A case for partial connectivity in large wireless multi-hop networks

Olivier Dousse
LCA (Station 14)
EPFL
CH-1015 Lausanne, Switzerland
Email: olivier.dousse@epfl.ch

Massimo Franceschetti
Dept. of Electrical and Computer Engineering
University of California
San Diego, CA, USA
Email: massimo@ece.ucsd.edu

Patrick Thiran
LCA (Station 14)
EPFL
CH-1015 Lausanne, Switzerland
Email: patrick.thiran@epfl.ch

Abstract— **As in every other network, nodes of a wireless multi-hop network must be well connected. Unlike many other networks, some wireless multi-hop networks (e.g., sensor networks), do not necessarily require full connectivity of all the nodes of the networks. In this paper, we show the benefits of replacing this usual requirement by that of a partial, η-connectivity, where only a given fraction $\eta < 1$ of the nodes needs to be connected to the network. This feature is found in models neglecting interferences, taking interferences as noise or taking a more information theoretic approach. The paper summarizes and updates the results in [1].**

I. Introduction

Connectivity is the most critical property that any communication network must have. Similarly to wired and cellular networks, this property is usually understood as the *full connectivity* of the wireless multi-hop network: there is a (possibly multi-hop) path that connects any pair of nodes in the network. Full connectivity is however very expensive, as it forces the maximal power to be large enough to connect the most isolated node to the network, and yields rather poor performances for large networks. We can relax this requirement by declaring the network to be connected if a given fraction η of the nodes are connected to each other, which we will refer to as *η-connectivity*. Allowing $\eta < 1$ instead of imposing $\eta = 1$ brings indeed a considerable gain in scalability for network connectivity [1], [2]. It is due to a phase transition (percolation) that separates a sub-critical regime where all clusters of connected nodes in the network are small compared to the total number of nodes, and a super-critical regime where a giant cluster emerges, which contains a large fraction of the nodes. We can impose that this fraction is at least a given value $\eta < 1$, even for networks reaching an infinite size, because the phase transition occurs for finite values of the parameters of the network (node density, powers, lengths of the codewords). In contrast, if we request that $\eta = 1$, we need to let at least one such parameter to infinity when the number of nodes also tends to infinity. Full connectivity is required in some scenarios, where all users must be granted access to the network, but many applications of sensor networks only need the network to cover a given area of interest, with a sufficient number of sensors connected to the base station. An η-connectivity metric is well suited to this setting.

We consider three different models: the simplest Boolean model, with a circular connectivity range for each node (Section II), a model with interferences, where nodes connect if their signal to noise ratio exceeds some given threshold (Section III), and a more "information theoretic" model, where nodes connect if they can exchange data at a rate larger than some given rate (Section IV). This shows therefore that the high cost of full connectivity can be spared without prejudice.

On the contrary, trading full connectivity for a giant component makes it possible to construct a scheduling and routing scheme that matches the upper bound on the transport capacity of the network [3] and offers multiple paths between most nodes (Section V).

II. The Boolean model

The first and simplest model for tackling the connectivity issue is the Boolean model. The main assumption is that nodes have a *connectivity range* r, within which they can wirelessly connect to their neighbors. We assume furthermore that the range is the same for all nodes. Therefore, two nodes are directly connected if the distance between them is less than r.

We take now a finite area, 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).

As a first observation, it is clear that the probability that the network is fully connected is always smaller than 1, 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 1. 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 a wireless network that requires η-connectivity, one can say that a network is still *well connected*, if disconnected nodes may exist but always represent a small fraction (at most $1 - \eta$) 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 multihop networks, we would like most nodes to belong to the same giant connected component (which forms the network itself).

