Covering algorithms, continuum percolation and the geometry of wireless networks

Lorna Booth † Jehoshua Bruck ‡ Massimo Franceschetti ‡
Ronald Meester §

†*University of Utrecht* ‡*California Institute of Technology*
§*Free University of Amsterdam*

May 3, 2001

Abstract

Continuum percolation models where each point of a two-dimensional Poisson point process is the center of a disc of given (or random) radius r, have been extensively studied. In this paper, we consider the generalization in which a deterministic algorithm (given the points of the point process) places the discs on the plane, in such a way that

AMS 1991 subject classifications. Primary 60D05, 60K35, 82B26, 82B43, 94C99.

Key words and phrases. Covering algorithms, (continuum) percolation, wireless communication networks, phase transition.

each disc covers at least one point of the point process and that each point is covered by at least one disc. This gives a model for wireless communication networks, which was the original motivation to study this class of problems.

We look at the percolation properties of this generalized model, showing the almost sure non-existence of an unbounded connected component of discs for small values of of the density λ of the Poisson point process, for any covering algorithm. In general, it turns out not to be true that unbounded connected components arise when λ is taken sufficiently high. However, we identify some large families of covering algorithms, for which such an unbounded component does arise for large values of λ.

We show how a simple scaling operation can change the percolation properties of the model, leading to the almost sure existence of an unbounded connected component for large values of λ, for any covering algorithm.

Finally, we show that a large class of covering algorithms, that arise in many practical applications, can get arbitrarily close to achieving a minimal density of covering discs. We also show (constructively) the existence of algorithms that achieve this minimal density.

1 Introduction and motivation

Geometric covering algorithms have been extensively studied in the last 20 years, in the context of computational geometry and combinatorial optimization (see the survey by Agarwal and Sharir (1998)). More recently, distributed versions of these algorithms have been proposed in the context of wireless network architectures (see Gerla and Tsai 1995).

Continuum percolation models (also referred to in the literature as Pois-

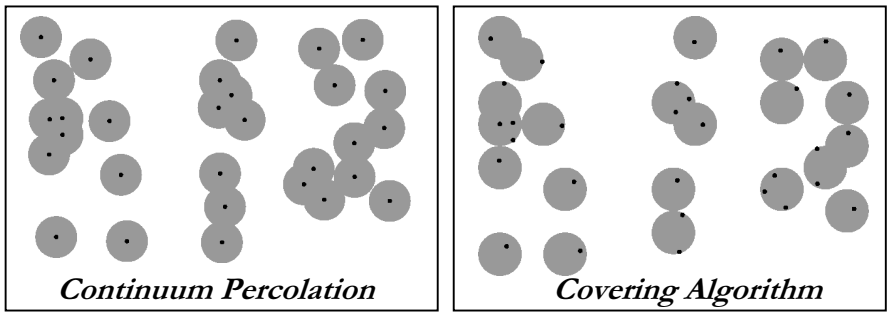

Figure 1: **Two different disc coverings of a random point process.** *A continuum percolation model places a disc at each random point; in our generalized model a covering algorithm places (possibly fewer) discs to cover all the points, according to a deterministic rule. We are interested in the almost sure existence of an unbounded connected component of discs, for a given density of points λ.*

son Boolean models) were introduced by Gilbert (1961) to model wireless networks of radio transmitting stations, and they have been considerably extended by mathematicians since then. In these models discs of a given (or random) radius r are centered at each point of a two-dimensional Poisson process X. The a.s. (almost sure) existence of unbounded connected components of discs, for a given density λ of the point process is often considered.

We consider the generalization in which a deterministic algorithm (given the points of the point process) places the discs on the plane, in such a way that each disc covers at least one point of X, and each point is covered by at least one disc (see Figure 1 for a visual example).

Our aim is twofold: on one side we explore the mathematics of the new model and we answer some very natural questions that arise from a pure mathematical point of view. On the other side, we note that most of our results can be applied to rigorously model geometric properties of wireless communication networks.

Random graphs are the natural tool that is often used to model communication networks. In such graphs vertices represent communication endpoints and edges represent two-way channels. In the standard model of Erdős and Rényi (1959,1960,1961a,b), each pair of vertices has some probability (the same for all pairs of vertices, regardless of their separation) of being joined by an edge. Therefore, for each natural number n, there is a a probability space consisting of all graphs with exactly n vertices. Erdős and Rényi proved that many interesting properties of random graphs occur a.s. as $n \to \infty$. Their model of random graphs, however, is not suited to accurately represent networks of short-range radio transmitting stations. This motivated Gilbert (1961) to propose an alternative model in which the range of the transmitters is a parameter. In his paper, he constructed a random network by considering a two-dimensional Poisson point process and joining each pair of points by an edge, if discs of radius r centered at those two points intersect. He was the first to introduce the concept of continuum percolation, identifying a phase transition behavior, i.e., the existence of a critical value λ_c for the density of the Poisson point process, at which an unbounded connected graph a.s. forms and the network can provide some long distance communication. His results were later extended, from a pure mathematical standpoint, by Hall (1985), Menshikov (1986), Roy (1990), Meester and Roy (1994), and others, leading to a theory of random coverage processes (see the books by Hall (1988) and Meester and Roy (1996)). In a more applied framework, Gupta and Kumar (1998, 2000) recently used a similar model to determine the throughput capacity of a wireless network.

Our model of communication refines the one introduced by Gilbert, considering a wireless backbone that routes data packets through the network. In our model, we differentiate between base stations and clients: clients com-

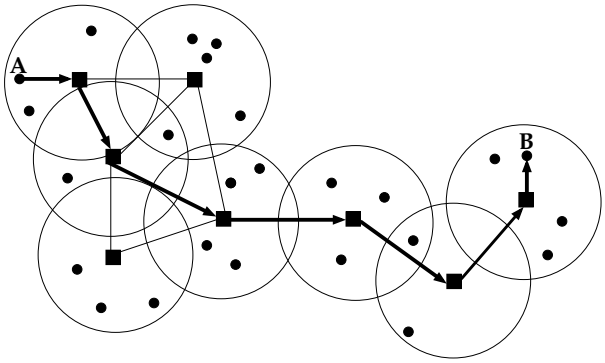

Figure 2: **The Wireless Backbone.** *A connected component of discs forms the wireless backbone of the network. If a client A wants to communicate to a client B, it connects to the closest base station and its message is routed through the backbone in a multi-hop fashion, until it reaches client B.*

municate between each other by connecting to base stations that forward their messages to their destinations (see Figure 2). If a client is within a given distance to a base station, it can connect to it and we say that the client is covered by the base station. Hence, the network appears as a set of circular cells (base stations broadcast domains) that cover a set of points (clients). A covering algorithm decides where to place the cells, according to the distribution of the clients. The algorithm can be a distributed, self-organizing one, in a model where the entire population of clients elects 'cluster heads' and divides itself into subsets that are covered by the cluster heads (see Gerla and Tsai 1995); or a more centralized one, in a model where the clients are mobile and the base stations are static. In the latter case, the base stations could be laid on a fixed grid and the covering algorithm could determine the subset of them that need to be turned on, at any given time, to provide coverage (see Franceschetti, Cook, and Bruck 2001). The algorithm would typically try to minimize the number of base stations that

need to be turned on to cover all the clients, or, in the dynamic case, would try to minimize the base stations movement needed to cover all the clients.

As in Gilbert's model, we assume a completely wireless network, hence, base stations can connect to other base stations only up to a limited distance. We assume two base stations to be connected only if the corresponding discs overlap (we will modify this requirement later in the paper). Therefore, if two clients are in the same connected component of overlapping discs, they can communicate, because they are reached by a connected path of base stations of that component.

In percolation theory one is interested in unbounded connected components. In our setting, unbounded connected components are of interest because they represent the almost fully connected state of the wireless network. If there exists a.s. an unbounded connected component, then the network is largely connected, since most of the clients fall inside the unbounded connected component and are able to communicate. In this case, few additional 'bridge' stations can be added to connect isolated components to the unbounded one and achieve the full connectivity of the network.

We informally summarize our main results in the next section; in Section 3 we introduce some notation and definitions; Sections 4 is devoted to existence and non-existence results for unbounded components, for different covering algorithms; Section 5 considers the effect of varying the base stations communication radius; Section 6 discusses existence of optimal covering algorithms; Section 7 contains some open problems.

2 Summary of results

Our results can be grouped into four categories: non-existence results, existence results for different classes of covering algorithms, scaling results, and results concerning the optimality of certain algorithms. In the following, we let X be a two-dimensional Poisson point process of density λ. The points of X represent the clients that are covered by base stations.

Non-existence results. Our first results regard the non-existence of an unbounded connected component of covering discs. We show (Theorem 4.1) that for any algorithm covering all the points of X by discs of radius r, there exists a $\lambda_0 > 0$ such that for all $0 < \lambda \leq \lambda_0$, P_λ(there is an infinite component) $= 0$. Then we show that the symmetric result, i.e., the a.s. existence of an unbounded connected component for large values of λ, is not generally true, but depends on the type of covering algorithm. It is known that a covering that places a disc centered at each point of X forms a.s. an unbounded connected component for large values of λ (Gilbert 1961). In order to show that this result does not generalize to all coverings, we specify a covering algorithm that does not form an unbounded connected component for any value of λ.

Existence results. We proceed by identifying different families of covering algorithms that form an unbounded connected component a.s. for large values of λ. One of such coverings, that is practical for our applications, is a *grid covering*. We show (Theorem 4.3) that for any algorithm covering all the points of X by discs of radius r centered at the vertices of a grid, there exists a $\lambda_1 < \infty$, such that for all $\lambda > \lambda_1$, P_λ(there is an infinite connected component) $= 1$.

Another large family of coverings that we consider are the *finite horizon coverings*. In these coverings, the position of the covering discs in a certain region does not depend on the points of the point process that are further away than some fixed distance (the horizon) of that region (a precise definition follows later). We show (Theorem 4.4) that for any finite horizon covering which satisfies a very natural condition, there exists a $\lambda_1 < \infty$, such that for all $\lambda > \lambda_1$, P_λ(there is an infinite connected component) $= 1$.

A third class of algorithms that we examine are the *shift invariant coverings*. These algorithms are defined by the requirement that the covering commutes with shifts of the points. We ask if any shift invariant algorithm forms a.s. an unbounded connected component for large values of λ, and we answer this question negatively, by constructing a shift invariant algorithm that does not exhibit this property.

Scaling results. We then introduce a further extension of our model. We note that when we consider overlapping discs as connected components, then we implicitly assume, in our model of a wireless network, that the maximum radius of communication between two base stations is twice as large as the maximum radius of communication between clients and base stations. This observation leads to the natural question of what would happen if the ratio between the two radii is different from two.

In the standard Poisson Boolean model, that places a disc centered at each point of a Poisson point process X, considering a different radius for connectivity corresponds to a simple scaling operation, hence it does not change the basic properties of the model. In our extended model, however, this leads to more interesting results. Call r the radius of the discs used to cover the points of X and R the maximum distance sufficient to connect

disc centers. We show (Theorem 5.1) that:

- If $R/r \leq 1$, then, for any grid G, there is a covering algorithm that places discs only at the vertices of G, and a.s. does not form an unbounded connected component, for any value of λ.
- If $1 < R/r < 2$, then, for some given dense grid G, there is a covering algorithm that places discs only at the vertices of G, and a.s. does not form an unbounded connected component, for any value of λ.
- If $R/r = 2$, then, for any grid G, any covering algorithm that places discs only at the vertices of G forms a.s. an unbounded connected component for large values of λ.
- If $R/r > 2$, then any algorithm forms a.s. an unbounded connected component for large values of λ, even if it is not grid-based.

