λ^* such that if the actual density λ of the point process is greater than λ^* , then the probability that the origin belongs to an unbounded occupied connected component is strictly positive. Furthermore, above the threshold, the unbounded occupied connected component is unique a.s..

In the following theorem, we show that the distance between the origin and the unbounded component along a given direction can be bounded from above by an exponential random variable.

Theorem 7.3 *Let D denote the coordinate of the first point on the positive part of the horizontal axis that belongs to the unbounded component ($D = 0$ if the origin is already in the unbounded component). There exist c_3 and $c_4 \in \mathbb{R}^+$ such that*

$$\mathbb{P}(D > x) \leq c_3 \exp(-c_4 x).$$

We prove this theorem by mapping the continuous model onto a (discrete) bond percolation model, and applying Theorem 7.1.

7.4.1 Mapping

We use a mapping similar to that used in Section 6.3. We declare each edge $\langle x, y \rangle \in \mathbb{L}^2$ *open* if the conditions below are fulfilled in the Boolean model (we assume without loss of generality that $x_1 \leq y_1$ and $x_2 \leq y_2$, see Figure 7.3):




Figure 7.3: A horizontal edge $\langle x, y \rangle$ that fulfills the three conditions.

- There exists an occupied crossing in the rectangle $[x_1d, (y_1 + 1)d] \times [x_2d, (y_2 + 1)d]$ in the direction of its longest edge.
- There exists a vertical and an horizontal crossing in the square $[x_1d, (x_1 + 1)d] \times [x_2d, (x_2 + 1)d]$.
- There exists a vertical and an horizontal crossing in the square $[y_1d, (y_1 + 1)d] \times [y_2d, (y_2 + 1)d]$.

These conditions depend on a parameter d , that determines the size of the rectangle. We know that if the Boolean model is supercritical, the probability that the three conditions are fulfilled tends to one when d tends to infinity (see e.g. Corollary 4.1 in [MR96]). However, in this mapping, the status of two contiguous edges are *not* independent. We constructed thus a dependent bond percolation model that has the desired property that any of its connected components corresponds to a connected component in the Boolean model. Therefore, if the first point on the horizontal axis that belongs to the infinite cluster has the horizontal coordinate N in the discrete model, then there exists a point of the Boolean model in the interval $[Nd, (N + 1)d]$ that belongs to the infinite cluster. It follows from this observation that if Theorem 7.1 is valid in this dependent bond percolation model, so is Theorem 7.3.

7.4.2 Domination of the dependent model

To validate Theorem 7.1 in the above dependent bond percolation model, we use the results in [LSS97]. Let $(N, 0)$ be the first point of $\mathbb{Z}^+ \times \{0\}$ (i.e. the one with the smallest coordinate) that belongs to the infinite connected cluster. Consider the 1-dependent bond percolation model defined above, where edge e is open with some probability $0 \leq p' \leq 1$ independently of edges separated by distance greater than 1. The claim is that $\mathbb{P}(N \leq n) \geq \alpha(n)$ for the 1-dependent model is implied when it is true for the independent model. That is, the probability to hit the infinite open cluster within n steps is at least as great on the dependent graph as on the independent graph.

We show this by employing Theorem 0.0 in [LSS97] which states that the k -dependent random field of $\{0, 1\}$ random variables $(X_s)_{s \in \mathbb{Z}^2}$ stochastically

dominates a translation invariant product random field $(Y_s)_{s \in \mathbb{Z}^2}$ of density $p'' > 0$, whenever p' is sufficiently large. Furthermore, $p'' \rightarrow 1$ when $p' \rightarrow 1$. If we consider the indicator function $1_{\{N \leq n\}}$ for the event $\{N \leq n\}$, where $1_{\{N \leq n\}}$ is an increasing function on our state space Ω , the theorem tells us that by stochastic domination

$$\int 1_{\{N \leq n\}} dP_{p'} \geq \int 1_{\{N \leq n\}} dP_{p''}. \quad (7.4)$$

The result of integration gives $P_{p'}(N \leq n) \geq P_{p''}(N \leq n)$. Finally, by increasing the mapping parameter d , one can make p' large enough, so that $p'' > 1/2$ and Theorem 7.1 becomes valid for the field $(Y_s)_{s \in \mathbb{Z}^2}$.

7.5 Conclusion

In this chapter, we showed that the distance traveled by an object moving along a straight line in a supercritical percolation model until it hits the giant component is bounded from above by a (shifted) exponential variable. It was already known that the distance before hitting *any* connected component is exponentially distributed. With the results of this chapter, we now have a lower and an upper bound for the distribution of the distance traveled before hitting the giant component.

This distribution is very relevant in the context of detection of an intruder in wireless sensor networks, as outlined in the introduction, as well as that of connectivity in mobile ad hoc networks. Indeed, finite connected components are essentially disconnected parts of the network, as they cannot exchange data with the vast majority of the other nodes (those who belong to the giant component). Therefore, from the point of view of a mobile user, the value of interest is the time before connecting to the giant component. The bound computed in this work provides some guarantee about connectivity in this very simple mobile scenario.

Chapter 8

An uncoordinated power saving mechanism

In this chapter, we consider a wireless multi-hop network with a particular energy saving mechanism. To reduce their consumption, nodes periodically turn off their radio device and therefore disappear from the wireless network. The on/off transitions happen at random times, independently for each node. To describe this scenario, we introduce an extension of the static Boolean model called *blinking Poisson Boolean model*. We are interested in the delays introduced in the network by this mechanism, and study the propagation speed of a broadcast message in this new model. As we will see in the sequel, this problem is closely related to first passage percolation.

The motivation for this study is the following. We assume a wireless sensor network, made of a large number of nodes with limited processing and communication capabilities. Each node performs some sensing of a particular confined area, and sends the result to a data collecting node (called sink) in a multi-hop fashion, using other nodes as relays. These networks need to provide at the same time a good *coverage* of the area to be monitored and a good *connectivity* of the network.

Since sensors are usually cheap devices, compared to their deployment costs, it is not necessary to have full connectivity. As we have seen in Section 4.1, we can set the node density so that a large fraction of the sensors are well connected. At the contrary, energy consumption is often a much more critical variable, because of the limited battery that can be put in a sensor as well as of the cost of replacing a node that has failed. Energy is consumed by sensors in their sensing, processing and communicating tasks. Sensing has to be done at a periodicity dictated by the monitored event. The energy it consumes can be reduced if the area covered by a single sensor is small (and therefore if the number of sensors is large enough). Processing and communication energy consumption depends on the hardware, but also on the way data is aggregated and medium is accessed. This consumption is the one offering probably the largest potential for reduction. Indeed, nodes spend a considerable amount of energy in listening to their neighbors, and as long as none of these neighbors transmits any data, this energy is simply wasted. Data collection and medium access control (MAC) schemes need therefore to incorporate energy saving as a

primary goal, and proposals for such algorithms have recently emerged in the literature, which show indeed that significant energy savings can be achieved. In particular, most proposals for energy saving MAC schemes for sensor networks introduce a sleeping mode for nodes, during which practically no energy is spent [SC01, WHE02].

Introducing a sleeping mode does however come at some cost. A solution is to use a Time Division Multiple Access (TDMA) scheme, but this requires nodes to synchronize with each other quite tightly, which can be a quite complex task in large networks with random node locations and imperfect (drifting) clocks. Letting the nodes set their wake-up and sleeping times in a decentralized fashion reduces this complexity, but this increases the delay (also called latency) to transfer information between the sink and a distant node. Pushing the decentralization to an extreme where nodes go to sleep independently from each other, which is the solution we adopt in this model, eliminates the complexity of having synchronized clusters of nodes, but at the same time raises concerns about an increase of the latency of the network. More importantly, it will not only increase the average latency itself, but it will also increase the variance of this latency. For some applications, such as spatial data collection for statistical purposes, this is certainly acceptable, but not for many others that are much more time-critical. A typical and important example of such a scenario is the use of sensor network for monitoring an area and sending an alarm when an abnormal event is sensed occurs (such as in the scenario considered in Chapter 7). Even if some fixed amount of latency can be tolerated, a highly variable latency due to the random position of the nodes, the random radio range, the non-synchronized or even random sleeping and active periods is much more problematic, and must therefore be studied.

In this chapter, we will show that the ratio between the time elapsed and the distance traveled by the broadcast message tends to a deterministic constant. Moreover, the asymptotic shape of the region flooded by this broadcast is circular.

8.1 Decentralized energy saving mechanism

The general strategy we adopt to save energy and yet bound the latency of the network is as follows.

In the absence of any incoming event or message from a neighbor, nodes switch between their sleeping and active phases independently from each other. To save energy, we assume that the sleeping times are much longer than the active times.

Once a node has sensed an event, it stays active, and keeps sending repeatedly this information to its immediate neighbors, which can be reached within a single hop. The neighbors will only be able to hear the message when they turn their radio on. When all the neighbors have heard that message, the sensing node can turn off its radio, and resume its regular on/off schedule. In order to avoid synchronization, the phase of the on/off schedule is drawn according to the stationary distribution. The other nodes will then remain active until all their neighbors have received the message, and so on, until the sink eventually receives the message.

Remember that nodes are scattered over the domain that needs to be moni-

tored, according to a Poisson process. We suppose that there is a large number of nodes, so that the sensing radius can be set to a small value without impacting the coverage of the network. However, since nodes switch between the two phases independently from each other, there may be quite few active nodes at any given time t , so that the set of simultaneously active nodes at any particular time t is always disconnected. This set of simultaneously active nodes represents however only a snapshot of the network at this particular time. By waiting long enough, a sufficiently large number nodes will have passed by an active phase, so that the superposition of all the different snapshots of the network at all times within an interval becomes supercritical. On this cumulative graph, there is a path between any sensor and the sink with high probability. While it is relatively easy to show that indeed the network will pass by a supercritical phase if we wait long enough (as we will show in Proposition 8.2), it is much more difficult to find how long we have to wait before this path is almost surely present, that is, to compute the latency. Clearly, the larger the distance between the sensing node and the sink, the larger the latency. But how does it depend on this distance? The bounds on the latency, which we establish in Theorem 8.1, show that it is linear, under the conditions stated in the introduction. The simulations, run under more general conditions (including propagation delays on the links, and a random connectivity radius) confirm this finding.

8.2 Blinking Poisson Boolean model

In order to take into account the alternation between the sleeping and active modes, we introduce a dynamic Poisson Boolean model, where the dynamics is due to the “blinking” of the nodes. As far as we know, only few dynamic percolation models have been researched mathematically. In the lattice, dynamic bond percolation have been studied by Häggström, Peres and Steif [HPS97, PS98]. Poisson Boolean models with moving points have been studied by van den Berg et al. [vdBMW97].

The *blinking Poisson Boolean model* $\mathcal{B}(\lambda, R, Z)$ is a simple modification of the static Poisson Boolean model. As in the static Poisson Boolean model, the positions of the nodes are determined by a Poisson point process with intensity λ . At each node, we attach a disk whose radius is distributed as R (deterministic or i.i.d. random radii). The dynamics follows from the assumption that the nodes alternate between on-state and off-state with periods determined by the stationary i.i.d. Bernoulli processes $Z^{(i)}$, ($Z^{(i)} = 1$ means that Node i is *on* and $Z^{(i)} = 0$ means that Node i is *off*). The distributions of these processes are equal to the distribution of process Z .

We assume that the lengths of the on/off periods are independent, with the off periods either constant or exponentially distributed with mean t_{off} , and the on periods distributed according to an arbitrary distribution with mean t_{on} . The only purely technical assumption we require for the on periods is that they have always nonzero length, which we can write as

$$\lim_{\delta \rightarrow 0} \mathbb{P} \left[\min_{t \in [s, s+\delta]} Z_t = 1 \mid Z_s = 1 \right] = 1. \quad (8.1)$$

This is always the case in practice. On the other hand, assuming that $\delta < t_{\text{off}}$,

the off periods satisfy

$$\mathbb{P}\left[\max_{t \in [s, s+\delta)} Z_t = 0 \mid Z_s = 0\right] = \begin{cases} e^{-\delta/t_{\text{off}}}, & \text{if exp-distributed,} \\ 1 - \frac{\delta}{t_{\text{off}}}, & \text{if constant.} \end{cases} \quad (8.2)$$

The stationary distribution of Z is given by

$$\pi_{\text{off}} \doteq \mathbb{P}(Z = 0) = \frac{t_{\text{off}}}{t_{\text{off}} + t_{\text{on}}} \quad \text{and} \quad \pi_{\text{on}} \doteq \mathbb{P}(Z = 1) = \frac{t_{\text{on}}}{t_{\text{off}} + t_{\text{on}}}.$$

Assume that the node density is supercritical, that is, $\lambda > \lambda^*$. From the point of view of sensor networks, it is interesting to know whether the network remains connected all the time, despite the alternation between sleeping and active states. If there is no unbounded connected component of active nodes at any snapshot, then another question is how long one has to wait in order to get a completely connected network in the cumulative coverage process. The two situations are visualized in Figure 8.1 and answers are given by the following propositions.