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 Pisztora [4] showed that for a large but finite area, the fraction of connected nodes is always close to the deterministic function $\theta(\pi\lambda r^2)$. If we let the area of the network tend to infinity (and thus also the number of nodes), then the fraction η of connected nodes tends to the constant $\theta(\pi\lambda r^2)$. This result matches the previous 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 may occur with positive probability. Suppose that the area on which the network is deployed is a square $[0, \sqrt{n/\lambda}] \times [0, \sqrt{n/\lambda}]$, so that the expected number of nodes in the network is n. Following Penrose [5] and Gupta and Kumar [6], one can state that the network is fully connected w.h.p. (i.e, asymptotically almost surely as $n \to \infty$) if and only if the range $r(n)$ is such that $\pi\lambda r^2(n) - \log n \to \infty$. Hence the average node degree must grow approximately like $\log n$, when the number of nodes in the network increases, contrary to the emergence of a giant cluster that appears with high probability as soon as the average node degree exceeds N^*. Moreover, it is the *most isolated node* that determines the critical range for full connectivity. When the network is strongly super-critical, but not yet full connected, the nodes that do not belong to the giant clusters are isolated with high probability [5], [1]. The network becomes fully connected when the last node joins the network. This means that full connectivity is not really a global property of the network, but rather depends on the "level" of isolation of the most isolated node. This also explains why the range of the nodes has to increase, even though the node density remains constant.

III. Model taking Interferences into account

We now modify the Boolean model of the previous section to account for interferences. We keep the Poisson distribution of the nodes, but assume that two nodes are neighbors if they can communicate despite the interferences from all other nodes. More precisely, there is a link between Nodes i and j if

$$\frac{Pl(||x_i - x_j||)}{N_0 + \gamma \sum_{k\neq i,j} Pl(||x_k - x_j||)} > \beta \quad \text{and}$$
$$\frac{Pl(||x_j - x_i||)}{N_0 + \gamma \sum_{k\neq i,j} Pl(||x_k - x_i||)} > \beta, \tag{1}$$

where P is the emitting power of all nodes, $l(d)$ is the attenuation of the signal over distance d, which we assume to have the form $l(d) = \min\{1, d^{-\alpha}\}$ with $\alpha > 2$, N_0 is the ambient noise, γ is a coefficient that weights the efficiency of a CDMA system (degree of orthogonality of the codewords) and β is the minimal Signal to Noise and Interferences Ratio (SNIR) required for successful decoding.

Since we take interferences into account, it is not enough to increase the range (i.e. the emitting power) of the nodes to keep the network connected. Increasing the emitting power P would also increase interferences, and at the limit $P \to \infty$, one can see in (1) that the SNIR converges to a maximum.

However, it has been proven in [7], [8] that percolation can still occur in this model, for appropriate values of λ, P, N_0, γ and β. We assume in the sequel that λ and N_0 are fixed, that P is a parameter, as well as either β or γ. Varying β amounts to adapting the data rate, for example as in IEEE 802.11, whereas varying γ corresponds to changing the lengths of the codewords in a CDMA system.

Let us consider the case where the network spreads over the whole plane first, and denote by $\theta(P, \beta)$ (respectively, $\theta(P, \gamma)$) the probability that a given node belongs to an infinite connected component (i.e. is connected to the network), when γ (resp., β) is fixed. The following proposition can be inferred from the results in [7], [8].

Proposition 1: (i) For any value of P, there exists a critical value $\beta^*(P)$ such that if $\beta < \beta^*(P)$, then there is an infinite connected component a.s., and $\theta(P, \beta) > 0$.
(ii) If $N_0 = 0$, then for any value of P, there exists a critical value $\gamma^*(P)$ such that if $\gamma < \gamma^*(P)$, then there is an infinite connected component a.s., and $\theta(P, \gamma) > 0$.

Although it is not formally proven as in [4], we can also reasonably conjecture that that for a finite network, the fraction of connected nodes tend to θ when the network area grows.

For full connectivity, we must study more precisely the shape of the functions $\theta(P, \beta)$ and $\theta(P, \gamma)$. Simulation based evaluations of $\theta(P, \beta)$ are presented on Figure 1. Observe that if β is fixed and P arbitrarily large, the fraction of connected nodes $\theta(P, \beta)$ is bounded above by a constant $\overline{\theta}(\beta) < 1$. In fact, letting $P \to \infty$ is equivalent to let $N_0 \to 0$ in (1). Therefore, at the limit, we obtain a purely interference-limited network. It can be shown that such a network always contain a non-zero fraction of disconnected nodes [8]. A similar observation holds if we fix β and let γ vary instead.

Now, to have full connectivity of the n-nodes network, we must have $\theta(P, \beta) \to 1$ when $n \to \infty$. The only way to achieve $\eta = 1$ is to have $\beta \to 0$, as shown on Figure 2, or vice-versa, to have $\gamma \to 0$. The rate on each link has therefore to decrease with the number of nodes.

