

Sentry selection in wireless networks

Paul Balister, Béla Bollobás, Amites Sarkar, and Mark Walters

Abstract Let \mathcal{P} be a Poisson process of intensity one in the infinite plane \mathbb{R}^2 . We surround each point x of \mathcal{P} by the open disc of radius r centred at x . Now let S_n be a fixed disc of area n , and let $C_r(S_n)$ be the set of discs which intersect S_n . Write E_r^k for the event that $C_r(S_n)$ is a k -cover of S_n , and F_r^k for the event that $C_r(S_n)$ may be partitioned into k disjoint single covers of S_n . We prove that $\mathbb{P}(E_r^k \setminus F_r^k) \leq \frac{c_k}{\log n}$, and that this result is best possible. We also give improved estimates for $\mathbb{P}(E_r^k)$. Finally, we study the obstructions to k -partitionability in more detail. As part of this study, we prove a classification theorem for (deterministic) covers of \mathbb{R}^2 with half-planes that cannot be partitioned into two single covers.

Paul Balister

Department of Mathematical Sciences, University of Memphis, Memphis, TN 38152, USA,
e-mail: pbalistr@memphis.edu

Supported by NSF grant CCF-0728928.

Béla Bollobás

Trinity College, Cambridge CB2 1TQ, UK, and Department of Mathematical Sciences,
University of Memphis, Memphis, TN 38152, USA, e-mail: B.Bollobas@dpmms.cam.ac.uk
Supported by ARO grant W911NF-06-1-0076, and NSF grants CNS-0721983, CCF-
0728928 and DMS-0906634.

Amites Sarkar

Department of Mathematics, Western Washington University, Bellingham, WA 98225,
USA, e-mail: amites.sarkar@wwu.edu

Mark Walters

School of Mathematical Sciences, Queen Mary, University of London, Mile End Road,
London E1 4NS, UK, e-mail: M.Walters@qmul.ac.uk

1 Introduction

Questions involving coverage arise frequently in engineering. Here is a typical example. A network of wireless sensors is deployed in a large area (the sensing region). Each sensor can detect any event occurring within distance r of itself. If the sensor locations are random (we will model them as points of a unit intensity Poisson process), how large should r be to ensure that any event occurring anywhere in the sensing region will be detected by at least one sensor? A good way to visualize, and indeed analyze, this problem is to imagine that we place discs of radius r around each sensor. Now we simply ask for the smallest r for which these discs cover the sensing region. A natural variant of this problem is to ask for the smallest r for which these discs k -cover the region, i.e., we require each point in the sensing region to be within range of at least k sensors. This may be more useful in applications as it gives a degree of tolerance in case of unreliable sensors, or the ability to locate an event by triangulation. It is worth noting that for fixed k and large sensing regions, k -coverage usually occurs for r only slightly more than that which is necessary for (single) coverage.

These problems have received much attention [5, 6, 8, 9, 10]. In this paper we consider a slightly different problem. Suppose that we wish to devise a rota system so that each sensor can *sleep* for most of the time, for example, to extend battery life. A natural way of doing this would be to partition the set of sensors into k groups, and arrange that only the sensors in group ℓ are active in the ℓ^{th} time slot. After k time slots have expired, we repeat the process. In order to detect an event occurring anywhere and at any time, it is necessary that the sensors in each group themselves form a single cover of the sensing region. Thus our question becomes: for fixed k , how large should r be to ensure that the sensors can be partitioned into k groups, each of which covers the sensing region? We call this the problem of *sentry selection*, since each of the groups is a group of sentries keeping watch over the region while the others are sleeping.

