

On the Throughput Scaling of Wireless Relay Networks

Olivier Dousse, *Member, IEEE*,
Massimo Franceschetti, *Member, IEEE*, and
Patrick Thiran, *Member, IEEE*

Abstract—The throughput of wireless networks is known to scale poorly when the number of users grows. The rate at which an arbitrary pair of nodes can communicate must decrease to zero as the number of users tends to infinity, under various assumptions. One of them is the requirement that the network is fully connected: the computed rate must hold for any pair of nodes of the network. We show that this requirement can be responsible for the lack of throughput scalability. We consider a two-dimensional (2-D) network of extending area with only one active source-destination pair at any given time, and all remaining nodes acting only as possible relays. Allowing an arbitrary small fraction of the nodes to be disconnected, we show that the per-node throughput remains constant as the network size increases. As a converse bound, we show that communications occurring at a fixed nonzero rate imply a fraction of the nodes to be disconnected. Our results are of information theoretic flavor, as they hold without assumptions on the communication strategies employed by the network nodes.

Index Terms—Ad hoc networks, information theory, percolation theory, scaling laws, sensor networks, wireless networks.

I. INTRODUCTION

A completely wireless network consists of n nodes that communicate over a common wireless channel. A natural question that arises in such systems is how the throughput scales with the number n . Typically, there are two ways of letting 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 first paper [1] to address these problems considered the dense network case, and a traffic scenario where each node generates packets for a destination nonvanishingly far away. Using a network theoretic approach based on multihop communication, it showed a lower bound on the per-node rate of $\Omega(1/\sqrt{n})$ bit/s, if nodes are arbitrarily located; and a lower bound of $\Omega(1/\sqrt{n \log n})$ bit/s if nodes are randomly located, see also [2]. Note that these results optimistically rely on point to point connections delivering infinite power as nodes tend to be closer

to each other, which is a physical impossibility. When the physical constraint of bounded power is enforced, the results are corrected by introducing an additional $1/\sqrt{n}$ factor, see [3].

For extended networks, the works in [4] and [5] present information theoretic bounds of $\Theta(1/\sqrt{n})$ bit/s per node, for arbitrarily located nodes, assuming some natural power attenuation law in the wireless medium [6], which can be bounded without affecting the final result. Finally, using percolation theory arguments, it has been shown [7] that it is possible to achieve the same $\Omega(1/\sqrt{n})$ rate with randomly located nodes.

All works aforementioned consider an all-to-all communication scenario. Instead, when nodes are located at random, but there is only one active source-destination pair at any given time, while the remaining nodes act as possible relays, it has been shown that the transmission rate in dense networks can grow at most as $O(\log n)$, under the assumption that around each of the two active nodes there is a dead zone of finite nonzero radius without any node [8]. 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 [9].

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 [8]. 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 correspondence, 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 nonvanishing rate can be achieved in two-dimensional (2-D) extended networks and in the relay scenario of [9]. On the contrary, it turns out that in the 1-D case, a nonzero rate is impossible even if we allow an arbitrary large fraction of nodes to be disconnected.

Finally, we want to spend few words on the intuition behind these results. The original result of [8] 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 factor of the total power. Our result for extended relay networks of constant density can be seen as a consequence of percolation theory: by choosing the constant density of the nodes λ larger than a critical value λ_c , a giant connected component forms. Inside this component every pair of nodes can communicate at a constant bit rate. Percolation theory tells us even more: it follows from a result by Penrose and Pisztor [10] that this 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 counterbalanced 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 nonzero 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 proven 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 [9]; and that they diverge as the density increases, in agreement with [8]. Moreover, it is interesting that allowing a fraction of the nodes to

Manuscript received March 15, 2005; revised December 16, 2005. This work was supported in part by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation by Grant 5005-67322. The material in this correspondence was presented at INFOCOM 2005, Miami, FL, March 2005.

O. Dousse is with the Deutsche Telekom Laboratories, 10587 Berlin, Germany (e-mail: olivier.dousse@telekom.de).

P. Thiran is with the School of Computer and Communication Sciences, École Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland (e-mail: patrick.thiran@epfl.ch).

M. Franceschetti is with the Department of Electrical and Computer Engineering, University of California at San Diego, La Jolla, CA 92093 USA (e-mail: massimo@ece.ucsd.edu).