Figure 8.1: Snapshots of the Blinking Poisson Boolean model at some particular time t . The black disks are the coverage area of the active nodes, the gray ones of the sleeping nodes. On the left $\lambda > \pi_{\text{on}}\lambda > \lambda^*$, i.e., there is always an unbounded active cluster. On the right, $\pi_{\text{on}}\lambda < \lambda^* < \lambda$, i.e., all the active clusters are finite a.s.

Proposition 8.1 *Assume $\lambda > \lambda^*$. If the $Z^{(i)}$ are stationary on/off processes with exponential or constant off times and on periods satisfying (8.1), then*

$$\mathbb{P}(\text{an infinite cluster exists for all } t \geq 0) = 1, \text{ if } \lambda^* < \lambda\pi_{\text{on}},$$

$$\mathbb{P}(\text{there is no infinite cluster for all } t \geq 0) = 1, \text{ if } \lambda^* > \pi_{\text{on}}\lambda.$$

Proof: A slightly modified proof from [HPS97]. Let us assume first $\lambda\pi_{\text{on}} > \lambda^*$. Let $\epsilon > 0$ such that $(1 - \epsilon)\pi_{\text{on}} > \lambda^*/\lambda$ and take $\delta > 0$ such that

$$\mathbb{P}\left[\min_{t \in [0, 0+\delta)} Z_t = 1 \mid Z_0 = 1\right] > 1 - \epsilon.$$

Then $p_\delta \doteq \mathbb{P}\left(\min_{t \in [0, \delta]} Z_t^{(i)} = 1\right) > (1 - \epsilon)\pi_{\text{on}} > \lambda^*/\lambda$ for an arbitrary i . Since events $\{\inf_{t \in [0, \delta]} Z_t^{(i)} = 1\}$, $i = 1, 2, \dots$, are mutually independent, we can consider a thinning of a Poisson process where we take only the nodes which are active the whole interval $[0, \delta]$. These nodes are distributed according to a Poisson process with intensity $p_\delta \lambda > \lambda^*$. Thus

$$\mathbb{P}(\text{an infinite cluster exists for all } t \in [0, \delta]) = 1.$$

The argument can be repeated for the intervals $[k\delta, (k+1)\delta]$ with integer k . Denote E_k the event that an infinite cluster exists for all $t \in [k, (k+1)\delta]$, and E_k^c its complement. Then

$$\mathbb{P}\left(\bigcap_k E_k\right) = 1 - \mathbb{P}\left(\bigcup_k E_k^c\right) \geq 1 - \sum_k \mathbb{P}(E_k^c) = 1.$$

If $\lambda\pi_{\text{on}} < \lambda^*$, then we consider the nodes which are not sleeping during the whole interval $[0, \delta]$. By Equation (8.2), for small enough δ ,

$$\mathbb{P}\left(\max_{t \in [0, \delta]} Z_t = 1\right) < \pi_{\text{on}} + \pi_{\text{off}}\delta/t_{\text{off}} < \lambda^*/\lambda.$$

Thus the intensity of those points is less than the critical intensity and there is a.s. no unbounded cluster. Again, the countable additivity completes the proof. \square

Even if $\pi_{\text{on}}\lambda < \lambda^*$ we can still have some form of connectivity if we take into account the cumulative coverage process. Let us define the cumulative Blinking Boolean model as the area which has been covered by some active disk within $[0, t]$.

Proposition 8.2 *Assume $\lambda > \lambda^*$. If the $Z^{(i)}$ are i.i.d. stationary on/off processes with exponential off-times, then the cumulative connectivity graph on $[0, t]$ has almost surely an infinite connected component whenever $t > t_{\text{off}} \log \frac{\lambda - \lambda^*}{\lambda - \lambda\pi_{\text{on}}}$. If the off periods last constant time then the condition is $t > t_{\text{off}} \frac{\lambda^* - \lambda\pi_{\text{on}}}{\lambda - \lambda\pi_{\text{on}}}$.*

Proof: The probability that process Z visits at least once the on-state within $[0, t]$ is

$$\begin{aligned} & \mathbb{P}\left(\max_{s \in [0, t]} Z_s = 1\right) \\ &= \mathbb{P}(Z_0 = 1) + \mathbb{P}\left[\max_{s \in [0, t]} Z_s = 1 \mid Z_0 = 0\right] \mathbb{P}(Z_0 = 0) \\ &= 1 - \pi_{\text{off}} \mathbb{P}\left[\max_{s \in [0, t]} Z_s = 0 \mid Z_0 = 0\right]. \end{aligned}$$

The points of the Poisson process, where the corresponding on/off processes visit the on-state during $[0, t]$, form a thinned Poisson point process with intensity $1 - \pi_{\text{off}} \mathbb{P}[\max_{s \in [0, t]} Z_s = 0 \mid Z_0 = 0]$. The corresponding Poisson Boolean model has an unbounded cluster a.s. if $\lambda(1 - \pi_{\text{off}} \mathbb{P}[\max_{s \in [0, t]} Z_s = 0 \mid Z_0 = 0]) > \lambda^*$. In order to complete the proof, apply Equation (8.2). \square

The next section shows how these results are useful for determining the coverage and connectivity of the network.

8.3 Sensing coverage and transmission connectivity




Figure 8.2: A sensing coverage on the left and the corresponding transmission network on the right when transmission radius is 5 times bigger than the sensing radius.

We assume that the radio equipments are alternating between sleep and active states with mean sojourn times t_{off} and t_{on} and stationary distributions π_{off} and π_{on} . The sensing apparatus are either active all the time or they also interchange according to some on/off processes, with parameters t_{off}^s and t_{on}^s and stationary distributions π_{off}^s and π_{on}^s .

For the sensing model, we let the radii be random, but for the transmission model the radii are assumed to be constant and denoted by $r = 2r_b$. Note that if two disks of radius r_b overlap, then the centers are at most r apart. Thus the radio connectivity graph is determined by Boolean model with disks of radius r_b .

8.3.1 Sensing area

For some networks, the covered fraction of the area is the key property, for others the probability of unwatched routes through the network. Thus, depending on the application, either the mean area coverage or the existence of a percolation cluster are the characteristics which give the conditions for the minimal density of the sensors.

Assume that each sensor can monitor a disk whose radius is distributed as R . As explained in the previous section, the mean fraction of area covered is given by $1 - \exp(-\lambda\pi_{\text{on}}^s\mathbb{E}(R^2))$. To ensure a covered fraction at least equal to p_A , the average number of active sensors per unit area $\lambda\pi_{\text{on}}^s$ has to satisfy $\lambda\pi_{\text{on}}^s \geq \frac{-1}{\pi\mathbb{E}(R^2)} \log(1 - p_A)$.

Regarding unwatched routes, percolation theory gives the critical intensities guaranteeing that there is an unbroken net guarding any passage through the network. To ensure a connected covered area at all instants, it is enough to have $\lambda\pi_{\text{on}}^s > \lambda^*$, according to Proposition 8.1.

If alarms are triggered by, for example, slowly moving events, then it is enough to obtain connectivity over a longer time period. In this case, Proposition 8.2 gives the sufficient conditions with respect to the intensity and the

mean lengths of the on/off periods. More details on the coverage properties of sensor networks can be found in [LT03, LBD⁺05].

8.3.2 Radio connectivity

Connectivity can be seen as the probability that an arbitrary node is connected to most of the others.

If we assume that the radio range of the devices is significantly larger than their sensing range, the conditions on the node density given in the above section lead to a highly super-critical radio connectivity graph (see Figure 8.2). In fact, almost all the nodes are connected in this case, and there exist highly redundant routes between nodes. This redundancy is the motivation for letting the nodes turn off their radio device sporadically.

In our mechanism, at each time instant, the number of nodes with active radio device is $\pi_{\text{on}}\lambda$. This defines a new static Poisson Boolean model $\mathcal{B}(\pi_{\text{on}}\lambda, r_b)$. Depending on the value of π_{on} , the new model can be either super-critical, either sub-critical, as shown in Proposition 8.1. In order to spare as much battery as possible, in this work, we choose π_{on} so low that the resulting process is sub-critical. Even in this case, messages can be carried from almost any node of the network to almost any other node. More precisely, messages can be exchanged between any two nodes that belong to the infinite cluster of $\mathcal{B}(\lambda, r_b)$. The transmission latency with respect to the distance between source and destination nodes is studied in detail in Section 8.4. The case where π_{on} is high enough to keep the network super-critical at all times will be addressed in future research.

8.4 Latency

Assume that the node sensing the incoming event is placed at the origin. This node starts sending an alarm message at time 0. All the active nodes, within transmission radius $2r_b$ or less, receive the message and they also begin to broadcast over their own transmission areas. Assuming no propagation delays, at $t = 0$ the message has spread to the cluster containing the origin (see Figure 8.3a). A sleeping node inside the set that was already covered by the broadcast message and that changes its state is called a *bridge*. After the bridge has started its broadcast, all the new active nodes who receive the message, either directly from the bridge or via a multi-hop path, are added to the original cluster (see Figure 8.3b). If the origin belongs to the infinite cluster of $\mathcal{B}(\lambda, r_b)$ the process continues forever, otherwise it stops after finitely many steps.

8.4.1 Linear spreading of alarm messages

In this section we state our main result. Assume that alarm messages are transmitted over a blinking Poisson Boolean model $\mathcal{B}(\lambda, r_b, Z)$. Remember that it means assuming a constant transmission range $2r_b$ for each sensor. First, we show that if an alarm occurs in a random place inside the infinite connected component of $\mathcal{B}(\lambda, r_b)$, the latency is asymptotically linear with respect to the distance to the sink. Secondly, the maximum distance from where the alarm message can be heard is also behaving asymptotically linearly.




Figure 8.3: Message spreading is shown in gray area. In this picture the disk radii are $r_{\text{rad}} = 2r$. White nodes are sleeping, black nodes active and a bridge is indicated by the box.

These results allow us to dimension the transmission part of a sensor network. First of all, they give us tools to tune up the alternation between sleep and active phase so that the speed of alarm detection meets the predefined requirements. In addition, if one wants to analyze the interference due to separate alarms appearing about the same time, maximum message distance can be used to approximate the area where the messages collide.

Assume that two nodes, located at X and Y , belong to the infinite cluster of $\mathcal{B}(\lambda, r_b)$. Denote by $T(X, Y)$ the time it takes to transmit an alarm message from X to Y . The first part of Theorem 8.1 shows that $T(X, Y)$ is asymptotically linear in distance $|X - Y|$. In the second scenario, we fix only the source X , and denote by G_t^X the nodes which have received the message at time t . G_t^X is called the *message cluster*. Then $\max |G_t^X| = \max\{|y - X| : y \in G_t\}$ measures the maximal transmission distance. This value is shown to be also asymptotically linear.

Theorem 8.1 *If two nodes, located at X and Y , belong to the infinite cluster of the Poisson Boolean model $\mathcal{B}(\lambda, r_b)$, resulting from a blinking Poisson Boolean model $\mathcal{B}(\lambda, r_b, Z)$ with $\lambda > \lambda^* > \pi_{\text{on}}\lambda$ and Z stationary on/off process with exponential or constant off-times, then there is a finite strictly positive constant η such that*

$$(1 - \epsilon)\eta \leq \frac{T(X, Y)}{|X - Y|} \leq (1 + \epsilon)\eta \quad (8.3)$$

for any $\epsilon > 0$ whenever $|X - Y|$ is large enough. Moreover, there are finite strictly positive constants $\underline{\mu}$ and $\bar{\mu}$ such that

$$\underline{\mu} \leq \frac{\max |G_t^X|}{t} \leq \bar{\mu}, \quad (8.4)$$

whenever t is large enough.

We will prove this theorem in the two following subsections. The linearity of the transmission time between a given pair of points is based on Liggett's sub-additive ergodic theorem. The linear growth of the message cluster is proved by coupling it with a continuum growth model.

It is important to notice that the constant η only depends on the parameters of the network — namely $\lambda, r, t_{\text{on}}$ and t_{off} — but not on the random disposition of the nodes. This value can thus be estimated by simulation, given a set of parameters, and used to predict the performance of the network before its deployment.

Proof of Equation (8.3)

We consider first passage percolation in the random graph $\mathcal{G}(\lambda, 2r_b)$. Assume that the random variables T_i satisfy

$$\begin{aligned}\mathbb{P}(T_i = 0) &= \pi_{\text{on}}, \\ \mathbb{P}(T_i > t) &= \pi_{\text{off}} \mathbb{P}(Z_s = 0 \forall s \in [0, t] \mid Z_0 = 0),\end{aligned}$$

i.e., T_i is the time until a node inside the message cluster turns active. If needed, the propagation delays could also be included in the random variables T_i . To each oriented edge (X_i, X_j) of the random graph, we attach a time coordinate (or “delay”) T_j . It is easy to couple the models in such a way that the message transmission time in a blinking Poisson Boolean model and the first passage time in the weighted graph are equal, i.e.,

$$T(X, Y) = \inf_{w(X, Y)} \left\{ \sum_{X_i \in w(X, Y)} T_i \right\}$$

where $w(X, Y)$ is an arbitrary path joining X and Y . However, notice that the other paths, except the fastest connection, may arise later in the static delay model. In the blinking model, when a node turns active inside the message cluster more than one new link can join the connectivity graph.

