Continuum percolation with holes

A. Sarkar[a,*], M. Haenggi[b]

[a] *Western Washington University, Bellingham WA 98225, USA*
[b] *University of Notre Dame, Notre Dame IN 46556, USA*

Abstract

We analyze a mathematical model of a cognitive radio network introduced in Yemini et al. (2016). Our analysis reveals several surprising features of the model. We explain some of these features using ideas from percolation theory and stochastic geometry.

Keywords: percolation, Poisson process, Gilbert model, cognitive networks, simultaneous connectivity

1. Introduction

Percolation on the standard disc graph (Gilbert's disc model) has been a well-studied topic since the seminal work of Gilbert (1961). It has applications in wireless ad hoc or sensor networks (Haenggi, 2012), where it is assumed that the network is composed of a single class of transceivers with a fixed transmission radius. In an important emerging class of networks, the so-called cognitive networks, however, there exist two classes of transceivers, where the so-called secondary users are only allowed to be active if they are not too close to any of the primary users (Lee and Haenggi, 2012). In these networks, the primary users are allowed unrestricted access to their licensed radio spectrum, while the secondary users are prohibited from causing harmful interference to the primary users, i.e., they need to respect a guard zone around the primary users.

We focus on percolation in the network formed by the secondary users. Assuming that primary and secondary users form independent Poisson point processes, the subset of secondary users who are allowed to be active is a *Poisson hole process*, since the guard zones around the primary users create holes in the point process of active secondary users. This point process was introduced in Lee and Haenggi (2012) and further studied in Yazdanshenasan et al. (2016).

The problem of joint percolation in both the primary and secondary networks was proposed and studied in Yemini et al. (2016). Our main contribution in this paper is threefold. First, we introduce a re-parametrization of the problem, reducing the number of parameters from the five in (Yemini et al., 2016) to

*Corresponding author
Email addresses: amites.sarkar@wwu.edu (A. Sarkar), mhaenggi@nd.edu (M. Haenggi)

three. This enables us to summarize the behavior of the full model in a single plot, from which one can easily read off information about the original model. Second, we present simulation results on the critical radius for the existence of a left-right crossing, which approximates the critical radius for percolation; we also indicate the necessary steps towards a better approximation of the latter parameter. Finally, these simulation results, shown in Figure 1, suggest several mathematical results on the dependence of the critical radius on the other parameters of the model. We state these results, together with sketches of some of their proofs, in the last main section. Rigorous proofs will appear in a forthcoming paper.

2. Mathematical background

We will use several facts and methods from continuum percolation. Most of these relate to the *Gilbert graph*, which was first defined and studied in Gilbert (1961), and we include a brief description of that model here. We also provide a short explanation of the methods that have been used to study the model, some of which were used in (Yemini et al., 2016), and some of which we will use ourselves. For more information, and rigorous proofs, the reader is encouraged to consult the books by Meester and Roy (1996), Penrose (2003), Bollobás and Riordan (2006), Haenggi (2012), and the survey article by Balister et al. (2009).

In Gilbert's model, we start with a Poisson process $\mathcal{P}$ of intensity one in the plane. These points form the vertices of an infinite graph G_r. The edges of G_r are obtained by joining two points of $\mathcal{P}$ if they lie at (Euclidean) distance less than r, where r is a fixed parameter.

The main quantity of interest for the Gilbert model is the *critical radius for percolation*. To define this, imagine fixing $\mathcal{P}$ and slowly increasing r, starting from $r = 0$. Initially, the graph G_r will consist of small components, whose vertices happen to lie close together, and isolated vertices. (Here, we use standard graph-theoretic terminology, so that a *component*, by definition, means a connected component.) As we increase r, these components will grow and merge, and at some point an infinite component I, containing a positive fraction $\theta(r) > 0$ of all the vertices in G_r, will appear. When this happens, we say that *percolation* occurs, or that the model G_r *percolates*. The fraction $\theta(r)$ of vertices in I can also be interpreted as the probability that a fixed vertex of G_r belongs to I, and, as r increases, $\theta(r)$ will naturally increase towards 1.