For the discs to be partitioned into k single covers (k -partitionability), it is clearly necessary that they k -cover. However, it is important to note that a k -cover of an arbitrary set cannot always be partitioned into k single covers. For instance, let S be the set of all subsets of $A = \{1, 2, \dots, n\}$ of size k . The n sets $S_i = \{B \in S : i \in B\}$, $1 \leq i \leq n$, form a k -cover of S which cannot even be partitioned into two single covers if $n \geq 2k - 1$. This example shows that a solution to our problem must make some use of its geometric setting. Also, even restricting ourselves to discs of equal radii, it is possible to construct k -covers of the plane that are not $\lceil(2k+2)/3\rceil$ -partitionable (see Section 9). Thus we shall also make use of the probabilistic setting.

Let us formalize our problem. Consider a Poisson process \mathcal{P} of intensity one in the infinite plane. Thus for any bounded measurable region A , the number of points in $A \cap \mathcal{P}$ is given by a Poisson random variable with mean equal to the area $|A|$, and is independent of the number and location of

points of \mathcal{P} in any disjoint region. We surround each point x of \mathcal{P} by the open disc $D_r(x)$ of radius r centred at x . Now for definiteness, let the sensing region S_n be the open disc of area n centred at the origin. (For all the results in this paper, it is enough that S_n be an open connected region of area n such that boundary ∂S_n is of length $O(n^{1-\delta})$ for some $\delta > 0$ and ∂S_n does not intersect any circle $\partial D_r(x)$ in more than a bounded number of points.) Consider the set $C_r(S_n)$ of discs $D_r(x)$, $x \in \mathcal{P}$, which intersect S_n . We wish to choose $r = r(n)$ as small as possible so that, with high probability, we may partition $C_r(S_n)$ into k classes such that each point $y \in S_n$ is contained in a disc from each class. Here, and throughout the paper, the phrase “with high probability”, abbreviated to **whp**, means “with probability tending to one as $n \rightarrow \infty$ ”. Note that some of the sensors defining $C_r(S_n)$ may lie outside of S_n . This slightly unusual method of dealing with the boundary is chosen to simplify the analysis of boundary effects, while remaining applicable to real-life situations.

The basic disc model we are considering was introduced by Gilbert [4] in 1961 in the context of percolation. His subsequent paper [5] on coverage contains the following important observation. In order for a family of discs of equal radii to cover a region without boundary, say a large torus, it is not only necessary but also sufficient that every intersection of two disc boundaries is contained in a third disc, provided there is at least one such intersection. Note that the intersection itself is not contained in either of the first two discs, since they are *open*. This fact, and its generalization to k -coverage, has been used in almost all subsequent work on the problem: it enabled Hall [6] to prove sharp results for k -coverage in d dimensions (see below). We use this observation in the following general form.

Lemma 1. *If S is a bounded open connected planar region, then S is k -covered by $C_r(S)$ provided every intersection point of the boundaries of two discs in $C_r(S)$ that lies in S , and every intersection of a boundary of some disc in $C_r(S)$ with the boundary of S , is k -covered, and at least one such intersection (of either type) exists.*

Proof. Suppose these conditions hold, but that S is not k -covered, and let R be a connected component of the subset of S that is not k -covered. Then the boundary ∂R consists only of points from ∂S or from some $\partial D_r(x)$, $x \in \mathcal{P}$. However, no intersection point lies in ∂R as these would then fail to be k -covered. Thus each component of ∂R is equal to either a component of ∂S or the whole of some $\partial D_r(x)$. We may remove $D_r(x)$ from $C_r(S)$ whenever $\partial R \supseteq \partial D_r(x)$ without removing or uncovering any intersection points since no other disc can intersect $D_r(x)$ without generating an intersection point on $\partial D_r(x)$. (Here we use the fact that the discs all have the same radii.) Thus we may assume $\partial R \subseteq \partial S$. But in this case any point in R may be joined to any point in S by a path not meeting ∂R . Thus $R = S$ and there are no intersection points inside S or on the boundary of S .

Throughout the paper we will ignore events of probability zero, so that, for instance, we will assume that an intersection of two disc boundaries does not lie on the boundary of a third disc.