Communicated by R. Srikant, Guest Editor.

Digital Object Identifier 10.1109/TIT.2006.874537

be disconnected does not change the scaling law in [4], [5], as in an all-to-all communication scenario the $1/\sqrt{n}$ bottleneck is due to the cost of relaying packets for other nodes.

The rest of the correspondence is organized as follows. We consider two different connectivity models, the Boolean model and the information theoretic relay model. Section II summarizes the assumptions made for both models. We begin with the Boolean model in Section III, to stress the importance of dimensionality on the connectivity, and to set up preliminary results that will be useful for the information theoretic relay model. Section IV contains the two main results of our correspondence, a lower bound (Theorem 2) and upper bound (Theorem 3) on the fraction of nodes that can communicate at a given rate. The lower bound is proven using percolation theory and the computation of the achievable throughput along the shortest path between the source and destination in Appendix I. The upper bound is proven by establishing an ergodic property of shot noises in Appendix II. Finally, Section V concludes the correspondence.

II. NETWORK MODELS

We consider 1- and 2-D 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.

Connectivity results strongly depend on how we define the word “connected.” Throughout this correspondence, we will look at the two following connectivity models:

A. Boolean Connectivity

In the first -and simplest- model, we assume that two nodes are *directly connected* if the distance between them is less than a given distance (or *range*) r . Two nodes are *connected* if there exists a path of directly connected nodes joining them.

It turns out that this definition of connectivity leads to a well known model in stochastic geometry called Boolean model [11]. In fact, if we center a ball (segment or disk, depending on the dimensionality of the network) of radius $r/2$ on each node, we end up with a Boolean model where the clusters formed by overlapping balls correspond exactly to the connected components in our network.

B. Information Theoretic Connectivity

In this model, we adopt a more practical point of view, and 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 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 (as, e.g., in [12]), we assume that l is decreasing and such that

$$\int_{\mathbb{R}} xl(x)dx < \infty.$$

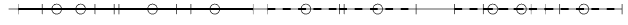


Fig. 1. 1-D Boolean model. Each node is the center of a segment of length r . The origin (on the left end) is connected to all nodes before the first gap (plain segments). All nodes on the right are disconnected (dashed segments).

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.

III. BOOLEAN CONNECTIVITY

In this section, we review connectivity results in the Boolean model. These results enlighten the fundamental difference between 1-D and 2-D connectivity and are then applied to prove our main theorems in the information theoretic model.

A. 1-D Network

In one dimension, connectivity is broken whenever there is a “gap” somewhere on the line (see Fig. 1). Here, in the Boolean model, connected components end when one finds two consecutive nodes separated by a distance greater than r . As the intervals between nodes are independently and exponentially distributed, when the network size increases, one finds arbitrarily many gaps (intervals longer than r) w.h.p. (see, e.g., [13]–[15]).

Moreover, if we look at the number of nodes connected to the origin, we observe that this number is a random variable with finite expectation. In fact, at each side of the origin, the number of intervals we have to look at before we find a gap, is a geometrical random variable. Therefore, the probability that the number of nodes connected to the origin is finite is equal to one. Now, if we compute the fraction of nodes that are connected to the origin, we have to divide this finite number by the total number of nodes n . As n goes to infinity, the fraction tends thus to zero a.s. This means that the network is completely disconnected w.h.p, regardless of the value of r .

B. 2-D Network

Contrary to the 1-D case, here we can have isolated nodes without disconnecting the whole network. The intuition is that if a gap arises, there can be a path that avoids it, and keep long-range connectivity (see Fig. 2).

This intuitive observation is confirmed by percolation theory: if we consider a Boolean model that spreads over the whole plane, there exists a critical radius r^* , such that if $r > r^*$, there exists one unbounded connected component a.s. We call *percolation probability* the probability $0 < \theta(r) < 1$ that an arbitrary node belongs to this unbounded component. We have also that $\theta(r)$ tends to 1 when r tends to infinity. On the contrary, if $r < r^*$, all connected components remain a.s. bounded, and the percolation probability is therefore equal to zero. This latter case (also called *subcritical case*) is similar to the 1-D case. See [11] for an extensive treatment of this subject.

