Percolation in multi-hop wireless networks [*]

Massimo Franceschetti[1], Lorna Booth, Matthew Cook[2],
Ronald Meester[3], Jehoshua Bruck[2]

[1] *University of California at Berkeley,* [2] *California Institute of Technology,*
[3] *Vrije Universiteit Amsterdam*

Abstract

Models of wireless ad-hoc and sensor networks are often based on the *geometric disc abstraction*: transmission is assumed to be isotropic, and reliable communication channels are assumed to exist (apart from interference) between nodes closer than a given distance. In reality communication channels are unreliable and communication range is generally not rotationally symmetric. In this paper we examine how these issues affect network connectivity. We compare networks of geometric discs to other simple shapes and/or probabilistic connections, and we find that when transmission range and node density are normalized across experiments so as to preserve the *expected number of connections* (ENC) enjoyed by each node, discs are the "hardest" shape to connect together. In other words, anisotropic radiation patterns and spotty coverage allow an unbounded connected component to appear at lower ENC levels than perfect circular coverage allows. This indicates that connectivity claims made in the literature using the geometric disc abstraction in general hold also for the more irregular shapes found in practice.

Keywords: wireless ad-hoc networks, continuum percolation, Boolean model, random connection model, Poisson processes.

1 Introduction and motivation

Ad-hoc wireless networks are usually defined as a collection of computers equipped with radio transmitters and communicating in a multi-hop fashion, routing each other's data packets. Recently, there has been a growing interest in modelling global properties of these networks as a function of the node population, and basic results on connectivity and throughput capacity have been obtained [3] [5] [6] [7] [9]. These works typically assume that nodes are randomly located and that they are able to reliably relay messages to sufficiently close neighbors.

A first stochastic model that exploited these ideas appeared in the 1961 paper of Gilbert's [4], who studied the problem of multi-hop connectivity of wireless broadcasting stations, establishing the foundations of the theory of Continuum Percolation [11]. In his formulation, points of a two-dimensional Poisson point process represent wireless transmitting stations of range $2r$ and he asks if the system can provide some long-distance communication. He shows the existence of a critical value λ_c for the density of the transmitters, such that, for $\lambda > \lambda_c$, an unbounded connected component of transmitters forms (i.e., the network *percolates*) with probability one, and so long-distance multi-hop communication is possible. In the past forty years Gilbert's model has received much attention from mathematicians and physicists and, recently, extensions relevant to wireless network design have also been proposed [2] [16].

In this paper we consider a further extension that is useful to model more realistic, unreliable and non-rotationally symmetric communication channels. Our motivation is made clear by looking at Figure 1, which depicts the result of an experiment with real wireless links. The probability of correct reception of a packet in the vicinity of a transmitter and in absence of interference is estimated and plotted from real data. The spatial distribution of packet losses significantly deviates from the idealistic assumption of reliable communication between nodes closer than a given distance. Following this observation, we consider a general random connection model where each pair of nodes can be connected according to some (probabilistic) function of their (random) position. Consider, for example, placing nodes at random on the plane and assuming there is an unreliable link between each pair of sufficiently close nodes. Gilbert's question of multi-hop connectivity can then be rephrased as follows: can you find an infinite path in the resulting graph, if each time you need to traverse a new edge you flip a coin (with success probability depending on the length of the edge) to decide if it is possible to do so?

It turns out that if we view results in terms of the average number of edges that can be traversed per node, the simplistic disc model used in continuum percolation, where edges in the connectivity graph correspond to reliable links and can always be traversed, is the "hardest" to percolate among many random connection models. In other words, the disc model forms an unbounded connected component when the expected number of connections (ENC) per node is higher than that required by many random connection models. This indicates that in principle real networks can exploit the presence of unreliable connections to achieve connectivity more easily, if they can maintain the average number of *functioning* connections. Numerical simulations show similar results in the case of anisotropic radiation patterns, where the disc appears the "most compact", and hence the hardest to percolate, among all shapes of equivalent effective area, whether convex or not, and whether