2 Results

Let $n, r \in \mathbb{R}$. For $k \in \mathbb{N}$, write E_r^k for the event that $C_r(S_n)$ is a k -cover of S_n , and F_r^k for the event that $C_r(S_n)$ may be partitioned into k single covers of S_n . Our main result is the following.

Theorem 1. *With $r, n \in \mathbb{R}$, $k \in \mathbb{N}$,*

$$\mathbb{P}(E_r^k \setminus F_r^k) \leq \frac{c_k}{\log n}.$$

In Section 5 we prove this theorem, and in Section 7 we show that this result is best possible (up to the value of c_k). In Section 6 we prove two hitting time versions of Theorem 1: if we fix n and slowly increase r , or if we fix r and add points uniformly at random to a given area, then with high probability, k -partitionability occurs as soon as we have k -coverage. The proofs identify the principal obstructions to k -partitionability in a k -cover as certain small non-partitionable k -covered configurations. These configurations involve a small area which is covered by $k - 2$ common discs, but for which the remaining discs form a 2-cover which is not 2-partitionable. Since these configurations are very small, the curvature of the discs forming them is negligible, so that our obstructions are essentially 2-covers with half-planes which cannot be partitioned into two single covers. It is therefore of interest to classify such configurations. Such a classification is achieved in Theorem 9, whose proof occupies Section 8.

For $C_r(S_n)$ to admit a k -partition, it must certainly form a k -cover of S_n . Hall [6] proved the following result, which shows that the area of the sensing regions required for k -coverage is only slightly more than that required for single coverage. In particular, we do not need k times the disc area for single coverage to ensure k -coverage, as one might naively expect.

Theorem 2 (Hall [6]). *Let $k \in \mathbb{N}$ and let r be given by*

$$\pi r^2 = \log n + k \log \log n + f(n).$$

*Then $C_r(S_n)$ is a k -cover of S_n **whp** if and only if $f(n) \rightarrow \infty$ as $n \rightarrow \infty$.*

In Section 4 we strengthen this result by proving the following.

Theorem 3. *Fix $k \in \mathbb{N}$ and let r be given by*

$$\pi r^2 = \log n + k \log \log n + f(n).$$

Then

$$\mathbb{P}(E_r^k) = e^{-e^{-f(n)+o(1)}/(k-1)!} + O\left(\frac{1}{\log n}\right) \quad \text{as } n \rightarrow \infty.$$

In particular, if $f(n) = t$ is a constant, then $\mathbb{P}(E_r^k) \rightarrow e^{-e^{-t}/(k-1)!}$ as $n \rightarrow \infty$. After proving Theorem 3, we discovered that this corollary can be read out of results of Janson [7].

3 Thinly covered regions

Call a point ℓ -thinly covered (or just *thinly covered* if ℓ is clear from the context) if it is not ℓ -covered by the discs $D_r(x)$, $x \in \mathcal{P}$. An ℓ -thinly covered region is a connected component of the set of points in S_n that are ℓ -thinly covered. Note that if an ℓ -thinly covered region is not $(\ell - 1)$ -covered then it will contain several *atomic regions*, i.e., components of $S_n \setminus \bigcup_{x \in \mathcal{P}} \partial D_r(x)$.

Let X° be the number of intersection points of the boundaries of the discs of $C_r(S_n)$ that lie inside S_n , and let X^∂ be the number of intersection points of ∂S_n with the boundaries of the discs of $C_r(S_n)$. Write $X = X^\circ + X^\partial$ for the total number of *intersection points*. Let X_k (respectively, X_k° , X_k^∂) be the number of intersection points (respectively, *interior intersection points*, *boundary intersection points*) that are k -thinly covered. Lemma 1 can now be restated as saying that if $X_k = 0$ and $X > 0$ then S_n is k -covered.