Without losing generality, we consider the first passage percolation in the direction of the x -axis. For any $(x, 0) \in \mathbb{R}^2$, we denote the index of the nearest node in the infinite cluster C_∞ by $i(x) = \operatorname{argmin}_i \{|(x, 0) - X_i| : X_i \in C_\infty\}$. Let $\tilde{X}(x) = X_{i(x)}$ and $\tilde{T}(x, y) = T(\tilde{X}(x), \tilde{X}(y))$ (see Figure 8.4). Next define the collection of indexed variables by $T_{m, n} = \tilde{T}(mx, nx)$, for some constant $x > 0$.

Using Liggett’s sub-additive ergodic theorem (Theorem 8.2), we can prove the following proposition.

Proposition 8.3

$$\lim_{n \rightarrow \infty} \frac{\tilde{T}(0, nx)}{n} = \lim_{n \rightarrow \infty} \frac{T_{0, n}}{n} = \eta(x) \quad a.s.$$

where $\eta(x) = \inf_{n \geq 1} \mathbb{E}(\tilde{T}(0, nx)) / n$.

Theorem 8.2 [Lig85, Liggett’s sub-additive ergodic theorem] *Let $\{T_{m, n}\}$ be a collection of random variables indexed by integers satisfying $0 \leq m < n$. Suppose $\{T_{m, n}\}$ has the following properties:*

- (i) $T_{0, n} \leq T_{0, m} + T_{m, n}$.
- (ii) For each n , $\mathbb{E}(|T_{0, n}|) < \infty$ and $\mathbb{E}(T_{0, n}) \geq cn$ for some constant $c > -\infty$.

(iii) The distribution of $\{T_{m,m+k} : k \geq 1\}$ does not depend on m .

(iv) For each $k \geq 1$, $\{T_{nk,(n+1)k} : n \geq 0\}$ is a stationary sequence.

Then:

(a) $\eta \doteq \lim_{n \rightarrow \infty} \mathbb{E}(T_{0,n})/n = \inf_{n \geq 1} E(T_{0,n})/n$.

(b) $T \doteq \lim_{n \rightarrow \infty} T_{0,n}/n$ exists a.s.

(c) $\mathbb{E}(T) = \eta$

Furthermore, if $k \geq 1$, $\{T_{nk,(n+1)k} : n \geq 0\}$ are ergodic, then

(d) $T = \eta$ a.s.

As $T_{0,m}$ is the first passage time from $\tilde{X}(0)$ to $\tilde{X}(mx)$, and $T_{m,n}$ the passage time from $\tilde{X}(mx)$ to $\tilde{X}(nx)$, it is clear that $T_{0,n}$ is at most $T_{0,m} + T_{m,n}$. Condition (i) is thus verified.




Figure 8.4: First passage percolation paths. The routes correspond to the fastest paths between points $\tilde{X}(0)$, $\tilde{X}(x)$ and $\tilde{X}(y)$.

As a first passage time cannot be negative, we have $\mathbb{E}(T_{m,n}) \geq 0, \forall m, n$. To compute an upper bound of $\mathbb{E}(T_{m,n})$, we consider the shortest path (in distance) from $\tilde{X}(mx)$ to $\tilde{X}(nx)$ in $\mathcal{B}(\lambda, r_b)$. A sub-optimal strategy is to follow this path, and wait at each step that the next hop becomes active. In the worst case, the message has to wait in average t_{off} seconds at each step. We have thus

$$\mathbb{E}(T_{m,n}) \leq t_{\text{off}} \mathbb{E}\left(L_{\tilde{X}(mx), \tilde{X}(nx)}\right),$$

where $L_{\tilde{X}(mx), \tilde{X}(nx)}$ denotes the length in hops of the shortest path between $\tilde{X}(mx)$ and $\tilde{X}(nx)$. Proposition 8.4 will ensure that the latter expected value is always finite, and thus that Condition (ii) is verified. To prove this proposition, we need the following two preliminary lemmas.




Figure 8.5: A good rectangle associated with a horizontal edge. It contains a left-right crossing and a top-bottom crossings in both sub-squares at its ends

Lemma 8.1 [Gri99, page 295] *In an independent bond percolation model of open edge density $\tilde{p} > 1/2$, we denote by A_n the event that there exists an open path in the rectangle $R_n = [0, 4n] \times [0, n]$ that joins its left and right borders (left-right crossing). There exist constants $\alpha < \infty$ and $\gamma > 0$ such that $\mathbb{P}_{\tilde{p}}(A_n) \geq 1 - \alpha n e^{-\gamma n}$.*

Lemma 8.2 *In $\mathcal{B}(\lambda, r_b)$, we denote by C_n the event that the origin is surrounded by an occupied circuit (i.e. a circuit entirely included in the occupied region) that is contained in the frame $F_n = ([-n, n] \times [-n, n]) \setminus ([-n/2, n/2] \times [-n/2, n/2])$. There exist constants $\alpha < \infty$ and $\gamma > 0$ such that $\mathbb{P}(C_n) \geq 1 - \alpha n e^{-\gamma n}$.*

Proof: We use a renormalization argument to map the Boolean model to a discrete model, in a similar way as in Section 6.3. We start by constructing a square lattice over the plane, with edge length d . For each edge (x, y) of the lattice, with $y = x + (d, 0)$ or $y = x + (0, d)$, we consider the rectangle $[x_1 - d/4, y_1 + d/4] \times [x_2 - d/4, y_2 + d/4]$, as depicted in Figure 8.5. We call a horizontal rectangle *good* if there exist in $\mathcal{B}(\lambda, r_b)$ an open cluster that crosses it from left to right, an open cluster that crosses $[x_1 - d/4, x_1 + d/4] \times [x_2 - d/4, x_2 + d/4]$ from bottom to top, and an open cluster that crosses $[y_1 - d/4, y_1 + d/4] \times [y_2 - d/4, y_2 + d/4]$ from bottom to top. We define *good* vertical rectangles in the same way, except that we exchange left-right with top-bottom. As $\mathcal{B}(\lambda, r_b)$ is supercritical, for any probability $p < 1$, one can choose d large enough so that a rectangle is good with probability at least p (see [MR96, Corollary 4.1]).

We then declare an edge *open* if it is surrounded by a good rectangle, and closed otherwise. We obtain thus a dependent bond percolation model. However, if two edges have no common vertex, their states are independent. Our model is thus a 1-dependent percolation model, which is known to percolate if p is large enough. More precisely, one can find a product measure $\mu_{\tilde{p}}$ on this model, where each edge is open with probability $\tilde{p} > 1/2$, that is stochastically dominated by our 1-dependent measure μ_p [Gri99].

We assume without loss of generality that the vertices of the lattice have the form (i, j) , for $i, j \in \mathbb{Z}$. In the bond percolation model with product measure, we look for left-right crossings in the rectangles $[-k, k] \times [-k, -k/2]$ and $[-k, k] \times [k/2, k]$ (their actual size is thus $2kd \times kd/2$) for some $k \in \mathbb{N}$. The probability




Figure 8.6: a) A frame that surrounds the origin. The four crossings form a circuit. b) The shortest path from X to Y . It is made of a chain of balls, that cannot overlap any other ball than their predecessor and successor.

that such a crossing exists is given by Lemma 8.1. The same is true for top-bottom crossings in $[-k, -k/2] \times [-k, k]$ and $[k/2, k] \times [-k, k]$. The probability that there is a crossing in each of these four rectangles simultaneously is bounded by

$$\mathbb{P}_{\mu_{\bar{p}}}(4 \text{ crossings}) \geq 1 - 4\alpha' n e^{-\gamma' k}. \quad (8.5)$$

As these crossings overlap at each corner, the origin is surrounded by an open circuit that is contained in

$$([-k, k] \times [-k, k]) \setminus ([-k/2, k/2] \times [-k/2, k/2])$$

(see Figure 8.6a). Since the product measure $\mu_{\bar{p}}$ is dominated by the 1-dependent measure, the crossings appear with higher or equal probability than in the independent case. Therefore,

$$\mathbb{P}_{\mu_p}(4 \text{ crossings}) \geq \mathbb{P}_{\mu_{\bar{p}}}(4 \text{ crossings}).$$

Moreover, our construction is such that the existence of an open circuit in the discrete model implies that $\mathcal{B}(\lambda, r_b) \cap ([-kd, kd] \times [-kd, kd]) \setminus ([-kd/2, kd/2] \times [-kd/2, kd/2])$ contains a component that surrounds the origin, which is the event we want. The final result is obtained by letting $\alpha = 4\alpha'$ and $\gamma = \gamma'/d$ in (8.5). \square

Proposition 8.4 *If X and Y are two points of a Boolean model, located at finite distance, that belong to the same cluster, and let L_{XY} be the number of hops in the shortest path between them. Then $\mathbb{E}(L_{XY})$ is finite.*

Proof: We assume without loss of generality that $X = (x, y)$ and $Y = (-x, -y)$, and consider a frame F_n , as defined in Lemma 8.2, with $n > 2\sqrt{x^2 + y^2}$. If this frame contains an occupied circuit, then the shortest path from X to Y is included in the square $[-n, n] \times [-n, n]$.

On the other hand, the shortest path is made of a chain of balls of radius r , and it is impossible that a ball overlaps more than two other balls of the

path. Otherwise, one could remove one ball and shorten the path. Therefore, a ℓ hops path must contain at least $\ell/2$ disjoint balls (see Figure 8.6b). The surface occupied by the path is thus at least $\pi r^2 \ell/2$. As the path is contained in a square of surface $4n^2$, the length of the path cannot exceed $\ell \leq \frac{8n^2}{\pi r^2}$. Therefore, $\mathbb{P}\left(L_{XY} \leq \frac{8n^2}{\pi r^2}\right) \geq \mathbb{P}(C_n)$. Combining this with Lemma 8.2 gives $\mathbb{P}\left(L_{XY} > \frac{8n^2}{\pi r^2}\right) \leq 1 - \mathbb{P}(C_n) \leq \alpha n e^{-\gamma n}$. Thus, we can find m large enough so that for $k \geq m$, $\mathbb{P}(L_{XY} > k) \leq \alpha \lfloor r\sqrt{k\pi/8} \rfloor e^{-\gamma \lfloor r\sqrt{k\pi/8} \rfloor}$. We can finally upper bound the expected value of the length of the shortest path:

$$\begin{aligned} \mathbb{E}(L_{XY}) &= \sum_{k=0}^{\infty} \mathbb{P}(L_{XY} > k) \\ &= \sum_{k=0}^{m-1} \mathbb{P}(L_{XY} > k) + \sum_{k=m}^{\infty} \mathbb{P}(L_{XY} > k) \\ &\leq m + \sum_{k=m}^{\infty} \alpha \lfloor r\sqrt{k\pi/8} \rfloor e^{-\gamma \lfloor r\sqrt{k\pi/8} \rfloor} < \infty. \end{aligned}$$

□

Conditions (iii) and (iv) are clearly verified, as $T_{m,n}$ is defined in a stationary way. The following lemma is to prove that the sequence $\{T_{n,n+1}\}$ is ergodic. In fact, we show that it is mixing (i.e., roughly speaking, asymptotically independent), which is a stronger property.