Note that the latter case has a high practical value, because it states that if base stations can communicate at a distance larger than twice the maximum communication distance to the clients, an unbounded connected component forms a.s. for large values of the density of the clients, regardless of the covering algorithm used to build the cellular network.

Optimality Results. Finally, we show (constructively, in Theorem 6.3) the existence of algorithms that are optimal in achieving a minimal density of covering discs. We also show that a certain class of practical algorithms can achieve densities arbitrarily close to the optimal.

3 Notation and definitions

Let $\mathbb{R}^2$ be the Euclidean plane, let $\mathcal{B}^2$ be the σ-algebra of Borel sets in $\mathbb{R}^2$ and let $\ell(\cdot)$ be Lebesgue measure in $\mathbb{R}^2$. Let N be the collection of all counting measures on $(\mathbb{R}^2, \mathcal{B}^2)$, which assign finite measure to bounded Borel sets and for which the measure of a point is at most 1. In this way, N can be identified with the set of all configurations of points in $\mathbb{R}^2$, without limit points. Let $\mathcal{N}$ be the σ-algebra of N generated by sets of the form $\{\nu \in N : \nu(A) = k\}$, for all integers k and bounded Borel sets A. A (planar) point process X is defined as a measurable mapping from a probability space $(\Omega, \mathcal{F}, P)$ into $(N, \mathcal{N})$. For $A \in \mathcal{B}^2$, we denote by $X(A)$ the random number of points inside A. In this paper, X will always be a Poisson process with density $\lambda > 0$. We sometimes abuse notation and write $x \in \nu$, for $x \in \mathbb{R}^2$ and $\nu \in N$, to express that x is one of the points of ν.

We define a shift operation $T_t : \mathbb{R}^2 \to \mathbb{R}^2$ as a translation in $\mathbb{R}^2$ over the vector $t \in \mathbb{R}^2$, such that $T_t(x) = t + x$ for all $x \in \mathbb{R}^2$. The shift T_t induces in a natural way a shift transformation on N, which we also denote by T_t. Let, for all $x \in \mathbb{R}^2$ and $r \geq 0$, $D(x, r)$ be the disc of radius r centered at x: $D(x, r) = \{y \in \mathbb{R}^2 : |y - x| \leq r\}$. A circle of radius r centered at x is the set $\{y \in \mathbb{R}^2 : |y - x| = r\}$.

We call two discs D_i, D_j *adjacent* if $D_i \bigcap D_j \neq \emptyset$. We write $D_i \leftrightarrow D_j$ if there exists a sequence $D_{i_1}, D_{i_2}, \ldots, D_{i_k}$ of discs such that $D_{i_1} = D_i$, $D_{i_k} = D_j$, and D_{i_l} is adjacent to $D_{i_{l+1}}$ for $1 \leq l < k$. A *(connected) component* or *cluster* is a set $\{D_i : i \in J\}$ of discs which is maximal with the property that $D_i \leftrightarrow D_j$ for all $i, j \in J$. We identify a component with the set of centers of the discs in it.

We now formally define a covering algorithm: A *covering algorithm* $\mathcal{A}$

with discs of radius r, is a measurable mapping $\mathcal{A} : N \to N$ with the following properties:

1. for all $x \in \mathcal{A}(\nu)\, \exists y \in \nu : y \in D(x, r)$,
2. for all $y \in \nu\, \exists x \in \mathcal{A}(\nu) : y \in D(x, r)$.

We define the *occupied region* C of $\mathcal{A}(\nu)$ as the union $\bigcup_{x \in \mathcal{A}(\nu)} D(x, r)$.

In this paper, we examine different classes of covering algorithms, which we define as follows:

1. *Grid Algorithms.* Let $G \subset \mathbb{R}^2$ be the set of all vertices of a two dimensional lattice. A grid algorithm $\mathcal{A}$ constraints the covering discs to be centered at the vertices of G. That is, $x \in \mathcal{A}(\nu)$ implies $x \in G$.
2. *Finite Horizon Algorithms.* Let $B_n(x)$ be the box of size $n \times n$ centered at x, and let, for all $\nu \in N$, $\nu|_{B_n(x)}$ denote the restriction of ν to $B_n(x)$. In other words, $\nu|_{B_n(x)}$ can be identified with the set of points $\{\nu \cap B_n(x)\}$. We say that a covering algorithm $\mathcal{A}$ has finite horizon if there exists a constant $h \geq 0$ (the *horizon*), so that whenever $\nu|_{B_{n+2h}(x)} = \nu'|_{B_{n+2h}(x)}$, we have $\mathcal{A}(\nu)|_{B_n(x)} = \mathcal{A}(\nu')|_{B_n(x)}$, for all n and x. In words, this means that changing ν outside $B_{n+2h}(x)$ does not change the covering inside $B_n(x)$.
3. *Shift Invariant Algorithms.* A shift invariant algorithm $\mathcal{A}$ is defined by the property that $T_t(\mathcal{A}(\nu)) = \mathcal{A}(T_t(\nu))$, for all t. In words, this means that the covering algorithm commutes with shifts of the points.
4. *n-Square Algorithms.* An n-square algorithm is obtained as follows. Partition the plane into boxes of size $n \times n$. For each such box B_n, the covering of the points inside B_n should use the minimal number of

discs possible, it should be independent of point configurations outside B_n, and also commute with horizontal and vertical shifts of n. Note that such an algorithm is automatically of finite horizon.

Suppose now that we want to cover the points of X by the covering algorithm $\mathcal{A}$, that is, we consider the measurable map $\mathcal{A} \circ X : \Omega \to N$. This Boolean model is denoted by $(X, \mathcal{A}) = (X, \lambda, r, \mathcal{A})$, where λ is the density of X, and r the radius of the covering discs. The law of this process is denoted by $P_{\lambda,r}$. The standard Poisson Boolean model that places a disc of radius r, centered at each point of X is obtained when we take $\mathcal{A}$ to be the identity, and is denoted by (X, λ, r). In this model there exists $\lambda_c(r)$ such that for $\lambda \leq \lambda_c(r)$ we have no infinite cluster a.s., while for $\lambda > \lambda_c(r)$ there is an infinite cluster with probability 1. We often denote $\lambda_c(1)$ by λ_c and scaling implies that $\lambda_c(r) = \lambda_c(1)/r^2$ (see Meester and Roy (1996) for more details).

Next, we define the *density* of $(X, \mathcal{A})$. Let $N_{(X,\mathcal{A})}(n)$ be the (random) number of discs centered inside the box $B_n(0)$. The density of $(X, \mathcal{A})$ is given by

$$\lim_{n\to\infty} N_{(X,\mathcal{A})}(n)/n^2,$$

whenever this limit (i) exists a.s. and (ii) is an a.s. constant.

Finally, we introduce one more piece of terminology. If $(X, \mathcal{A})$ contains an unbounded component of discs with positive probability, we say that $(X, \mathcal{A})$ *percolates*.

4 Existence of percolation

In this long section we think of r as being fixed, while λ varies. Accordingly, we sometimes write $P_\lambda = P_{\lambda,r}$. We also use P to mean $P_{1,1}$. The expectation under P we denote by E.

Our first result deals with the non-existence of percolation for small values of λ.

Theorem 4.1 *For any covering algorithm $\mathcal{A}$, there exists a $\lambda_0(r) > 0$ such that for all $0 < \lambda \leq \lambda_0$, $(X, \lambda, r, \mathcal{A})$ does not percolate.*

Proof of Theorem 4.1. Assume that, with positive probability, there is an unbounded connected component of covering discs for $(X, \lambda, r, \mathcal{A})$. Then with positive probability, there is an unbounded connected component in the Poisson Boolean model $(X, \lambda, 2r)$. That is because two intersecting covering discs in $(X, \lambda, r, \mathcal{A})$ cover points that are at a distance of at most $4r$ to each other; and the Poisson Boolean model $(X, \lambda, 2r)$ places discs of radius $2r$ at each of the covered points. We then choose $\lambda_0 = \lambda_c/(2r)^2$, so that $(X, \lambda, 2r)$ does not form an unbounded connected component a.s. for $\lambda \leq \lambda_0$. □

A symmetric result to Theorem 4.1, i.e., the a.s. existence of percolation for large values of λ, depends on the type of covering algorithm used:

Proposition 4.2 *There exists a covering algorithm $\mathcal{A}$, such that for all λ, $(X, \lambda, r, \mathcal{A})$ does not percolate.*

Proof of Proposition 4.2. The proof is constructive. Draw circles of radii $\{3kr, k \in \mathbb{N}\}$ around the origin, and notice that a.s. no Poisson point falls on any of these circles. Then cover the Poisson points, with discs of radius r, without intersecting these circles. Notice that the circles divide the plane into finite annuli and, since each cluster of discs resides in at most one of these finite annuli, each cluster must be bounded, whatever the value of λ.

□

We now look at families of algorithms that do percolate for large values of λ. We first consider grid algorithms.

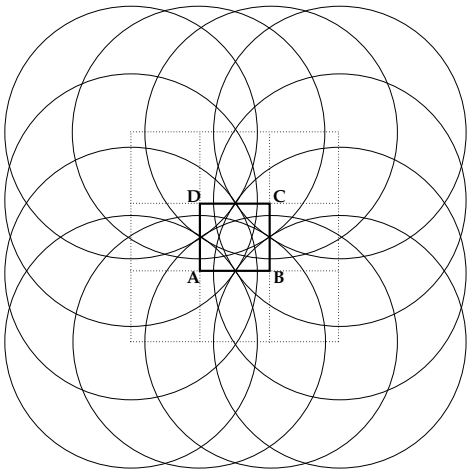

Figure 3: **Mapping to the site percolation model.** *The grid discs partition the square $ABCD$ into a finite number of small areas.*

Theorem 4.3 *For any grid covering algorithm $\mathcal{A}$, there exists a $\lambda_1 < \infty$, such that $(X, \lambda, r, \mathcal{A})$ percolates for all $\lambda > \lambda_1$.*

Proof of Theorem 4.3. The proof relies on a construction that maps the covering discs to a discrete site-percolation model. We illustrate the idea by considering a square lattice and a distance between two neighboring lattice vertices of one. Call a covering disc centered at a lattice vertex a grid disc. Clearly, the radius of a disc must be $r \geq \frac{\sqrt{2}}{2}$, in order to being able to cover all possible points on the plane by using only grid discs. For any $r \geq \frac{\sqrt{2}}{2}$, the number of grid discs that intersect a lattice square $ABCD$ is finite and it partitions the square into some number k_r of small regions A_i (see Figure 3). If at least one point of the Poisson point process falls into each region A_i, then the entire square $ABCD$ must be covered by grid discs. Now view each lattice square as a site of a site percolation model. Call the site occupied if there is at least one point of the Poisson process

situated inside each region A_i, for $i = 1 \ldots k_r$. Note that the occupancy of a site is independent of the occupancy of other sites and the probability of a site being occupied is given by: $p = \prod_{i=1}^{k_r}(1 - e^{-\lambda \ell(A_i)})$. Moreover, if two adjacent sites are both occupied, then the corresponding covering discs form a connected component. Thus, if there is an unbounded component of occupied adjacent sites, then there is an unbounded connected component of covering discs. Next, we choose λ large enough so that $p > p_c$, where p_c is the critical probability for site percolation on a square lattice. The a.s. existence of an unbounded connected component of covering discs immediately follows. □

Next, we consider finite horizon algorithms $\mathcal{A}$ with the property that the restriction of $\mathcal{A}$ to any box of size $n \times n$ contains at most $k = k(n)$ discs, for any value of λ. Note that this really is a weak requirement, since we can *completely cover* the box using at most $\alpha\lceil(n/r)\rceil^2$ discs, for some $\alpha \leq 1$. Any 'sensible' algorithm should therefore satisfy the requirement. Note that $k(n) < \infty$ for some n, immediately implies that $k(m) < \infty$, for any m, as we can cover an $m \times m$ square by a finite number of $n \times n$ squares.