2
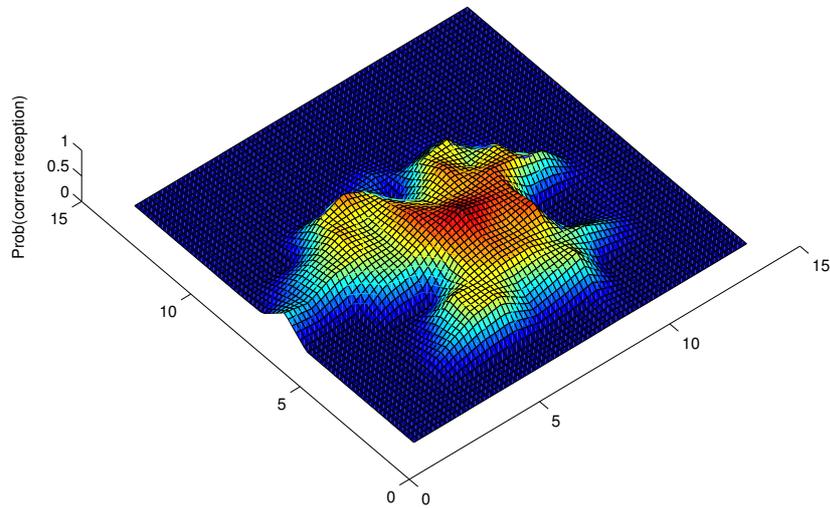
Figure 1: **Real wireless links.** *The figure depicts the estimated probability distribution for correct reception of packets in a real wireless ad-hoc network. The experiment was carried out using 168 nodes (Berkeley Rene motes running the Tiny-OS operating system) placed on the ground in an open space, forming a 12x14 grid, with grid spacing of 2 feet. Only one node transmitted broadcast messages, and the figure depicts the fraction of messages received by the surrounding nodes. Values for non-grid positions are interpolated. Data is available on-line at* `http://localization.millennium.berkeley.edu`

probabilistic or not.[1]

In the next section we present results on noisy communication channels. In Section 3 we present results on anisotropic radiation patterns. Section 4 concludes the paper.

2 Percolation with noisy links

We model imperfect links by considering a random connection model where each pair of points (x_i, x_j) of a Poisson point process of density λ is connected with probability $g(x_i - x_j)$, for some given function $g : \mathbb{R}^2 \to [0, 1]$. We may, for example, pick a function g such that the probability of existence of a link between a transmitter and a receiver decreases as the two points get further away. For generality, however, we prefer to let g be an arbitrary function.

In this section we make the assumption that $g(x)$ only depends on the modulus $\|x\|$, which means that we can view g as a function from $\mathbb{R} \to [0, 1]$; this is convenient when visualizing g with a graph. In order to avoid a trivial model, we assume that the *effective area* $e(g) = \int_{x \in \mathbb{R}^2} g(x) dx$ satisfies $0 < e(g) < \infty$. We call H the class of functions that satisfy this requirement. The two cases $e(g) = 0$ and $e(g) = \infty$ are not interesting because, since $\lambda e(g) = ENC$, nodes have on average respectively 0 or infinitely many neighbors in those cases. Instead, for any function $g \in H$ there is a critical value $\lambda_c(g)$ that ensures connectivity "almost surely" (a.s.) in the sense of [11], i.e., with probability one. This is defined as

$$0 < \lambda_c(g) = \inf\{\lambda : \exists \text{ infinite connected component a.s.}\} < \infty.$$

When $\lambda > \lambda_c$ we say that the random connection model *percolates*. Note that the considered model generalizes standard continuum percolation, where Poisson points are connected with probability one, if discs of radius r centered at each point overlap, as this can be seen as a random connection model with a connection function

$$g(x) = \begin{cases} 1 & \text{if } \|x\| \leq 2r \\ 0 & \text{if } \|x\| > 2r. \end{cases} \tag{1}$$

We are interested in how the percolation properties of the model change when we change the form of the connection function, while preserving its effective area. We start by considering the following transformation.

2.1 Squishing and squashing

Given a function $g \in H$ and $0 < p < 1$, define g_p^{squash} by $g_p^{squash}(x) = p \cdot g(\sqrt{p}x)$. This function, as illustrated in Figure 2, is a version of g in which probabilities are reduced by a factor of p but the function is spatially stretched so as to maintain the original effective area. Therefore, the ENC of each point remains the same, but these connections have a 'wider range' of lengths.

---

[1]Note that aligned ellipses will percolate at the same density as discs of the same area, since an affine transformation of the plane can convert one shape into the other without affecting either connectivity or the Poisson process.
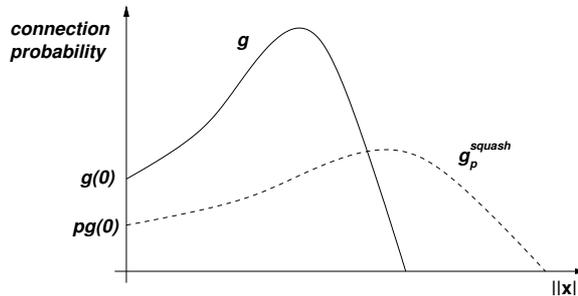
Figure 2: **Squishing and squashing.** *The function g is squished and squashed to give the function* g_p^{squash}.

Note that considering the squashed version of a rectangular connection function such as the one in equation (1), corresponds to connecting Poisson points with probability p if discs of radius $r\sqrt{p}$ centered at the points overlap.

We have the following theorem.

Theorem 2.1 *For all $g \in H$ and $0 < p < 1$ we have,*

$$\lambda_c(g) \geq \lambda_c(g_p^{squash}).$$

Proof of Theorem 2.1 We are to compare the critical densities associated with the connection functions g and g_p^{squash}. We do this by relating both connection functions to a third connection function of greater effective area, namely $h_p(x) = g(\sqrt{p}x)$.

