

Continuum percolation with unreliable and spread-out connections*

Massimo Franceschetti¹, Lorna Booth, Matthew Cook²,
Ronald Meester³, Jehoshua Bruck²

¹*Dept. of Electrical Engineering and Computer Sciences, UC Berkeley*

²*Dept. of Computation and Neural Systems, CalTech*

³*Dept. of Mathematics, Vrije Universiteit Amsterdam*

Abstract

We derive percolation results in the continuum plane that lead to what appears to be a general tendency of many stochastic network models. Namely, when the selection mechanism according to which nodes are connected to each other, is sufficiently spread out, then a lower density of nodes, or on average fewer connections per node, are sufficient to obtain an unbounded connected component. We look at two different transformations that spread out connections and decrease the critical percolation density while preserving the average node degree. Our results indicate that real networks can exploit the presence of spread-out and unreliable connections to achieve connectivity more easily, provided they can maintain the average number of *functioning* connections per node.

Keywords: continuum percolation, random connection model, Poisson processes, stochastic networks, unreliable connections.

*This work was partially supported by the Lee Center for Advanced Networking at Caltech, and by DARPA grant number F33615-01-C-1895 at Berkeley. Contact author address: Massimo Franceschetti Dept. of Electrical Engineering and Computer Sciences, University of California at Berkeley, 333 Cory Hall, Berkeley, CA, 94720-1770. Phone 510-643-5806. Email massimof@EECS.berkeley.edu

1 Introduction

Stochastic network models have received much attention in the past few years for their ability to describe physical, chemical, biological, engineering and social structures. One of the early models originally appeared in the 1961 paper of Gilbert's [6], motivated by multi-hop communication of wireless broadcasting stations and established the foundations of the theory of Continuum Percolation. In Gilbert's 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. The monograph by Meester and Roy [8] describes many mathematical extensions of Gilbert's model that later appeared in the literature. In physics, continuum percolation is applied to the study of clustering behavior of particles in continuum systems and is relevant in phenomena like conduction in dispersions, flow in porous media, elastic behavior of composites, sol-gel transition in polymers, aggregation in colloids, and the structure of liquid water, to name a few, see for example the works [12, 13, 14] and references therein. In engineering, interest relies in the connection of continuum percolation to the recent developments of radio packet, ad-hoc networks [1, 3, 4, 11, 16].

In this paper, we study a generalization of Gilbert's model in which each pair of Poisson points can be connected according to some (probabilistic) function of their (random) position, and we look at how the critical density value λ_c varies with the shape of the connection function. We find what appears to be a general tendency in these kind of random connection models, namely, when the selection mechanism with which nodes are connected to each other is sufficiently spread-out, then a lower density of nodes, or equivalently fewer connections per node, will suffice to obtain an unbounded connected component. We show this by considering two different transformations that spread out the connection function, decreasing its critical density value λ_c , while preserving the average node degree. Our results indicate that real networks can exploit the presence of spread out, unreliable connections, to achieve connectivity more easily, provided they can maintain on average few *functioning* connections per node.

There are a number of related results in the literature that support this general tendency of obtaining connectivity by using fewer, spread out connections. Penrose [10] has shown that as a connection function of area 1 gets more spread-out, and the probability of connection between each pair of nodes tends to zero, its critical density for percolation converges to one, and so does the average number of connections per node. This can be seen as the limiting case of our Theorem 2.1, while our Theorem 2.4 shows convergence to one for a different kind of spreading the connection function. Meester, Penrose, and Sarkar [9] proved a similar limiting result as the dimension of the space spreads out to infinity. Dubhashi, Häggström, and Panconesi [5] proved that for a certain class of random

geometric graphs a giant component forms by selecting only few, among a large number of potentially far away neighbors. On the contrary, when connections do not spread out, and nodes connect only to nearest neighbors, then few links are not enough to obtain an unbounded component, see Häggström and Meester [7], while no finite number of nearest neighbors connections suffices for connectedness, see Xue and Kumar [15]. Finally, we point out that, independently and simultaneously, Balister, Bollobás, and Walters [2], proved a similar result as our Theorem 2.4, using a somewhat different approach¹.