Another version of this theorem requires an upper bound on the density of discs centers to exist, along with stationarity. The proof is very similar and we will mention the small adaptation necessary after the proof of the main theorem.

Theorem 4.4 *Let $\mathcal{A}$ be a finite horizon covering algorithm with the property that the restriction to any box of size $n \times n$ contains at most $k(n)$ discs, for some $k(n) < \infty$. Then there exists $\lambda_1 < \infty$, so that $(X, \lambda, r, \mathcal{A})$ percolates for all $\lambda > \lambda_1$.*

At first sight, the statement of the theorem is counterintuitive, since we claim that we force percolation by *restricting* the number of discs. The point is that by restricting the number of discs (independently of λ), the requirement of covering all points with this restricted number of discs makes percolation unavoidable. Before we prove Theorem 4.4, we first state and prove a preliminary geometric lemma.

Lemma 4.5 *Consider a collection of discs of radius r, with the property that at most $k(n) < \infty$ discs intersect any box of size $n \times n$. Then there exists an $\epsilon = \epsilon(n, r) > 0$ with the following property: if there are, either, at least two components of discs that intersect the boundaries of both $B_{n+r}(x)$ and $B_{n+2r}(x)$, or a component wholly contained in $B_{n+2r}(x)$, then there is a disc of radius ϵ, contained in $B_{n+3r}(x)$, which is not intersected by any disc.*

Proof of Lemma 4.5. We write $B_n = B_n(x)$. All discs that intersect B_{n+3r} must be centered inside B_{n+5r}. Therefore, at most $k = k(n + 5r)$ discs intersect B_{n+3r}. Let C be a component that intersects the boundaries of both B_{n+r} and B_{n+2r}. The number of discs in C that intersect B_{n+2r} is denoted by l. Note that $l \leq k$.

Consider a section AB of the perimeter of C, from the boundary of B_{n+r} to the boundary of B_{n+2r}, which does not intersect either of these boundaries except at its ends (see Figure 4). This section has length at least $r/2$, and consists of parts of the boundaries of at most l discs, each of which appears only once. This latter fact follows from the observation that, since the distance between the boundaries of the two boxes is only $r/2$, any disc that contributes to an arc in AB, must overlap the boundary of at least one box. Moreover, note that were AB to contain two disjoint arcs from

the same disc, then any disc overlapping that disc in order to make these arcs disjoint, must overlap the boundary of the other box. It follows that at least one arc in AB is of length at least $r/2l$. Call this arc a.

Note that there are at most $k-1$ discs intersecting B_{n+3r}, if we do not count the one that has a as a part of its boundary, and none of them intersect a, except at its end points. If we divide a into k arcs of equal size, then each of these discs will be nearest to one of these smaller arcs. One of the smaller arcs (of size at least $r/2kl$) will, however, have no disc which has it as its nearest neighbor. This means that the space left by discs tangent to the ends of this smaller arc cannot be covered, and we can choose ϵ so small that a disc of radius ϵ fits into this space (see Figure 4). The value of ϵ that we have to choose only depends upon r and n.

The same argument applies to a component wholly contained in B_{n+2r}, by considering its perimeter rather than the edge between B_{n+r} and B_{n+2r}. □

Proof of Theorem 4.4. Let $t, u \in \mathbb{Z}$, and denote the box of size $n \times n$ centered at (tn, un) by $B_n(tn, un)$, as before. Let ϵ be chosen as in Lemma 4.5. We say that the vertex (t, u) is a *neighbor* of (t', u') if the boxes $B_n(tn, un)$ and $B_n(t'n, u'n)$ share an edge or corner. We call a vertex (t, u) *good* if all discs of radius ϵ contained in $B_{n+3r}(tn, un)$ contain at least one point of the Poisson process. Denote the event that (t, u) is good by $G(t, u)$. It is clear that when $\lambda \to \infty$, the probability of $G(t, u)$ converges to 1. It is also clear that $G(t, u)$ and $G(t', u')$ are independent whenever $\max\{|t - t'|, |u - u'|\} \geq \lceil \frac{(3r+h)}{n} \rceil$, for n larger than $2r$, and where h is the horizon of the covering. Hence, the configuration of good sites is formed through a discrete, finite-range dependent percolation process, and it fol-

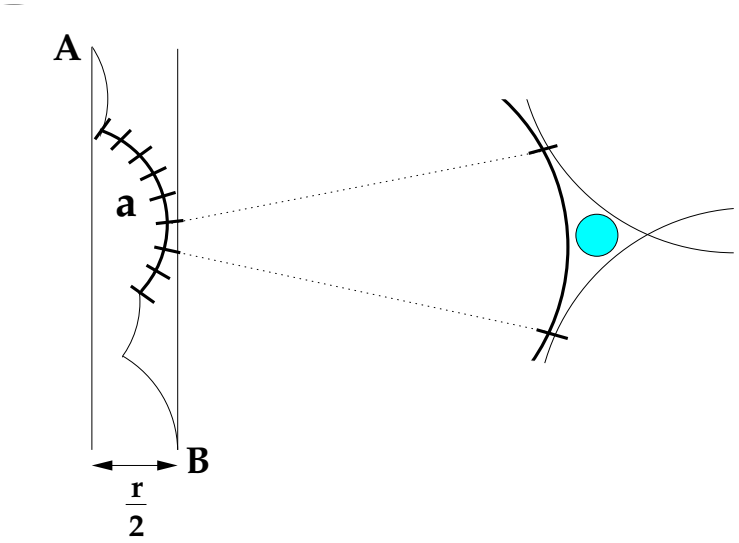

Figure 4: **Subdividing the arc.** An edge AB of a component C connecting the two boundaries of B_{n+r} and B_{n+2r}, has length at least $r/2$. This edge contains an arc a of length at least $r/2l$. Arc a is divided up into k sections, and by one of these we can place a small disc of radius ϵ that is not contained in any cluster.

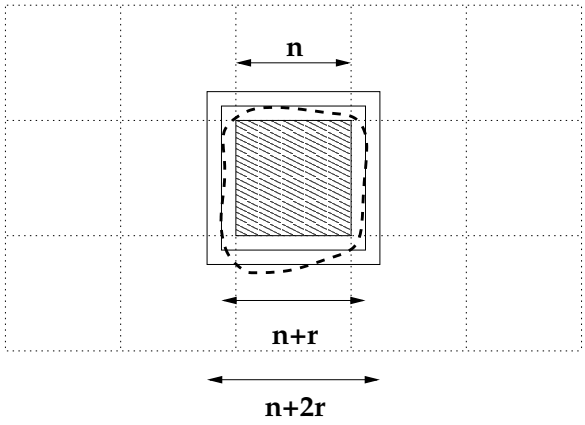

Figure 5: **A good square.** There is only one component of discs (represented by the dashed line) that intersects both B_{n+2r} and B_{n+r}. This component must reach to within 2ϵ of all edges of B_{n+r}, and will therefore intersect a component of an adjacent good square.

lows then from Kesten (1980) that for λ high enough, the good vertices percolate, i.e., contain an infinite component of good squares with probability one. What does this mean for our covering? Consider a good square B_n. By Lemma 4.5 any component cannot be wholly contained in B_{n+2r}, therefore, a component that covers points inside B_{n+2r}, must also intersect the boundary of B_{n+2r}. Also by Lemma 4.5, there can be only one component that intersects the boundaries of both B_{n+2r} and B_{n+r}. For n larger than 2ϵ such component exists and must reach to within 2ϵ of all edges of B_{n+r} (see Figure 5). Thus, the components associated with adjacent good squares must overlap, and we must have an infinite component of discs with probability one. □

The alternative assumption for Theorem 4.4 is that a $\delta < \infty$ exists, such

that:

$$\limsup_{n\to\infty} \frac{\text{number of discs in } B_n}{n^2} < \delta.$$

We also require that we have stationarity under some shift, say of T_m.

In this case, we first choose $\gamma > 0$ sufficiently small that $1-\gamma$ is strictly above the critical point for the finite range discretized process to percolate. We then find n, a multiple of m, so large that the probability that B_{n+5r} is intersected by more than $(\delta+1)(n+5r)^2$ discs is less than γ, uniformly in λ. We then use $(\delta+1)(n+5r)^2$ as our k in the lemma, and find an ϵ such that, if we have at most $(\delta+1)(n+5r)^2$ discs intersecting B_{n+5r} in the way described in the lemma, then we must have a disc of radius ϵ empty of Poisson points.

In the proof of the theorem we then call B_n good if B_{n+3r} has both no disc of radius ϵ empty of Poisson points in it and B_{n+5r} contains at most $(\delta+1)(n+5r)^2$ points. Other boxes are called good analogously. If λ is high enough, these boxes percolate, and we again have an infinite component of discs.

We have seen that both finite horizon algorithms under a bounded density condition and grid algorithms necessarily percolate for high enough λ. It is natural to ask whether the same holds for shift invariant algorithms. It turns out that for these algorithms, large values of the density λ of the points do not guarantee the a.s. existence of an unbounded connected component. This is shown by describing a shift invariant covering algorithm that does not form an unbounded connected component for all λ, thus proving the following theorem:

Theorem 4.6 *There exists a shift invariant covering algorithm $\mathcal{A}$ of all the points of X by discs of radius r, such that for all λ, $(X, \lambda, r, \mathcal{A})$ does not*

percolate.

The proof of Theorem 4.6 is constructive and rather technical. The covering we describe will have density λ.

Without loss of generality, consider covering discs of radius $r = 1$. The main idea is the following: given a realization of the Poisson point process, we first build a large structure of circles similar to those obtained in continuum fractal percolation models. We do this by placing circles of radii $r = 18^m$, with $m \in \mathbb{N}$, where we see certain configurations of points in the plane. The resulting structure is composed of clusters that are finite, but contain every bounded region of the plane. We then derive a shift invariant covering of all the points of X by discs of radius $r = 1$, leaving an empty space near to the boundaries of these clusters.

We illustrate the proof taking the density λ to be 1. The proof for the covering of a Poisson process of another density follows in the same way (in the proofs to follow, only the values of the ϵ's change).

We define a *potential-point* to be a Poisson point with at least one other point in the half-disc of radius $\frac{1}{2}$ to the right of it, and no points in the disc of radius 1 of it, except in the aforementioned half-disc (see Figure 6).