Lemma 8.3 *The sequence $\{T_{n,n+1}\}$, $n \geq 0$ is mixing.*

Proof: We compute $T_{0,1}$ by the following construction: we consider the square B_k of edge length kx centered at the origin. We denote by C_k the largest occupied connected component of $B_k \cap \mathcal{B}(\lambda, r_b)$, and by

$$\tilde{X}^{(k)}(y) = X_{\text{argmin}_i \{|(y,0) - X_i| : X_i \in C_k\}}$$

the closest point to $(y, 0)$ of C_k . We then define $T_{0,1}^{(k)}$ as the transmission time from $\tilde{X}^{(k)}(0)$ to $\tilde{X}^{(k)}(x)$. We observe that when k goes to infinity, as $\mathcal{B}(\lambda, r_b)$ is supercritical, the largest occupied component in B_k is $C_\infty \cap B_k$. Moreover, as a consequence of Proposition 8.4, the shortest path between $\tilde{X}^{(k)}(0)$ and $\tilde{X}^{(k)}(x)$ is finite. Thus we have that $\lim_{k \rightarrow \infty} T_{0,1}^{(k)} = T_{0,1}$ almost surely and therefore,

$$\lim_{k \rightarrow \infty} \mathbb{P}\left(T_{0,1}^{(k)} < t\right) = \mathbb{P}(T_{0,1} < t), \quad \forall t \in \mathbb{R}.$$

We consider now the translation $S_{(x,0)}$ over the vector $(x, 0)$ in \mathbb{R}^2 . Clearly we have $S_{(x,0)}^n(T_{0,1}) = T_{n,n+1}$. Similarly, we define $T_{n,n+1}^{(k)} \doteq S_{(x,0)}^n(T_{0,1}^{(k)})$. The same property is true for the translated variables

$$\lim_{k \rightarrow \infty} \mathbb{P}\left(T_{n,n+1}^{(k)} < t'\right) = \mathbb{P}(T_{n,n+1} < t').$$

The same is true for the combination of two events

$$\lim_{k \rightarrow \infty} \mathbb{P}\left((T_{n,n+1}^{(k)} < t) \cap (T_{m,m+1}^{(k)} < t')\right) = \mathbb{P}((T_{n,n+1} < t) \cap (T_{m,m+1} < t')).$$

Finally, we can show that the sequence $\{T_{n,n+1}\}$ is mixing by setting $m = n+2k$:

$$\begin{aligned} & \lim_{k \rightarrow \infty} \mathbb{P}((T_{n,n+1} < t) \cap (T_{n+2k,n+2k+1} < t')) \\ &= \lim_{k \rightarrow \infty} \mathbb{P}\left((T_{n,n+1}^{(k)} < t) \cap (T_{n+2k,n+2k+1}^{(k)} < t')\right) \\ &= \lim_{k \rightarrow \infty} \mathbb{P}\left(T_{n,n+1}^{(k)} < t\right) \mathbb{P}\left(T_{n+2k,n+2k+1}^{(k)} < t'\right) \\ &= \mathbb{P}(T_{n,n+1} < t) \mathbb{P}(T_{n+2k,n+2k+1} < t'), \quad \forall t, t' \in \mathbb{R}. \end{aligned}$$

The second equality follows from the fact that $T_{n,n+1}^{(k)}$ and $T_{n+2k,n+2k+1}^{(k)}$ are independent, as they depend on the realization of the blinking Boolean model on two disjoint squares. \square

Now we have seen that $T_{m,n}$ satisfies all the conditions of Theorem 8.2 and thus proved Proposition 8.3. Proposition 8.5 presented in the next section ensures that $\eta(x) > 0$.

Finally, we should show that $\Delta^{nx} = \tilde{X}(nx) - (nx, 0)$ does not play any role asymptotically and that the discrete limit can be replaced by a continuous one. Although both claims are quite evident, we give a short sketch of their proof.

Lemma 8.4 $\lim_{y \rightarrow \infty} \frac{T(\tilde{X}(0), \tilde{X}(y))}{|\tilde{X}(0) - \tilde{X}(y)|} = \eta$ a.s.

Proof:(Sketch) For independent bond percolation with $p > p^*(\mathbb{Z}^d)$, $d \geq 2$, there exists $\alpha > 0$ such that $\mathbb{P}(B_k \cap C_\infty = \emptyset) \leq e^{-\alpha k}$ (see e.g. [GM90]). Using Borel-Cantelli and the same mapping from Poisson Boolean model to discrete percolation as in Lemma 8.2, shows that $\Delta^{nx}/nx \rightarrow 0$ a.s.

Consider rational $t_n = nq = nk/m$, $k, m \in \mathbb{Z}^+$ and denote $\bar{T}_t = \frac{\tilde{T}(0, tx)}{t}$. Then $\bar{T}_{nk} \rightarrow \eta(x)$, since $\{\bar{T}_{nk}\}$ is a subsequence of $\{\bar{T}_n\}$ which converges to $\eta(x)$ by Proposition 8.3. Also by Proposition 8.3, $\frac{\tilde{T}(0, nqx)}{n} \rightarrow \eta(qx)$ a.s. Thus $\bar{T}_{nq} \rightarrow \eta(qx)/q$ a.s. Since $\{\bar{T}_{nk}\} \subset \{\bar{T}_{nq}\}$, they have the same limit, i.e., $\eta(qx)/q = \eta(x)$. Thus

$$\lim_{t_n \rightarrow \infty} \frac{\tilde{T}(0, t_n x)}{t_n x} = \eta(1) \doteq \eta \quad \text{a.s.}$$

\square

Proof of Equation (8.4)

The lower bound in Inequality (8.4) follows directly from (8.3). The next proposition gives the upper bound. Assuming a source at the origin, let G_t denote the nodes which have received the message at time t and S_t the area covered by these nodes if a disk of radius $2r_b$ is attached to each of them. Naturally, $G_t \subset S_t$.

Proposition 8.5 Consider the blinking Poisson Boolean model $\mathcal{B}(\lambda, r_b, Z)$ with $\lambda > \lambda^*$. If $\pi_{\text{on}} \lambda < \lambda^*$, then there is $\mu > 0$ such that almost surely $\frac{\max_t |S_t|}{t} \leq \mu$ for all sufficiently large t .

The claim is proved by showing that S_t can be bounded above by a continuum growth model (see Appendix C.1) which is driven by a Poisson point

process with intensity $\lambda\pi_{\text{off}}/t_{\text{off}}$ and an exponentially bounded disk size distribution.

Set $T_0 = 0$. At time epochs $T_i > 0$, $i = 1, 2, \dots$ one of the off-nodes inside $S_{T_{i-1}}$, denoted by X_i , changes its state and S_t possibly grows. $C(X_i) \doteq S_{T_i} \setminus S_{T_{i-1}}$ determines the new area which receives the message at time T_i (shown in light gray in Figure 8.7). We will first show that $C(X_i)$ can be bounded by disks $B(X_i, D_i)$ with i.i.d. D_i . This follows from the cluster size distribution in a sub-critical regime of a Poisson Boolean model.

Let $A \rightsquigarrow B$ denote the event that sets A and B intersect the same cluster.

Lemma 8.5 [MR96, Theorem 2.4] *Consider Poisson Boolean model $\mathcal{B}(\lambda, R)$ where R satisfies $0 \leq R \leq r$ for some $r < \infty$. Assuming $\mathbb{E}(\text{diam}(W(\{0\}))) < \infty$ then there exist positive constants C_1 and C_2 , depending on λ and the dimension d , such that*

$$\mathbb{P}(S \rightsquigarrow B(0, m)^c) < C_1 \exp(-C_2 m)$$

for an arbitrary bounded set S .

Lemma 8.6 *If $\pi_{\text{on}}\lambda < \lambda^*$, then $C(X_i) \subseteq B(X_i, D_i)$ a.s., where D_i are i.i.d. with*

$$\mathbb{P}(D_i > a) < C_1 \exp(-C_2 a) \quad (8.6)$$

for some positive constants C_1 and C_2 .

Proof:

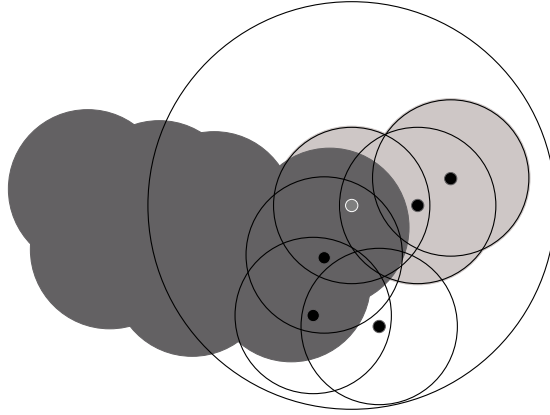


Figure 8.7: Conditioned (light gray), unconditioned increments (black circles) and the coupling disk (the largest circle). The radii of the disks are $r = 2r_b$.

Let Φ be an arbitrary point process. For a Boolean model with disks centered according to Φ and radii r_b , let $W_{r_b}(A, \Phi)$ denote the union of the occupied clusters intersecting set A and $\tilde{W}_{r_b}(A, \Phi) \doteq \{\mathbf{x} : |\mathbf{x} - (A \cup W_{r_b}(A, \Phi))| \leq r_b\}$ the set where the message is heard.

When S grows at $t = T_i$, all the possible new points, except the bridge X_i , are outside $S_{T_{i-1}}$. In $S_{T_{i-1}}^c$ the active nodes are distributed according to a stationary Poisson point process with intensity $\lambda\pi_{\text{on}}$ which we denote by Φ_i . Thus

$$C(X_i) = \tilde{W}_{r_b}(B(X_i, r_b), \Phi_i(S_{T_{i-1}}^c)) \setminus S_{T_{i-1}}.$$

This set can be drawn without any information about nodes in $S_{T_i}^c$. Thus the increments $C(X_i)$ can be determined using a sequence of i.i.d. Poisson processes Φ_i . Moreover,

$$C(X_i) \subseteq \tilde{W}_{r_b}(B(X_i, r_b), \Phi_i)$$

since neglecting the conditions with respect $S_{T_{i-1}}$ gives naturally a larger set (see Figure 8.7).

Finally, let $D_i = \text{diam}(W_{r_b}(B(X_i, r_b))) + 4r_b$ so that

$$C(X_i) \subseteq B(X_i, D_i).$$

Applying Lemma 8.5 yields that D_i satisfies (8.6). \square

Next we show that the process indicating when and where a bridge appears can be stochastically bounded by a Poisson process in \mathbb{R}^3 . In other words, the Poisson point process includes all the bridges (plus infinitely many more), it preserves the opening order, and each bridge opens earlier than it was originally scheduled.

Lemma 8.7 *If the sleeping periods are either constant or exponentially distributed, then the bridge process can be coupled with a stationary Poisson point process in \mathbb{R}^3 of intensity $\lambda\pi_{\text{off}}/t_{\text{off}}$.*

Proof: Let $\{X_i^0\}$ denote the bridges in S_0 with opening times $\{T_i^0\}$. At time $t = 0$, the X_i are distributed according to Poisson point process with intensity $\lambda\pi_{\text{off}}$. If $U_i \sim \text{Uniform}(0, t_{\text{off}})$, then points $(X_i^0, U_i) \in S_0 \times [0, t_{\text{off}}]$ are distributed according to a Poisson point process with intensity $\lambda\pi_{\text{off}}/t_{\text{off}}$. Thus if the T_i^0 are uniformly distributed then we know that they all appear according to a Poisson process on $S_0 \times [0, t_{\text{off}}]$. On the other hand, if the T_i^0 are exponentially distributed, they can be coupled with the U_i by

$$T_i^0 = -t_{\text{off}} \log(1 - U_i/t_{\text{off}}) \geq U_i$$

and thus with the Poisson point process with intensity $\lambda\pi_{\text{off}}/t_{\text{off}}$.

The same reasoning holds for any new set added to the cluster. \square

Proof: (For Proposition 8.5) By Lemma 8.6, we bound each new set added to the message cluster by disks with i.i.d. radii with cumulative distribution satisfying (8.6). By Lemma 8.7, each of the bridges originating an increase of the message set are included in a Poisson point process in \mathbb{R}^3 with intensity $\lambda\pi_{\text{off}}/t_{\text{off}}$. Thus applying Theorem C.1 (shown in Appendix C.1) completes the proof. \square

8.4.2 Duration of a transmission phase

When a node receives or generates a message, it keeps transmitting until its neighbors have received the message. From an energy consumption point of view, it is important to know how long this transmitting phase will last. Assuming that the node is connected to the infinite cluster, then — in principle — it would be enough that the “optimally” located neighboring nodes receive the message. However, in order to maximize the speed at which messages travel, we define the broadcast duration as the time until all the neighboring nodes have received the message with probability p_b . Furthermore, we assume here that the emitter ignores its neighborhood and receives no feedback from the

receivers. Although it would be easy to achieve better conditions by designing an appropriate protocol, we only consider this simple mechanism.

Given the transmission radius $r = 2r_b$ of the sensors, there is a Poisson distributed number of sensors inside its transmission range. In a worst case scenario, we assume that all the nodes are in sleeping mode at the time when the broadcast starts. If the sleeping periods are of constant length, then the natural and 100% safe broadcasting duration is $t_{\text{off}} + \Delta$, where Δ is the propagation delay. Otherwise, we assume that the off periods of the neighbors, T_k , $k = 1, \dots, N$, are exponentially distributed and $N \sim \text{Poisson}(\lambda\pi r^2)$. In order to be p_b sure that every node receives the message, the broadcasting time t_b has to satisfy

$$\begin{aligned} \mathbb{P}(\max\{T_k\}_{k=1}^N < t_b) &= \sum_{k=0}^{\infty} \frac{(\lambda\pi r^2)^k}{k!} e^{-\lambda\pi r^2} (1 - e^{-t_b/t_{\text{off}}})^k \\ &= \exp\left(\lambda\pi r^2 e^{-t_b/t_{\text{off}}}\right) < p_b. \end{aligned}$$

Thus, if the broadcasting time t_b satisfies

$$t_b \geq t_{\text{off}} \log\left(\frac{\lambda\pi r^2}{-\log p_b}\right) + \Delta,$$

then it is p_b sure that every node in its neighborhood gets the message.

8.4.3 Simulation studies

We have performed a series of simulations to validate the results concerning latency presented in Section 8.4.1. These numerical studies clearly agree with the analytical results on the linear spreading rate. Although Theorem 8.1 is stated only for constant transmission radii, the simulations suggest that it is also valid for random transmission radii.