Consider a realisation $\mathcal{G}$ of a random connection model with density λ and connection function h_p. On $\mathcal{G}$, we can perform independent *bond* percolation with the same parameter p, by removing any connection (independent of its length) with probability $1 - p$, independently of all other connections. The resulting random graph can now effectively be viewed as a realisation of a random conection model with density λ and connection function $ph_p(x) = g_p^{squash}$. On the other hand, we can also perform independent *site* percolation on $\mathcal{G}$ with connection function h_p, by removing any vertex of $\mathcal{G}$ (together with the connections emenating from it) with probability $1 - p$, independently of all other vertices. This results in a realisation of a random connection model with density $p\lambda$ and connection function h_p, which can be seen (by scaling) as a realisation of a random connection model with density λ and connection function g.

On any graph G it is well known that the site percolation critical value $p_c^{site}(G)$ and the bond percolation critical value $p_c^{bond}(G)$ are related as

$$p_c^{site}(G) \geq p_c^{bond}(G).$$

In words, this means that if site percolation with parameter p occurs, then also bond percolation with the same parameter occurs. In the above construction, we apply this to $\mathcal{G}$. If site percolation occurs on $\mathcal{G}$, or equivalently, if a random connection model with density λ
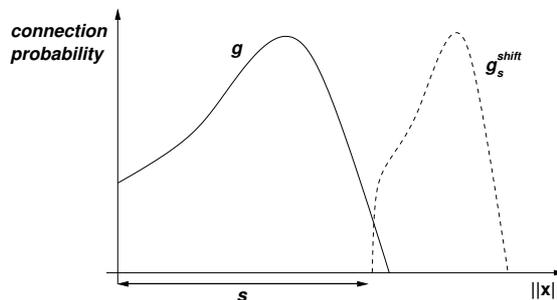
5

Figure 3: **Shifting and squeezing** The function g is shifted and squeezed to give the function g_s^{shift}.

and connection function g percolates, then also bond percolation occurs, or equivalently, a random connection model with density λ and connection function g_p^{squash} percolates. This proves the theorem. $\square$

The theorem above has a certain depth. Essentially, it states that unreliable links are at least as good at providing connectivity as reliable links, if the ENC per node is the same in each case. Another way of looking at this is that the longer links introduced by stretching the connection function are making up for the increased unreliability of the connections.

In some related work Penrose [13] has shown that as a connection function of effective area 1 gets more spread out, its critical density for percolation converges to 1. This can be seen as the limiting case of our Theorem 2.1. Meester, Penrose, and Sarkar [12] proved a similar result as the dimension tends to infinity. Philosophically, the idea that some longer connections help to reach percolation at a lower density of points is also related to the small world networks described, for example, in the paper by Watts and Strogatz [17].

2.2 Shifting and squeezing

Another transformation of g that we consider is $g_s^{shift}(x)$. Here we 'shift' the function g outwards (so that a disc becomes an annulus, for example) by a distance s, but squeeze the function after that, so that it has the same effective area. See Figure 3 for an illustrating example.

Note that the shifting and squeezing transformation maintains on average the same number of bonds per node, but makes them all longer. This contrasts with the squishing and squashing transformation which gives more of a mixture of short and long edges. A natural question to ask is the following. Is the stretching of the edges enough to help the percolation process, or do we need the mixture provided by squishing and squashing? It turns out that the longer edges introduced by the shifting and squeezing transformation are enough to help the percolation process, i.e., long bonds are more useful for percolation than short ones at a given density of points. However, it must be noted that a scaling of the whole picture will make the bonds longer while changing the density of points. Naturally, this process does not change the connectivity, but is essentially different in that it changes the effective area. We also note that, at first sight, it may seem that the considered transformation does not have an immediate practical application to wireless networks; nevertheless it is a first step in
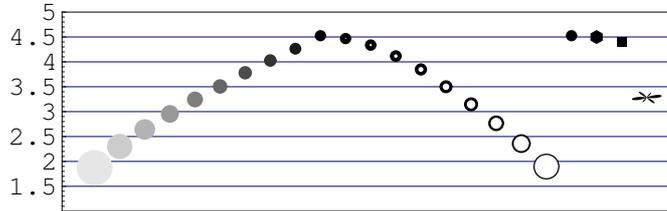
6

Figure 4: **Percolation thresholds of various shapes.** The numbers on the left represent the ENC values needed for percolation of the different shapes.

understanding how to compare the percolation properties of arbitrary connection functions of the same effective area, which is of great practical interest. Moreover, as discussed above, it is critical to determine how useful long edges are for connectivity.