Given a decreasing sequence of positive numbers, $b_1 = \frac{1}{2}, b_2, b_3 \ldots$, an *m-point*, for $m \in \mathbb{N}$, is a potential-point which has its nearest neighboring point between b_m and b_{m+1} away.

We start by proving a few lemmas.

Lemma 4.7 *For $\epsilon > 0$ sufficiently small, there exists a sequence $b_1 = \frac{1}{2}, b_2, b_3, \ldots$, such that the density of m-points is exactly $\epsilon 18^{-2m}$, for each $m \in \mathbb{N}$.*

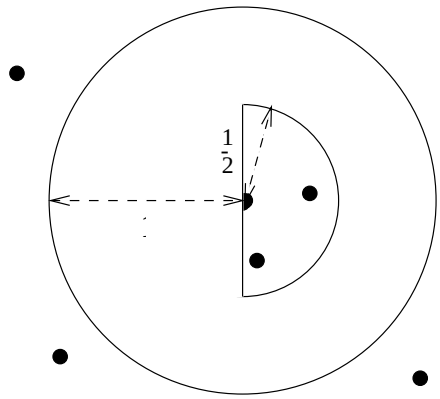

Figure 6: **A potential point** A potential point is a Poisson point with at least one other point in the half-disc of radius $\frac{1}{2}$ to the right of it, and no points in the disc of radius 1 of it, except in that half-disc.

Proof of Lemma 4.7. The density of potential points, λ_p, is calculable. Choose $\epsilon > 0$ so that

$$\sum_{m=1}^{\infty} \epsilon 18^{-2m} = \epsilon \frac{1}{18^2 - 1} \leq \lambda_p.$$

We can now define b_i inductively, i.e. given b_m we choose b_{m+1} so that the density of m-points is exactly $\epsilon 18^{-2m}$. □

We consider a circle of radius 18^m around every m-point. Call such a circle an *m-circle.*

Lemma 4.8 *Every bounded region of the plane is a.s. wholly contained in some m-circle, for some m.*

We will use Proposition 3.1 of Meester and Roy (1994). This states:

Proposition 4.9 (Meester and Roy, 1994) *Let S be a stationary point process in $\mathbb{R}^d$, and let ρ be a non-negative random variable. If $E(\rho^d) = \infty$, then in the Boolean model (S, ρ) the occupied component is a.s. $\mathbb{R}^d$.*

Proof of Lemma 4.8. Let S be the random collection of m-points, for all m. S is stationary. Note that the radii of points in S are independent. Let ρ be the radius distribution of the circles. Then (S, ρ) is a Boolean model, and the occupied component corresponds to all areas contained in some circle. In addition

$$E(\rho^2) = \sum_{m=1}^{\infty} (18^m)^2 \epsilon 18^{-2m} = \infty,$$

so we can apply Proposition 4.9 and conclude that every bounded region is contained in some circle.

□

Now that we have shown that our circles a.s. contain any bounded region, we want to show that any cluster of intersecting circles is a.s. finite. To do this, we give a slight variation of a proof for fractal percolation that is in the book by Meester and Roy (1996) (Theorem 8.1). Again, we proceed by proving a series of lemmas but first we need a couple of definitions.

We define the *sets of possible dependence* to the square $[0, 18^m]^2$ as the rectangles $(I_1 \times I_2 : I_i \in \{[-4 \times 18^{m-1}, 0], [0, 18^m]\}) \backslash [0, 18^m]^2$ (see Figure 7). Sets of possible dependence to other squares of the same size are the natural translations of this. Call a set of sets of possible dependence to a certain square, a *known region*. Define A_m to be the number of m-points in the square $[0, 18^m]^2$, and $A_{>m}$ to be the number of k-points, with $m < k \in \mathbb{N}$, within distance 4 of the square. Let C_m^J be the number of m-points in the known region J to the square.

Lemma 4.10 *For any $\delta > 0$, we can find m' and, uniformly in $m > m'$, $\epsilon > 0$, sufficiently small that:*

$$P(A_m > 0 | A_{>m} = 0 \cap C_m^J = 0) \leq \delta,$$

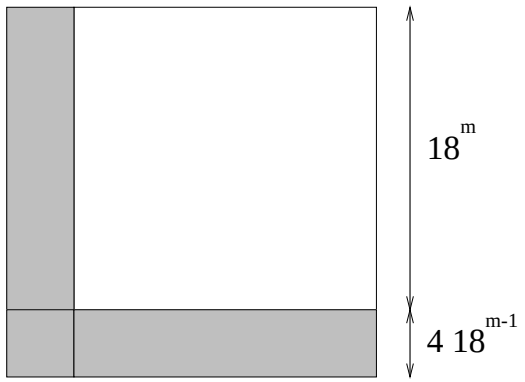

Figure 7: **Sets of possible dependence.**

for any known region, J, to the square $[0, 18^m]^2$.

Proof of Lemma 4.10. We are interested in

$$\begin{aligned} P(A_m > 0 | A_{>m} = 0 \cap C_m^J = 0) &= \frac{P(A_m > 0 \cap A_{>m} = 0 \cap C_m^J = 0)}{P(A_{>m} = 0 \cap C_m^J = 0)} \\ &\leq \frac{P(A_m > 0)}{P(A_{>m} = 0 \cap C_m^J = 0)} \\ &\leq \frac{P(A_m > 0)}{1 - P(A_{>m} > 0) - P(C_m^J > 0)}, \end{aligned}$$

if $E(A_{>m}) + E(C_m^J) < 1$. We recall that, for a non-negative integer-valued random variable N, $P(N > 0) \leq E(N)$, obtaining:

$$\begin{aligned} P(A_m > 0 | A_{>m} = 0 \cap C_m^J = 0) &\leq \frac{P(A_m > 0)}{1 - E(A_{>m}) - E(C_m^J)} \\ &\leq \frac{E(A_m)}{1 - E(A_{>m}) - E(C_m^J)}. \end{aligned}$$

Now, $E(A_{>m}) = (18^{2m} + 4(18^m + 1))18^{-2m}\frac{\epsilon}{18^2-1}$, which is the area of the region within distance 4 of the square multiplied by the total density of k-points, for all $k > m, k \in \mathbb{N}$. Noting that the maximum area of the known region is $(18^m(1+4/18))^2 - (18^m)^2$, we see that $E(C_m^J) \leq ((18^m(1+$

$4/18))^2 - (18^m)^2)\epsilon 18^{-2m}$. Thus, the bound becomes:

$P(A_m > 0 | A_{>m} = 0 \cap C_m^J = 0) \leq$

$$\frac{\epsilon}{1 - (18^{2m} + 4(18^m + 1))18^{-2m}\frac{\epsilon}{18^2-1} - ((18^m(1+4/18))^2 - (18^m)^2)\epsilon 18^{-2m}},$$

which we can make uniformly less than δ by choosing ϵ small enough. $\square$

Let $H(18^n) = [-\frac{3}{2} \times 18^n, \frac{3}{2} \times 18^n]^2 \backslash (-\frac{1}{2}18^n, \frac{1}{2}18^n)^2$. Define $I_{H(18^n)}$ to be the maximal connected cluster of (possibly partial) circles and the boundary of the box $[-\frac{1}{2}18^n, \frac{1}{2}18^n]^2$, fully contained in $H(18^n)$, and define $O_{H(18^n)}$ to be the maximal connected cluster of circles and the boundary of $[-\frac{3}{2} \times 18^n, \frac{3}{2} \times 18^n]^2$, fully contained in $H(18^n)$.

Let $G(18^n)$ be the event that there is a *gap* in $H(18^n)$, i.e. the minimal distance between $I_{H(18^n)}$ and $O_{H(18^n)}$ is at least 18.

Lemma 4.11 *For $\epsilon > 0$ sufficiently small,*

$$\lim_{n\to\infty} P(G(18^n)) \geq \frac{1}{2}.$$

We prove Lemma 4.11 in two parts. Let $G_1(18^n)$ be the event that there is a gap between $I_{H(18^n)}$ and $O_{H(18^n)}$ in $H(18^n)$, when we consider circles only of radius 18^{n-1} or less. The size 18^{n-1} is chosen because it is a convenient size comparable to the size of $H(18^n)$. Let $G_2(18^n)$ be the event that no circles of radius 18^n or greater intersect $H(18^n)$ at all. Clearly, if $G_1(18^n)$ and $G_2(18^n)$ both occur, then $G(18^n)$ does also. Thus, Lemma 4.11 follows from the following two lemmas.

Lemma 4.12 *For $\epsilon > 0$ sufficiently small,*

$$\lim_{n\to\infty} P(G_1(18^n)) \geq \frac{3}{4}.$$

Lemma 4.13 *For $\epsilon > 0$ sufficiently small,*

$$\lim_{n\to\infty} P(G_2(18^n)) \geq \frac{3}{4}.$$

Proof of Lemma 4.12. This proof closely follows that of Theorem 8.1 of Meester and Roy (1996) for fractal percolation, except for a number of extra technicalities. We are going to show that we dominate a version of the process which has more independence.

We first divide $H(18^n)$ into 8×18^2 sub-squares of size 18^{n-1} in the obvious way. We call two squares of the same size neighbors if they share an edge or corner.

Suppose for a moment that the probability that a square of size 18^m contains an m-point, is uniformly δ, for all m, independently of the occurrence of k-points anywhere, with $k \neq m, k \in \mathbb{N}$, and independently of the occurrence of m-points outside the square. We give our proof initially under this assumption, and then compare our original process with this.

We consider the 8×18^2sub-squares of $H(18^n)$, and in an order such that for any two squares, B and B' say, B is considered before B' if B is neither to the right of nor above B'. We examine, in this order, these sub-squares of size 18^{n-1}, looking for $(n-1)$-points, in the following inductive fashion:

- to begin, all squares are declared to be neither corrupt nor bad.
- if a sub-square is not corrupt, then we examine the whole of it, looking for $(n-1)$-points. If it contains any, then we call it *bad*, and its neighboring squares of the same level *corrupt*.
- if a sub-square is corrupt then we do not examine it.

Squares are bad if they contain centers of circles of comparable size to themselves, which then may extend into the corrupt squares. Corrupt squares may or may not contain $(n-1)$-points. We are careful not to find out this information, as it might tell us something about the distribution of points in the squares we have not yet considered.

We can then divide up each of the *good* squares (those that are neither bad nor corrupt) into 18^2 pieces, obtaining at most $8 \times 18^{2\times 2}$ squares of size 18^{n-2}, and we examine those, in an order such those squares nearer the bottom-left hand corner are considered first, looking for $(n-2)$-points, in the same inductive fashion as above. We end up declaring each of the squares of size 18^{n-2} that are sub-squares of good squares in $H(18^n)$ to be good, bad or corrupt.

We divide up each of the good squares of size 18^{n-2} into 18^2 squares of size 18^{n-3}, and use the same procedure to declare each good, bad or corrupt. We can then divide up each of the good squares, and repeat this procedure, while we still have good squares, and to a minimum square size of 18.

We then work backwards through the squares, starting with the smallest, to declare each either *dreadful* or not. A square of size 18 is dreadful if it is bad. In an inductive fashion, a square of size 18^m is dreadful if it is either a) bad or b) good but contains 2 or more dreadful squares of size 18^{m-1}. We call $H(18^n)$ dreadful if it contains any dreadful squares of size 18^{n-1}.

