Local vs. Global Scalability in Ad Hoc and Sensor Networks

András Faragó

The University of Texas at Dallas,
Richardson, Texas 75080, USA
farago@utdallas.edu

Abstract. We address large, random network topologies that are typical in ad hoc and sensor networks. In these systems we need at least two different types of scalability. First, we want that with growing network size the topology remains connected, so that communication is possible between nodes. Second, it is also necessary that the individual nodes can operate with limited energy and complexity, which requires that the number of neighbors of any node remains bounded. Unfortunately, these global vs. local scalability requirements conflict with each other, as it is known, under very general conditions, that full connectivity can only be achieved with infinitely growing node degrees. Therefore, it is important to quantify how large part of the random topology can be still expected to belong to a connected component if the nodes are confined to some bounded degree. We investigate this issue in a model that is more general than previously investigated random wireless network topology models. In our general model we derive an asymptotically optimal trade-off between node degrees and the fraction of nodes that form a connected component.

Keywords: Scalability, network topology, connectivity, node degrees.

1 Introduction

Large wireless networks with random topology and no supporting infrastructure, such as ad hoc and sensor networks, are expected to be an important part of the future communications landscape. The network topology of these systems is often modeled by various random graph models, most typically by geometric random graphs. (For general background on such models see, e.g., the books [5], [10].)

Due to the randomness of the network topology, it is not at all guaranteed that any two nodes can send messages to each other, since the random graph that represents the network topology may not be connected. If we want to ensure that all nodes can reach each other, then, as a minimum requirement, we have to make sure at least that the network topology (which is usually represented by an undirected graph) is connected.

The connectivity requirement, however, is not as innocent as it may look, due to random node positions and limited wireless transmission ranges. It turns out (see,

P. Mueller, J.-N. Cao, and C.-L. Wang (Eds.): Infoscale 2009, LNICST 18, pp. 26–45, 2009.

e.g., Gupta and Kumar [6,7]) that in typical cases, such as placing the nodes in a planar disk independently and uniformly at random, the price of connectivity is very high: the transmission range needs to be set such that it asymptotically results in an infinitely growing number of neighbors.

This phenomenon is a serious threat to scalability in these networks. One might hope at this point that for different geometric random graph models the situation may perhaps improve. For example, one may try different deployment domains, different probability distributions, different distance metrics, etc. Unfortunately, however, it has been proven in a very general model that none of these can relieve the scalability bottleneck, see Faragó [3]. It appears that unbounded node degrees are unavoidable whenever full connectivity is required in the limit in a random, geometrically induced topology. This is, of course, bad news for scalable implementation, since a node with finite processing power cannot be expected to handle an unbounded neighborhood with bounded delay.

It is therefore of keen importance whether better scalability can be achieved if we are willing to give up full connectivity and substitute it with the milder requirement of *partial connectivity*. This means, as a price for keeping the node degrees bounded, we accept that only most, but not all, nodes are in a connected component. The motivation is that in many potential applications, such as a network of randomly placed sensors, it is acceptable to have only a majority (say, 99%) of nodes in a connected component and the rest are possibly disconnected.

We investigate the fundamental limits related to such partial connectivity, under very general modeling assumptions. Building on the work we have started in [3,4] on connectivity issues, now we explore the asymptotically optimal trade-off between the fraction of nodes that can be kept in a connected component as a function of the bound on the expected node degrees.

We look for the possibly most general conditions under which we can still prove such a trade-off. It turns out that aiming at generality does pay off: we are able to prove that certain very mild conditions on the otherwise unrestricted model already suffice for the proof. The level of generality also makes the proofs much more transparent than the usual stochastic geometry based analysis of random geometric graphs or percolation models, and allows to easily answering questions that would otherwise be quite hard. To illustrate it, a motivating example is presented in the next section.

2 A Motivating Sample Problem

In this example we model a mobile wireless ad hoc network. The initial position of each node is chosen in the following way. Let P be a probability measure over a planar domain D. First we choose k pivot points independently at random, using P. Then the actual node positions are generated such that each potential node is chosen independently at random from P, but it is kept only if it is within a given distance d_0 to at least one of the random pivot points, otherwise it is discarded. Note that this way of generating the nodes makes them dependent, as the non-discarded ones cluster around the random pivot points, thus modeling a clustered, non-independent node distribution.

The mobility of the nodes in this example is modeled in the following way. Over some time horizon T_n, that may depend on n, the number of nodes, each node moves along a random curve from its initial position with a constant speed v_0. The curve is chosen from a set $\mathcal{C}$ of available potential trajectories in D. For simplicity, it is assumed that each curve can be identified by a real parameter. This parameter is chosen using a probability distribution $Q_{x,y}$ that depends on the initial position (x, y) of the node. Then the randomly obtained curve is shifted so that its start point coincides with the random initial position of the node and then the node will move along this random trajectory. It is assumed that $\mathcal{C}$ and D are such that the shifted curves still remain in the domain.

Let $d(x, y)$ be a nonnegative real valued function over $D \times D$, with the only restriction that $d(x, x) = 0$ holds for any x. This function is intended to measure "radio distance" in D. The assumption is that whenever $d(x, y)$ is small enough, then two nodes positioned at x and y can receive each others' transmissions. The function $d(x, y)$, however, does not have to satisfy the usual distance axioms, it may reflect complex radio propagation characteristics, such as expected attenuation and fading, it may account for the heterogeneity of the terrain, for propagation obstacles etc. We may also include random effects, making $d(x, y)$ a random variable, reflecting special conditions of interest, such as the random presence of eavesdroppers that can trigger the inhibition of certain links. We assume, however, that if there is randomness in $d(x, y)$, then it is independent of the other random variables in the model.

We now define the links of the network, as follows. Consider two nodes with initial position vectors $X_1(0), X_2(0)$, respectively. As they move along their random trajectories, their positions at time t is denoted by $X_1(t), X_2(t)$, respectively. The two nodes are considered connected by a link, if there is a closed subinterval of length at least t_n within the time horizon $[0, T_n]$, such that $d(X_1(t), X_2(t)) \leq r_n$ holds for every time t within the subinterval, with the possibly complicated radio distance. Here t_n and r_n are parameters that may also depend on the number n of nodes. The motivation for this link definition is that the nodes should be within range at least for the time of sending a packet.