The effect of the shifting and squeezing transformation and of the squishing and squashing one can be visually seen by examining the left-hand side of Figure 4, where we report simulation results for a connection function of the type described by equation (1). Each shape shown has effective area 1 and is positioned so that the height of the center of the shape is at the critical ENC value needed for percolation of that shape. For example, the solid disc at the top is at a height of around 4.5, indicating that disc percolation occurs whenever each disc has over 4.5 neighbors on average, in agreement with numerical results such as in [14]. The bottom of the graph is at ENC of 1, which is the theoretical minimum, attained by very large, diffuse shapes [13]. The ascending shapes on the left are discs with probabilistic connections to other points touching the disc, with probability varying linearly from 0.1 for the large light gray circle to 1.0 for the non-probabilistic disc at the top. Then the descending shapes are annuli whose inner radius varies linearly from 0.0 of the outer radius for the solid disc at the top (the same disc as the last one in the sequence of ascending shapes) to 0.9 for the thin ring at the bottom. The error bars for each shape's height are less than the distance between adjacent ticks shown on the vertical axis, except for the bottommost shapes, where the finiteness of the simulation size more easily interferes with the long range connectivities that are common in components formed by those shapes. The remaining shapes depicted on the right-hand side of the figure, along with the technique we used to determine the ENC value needed for percolation of each shape, are described in Section 3.

We now prove a theorem for shifting and squeezing that shows that it is possible to decrease the percolation threshold of the random connection model by taking a sufficiently large shift. Consider the following percolation process. We start with a Poisson process in two dimensions, with density $\lambda > 0$. Denote by A_r the annulus with inner radius r and with area 1, so that A_0 is just the disc of unit area. Now consider, for each point x of the Poisson process, the set $A_r(x) := x + A_r$, that is, the annulus with inner radius r centered at x. We draw undirected edges between x and all points in $A_r(x)$. This gives a random graph, and we are interested in the critical value of this process in terms of percolation, and we denote this critical value by $\lambda_c(r)$.

Theorem 2.2 *For all $r \geq 0$ and $1 < \lambda^* < \lambda_c(r)$, there exists a finite $r^* > r$ such that A_{r^*}*

7

percolates, for all $\lambda > \lambda^*$.

The theorem immediately follows from the following two results.

Lemma 2.3 $\lambda_c(r) > 1$, *for all $r \geq 0$.*

Proof of Lemma 2.3 Fix $r \geq 0$. We compare the percolation process to a Galton-Watson branching process with a Poisson-λ offspring distribution as follows. The 0-th generation point is placed at the origin. Its children are distributed uniformly and independently over an annulus of area 1, centered at the origin. Subsequent children of any point x are also distributed uniformly and independently over an annulus of area 1 centered at x, but if such a child happens to fall into one of the annuli that has been considered before, it is discarded. Note that the overlap between the annulus centred at x and the previously considered annuli is uniformly bounded below by some number $c(r) > 0$, namely the intersection with the annulus of the parent of x. This means that the offspring of any point (apart from the origin) is stochastically dominated by a Poisson random variable with parameter $\lambda(1-c(r))$. Hence, there is a $\lambda_0 > 1$ such that $\lambda_0(1-c(r)) < 1$ and for this λ_0, the percolation process is dominated by a subcritical branching process, and hence dies out. This means that infinite components cannot exist for λ_0, which shows that $\lambda_c(r)$ is strictly larger than 1. $\qquad\square$

Theorem 2.4 $\lim_{r\to\infty} \lambda_c(r) = 1$.

Proof of Theorem 2.4 The proof of this theorem is a little more involved. To ease the presentation, we do not give all details here, but hopefully enough to make the present proof self-contained.

The proof of proceeds via a suitable renormalisation and comparison with a discrete directed percolation process and a supercritical branching process. We first describe a supercritical spatial branching process which is, in some sense to be made precise below, the limiting object of our percolation process as $r \to \infty$.

A spatial branching process. Consider an ordinary Galton-Watson branching process with Poisson-λ offspring distribution, where $\lambda > 1$. Note that this process is supercritical, and hence there is a positive probability that the process does not die out. We add a geometric element to this process as follows: The 0-th generation point is placed at the origin, say. The children of any point x of the process are distributed uniformly and independently over the circumference of a circle with radius 1, centered at x.

A sequential construction of the percolation process. We now describe a way to construct a percolation cluster in our percolation process, which looks very much like the branching process just described. One of the aspects of this construction is that we create the point process along the way, so at the beginning of the construction, we think of the plane as being completely empty. The density of the underlying Poisson process is the same $\lambda > 1$ as above.

We start with a point in the origin, and consider the annulus $A_r = A_r(0)$. We now 'fill' A_r with a Poisson process, that is, we take a Poisson-μ random number of points, and distribute these uniformly (and independent of each other) over A_r. These points are the

points that are directly connected to the origin. If there are no points in A_r we stop the process; if there are points in A_r we denote these by $y_1, y_2, \ldots, y_r$, ordered by modulus, say. In order to decide about the connections from y_1, we consider $A_r(y_1)$ and 'fill' this annulus with an independent Poisson process. The (random) points that we obtain in $A_r(y_1)$ are the points that are directly connected to y_1 but not to 0.