All the simulations were run with intensity $\lambda = 3$. The sleeping periods were drawn from an exponential distribution with mean t_{off} . The size of the area which is reached by the message at time t is measured by $\max |G_t| = \text{argmax}\{|x| : x \in G_t\}$ and the spreading rate is estimated by

$$\hat{\mu} = \frac{\mathbb{E}(\max |G_t|)}{t}.$$

Simulations of the growth of the message cluster with varying period lengths t_{on} and t_{off} are shown in Figures 8.8 and 8.9. In principle, only the ratios $t_{\text{on}}/(t_{\text{off}} + t_{\text{on}})$ and $t_{\text{off}}/(t_{\text{off}} + t_{\text{on}})$ matter. By a simple time scaling argument t_{off} (or t_{on}) can be taken as the time unit. This mapping, of course, changes the actual value of μ . For fixed t_{off} , decreasing the lengths of the active periods does not really worsen the performance when t_{on} is small enough. This is natural, since for very small values of t_{on} , almost all the sensors are sleeping when the message arrives in their range so that the time to wait for a wakeup determines the transmission speed.

When t_{on} is fixed, increasing t_{off} naturally decreases the message spreading velocity. However, t_{off} is the determinant factor for energy saving, especially if switching on the radio is costly. Therefore, the lower part of Figure 8.9 presents




Figure 8.8: Maximum connection distance from the origin for different active periods. Ten independent simulations with parameters $\lambda = 3$, $r_b = 1$, $t_{\text{off}} = 1$, and $t_{\text{on}} = 0.1, 0.2, 0.3$




Figure 8.9: Estimation of growth rates. Parameters $\lambda = 3$, $r_b = 1$. 100 independent simulations per estimate.

the real trade-off in this mechanism. We observe that the curve decreases very fast at the beginning, meaning that allowing the nodes to sleep for some time costs a lot of latency. But afterward, increasing the sleeping period has less impact, as the curve becomes flatter.




Figure 8.10: Maximum connection distance from the origin for the models with constant or Uniform(0.5, 1.5) distributed radii. Ten independent simulations with parameters $\lambda = 3$, $t_{\text{off}} = 1$, $t_{\text{on}} = 0.1$

The effect of allowing random radii can be seen in Figures 8.10 and 8.11. The main difference between constant and random radii models is that in the latter longer jumps are possible. The other difference is that the mean covered area is larger in the random case, since $\mathbb{E}(R^2) \geq (\mathbb{E}(R))^2$. This naturally results in better connectivity. However we observe that the qualitative behavior is similar. Actually, most of the proofs presented in this chapter can be easily extended to the random radii model, and we conjecture that all our results hold for this model.

Figure 8.12 shows the impact of propagation delays on the message spreading. The asymptotic behavior is qualitatively conserved, as predicted. The quantitative impact is also tiny, if we consider the realistic case where delays are short compared to t_{on} and t_{off} .




Figure 8.11: Models with $r_b = 1$ and $R \sim \text{Uniform}(0.5, 1.3)$. Estimated maximum growth rates. Parameters $\lambda = 3$, $t_{off} = 1$, $t_{on} = 0.1$




Figure 8.12: Model with propagation delays. Maximum connection distance from the origin. Ten independent simulations with parameters $\lambda = 3$, $r_b = 1$, $t_{off} = 1$, $t_{on} = 0.2$ and $\Delta = 0, 0.05, 0.5, 5$.




Figure 8.13: Model with propagation delays. Estimated growth rates. Parameters $\lambda = 3$, $r_b = 1$, $t_{\text{off}} = 1$, delays 0.05, 0.5 and 5. Estimates based on 100 independent simulations.

8.5 Conclusion

The blinking Poisson Boolean model is suited for sensor networks where nodes switch between a sleeping and an active phase. Even though their switching on/off schedules are not coordinated at all, their positions are random, and the average durations t_{on} and t_{off} are such that the number of active nodes at any particular time is so low that the network is always disconnected, we have proved that any message (alarm) generated by a sensor will reach the sink in a time proportional to the distance between the sensor and the sink. The value of the rate of this linear growth does not depend on the random locations of the nodes, but only on the parameters λ (node density), r (connectivity range), t_{on} and t_{off} (average active and sleeping durations).

Chapter 9

Conclusion

Throughout this dissertation, we have been applying percolation theory to investigate the structure and properties of wireless multi-hop networks. When the number of nodes becomes large, the network must be considered as a macroscopic object, that acquires properties that are not observable at a smaller scale.

We saw in Chapter 2 that the dimensionality of the network plays a crucial role in the connectivity of the network. In one dimension, the use of base stations is necessary for keeping the network connected. On the contrary, in two dimensions, the percolation phenomenon implies the existence, under some conditions, of a unique unbounded cluster of connected nodes. This cluster forms the actual network, whereas other (bounded) clusters provide no useful contribution, and are said to be disconnected.

We proved in Chapter 3 that such a cluster may exist even in very pessimistic conditions, namely when all nodes emit at the same time. Moreover, the fraction of nodes belonging to the unbounded cluster can be made arbitrarily close to one by adapting the parameters, independently of the network size. In other words, the fraction of connected nodes remains constant when the network size increases. This *partial connectivity* is thus a perfectly scalable property, as opposed to *full connectivity*, which is very costly in terms of network resources. We showed in Chapter 5 that partial connectivity also scales under a very general information theoretic definition of connectivity.

The quality of connectivity has however a cost. We have seen in Chapter 4 that there is a clear trade-off between connectivity and throughput between sources and destinations. Indeed, long range links cannot achieve a high data rate because of the background noise and interferences.

The result presented in Chapter 6 is a good example of this trade-off. We showed that a throughput per node of order $1/\sqrt{n}$ is achievable in the random networks scenario introduced by Gupta and Kumar [GK00]. This throughput can be achieved by adjusting the connectivity at the right level. In the proposed scheme, we require that only a fraction of the nodes participate to the relaying of the packets, allowing for some disconnected clusters. This allows to keep the node's connectivity range bounded, thus preventing excessive interference. However, we showed that the giant cluster has sufficient regularity properties to allow for deterministic routing. This choice optimizes the overall throughput of the network, achieving a better bound than in previous works.

We studied further properties of the unbounded cluster in Chapters 7 and 8.

We looked first at the distance between an arbitrary location and the giant component, and proved that this distance is bounded from above by an exponential variable. This result can be applied to sensor networks, for bounding the time before a mobile object hits a network of interconnected sensors. In Chapter 8, we looked at the unbounded cluster in a multi-hop network using first passage percolation. This approach differs from the standard percolation approach because it does not only look at connectivity, but at the shortest paths length in a weighted graph. The result we obtain in this chapter confirms that the giant component exhibits some regularity properties (from a macroscopic viewpoint), since these lengths are proportional to the actual distance traveled. We applied this result to the case where nodes implement a uncoordinated energy saving mechanism, and proved that messages travel with an asymptotically constant speed in such networks.

9.1 Future research

There are still many possible promising applications of percolation theory to wireless multi-hop networks. In this section, we give three directions, that could be natural extensions of this work.

Synchronization and stability

In Chapters 6 and 8, we intensively used the properties of the giant cluster to show our results, once in the context of throughput and once in the context of propagation delay. The structure of the giant cluster can however still be explored in other contexts. For instance, it would be very interesting to study the behavior and stability of a synchronization mechanism: assume that nodes exchange locally their current time, in order to synchronize their reference clocks. What would be the critical density for having coherent time throughout the whole network? Intuitively, we guess that below a certain threshold, even if the network is still super-critical, the interaction between nodes may be too weak to maintain a stable state. Several islands, each having its own time, may appear. There should thus be a new threshold, below which the protocol is unstable.

Furthermore, in the stable case, there are still many questions that can be addressed using percolation. For example, what is the maximum time difference between two nodes located at a certain distance from each other? It would also be interesting to check whether this difference is proportional to the distance. Clock synchronization is only an example, and the same questions apply to any protocol that intends to establish an kind of consensus in the network.

Mobility

In Chapter 7, we looked at a scenario with some (very limited) mobility. An open question is what happens if all sensors are moving. What is the impact on the detection time? It is already well known that under mild mobility models (including constant speed straight line motion addressed in this work), the node distribution is still Poisson at any instant, and the giant cluster exists a.s. at any instant as well. The same problem could thus be addressed in the fully mobile case, and we expect similar results.

Moreover, the choice of mobility model may affect the results. It would thus be interesting to study other mobility models, like Brownian motion, which would be a good counterpart to straight line motion.

More generally, connectivity in mobile multi-hop networks is not well understood yet, and few papers address this issue. In particular, the computation of the distribution of the connected and disconnected epoch of the nodes is still an open problem.

Directed connectivity graphs

In this work, we assumed isotropic wave propagation. We argued that this assumption is very conservative, in terms of connectivity, since non-circular shapes percolate better than disks. Therefore, it would be clever to use of directional antennas in multi-hop network, to *intentionally* make wave propagation non-isotropic, even with random directions.

It would thus be interesting to study connectivity with non-isotropic attenuation functions. The first option for doing that is to use the Boolean model with non-circular grains. However, this approach does not model the antenna's directive gain appropriately. We think that there is a need for a new model here.

Furthermore, the use of directional antennas certainly leads to a large proportion of unidirectional links in the connectivity graph. Such links are usually considered useless, because of protocol issues (for example acknowledgments). The study of the size of the loops in directed graphs would thus help deciding on the usability of those links. Very simple protocols without acknowledgments could also make use of unidirectional links, for broadcasting or synchronization purpose for example. Percolation in directed graph is thus a valuable research topic, and has, to our knowledge, received very little attention so far.

Appendix A

Appendices for Part I

A.1 Derivation of the lower bound to the probability of connection in one-dimensional hybrid networks

We start from the expression

$$P_{cb}^+ \geq 2(1 - e^{-\lambda r}) e^{-\lambda(L-4r)e^{-\lambda r}/2} - 2\lambda r e^{-\lambda r} - 1(1 - e^{-\lambda r})^2 e^{-\lambda(L-2r)e^{-\lambda r}}.$$

The right-hand side can be put in the following form:

$$P_{cb}^+ \geq 1 - \left(1 - (1 - e^{-\lambda r})e^{-\lambda(\frac{L}{2}-r)e^{-\lambda r}}\right)^2 - 2(1 - e^{-\lambda r})e^{-(\frac{\lambda L}{2}-2r)e^{-\lambda r}} \left[e^{-\lambda r e^{-\lambda r}} - 1\right] - 2\lambda r e^{-\lambda r}. \quad (\text{A.1})$$

Observing that for any $x > 0$ and $\alpha > 0$, we have

$$1 - (1 - \alpha e^{-x})^2 \geq 1 - (1 - \alpha(1 - x))^2,$$

we can transform the first line of (A.1) as follows:

$$\begin{aligned} & 1 - \left(1 - (1 - e^{-\lambda r})e^{-\lambda(\frac{L}{2}-r)e^{-\lambda r}}\right)^2 \\ & \geq 1 - \left(1 - (1 - e^{-\lambda r})(1 - \lambda(\frac{L}{2} - r)e^{-\lambda r})\right)^2 \\ & = 1 - e^{-2\lambda r} \left(1 + \lambda(\frac{L}{2} - r)(1 - e^{-\lambda r})\right)^2 \\ & = 1 - e^{-2\lambda r} \left(1 + \frac{\lambda L}{2} - \left(\lambda r + \frac{\lambda(L-2r)}{2}(1 - e^{-\lambda r})\right)\right)^2 \end{aligned}$$

By developing the square, and rearranging the terms, we obtain a bound to the first line of (A.1)

$$\begin{aligned}
& 1 - \left(1 - (1 - e^{-\lambda r})e^{-\lambda(\frac{L}{2}-r)e^{-\lambda r}}\right)^2 \\
& \geq 1 - \left(1 + \frac{\lambda L}{2}\right)^2 e^{-2\lambda r} + 2\lambda r e^{-\lambda r} + \lambda^2 r(L - 3r)e^{-2\lambda r} + 2\lambda^2 r^2 e^{-2\lambda r} \\
& \quad + \lambda(L - 2r)(1 - e^{-\lambda r})e^{-2\lambda r} \left(1 + \frac{\lambda(L - 2r)(1 + e^{-\lambda r})}{2}\right). \quad (\text{A.2})
\end{aligned}$$

The second line of (A.1) can be expanded using Taylor series:

$$\begin{aligned}
& 1 - \left(1 - (1 - e^{-\lambda r})e^{-\lambda(\frac{L}{2}-r)e^{-\lambda r}}\right)^2 \\
& \geq 2(1 - e^{-\lambda r}) \left(1 - \lambda\left(\frac{L}{2} - 2r\right)e^{-\lambda r}\right) \left(1 - (1 - \lambda r e^{-\lambda r} + \frac{(\lambda r)^2}{2}e^{-2\lambda r})\right) \\
& \quad - 2\lambda r e^{-\lambda r}
\end{aligned}$$

Developing and rearranging the terms, we obtain

$$\begin{aligned}
& 1 - \left(1 - (1 - e^{-\lambda r})e^{-\lambda(\frac{L}{2}-r)e^{-\lambda r}}\right)^2 \\
& \geq -2\lambda r e^{-2\lambda r} - \lambda^2 r(L - 3r)e^{-2\lambda r}(1 - e^{-\lambda r}) \\
& \quad + \lambda^3 r^2 \left(\frac{L}{2} - 2r\right)e^{-3\lambda r}(1 - e^{-\lambda r}) \quad (\text{A.3})
\end{aligned}$$

Combining (A.2) and (A.3) we obtain

$$\begin{aligned}
P_{cb}^+ & \geq 1 - \left(1 + \frac{\lambda L}{2}\right)^2 e^{-2\lambda r} + \lambda^2 r(L - 3r)e^{-2\lambda r} + 2\lambda^2 r^2 e^{-2\lambda r} \\
& \quad + \lambda(L - 2r)(1 - e^{-\lambda r})e^{-2\lambda r} \left(1 + \frac{\lambda(L - 2r)(1 + e^{-\lambda r})}{2}\right) \\
& \quad + \lambda^3 r^2 \left(\frac{L}{2} - 2r\right)e^{-3\lambda r}(1 - e^{-\lambda r}) \\
& \geq 1 - \left(1 + \frac{\lambda L}{2}\right)^2 e^{-2\lambda r}
\end{aligned}$$

A.2 Percolation in the STIRG model with bounded support attenuation function

The existence of a percolation phenomenon in the STIRG model was proved for the first time in [DBT03]. The proof presented in this paper, however, requires the assumption that the attenuation function $L(\cdot)$ has a bounded support. Furthermore, the existence of a giant component in $\mathcal{H}(\gamma, \lambda)$ for non-zero γ was proved only for node densities higher than a certain value λ' . We know from Section 3.2.3 that there exist non-zero values of γ such that the model percolates whenever $\lambda > \lambda^*$, where λ^* is the percolation threshold for the Boolean model $\mathcal{H}(0, \lambda)$.