2 The random connection model

We consider 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 $g : \mathbb{R}^2 \rightarrow [0, 1]$. We may, for example, pick a function g such that the probability of existence of a link between two points decreases as they get further away. For generality, however, we prefer to let g be an arbitrary function. We make the additional assumption that $g(x)$ only depends on the modulus $|x|$, which means that we can view g as a function from $\mathbb{R} \rightarrow [0, 1]$; 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 $\lambda e(g)$ is the expected number of connections per node (ENC), nodes have on average respectively 0 or infinitely many neighbors in those cases. Instead, it is well known that for any function $g \in H$ there is a critical value $\lambda_c(g)$ that ensures connectivity almost surely (a.s.), i.e., with probability one [8]. 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

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

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

¹Note to the Editor: the authors in [2] acknowledge our simultaneous, independent discovery in their paper. A discussion about the difference in the two approaches is given in Section 2.2




Figure 1: **Squishing and squashing.** The function g is squished and squashed to give the function g_p^{squash} .

2.1 Squishing and squashing transformation

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 1, 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. Note that an affine transformation of the plane, for example converting discs into aligned ellipses of the same area, would also introduce a wider range of connection lengths, but would not affect the connectivity threshold.

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 larger 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 connection 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 emanating 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 .




Figure 2: **Shifting and squeezing** The function g is shifted and squeezed to give the function g_s^{shift} .

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 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

We point out that 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. Note also that the considered transformation spreads out connections, but the resulting graph still presents a mixture of short and long links at all scale lengths. Hence, it remains to be established, whether longer connections alone are enough to help the percolation process, or do we need the mixture of all scale lengths provided by the squishing and squashing transformation. It turns out that long bonds are more useful for percolation than short ones at a given density of points. We show this by introducing another transformation that maintains on average the same number of bonds per node, but makes them *all* longer.

2.2 Shifting and squeezing transformation

We call this transformation $g_s^{shift}(x)$. 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 2 for an illustrating example; we do not bother to give the precise technical definition.




Figure 3: **Percolation thresholds of various shapes.** The numbers on the left represent the ENC values needed for percolation of the different shapes.

From numerical simulations it appears as the effect of the shifting and squeezing transformation is similar to that of the squishing and squashing one. This would suggest that a monotonicity result along the same lines of Theorem 2.1 should hold in this case. However, we will only prove a weaker, limiting result. First, let us look the simulation results that are presented in Figure 3. 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 4.51, indicating that disc percolation occurs whenever each disc has over 4.51 neighbors on average, in agreement with numerical results such as in [12]. The bottom of the graph is at ENC of 1, which is the theoretical minimum, attained by very large, diffuse shapes [10]. The ascending shapes on the left represent the squishing and squashing transformation applied to a connection function of the type described by equation (1). These 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 represent the shifting and squeezing transformation that transforms discs into 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. Finally, simulation results for some more irregular geometric shapes of the same area are grouped on the right-hand side of the figure.

We now turn to our analytic results for the shifting and squeezing transformation. We show that it is possible to decrease the percolation threshold of the random connection model by taking a sufficiently large shift. We also show that in the limit of very spread-out connections, there is no difference between diffuse discs and thin annuli, i.e., between squishing and squashing, and shifting and squeezing, as they both converge to one (see also [2] [10]).

Results hold for any $g \in H$, but to ease the presentation we focus only on annuli shapes.

The general case is obtained following the same proof steps. Finally, we want to spend few words on our proof technique. We proceed comparing our limiting percolation process to a branching process where the population may increase infinitely as soon as the expected offspring is larger than one. As connections are more spread out, the two processes become tightly coupled and, in conjunction with a renormalisation argument, this explains why percolation is achieved as the expected node degree exceeds one. Comparison with branching processes is quite standard in percolation theory [8], and the reader can also refer to [7] for a similar application of this technique. Still, there are some technical difficulties that we must overcome in the proof. Another possible comparison is with ordinary random graphs, that exhibit a giant component as the average node degree exceeds one. This approach is more combinatorial in flavor and has been exploited in [2]. The two methods really are very different. Roughly speaking, the method in [2] gives more precise results about the rate of convergence for this particular case, while our method is easier to generalise to other connection functions.

We formally proceed as follows. 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^*} 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. \square

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

Proof of Theorem 2.4 The proof of this theorem is more involved. It 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 $r \rightarrow \infty$.