Now note that we make a mistake by doing this, since the region $A_r(0) \cap A_r(y_1)$ is not empty, and this region has now been filled twice, and therefore the intensity of the Poisson process in the intersection is 2λ, not λ. For the moment we ignore this problem; we come back to this in a few moments. We now continue in the obvious way, each time 'filling' the next annulus with a Poisson process, and each time making a mistake as just observed.

Comparison between branching process and percolation process. Ignoring the mistakes we make, this process is similar to the spatial branching process described before. We can actually couple the two processes (still ignoring mistakes) by insisting that the offspring of the branching process also be the points of the percolation process. If a point in the branching process is placed at a certain position at distance 1 from its parent, then the point in the percolation process is located at the same angle, and uniformly distributed over the width of the annulus. Since $\lambda > 1$, the percolation process would continue forever with positive probability, thereby creating an infinite percolation component.

However, we have to deal with the mistakes we make along the way. We have two tools at our disposal that can be helpful now. First, it should be noted that the overlap between the various annuli gets smaller as $r \to \infty$. Secondly, we will only use the coupling between the spatial branching process and the percolation process for a uniformly bounded number of generations. So fix the upper bound M of points involved, and $\epsilon > 0$. We can then choose r so large, that the probability that any point of the Poisson process in the first M annuli of the process falls into any intersection of two different annuli is at most ϵ.

In fact, we can do even more. If we start not with one point in the origin, but with a collection of N points such that no two such points are within distance δ of each other, then we can make r so large that the probability of a mistake in the union of these N independent branching processes (each up to M annuli involved) is again at most ϵ.

The renormalisation. We now describe the renormalisation anticipated before. Divide the positive quadrant into boxes of size $L \times L$, where we choose L in a moment. The box with lower leftmost point (iL, jL) is denoted by $B_L(i, j)$. We now choose the various quantities as follows (explanation follows after the list). First let ϵ and δ be given positive numbers, and let μ be as before.
1. First choose N so large that the probability that at least one out of a collection of N independent spatial branching processes survives for ever, is at least $1 - \epsilon$.
2. Then we choose L so large that the probability that any box $B_L(i, j)$ contains a collection of N points such that no two points of the collection are within distance δ of each other and the boundary of the box, is at least $1 - \epsilon$. We call such a collection of points a *good* collection.
3. Then we choose M so large that in the spatial branching process (which, we recall, uses circles of radius 1) it is the case that if we start with a good collection of points in $B_L(0, 0)$, the probability that the total progeny of this collection contains a good collection in both $B_L(1, 0)$ and $B_L(0, 1)$ without using more than M annuli, is at least $1 - \epsilon$.

4. Finally, we take r so large that the probability of a mistake involving one of the first M annuli of the sequential construction process (not the spatial branching process) is at most ϵ. When we compare between the spatial branching process and the percolation process, there is another mistake that we need to say something about. In (3) we require that the progeny in the branching process has enough points in certain regions. Now the coupling between the branching process and the percolation process is not perfect, since the annuli have a certain width, and two offspring of the same parent will not be at the exact same distance from the parent. Therefore, there will be some mistake when we compare the two processes. However, if we consider the percolation process for at most M annuli, and rescale the boxes by a factor of r, then the probability that there is a point which ends up in the wrong box, can be made arbitrarily small by taking r large.

Perhaps it is necessary to spend a few words on the reasons that the various choices are indeed possible. For (1), it is clear that an N as described exists, given the fact that each individual process has positive probability to survive. (2) is clear, but (3) needs a little reflection. Indeed, a single line of descent in the spatial branching process can be seen as a random walk with zero drift. Since the number of descendants grows to infinity, so does the number of independent random walks, and the claim then follows from the fact that two-dimensional zero-drift random walk is recurrent. The choice of r in (4) has been commented on before and is possible since the overlap between two annuli gets smaller when the radii increase to infinity.

Now we sketch the actual renormalisation. We start with the boxes $B_L(i, j)$ defined before, but rescale according to (4) above. To explain what this means, note that the spatial branching process has distance 1 between a parent and child, and the choice of N, L and M are in terms of this process. When we want to make the probability of mistakes small in the first M annuli of the process, we have to take r large in (4), and as a result, all stepsizes should also be multiplied by a factor of r. When we do that, it follows from the choice of r in (4) that the spatial branching process and the percolation process look alike during the first M annuli of the processes, in the sense that if a point in the branching process ends up in a certain box $B_L(i, j)$, then the corresponding point in the percolation process ends up in the corresponding box $B_{rL}(i, j)$ (the box with side length rL whose lower left corner is at (irL, jrL) and vice versa.

Now order the vertices of the positive quadrant in such a way that the modulus is non-decreasing. We look at vertices (i, j). We call the vertex (i, j) *open* if the following two things happen in the percolation process:

1. The box $B_{rL}(i, j)$ contains a good collection of points.

2. The progeny of this good collection, restricted to the first M annuli of the process contains a good collection in both $B_{rL}(i + 1, j)$ and $B_{rL}(i, j + 1)$.