Now the question is this: for given P, D, $\mathcal{C}$, $Q_{x,y}$ and $d(x, y)$, and for the described way of dependent node generation, can we somehow choose the model parameters, such that the expected node degrees remain bounded by, say, 4 and still at least 99% of nodes belong to a connected component?

We believe that it would be rather hard to answer this question with a direct analysis for arbitrary complex choices of P, D, $\mathcal{C}$ $Q_{x,y}$ and $d(x, y)$. On the other hand, in view of our general results that we build up in the subsequent sections, it becomes quite straightforward.

3 On a Bottleneck of Scalability

Before presenting our results, let us briefly analyze how the considered issue imposes a *general* scalability bottleneck for large wireless networks.

It is well known that network connectivity is of primary concern in large, random ad hoc networks, since a connected network topology is a fundamental condition

for successful operation. Another, less obvious, important reason for taking a closer look at connectivity is its role in the limiting behavior of the *transport capacity* of ad hoc networks.

Reseach interest in the fundamental limits on the transport capacity of ad hoc networks had an upsurge after the the seminal paper of Gupta and Kumar [6] on the capacity of wireless networks. Many results followed on how much traffic can an ad hoc network support in different modes of operation, which values of local parameters can guarantee global connectivity, what can topology control principally achieve etc.

One of the key results of Gupta and Kumar [6] considers the achievable throughput per source-destination (S-D) pair in a large network, where n nodes are placed independently and uniformly at random in a planar disk of unit area. The nodes have the same (positive) transmission radius r which, in principle, can be chosen arbitrarily. The nodes can transmit with a fixed maximum rate (bit/sec), but they are allowed to divide the channel into subchannels without any constraint, in any domain (e.g., frequency, time, code). The nodes communicate using arbitrary protocols for channel access and routing, with the only restriction that there is a minimal requirement of interference avoidance in the same (sub)channel to ensure successful receptions. The considered traffic pattern is that each node has a randomly chosen destination in the network. The authors prove that in this general model the achieveable throughput per S-D pair is $\Theta\left(\frac{1}{\sqrt{n \log n}}\right)$.

Regarding scalability, the key message of this result is that the achievable throughput per S-D pair tends to zero as the network size grows to infinity. We call this the *vanishing throughput effect*. It means, the network is fundamentally not scalable, since it becomes unable to usefully operate when it grows very large.

To better understand a fundamental reason for the vanishing throughput and its key relationship to asymptotic connectivity, let us briefly discuss what causes it and how in the Gupta-Kumar model [6].

The transmission radius r of nodes has two opposite effects. If r is small, then more hops are needed to deliver a packet to its destination, since in each hop the packet can advance at most a distance of r towards its destination. Therefore, small r causes a growing burden of nodes to serve as relays, which decreases the end-to-end throughput. This would justify choosing r as large as possible. On the other hand, large transmission radius increases interference which causes the throughput to decrease again, so interfence reduction would require to choose r as small as possible.

The detailed analysis of the above conflicting tendencies in [6] shows that the forwarding burden is proportional to $1/r$, since the average route hop-length is proportional to $1/r$. On the other hand, the interference increases quadratically with the radius, due to the area involved, which is proportional to r^2. Balancing the two effects yields the conclusion that r has to be chosen as small as possible, as the interference reduction, due to its quadratic nature, brings more benefit than what is lost by longer routes. Specifically, the joint effect results in the formula

$$\lambda(n) = O\left(\frac{1}{nr(n)}\right) \tag{1}$$

where n is the number of nodes, $\lambda(n)$ is the throughput in bit/sec for each S-D pair and $r(n)$ is the transmission radius (as a function of n). It is clear from (1) that if we want to avoid $\lambda(n) \to 0$, then $r(n)$ has to tend to 0 at least as fast as $O(1/n)$, so that $nr(n)$ does not grow to infinity.

There is, however, an effect that does not allow choosing the transmission radius arbitrarily small. This is the requirement that the network topology must be *connected*, since otherwise communication between certain endpoints becomes impossible. The discussed model uses an earlier result of the same authors [7] about the needed transmission radius for asymptotic connectivity. They show that if the nodes are placed uniformly at random in a unit disk, then the network is connected with probability approaching 1 if and only if the transmission radius satisfies

$$\pi r^2 = \frac{\ln n + c(n)}{n} \tag{2}$$

with $c(n) \to \infty$. According to (2), it is necessary for connectivity that

$$r > \sqrt{\frac{\ln n}{\pi n}}$$

holds for all large enough values of n. Combining it with (1), we obtain

$$\lambda(n) = O\left(\frac{1}{n\sqrt{\frac{\ln n}{n}}}\right) = O\left(\frac{1}{\sqrt{n \ln n}}\right)$$

clearly showing the vanishing throughput effect.

According to the above discussion, the lower bound on the transmission radius, enforced by the network connectivity requirement, can be viewed as an important factor in the vanishing throughput effect, since without it the transmission radius could be chosen small enough to gain arbitrarily more from reduced interference than what is lost by longer hop-distances.

To obtain a graph theoretical view of the network topology, we can translate (2) into the expected node degrees. If the nodes are placed uniformly at random in a unit disk, then πr^2 is the expected number of nodes that fall in the range of a node that is not at the border of the disk, so that its range is fully in the domain. This is also the probability that a random node falls in this range, so the expected number of neighbors is $n\pi r^2$ and by (2) we have $n\pi r^2 > \ln n$. Since only a vanishing fraction of nodes are close to the border, we can asymptotically ignore the border effect and reformulate (2) in the following way: for the Gupta-Kumar model, *connectivity requires that the node degrees grow to infinity at least logarithmically with the network size.*

It is worth mentioning that infinitely growing node degrees alone, even without the capacity analysis, can kill scalability, since the processing requiments of a node are likely to grow at least proportionally with the size of the neighborhood.

Thus, we can conclude that an important factor in the vanishing throughput phenomenon is that the connectivity requirement does not allow the transmission

radius to shrink too fast, which, if translated to a graph view, means that the node degrees must tend to infinity, whenever *full* connectivity is required. This naturally leads to the idea of relaxing the full connectivity requirement and it makes important to investigate the trade-off between the fraction of nodes that can be kept in a connected component versus the expected node degrees.

4 Random Graph Models

Let us first explain what we mean by random graphs and a random graph model in the most general sense.