Under our temporary independence assumption, a square of size 18 is dreadful with probability δ, conditioned on the fact that it has not been declared corrupt before being checked. Thus, the probability that it is dreadful is at most δ. Then the probability p_m that a square of size 18^m is dreadful, is the probability it is a) bad or b) good but contains 2 or more dreadful squares of size 18^{m-1}. The probability that it is bad, is, as for a square of

size 18, at most δ. The probability that it is good is at most $1-\delta$. It may be that some of the sub-squares of size 18^{m-1} of this square are corrupt, due to being neighbors of bad squares of the same size outside this square. Let N be the number of such squares. As the probability of a square being dreadful is maximal when $N=0$, it follows that:

$$p_m \leq \delta + (1-\delta)((1-p_{m-1})^{18^2} - 18^2 p_{m-1}(1-p_{m-1})^{18^2-1}) .$$

Letting

$$f(p,\delta) = \delta + (1-\delta)(1-(1-p)^{18^2} - 18^2 p(1-p)^{18^2-1}),$$

this becomes $p_m \leq f(p_{m-1},\delta)$. Note that $p_1 = \delta$. If we can now show, for all δ, that

$$0 \leq p \leq b(\delta) \text{ implies } 0 \leq f(p,\delta) \leq b(\delta),$$

for some $b(\delta) \geq \delta$, it will follow that $p_m \leq b(\delta)$ for all m. We need $b(\delta)$ to be a bound that tends to 0 with δ.

Note that $f(p,\delta)$ is continuous in p and δ, that $f(0,\delta) = \delta$ and that $\frac{\partial f}{\partial p}(0,\delta) = 0$. It follows that $f(p,\delta) = p$ has a solution in $[0,1]$ for δ sufficiently small. The smallest such solution we call $b(\delta)$. Notice that $b(\delta) \geq \delta$, as $f(0,\delta) = \delta$ and $\frac{\partial f}{\partial p}(p,\delta) \geq 0$ on $[0,1]$. $\frac{\partial f}{\partial p}(p,\delta) \geq 0$ on $[0,1]$ also tells us that $0 \leq p \leq b(\delta)$ implies $0 \leq f(p,\delta) \leq b(\delta)$. Since $f(0,0) = 0$, $\frac{\partial f}{\partial p}(0,\delta) = 0$ and f is continuous, $\lim_{\delta \downarrow 0} b(\delta) = 0$.

We have shown that the probability that a square of size 18^m is dreadful is bounded from above, uniformly in m, by a function of δ that we can make arbitrarily small by choosing δ small enough. It then follows that the probability that $H(18^n)$ is dreadful (i.e. contains any dreadful squares of size 18^{n-1}) can be made as small as we like by choosing δ sufficiently small. We choose δ so that this probability is less than $\frac{1}{4}$.

We can now come back to our original process, and give up the independence assumption. We make the following comments:

- The probability that a square of size 18^m contains an m-point, when we come to check it, given any of the information we already have found, is at most δ, by Lemma 4.10, and because we never consider a corrupt square.
- By choosing ϵ sufficiently small we can make the bound δ as small as we need.

Let us consider what it means for $H(18^n)$ not to be dreadful. We argue that, in this case, we cannot have a connection by circles, of the appropriate sizes, from the inside to the outside of the box.

We first note that any m-point in $H(18^n)$ is either in a bad or a corrupt box of size 18^m. An m-circle may thus only intersect a box of size 18^m, if that box is either bad, corrupt, or neighbors a corrupt square of the same size, unless it neighbors the boundary of $H(18^n)$.

We now give a series of definitions. Call a box of size 18^m *dodgy* if it is either bad, corrupt, neighbors a corrupt square of the same size, is dreadful, neighbors a dreadful square, or neighbors the border of $H(18^n)$. Call it *clean* if it is not dodgy. Call two boxes of size 18^m adjacent if they share an edge. Define an *m-circuit* to be a series of boxes of size 18^m, $(B_1, B_2, B_3, \ldots, B_k)$ where B_i is adjacent to B_{i+1} for $i = 1, 2, \ldots, k-1$ and B_k is adjacent to B_1. We also require that the circuit cuts off the origin from infinity.

A sub-circuit of boxes of size 18^j,$(B_1^j, B_2^j, B_3^j, \ldots, B_{k_j}^j)$, of a circuit of boxes of size 18^m (with $j < m$),$(B_1^m, B_2^m, B_3^m, \ldots, B_{k_m}^m)$, is a circuit of boxes of size 18^j inside the circuit of boxes of size 18^m, such that there exists $0 = i_1, i_2, i_3, \ldots, i_{k_m-1}, i_{k_m} = k_j$, such that $B_{i_l+1}^j, B_{i_l+2}^j, B_{i_l+3}^j, \ldots, B_{i_{l+1}}^j$

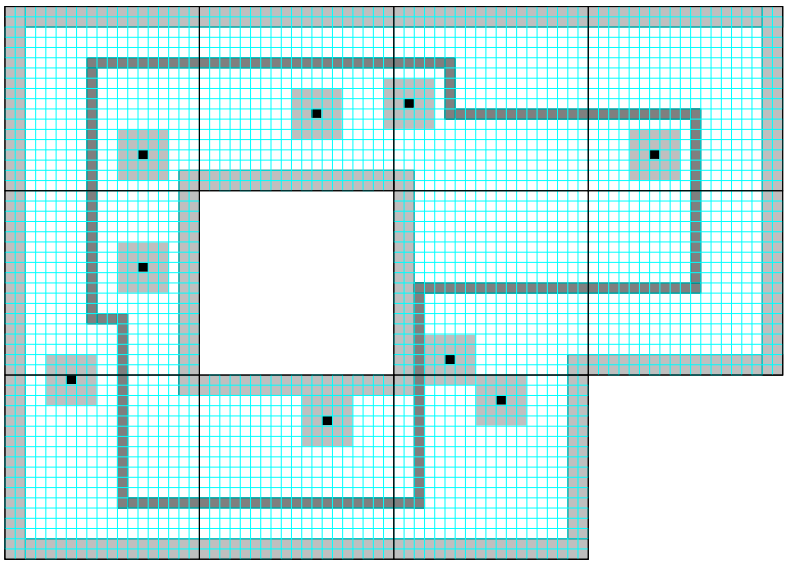

Figure 8: **Avoiding dodgy squares.** A circuit of squares of size 18^m contains a circuit of squares of size 18^{m-1}, avoiding dodgy squares.

are contained in B_l^m. This means that the first few boxes of size 18^j are contained in B_1^m, the next few are in B_2^m, and so on.

We say that the property E_m holds if any m-circuit consisting of clean boxes of size 18^m contains a sub-circuit of boxes of size 18 that are not intersected by k-circles, for all $k \leq m$. Our objective is to show that E_m holds for $m = 1, 2, \ldots, n-1$, by induction on m.

If $m = 1$, we have a 1-circuit of clean boxes. Therefore there can be no 1-circle intersecting any of these boxes, and E_1 holds.

The inductive hypothesis is that E_{m-1} holds. In order to perform the induction step and prove that E_m also holds, consider an m-circuit of clean boxes of size 18^m, as depicted by in Figure 8. In this figure, the smallest squares are of size 18^{m-1}. If we show the existence of a sub-circuit of clean

boxes of size 18^{m-1} inside the m-circuit, then, by the inductive hypothesis, the occurrence of E_m will follow.

By construction, all the boxes of size 18^m in our circuit are clean. This tells us that none of them intersects an m-circle. It also means that there is at most one dreadful square inside each box of size 18^m in the circuit. These are depicted in Figure 8 as black squares. A dreadful square may cause a 5×5 block of squares of size 18^{m-1} to be dodgy (grey 5×5 blocks in Figure 8). This can happen because the dreadful square can neighbor a corrupt square, which also has neighbors. Finally, any of the squares of size 18^{m-1} neighboring the edge of the circuit, or the neighbors of these squares, may be dodgy, due to the proximity of dreadful squares just outside the circuit. This latter case is depicted in Figure 8 by the two grey circuits of width $2 \times 18^{m-1}$, along the edges of the m-circuit. These are all the possibilities for dodgy squares of size 18^{m-1} in our circuit.

By considering all possible arrangements of the dodgy squares we see that there must be a circuit of squares of size 18^{m-1} inside our circuit that avoids those dodgy squares (see Figure 8). Then, by the inductive hypothesis, there must also be a sub-circuit of squares of size 18, so E_m holds.

H_{18^n} is a circuit of boxes of size 18^n, and if all are clean, by the argument above, it follows that we must have a sub-circuit of boxes of size 18 that are not intersected by any m-circles, for $m = 1, 2, \ldots, n$, which is what we wanted to prove. □

Proof of Lemma 4.13. Let X be the number of circles of radius at least 18^n that intersect $[-\frac{3}{2} \times 18^n, \frac{3}{2} \times 18^n]^2$. A circle can only touch the boundary of the box if the distance between the center of the circle and the center of the box is within $3 \times 18^n \times \frac{1}{\sqrt{2}}$ of the circle's radius (see Figure 9). Now

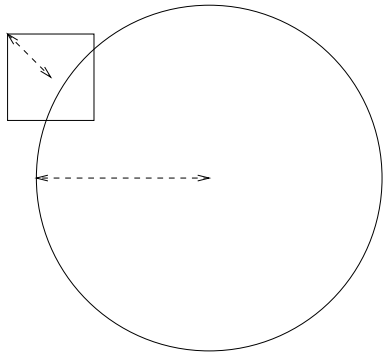

Figure 9: **Intersection between disc and box.** A circle of radius at least 18^n can touch the boundary of the box $[-\frac{3}{2} \times 18^n, \frac{3}{2} \times 18^n]^2$ only if the distance between the center of the circle and the center of the box is within $3 \times 18^n \times \frac{1}{\sqrt{2}}$ of the circle's radius.

write:

$$\begin{aligned} P(X > 0) &\leq E(X) \\ &\leq \sum_{b=n}^{\infty} \epsilon 18^{-2b} \pi ((18^b + \frac{3 \times 18^n}{\sqrt{2}})^2 - (18^b - \frac{3 \times 18^n}{\sqrt{2}})^2). \end{aligned}$$

We can make this less than $1/4$ for all n by our choice of ϵ, implying that the probability of a large disc intersecting the box is less than $1/4$. □

We can now prove the finiteness of our clusters. We define a *thickened cluster* as a maximal connected component of points strictly within distance 9 of a disc boundary.

Lemma 4.14 *All thickened clusters are finite, for $\epsilon > 0$ sufficiently small.*

Proof of Lemma 4.14.

By Lemma 4.11 there is a gap between $[-\frac{18^m}{2}, \frac{18^m}{2}]^2$ and $[-\frac{3 \times 18^m}{2}, \frac{3 \times 18^m}{2}]^2$, with probability at least $1/2$, by our choice of ϵ, for every m.

We now need to show that such a gap exists a.s. , for some m. We need to be careful in doing this, so that the negative information gained by the knowledge that there is no such gap, for a certain m, does not prejudice our attempts to find one in a later m. We proceed as follows.