It should be noted that we stop the percolation scanning procedure from any point which falls into one of the two 'target boxes'. So we do not yet check offspring of these points; this will be done at a later step of the algorithm.

We now consider the points of the first quadrant one by one, in the given order. The probability that $(0, 0)$ is open can be made as close to one as desired, by appropriate choice

of the parameters. In particular, we can make this probability larger than $p_c + \epsilon$, where p_c is the critical value of directed two-dimensional independent site percolation on the square lattice.

If the origin is not open, we terminate the process. if it is open, we consider the next vertex, $(0, 1)$ say. The corresponding box $B_{rL}(0, 1)$ contains a good collection, and we can now start all over again with this good collection of points, and see whether or not we can reach $B_{rL}(1, 1)$ and $B_{rL}(0, 2)$ in the same way as before. Note that there is one additional problem now, since we have to deal with overlap with annuli from previous steps of the algorithm, that is, with annuli involved in the step from $(0, 0)$ to $(0, 1)$. This is easy though: since we have bounded the number of annuli involved in step of the procedure, there is a uniform upper bound on all annuli which have any effect on any given step of the algorithm. Therefore, the probability of a mistake due to any of the previous annuli can be made arbitrarily small by taking r even larger, if necessary. This shows that we can make the probability of success each time larger than $p_c + \epsilon$, no matter what the history of the process is. This implies that the current percolation process dominates independent site percolation with parameter $p_c + \epsilon$, and is therefore supercritical. □

It is not hard to see that this proof can be generalised to any g_s^{shift}. In the general case, the offspring of a particle is distributed according to an inhomogeneous Poisson process, depending on the connection function. Details are left to the reader.

3 Anisotropic radiation patterns

We now discuss another extension of the connectivity model, namely non-rotationally symmetric transmission ranges. In practice, the communication range of a radio transmitter can not be perfectly rotationally symmetric. The imperfection is due to the characteristics of the transmitting antenna as well as the surrounding environment. It is therefore useful to examine how percolation properties are influenced by the shape of the transmitter footprint.

We consider centrally symmetric shapes that are not necessarily rotationally symmetric, i.e. shapes that are identical to themselves when rotated by 180 degrees. Accordingly, we let $B \subset H$ be the set of all connection functions $b_C(x)$ of the form $b_C(x) = 1$ if x is inside some convex centrally symmetric shape C of area 1 and 0 otherwise. Let D be the disc of area 1 and let S be the square of area 1.

Centrally symmetric shapes allow us to consider only bi-directional links: a particularly nice property of centrally symmetric shapes is that two shapes overlap each other's centers if and only if the shapes scaled down by a factor of two in each direction overlap each other. If a node falls inside the shape of another node, then the reverse is true as well. This property does not hold for non-centrally symmetric shapes, see Figure 5.

Jonasson [10] has shown that if any convex shape of area 1 percolates at certain density, then triangles (that are not centrally symmetric) of area 1 will also do so. Roy and Tanemura [15] strengthened this result to show a strict inequality between the critical density of triangles and that of any other given convex shape of the same area. Jonasson [10] has also shown that the convex shape with the highest critical density will be centrally symmetric. It is natural to ask which of the centrally symmetric shapes percolate most and least easily. We believe the answers are the square and the disc.
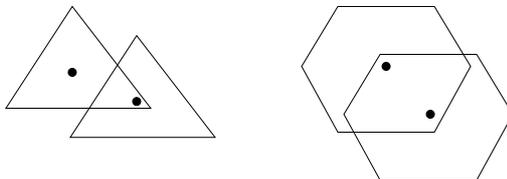
Figure 5: **Central symmetry and asymmetry of shapes.** Left-hand side, centrally asymmetric shapes: the lower point is contained in the triangle around the higher point but not vice versa. Right-hand side, centrally symmetric shapes. The lower point is contained in the hexagon around the higher point and vice versa.

Conjecture 3.1 *For all $b_C \in B$ we have*

$$\lambda_c(b_S) \leq \lambda_c(b_C) \leq \lambda_c(b_D).$$

There is an important practical consequence that follows if Conjecture 3.1 is true. If the disc is the centrally symmetric shape that percolates at the highest density, it follows from Jonasson [10] that this is also the highest percolating density shape over all the convex shapes. This means that percolation theory results on the existence of an unbounded connected cluster, derived in the standard model where overlapping discs are connected with probability one, are robust. If an ideal model, where transmission footprints are perfect discs, allows long-distance multi-hop communication, then long-distance multi-hop communication is also possible, under the same density conditions, in a less idealistic model where transmission footprints can have any convex shape with the same effective area as the disc.