Lemma 2. Fix $k \in \mathbb{N}$ and write

$$\pi r^2 = \log n + k \log \log n + t,$$

where $t = t(n) = o(\log n)$. Then as $n \rightarrow \infty$,

$$\mathbb{E}(X^\circ) = 4\pi r^2 n, \quad \mathbb{E}(X^\partial) = n^{1-\Omega(1)}, \quad \mathbb{P}(X = 0) = n^{-\omega(1)},$$

$$\mathbb{E}(X_k) = \frac{(4 + o(1))e^{-t}}{(k-1)!} \quad \text{and} \quad \mathbb{E}(X_k^\partial) = n^{-\Omega(1)}.$$

Proof. A fixed disc boundary $\partial D_r(x)$ with centre $x \in \mathcal{P}$ intersects twice each disc boundary whose centre lies within distance $2r$ of x . Therefore, the expected number of intersections involving $\partial D_r(x)$ is $8\pi r^2$. The density of intersections in the infinite plane is thus $4\pi r^2$ (since each intersection is counted twice), and so $\mathbb{E}(X^\circ) = 4\pi r^2 |S_n| = 4\pi r^2 n$. For any intersection point $u \in \partial S_n$, there must be a point $x \in \mathcal{P}$ at distance r from u , so $x \in \partial D_r(u)$. Since each such x gives rise to a bounded number of intersection points $u \in \partial S_n$, and the total area A swept out by $\partial D_r(u)$ as u moves around ∂S_n is at most $2\pi r |\partial S_n| = O((\log n)^{1/2} n^{1-\delta}) = n^{1-\Omega(1)}$, we have $\mathbb{E}(X^\partial) = n^{1-\Omega(1)}$. The area A includes all points within distance r of ∂S_n when the diameter $\text{diam}(S_n)$ is greater than $2r$, however $\text{diam}(S_n) = \Omega(\sqrt{n})$, so this holds for sufficiently large n . If $\text{diam}(S_n) = \|x - y\|$ with $x, y \in \partial S_n$,

then any line perpendicular and crossing the line segment from x to y must intersect ∂S_n , and hence intersects A in an interval of length at least $2r$. Thus $|A| \geq 2r \operatorname{diam}(S_n)$ and the probability that there is no boundary intersection point is $\mathbb{P}(\operatorname{Po}(|A|) = 0) = e^{-|A|} \leq e^{-\Omega(r\sqrt{n})} = n^{-\omega(1)}$. Hence $\mathbb{P}(X = 0) = n^{-\omega(1)}$.

A fixed intersection u needs at least k other points of \mathcal{P} within distance r to be k -covered. Conditioning on the locations of the point(s) of \mathcal{P} giving rise to u does not change the distribution of the remaining points of \mathcal{P} . (This is a key property of a Poisson process.) Therefore, the probability p that u is not k -covered may be estimated by

$$\begin{aligned} p &= \mathbb{P}(\operatorname{Po}(\pi r^2) < k) \\ &= \sum_{i=0}^{k-1} e^{-\pi r^2} (\pi r^2)^i / i! \\ &= e^{-\log n - k \log \log n - t} (\pi r^2)^{k-1} \frac{1}{(1 - O(k/\pi r^2))(k-1)!} \\ &= (\pi r^2 n)^{-1} e^{-t} (\pi r^2 / \log n)^k \frac{1}{(1 - O(k/\pi r^2))(k-1)!} \\ &= (\pi r^2 n)^{-1} e^{-t} \frac{1 + o(1)}{(k-1)!}, \end{aligned}$$

where the third line follows by comparison with a geometric series, and the last line uses the fact that $t = o(\log n)$ and hence $\pi r^2 = (1 + o(1)) \log n$. Consequently $\mathbb{E}(X_k) = (4\pi r^2 n + n^{1-\Omega(1)})p = (4 + o(1))e^{-t}/(k-1)!$ and $\mathbb{E}(X_k^2) = n^{1-\Omega(1)}p = n^{-\Omega(1)}$ as required.

Note that for $\pi r^2 = (1 + o(1)) \log n$, $\mathbb{E}(X_k)/\mathbb{E}(X_{k-1}) = \frac{1+o(1)}{k-1} \log n = \Theta(\log n)$, so one would expect many regions to fail to be k -covered before any region fails to be $(k-1)$ -covered.