Finally, instead of letting all nodes emit with power P at all times, which results in strong interferences, we could use a random TDMA scheme, where each node picks a time slot at random, out of t time slots. Only a fraction $1/t$ of the nodes are emitting at a given time, making the average interference level t times lower. For fixed P, γ and β, one can now increase the number of time slots t to improve the network connectivity.

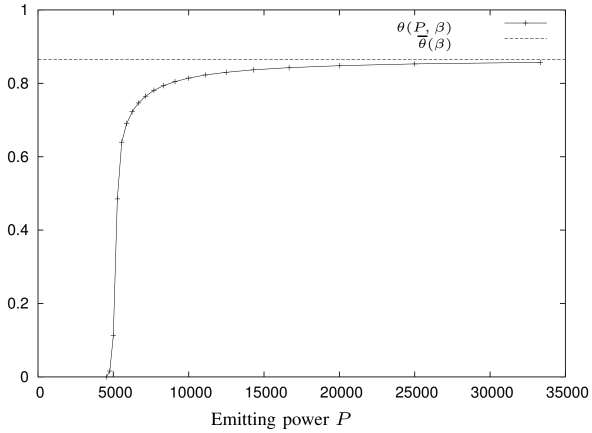

Fig. 1: The fraction of connected nodes, as a function of the emitting power P, for a fixed $\beta = 0.02$, with fixed $P \simeq 7.7 \cdot 10^3$. The function converges to a value $\bar{\theta}(\beta) < 1$.

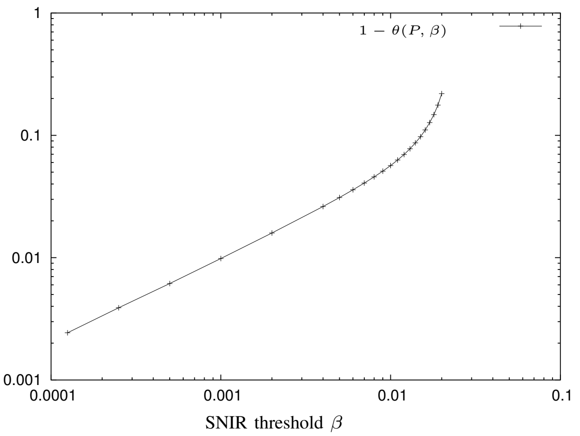

Fig. 2: The fraction of disconnected nodes, as a function of the SNIR threshold β, with fixed $P \simeq 7.7 \cdot 10^3$.

This setting has also been studied in [7], and it has been shown that the network can be made at least super-critical (i.e. partially connected) by setting t sufficiently large.

But whether we decrease β or γ, and/or increase t, we always end up by reducing the throughput per link. It turns out in this model that increasing the emitting power does not allow to reconnect the network. The only way to achieve full connectivity is to reduce the throughput.

IV. Information theoretic connectivity

In this section, we consider a slightly more sophisticated (and more realistic) definition of connectivity: two nodes are connected, if it is possible to transmit data from one node to the other at rate at least R, all other nodes of the network serving as relays. Nodes are still distributed according to a Poisson point process of intensity λ over an area of size A. This model has been first introduced by Liu and Srikant [9]. It has in fact many more parameters than the Boolean model, like the maximum emitting power of the nodes, the attenuation exponent or the ambient noise. Here, we assume them all fixed, but the node density λ, the network area A and the rate R.

We raise the same couple of questions as above. Is the network is fully connected ? Does it contain a giant connected component with most nodes? It turns out that the first question leads to asymptotic results similar to those obtained for the Boolean model:

Theorem 1 ([9]): For an attenuation function of the form $l(d) = d^{-\alpha}$, with $\alpha > 1$, and if $R(A) \geq \frac{c_2 \lambda^\alpha}{(\log \lambda A)^{\alpha-1}}$ where $c_2 = \frac{48P4^{\alpha-1}(1+\varepsilon)(2\alpha-1)}{(1-\varepsilon)^\alpha N_0(\alpha-1)}$ for some $\varepsilon > 0$, the network is disconnected w.h.p. when $A \to \infty$.

We observe that when the network size A increases, the rate must decrease to keep the network connected. Equivalently, if we fix the rate, then the emitting power of the nodes should increase. Again, this result is explained by a pure statistical effect, as the probability to find an arbitrarily isolated node in the network tends to 1 when $A \to \infty$. This argument is used in [9] to prove Theorem 1.