A spatial branching process. Consider an ordinary Galton-Watson branching process with Poisson- λ offspring distribution, where $\lambda > 1$. 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. We will then couple the two processes. 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- λ random number of points, and distribute these uniformly (and independent of each other) over A_r . These points 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, \dots, y_s , 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, in the same way as before. The (random) points that we obtain in $A_r(y_1)$ are directly connected to y_1 but not to 0. Now note that we make a mistake by doing this, in the sense that 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 instead of the desired 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 possibly making a mistake as just observed.

Comparison between branching process and percolation process. Ignoring the mistakes we make, the sequential construction described above is similar to the spatial branching process. 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 relative 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 $r \rightarrow \infty$. Secondly, we will only use the coupling between the spatial branching process and the percolation process for a uniformly bounded number of annuli, to build a renormalised process that dominates supercritical directed site percolation on a square lattice.

Renormalisation blocks. We now describe the renormalisation and the coupling that works for a uniformly bounded number of annuli. We first look at a construction for the spatial branching process, and then show that the same construction is achieved in the percolation process, with arbitrary large probability. 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 $B_L(i, j)$. Let ϵ and δ be given positive numbers, and let λ be as before. We choose various quantities as follows.

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 choose a preliminary 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, is at least $1 - \epsilon$. We call such a collection of N points a *good* collection. We might need a larger L in the next step, but the property specified here remains intact when L increases.
3. Then choose M so large that in the spatial branching process (which, we recall, uses circles of radius 1) the following is the case: if we start with any good collection of points in $B_L(0, 0)$, and we discard all further offspring of any point which falls in either $B_L(2, 0)$ or $B_L(0, 2)$, then the probability that the total progeny of this collection, restricted to the first M points, contains a good collection in both $B_L(2, 0)$ and $B_L(0, 2)$, is at least $1 - \epsilon$. The possibility of this choice requires a little reflection. We want to ensure that the N branching processes, after generating at most M points, will create a good collection of points in the two ‘target’ boxes $B_L(2, 0)$ and $B_L(0, 2)$, even if we discard all offsprings departing from points inside the two target boxes. Therefore, first we must ensure that the branching processes starting with the good collection are not stopped by points falling in either $B_L(2, 0)$ or $B_L(0, 2)$, and this can be achieved with high probability by taking L larger if necessary. After that, we must ensure that enough points end up in the two target boxes. Since a single line of descent of any point follows a simple two-dimensional random walk with zero drift, this random walk is recurrent, and after a sufficiently large number of steps, some descendent will end up in either $B_L(0, 2)$ or $B_L(2, 0)$. The probability that this happens for at least N lines of descent in each of the two target boxes and that the collection of ‘terminal’ points in each of the two target boxes contain a good set, can be made as high as we like by taking the allowed total number M of points large enough.
4. Finally, we choose a δ' small enough so that the probability that the distance between any two of the first M points generated by the initial N branching processes is smaller than

δ' , is at most ϵ .

Note that the construction described up to now has been in terms of the spatial branching process and it ensures that a good collection of points in $B_L(0, 0)$ can reach good collections in $B_L(0, 2)$ and $B_L(2, 0)$, in a bounded number of iterations, with probability at least $1 - \epsilon$. We now want to show that it is also possible to obtain the same, with high probability, in the sequential percolation process. To do this we will need to take the radius r of the annuli in the percolation process large enough. First of all, we note that if we fix an upper bound M of annuli involved, and $\epsilon > 0$, we can choose r so large that the probability that in the union of N sequential percolation processes, any point falls into an intersection of two among the first M annuli, is at most ϵ . This is because we start the N processes with annuli separated by at least δ , and evolve generating a total number of at most M annuli that are at distance at least δ' to each other. Hence, the total overlap between the annuli can be made as small as we want by taking r large.

The percolation process and the branching process now look alike in the first M steps, in the sense that if the branching process survives while generating M points, the percolation process also survives with high probability. To complete the construction we need something slightly stronger than this. We also need to make sure 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. Note that since the annuli have a certain width, two offspring of the same parent will not be at the exact same distance from the parent. Therefore, points can possibly end up in the wrong box. However, the probability that there is a point which ends up in the wrong box can again be made less than ϵ by taking r large. To explain why this is, note that the spatial branching process has distance 1 between a parent and child, and the choice of N , L , M and δ' are in terms of this process. When we couple the branching process with the percolation process and we take r large, we also have to scale the whole picture by a factor r . When we do this, the width of each annulus becomes smaller and tends to 0. Therefore, the probability of making a mistake by placing a point in the wrong box decreases to 0 as well.