From a rigorous mathematical perspective, Kolmogorov's 0-1 law on tail events implies that, for any fixed value of r, the probability that G_r percolates (and also $\theta(r) > 0$) is either zero or one. In other words, if we consider several different instances of $\mathcal{P}$, and simultaneously increase r in each of them, at the same rate, then percolation occurs *at the same time* in each instance. Consequently, if we define r_{crit} as

$$r_{\text{crit}} = \sup\{r : \theta(r) = 0\},$$

then, for $r < r_{\text{crit}}$, G_r does not percolate (almost surely), and, for $r > r_{\text{crit}}$, G_r

percolates (again, almost surely, i.e., with probability 1). As it happens, when $r = r_{\text{crit}}$, G_r does not percolate; this was established in Alexander (1996).

Given this, the next step is to obtain good bounds on r_{crit}. Currently the best known rigorous bounds, due to Hall (1985), are

$$0.833 < r_{\text{crit}} < 1.836.$$

These bounds are only slight improvements of Gilbert's original bounds from 1961, and were obtained using refinements of Gilbert's original methods. The lower bounds were obtained using branching processes, while the upper bounds come from comparison with a lattice percolation model, specifically, face percolation on a hexagonal lattice. More recently, Balister et al. (2005) used dependent percolation to show that, with confidence 99.99%,

$$1.1978 < r_{\text{crit}} < 1.1989.$$

(In detail, Balister et al. showed that, subject to a certain bound on a certain multidimensional integral, the stated bounds on r_{crit} hold; the integral itself was estimated using Monte Carlo methods, resulting in the stated confidence level.)

For more complicated models, such as the secrecy graph model (Sarkar and Haenggi, 2013), and the model considered in Yemini et al. (2016), comparison with (dependent or independent) lattice percolation remains the main tool for bounding the various thresholds (indeed, it is used extensively in (Yemini et al., 2016)). These comparisons work by superimposing an appropriately-sized lattice on the plane, and declaring a face F of the lattice "open" if F contains a point of $\mathcal{P}$. If the lattice spacing has been chosen correctly, then face percolation in the lattice implies percolation in the original model. Therefore, we can use classical bounds for lattice percolation thresholds to deduce that percolation occurs in the original model, for certain parameter values. The method can also be used to show that, for certain other parameter values, percolation does not occur; occasionally, one has to use dependent percolation to make the comparisons work, and this usually results in very weak bounds. Recent innovations include the *rolling ball method* of Balister and Bollobás (2016), and the *high confidence* method introduced in Balister et al. (2005), referred to above. Both these newer methods can also be adapted to other models; for instance they were used in Sarkar and Haenggi (2013) to study the secrecy graph.

The Gilbert model is primarily a model of a random geometric *graph*. However, there is a related *coverage process*, which we will make heavy use of in this paper. To define this coverage process $\mathcal{C}_r$, known as the *Gilbert disc model*, we start with a unit-intensity Poisson process $\mathcal{P}$ as before, but this time we place an open disc $B(p, r)$ of radius r around each point $p \in \mathcal{P}$. The connection between the Gilbert disc model and the Gilbert graph G_r is that graph-theoretic components in G_r correspond exactly to topological components of $\mathcal{C}_{r/2}$. If $\mathcal{C}_{r/2}$ has an infinite (topological) component, we extend our earlier terminology by saying that $\mathcal{C}_{r/2}$ percolates, which, by the above, occurs if and only if G_r percolates as well.

There are several quantities related to the Gilbert disc model which can be conveniently expressed in terms of the connection radius r. First, there is the *average coverage level* $\alpha_r = \pi r^2$, which represents both the average number of times a point of $\mathbb{R}^2$ is covered by $\mathcal{C}_r$ and also the average degree in G_r. Then there is the *reduced coverage level* $\alpha'_r = \alpha_{r/2} = \alpha_r/4$, representing the average coverage level of the "reduced" Gilbert model $\mathcal{C}_{r/2}$. Finally, there is the *covered area fraction* β_r, defined to be the proportion of $\mathbb{R}^2$ which is covered by $\mathcal{C}_r$, and which can also be interpreted as the probability that a fixed point $x \in \mathbb{R}^2$ is covered by $\mathcal{C}_r$. Since x is not covered by $\mathcal{C}_r$ if and only if no $p \in \mathcal{P}$ lies within distance r of x, which happens with probability $e^{-\pi r^2}$, we see that $\beta_r = 1 - e^{-\pi r^2}$. (There is also the covered area fraction for the reduced model $\beta_{r/2}$, but we will not use this parameter in our analysis.) To summarize:

$$\text{Reduced coverage level} = \alpha'_r = \frac{1}{4}\pi r^2$$
$$\text{Covered area fraction} = \beta_r = 1 - e^{-\pi r^2}$$

The main facts we will use about $\mathcal{C}_r$ are the precise results from Balister et al. (2010) on the distribution of the regions left uncovered by $\mathcal{C}_r$ when r is large. Since these results are somewhat technical, we will postpone detailed discussion of them until we need them (in Section 7).