However, if we only require partial connectivity, with an arbitrarily low fraction $1-\eta$ of disconnected nodes, the asymptotic behavior of R dramatically changes: R can be kept constant when the network size becomes infinite. The following theorem has been proven in [10], when the network area A is infinite.

Theorem 2: (i) For any $0 < \check{\theta} < 1$, there exists a rate $R > 0$ independent of n, such that a fraction at least $\check{\theta}$ of the nodes can exchange data at rate R w.h.p.
(ii) For any rate $R > 0$, the fraction of nodes that can send data to some destination at that rate w.h.p.

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

where I is the *shot-noise* defined by $I = \sum_{x \in N} l^2(||x||)$, with N a Poisson point process of unit density over $\mathbb{R}^2$.

In other words, when the network size tends to infinity, a constant fraction of the nodes have the property that a path with throughput R can be established between any pair of them, even the two nodes are far away from each other. The converse says that for any rate R, there will be a non-zero fraction of nodes that cannot communicate with any other node with rate at least R. Again, with this model of connectivity, the results have the same flavor: requiring full connectivity makes the network performance drop when the number of nodes is large. On the contrary, if we only require partial connectivity (even with a very small fraction of disconnected nodes), the network performance scales perfectly.

V. Number of paths and throughput scaling

The previous sections compared two metrics of network connectivity, and showed that the price to pay for full connectivity is scalability. In this section, we show that a similar trade-off exists in term of throughput scaling, when the traffic matrix is uniform. In the previous section, we considered the case where only one link is active at a time. Although this assumption is realistic for lightly loaded sensor networks, it is also interesting to consider the other extreme case, where all nodes want to transmit data at the same time. In this case,

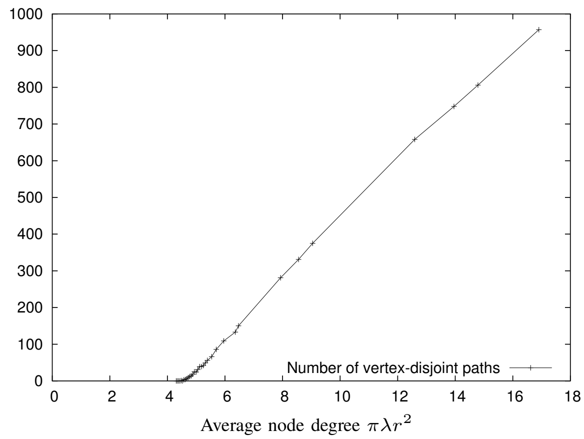

Fig. 3: Number of vertex-disjoint paths in a Boolean model.

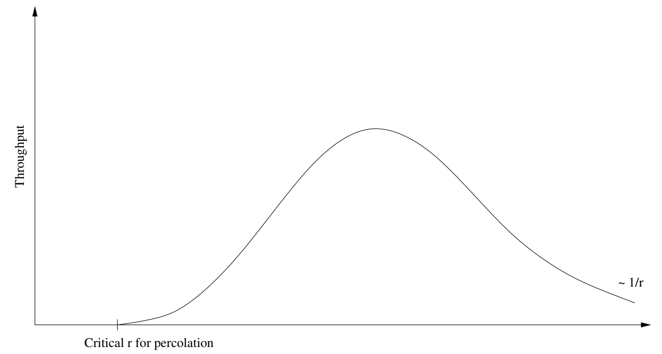

Fig. 4: Horizontal throughput as a function of the average node degree.

if *all* pairs of sending and receiving nodes need to be granted the same minimal rate, it is known since Gupta and Kumar's result [11] that the latter must decrease as $1/\sqrt{n}$ for n nodes. On the other hand, if we only require that a fraction η of the nodes (more precisely the nodes in the giant component in the network) be granted a minimal throughput, the latter can be much larger. In this context, the throughput is directly related to the number of vertex-disjoint paths that can be established in the percolation giant cluster, under the protocol model, as defined in [11] and the model presented in Section III.

We assume that the network is deployed on a squared area $[0, \sqrt{n}] \times [0, \sqrt{n}]$, with a node density $\lambda = 1$, and thus an average number of nodes equal to n. We want to compute the throughput that can be achieved from the nodes on the left border to the nodes close to the right border of the square, under a uniform traffic matrix for all other nodes in the square.