In our extended network model, we only look at the box $B_n = [0, \sqrt{n}] \times [0, \sqrt{n}]$ and we make use of the following result by Penrose and Pisztor [10].

Theorem 1 (Penrose and Pisztor, 1996): Suppose $r > r^*$, and $0 < \varepsilon < 1/2$. Let $E(n)$ be the event that

- there is a unique connected cluster C inside B_n containing more than $\varepsilon\theta(r)n$ points;
- $(1 - \varepsilon)\theta(r)n \leq \text{Card}(C) \leq (1 + \varepsilon)\theta(r)n$.

Then there exist constants $c_1 > 0$ and $n_0 > 0$ such that

$$\mathbb{P}[E(n)] \geq 1 - \exp(-c_1\sqrt{n}), \quad n \geq n_0.$$

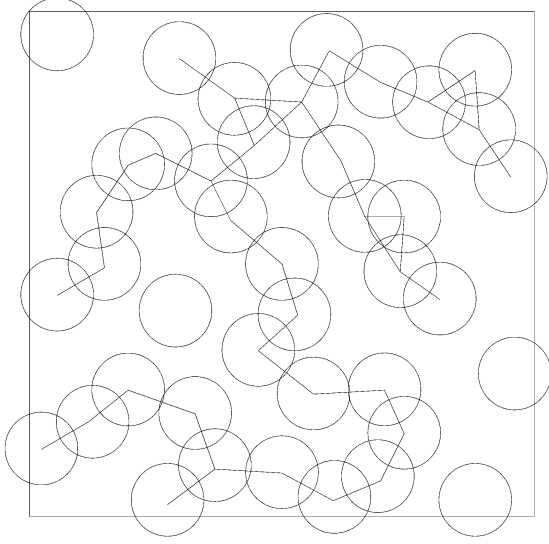


Fig. 2. 2-D Boolean model. Although there are some isolated nodes, the origin is connected with a large fraction of the nodes.

We deduce from the above theorem the following corollary.

Corollary 1: For any $0 < \tilde{\theta} < 1$, there exists an $r < \infty$ such that there exists a connected cluster that contains a fraction of nodes larger than $\tilde{\theta}$.

Proof: $\theta(r)$ is continuous for $r > r^*$ and tends to 1 when r tends to infinity, see [11]. Thus, given a fraction $\tilde{\theta}$, one can choose r large enough so that $\theta(r) = \tilde{\theta}/(1 - \varepsilon)$ for some $\varepsilon > 0$. Theorem 1 ensures that the number of nodes inside C is at least $(1 - \varepsilon)\theta(r)n$ w.h.p. Thus the fraction of nodes inside this cluster is larger than $\tilde{\theta}$. \square

This result makes a great difference with the 1-D case, as now, with a fixed r , we can have nonzero fraction of connected nodes for an arbitrarily large network.

However, if we require *full connectivity*, i.e., if we impose $\theta(r) = 1$, then the price to pay is very high: we know that this is not possible with fixed r , because there will be a.s. a node that has no neighbor within distance r . Actually, Penrose [16], and independently Gupta and Kumar [17], showed that r has to increase like $\log n$ to achieve full connectivity.

This big difference (constant vs increasing r as a function of n) follows from the fact that if we want to connect the most isolated node in the network, as the network size increases, we find worse and worse cases, for pure statistical reasons. In practice, it is not much of a problem if a tiny fraction of the nodes are disconnected, especially as it then allows r to no longer increase with n .

This latter result matches the intuition that as the node density remains constant, the quality of the connectivity should not change with the size of the network. In fact, the explanation behind the result in [16], is of statistical nature.

IV. INFORMATION THEORETIC CONNECTIVITY

In this section, we use results from Boolean connectivity to obtain new results under the information theoretic definition of connectivity.

A. 1-D Network

For 1-D networks similar results hold 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 [9], 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 $[x, n]$ 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 Section IV-B.

B. 2-D 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), the asymptotic behavior of the rate dramatically changes, and stays constant when n tends to infinity.

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

Proof: To prove this theorem, we use Corollary 1. Given $\tilde{\theta}$, this ensures that there exists r such that under the Boolean model, there exists a connected cluster of size greater than $n\tilde{\theta}$ w.h.p.