3. The work of Yemini et al.

Yemini et al. (2016) consider the following five-parameter model of two interlinked random geometric graphs, which they term the *heterogeneous* model. In their model, there are two networks, a *primary* network, and a *secondary* one. The primary network consists of *primary nodes*, distributed according to a Poisson process $\mathcal{P}_1$ of intensity λ_p in the plane; two nodes x and y of the primary network are connected if they lie at Euclidean distance at most (or, equivalently, less than) D_t. The secondary network consists of *secondary nodes*, distributed according to a Poisson process $\mathcal{P}_2$ of intensity λ_s in the plane, where $\mathcal{P}_1$ and $\mathcal{P}_2$ are independent point processes. Two nodes u and v of the secondary network are connected if they lie at Euclidean distance at most d_t from each other, and if *also* there is no primary node within distance D_f of either u or v. The parameters D_t and d_t are the *transmission ranges* of the primary and secondary networks respectively, while D_f is the radius of the *guard zone* of the primary nodes.

This is a complicated and very general model, so one should perhaps not expect a complete analysis of every aspect of it. The first result in (Yemini et al., 2016) concerns the region of the parameter space for which both networks percolate, which the authors call the *simultaneous connectivity region*. The result is that this region forms a connected subset of $\mathbb{R}^5$. After that, the authors modify an argument from Meester and Roy (1996) (who had in turn modified the classic "trifurcation" argument in Burton and Keane (1989)) to show that, with probability one, there is at most one infinite component in the primary

network, and at most one infinite component in the secondary network. Then, they use a "Peierls" (circuit-counting) argument to show that, for all fixed D_t and d_t, and for λ_p and λ_s above the respective separate percolation thresholds, there exists $D_f > 0$ such that both networks percolate; intuitively this follows by taking D_f sufficiently small so that the primary network has a negligible effect on the secondary one. Finally, they use comparison with site percolation to derive necessary conditions and (much stronger) sufficient conditions for both networks to percolate.

4. Re-parametrization

Our first remark is that, for percolation to occur in the primary network, it is necessary and sufficient that

$$D_t > {\lambda_p}^{-1/2} r_{\text{crit}}. \tag{1}$$

This condition, which appears in the form $\lambda_p > D_t^{-2}\lambda_c(1)$ in each theorem of Sections 5 and 6 of Yemini et al. (2016), is separate from the rest of the model. In other words, to check whether percolation occurs in the primary network, we just check whether (1) is satisfied, and this just depends on D_t and λ_p. Consequently, we can ignore the parameter D_t, as well as the issue of percolation in the primary network. What is at stake is percolation in the secondary network.

Second, there is no loss of generality in assuming that $\lambda_p = 1$. This is because the model with parameters

$$(\lambda_p, \lambda_s, D_t, d_t, D_f) \tag{2}$$

can be re-scaled to have parameters

$$(1, \lambda_s/\lambda_p, D_t\sqrt{\lambda_p}, d_t\sqrt{\lambda_p}, D_f\sqrt{\lambda_p}). \tag{3}$$

In detail, given the model with the parameters in (2), we can magnify the plane $\mathbb{R}^2$ by a linear factor of $\sqrt{\lambda_p}$ in both directions. This has the effect of dividing both intensities by λ_p and multiplying all distances (including the transmission ranges and guard zone radii) by the factor $\sqrt{\lambda_p}$, yielding the parameters in (3). This special case therefore captures all the essentials of the original model and is the only case we need consider.