We begin by looking whether there is a gap between the boundaries of $[-\frac{18}{2}, \frac{18}{2}]^2$ and $[-\frac{3\times 18}{2}, \frac{3\times 18}{2}]^2$. In order to do this we search for the centers of the circles that might intersect this area, in order of increasing size. Either there is no gap, in which case a (random) $K_1 < \infty$ exists, so that these boundaries are connected by circles of radius up to 18^{K_1}, or there is a gap (with probability of at least $1/2$). In the latter case we would be satisfied. In the first case we can find an M_1 so large that $[-\frac{18^{M_1}}{2}, \frac{18^{M_1}}{2}]^2$ and $[-\frac{3\times 18^{M_1}}{2}, \frac{3\times 18^{M_1}}{2}]^2$, cannot be overlapped by any circle of radius up to 18^{K_1} that could also have overlapped $[-\frac{3\times 18}{2}, \frac{3\times 18}{2}]^2$. We know nothing about larger circles.

We have no information about the circles that may connect $[-\frac{18^{M_1}}{2}, \frac{18^{M_1}}{2}]^2$ and $[-\frac{3\times 18^{M_1}}{2}, \frac{3\times 18^{M_1}}{2}]^2$; hence, the probability that there is a gap between them is again at least $1/2$. We search for a connection between the boundaries of $[-\frac{18^{M_1}}{2}, \frac{18^{M_1}}{2}]^2$ and $[-\frac{3\times 18^{M_1}}{2}, \frac{3\times 18^{M_1}}{2}]^2$, again starting by looking at the smallest circles. Either there is a gap, or there exists $K_2 < \infty$ such that there is a connection from one of these boundaries to the other using circles of size up to 18^{K_2}. In this latter case, we can find an $M_2 > M_1$ so large that $[-\frac{18^{M_2}}{2}, \frac{18^{M_2}}{2}]^2$ and $[-\frac{3\times 18^{M_2}}{2}, \frac{3\times 18^{M_2}}{2}]^2$, cannot be overlapped by any circle of radius up to 18^{K_2} that could also have overlapped $[-\frac{3\times 18^{M_2}}{2}, \frac{3\times 18^{M_2}}{2}]^2$.

We now search again for a gap, this time between $[-\frac{18^{M_2}}{2}, \frac{18^{M_2}}{2}]^2$ and $[-\frac{3\times 18^{M_2}}{2}, \frac{3\times 18^{M_2}}{2}]^2$, and repeat. At every stage we have a probability of at least $1/2$ of there being a gap, independently of the previous times. If there

is no gap (which happens with a probability of at most 1/2) then we find this out at some time and search in a larger annulus. It follows that there is almost surely a gap.

□

Finally, we can now give a proof of Theorem 4.6, by describing a shift invariant covering algorithm that a.s. never forms an unbounded component of covering discs, for all λ.

Proof of Theorem 4.6. As this covering should be a deterministic function of the points, we first calculate λ in our realization, via

$$\lambda = \lim_{n \to \infty} \frac{\text{number of points in } B_n(0)}{n^2}$$

if this limit exists and is constant. This happens with probability 1. Otherwise we take $\lambda = 1$. Fix ϵ at half the supremum of all values of ϵ that allow all our proofs to work at this particular value of λ. We construct *smooth* curves based upon the finite circle clusters. Consider some maximal set of clusters such that, if we take the locus of points at a maximal distance of 4 from the points in the clusters, then this forms a connected set, and run a disc of radius two around the outside of this set (see Figure 10). The disc traces out a kind of sausage shape around the clusters. We note that all such sausages must be finite by Lemma 4.14.

We take the inside edge of this sausage as our curve, and note that a covering disc (of radius 1) can get arbitrarily close to any point of it without touching it.

We construct these smooth curves for each set of sufficiently close clusters, noting that they surround every region, are always finite, and never come within distance 18 of each other.

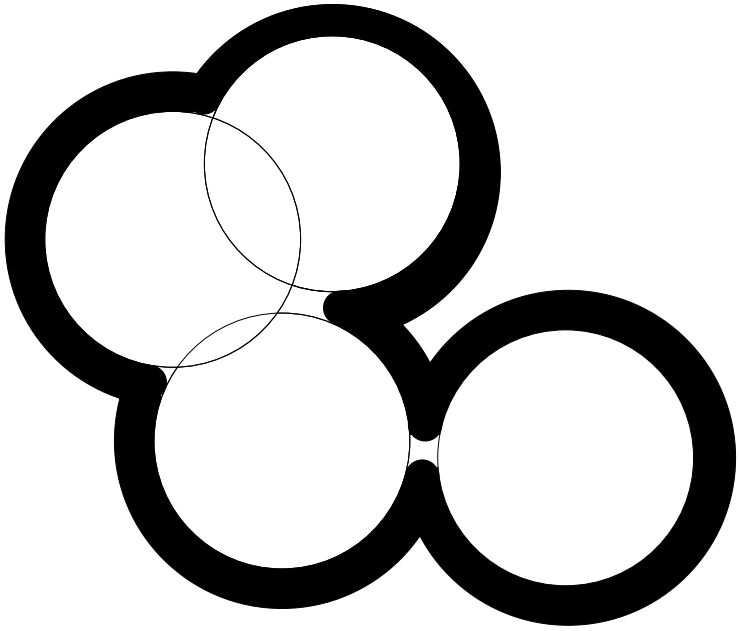

Figure 10: **Sausages.** By sliding a disc of radius 2 along the boundary of some cluster, we trace a kind of sausage shape.

We finally cover our Poisson points as follows:

- if a point is at a distance more than 2 from every smooth curve, then we center a covering disc at the point.
- if a point is within distance 2 of a smooth curve, then we place a disc so that its edge covers the point, and so that the center of the disc is at the maximum distance away from the smooth curve. If there are a number of such possible positions, we choose the leftmost.

It immediately follows that, for any given value of λ, a.s. there is no percolation. □

5 Scaling

In this section we consider an extension that is useful to model the transmission power in wireless communication networks. We look at the percolation properties of our model, for different values of the connectivity range of the base stations and of the clients.

Let r be the clients connectivity range and let R be the base stations connectivity range. It follows that discs of radius r are used to cover the points of X and two disc centers are considered connected, if their distance is less than, or equal to R (see Figure 11). We are interested in the a.s. existence of an unbounded connected component of disc centers, for large values of the density λ of the Poisson point process.

Our result is the following.

Theorem 5.1 (The Scaling Theorem) *Let $G \subset \mathbb{R}^2$ be the set of all vertices of a square lattice in which the distance between two neighboring*

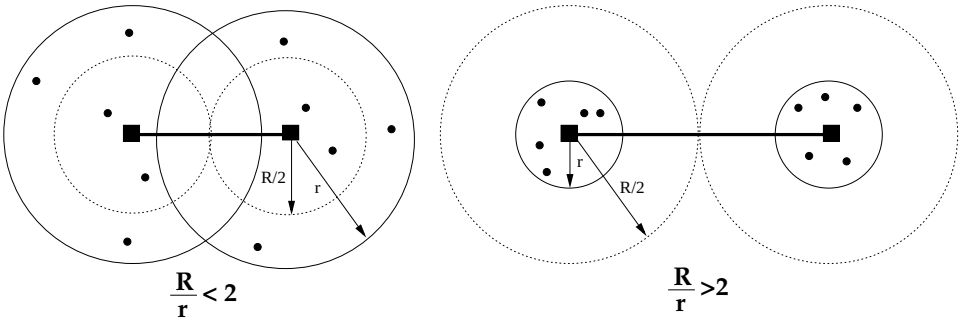

Figure 11: **Scaling.** Points are covered by solid line discs of radius r. Discs centers are considered connected if their distance is at most R.

lattice vertices is δ. Call two disc centers connected if their distance is at most R. We have:

- *CASE 1. If $\frac{R}{r} \leq 1$ then, for any $\delta > 0$, there exists a grid covering algorithm $\mathcal{A}$ that places discs only at the vertices of G, such that, for all λ, $(X, \lambda, r, \mathcal{A})$ does not percolate.*

- *CASE 2. If $1 < \frac{R}{r} < 2$ then, there exists a $\delta > 0$, depending on $\frac{R}{r}$, such that there exists a grid covering algorithm $\mathcal{A}$ that places discs only at the vertices of G and, for all λ, $(X, \lambda, r, \mathcal{A})$ does not percolate.*

- *CASE 3. If $\frac{R}{r} = 2$ then, for any $\delta > 0$, for any grid covering algorithm $\mathcal{A}$, there exists a $\lambda_1 < \infty$, such that, for all $\lambda > \lambda_1$, $(X, \lambda, r, \mathcal{A})$ percolates.*

- *CASE 4. If $\frac{R}{r} > 2$ then, for any covering algorithm $\mathcal{A}$, there exists a $\lambda_1 < \infty$, such that, for all $\lambda > \lambda_1$, $(X, \lambda, r, \mathcal{A})$ percolates.*

Note that Case 4 of the theorem states that in a wireless network in which base stations can communicate at a distance larger than twice the

maximum communication distance to the clients, an unbounded connected component forms a.s. for large values of the density of the clients, regardless of the covering algorithm used to build the cellular network.

Before giving a proof of Theorem 5.1, we discuss an intuitive interpretation of the Theorem. Consider a fixed value of R and let r approach zero. In the limit for $r \to 0$, a covering algorithm needs to place a disc at each point of X, therefore, any covering algorithm behaves as the standard Poisson Boolean model $\{X, \lambda, R/2\}$. For this model it is known that an unbounded connected component a.s. forms, for large values of the density λ.

What Theorem 5.1 states is that, when r is small, the covering algorithm is constrained to place the discs almost as a Poisson point process, therefore an unbounded connected component a.s. forms, for large values of λ. On the contrary, when r is large, a covering algorithm has more freedom in placing the covering discs and percolation can be avoided.

Note that (surprisingly) we do not need $r \approx 0$, and the covering process to behave exactly as a Poisson Boolean model, to obtain the percolation property, but as long as r is small enough that the ratio $\frac{R}{r}$ is greater than 2, Case 4 of the theorem applies, and the result on the existence of an unbounded connected component holds for any covering algorithm.

Proof of Theorem 5.1.

Case 1. We can restrict our attention to $\frac{R}{r} = 1$. That is because if a grid covering algorithm does not form an unbounded connected component when $\frac{R}{r} = 1$, then it does not form such a component when $\frac{R}{r} < 1$ either.

Note that, for $\frac{R}{r} = 1$, two disc centers are considered connected if and only if the corresponding discs of radius r cover each other's centers. Moreover, in order to being able to cover all points on the plane by using only

grid discs of radius r, the grid spacing δ must be at most $\sqrt{2}r$.

We now consider all values of the grid spacing $\delta \leq \sqrt{2}r$, subdivided into intervals.

- For $r < \delta \leq \sqrt{2}r$, any grid covering algorithm places discs on the plane that do not touch each other's centers.

- For $\frac{r}{\sqrt{2}} < \delta \leq r$, consider the tiling of the plane depicted in the left part of Figure 12. Discs of this tiling do not cover each other's centers, therefore, any grid covering algorithm that covers all the points of X using only the grid discs depicted in the left part of Figure 12 does not form an unbounded connected component, a.s. , for any value of λ.

- For $2r < \delta \leq \frac{r}{\sqrt{2}}$, consider the tiling depicted in the right part of Figure 12. Discs of this tiling do not cover each other's centers, therefore, any grid covering algorithm that covers all the points of X using only the grid discs depicted in the right part of Figure 12 does not form an unbounded connected component, a.s., for any value of λ.