In full generality, by a *random graph* on a fixed number of nodes (n) we mean a random variable that takes its values in the set of all undirected graphs on n nodes. We use the notation G_n for a random graph on n nodes. At this point, it is still completely general, it can be generated by any mechanism, with arbitrary dependencies among its parts, it is just *any* graph-valued random variable, taking its values among undirected graphs on n nodes.

A *random graph model* is given by a sequence of graph valued random variables, one for each possible value of n:

$$\mathcal{M} = (G_n;\ n \in \mathbf{N}).$$

Let us now consider some important general parameters and properties. Let G_n be any random graph on n nodes and denote by $e(G_n)$ the number of edges in the graph. We characterize the degrees of G_n by the expected degree of a randomly chosen vertex, which we call he *expected average degree* of G_n. It is denoted by $\overline{d}(n)$ and defined by

$$\overline{d}(n) = \frac{2\mathrm{E}(e(G_n))}{n}$$

based on the fact that the actual average degree in any graph G on n nodes is $2e(G)/n$. Often the expected degree of each individual node is also equal to $\overline{d}(n)$, but in a general model it may not hold. (Note that even if the expected degree of each node is equal to the expected average degree, it does not mean that the *actual* random degrees are also equal.)

Ideally, we would like a random graph model in which $\overline{d}(n)$ remains constant and the model is *asymptotically almost surely (a.a.s.)* connected, meaning

$$\lim_{n\to\infty} \Pr(G_n \text{ is connected}) = 1.$$

Note: Whenever we write down a limit, such as the one above, we also assume that the limit exists.

Since, as mentioned in the Introduction, asymptotic connectivity is not possible in most models without unbounded degrees, therefore, one may hope that if less than full connectivity is required, then there is a better chance to keep the node degrees bounded. To this end, let us define a weaker version of connectivity.

Definition 1. (β-connectivity) *For a real number $0 \le \beta \le 1$, a graph G on n nodes is called β-connected if G contains a connected component on at least βn nodes.*

When we consider a sequence of graphs with different values of n, then the parameter β may depend on n. When this is the case, we write β_n-connectivity. Note that even if $\beta_n \to 1$, this is still weaker than full connectivity in the limit. For example, if $\beta_n = 1 - 1/\sqrt{n}$, then we have $\beta_n \to 1$, but each graph on n nodes can still have $n - \beta_n n = \sqrt{n}$ nodes that are not part of the largest connected component.

Let us now introduce a model class that reflects a typical feature of geometric random graph models. This feature is that in geometric random graphs the primary random choice is picking random nodes from some domain and then the edges are already determined by some geometric property (typically some kind of distance) of the random nodes. We elevate this approach to an abstract level that includes many special cases of interest. Based on this high level of abstraction, we call it *abstract geometric random graph model.*

The most general version of our abstract geometric model is built using the following components:

- **Node variables.** The nodes are represented by an infinite sequence $X_1, X_2, \ldots$ of random variables, called *node variables.* They take their values in an arbitrary (nonempty) set S, which is called the *domain* of the model. In most practical cases the domain is a simple subset of the Euclidean plane or of the 3-dimensional space. In general, however, S can be any abstract set from which we can choose random elements[1]. When we want to generate a random graph on n nodes, then we use the first n entries of the sequence, that is, $X_1, \ldots, X_n$ represent the nodes in G_n. It is important to note that we do not require the node variables to be independent.
- **Edge functions.** We denote by $Y_{ij}^{(n)} \in \{0, 1\}$ the indicator of the edge between nodes X_i, X_j in the random graph G_n. Since loops are not allowed, we always assume $i \neq j$, without repeating this condition each time. The (abstract) geometric nature of the model is expressed by the requirement that the random variables $Y_{ij}^{(n)}$ are determined by the nodes $X_1, \ldots, X_n$, possibly with additional independent randomization. Specifically, we assume that there exist functions $f_{ij}^{(n)}$, such that

$$Y_{ij}^{(n)} = f_{ij}^{(n)}(X_1, \ldots, X_n, \xi_{ij})$$

  where ξ_{ij} is a random variable that is uniformly distributed on $[0, 1]$ and is independent of all the other defining random variables of the model (i.e., the node variables and all the other ξ_{kl} variables). Henceforth the role of ξ_{ij} is referred to as *independent randomization*[2]. The undirected nature of the graph is expressed by the requirement $Y_{ij}^{(n)} = Y_{ji}^{(n)}$, which can simply be enforced by

[1] To avoid mathematical complications that would only obscure the main message, we assume that all considered sets, functions etc. are measurable with respect to the used probability measures and all considered expected values exist. This is satisfied in in every practically relevant model.

[2] Note that the specified distribution of ξ_{ij} does not impose a restriction, since the functions $f_{ij}^{(n)}$ are arbitrary.

computing all values for $i < j$ only and defining the $i > j$ case by exchanging i and j.

Regarding the abstract geometric random graph model in the presented very general form, it is clear that allowing *totally arbitrary* node variables and edge functions offers little hope for meaningful analysis. Therefore, next we introduce some restricting conditions. Later we are going to see that one has to make only surprisingly mild restrictions to meaningfully analyze the trade-off between node degrees and β-connectivity.

Up to now we allowed that an edge in G_n can depend on all the nodes, and the dependence expressed by the $f_{ij}^{(n)}$ functions can be arbitrary and different for each edge. To get a little closer to the usual geometric random graph model, let us introduce the following property, called locality. Informally, it restricts the dependence of an edge to its endpoints, in a homogeneous way, but still via an *arbitrary* function.

Definition 2. (Locality) *An abstract geometric random graph model is called* local*, if for every n and $i, j \leq n$ the existence of an edge between X_i, X_j depends only on these nodes. Moreover, the dependence is the same for every i, j, possibly with independent randomization. That is, there are functions $f^{(n)}$ such that the edge indicators are expressible as*

$$Y_{ij}^{(n)} = f^{(n)}(X_i, X_j, \xi_{ij})$$

where ξ_{ij} represents the independent randomization.

Our second condition called *name invariance* refers to the joint distribution of nodes. If we allow totally arbitrary joint distribution, then it offers little chance for meaningful analysis. On the other hand, restricting ourselves only to independent, identically distributed (i.i.d.) node variables would exclude important cases, such as clustering. Therefore, we introduce a condition that allows more general than i.i.d. node variables, but still makes meaningful analysis possible. To introduce it, let us first recall a useful concept from probability theory, called exchangeability.