From now on, we will write $\lambda = \lambda_s, d = d_t$ and $D = D_f$ for the remaining three parameters. The model can be described as follows. The primary nodes land in the plane, and each primary node creates a "hole" of radius D around itself. These holes cover a fraction

$$\beta_D = 1 - e^{-\pi D^2}$$

of the plane. Then the secondary nodes land. Although these have intensity λ, only the ones landing outside the holes can transmit. Therefore, the average

intensity of *active* secondary nodes, i.e., secondary nodes which can actually transmit, is given by

$$\lambda' = \lambda(1 - \beta_D) = \lambda e^{-\pi D^2}.$$

The active secondary nodes do not form a Poisson process but a Poisson hole process as introduced in Lee and Haenggi (2012). Our analysis will be in terms of the parameter

$$\alpha_{\text{sec}}(d, D, \lambda) = \lambda' \alpha'_d = \frac{1}{4}\lambda\pi d^2 e^{-\pi D^2},$$

which represents the average coverage level when discs of radius $d/2$ are placed around each active secondary node. (This corresponds to the reduced coverage level introduced earlier, except that the underlying point process is no longer Poisson.) Specifically, if d_{crit} is the critical radius for percolation in the model (whose existence follows from Kolmogorov's 0-1 law), then we write

$$\alpha_\lambda(\beta) = \alpha_{\text{crit}}(D, \lambda) = \alpha_{\text{sec}}(d_{\text{crit}}, D, \lambda) = \frac{1}{4}\lambda\pi d_{\text{crit}}^2 e^{-\pi D^2} = \frac{1}{4}\lambda\pi d_{\text{crit}}^2(1 - \beta).$$

Our goal will be to understand the dependence of $\alpha = \alpha_{\text{crit}}(D, \lambda)$ on $\beta = \beta_D$, for different values of λ. In other words, we will essentially consider α as a function $\alpha_\lambda(\beta)$ of β. For fixed λ, the original formulation of the problem asked for d_{crit} as a function of D. Our formulation is equivalent, since β can be calculated from D using the equation $\beta = 1 - e^{-\pi D^2}$, and d_{crit} can be calculated from α, β and λ using the equation displayed above. Thus, knowing the dependence of α on β also indicates the dependence of d_{crit} on D, for a fixed value of λ. In addition, our reformulation allows us to observe and explain some new phenomena.

First we present some simulation results, followed by some theorems suggested by these results.

5. Simulation method

In the simulation, we focus on left-right (LR) crossings in a large square region of the network. For given values of λ and β (or D), we adjust the parameter α until there exists an LR crossing in between 45 and 55 out of 100 realizations of the model. This experiment is repeated 50 times, and the average of the 50 resulting values of α is taken. Overall, this simulation procedure provides a good trade-off between accuracy and simulation time.

For the initial guess for d_{crit} (and thus α), we use

$$d_{\text{init}} = \frac{r_{\text{crit}}}{\sqrt{\lambda}} e^{\pi D^2/2}(1 + \lambda D^2/12).$$

Depending on the fraction of LR crossings found in 100 network realizations, the initial guess for d_{crit} is adjusted up or down, until it falls in the 45-55% range. The network is simulated on a square of side length $2s$, where $s = \max\{24, 10d_{\text{init}}\}$.

This simulation process is relatively efficient for most of the parameter space but gets time-consuming as $\lambda \to 0$ and $\beta \to 1$ (for $D > 1.3$). Although the

existence of an LR crossing does not guarantee percolation, the resulting values for α give some insight into the behavior of $\alpha_\lambda(\beta)$ and also provide supporting evidence for various conjectures on its shape. In principle, we can get arbitrarily good approximations to the actual curves using the high-confidence method of Balister et al. (2005), except that this would take a very long time, since the method necessitates a laborious computational process for each value of λ and β; in any case, we expect these results to strongly resemble those in Figure 1 below.

6. Results

The simulation results for $\lambda = 1/8, 1/4, 1/2, 1, 2$ and 4 are shown in Figure 1. For each value of λ, the percolation region lies above the curve. Note that, when $D = 0$, so that $\beta_D = 0$, we recover the original Gilbert model, for which

$$\alpha_{\text{crit}} := \alpha_\lambda(0) = \frac{1}{4}\pi r_{\text{crit}}^2 \approx 1.13.$$

Note also that when $\lambda = 0$ there is no secondary network, so that we cannot meaningfully talk about a critical radius d_{crit} in this case.

Several features are immediately apparent from Figure 1. Here are some of them.

Observation 1. For each λ, and every $\beta \in [0, 1)$, $\alpha_\lambda(\beta) \geq \alpha_{\text{crit}}$.

Observation 2. For each fixed β, $\alpha = \alpha_\lambda(\beta)$ is monotonically increasing with λ.