Next we need a simple lemma bounding the tail of the Poisson distribution.

Lemma 3. Fix $c > 0$, and set

$$c_- = ce^{-1-1/c}.$$

Then

$$\mathbb{P}(\operatorname{Po}(c_- \log n) \geq c \log n) = n^{-1-\Omega(1)}.$$

Proof. Write $k = \lceil c \log n \rceil$ and $\lambda = c_- \log n$. Then, by comparison with a geometric series, the probability can be estimated as

$$\begin{aligned}
\sum_{l=k}^{\infty} e^{-\lambda} \frac{\lambda^l}{l!} &< e^{-\lambda} \frac{1}{1-(\lambda/k)} \frac{\lambda^k}{k!} \\
&< e^{-\lambda} \frac{1}{1-(\lambda/k)} \left(\frac{\lambda e}{k}\right)^k \\
&= \frac{1}{1-(c_-/c)} n^{c(\log(c_-/c)+1)-c_-+o(1)},
\end{aligned}$$

which is $n^{-1-\Omega(1)}$ provided

$$c_- < c \quad \text{and} \quad c \log(c_-/c) + c - c_- < -1.$$

However this holds for c_- as in the statement of the theorem.

We use this lemma to show that **whp** the intersections are not too clustered.

Lemma 4. *If $\pi r^2 = (1+o(1)) \log n$, then with probability $1-n^{-\Omega(1)}$ there are at most $677(\log n)^2$ intersections within distance $2r$ of any other intersection.*

Proof. Fix two points $x, y \in \mathcal{P}$ (or one point $x \in \mathcal{P}$ and ∂S_n) giving rise to an intersection u . Conditioning on the locations of x and y , the distribution of the remaining points in \mathcal{P} is still given by a Poisson process. Any intersection within distance $2r$ of u is defined by (at most) two points of \mathcal{P} , each within distance $3r$ of u . By Lemma 3 (with $c = 26$, which gives $c_- > 9$),

$$\mathbb{P}(\text{Po}(9\pi r^2) \geq 26 \log n) = n^{-1-\Omega(1)}.$$

By Lemma 2, the expected number of intersection points in (or on the boundary of) S_n is $4\pi r^2 n + n^{1-\Omega(1)} = n^{1+o(1)}$. It follows that the expected number of intersection points with more than $26 \log n + 2$ points of \mathcal{P} within distance $3r$ is $n^{-\Omega(1)}$. Thus with probability $1 - n^{-\Omega(1)}$, there is no intersection point with more than $26 \log n + 2$ points of \mathcal{P} within distance $3r$, and hence no intersection point with more than $677(\log n)^2$ intersections within distance $2r$ (for sufficiently large n).

Typically, most of S_n is very heavily covered by the discs of $C_r(S_n)$. It seems reasonable that these areas are unlikely to cause any obstruction to k -partitionability. Thus most of the work is involved in partitioning the discs covering lightly covered regions. Fortunately these regions occur in well-separated groups.

Let $\ell = 3k \log \log n$ and define ϵ by

$$\epsilon = \frac{2\ell \log \log n}{\sqrt{\log n}} = \frac{6k(\log \log n)^2}{\sqrt{\log n}}. \quad (1)$$

Note that $\epsilon \rightarrow 0$ as $n \rightarrow \infty$.

On each disc $D_r(x)$ in $C_r(S_n)$ define seven *marked points* equally spaced around its boundary $\partial D_r(x)$, say at angles $2\pi j/7$, $0 \leq j < 7$, from the x -axis.

For technical reasons we shall need to include these as well as intersection points in the following lemma.