Definition 3. (Exchangeable random variables) *A finite sequence $\xi_1, \ldots, \xi_n$ of random variables is called* exchangeable *if for any permutation σ of $\{1, \ldots, n\}$, the joint distribution of $\xi_1, \ldots, \xi_n$ is the same as the joint distribution of $\xi_{\sigma(1)}, \ldots, \xi_{\sigma(n)}$. An infinite sequence of random variables is called exchangeable if every finite initial segment of the sequence is exchangeable.*

Exchangeability can be equivalently defined such that for any $k \geq 1$ among the random variables, say, $\xi_{j_1}, \ldots, \xi_{j_k}$, their joint distribution is always the same (for a given k), it does not depend on which particular set of k indices is selected. Note that i.i.d. random variables are always exchangeable, but the converse generally does not hold, so this is a larger family.

Now let us introduce the condition that we use to restrict the arbitrary dependence of node variables.

Definition 4. (Name invariance) *An abstract geometric random graph model is called* name invariant*, if its node variables are exchangeable.*

We call it the *name invariance* of the model because it means the names (the indices) of the nodes are irrelevant in the sense that the joint probabilistic behavior of any fixed number of nodes is invariant to renaming (reindexing) the nodes. In particular, it also implies that the node variable are identically distributed, but they do not have to be independent.

Name invariance is naturally satisfied with the most frequently used random node choices, such as uniform independent random points in a planar domain, or a Poisson point process in the plane, or in higher dimension. We allow, however, much more complex node generation (over an arbitrary set!) since dependencies are not excluded by name invariance.

A simple example for a dependent, yet still name invariant, node generation process is a "clustered uniform" node generation. As an example, let S be a sphere in 3-dimensional space, i.e., the surface of a 3-dimensional ball. Let R be the radius of the ball. Let us first generate a pivot point Y uniformly at random from S. Then generate the nodes $X_1, X_2, \ldots$ uniformly at random and independently of each other from the neighborhood of radius $r \ll R$ of the random pivot point Y (within the sphere). It is directly implied by the construction that exchangeability holds. Moreover, any particular X_i will be uniformly distributed over the *entire* sphere, since Y is uniform over the sphere. On the other hand, the X_i are far from independent of each other, since they cluster around Y, forcing any two of them to be within distance $2r$. The setting can be generalized to applying several pivot points and non-uniform distributions, creating a more sophisticated clustering.

5 Example Models

Before turning to the results, let us present some example models to show the usefulness and comprehensiveness of the generalization provided by our abstract geometric random graphs. Since the results will apply to local and name invariant models, we restrict ourselves to such models in the examples.

Geometric random graphs. All the usual geometric random graph models fit naturally in our general framework. For example, the base set S can be chosen as a unit disk or square in the plane or a unit ball or cube (or any other domain) in higher dimension. Let us choose i.i.d. points $X_1, X_2, \ldots$ from S, according to some probability distribution. Let $\rho(x, y)$ denote the distance of the points $x, y \in S$, it can be any distance function. Finally, let $r > 0$ be a radius (possibly depending on n). Then the edge function

$$f^{(n)}(X_i, X_j, \xi_{ij}) = \begin{cases} 1 & \text{if } \rho(X_i, X_j) \leq r \\ 0 & \text{if } \rho(X_i, X_j) > r \end{cases} \tag{3}$$

defines a geometric random graph in the usual sense. (The independent randomization is not used here, so the edge function does not depend on ξ_{ij}.) It is clear that

this includes all the usual geometric random graph models, allowing any metric space as the basis. Moreover, we can also use non-independent points, such as the "clustered uniform" example in the previous section, as long as the distribution is exchangeable.

Erdős-Rényi random graphs. The by now classical random graph model of Erdős and Rényi (see, e.g., [1,8]), where each possible edge is included independently with some probability p is also included as a direct special case. We can set $S = \{1, \ldots, n\}$ and for $X_i, X_j \in S$

$$f^{(n)}(X_i, X_j, \xi_{ij}) = \begin{cases} 1 & \text{if } \xi_{ij} \leq p \\ 0 & \text{if } \xi_{ij} > p \end{cases}$$

Note that now the edge function depends only on the independent randomization, so indeed each edge is included independently with probability p.

A geometric but non-metric example: battery levels. In the geometric random graph models ρ satisfies the triangle inequality. This, however, cannot capture all situations that occur in ad hoc or sensor networks. As an example, assume the nodes are located in the plane. Let x_i, y_i be the coordinates of the i^{th} node. Furthermore, we also characterize a node with its battery level $E_i > 0$. E_i represents the remaining energy, assuming the node is not fully out of energy. Thus, a node is represented by a triple $X_i = (x_i, y_i, E_i)$. Let $d(E_i)$ be the distance over which a node can communicate, given its energy level E_i. (The function $d(E_i)$ can be derived from the physical characteristics of the node and from radio propagation conditions.) Now, a possible example of a "distance" function is

$$\rho(X_i, X_j) = \frac{\sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}}{\min\{d(E_i), d(E_j)\}}$$

If we take $r = 1$ and use the above ρ function in (3), then it expresses the condition that a link exists if and only if its end nodes are at most at a distance that can be bridged by the energy levels of both nodes. Note that the above function ρ does not satisfy the triangle inequality, so it does not lead to a geometric random graph model in the usual sense. On the other hand, it still fits in our framework, as in (3) we did not require the triangle inequality to hold for ρ.

Another non-metric example: link blocking. We can capture some features of traffic dependent network characteristics, as well. Let each node i be characterized by a triple $X_i = (x_i, y_i, \lambda_i)$, where x_i, y_i are planar coordinates and λ_i is the traffic demand of the node. Let B_{ij} be the blocking probability of the link (i, j), given that the link exists. We may compute B_{ij} as a function of λ_i, λ_j from some traffic model. For example, if we use Erlang's well known formula, assuming a capacity of C units on the link and its load is taken as the sum of its end nodes' traffic load $\lambda_i + \lambda_j$, then we obtain

$$B_{ij} = \frac{(\lambda_i + \lambda_j)^C / C!}{\sum_{i=0}^{C} (\lambda_i + \lambda_j)^i / i!}.$$

(Of course, we may use other traffic models, as well, this is just an example.) Now we can take the "distance" function

$$\rho(X_i, X_j) = \frac{1}{1 - B_{ij}} \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}$$

and use it in (3) with some radius r. We can observe that for small blocking probability ($B_{ij} \ll 1$) $\rho(X_i, X_j)$ will be approximately the same as the Euclidean distance. On the other hand, as B_{ij} approaches 1, the factor $\frac{1}{1-B_{ij}}$ tends to infinity and, therefore, high blocking probability makes the existence of the link in the model less likely, even if the physical distance is small. This example also violates the triangle inequality, so it is not a geometric random graph.