Observation 3. For each fixed λ, the curve $\alpha = \alpha_\lambda(\beta)$ is unimodal.
Also, as λ increases, the peak of $\alpha_\lambda(\beta)$ rises and moves further to the right.

The next two observations relate to the asymptotic behavior of the model and are, in fact, the only observations we can fully explain. The first (Observation 4) is relevant when the primary nodes are surrounded by large guard zones, so that most secondary nodes are inactive, while the second (Observation 5) is relevant when the intensity of secondary nodes is very low.

Observation 4. For each fixed λ, $\alpha_\lambda(\beta) \to \alpha_\lambda(0) = \alpha_{\text{crit}}$ as $\beta \to 1$.
(Note: it took several hours of computer time to confirm this.) Naturally, as $D \to \infty$ with λ and d fixed, the secondary model does not percolate. More precisely, from the observation, for fixed λ, $d_{\text{crit}}^2(1 - \beta) \to c(\lambda)$ as $\beta \to 1$ (equivalently, as $D \to \infty$), where $\frac{1}{4}\lambda\pi c(\lambda) = \alpha_{\text{crit}}$. Consequently, as $\beta \to 1$ with λ fixed, $d_{\text{crit}} \to \infty$.

Observation 5. We have $\alpha_\lambda(\beta) \to \alpha_{\text{crit}}$ uniformly as $\lambda \to 0$.
Naturally, if $\lambda \to 0$ with D and d fixed, the secondary model does not percolate. More precisely, from the observation, for fixed β, $\lambda d_{\text{crit}}^2 \to c(\beta)$ as $\lambda \to 0$, where $\frac{1}{4}\pi c(\beta)(1 - \beta) = \alpha_{\text{crit}}$. Consequently, as $\lambda \to 0$ with β fixed, $d_{\text{crit}} \to \infty$.

We can prove some, but not all, of these observations. In the next section, we explain as many of them as we can.

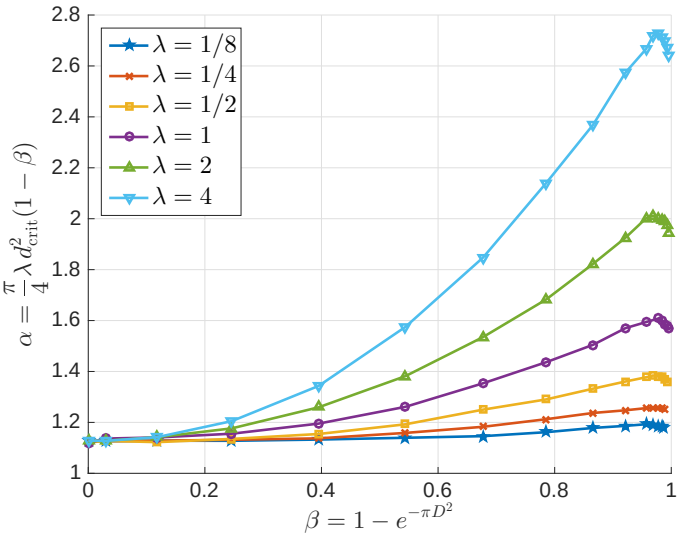

Figure 1: Critical average coverage levels $\alpha_\lambda(\beta)$ for different values of λ.

7. Analysis

For Observations 1 and 2, we have only heuristic explanations. First, fix $\beta \in [0,1)$ and $\lambda > 0$. The holes $B(p,D)$ around each point $p \in \mathcal{P}_1$ of the primary network constrain the locations of the active secondary nodes $q \in \mathcal{P}_2 \setminus \bigcup_{p\in\mathcal{P}_1} B(p,D)$, causing the discs $B(q,d/2)$ to overlap more than in the basic Gilbert model, without aiding percolation. This overlap, in turn, increases the average coverage level required for percolation, explaining Observation 1. The overlap effect will be more pronounced for large λ, explaining Observation 2.

An extreme illustration of this phenomenon is provided when $-\lambda(\log(1-\beta)^{-1}) \to c$. (This case is considered in detail in Theorem 1 below.) In this scenario, the active secondary nodes are clustered together in the tiny regions left vacant by the primary node guard zones, with, on average, c active secondary nodes per vacant region. The discs around the secondary nodes in each tiny cluster will almost exactly coincide, so that the critical value of α increases from α_{crit} by a factor of roughly, but not exactly, c. (The exact factor is given in Theorem 1.)