For the remaining values of δ, we can use the same tiling of the two cases depicted in Figure 12, scaled by the appropriate factor.

Case 2. In this case, two disc centers are considered connected if and only if the corresponding discs of radius r overlap by a region of measure at least $\epsilon > 0$, where the value of ϵ depends on the ratio $\frac{R}{r}$.

We follow a similar construction as that used to prove Proposition 4.3. Draw circles of radii $\{3kr, k \in \mathbb{N}\}$ around the origin, and notice that a.s. no Poisson point falls on any of these circles. Then cover the Poisson points, each with a disc of radius r, without intersecting these circles. Notice that the circles divide the plane into finite annuli, whose boundaries are not

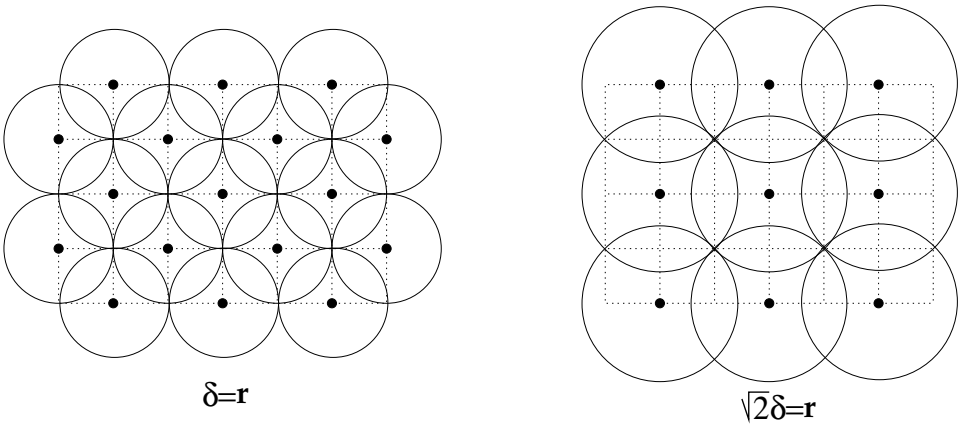

Figure 12: **Theorem 5.1, Case 1.** Two tilings of the plane by discs centered on a grid that do not cover each other's centers.

covered by discs. We now approximate this covering using a grid covering. Consider a square grid G and move each disc of the above covering to the nearest vertex of G that still allows to cover its corresponding Poisson point. Note that each disc needs to be translated by at most $\sqrt{2}\delta$. That is because a Poisson point is covered by a disc centered within r from it, and there is always a grid vertex, within radius r from the Poisson point, that is also within $\sqrt{2}\delta$ from this center (see Figure 13). By this translation, some discs may intersect the boundaries of the annuli, that were previously untouched. We then take the grid size δ so small, that any two discs that intersect these boundaries do not overlap by an area of measure greater, or equal to ϵ, and are therefore not connected.

It immediately follows that, for any given value of the density λ, a.s. there is not any unbounded connected component for this covering.

Case 3. This case is proven by Theorem 4.3.

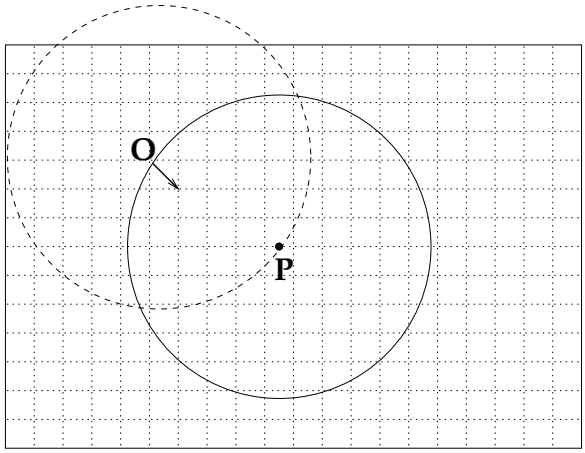

Figure 13: **Theorem 5.1, Case 2.** A Poisson point P is covered by a disc centered at point O, that is within r from P. The covering disc can be moved to a nearby grid vertex, that is inside the solid disc and is within $\sqrt{2}\delta$ from O, and still covers point P.

Case 4. In this case, two disc centers are considered connected if and only they are at a distance of at most $R - 2r$ (see right hand side of Figure 11).

We construct a mapping from the covering discs to a discrete site percolation model. Consider a partition of the plane into boxes of side length $\frac{\epsilon}{\sqrt{2}}$, with $0 < \epsilon \leq \frac{R}{2} - r$. Note that if some point of X falls inside a box of the partition, then it must be covered by a disc of radius r, and therefore the entire box is covered by a disc of radius $\frac{R}{2}$ (see Figure 14).

Consider now each $\epsilon \times \epsilon$ box as a site of a site percolation model. Call the site occupied, if there is at least a point of X situated inside the box. Clearly, the occupancy of a site is independent of the occupancy of other sites, and, if two adjacent sites are both occupied, then the corresponding covering discs form a connected component. Next, we choose λ large enough that the probability of a site being occupied is larger than p_c, where p_c is the critical probability for site percolation on a square lattice. The a.s. existence of an

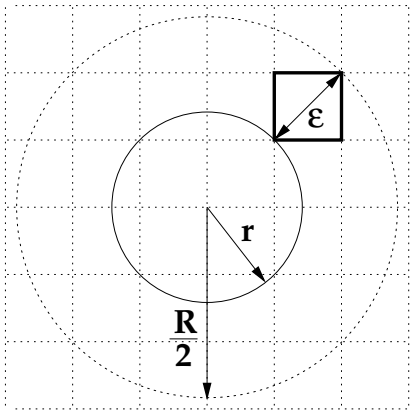

Figure 14: **CASE 4.** If some point of X falls inside a box of side length $\frac{\epsilon}{\sqrt{2}}$, then the entire box is covered by a disc of radius $\frac{R}{2}$.

unbounded connected component of covering discs immediately follows. □

6 Optimal algorithms

In this section we explore the notion of an optimal algorithm, i.e. one which uses as few discs as possible. We first consider n-square algorithms, and show that they are nearly optimal, then we describe an algorithm which really is optimal.

Fix the density of points, λ, and r, and extend the definition of the density of a covering, $\mathcal{A}$, to be

$$\delta_{\mathcal{A}} = \lim_{n\to\infty} \frac{\text{number of discs centered in } B_n}{n^2}$$

if this exists and is a constant a.s., and ∞ otherwise. We then define the *optimal density* to be

$$\delta^{opt} = \inf_{\mathcal{A}} \delta_{\mathcal{A}},$$

where the infimum runs over all covering algorithms. We are interested in whether there exists an *optimal algorithm*, $\mathcal{A}^{opt}$, for which, a.s. , $\delta^{opt} = \delta_{\mathcal{A}^{opt}}$. This would be an algorithm which uses the minimal number of discs possible. We would not expect such an algorithm to have a finite horizon.

First we show that n-square algorithms can get as close as we like to the optimal density. Define δ_n to be the density of discs under an n-square algorithm. By ergodicity $\delta_n < \infty$ exists. Notice that δ_n does not depend upon the particular n-square algorithm we choose.

Theorem 6.1 *Given $\epsilon > 0$, there exists n_ϵ such that $\delta_{n_\epsilon} < \delta^{opt} + \epsilon$, and hence $\delta^{opt} = \inf_n \delta_n$.*

Proof of Theorem 6.1 We prove this theorem by contradiction, so suppose that we can find an ϵ such that there is no n-square covering with density between δ^{opt} and $\delta^{opt} + \epsilon$. We can find another covering, $\hat{\mathcal{A}}$ say, with density $\delta_{\hat{\mathcal{A}}} \in [\delta^{opt}, \delta^{opt} + \epsilon/4]$, by the definition of δ^{opt}.

Choose $\gamma > 0$ such that $(1-\gamma)(\delta^{opt} + \epsilon/2) + \gamma(1+\epsilon/2)/r^2 < \delta^{opt} + \epsilon$. Note that the number of discs necessary to cover an n-square is at most $\lceil n/r \rceil^2$.

As $\delta_{\hat{\mathcal{A}}} = \lim_{n\to\infty} \frac{\text{number of discs centered in } B_n}{n^2}$, a.s, we can choose $n_\epsilon > 1$ sufficiently large that a) the number of discs centered in B_{n_ϵ} is less than $n_\epsilon^2(\delta^{opt} + \epsilon/2)$ with probability larger than $1-\gamma$, and b) $\lceil n/r \rceil^2 (1/n)^2 \leq (1+\epsilon/2)/r^2$.

Given a finite set of points there are a finite number of distinct possible coverings of those points, where we call two coverings *distinct* if there exist a set of points covered by one disc in one of the coverings but covered by two or more in the other. Coverings that are not distinct are *equivalent*. We will occasionally work with the equivalence classes of theses coverings.

We now define a covering of the box B_{n_ϵ} which is based upon $\hat{\mathcal{A}}$ but is independent of the points outside B_{n_ϵ}. Given a point configuration, π, in B_{n_ϵ} there is a finite set of equivalent classes of coverings of these points. Let S_π be the subset of equivalence classes which occur with positive probability if we use $\hat{\mathcal{A}}$ to cover $\pi \cup X|'_{B^c_{n_\epsilon}}$, where $X|'_{B^c_{n_\epsilon}}$ is a Poisson process on $B^c_{n_\epsilon}$ independent of π and X. In each equivalence class all coverings use the same number of discs so we can choose an equivalence class from S_π in which the number of discs used is minimal, according to some deterministic rule. Choose a covering from this class, according to some other deterministic rule. This is the covering we use to cover B_{n_ϵ}. Note that this is independent of the actual point process outside B_{n_ϵ}. The expected number of discs required to cover the points in B_{n_ϵ} under this algorithm can be at most that under $\hat{\mathcal{A}}$.

We now divide up the plane into squares of size n_ϵ, and cover each square independently using the same algorithm in each as we use on B_{n_ϵ}. For those squares for which this requires less than or equal to $n_\epsilon^2(\delta^{opt} + \epsilon/2)$ discs, we use this covering. In the other squares we cover optimally, which means that we have a density of at most $(1 + \epsilon/2)/r^2$ on these squares.

We have created an algorithm that covers each square of size n_ϵ independently, and which therefore cannot have a density less than δ_{n_ϵ}. However, the density of the covering is at most $(1-\gamma)(\delta^{opt}+\epsilon/2)+\gamma(1+\epsilon/2)/r^2 < \delta^{opt}+\epsilon$, and we have a contradiction. □

Next, we prove a proposition that extends the previous theorem and will be useful later.

Proposition 6.2 *Let δ_n be the density of discs under an n-square algo-*

rithm. Then

$$\lim_{n\to\infty} \delta_n = \delta^{opt} \ .$$

Proof of Proposition 6.2 We know that $\delta^{opt} = \inf_n \delta_n$, and that $\delta_n n^2$ is the expected number of discs needed to cover an n-square. For the sake of contradiction suppose that there exist $\epsilon > 0$ and a sequence $\{t_1, t_2, \dots\}$, $\lim_i t_i = \infty$ such that $\delta_{t_i} > \delta^{opt} + \epsilon$ for all i. However we can choose α so that $\delta_\alpha < \delta^{opt} + \epsilon/3$. We can also choose i so large that $\left(1 - \left\lfloor \frac{t_i}{\alpha} \right\rfloor^2 \left(\frac{\alpha}{t_i}\right)^2\right) \lambda < \epsilon/3$. The reason we need this will become clear shortly.