Log-normal shadowing. A typical phenomenon in the radio environment is *fading*. An example of fading is a relatively slow random fluctuation in the signal strength, which occurs even if the locations are fixed. Measurements show that this random variation can be accurately modeled by a log-normal distribution (see, e.g., [9]). Hence the name *log-normal shadowing*, which is widely used for this phenomenon. A way to capture it in our model is this. Let us characterize a node i by a triple $X_i = (x_i, y_i, \eta_i)$, where x_i, y_i represent a random position in the plane and each η_i is an infinite sequence of independent, log-normally distributed random variables:

$$\eta_i = (\eta_j^{(i)};\ j = i, i+1, i+2, \ldots).$$

The "distance" is defined as

$$\rho(X_i, X_j) = \eta_b^{(a)} \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}$$

where $a = \min\{i, j\}$ and $b = \max\{i, j\}$. (The reason for we need an infinite sequence of log-normal random variables is that this way we can have independent log-normal shadowing for every link.) This distance can express the fact that from the radio communication point of view we really perceive an "effective distance", which is a log-normally modulated random variant of the physical distance. Using this ρ in (3) leads again to a random graph that is not geometric, as ρ does not satisfy the distance axioms.

6 Threshold Function for Partial Connectivity

Let us now define a concept that will characterize the trade-off between node degrees and the type of partial connectivity that we introduced as β-connectivity in Definition 1. The set of nonnegative real numbers, extended with ∞, will be denoted by $\mathbf{R}_0^\infty$. Real functions are also extended to ∞ by $f(\infty) = \lim_{x\to\infty} f(x)$, whenever the limit exists (it will always exist in our cases). The value of β is always assumed to be in $[0, 1]$.

Before the formal definition let us explain the concept informally. We define a threshold function for β-connectivity, such that whenever β is above the threshold,

then it is impossible to achieve a.a.s. β-connectivity for any model in the considered family of random graph models. On the other hand, if β is below the threshold, then this is not the case anymore, that is, there is at least one model in the family that achieves a.a.s β-connectivity with this β. Now let us present the formal definition. Recall that the expected average degree in a random graph G_n is defined as $\overline{d}(n) = 2\mathrm{E}(e(G_n))/n$.

Definition 5. (Threshold for β-connectivity) *Let $\mathcal{F}$ be a family of random graph models. For any model $\mathcal{M} \in \mathcal{F}$ let G_n denote the random graph on n nodes generated by $\mathcal{M}$ and set*

$$D_{\mathcal{M}} = \limsup_{n\to\infty} \overline{d}(n).$$

A function $f : \mathbf{R}_0^\infty \mapsto [0,1]$ is called a β-connectivity threshold function for $\mathcal{F}$ if the following two conditions are satisfied:

(i) For any model $\mathcal{M} \in \mathcal{F}$ and for every $\beta > f(D_{\mathcal{M}})$

$$\lim_{n\to\infty} \Pr(G_n \textit{is } \beta\textit{-connected}) < 1$$

holds, where G_n is generated by $\mathcal{M}$.

(ii) If β is below the threshold, then (i) does not hold anymore, in the following sense. For every $\epsilon > 0$ there exists a model $\mathcal{M}_0 \in \mathcal{F}$ and a

$$\beta \leq f(D_{\mathcal{M}_0}) - \epsilon$$

such that

$$\lim_{n\to\infty} \Pr(G_n \textit{is } \beta\textit{-connected}) = 1$$

where G_n is generated from $\mathcal{M}_0$.

The importance of this concept is the following. If for a considered class $\mathcal{F}$ of random graph models we can find out what the corresponding β-connectivity threshold function is, then we can tell precisely what range of expected average degrees allow a.a.s. β-connectivity for a given β. Or, conversely, if we know the (asymptotic) expected average degree for a particular model $\mathcal{M}$ in the considered class, then we can decide what level of connectivity can be asymptotically achieved for this model.

7 Results

Now we show that for the quite general class of abstract geometric random graph models we can find the precise β-connectivity threshold function, if we assume that the models satisfy the conditions of locality and name invariance. The previously presented examples all satisfy these conditions, so they show that even with these restrictions we can still include many complex and practically important models.

Theorem 1. (Threshold function for local and name invariant abstract geometric graphs) *Let $\mathcal{F}$ be the family of local and name invariant abstract geometric random graph models For any model $\mathcal{M} \in \mathcal{F}$ set*

$$D_{\mathcal{M}} = \limsup_{n\to\infty} \overline{d}(n).$$

Then the β-connectivity threshold function for $\mathcal{F}$ is

$$f(D_{\mathcal{M}}) = 1 - \mathrm{e}^{-D_{\mathcal{M}}}.$$

The proof is based on another theorem, which is interesting on its own right. As a further notation, the (random) number of isolated nodes in G_n is denoted by I_n.

Theorem 2. (Lower bound on the expected number of isolated nodes) *The expected number of isolated nodes in a local and name invariant abstract geometric random graph G_n always satisfies*

$$\mathrm{E}(I_n) \geq n\left(1 - \frac{\overline{d}(n)}{n-1}\right)^{n-1}. \tag{4}$$

Proof of Theorem 2. First we note that since our model is abstract and does not involve any real geometry, one has to be careful to avoid using such intuition that may appeal geometrically, but does not follow from the abstract model.

As a first step, observe the following: name invariance implies that for any function g of the node variables and for any permutation σ of $\{1,\ldots,n\}$ we have

$$\mathrm{E}(g(X_1,\ldots,X_n)) = \mathrm{E}(g(X_{\sigma(1)},\ldots,X_{\sigma(n)})).$$

Since the probability that a particular node has any given degree k is also expressible by such a function, therefore, the probability distribution of the node degree must be the same for all nodes (but the degrees, as random variables, may not be independent). As a consequence, the expected degree of each node is the same, which then must be equal to the expected average degree $\overline{d}(n)$.