When λ is very small, the locations of the active secondary nodes still form an approximate Poisson process, regardless of the value of β, so that our model approximates the basic Gilbert model (albeit scaled by a factor of $\lambda^{-1/2}$). This explains Observation 5.

For Observations 3 and 4, we need to describe the geometry of the union of discs $\bigcup_{p\in\mathcal{P}_1} B(p,D)$ in some detail, following the approach outlined in Section 2.3 of (Balister et al., 2009). (The technical details are contained in Section 4 of (Balister et al., 2010)). The idea is to consider the boundaries $\partial B(p,D)$ of the discs $B(p,D)$, rather than the discs themselves. Consider a fixed disc

boundary $\partial B(p, D)$. This boundary intersects the boundaries $\partial B(p', D)$ of all discs $B(p', D)$ whose centers p' lie at distance less than $2D$ from p. There is an expected number $4\pi D^2$ of such points $p' \in \mathcal{P}_1$. Each such p' contributes two intersection points $\partial B(p, D) \cap \partial B(p', D)$, and each intersection is counted twice (once from p and once from p'). Therefore we expect $4\pi D^2$ intersections of disc boundaries per primary node, and, since the intensity of the primary nodes (in our reformulation) is 1, we expect $4\pi D^2$ intersections of disc boundaries per unit area over the entire plane. Note that these intersections *do not* form a Poisson process, since they are constrained to lie on various circles.

The next step is to move from intersections to regions. The disc boundaries partition the plane into small "atomic" regions. Drawing all the disc boundaries in the plane yields an infinite plane graph, each of whose vertices (disc boundary intersections) has four curvilinear edges emanating from it. Each such edge has two vertices, so (by double counting) there are almost exactly twice as many edges as vertices in any large region R. It follows from Euler's formula for plane graphs (Bollobás, 1998) that the number of atomic regions in R is asymptotically the same as the number of intersection points in R. (Euler's formula states that, for a plane graph with V vertices, E edges and F faces (including the outer face), we have $V - E + F = 2$. Thus, when V, E and F all tend to infinity, and $V \sim 2E$, we must have $F \sim V$.)

An alternative argument is that each atomic region has a "leftmost" intersection point, and each intersection point is the leftmost point of exactly one atomic region, so that there is a one-to-one correspondence between intersections and regions. Moreover, each vertex borders four atomic regions, so that the average number of vertices bordering an atomic region is also four. Note that this last figure is just an average, and that many atomic regions will have less than, or more than, four vertices on their boundaries.

The third step is to return to the discs themselves and calculate the expected number of *uncovered* atomic regions per unit area. It is most convenient to calculate this in terms of uncovered intersection points. A fixed intersection point is uncovered by $\bigcup_{p \in \mathcal{P}_1} B(p, D)$ with probability $e^{-\pi D^2}$ (using the independence of the Poisson process), so we expect $4\pi D^2 e^{-\pi D^2}$ uncovered intersections, and so $\pi D^2 e^{-\pi D^2}$ uncovered regions, per unit area in R. (It turns out that the average number of sides of an *uncovered* region is also four.) Consequently, for large D, these uncovered regions are rare, and, moreover, their distribution in the plane is approximately Poisson. This last assertion can be made precise, and proved using the Chen-Stein method of Poisson approximation (see (Balister et al., 2010) for details).

How large are these uncovered atomic regions? To answer this, recall that the uncovered area fraction in $\mathbb{R}^2$ is $e^{-\pi D^2}$. Since the uncovered regions form an approximate Poisson process of intensity $\pi D^2 e^{-\pi D^2}$, the expected area of each uncovered region is $(\pi D^2)^{-1} = (-\log(1-\beta))^{-1}$. Again, this is just an average; some uncovered regions will be much larger than this, and others will be much smaller.