Figure A.1: Mapping

The assumptions in [DBT03] make indeed the model much simpler, since the dependence between edges has a limited range; only nodes in a bounded neighborhood of the receiver can interfere with a given transmission. The proof presented initially in [DBT03] makes extensive use of this finite range dependency. In this appendix, we present an alternative proof for the same result, that is particularly simple and elegant. This proof was proposed by Artyom Sapozhnikov, and requires the following assumption on $L(\cdot)$:

1. $L(\mathbf{x}) \leq 1$ for all $\mathbf{x} \in \mathbb{R}^2$.
2. $L(\mathbf{x}) = 0$ for all $\mathbf{x} \in A^c$, with A a bounded subset of \mathbb{R}^2 .
3. There exists $d > 0$ such that $\|\mathbf{x}\| \leq d$ implies $L(\mathbf{x}) \geq m$, for some $m > \frac{\beta N_0}{P}$

We prove then the following theorem.

Theorem A.1 *Suppose that the attenuation function $L(\cdot)$ satisfies the assumptions [1–3]. There exist a density $\lambda' < \infty$ and a function $\gamma^*(\lambda) > 0$ such that if $\lambda > 0$ and $\gamma < \gamma^*(\lambda)$, then $\mathcal{H}(\gamma, \lambda)$ contains an unbounded connected component.*

Proof: (sketch) We start by mapping the STIRG model onto a site percolation model as follows. We set $D \doteq \text{diam}(A)$, where A is the bounded set of Assumption [2], and

$$s \doteq 1 + \left\lceil \frac{2\sqrt{5}D}{d} \right\rceil.$$

Thus we have

$$\|\mathbf{x}\| > D \Rightarrow L(\mathbf{x}) = 0.$$

Next, we divide the plane into squares of size $ds/\sqrt{5}$. Then, each square can be re-divided into s^2 little squares of size $d/\sqrt{5}$, as depicted in Figure A.1.

We declare each big square *open* if it fulfills the two following conditions:

- C1 Each little square of the central column and of the central row contains at least one point (see Figure A.1).
- C2 The square contains no more than M points.

We will specify the value of M later. The attentive reader will notice that the two conditions are incompatible if $M < 2s - 1$.

The probability that Condition C1 is fulfilled (denoted here by $\mathbb{P}(\text{C1})$) can be made arbitrarily close to one by increasing the node density λ . Then, once we chose λ , we can make $\mathbb{P}(\text{C2})$ arbitrarily close to one by increasing the value of M . Since $\mathbb{P}(\text{C1} \cap \text{C2}) \geq \mathbb{P}(\text{C1}) + \mathbb{P}(\text{C2}) - 1$, we can choose λ and M such that $\mathbb{P}(\text{C1} \cap \text{C2}) > p_{\text{site}}$, where p_{site} is the critical threshold for site percolation.

By construction, the squares are declared open independently from each other, since Conditions C1 and C2 regard the realization of the Poisson point process inside each square only. Thus we can choose λ and M such that the squares form a super-critical site percolation model.

To conclude the proof, we have to show that if there is an unbounded cluster of open squares in the site percolation model, then $\mathcal{H}(\gamma, \lambda)$ contains an unbounded connected component for some strictly positive value of γ . Let us choose

$$\gamma = \frac{1}{2\beta(M-1)} \left(m - \frac{\beta N_0}{P} \right). \quad (\text{A.4})$$

We have clearly $\gamma > 0$ because of Assumption [3]. We consider now two open adjacent squares, and the sequence of little squares joining the two centers, as depicted in Figure A.2. As there are at most $2M$ nodes in the two big squares, the interference inside the little squares is at most $(2M-2)P$, since the emitter and receiver are counted in the $2M$ nodes, and since $L(\mathbf{x}) \leq 1$. On the other hand, since each little square contains at least one node, and since the distance between two nodes X_i and X_j in consecutive little squares is less than d , the SINR reads

$$\begin{aligned} \frac{PL(X_i - X_j)}{N_0 + \gamma \sum_{k \neq i, j} PL(X_k - X_j)} &\geq \frac{Pm}{N_0 + \gamma(2M-2)P} \\ &= \frac{Pm}{N_0 + \frac{P}{\beta} \left(m - \frac{\beta N_0}{P} \right)} \\ &= \beta, \end{aligned}$$

This means that there is a link between the nodes of each pair of adjacent little squares of the sequence, joining the two centers. Therefore, the centers of all adjacent open big squares are connected in $\mathcal{H}(\gamma, \lambda)$, and an unbounded cluster of open squares implies an unbounded connected component in $\mathcal{H}(\gamma, \lambda)$. \square

A little modification of the above proof allows to show that percolation occurs also in the TDMA scheme presented in Section 3.4.

Corollary A.2 *Suppose that the attenuation function $L(\cdot)$ satisfies the Assumptions [1-3]. There exist a density $\lambda' < \infty$ and a function $t^*(\lambda) > 0$ such that if $\lambda > 0$ and $t > t^*(\gamma, \lambda)$, then $\mathcal{H}^t(\gamma, \lambda)$ contains an unbounded connected component.*

Proof: The result is obtained by replacing Condition C2 with the following:

C2' During each time slot, the square contains no more than M active points.

An important difference is that here, since γ is fixed, we cannot choose M . Instead, we compute M similarly as in (A.4)

$$M = 1 + \left\lfloor \frac{1}{2\beta\gamma} \left(m - \frac{\beta N_0}{P} \right) \right\rfloor.$$




Figure A.2: Two adjacent open squares

The new condition C2' is in fact t times Condition C2 applied to a point process of intensity λ/t . Therefore, the probability that C2' is verified can be computed as follows:

$$\begin{aligned}
 \mathbb{P}(C2') &= \mathbb{P}(C2)^t \\
 &= \mathbb{P}(N \leq M)^t \\
 &\geq \mathbb{P}(N \leq 1)^t \\
 &= \left[\exp\left(\frac{-\lambda d^2 s^2}{5t}\right) + \frac{\lambda d^2 s^2}{5t} \exp\left(\frac{-\lambda d^2 s^2}{5t}\right) \right]^t \\
 &= \left(1 + \frac{\lambda d^2 s^2}{5t}\right)^t \exp\left(\frac{-\lambda d^2 s^2}{5}\right).
 \end{aligned}$$

We observe that this expression tends to one when t goes to infinity. Therefore, $\mathbb{P}(C2')$ can be made arbitrarily close to one by choosing t sufficiently large. The remainder of the proof is identical to the proof of Theorem A.1. \square

A.3 Invariance with respect to scaling

We consider a function that fulfill the following property:

$$\frac{l(ax)}{l(ay)} = \frac{l(x)}{l(y)} \quad \forall x, y, a > 0$$

It follows from the above equality that

$$\frac{l(ax)}{l(x)} = \frac{l(ay)}{l(y)} = c(a),$$

because x and y can be arbitrary chosen. Thus we have

$$l(ax) = c(a)l(x), \quad \forall a, x \neq 0.$$

By letting $x = 1$ in this last equality, we find that

$$c(a) = \frac{l(a)}{l(1)},$$

and thus

$$l(ax) = \frac{l(a)l(x)}{l(1)} \quad \forall a, x \neq 0$$

Let us consider now a transform of l :

$$L(x) \doteq \log(l(e^x))$$

This new function is affine:

$$\begin{aligned} L(x+y) &= \log[l(e^{x+y})] \\ &= \log[l(e^x e^y)] \\ &= \log\left[\frac{l(e^x)l(e^y)}{l(1)}\right] \\ &= \log[l(e^x)] + \log[l(e^y)] - \log[l(1)] \\ &= L(x) + L(y) - L(0), \end{aligned}$$

and thus

$$L(x) = \alpha x + L(0)$$

for some constant α . We finally compute then the original function $l(x)$:

$$\begin{aligned} l(x) &= \exp\{L[\log(x)]\} \\ &= \exp[\alpha \log(x) + L(0)] = l(1)x^\alpha. \end{aligned}$$

Appendix B

Appendices for Part II

B.1 Shot-noise

We start this section by introducing some notation. For a set $A \subset \mathbb{R}^d$ and a point $\mathbf{x} \in \mathbb{R}^d$, we define the set $A + \mathbf{x}$ as follows:

$$\mathbf{y} \in A + \mathbf{x} \quad \text{iff} \quad \mathbf{y} - \mathbf{x} \in A.$$

Then we define the shift operator $S_{\mathbf{x}}$ as follows: for a random measure N and a Borel set A ,

$$S_{\mathbf{x}}N(A) = N(A + \mathbf{x}).$$

We define now formally the shot-noise process built on a two-dimensional Poisson point process, and prove an ergodic result in Theorem B.1.

Let Φ be a Poisson point process of intensity λ in \mathbb{R}^d , and denote by $\{X_k\}$ the points of this process. We define the *shot-noise* as

$$I = \int_{\mathbb{R}^d \setminus \{0\}} f(\mathbf{x})\Phi(d\mathbf{x}) = \sum_k f(X_k)$$

for some positive function f such that

$$\int_{\mathbb{R}^d} f(\mathbf{x})d\mathbf{x} < \infty.$$

Given a constant $M > 0$, we define by Y the random counting measure that counts the number of points of Φ such that the value of the shot-noise at this point is less than M :

$$Y(A) = \int_A 1_E(S_{\mathbf{x}}\Phi)\Phi(d\mathbf{x})$$

where E is the event

$$E = \{I < M\}.$$

The mean density of the new point process Y is clearly equal to the intensity of Φ , namely λ , times the Palm probability $\mathbb{P}_0(E)$ of the event E . Because of the properties of the Poisson point process, the Palm probability of this event is equal to the regular probability $\mathbb{P}(E)$. The mean density of Y is therefore equal to $\lambda\mathbb{P}(E)$.

Before stating the main theorem of this Appendix, we need the following definition:

Definition B.1 A sequence $\{A_n\}$ of bounded Borel sets in \mathbb{R}^d is a convex averaging sequence if

1. each A_n is convex
2. $A_n \subset A_{n+1}$ for $n = 1, 2, \dots$
3. $r(A_n) \rightarrow \infty$, where $r(A) = \sup\{r : A \text{ contains a ball of radius } r\}$.

Theorem B.1 For a convex averaging sequence $\{A_n\}$ of Borel sets in \mathbb{R}^d , as $n \rightarrow \infty$ we have

$$\frac{Y(A_n)}{\ell(A_n)} \rightarrow \lambda\mathbb{P}(E),$$

where $\ell(A_n)$ denotes the area of A_n .

This theorem follows directly from Theorem B.2 stated below, and the observation that the point process Y is stationary and ergodic, since the underlying point process Φ is stationary and ergodic. Then we use Proposition 10.3.III in [DVJ88] to prove that Y is also metrically transitive, since it is ergodic.

Theorem B.2 (Corollary 10.2.V in [DVJ88]) Let ξ be a stationary and metrically transitive random measure on \mathbb{R}^d with finite mean density m , and $\{A_n\}$ a convex averaging sequence of Borel sets on \mathbb{R}^d . Then as $n \rightarrow \infty$,

$$\xi(A_n)/\ell(A_n) \rightarrow m \text{ (a.s. and in } L_1 \text{ norm)}.$$

Appendix C

Appendices for Part III

C.1 Continuum growth model

Continuum growth model is a model for a spreading infection. Assuming that an infection has spread to set S_t at time t , the time until the next outburst occurs somewhere in S_t is exponentially distributed with parameter $\lambda v_d(S_t)$ and the location of the outburst is uniformly distributed over S_t . Outburst are assumed to be balls with i.i.d. random radii. This model was first studied by Deijfen [Dei03] (see also Deijfen, Häggström and Bagley [DHB04]).