Let us pick a node X_i. We derive a lower bound on the probability that X_i is isolated, i.e., its degree is 0. Due to the above symmetry considerations, it does not matter which node is chosen, so we can take $i = 1$. Let $\mathcal{I}_n$ be the (random) set of isolated nodes in G_n. What we want to compute is a lower bound on $\Pr(X_1 \in \mathcal{I}_n)$. Then we are going to use the fact that

$$\mathrm{E}(I_n) = \mathrm{E}(|\mathcal{I}_n|) = \sum_{i=1}^{n} \Pr(X_i \in \mathcal{I}_n)$$

Note that, due to the linearity of expectation, this remains true even if the events $\{X_i \in \mathcal{I}_n\}$ are not independent, which is typically the case. Then, by the symmetry considerations, we can utilize that $\Pr(X_i \in \mathcal{I}_n)$ is independent of i, yielding $\mathrm{E}(I_n) = n\Pr(X_1 \in \mathcal{I}_n)$.

In order to derive a lower bound on $\Pr(X_1 \in \mathcal{I}_n)$, we need a fundamental result from probability theory, called *de Finetti's Theorem*[3]. This theorem says (in its

[3] It was first published in [2]. Being a classical result, it can be found in many advanced textbooks on probability. Interestingly, it seems that, despite its usefulness, it is rarely applied by the networking community.

simplest form that is already sufficient for our purposes) that if an infinite sequence $\xi_1, \xi_2, \ldots$ of 0-1 valued random variables is exchangeable, then the following hold:

(i) The limit

$$\eta = \lim_{N\to\infty} \frac{\xi_1 + \ldots + \xi_N}{N} \tag{5}$$

exists with probability 1. Note that exchangeability implies that all ξ_i have the same expected value, so in case they were independent, then the strong law of large numbers would apply and the limit would be the common expected value, with probability 1. Since, however, the ξ_i are not assumed independent (only exchangeable), therefore, the average may not tend to a constant, it can be a non-constant random variable in $[0, 1]$.

(ii) For any N and for any system $a_1, \ldots, a_N \in \{0, 1\}$ of outcomes with $s = \sum_{i=1}^{N} a_i$

$$\Pr(\xi_1 = a_1, \ldots, \xi_N = a_N) = \int_0^1 x^s (1-x)^{N-s} dF_\eta(x)$$

holds, where F_η is the probability distribution function of η.

(iii) The ξ_i are conditionally independent and identically distributed (conditionally i.i.d.), given η, that is,

$$\Pr(\xi_1 = a_1, \ldots, \xi_N = a_n \,|\, \eta) = \prod_{i=1}^{N} \Pr(\xi_i = a_i \,|\, \eta).$$

Informally, de Finetti's theorem says that exchangeable 0-1 valued random variables, even if they are not independent, can always be represented as a mixture of Bernoulli systems of random variables. It is important to note, however, that even though the statements (ii) and (iii) refer to finite initial segments of the sequence $\xi_1, \xi_2, \ldots$, it is necessary that the entire *infinite* sequence is exchangeable. For finite sequences the theorem may not hold, counterexamples are known for the finite case [11].

Let us now define the infinite sequence of 0-1 valued random variables

$$e_j = f^{(n)}(X_1, X_j, \xi_{1j}), \qquad j = 2, 3 \ldots$$

Of these, $e_2, \ldots, e_n$ are the indicators of the edges with one endpoint at X_1. But the function $f^{(n)}$ is defined for any $(x, y, z) \in S \times S \times [0, 1]$, so nothing prevents us to define the *infinite* sequence e_j; $j = 2, 3, \ldots$, by taking more independent and uniform $\xi_{1j} \in [0, 1]$ random variables.

Observe now that the sequence e_j; $j = 2, 3, \ldots$ is an infinite exchangeable sequence of 0-1 valued random variables. Only the exchangeability needs proof. If we take any k indices $j_1, \ldots, j_k$, then the joint distribution of $e_{j_1}, \ldots, e_{j_k}$ depends only the joint distribution of $X_{j_1}, \ldots, X_{j_k}$, plus the independent randomization. If we replace $j_1, \ldots, j_k$ by other k indices, then it will not change the joint distribution of the k node variables, due to their assumed exchangeability. The independent randomization also does not change the joint distribution, since the ξ_{1j}

are i.i.d., so it does not matter which k are taken. Furthermore, the locality of the model implies that each e_j depends on one X_j (besides X_1) so taking another k cannot change how many node variables will any subset of the e_j share. Thus, for any k, the joint distribution of $e_{j_1}, \ldots, e_{j_k}$ does not depend on which k indices are chosen, proving that e_j; $j = 2, 3, \ldots$ is an infinite exchangeable sequence of 0-1 valued random variables.

Now, by de Finetti's Theorem, there is a random variable $\eta \in [0, 1]$, such that the e_j are conditionally i.i.d., given η. Then we can write

$$\begin{aligned}
\Pr(X_1 \in \mathcal{I}_n) &= \Pr(e_2 = \ldots = e_n = 0) \\
&= \mathrm{E}(\Pr(e_2 = \ldots = e_n = 0 \,|\, \eta)) \\
&= \mathrm{E}\left(\prod_{j=2}^{n} (\Pr(e_j = 0 \,|\, \eta))\right) \\
&= \mathrm{E}\left(\prod_{j=2}^{n} (1 - \Pr(e_j = 1 \,|\, \eta))\right). \qquad (6)
\end{aligned}$$

Notice that $\Pr(e_j = 1 \,|\, \eta)$ is the probability that an edge exists between X_1 and X_j, conditioned on η. Consequently, $\xi = \Pr(e_j = 1 \,|\, \eta)$ is a random variable, depending on η. At the same time, it does not depend on j, as by de Finetti's theorem, the e_j are conditionally i.i.d., given η, so it does not matter which j is taken in $\xi = \Pr(e_j = 1 \,|\, \eta)$. Thus, we can continue (6) as

$$\Pr(X_1 \in \mathcal{I}_n) = \mathrm{E}\left(\prod_{j=2}^{n} (1 - \xi)\right) = \mathrm{E}\left((1 - \xi)^{n-1}\right). \qquad (7)$$

We can now observe that $\xi \in [0, 1]$ and the function $g(x) = (1 - x)^n$ is convex in $[0, 1]$, so we may apply Jensen's inequality. Jensen's well known inequality says that for any random variable ζ and for any convex function g the inequality $\mathrm{E}\big(g(\zeta)\big) \geq g\big(\mathrm{E}(\zeta)\big)$ holds, which is a consequence of the definition of convexity. Thus, we can further continue (7), obtaining

$$\Pr(X_1 \in \mathcal{I}_n) = \mathrm{E}\left((1 - \xi)^{n-1}\right) \geq (1 - \mathrm{E}(\xi))^{n-1}.$$