Having obtained the expected size ($(-\log(1-\beta))^{-1}$) and the approximate

spatial distribution (Poisson, with intensity $-(1-\beta)\log(1-\beta)$) of the uncovered regions when β is close to 1, we turn to more detailed results on the sizes of these regions. It is known (Calka et al., 2010) that the scaled area distribution of the uncovered regions in a high-intensity (i.e., $\beta \to 1$) Boolean model converges in law to the area distribution A of the typical cell of a *Poisson line process*. (The Poisson line process $\mathcal{P}_l$ is somewhat analogous to the Poisson point process. It is spatially homogeneous (stationary) and can be defined (Miles, 1964) using a one-dimensional Poisson process X and, for each $x \in X$, a uniform random variable Y_x on $[0, \pi)$ as follows: $x \in X$ gives the (signed) distance from a line $l \in \mathcal{P}_l$ to the origin, and Y_x gives the orientation of l. All Y_x are independent.) The exact distribution of A is unknown, but, normalizing so that $\mathbb{E}(A) = 1$, the second moment $\mathbb{E}(A^2) = \frac{1}{2}\pi^2$ and the third moment $\mathbb{E}(A^3) = \frac{4}{7}\pi^4$ were both obtained as early as 1945 by S.A. Goudsmit and D.G. Kendall respectively. Our analysis will in be terms of the moment generating function $M_A(t) = \mathbb{E}(e^{tA})$ of this normalized area A.

In the following theorem, we let $\beta \to 1$ with $\lambda = \lambda(\beta)$ depending on β. In other words, we study the limit of $\alpha_\lambda(\beta)$ as $\beta \to 1$ along a specified curve in the (β, λ) plane. In the proof, we use (almost) standard asymptotic notation, so that $f(x) \sim g(x)$ as $x \to c$ means $f(x)/g(x) \to 1$ as $x \to c$.

Theorem 1. *Suppose that $\beta \to 1$ with $\lambda(\beta) = -c\log(1-\beta)$, where $c > 0$ is a fixed constant. Then*

$$\alpha_{\lambda(\beta)}(\beta) \to \frac{c\alpha_{\text{crit}}}{1 - M_A(-c)}.$$

Proof. (Sketch) Suppose that the hypotheses of the theorem hold, and note that the disc radius D (corresponding to the covered area fraction β) satisfies $\pi D^2 = -\log(1-\beta)$. Write U for the random variable representing the number of active secondary nodes in an uncovered atomic region, and V for the random variable representing the area of an uncovered atomic region. Then, by the tower law of conditional expectation and basic properties of the Poisson process,

$$\mathbb{E}(U) = \mathbb{E}(\mathbb{E}(U|V)) = \mathbb{E}(\lambda(\beta)V) = \lambda(\beta)\mathbb{E}(V) = \frac{\lambda(\beta)}{\pi D^2} = \frac{-\lambda(\beta)}{\log(1-\beta)} = c.$$

We are interested in the uncovered regions which contain at least one active secondary node. The intensity $I^a(\beta)$ of these "active" uncovered regions is the intensity $I(\beta)$ of the uncovered regions multiplied by the probability $\mathbb{P}(U \neq 0)$ that the typical uncovered region contains at least one secondary node. This latter probability can be obtained in terms of the moment-generating function of A. Indeed, using the fact that $-\log(1-\beta)V$ converges in law to A (defined above), and writing f_V and f_A for the densities of V and A respectively, we have

$$\begin{aligned} I^a(\beta) &= I(\beta)\mathbb{P}(U \neq 0) = I(\beta)\left(1 - \int_0^\infty f_V(t)e^{-t\lambda(\beta)}\,dt\right) \\ &\sim I(\beta)\left(1 - \int_0^\infty f_A(t)e^{t\lambda(\beta)(\log(1-\beta))^{-1}}\,dt\right) = I(\beta)\left(1 - \int_0^\infty f_A(t)e^{-ct}\,dt\right) \\ &= I(\beta)(1 - M_A(-c)) = -(1-\beta)\log(1-\beta)(1 - M_A(-c)). \end{aligned}$$

The active uncovered regions form an approximate Poisson process with intensity $I^a(\beta)$. Therefore, d_{crit} is asymptotically given by $d_{\text{crit}} = I^a(\beta)^{-1/2} r_{\text{crit}} \approx 1.2 I^a(\beta)^{-1/2}$. If each active uncovered region had exactly one secondary node, then $\alpha_\lambda(\beta)$ would simply be $\alpha_{\text{crit}} \approx 1.13$ as before, since the secondary nodes would essentially form a scaled version of the basic Boolean model. However, instead, we have an asymptotic expected number $\mathbb{E}(U)(1 - M_A(-c))^{-1} = c(1 - M_A(-c))^{-1}$ of secondary nodes in each active uncovered region; this is the expected value of U conditioned on $U \neq 0$. This means that the average "overlap" factor tends to $c(1 - M_A(-c))^{-1}$, so that $\alpha_{\lambda(\beta)}(\beta) \to \frac{c\alpha_{\text{crit}}}{1 - M_A(-c)}$, as in the statement of the theorem. □