Then, we use Theorem 4 in Appendix I to show that along the shortest path in the Boolean model, the throughput

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

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

The simple TDMA construction of Theorem 4 described in Appendix I, along with the percolation theory result of Corollary 1 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 [4] and [18], that leads in our case to an achievable throughput

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

However, the asymptotic behavior remains the same even with our very simple scheme and only the preconstant 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.

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 3: 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

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

where I is the *shot-noise* defined by

$$I = \sum_{x \in N} l(\|x\|)$$

and N is a Poisson point process of unit density over \mathbb{R}^2 .

Proof: To prove Theorem 3, 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 [19, max-flow min-cut Th. 14.10.1]. 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 $x_i, i \neq 0$. 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 [20] and corresponds to the rate at which the emitter x_0 can send data to the rest of the nodes x_i

$$R = \frac{1}{2} \log \left(1 + \frac{P \sum_{i=1}^{\infty} l(\|x_i - x_0\|)}{N_0} \right).$$

The sum in the expression is a shot-noise process evaluated in x_0 , that we will denote by S later on

$$S(x_0) := \sum_{i=1}^{\infty} l(\|x_i - x_0\|).$$

From this result, we conclude that a necessary condition for achieving rate R from node x_0 is

$$S(x_0) \geq \frac{N_0}{P} \left(2^{2R} - 1 \right) := M.$$

We can then use Theorem 5 in Appendix II 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 1 in Appendix II), we conclude from this theorem that the fraction of nodes that do not fulfill the condition is equal to

$$\mathbb{P}(I < M) := 1 - \hat{\theta}$$

where

$$I = \sum_{x \in N} l(\|x\|)$$

and N is a Poisson point process of intensity 1. As they do not fulfill the necessary condition, the fraction $1 - \hat{\theta}$ of nodes cannot for sure 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 3 is not explicit, as it involves the cumulative distribution function of a shot-noise; Lemma 1 in Appendix II shows that it is always strictly smaller than one. To obtain an explicit bound, we can use Chernoff's inequality. We know from Campbell's theorem that (see [21, p. 28])

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

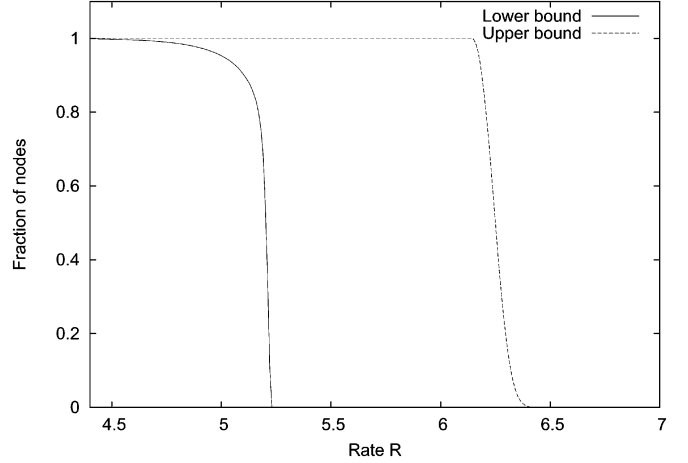


Fig. 3. Upper and lower bounds on the fraction of nodes that can achieve a given rate R . The lower bound has been improved using the multiple relay channel result given in (1), and the upper bound has been computed using the bound presented in Section IV-B2.

Therefore

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

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

C. Discussion

We plot in Fig. 3 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 correspond to the results in [9].

At the other end, 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 nonzero 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.

V. CONCLUSION

In wireless random networks with a finite spatial density of nodes, the price to pay for full connectivity is high: it makes the throughput of any node vanish when the network size gets large. It is impossible to offer a nonzero rate to every node of an extended network when the number of nodes tends to infinity. 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), Theorem 3 shows that 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 nonzero, but arbitrary small fraction θ of nodes to be disconnected, then Theorem 2 shows that it is possible to find a rate $R(\theta)$ that any other pair of nodes can enjoy in an arbitrarily large network. The theorem is proven by continuum percolation techniques, and therefore holds for 2-D networks, but not for 1-D 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 3) and sufficient (Theorem 2) to have nodes communicating at a nonzero, 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.