Note that $\mathrm{E}(\xi) = \mathrm{E}(\Pr(e_j = 1 \,|\, \eta)) = \Pr(e_j = 1)$ is the probability that an edge exists between X_1 and X_j. By name invariance, this is the same probability for any two nodes, let p_n denote this common value. Thus,

$$\Pr(X_1 \in \mathcal{I}_n) \geq (1 - p_n)^{n-1}$$

follows. We know that there are $n - 1$ potential edges adjacent to each node, each with probability p_n. Therefore, despite the possible dependence of edges, the linearity of expectation implies the expected degree of each node under our conditions is $(n - 1)p_n$, which is also equal to $\overline{d}(n)$. We can then substitute $p_n = \overline{d}(n)/(n - 1)$, which yields

$$\Pr(X_1 \in \mathcal{I}_n) \geq \left(1 - \frac{\overline{d}(n)}{n-1}\right)^{n-1},$$

implying

$$\mathrm{E}(I_n) = n \Pr(X_1 \in \mathcal{I}_n) \geq n \left(1 - \frac{\overline{d}(n)}{n-1}\right)^{n-1},$$

completing the proof.

♠

Now, using Theorem 2, we can prove the threshold claimed in Theorem 1.

Proof of Theorem 1. Fix any model $\mathcal{M} \in \mathcal{F}$. Since $D_{\mathcal{M}} = \limsup_{n\to\infty} \overline{d}(n)$ and

$$\left(1 - \frac{D_{\mathcal{M}}}{n-1}\right)^{n-1} \to \mathrm{e}^{-D_{\mathcal{M}}},$$

therefore, there must exist a sequence $a_n \to 1$, such that

$$\left(1 - \frac{\overline{d}(n)}{n-1}\right)^{n-1} \geq a_n \mathrm{e}^{-D_{\mathcal{M}}}.$$

Hence, by Theorem 2,

$$\mathrm{E}(I_n) \geq a_n \mathrm{e}^{-D_{\mathcal{M}}} n \tag{8}$$

holds for every n.

Now fix a β with $1 \geq \beta > 1 - \mathrm{e}^{-D_{\mathcal{M}}}$. (We can assume $D_{\mathcal{M}} < \infty$, since otherwise there is no such β.) We are going to show that $\Pr(G_n$ is β-connected) cannot tend to 1.

Set $s_n = \Pr(I_n \leq (1-\beta)n)$, i.e. s_n is the probability that at most $(1-\beta)n$ nodes are isolated. Then $\Pr(G_n$ is β-connected) $\leq s_n$ must hold, since β-connectivity implies that there may be at most $(1-\beta)n$ isolated nodes. Consider now the random variable $\gamma_n = n - I_n$, which is the number of non-isolated nodes. Then $\gamma_n \geq 0$ and $\mathrm{E}(\gamma_n) = n - \mathrm{E}(I_n)$. Therefore, (8) implies that

$$\mathrm{E}(\gamma_n) \leq (1 - a_n \mathrm{e}^{-D_{\mathcal{M}}})n$$

holds for every n. Furthermore, by the definition of γ_n, the events $\{I_n \leq (1-\beta)n\}$ and $\{\gamma_n \geq \beta n\}$ are identical. Thus, we can write, using the well known Markov inequality for nonnegative random variables:

$$s_n = \Pr(I_n \leq (1-\beta)n) = \Pr(\gamma_n \geq \beta n)$$

$$\leq \frac{\mathrm{E}(\gamma_n)}{\beta n} \leq \frac{(1 - a_n \mathrm{e}^{-D_{\mathcal{M}}})n}{\beta n} = \frac{1 - a_n \mathrm{e}^{-D_{\mathcal{M}}}}{\beta}.$$

Since $a_n \to 1$, β is constant and $\beta > 1 - \mathrm{e}^{-D_{\mathcal{M}}}$, we can conclude that $\limsup_{n\to\infty} s_n < 1$ must hold. This, together with $\Pr(G_n$ is β-connected) $\leq s_n$ yields that the probability of β-connectivity cannot tend to 1.

Next we prove the other side, i.e., that with β below the threshold by an arbitrarily small fixed $\epsilon > 0$, a.a.s. β-connectivity does occur at least for some model $\mathcal{M}_0 \in \mathcal{F}$. Let us chose the model $\mathcal{M}_0$ as follows. Generate G_n such that each edge is added with some fixed probability p, independently of the others. This is a classical Erdős-Rényi random graph model (see Section 5), which is part of the model family $\mathcal{F}$, since it clearly satisfies name invariance and locality. We use the following result from [8] about the Erdős-Rényi model: If $np = c$ for a constant $c > 1$, then there is a sequence $h_n \to 0$, such that G_n a.a.s. contains a connected component of size $(1 + h_n)\beta_0 n$, where β_0 is the unique root of the equation

$$\beta_0 + \mathrm{e}^{-\beta_0 c} = 1. \tag{9}$$

For this model we have

$$\overline{d}(n) = (n-1)p = \frac{n-1}{n}c$$

implying $D_{\mathcal{M}} = c$.

Now let us fix an arbitrary $\epsilon > 0$. Since the root $\beta_0 = \beta_0(c)$ of equation (9) is less than 1 for any fixed c, but approaches 1 as $c \to \infty$, therefore, due to continuity, we can choose c such that

$$\beta_0 = 1 - \mathrm{e}^{-\beta_0 c} = 1 - \mathrm{e}^{-c} - \epsilon/2$$

holds. Then, by the above cited result, G_n a.a.s. has a connected component of size $(1 + h_n)\beta_0 n$ with $h_n \to 0$. To compensate the effect of $h_n \to 0$, we can slightly decrease β_0 to some $\beta < \beta_0$. Let us choose $\beta = 1 - \mathrm{e}^{-c} - \epsilon < \beta_0$. Then we have proved that G_n, with the appropriately chosen c, a.a.s. contains at least βn connected nodes. Moreover, due to $D_{\mathcal{M}} = c$,

$$\beta = 1 - \mathrm{e}^{-D_{\mathcal{M}}} - \epsilon$$

holds, which completes the proof.