We support Conjecture 3.1 with computer simulations that we performed using different shapes of unit area. We report the size of the largest and second largest clusters found in simulations for discs and squares in Figure 6. For each shape, we ran a total of 7000 experiments, using 100000 randomly placed discs for each experiment. The density of the shapes in the experiments varied from a minimum of 0.25 to a maximum of 0.32, with a 0.00001 incremental step. For each density, the number of vertices in the largest and second largest clusters were recorded. Each data point shown in Fig. 6 corresponds to an average within a sliding window of 100 consecutive experiments. Results show that the sizes of the largest two clusters of discs diverge at a higher value of the density than do the clusters of squares. If we shift the plot obtained for squares by 3% (see Figure 7), we find a striking matching of the curves. Accordingly, we conclude that squares seem to percolate at a density approximately 3% smaller than discs. This implies that the average number of connections needed for percolation (CNP) for squares is similarly reduced from the CNP for discs. Just as the curve for squares in Figure 7 lined up after being shifted by 3%, we can take the curve for any shape and see how much it needs to be expanded to achieve a good fit with the curve for discs. Since the percolation threshold for discs is known (experimentally), we assume that the percolation threshold for the other shape is at the same point in the graph, and so we divide the threshold for discs by the expansion factor needed for lining up the
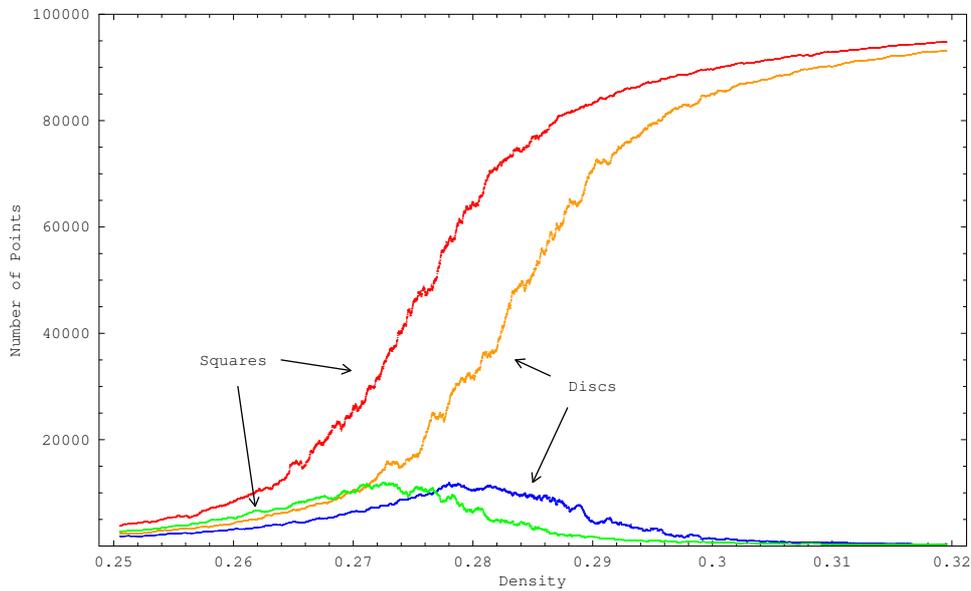
Figure 6: **Size of larger clusters**. We report the size of the largest and second largest clusters of discs and squares of area one. The size of the largest cluster of discs tends to diverge at a smaller value of the density than for squares.
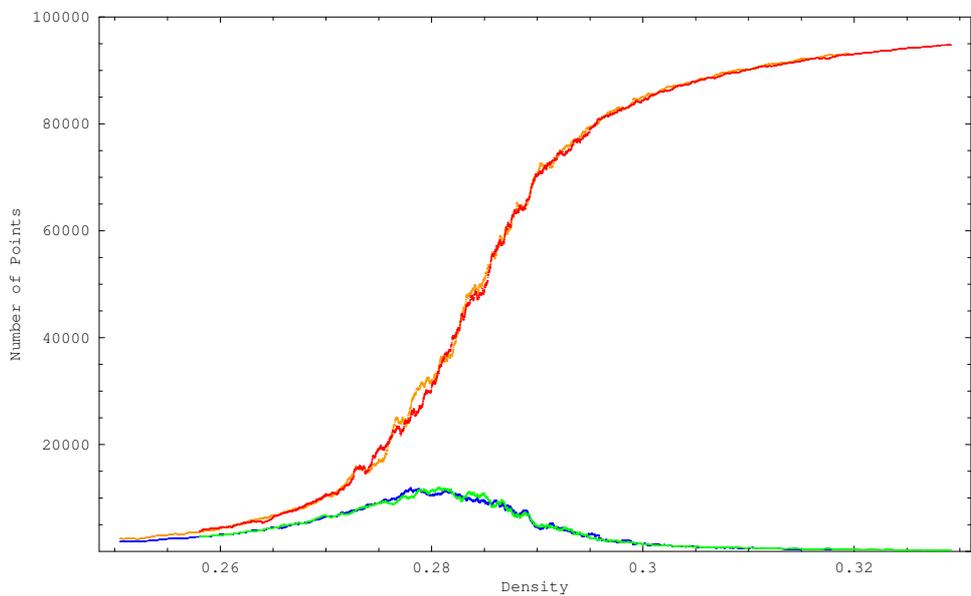