Formally, Continuum growth model is defined as follows. We start with a stationary Poisson point process $N = (X_k, T_k)_k \in \mathbb{R}^{d+1}$ with $X_k \in \mathbb{R}^d$. The Poisson process has intensity λ and the points laying inside a set G are denoted by $N(G) = G \cap N$. For each X_k we attach a ball $B(X_k, D_k) = B_k \subset \mathbb{R}^d$ which is centered at X_k and has a random radius D_k . Assume the radii D_k i.i.d. with a common cumulative distribution function F .




Figure C.1: Continuum growth model in \mathbb{R}^2 . S_t is shown by the shaded area.

Let us enumerate points of N as follows. Let $X_0 = 0$, $T_0 = 0$, $D_0 = \gamma$, and

$S_0 = B(0, \gamma) = B_0$. Given (X_k, T_k, D_k) , $k = 0, \dots, n$,

$$T_{n+1} = \inf_k \{T_k : T_k > T_n, (X_k, T_k) \in N(\bigcup_{i=0}^n B_i \times \mathbb{R})\}.$$

For each T_n there is (a.s. unique) X_n and a random radius D_n .

Continuum growth process $S_t \subset \mathbb{R}^d$ is a Markov process which is constructed from the sequence of the B_i :

$$S_t = \bigcup_{i=0}^{n(t)} B_i,$$

where $n = \inf_k \{k : T_k \geq t\}$. Figure C.1 shows an example where $S_t \subset \mathbb{R}$.

For our purposes, the main property of Continuum growth model is that the size of the infected area grows asymptotically linearly.

Theorem C.1 [DHB04] Fix $d \geq 1$ and consider the d -dimensional continuum growth model with rate λ . Assume that

$$\int_0^\infty e^{-\varphi r} dF(r) < \infty \tag{C.1}$$

for some $\varphi < 0$ and let $S_0 \subset \mathbb{R}^d$ be arbitrary but bounded with strictly positive Lebesgue measure. Then there exists a real number $\mu > 0$ such that for any $\epsilon \in (0, \lambda\mu^{-1})$, almost surely

$$(1 - \epsilon)B(0, \lambda\mu^{-1}) \subseteq \frac{S_t}{t} \subseteq (1 + \epsilon)B(0, \lambda\mu^{-1})$$

for all sufficiently large t . Moreover, the time constant μ is given by

$$\mu = \lim_{n \rightarrow \infty} \frac{\mathbb{E}(\tilde{T}(\mathbf{n}))}{n} = \lim_{n \rightarrow \infty} \frac{\tilde{T}(\mathbf{n})}{n},$$

where $\mathbf{n} = (n, 0, \dots, 0)$ and $\tilde{T}(x) = \inf\{t : B(x, \gamma) \subseteq S_t\}$.

Notations

| Expression | Definition | Page |
|------------------------------|---|------|
| $\{X_i\}_{i=0}^{\infty}$ | The points of a Poisson point process. | 12 |
| $B(\mathbf{x}, r)$ | The closed ball centered on \mathbf{x} of radius r | 12 |
| $\mathcal{B}(\lambda, R)$ | The union of all balls of a Poisson Boolean model (also referred as <i>occupied region</i>). λ denotes the intensity of the underlying Poisson point process, and R is the reference random variable that give the distribution of the radii of the disks. | 12 |
| $\mathcal{B}(\lambda, r_b)$ | Same as $\mathcal{B}(\lambda, R)$, but with constant disk radii r_b . | 12 |
| $\mathcal{B}(\lambda, R, Z)$ | Blinking Poisson Boolean model. Based on $\mathcal{B}(\lambda, R)$, disks appear and disappear according i.i.d. on/off processes. The distribution of these processes is equal to the distribution of Z . | 115 |
| $\mathcal{G}(\lambda, r)$ | The connectivity graph associated with the Boolean model $\mathcal{B}(\lambda, r/2)$ | 12 |
| λ^* | Critical node density for percolation in the Boolean model | 15 |
| p^* | Percolation threshold in the independent bond percolation model. | 104 |
| $W(A)$ | Union of the connected components of $\mathcal{B}(\lambda, R)$ having non-empty intersection with the set A . | 12 |
| C_{∞} | The unbounded component in $\mathcal{B}(\lambda, R)$ | 15 |
| \mathcal{U} | The set of nodes belonging to the unbounded cluster | 15 |
| $\theta(\cdot)$ | Percolation function. In the Boolean model, gives the probability that the origin belongs to the unbounded cluster as a function of the average number of neighbors $\lambda\pi\mathbb{E}(R^2)$. | 15 |
| $\mathcal{C}_i(p)$ | The set of the closed curves surrounding Node i , which are included in $B(X_i, p)$ | 17 |
| $\mathcal{E}(\mathcal{C})$ | The set of edges of $\mathcal{G}(\lambda, r)$ that are cut by the closed curve \mathcal{C} . | 17 |
| $N_i(p)$ | The size of the minimal cutset within a circle of radius p centered on X_i . | 17 |
| N_i^{∞} | Minimal number of edges to remove in order to disconnect Node i from \mathcal{U} . | 17 |

| Expression | Definition | Page |
|----------------------------------|---|---------|
| \mathcal{S}_i | <i>Domain of Node i.</i> The set of the nodes that are still connected to Node i when the N_i^∞ edges have been removed (see definition of N_i^∞ above) | 20 |
| \mathcal{I}_i | <i>Island of Node i.</i> The largest domain that contains i . | 22 |
| $L(\mathbf{x})$ | Attenuation function. \mathbf{x} is a vector of \mathbb{R}^2 , representing distance traveled by the signal. | 32 |
| $l(x)$ | Isotropic attenuation function. x is a real number representing the distance. | 34 |
| $I(\mathbf{x})$ | Poisson shot-noise process. $I(\mathbf{x}) = \sum_k L(X_k - \mathbf{x})$. | 40 |
| β_{ij} | Signal-to-interference-plus-noise ratio (SINR) at Node j of the signal received from Node i . | 32 |
| β | SINR threshold for successful decoding. | 32 |
| N_0 | Power of the ambient noise at any receiver | 32 |
| γ | Inverse processing gain of the CDMA system. Weights the interference term in the SINR expression. | 32 |
| $\gamma^*(\lambda)$ | Maximum value of γ for percolation, when the node density is λ . This function is only defined for $\lambda > \lambda^*$. | 36 |
| $\mathcal{H}(\gamma, \lambda)$ | The connectivity graph associated with the STIRG model | 32 |
| $\mathcal{H}^t(\gamma, \lambda)$ | The same as above, but with a t -slots TDMA scheme | 56 |
| \mathbb{L}^2 | The set of edges joining adjacent vertices of \mathbb{Z}^2 . $\mathbb{L}^2 = \{ \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}^4 : x_1 - y_1 + y_2 - x_2 = 1 \}$ | 39, 104 |
| \mathcal{L} | The two-dimensional square lattice with edge length d . $\mathcal{L} = d \cdot \mathbb{L}^2$, i.e. $\mathcal{L} = \{ \langle \mathbf{v}, \mathbf{w} \rangle \in \mathbb{R}^4 : \mathbf{v} = d\mathbf{x}, \mathbf{w} = d\mathbf{y} \text{ and } \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{L}^2 \}$ | 39 |
| \mathcal{L}' | Dual lattice of \mathcal{L} , obtained by shifting \mathcal{L} by a vector $(d/2, d/2)$. | 43 |
| $O(\cdot)$ | $f = O(g)$ if $\limsup_{n \rightarrow \infty} \frac{f(n)}{g(n)} < +\infty$. | 73 |
| $\Omega(\cdot)$ | $f = \Omega(g)$ if $g = O(f)$. | 73 |
| $\Theta(\cdot)$ | $f = \Theta(g)$ if $f = O(g)$ and $g = O(f)$. | 73 |

Publications

Conference papers

- Olivier Dousse, Patrick Thiran and Martin Hasler. Connectivity vs Capacity in Dense Ad Hoc Networks. In *Proc. IEEE Infocom*, New York, June 2002.
- Olivier Dousse, François Baccelli and Patrick Thiran. Impact of Interferences on Connectivity in Ad Hoc Networks. In *Proc. IEEE Infocom*, San Francisco, April 2003.
- Olivier Dousse and Patrick Thiran. Connectivity vs Capacity in Dense Ad Hoc Networks. In *Proc. IEEE Infocom*, Hong Kong, March 2004.
- Olivier Dousse, Petteri Mannersalo and Patrick Thiran. Latency of wireless sensor networks with uncoordinated power saving mechanisms. In *Proc. Mobihoc*, Tokyo, May 2004.
- Massimo Franceschetti, Olivier Dousse, David Tse and Patrick Thiran. Closing the gap in the capacity of random wireless networks. In *Proc. ISIT*, Chicago, IL, July 2004.
- Olivier Dousse, Massimo Franceschetti and Patrick Thiran. The Costly Path from Percolation to Full Connectivity. In *Proc. Allerton Conference*, Monticello IL, September 2004.
- Olivier Dousse, Massimo Franceschetti and Patrick Thiran. Information theoretic bounds on the throughput scaling of wireless relay networks. In *Proc. IEEE Infocom*, Miami, March 2005.
- Benyuan Liu, Peter Brass, Olivier Dousse, Philippe Nain and Don Towsley. Mobility Improves Coverage of Sensor Networks. In *Proc. Mobihoc*, Urbana-Champaign, IL, May 2005.
- Olivier Dousse, Hung X. Nguyen and Patrick Thiran. Deux applications de processus ponctuels aux réseaux de communication. In *Proc. GRETSI'05 (20ème colloque GRETSI sur le traitement du signal et des images)*, Louvain-la-Neuve, Belgium, September 2005.
- Olivier Dousse, Christina Tavoularis and Patrick Thiran. Delay of Intrusion Detection in Wireless Sensor Networks. *Submitted to Infocom 2006*.

Journal papers

- Olivier Dousse and Patrick Thiran. Connectivity of self-organized ad hoc wireless networks. *IEEE Intelligent Systems*, 18(4):83–86, April 2003.
- Olivier Dousse, François Baccelli and Patrick Thiran. Impact of Interferences on Connectivity in Ad Hoc Networks. *IEEE/ACM Trans. on Networking*, 13(2):425–436, April 2005.
- Massimo Franceschetti, Olivier Dousse, David Tse and Patrick Thiran. Closing the gap in the capacity of random wireless networks. *Submitted to IEEE Trans. on Information Theory*.
- Olivier Dousse, Massimo Franceschetti and Patrick Thiran. Information theoretic bounds on the throughput scaling of wireless relay networks. *Submitted to IEEE Trans. on Information Theory and IEEE/ACM Trans. on Networking joint issue*.
- Olivier Dousse, Massimo Franceschetti, Nicolas Macris, Ronald Meester and Patrick Thiran. Percolation in signal-to-interference-ratio connectivity graphs. *In preparation, to be submitted to Annals of Applied Probability*.

Curriculum Vitæ

Olivier Dousse was born in 1974 in Monthey, Switzerland. He received the M.S. degree in Physics from École Polytechnique Fédérale de Lausanne (EPFL), Switzerland in April 2000. The same year, he joined the Pre-Doctoral School in Communication Systems at EPFL. In 2001, he was awarded a prize for being one of the three best student in the pre-doctoral program.

In September 2001, he joined the Laboratory for Computer communication and Applications (LCA), and started his PhD thesis under the supervision of Professor Patrick Thiran. There, he participated to the National Center of Competence in Research on Mobile Information and Communication Systems (NCCR-MICS). In summer 2003, he spent two months at U.C Berkeley as a visiting scholar at the Robotics Lab (EECS Department).

In 1998 and 1999, he worked in for the Laboratory for Experimental Physics (EPFL) as an assistant for the practical works in physics, where he designed and set up new experiments for students. From 1999 to 2004, he was a teaching assistant for the class of stochastic models for communication from 1999 to 2004. In 2004, he was promoted lecturer for the same class, and started teaching jointly with Prof. Thiran.

His current research interests are wireless networks and percolation theory. He is a reviewer for the IEEE/ACM Transactions on Networking, the ACM/Kluwer Wireless Networks Journal, Infocom, Sigmetrics, ISIT, Seconn and other conferences.

Some invited talks

- “Impact of interferences on connectivity of ad hoc networks”, *École Normale Supérieure*, Paris, May 2003.
- “Connectivity in ad hoc networks and percolation”, *University of California*, Berkeley, September 2003.
- “Connectivity in ad hoc networks and percolation”, *University of Massachusetts*, Amherst, October 2003.
- “Connectivity and interferences in ad hoc networks, a percolation approach”, *MICS Annual Workshop*, Monte Verità, October 2003.
- “The Costly Path from Percolation to Full Connectivity”, *Technical University*, Vienna, October 2004.
- “Asymptotic Properties of Wireless Multi-Hop Networks”, *University of Toronto*, Toronto, July 2005.