We start with the Boolean model, where each node has the same connecting range r. We try here to find an optimal value of r that maximizes the aggregate throughput from the left side of the network to the right side. We know that if the range r is too small, the network is sub-critical, hence we will not find a path crossing the network from left to right. Therefore, r has to be such that $\pi\lambda r^2 = \pi r^2 > N^*$. Furthermore, when r increases, the connectivity graph becomes richer, and the number of left-right crossings increases, as shown in Figure 3.

However, long ranges cause many interferences. To evaluate their impact on throughput, we use the protocol model [11]: each transmission occupies a footprint of area $\pi(\Delta r)^2$. The maximum number of simultaneous transmissions cannot be larger than $n/\pi\Delta^2 r^2 \lambda$. On the other hand, each transmission can transport data at a distance at most r. Thus, the total transport capacity is $n/\pi\Delta^2 r\lambda$. As the length of a left-right crossing is at least $\sqrt{n}$, the horizontal throughput across the network cannot be larger than $\sqrt{n}/\pi\Delta^2 r$.

Figure 4 sketches the shape of the throughput as a function of r. For small average node degrees, no path is found, because the network is sub-critical. Above the percolation threshold, the number of vertex-disjoint paths (and thus the throughput) slowly increases. This corresponds to the noise-limited regime. For high node degrees, many paths are found, but they interfere with each other, and the overall throughput decreases like $1/r$, for the reason mentioned above. This corresponds to the interference-limited regime. Therefore, there is an optimal value for the range r.

Let us now move to the physical model of Section III, with the random TDMA scheme with t time slots. In this model, the SNIR is at least β on each link. Thus, the throughput on each path is guaranteed, and it is enough to count the number of vertex-disjoint path to estimate the throughput. However, if we use TDMA, we have to divide the result by the number of time slots t. Figure 5 shows the throughput obtained for several values of t, with the other parameters of the model being such that the network is sub-critical when no TDMA is used ($t = 1$). This is the interference-limited case. When we increase t, interferences decrease, making more links available, and the number of paths increases. But when t becomes very large, the interferences are so low that the noise N_0 becomes dominant, and the connectivity graph saturates (no more paths are found). At this point, the throughput decreases like $1/t$.

Similarly to Section III, an alternative strategy would be to reduce the rate on the links (reduce β). As the required SNIR would be lower, one would obtain more links, but with lower rate. This strategy leads to a similar trade-off as the one of the TDMA scheme [1]. If the threshold β is too large, the network does not percolate, and no path is found. When the β is sufficiently small so that $l^{-1}(\beta N_0/P) < 1$, (1) yields that the actual range of the nodes is bounded by

$$||x_i - x_j|| \leq \left(\frac{\beta N_0}{P}\right)^{-\frac{1}{\alpha}} := r. \qquad (2)$$

The connectivity graph in this model is thus a subgraph of the connectivity graph in the Boolean model with the range r defined by (2). Therefore, we can bound the number of paths by counting the number of paths in the Boolean model.

We consider a slice of width r that cuts the network from top to bottom. Clearly, in the Boolean model, each vertex-disjoint path must have at least one node in this slice. The number of left-right crossings of the square is thus limited by the number of nodes in the slice, which increases linearly with r. Because of (2), we know that r increases like $\beta^{-1/\alpha}$. Hence the number

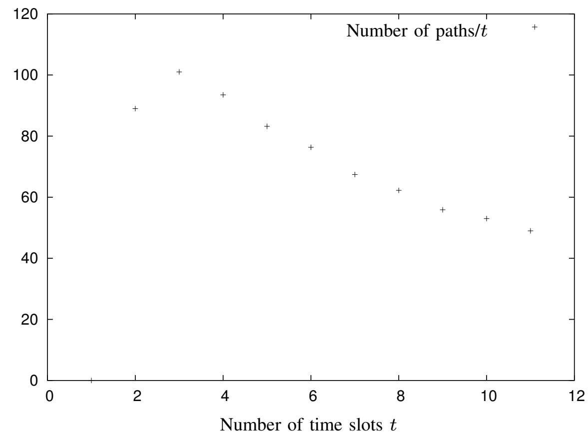

Fig. 5: The horizontal throughput in the signal-to-interference model: the normalized number of left-right crossings with TDMA. In this simulation, the optimal number of time slots is $t = 3$.