We therefore believe that these results reinforce the case for “partial connectivity” (or “ θ -connectivity”), where a fraction θ of the nodes randomly drawn from the network is connected, as opposed to the traditional full connectivity.

APPENDIX I PATH THROUGHPUT

In a Boolean model with fixed ball radius $r/2$ and spatial density 1, we consider the shortest path between two points x_1 and x_2 that belong to the same cluster. We show in the following theorem an achievable throughput along this path, that does not depend on the number of hops, but only on r .

Theorem 4: In a Boolean model with unit point density and ball radius $r/2$, 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 Fig. 4), 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.

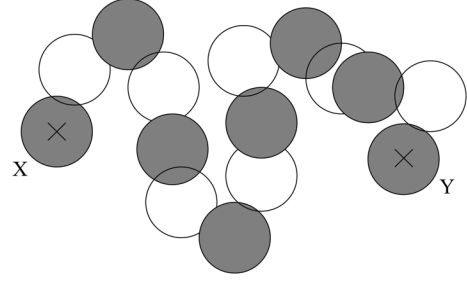


Fig. 4. A shortest path from X to Y in a Boolean model. Filled balls do not overlap, and neither do the empty balls.

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

$$\begin{aligned} & \sum_{i=1}^{\infty} Pl(x_{2i} - x_0) \\ & \leq P \left[6l(r) + 6l(\sqrt{3}r) + 6l(2r) \right. \\ & \quad \left. + 12l \left(\left(2 + \frac{\sqrt{3}}{2} \right) r \right) + \dots \right] \\ & \leq P \sum_{k=1}^{\infty} 6(k+1)l(kr) := PK(r). \end{aligned}$$

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

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

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 (2), since we used three time slots. \square

APPENDIX II SHOT-NOISE

We start this section by introducing some notation. For a set $A \subset \mathbb{R}^d$ and a point $x \in \mathbb{R}^d$, we define the set $A + x$ as follows:

$$y \in A + x \text{ iff } y - x \in A.$$

Then we define the shift operator S_x as follows: for a random measure N and a Borel set A

$$S_x N(A) = N(A + x).$$

We define now formally the shot-noise process built on a 2-D Poisson point process, and prove an ergodic result in Theorem 5.

Let N be a Poisson point process of intensity λ in \mathbb{R}^d . We define the *shot-noise* as

$$I = \int_{\mathbb{R}^d \setminus \{0\}} f(x) N(dx) = \sum_{x \in N} f(x)$$

for some positive function f such that

$$\int_{\mathbb{R}^d} f(x)dx < \infty. \quad (3)$$

Given a constant $M > 0$, we define by Y the random counting measure that counts the number of points of N such that the value of the shot-noise at this point is less than M

$$Y(A) = \int_A 1_E(S_x N)N(dx)$$

where E is the event

$$E = \{I < M\}.$$

The mean density of the new point process Y is clearly equal to the intensity of N , 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 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 5: 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 6 stated here, and the observation that the point process Y is stationary and ergodic, since the underlying point process N is stationary and ergodic. Then we use [22, Prop. 10.3.III] to prove that Y is also metrically transitive, since it is ergodic.

Theorem 6 ([22, Corollary 10.2.V]): 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 $n \rightarrow \infty$

$$\xi(A_n)/\ell(A_n) \rightarrow m \text{ (a.s. and in } L_1 \text{ norm)}.$$

Finally, we prove that $\mathbb{P}(E)$ is strictly positive.

Lemma 1: For any λ and M , we have $\mathbb{P}(E) > 0$.

Proof: As f verifies (3), we can find a number $L > 0$ such that

$$\int_{\mathbb{R}^d \setminus B(0,L)} f(x)dx < \frac{M}{2\lambda}$$

where $B(0, L)$ denotes the ball of radius L centered at the origin. Let us define

$$I_L = \sum_{x \in N \cap B(0,L)} f(x)$$

and

$$I^L = \sum_{x \in N \setminus B(0,L)} f(x)$$

so that $I = I_L + I^L$. Then we have $\mathbb{P}(I_L < M/2) > 0$ since $N \cap B(0, L)$ has nonzero probability to be empty. Furthermore, by Markov's inequality and Campbell's theorem we have

$$\begin{aligned} \mathbb{P}(I^L < \frac{M}{2}) &= 1 - \mathbb{P}(I^L \geq \frac{M}{2}) \\ &\geq 1 - \frac{2E[I^L]}{M} \\ &= 1 - \frac{2\lambda \int_{\mathbb{R}^d \setminus B(0,L)} f(x)dx}{M} \\ &> 0. \end{aligned}$$