Figure 7: **Shifting by 3%**. After shifting the curves of the size of the largest and second largest clusters of squares by 3% , we find a striking matching with the corresponding curves plotted for discs. This indicates that squares percolate at a density approximately 3% smaller than discs.

graphs to get an estimate of the percolation threshold for the other shape. Final results, depicted on the right-hand side of Figure 4, suggest that solid discs have the highest CNP of any shape. Note that the first three shapes on the right-hand side of the figure are a solid disc (repeated for comparison with the hexagon), a hexagon, and a square. The solid disc, at a CNP value of 4.5 is the highest of all the shapes depicted. It would appear that $2n$-gons have a CNP that increases with n, approaching the CNP for discs as n goes to infinity. The last shape on the right is an irregular lobed shape, inspired by the radiation pattern of a multidirectional antenna, and we see that its CNP is markedly lower than that of discs.

To the best of our knowledge, the numerical technique outlined above has never been used before to find the percolation threshold of two-dimensional shapes. As validation of the proposed procedure, we note that our results for squares are in agreement with those in [1]. In that paper numerical evidence is also given that allowing a random orientation of squares decreases their percolation threshold. This, of course, can not happen for discs, which remain the hardest shape to percolate.

4   Conclusions

We have studied the effects of unreliable communication and anisotropic radiation patterns on the connectivity of wireless ad-hoc networks. Following a percolation theory approach, we have shown that real networks can exploit the presence of unreliable connections and anisotropic radiation patterns, to achieve connectivity more easily, if they can maintain an average number of functioning connections. This result indicates that connectivity claims made in the literature, using the simplistic assumption of reliable communication between nodes closer than a given distance, continue to hold in more realistic scenarios. Finally, our study leads to the "discs are hardest" conjecture, a topic of future research.

References

[1] D. R. Baker, G. Paul, S. Sreenivasan, and E. Stanley (2002). Continuum percolation threshold for interpenetrating squares and cubes. *Physical Review E 66, 046236.*

[2] L. Booth, J. Bruck, M. Franceschetti, and R. Meester (2003). Covering algorithms, continuum percolation, and the geometry of wireless networks. *Annals of Applied Probability* **13**(2), to appear.

[3] O. Dousse, F. Baccelli, P. Thiran (2003). Impact of interferences on connectivity in ad-hoc networks. *IEEE/ACM Transactions on Networking*, to appear.

[4] E. N. Gilbert (1961). Random plane networks. *Journal of SIAM* **9**, 533-543.

[5] M. Grossglauser and D. Tse (2001). Mobility increases the capacity of ad hoc wireless networks. *ACM/IEEE Transactions on Networking* **10**(4), 477-486.

[6] P. Gupta and P. R. Kumar (1998). Critical power for asymptotic connectivity in wireless networks. *Stochastic Analysis, Control, Optimization and Applications: A Volume in*

Honor of W.H. Fleming, W. M. McEneaney, G. Yin, and Q. Zhang (Eds.), Birkäuser, Boston.

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

[8] O. Häggström and R. Meester (1996). Nearest Neighbour and Hard Sphere Models in Continuum Percolation. *Random Structures and Algorithms* **9**, 295-315.

[9] S. R. Kulkarni, and P. Viswanath (2001). A deterministic approach to throughput scaling in wireless networks. *Proceedings of the International Symposium on Information Theory (ISIT), 351.*

[10] J. Jonasson (2001). Optimization of shape in continuum percolation. *Annals of Probability* **29**(2), 624-635.

[11] R. Meester and R. Roy (1996). Continuum percolation. *Cambridge University Press.*

[12] R. Meester and M. D. Penrose and A. Sarkar (1997). The random connection model in high dimensions. *Statistics and Probability Letters* **35**, 145-153.

[13] M. D. Penrose (1993). On the spread-out limit for bond and continuum percolation. *Annals of Applied Probability* **3**(1), 253-276.

[14] S. Quintanilla, S. Torquato, and R. M. Ziff (2000). Efficient measurement of the percolation threshold for fully penetrable discs. *Journal of Physics A: Mathematical and General* **33**, L399-L407.

[15] R. Roy and H. Tanemura (2002). Critical intensities of Boolean models with different underlying convex shapes. *Advances in Applied Probability* **34**(1), 48-57.

[16] A. Wagner and V. Anantharam (2002). Designing a Contact Process. The Piecewise-Homogeneous Process on a Finite Set with Applications. *Annals of Applied Probability, submitted for publication.*

[17] D. J. Watts, and S. H. Strogatz (1998). Collective dynamics of small world networks. *Nature* **393**, 440-442.