We then cover the square B_{t_i} as follows. We first divide as much of the square as possible into squares of size α. Each of these we cover optimally. We have an area of $(t_i^2 - \alpha^2 \left\lfloor \frac{t_i}{\alpha} \right\rfloor^2)$ left, and each of the points in this area we cover with one disc. This gives us a covering with expected density,

$$\left\lfloor \frac{t_i}{\alpha} \right\rfloor^2 \left(\frac{\alpha}{t_i}\right)^2 \delta_\alpha + \left(1 - \left\lfloor \frac{t_i}{\alpha} \right\rfloor^2 \left(\frac{\alpha}{t_i}\right)^2\right) \lambda \leq \delta^{opt} + \epsilon/3 + \epsilon/3.$$

However, the minimal expected density for any algorithm covering the box B_{t_i}, $\delta_{t_i} > \delta^{opt} + \epsilon$, so we have a contradiction. □

Note that it is still not clear a priori that an optimal algorithm should exist. The existence of an optimal density, defined as the infimum over all attainable densities, does not have to be attainable itself. However, we have the following theorem:

Theorem 6.3 *There exists an optimal algorithm.*

We give first the algorithm that we claim is optimal, and then prove that this is so in a number of steps.

The algorithm

The basic idea is to recursively find coverings of points in boxes, which are part of optimal coverings of points in larger boxes, incrementally covering the whole plane. Consider the boxes $B_n, n \in \mathbb{N}$. In each box there is a finite number of distinct possible optimal coverings of the points in that box, and from now on we work with the equivalence classes of these coverings.

Let A_n be an optimal covering of B_n chosen according to some rule, and let $A_n|_m$ be the covering of the points of B_m induced by A_n, by which we mean the covering consisting of all discs of A_n that cover at least one point of B_m. Take B_1 and consider the sequence of coverings $A_n|_1, n \in \mathbb{N}$. There are only finitely many equivalent coverings of B_1, and at least one of these must appear infinitely many times in $A_n|_1, n \in \mathbb{N}$. Choose such an equivalence class and cover B_1 using a covering from this class. Let $I_1 \subset \mathbb{N}$ be the infinite set of indices, such that $A_n|_1, n \in I_1$ is in the chosen equivalence class. Let J_1 be the smallest element of this set that is larger than 1.

We can repeat this exercise to find a covering of B_{J_1}. We consider the sequence of equivalence classes of coverings $A_n|_{J_1}, n \in I_1$. Again this set is finite, so we can find an infinite set $I_2 \subseteq I_1$, such that all the coverings $A_n|_{J_1}, n \in I_2$ are in the same equivalence class. We fix a covering for B_{J_1} from this class, which is necessarily consistent with the covering we have already chosen for B_1. Let J_2 be the smallest element of I_2 that is larger than J_1.

We repeat this procedure in the natural way. Every time we fix a covering of a box corresponding to the lowest element of some index set, and take a new index set which is an infinite subset of the previous one. We find in this way a covering of the whole plane.

It is not obvious that this covering should be optimal. We have a se-

quence of boxes, $\{B_1, B_{J_1}, B_{J_2}, \ldots\}$ which contain coverings that are parts of optimal coverings of larger boxes, but there is no a priori guarantee that the limit $\lim_{i\to\infty} \frac{\text{number of discs in } B_i}{i^2}$ will be δ^{opt}. To show that it really is optimal we need a few definitions and a lemma.

For a bounded subset of the plane, V, let $\tilde{V}$ be the set $\{x : \exists y \in V \text{ such that } |x - y| \leq 2r\}$. Let N_V be the (random) number of discs in an optimal covering of the points in V. If $W \supseteq \tilde{V}$ is another such subset then, given an optimal covering algorithm of W, let $N_W|_V$ be the number of discs in the covering of W that also cover points in V.

Lemma 6.4 *Let A and B be bounded subsets of the plane such that $B \supseteq \tilde{A}$. Fix optimal covering algorithms for $\tilde{A}$ and B. Then for every point configuration $N_B|_A \leq N_{\tilde{A}}$.*

Proof of Lemma 6.4 Notice first that we can find a covering of B by covering all the points in $\tilde{A}$ optimally and then use an optimal covering of B to cover the points in $B/\tilde{A}$. This gives us:

$$N_B \leq N_{\tilde{A}} + N_B|_{B/\tilde{A}}.$$

Secondly, notice that any point in $B/\tilde{A}$ is at least $2r$ from any point in A. This implies that each disc used in an optimal covering of B can cover points in only one of these sets. Thus,

$$N_B|_A + N_B|_{B/\tilde{A}} \leq N_B.$$

Combining these inequalities gives us the result. □

Let us take a box, B_t, as A in the lemma, and B_{J_i} as B, where $B_{t+2r} \subset B_{J_i}$. Then the lemma says that the number of discs covering points in B_t under an optimal covering of B_{J_i} is at most the number of discs covering

points in B_{t+2r} under an optimal covering of this area. We note that this last quantity is at most the optimal number of discs covering B_t plus the number of points in B_{t+2r}/B_t (call this number P_t). Hence,

$$N_{B_j}|_{Bt} \leq N_{B_{t+2r}} \leq N_{B_t} + P_t.$$

Let M_t be the number of discs in the covering that we claim is optimal that cover points in B_t. Note that $M_t = N_{B_j}|_{B_t}$ for all $j > t+2r$. Therefore we can write:

$$\frac{N_{B_t}}{t^2} \leq \frac{M_t}{t^2} \leq \frac{N_{B_t} + P_t}{t^2}.$$

Notice that $\frac{P_t}{t^2}$ goes to 0 as t goes a.s. to infinity, as follows from, for example, Chebyshev's inequality.

It follows that to prove Theorem 6.3 it is enough to show:

Lemma 6.5

$$\lim_{t\to\infty} \frac{N_{B_t}}{t^2} = \delta^{opt}, \ \ a.s.$$

Proof of Lemma 6.5 We use a continuous sub-additive ergodic theorem. See Akcoglu and Krengel (1981), in particular Theorem 2.8, for more details.

Let N_R be the minimum number of discs required to cover all the points in the rectangle R, and suppose we have two disjoint rectangles, R_1 and R_2. Then we note that:

$$N_{R_1 \cup R_2} \leq N_{R_1} + N_{R_2}.$$

It then follows that

$$\lim_{t\to\infty} \frac{N_{B_t}}{t^2}$$

exists and is a constant, a.s. Proposition 6.2 together with the fact that $N_{B_t}/t^2, t > t'$ is bounded above, for any $t' > 0$, gives by dominated convergence that this limit is δ. □

7 Open Problems

We would like to mention a number of open problems:

- For which classes of algorithms does there exist a critical density? By this we mean a a critical value λ_c, such that percolation occurs for $\lambda > \lambda_c$ and does not occur for $\lambda < \lambda_c$.

- For which classes of algorithms is the infinite cluster unique? In other words, when do we have either 0 or 1 infinite cluster, a.s.?

- We have shown in Theorem 4.4 that if we have a stationary algorithm with a finite horizon and a bounded density of discs then we must have percolation for λ high enough. We have also shown in Theorem 4.6 that we can have a shift invariant covering algorithm with an unbounded density of discs and no finite horizon that does not percolate, even for high values of λ. Do we have percolation for λ high enough for a finite horizon algorithm with an unbounded density of discs? Do we have percolation for λ high enough if we have a bounded density of discs but no finite horizon?

Acknowledgments. The authors would like to thank Richard Gill, who suggested the intuitive interpretation of Theorem 5.1, reported in Section 5. Massimo Franceschetti would like to thank his fellow student Matthew Cook, for many interesting discussions, and for reading and providing feedback on some of the proofs.

References

[1] P.J. Agarwal and M. Sharir (1998). Efficient Algorithms for Geometric Optimization. *ACM Computing Surveys 30(4), 412-458.*

[2] M.A. Akcoglu and U. Krengel (1981). Ergodic Theorems for superadditive processes. *Journal für die Reine und Angewandte Mathematik 323, 53-67.*

[3] P. Erdős and A. Rényi (1959). On Random Graphs I. *Publicationes Mathematicae Universitatis Debreceniensis 6, 290-297.*

[4] P. Erdős and A. Rényi (1960). On the Evolution of Random Graphs. *Publications of the Mathematical institute of the Hungarian Academy of Sciences 5,17-61.*

[5] P. Erdős and A. Rényi (1961). On the Evolution of Random Graphs. *Bulletin of the International Institute of Statistics of Tokio 38, 343-347.*

[6] P. Erdős and A. Rényi (1961). On the Strength of Connectedness of a Random Graph. *Acta Mathematica Acadamiae Scientiarum Hungaricae 12,261-267.*

[7] M. Franceschetti, M. Cook and J. Bruck (2001). A Geometric Theorem for Approximate Disk Covering Algorithms. Submitted for publication. Available online at: `http://www.paradise.caltech.edu/ETR.html`.

[8] M. Gerla and J.T. Tsai (1995). Multicluster, mobile, multimedia radio network. *ACM/Baltzer Wireless Networks 1(3), 255-265.*

[9] E. N. Gilbert (1961). Random Plane Networks. *Journal of the Society for Industrial and Applied Mathematics 9(4), 533-543.*

[10] 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. McEneany, G. Yin, and Q. Zhang editors, Birkhäuser, Boston.*

[11] P. Gupta and P.R. Kumar (2000). The Capacity of Wireless Networks. *IEEE Transactions on Information Theory 46(2), 388-404.*

[12] P. Hall (1985). On Continuum Percolation. *The Annals of Probability 13(4), 1250-1266.*

[13] P. Hall (1988). Introduction to the Theory of Coverage Processes. *J. Wiley and sons, New York.*

[14] H. Kesten (1980). Percolation Theory for Mathematicians, *Birkhäuser.*

[15] R. Meester and R. Roy (1994). Uniqueness of Unbounded Occupied and Vacant Components in Boolean Models. *The Annals of Applied Probability 4(3), 933-951.*

[16] R. Meester and R. Roy (1996). Continuum Percolation. *Cambridge University Press.*

[17] M. V. Menshikov (1986). Coincidence of Critical Points in Percolation Problems. *Soviet Mathematics Doklady 24, 856-859.*

[18] R. Roy (1990). The Russo-Seymour-Welsh Theorem and the Equality of the Critical Densities and Critical Dual Densities for Continuum Percolation on $\mathbb{R}^2$. *The Annals of Probability 18, 1563-1575.*

Mathematical Institute
Utrecht University
PO Box 80010
3508TA Utrecht
The Netherlands
booth@math.uu.nl

California Institute of Technology
Moore Laboratory of Electrical Engineering
m.c. 136-93 Pasedena CA 91125
USA
massimo@paradise.caltech.edu
bruck@paradise.caltech.edu

Faculty of Exact Sciences
Free University of Amsterdam
De Boelelaan 1083a
1081 HV Amsterdam
The Netherlands
rmeester@cs.vu.nl