Since the two events are independent, their joint probability is still positive, and thus, $\mathbb{P}(I < M) > 0$. \square

REFERENCES

- [1] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Inf. Theory*, vol. 46, pp. 388–404, Mar. 2000.
- [2] S. Toumpis and A. Goldsmith, "Large wireless networks under fading, mobility, and delay constraints," in *Proc. IEEE INFOCOM*, Hong Kong, Mar. 2004.
- [3] O. Dousse and P. Thiran, "Connectivity vs capacity in ad hoc networks," in *Proc. IEEE INFOCOM*, Hong Kong, Mar. 2004.
- [4] L.-L. Xie and P. R. Kumar, "A network information theory for wireless communication: Scaling laws and optimal operation," *IEEE Trans. Inf. Theory*, vol. 50, pp. 748–767, 2004.
- [5] O. Lévêque and E. Telatar, "Upper bounds on the capacity of ad-hoc wireless networks," in *Proc. Inf. Theory Symp. (ISIT)*, Chicago, IL, Jul. 2004.
- [6] M. Franceschetti, J. Bruck, and L. Schulman, "A random walk model of wave propagation," *IEEE Trans. Antennas Propag.*, vol. 52, pp. 1304–1317, May 2004.
- [7] M. Franceschetti, O. Dousse, D. Tse, and P. Thiran, "Closing the gap in the capacity of random wireless networks," in *Proc. Inf. Theory Symp. (ISIT)*, Chicago, IL, Jul. 2004.
- [8] M. Gastpar and M. Vetterli, "On the capacity of wireless networks: The relay case," in *Proc. IEEE INFOCOM*, New York, Jun. 2002.
- [9] X. Liu and R. Srikant, "An information theoretic view of connectivity in wireless sensor networks," in *Proc. IEEE Secon*, Santa Clara, CA, Oct. 2004.
- [10] M. Penrose and A. Pisztor, "Large deviations for discrete and continuous percolation," *Adv. Appl. Prob.*, vol. 28, pp. 29–52, 1996.
- [11] R. Meester and R. Roy, *Continuum Percolation*. Cambridge, U.K.: Cambridge Univ. Press, 1996.
- [12] O. Dousse, F. Baccelli, and P. Thiran, "Impact of interferences on connectivity of ad hoc networks," *IEEE/ACM Trans. Netw.*, vol. 13, pp. 425–436, Apr. 2005.
- [13] D. Stoyan, W. Kendall, and J. Mecke, *Stochastic Geometry and its Applications*. Chichester, U.K.: Wiley, 1995.
- [14] O. Dousse, P. Thiran, and M. Hasler, "Connectivity in ad-hoc and hybrid networks," in *Proc. IEEE INFOCOM*, New York, Jun. 2002.
- [15] D. Miorandi and E. Altman, "Connectivity in ad-hoc networks: A queueing theoretical approach," in *Proc. Worksh. Model. Optimiz. Mobile, Ad Hoc, Wireless Netw.*, Cambridge, U.K., Mar. 2004.
- [16] M. Penrose, "The longest edge of the random minimal spanning tree," *Ann. Appl. Prob.*, vol. 7, pp. 340–361, 1997.
- [17] P. Gupta and P. R. Kumar, "Critical power for asymptotic connectivity in wireless networks," in *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming*, W. McEneaney, G. Yin, and Q. Zhang, Eds., Birkhäuser, 1998.
- [18] G. Kramer, M. Gastpar, and P. Gupta, "Capacity theorems for wireless relay channels," in *Proc. Allerton Conf. Commun., Control, Comp.*, Monticello, IL, Oct. 2003.
- [19] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
- [20] E. Telatar, "Capacity of multi-antenna Gaussian channels," *Eur. Trans. Telecommun.*, vol. 10, no. 6, pp. 585–595, 1999.
- [21] J. F. C. Kingman, *Poisson Processes*. Oxford, U.K.: Clarendon, 1993.
- [22] D. J. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes*. New York: Springer, 1988.