From this, Observation 4 follows easily. Writing $c(\beta, \lambda) = -\lambda(\log(1-\beta))^{-1}$, we see that if λ is fixed then $c \to 0$ as $\beta \to 1$; since moreover $1 - M_A(-c) = c + O(c^2)$, we have $\alpha_\lambda(\beta) \to \alpha_{\text{crit}}$ in this case. Theorem 1 is also consistent with Observation 3. It is less clear what happens close to the maximum of $\alpha_\lambda(\beta)$ for each fixed λ, so that most of Observation 3 remains unexplained. We hope to return to this question in a future paper.

8. Conclusions and open problems

In this paper we have analyzed a mathematical model of a cognitive radio network with two classes of user. This model, proposed in Yemini et al. (2016), is related to the *Poisson hole process* introduced in Lee and Haenggi (2012). We have simplified the original model, while retaining every one of its essential features, and then studied the new model both computationally and analytically.

Our simulation results have revealed several surprising phenomena, some of which are easier to explain than others. Rigorous proofs of Theorem 1 and Observations 4 and 5 will appear in a forthcoming paper. For now, the main open problems are to prove Observations 1, 2 and 3. We hope that our work will stimulate further research on this problem.

References

M. Yemini, A. Somekh-Baruch, R. Cohen, A. Leshem, On simultaneous percolation with two disk types, arXiv:1601.04471v2, 2016.

M. Yemini, A. Somekh-Baruch, R. Cohen, A. Leshem, Simultaneous connectivity in heterogeneous cognitive radio networks, Proc. IEEE Int. Symp. Information Theory (ISIT16) (2016) 1262–1266.

E. Gilbert, Random plane networks, Journal of the Society for Industrial Applied Mathematics 9 (1961) 533–543.

M. Haenggi, Stochastic Geometry for Wireless Networks, Cambridge University Press, 2012.

C.-H. Lee, M. Haenggi, Interference and outage in Poisson cognitive networks, IEEE Transactions on Wireless Communications 11 (4) (2012) 1392–1401.

Z. Yazdanshenasan, H. Dhillon, M. Afshang, P. H. J. Chong, Poisson hole process: Theory and applications to wireless networks, arXiv:1601.01090v1, 2016.

R. Meester, R. Roy, Continuum Percolation, Cambridge University Press, 1996.

M. Penrose, Random Geometric Graphs, Oxford Studies in Probability, Oxford University Press, 2003.

B. Bollobás, O. Riordan, Percolation, Cambridge University Press, 2006.

P. Balister, B. Bollobás, A. Sarkar, Percolation, connectivity, coverage and colouring of random geometric graphs, in: Handbook of Large-Scale Random Networks, Springer, 117–142, 2009.

K. Alexander, The RSW theorem for continuum percolation and the CLT for Euclidean minimal spanning trees, Annals of Applied Probability 6 (1996) 466–494.

P. Hall, On continuum percolation, Annals of Probability 13 (1985) 1250–1266.

P. Balister, B. Bollobás, M. Walters, Continuum Percolation with steps in the square or the disc, Random Structures and Algorithms 26 (4) (2005) 392–403.

A. Sarkar, M. Haenggi, Percolation in the secrecy graph, Discrete Applied Mathematics 161 (13–14) (2013) 2120–2132.

P. Balister, B. Bollobás, Percolation in the k-nearest neighbour model, in: Recent Results in Designs and Graphs: a Tribute to Lucia Gionfriddo, Aracne, 83–100, 2016.

P. Balister, B. Bollobás, A. Sarkar, M. Walters, Sentry selection in wireless networks, Advances in Applied Probability 42 (1) (2010) 1–25.

R. Burton, M. Keane, Density and uniqueness in percolation, Comm. Math. Phys. 121 (1989) 501–505.

B. Bollobás, Modern Graph Theory, Cambridge University Press, 1998.

P. Calka, J. Michel, K. Paroux, Refined convergence for the Boolean model, Advances in Applied Probability 41 (4) (2009) 940–957.

R. Miles, Random polygons determined by random lines in a plane I, Proc. Nat. Acad. Sci. USA 52 (1964) 901–907.