♠

It is worth mentioning that the definition of the treshold function and Theorem 1 directly imply that bounded expected average degrees in $\mathcal{F}$ exclude a.a.s. β_n-connectivity when $\beta_n \to 1$. Then, of course, a.a.s. full connectivity, which corresponds to $\beta = 1$, is also excluded.

Corollary 1. *Let β_n be a sequence in $[0, 1]$ with $\beta_n \to 1$. Then for any local and name invariant abstract geometric random graph model $\mathcal{M}$ it holds that if $D_{\mathcal{M}} < \infty$, then the random graphs generated by $\mathcal{M}$ cannot be a.a.s. β_n-connected.*

8 Solving the Sample Problem

Our results are derived through a more general approach than what is usual in this context. In particular, it is purely probabilistic, geometry is replaced by a higher level abstraction. To illustrate that despite the relative simplicity, the results can

have a surprising strength, let us apply them to the motivating sample problem of Section 2.

As noted in the exposition of the sample problem, it would be hard to solve it for arbitrary choices of $P, D, \mathcal{C}, Q_{x,y}$ and $d(x, y)$ with directly analyzing the stochastic geometry of the model (the notations are re-used from Section 2). On the other hand, we can easily check that it satisfies our general conditions, as shown below.

Let us choose the model domain S as a 3-dimensional phase space, in which each node is represented by a point such that the first two coordinates describe the initial position of the node and the last coordinate encodes which random trajectory was chosen from $\mathcal{C}$ for the node. Let $X_1, X_2, \ldots$ be the representations of the nodes in this phase space.

We can now check that, for any n, the joint distribution of $X_1, \ldots, X_n$ is invariant to re-indexing them. The reason is that both the initial positions and the trajectory choices are generated by processes in which the indices do not play any role. Therefore, the model is *name invariant*. Interestingly, this remains true despite having a lot of dependencies among the nodes: the initial positions of different nodes are not independent (due to clustering), and the trajectory of a given node is also not independent of its initial position, as it is drawn from a probability distribution that may depend on the location. Through this, the trajectories and initial positions of different nodes also become dependent, making their whole movement dependent. Yet, the model is still name invariant.

Let us now consider the links. As defined in Section 2, two nodes are considered connected if during their movement over the time horizon $[0, T_n]$ there is a subinterval of time, of length at least t_n, such that they remain within "radio distance" $\leq r_n$ during the entire subinterval. The radio distance, however, may be very different from the Euclidean distance, it may be described by an arbitrary function that may account for complex propagation characteristics, attenuation, obstacles, and it may also contain independent randomness.

Given some possibly complicated radio distance $d(x, y)$ and the node generation and movement process with possibly complex trajectories, it may not be easy to compute whether a link actually exists between two nodes according to the above definition. On the other hand, for us it is enough to note that once the phase space representations X_i, X_j of any two nodes are given, plus the realization of the independent randomness of the distance, they together determine whether a link exists between the two nodes or not. The reason is that the initial positions and the trajectories, given in the phase space representation, fully determine the movement of the nodes. Once this is known, it determines, along with the realization of the independent randomness of the distance function, whether the link definition is satisfied, i.e., if there is a subinterval of length $\geq t_n$ in $[0, T_n]$, such that the nodes stay within radio distance $\leq r_n$ during the entire subinterval. To actually compute it may not be easy for a sophisticated case, but for our purposes it enough to know that it is *determined* by the listed factors, without knowing anything about the other nodes. This implies that the model is *local*.

Thus, we have established that the problem can be described by a local and name invariant abstract geometric graph model, for any choice of the parameters.

Then, by Theorem 1, the threshold function for β-connectivity is

$$f(D_{\mathcal{M}}) = 1 - \mathrm{e}^{-D_{\mathcal{M}}}.$$

If we require that node degrees are bounded by, say, 4, then we have $D_{\mathcal{M}} \leq 4$, implying

$$f(D_{\mathcal{M}}) = 1 - \mathrm{e}^{-D_{\mathcal{M}}} \leq 1 - \mathrm{e}^{-4} < 0.9817.$$

Thus, the threshold in this case falls below 0.99, so by Theorem 1, it is impossible to achieve that asymptotically 99% of the nodes belong to a connected component, no matter how the other parameters are chosen.

Note that the direct application of our general results was able to cut through a lot of complexity that would otherwise arise if we wanted to reach the same conclusions by directly analyzing the stochastic geometry of such a model.

9 Conclusion

We have quantified the precise trade-off between expected node degrees and the fraction of nodes that belong to a connected component in a large wireless network topology, under very general conditions. Our conditions can be easily checked in most specific cases. Therefore, the approach can serve as a powerful method to explore the degree vs. partial connectivity trade-off in possibly complicated random network topology models, which would be otherwise hard to analyze directly.

Acknowledgment

The author is grateful for the support of NSF Grant CCF-0634848.

References

1. Bollobás, B.: Random Graphs. Cambridge University Press, Cambridge (2001)
2. de Finetti, B.: Funzione Caratteristica di un Fenomeno Aleatorio. Atti della R. Academia Nazionale dei Lincei, Serie 6, Classe di Scienze Fisiche, Mathematice e Naturale 4, 251–299 (1931)
3. Faragó, A.: On the Fundamental Limits of Topology Control in Ad Hoc Networks. Algorithmica 49, 337–356 (2007)
4. Faragó, A.: Scalability of Node Degrees in Random Wireless Network Topologies. In: IEEE JSAC (in press, 2009)
5. Franceschetti, M., Meester, R.: Random Networks for Communication. Cambridge University Press, Cambridge (2007)
6. Gupta, P., Kumar, P.R.: The Capacity of Wireless Networks. IEEE Trans. Information Theory 46, 388–404 (2000)
7. Gupta, P., Kumar, P.R.: Critical Power for Asymptotic Connectivity in Wireless Networks. In: McEneany, W.M., Yin, G., Zhang, Q. (eds.) Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming, pp. 547–566. Birkhauser, Boston (1998)

8. Janson, S., Luczak, T., Rucinski, A.: Random Graphs. Wiley, Chichester (2000)
9. Pahlavan, K., Krishnamurty, P.: Principles of Wireless Networks. Prentice-Hall, Englewood Cliffs (2002)
10. Penrose, M.: Random Geometric Graphs. Oxford University Press, Oxford (2003)
11. Stoyanov, J.M.: Counterexamples in Probability. Wiley, Chichester (1987)