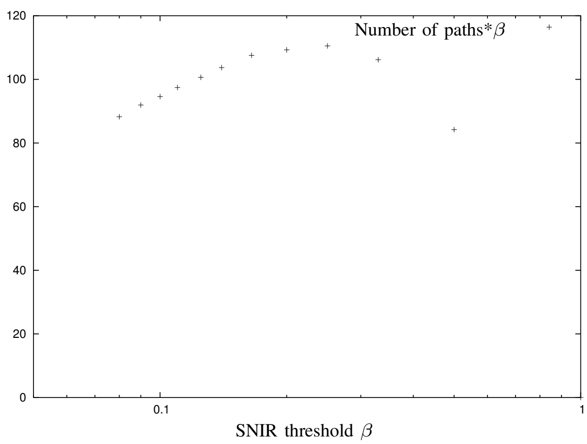

Fig. 6: The horizontal throughput in the signal-to-interference model: the normalized number of left-right crossings vs the SNIR threshold β.

of paths increases like $\beta^{-1/\alpha}$. On the other hand, according to Shannon's formula, the rate of the links is linear in β when β is small. Therefore, the actual horizontal throughput across the network is of the order of $\beta \cdot \beta^{-1/\alpha} = \beta^{1-1/\alpha}$, and tends therefore to zero when $\beta \to 0$. Consquently, there is an intermediate value of β that maximizes the aggregate throughput from left to right, as confirmed by the simulation results in Figure 6.

Conversely, if we keep β fixed, the only way to improve the fraction of connected nodes is to reduce γ, until the network becomes noise-limited, making then any improvement of the processing gain of the codes for interference cancellation useless. The price to pay for reducing γ is again throughput, as one needs longer codewords to make them more orthogonal in a CDMA system. We end up therefore with the same situation as before: by reducing γ, we also reduce the throughput on the links, but we make the network more connected, i.e. add more links. As we have seen earlier, 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 $\gamma = 0$. 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.

VI. Conclusion

Full connectivity does not scale with network size, whether we take the simple Boolean model, or the more sophisticated information theoretic connectivity model. In contrast, if we allow for a (possibly very small) fraction of disconnected nodes, then the range (respectively the rate) does not need to be adjusted when the number of nodes tend to infinity.

When several flows have to share the available bandwidth and 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 optimal throughput.

Acknowledgment

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 under grant number 5005-67322.

References

[1] O. Dousse, M. Franceschetti, and P. Thiran. The costly path from percolation to full connectivity. In *Proc. 42nd Annual Allerton Conference on Communication, Control and Computing*, Monticello, September 2004.

[2] O. Dousse, M. Franceschetti, and P. Thiran. On the throughput scaling of wireless relay networks. *IEEE Transactions on Information Theory (joint Issue ACM/IEEE Transactions on Networking)*, 52(6), June 2006.

[3] M. Franceschetti, O. Dousse, D. Tse, and P. Thiran. Closing the gap in the capacity of random wireless networks. In *Proc. of Information Theory Symposium (ISIT)*, Chicago, Illinois, July 2004.

[4] M. Penrose and A. Pisztora. Large deviations for discrete and continuous percolation. *Adv. Appl. Prob.*, 28:29–52, 1996.

[5] M. Penrose. The longest edge of the random minimal spanning tree. *Ann. Appl. Probability*, 7:340–361, 1997.

[6] 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. McEneany, G. Yin, and Q. Zhang, (Eds.) Birkhuser.

[7] O. Dousse, F. Baccelli, and P. Thiran. Impact of interferences on connectivity of ad hoc networks. *ACM/IEEE Transactions on Networking*, 13(2):425–436, April 2005.

[8] O. Dousse, M. Franceschetti, N. Macris, R. Meester, and P. Thiran. Percolation in the signal to interference ratio graph. *Journal of Applied Probability*, 43(2), June 2006.

[9] X. Liu and R. Srikant. An information theoretic view of connectivity in wireless sensor networks. In *Proc. IEEE Secon*, Santa Clara, CA, October 2004.

[10] O. Dousse, M. Franceschetti, and P. Thiran. Information theoretic bounds on the throughput scaling of wireless relay networks. In *Proc. IEEE Infocom*, Miami, FL, March 2005.

[11] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, March 2000.