Bibliography

- [BB01] F. Baccelli and B. Błaszczyszyn. On a coverage process ranging from the Boolean model to the poisson voronoi tessellation, with applications to wireless communications. *Adv. Appl. Prob.*, 33(2):293–323, 2001.
- [BBCF03] L. Booth, J. Bruck, M. Cook, and M. Franceschetti. Ad hoc wireless networks with noisy links. In *Proc. Information Theory Symposium (ISIT)*, Yokohama, Japan, 2003.
- [BBFM03] L. Booth, J. Bruck, M. Franceschetti, and R. Meester. Continuum percolation and the geometry of wireless networks. *Annals of Applied Probability*, 13(2):722–731, 2003.
- [BBM03] F. Baccelli, B. Błaszczyszyn, and P. Muhlethaler. A spatial reuse aloha mac protocol for multihop wireless mobile networks. In *Proc. Allerton Conference*, Monticello, IL, 2003.
- [BBT02] F. Baccelli, B. Błaszczyszyn, and F. Tournois. Spatial averages of coverage characteristics in large cdma networks. *Wireless Networks*, 8:569–586, 2002.
- [Bet02] C. Bettstetter. On the minimum node degree and connectivity of a wireless multihop network. In *Proc. MobiHoc*, Lausanne, Switzerland, June 2002.
- [BH96] R. J. C. Bultitude and D. A. Hughes. Propagation loss at 1.8 ghz on microcellular mobile radio channels. In *Proc. PIMRC'96*, October 1996.
- [CG04] C.-F. Chiasserini and M. Garetto. Modeling the performance of wireless sensor networks. In *Proc. IEEE Infocom*, Hong Kong, March 2004.
- [CR89] Y.-C. Cheng and T. G. Robertazzi. Critical connectivity phenomena in multihop radio models. *IEEE Trans. Commun.*, 37:770–777, July 1989.
- [CT91] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley, New York, 1991.
- [CT00] J.-H. Chang and L. Tassiulas. Energy conserving routing in wireless ad-hoc networks. In *Proc. IEEE Infocom*, pages 22–31, Jerusalem, 2000.

- [Dal71] D. J. Daley. The definition of a multi-dimensional generalization of shot noise. *J. Appl. Prob.*, 8:128–135, March 1971.
- [DBT03] O. Dousse, F. Baccelli, and P. Thiran. Impact of interferences on connectivity of ad hoc networks. In *Proc. IEEE Infocom*, San Francisco, April 2003.
- [Dei03] M. Deijfen. Asymptotic shape in a continuum growth model. *Advances in Applied Probability*, 35(2):303–318, 2003.
- [DHB04] M. Deijfen, O. Häggström, and J. Bagley. A stochastic model for competing growth on \mathbb{R}^d . *Markov Processes and Related Fields*, 10:217–248, 2004.
- [DVJ88] D. J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes*. Springer, 1988.
- [eGMPS04] A. el Gamal, J. Mammen, B. Prabhakar, and D. Shah. Throughput delay trade-off in wireless networks. In *Proc. IEEE Infocom*, Hong Kong, March 2004.
- [eGNP⁺02] A. el Gamal, C. Nair, B. Prabhakar, E. Uysal-Biyikoglu, and S. Zahedi. Energy-efficient scheduling of packet transmissions over wireless networks. In *Proc. IEEE Infocom*, New York, June 2002.
- [Eph02] A. Ephremides. Energy concerns in wireless networks. *IEEE Wireless Communications*, 9(4):48–59, 2002.
- [FBC⁺05] M. Franceschetti, L. Booth, M. Cook, R. Meester, and J. Bruck. Continuum percolation with unreliable and spread out connections. *Journal of Statistical Physics*, 118(3/4), February 2005.
- [FBS04] M. Franceschetti, J. Bruck, and L. Schulman. A random walk model of wave propagation. *IEEE Trans. Antennas Propagat.*, 52(5):1304 – 1317, May 2004.
- [Gil61] E. N. Gilbert. Random plane networks. *SIAM J.*, 9:533–543, 1961.
- [GK98] P. Gupta and P. R. Kumar. Critical power for asymptotic connectivity in wireless networks. *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming*, 1998. Edited by W.M. McEneaney, G. Yin, and Q. Zhang, (Eds.) Birkhäuser.
- [GK00] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Trans. Inform. Theory*, 46(2):388–404, March 2000.
- [GM90] G. R. Grimmett and J. Marstrand. The supercritical phase of percolation is well behaved. *Proc. Royal Soc. London Ser. A*, 430:439–457, 1990.
- [GM04] Chao Gui and Prasant Mohapatra. Power conservation and quality of surveillance in target tracking sensor networks. In *Proc. Mobicom*, Philadelphia, September 2004.

- [GP60] E. N. Gilbert and H. O. Pollak. Amplitude distribution of shot noise. *Bell Syst. Tech. J.*, 39:333–350, 1960.
- [Gri99] G. R. Grimmett. *Percolation*. Springer, 1999.
- [GT01] M. Grossglauser and D. Tse. Mobility increases the capacity of ad-hoc wireless networks. In *Proc. IEEE Infocom*, Anchorage, Alaska, April 2001.
- [GV02] M. Gastpar and M. Vetterli. On the capacity of wireless networks: the relay case. In *Proc. IEEE Infocom*, New York, June 2002.
- [Har87] E. Harrison. *Darkness at night: a riddle of the universe*. Harvard University Press, Cambridge, MA, 1987.
- [Has93] H. Hashemi. The indoor radio propagation channel. *Proceedings of the IEEE*, 81(7):943–968, July 1993.
- [HPS97] O. Häggström, Y. Peres, and J. Steif. Dynamic percolation. *Ann. IHP Probab. et. Statist.*, 33:497–528, 1997.
- [JVK04] A. Jovicic, P. Viswanath, and S. R. Kulkarni. Upper bounds to transport capacity of wireless networks. *IEEE Trans. Inform. Theory*, 50(11):2555–2565, November 2004.
- [KEW02] B. Krishnamachari, D. Estrin, and S. Wicker. The impact of data aggregation in wireless sensor networks. In *Proc. of International Workshop of Distributed Event Based Systems (DEBS)*, Vienna, Austria, 2002.
- [KGG03] G. Kramer, M. Gastpar, and P. Gupta. Capacity theorems for wireless relay channels. In *Proc. Allerton Conf. on Commun., Control, and Comp.*, Monticello IL, October 2003.
- [Kin93] J. F. C. Kingman. *Poisson Processes*. Clarendon Press, Oxford, 1993.
- [KKP03] G. Kesidis, T. Konstantopoulos, and S. Phooha. Surveillance coverage of sensor networks under a random mobility strategy. In *Proc. IEEE Sensors*, October 2003.
- [KV02] S. R. Kulkarni and P. Viswanath. A deterministic approach to throughput scaling in wireless networks. In *Proc. IEEE International Symposium on Information Theory*, November 2002.
- [LBD⁺05] B. Liu, P. Brass, O. Dousse, P. Nain, and D. Towsley. Mobility improves coverage of sensor networks. In *Proc. Mobihoc*, Urbana-Champaign, IL, May 2005.
- [Lig85] T. M. Liggett. An improved subadditive ergodic theorem. *The Annals of Probability*, 13(4):1279–1285, 1985.
- [LLM89] V. Lakshmikantham, S. Leela, and A. A. Martynyuk. *Stability analysis of nonlinear systems*. M. Dekker, New York, 1989.

- [LLT03] B. Liu, Z. Liu, and D. Towsley. On the capacity of hybrid wireless networks. In *Proc. IEEE Infocom*, San Francisco, April 2003.
- [LS04] X. Liu and R. Srikant. An information theoretic view of connectivity in wireless sensor networks. In *Proc. IEEE Secon*, Santa Clara, CA, October 2004.
- [LSS97] T. M. Liggett, R. H. Schonmann, and A. M. Stacey. Domination by product measures. *Annals of Probability*, 25(1):71–95, 1997.
- [LT03] B. Liu and D. Towsley. On the coverage and detectability of wireless sensor networks. In *Proc. WiOpt*, Sophia Antipolis, France, 2003.
- [LT04] O. Lévêque and E. Telatar. Upper bounds on the capacity of ad-hoc wireless networks. In *Proc. Information Theory Symposium (ISIT)*, Chicago, Illinois, July 2004.
- [MR96] R. Meester and R. Roy. *Continuum percolation*. Cambridge University Press, 1996.
- [Pen97] M. Penrose. The longest edge of the random minimal spanning tree. *Ann. Appl. Probability*, 7:340–361, 1997.
- [PP96] M. Penrose and A. Pisztor. Large deviations for discrete and continuous percolation. *Adv. Appl. Prob.*, 28:29–52, 1996.
- [PPT89] T. K. Philips, S. S. Panwar, and A. N. Tantawi. Connectivity properties of a packet radio network model. *IEEE Trans. Inform. Theory*, 35:1044–1047, September 1989.
- [PS98] Y. Peres and J. Steif. The number of infinite clusters in dynamical percolation. *Probab. Th. Rel. Fields*, 111:141–165, 1998.
- [PUBeG01] B. Prabhakar, E. Uysal-Biyikoglu, and A. el Gamal. Energy-efficient transmission over a wireless link via lazy packet scheduling. In *Proc. IEEE Infocom*, pages 386–394, Anchorage, Alaska, 2001.
- [QTZ00] S. Quintanilla, S. Torquato, and R. M. Ziff. Efficient measurement of the percolation threshold for fully penetrable discs. *J. Phys. A: Math. Gen.*, 33:L399–L407, October 2000.
- [RCM02] V. Ramasubramanian, R. Chandra, and D. Mosse. Providing a bidirectional abstraction for unidirectional adhoc networks. In *Proc. IEEE Infocom*, New York, June 2002.
- [SB02] P. Santi and D. M. Blough. An evaluation of connectivity in mobile wireless ad hoc networks. In *Proc. IEEE DSN*, pages 89–98, Washington DC, June 2002.
- [SB03] P. Santi and D. M. Blough. The critical transmitting range for connectivity in sparse wireless ad hoc networks. *IEEE Trans. Mobile Comput.*, 2(1):25–39, March 2003.

- [SBV01] P. Santi, D. M. Blough, and F. Vainstein. A probabilistic analysis for the range assignment problem in ad hoc networks. In *Proc. Mobicoc*, pages 212–220, Long Beach CA, October 2001.
- [SC01] A. Sinha and A. Chandrakasan. Dynamic power management in wireless sensor networks. *IEEE Design and Test of Computers Magazine*, 18(2):62–74, 2001.
- [SKM95] D. Stoyan, W. Kendall, and J. Mecke. *Stochastic Geometry and its Applications*. John Wiley & Sons, Chichester, 1995.
- [SSS03] S. Shakkottai, R. Srikant, and N. Shroff. Unreliable sensor grids: Coverage, connectivity and diameter. In *Proc. IEEE Infocom*, San Francisco, April 2003.
- [SWR98] S. Singh, M Woo, and C.S. Raghavendra. Power-aware routing in mobile ad hoc networks. In *Proc. ACM Mobicom*, pages 181–190, Dallas, TX, October 1998.
- [Tel99] E. Telatar. Capacity of multi-antenna gaussian channels. *Eur. Trans. Telecommun.*, 10(6):585–595, 1999.
- [TG04] S. Toumpis and A.J. Goldsmith. Large wireless networks under fading, mobility, and delay constraints. In *Proc. IEEE Infocom*, Hong Kong, March 2004.
- [vdBMW97] J. van den Berg, R. Meester, and D. G. White. Dynamic Boolean models. *Stochastic Processes and their Applications*, 69:247–257, 1997.
- [VS99] V. Veeravalli and A. Sendonaris. The coverage-capacity tradeoff in cellular cdma systems. *IEEE Trans. Veh. Technol.*, 48:1443–1451, 1999.
- [WHE02] Y. Wei, J. Heidemann, and D. Estrin. An energy-efficient MAC protocol for wireless sensor networks. In *Proc. IEEE Infocom*, New York, June 2002.
- [XK04a] Liang-Liang Xie and P. R. Kumar. A network information theory for wireless communication: Scaling laws and optimal operation. *IEEE Trans. Inform. Theory*, 50(5):748–767, 2004.
- [XK04b] F. Xue and P. R. Kumar. The number of neighbors needed for connectivity of wireless networks. *Wireless Networks*, 10(2):169–181, 2004.
- [XK05] Liang-Liang Xie and P. R. Kumar. On the path-loss attenuation regime for positive cost and linear scaling of transport capacity in wireless networks. 2005. submitted to IEEE trans. Inform. Theory.
- [YH03] W. Ye and J. Heidemann. Medium access control in wireless sensor networks. Technical Report ISI-TR-580, USC, 2003.
- [YKP04] Y. Yu, B. Krishnamachari, and V. Prasanna. Energy-latency tradeoffs for data gathering in wireless sensor networks. In *Proc. IEEE Infocom*, Hong Kong, March 2004.