Dynamic coupling with discrete percolation. We are now ready to show that the renormalisation described above dominates a supercritical directed site percolation process on a square lattice. Let us order the vertices corresponding to boxes in the positive quadrant in such a way that the modulus is non-decreasing. We look at vertices (i, j) . We call the vertex $(0, 0)$ *open* if the following two things happen in the percolation sequential construction:

1. The box $B_{rL}(0, 0)$ contains a good collection of points; we choose one such collection according to some previously determined rule.
2. The progeny of this chosen good collection, restricted to the first M annuli of the

process (and where we discard further offspring of points in any of the two target boxes $B_{rL}(0, 2)$ and $B_{rL}(2, 0)$) contains a good collection in both $B_{rL}(0, 2)$ and $B_{rL}(2, 0)$.

We now consider the vertices (i, j) associated with each box 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, $(0, 2)$ say. The corresponding box $B_{rL}(0, 2)$ contains a good collection, and we can choose any such good collection according to some previously determined rule. We start all over again with this good collection of points, and see whether or not we can reach $B_{rL}(2, 2)$ and $B_{rL}(0, 4)$ in the same way as before. If this is the case, we declare $(0, 2)$ open, otherwise we call it closed. Note that there is one last 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 $(0, 2)$. This is easy though: since we have bounded the number of annuli involved in each step of the procedure, there is a uniform upper bound on the number of annuli that 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 renormalised percolation process dominates independent site percolation with parameter $p_c + \epsilon$, and is therefore supercritical. Finally, it is easy to see that if the renormalised process percolates, so does the underlying percolation process, proving the result. \square

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.

References

- [1] F. Baccelli and B. Błaszczyszyn (2001). On a coverage process ranging from the boolean model to the Poisson Voronoi tassellation with application to wireless communications. *Advances in Applied Probability*, **33**(2), 293-323.
- [2] P. Balister, B. Bollobás, and M. Walters (2004). Continuum percolation with steps in an annulus. *Annals of applied probability, to appear*.
- [3] 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), 722-741.

- [4] O. Dousse, F. Baccelli, P. Thiran (2003). Impact of interferences on connectivity in ad-hoc networks. *IEEE/ACM Transactions on Networking*, to appear.
- [5] D. Dubashi, O. Häggström, and A. Panconesi (2004). Connectivity properties of bluetooth wireless networks. *Submitted*. <http://www.math.chalmers.se/~olleh/papers.html>
- [6] E. N. Gilbert (1961). Random plane networks. *Journal of SIAM* **9**, 533-543.
- [7] O. Häggström and R. Meester (1996). Nearest Neighbour and Hard Sphere Models in Continuum Percolation. *Random Structures and Algorithms* **9**, 295-315.
- [8] R. Meester and R. Roy (1996). Continuum percolation. *Cambridge University Press*.
- [9] R. Meester and M. D. Penrose and A. Sarkar (1997). The random connection model in high dimensions. *Statistics and Probability Letters* **35**, 145-153.
- [10] M. D. Penrose (1993). On the spread-out limit for bond and continuum percolation. *Annals of Applied Probability* **3**(1), 253-276.
- [11] T. K. Philips, S.S. Panwar, and A. N. Tantawi (1989). Connectivity properties of a packet radio network model. *IEEE Transactions on Information Theory* **35**(5), 1044-1047.
- [12] 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.
- [13] J. Quintanilla, and S. Torquato (1997). Clustering in a continuum percolation model. *Advances in applied probability*, **29**, 327-336.
- [14] G. Stell.(1996). Continuum theory of percolation. *J. Phys.: Condens. Matter* **8**(25A), A1-A17.
- [15] F. Xue and P. R. Kumar (2004). The number of neighbors needed for connectivity of wireless networks. *Wireless networks*, to appear.
- [16] A. Wagner and V. Anantharam (2002). Designing a Contact Process. The Piecewise-Homogeneous Process on a Finite Set with Applications. *Stochastic processes and their applications, submitted for publication*. <http://www.eecs.berkeley.edu/~ananth>