Lemma 5. *Assume $\pi r^2 = (1 + o(1)) \log n$, $\ell = 3k \log \log n$, and ϵ is defined as in (1). Then the expected number of ℓ -thinly covered points that are intersections or marked points and are within distance $4r + \epsilon$ of ∂S_n is $(\log n)^{-\omega(1)} \mathbb{E}(X_k)$. The expected number of pairs of ℓ -thinly covered points u and v , each of which are either an intersection point or a marked point, and with $\epsilon \leq \|u - v\| \leq 4r + 2\epsilon$ is also $(\log n)^{-\omega(1)} \mathbb{E}(X_k)$.*

Proof. The number of intersection and marked points within distance $4r + \epsilon$ of ∂S_n is $O(r^3 |\partial S_n|) = n^{1-\Omega(1)}$ and so, as in the proof of Lemma 2, the expected number of these that are ℓ -thinly covered is $n^{1-\Omega(1)} \mathbb{P}(\text{Po}(\pi r^2) < \ell)$. Now

$$\begin{aligned} \mathbb{P}(\text{Po}(\pi r^2) < \ell) &= e^{-\pi r^2} \frac{1}{1 - O(\ell/\pi r^2)} \frac{(\pi r^2)^{\ell-1}}{(\ell-1)!} \\ &\leq n^{-1+o(1)} (\log n)^{(3+o(1))k \log \log n} \\ &= n^{-1+o(1)+(3+o(1))k(\log \log n)^2 / \log n}, \\ &= n^{-1+o(1)}, \end{aligned}$$

where we have used the fact that $\pi r^2 = (1 + o(1)) \log n$ and $\ell = o(\log n)$. Thus the expected number of ℓ -thinly covered intersection or marked points within distance $4r + \epsilon$ of ∂S_n is $n^{-\Omega(1)}$, which is $(\log n)^{-\omega(1)} \mathbb{E}(X_k)$ by Lemma 2. For the second part note that such pairs (u, v) are defined by (at most) four points x_1, \dots, x_4 of \mathcal{P} , any two of which are at most $6r + 2\epsilon \leq 7r$ apart (for large n). Now x_1 lies within distance r of S_n , so there are $\text{Po}((1 + o(1))n)$ choices for x_1 . Conditioned on the choices of x_1, \dots, x_{i-1} , there are at most $\text{Po}(49\pi r^2)$ choices for x_i (using once again the fact that conditioning on the existence of certain points in a Poisson process leaves the remaining points occurring with an identical Poisson process distribution). Thus the expected number of ordered 4-tuples (x_1, \dots, x_4) of distinct points is bounded by $(1 + o(1))n(49\pi r^2)^3 = O(n(\log n)^3)$. Similarly, the number of ordered 3, 2, or 1-tuples is also $O(n(\log n)^3)$. Since any choice of tuple corresponds to a bounded number of choices of pair (u, v) , there are on average $O(n(\log n)^3) = O((\log n)^2 \mathbb{E}(X))$ pairs (u, v) of marked or intersection points with $\epsilon \leq \|u - v\| \leq 4r + 2\epsilon$.

For any such pair u, v , we have

$$\begin{aligned} |D_r(u) \cup D_r(v)| &\geq \pi r^2 + 2\epsilon \sqrt{r^2 - (\epsilon/2)^2} \\ &\geq \pi r^2 + \left(\frac{2}{\sqrt{\pi}} - o(1)\right) \epsilon \sqrt{\log n} \\ &\geq \pi r^2 + 2.2\ell \log \log n \end{aligned}$$

for sufficiently large n . Also $D_r(u) \cup D_r(v)$ must contain less than 2ℓ points of \mathcal{P} . Even conditioned on x_1, \dots, x_4 , the probability of this is at most

$$\begin{aligned}
\mathbb{P}(\text{Po}(\pi r^2 + 2.2\ell \log \log n) \leq 2\ell) &\leq e^{-\pi r^2 - 2.2\ell \log \log n} \frac{1}{1 - O(\ell/\pi r^2)} \frac{(\pi r^2)^{2\ell}}{(2\ell)!} \\
&\leq e^{-\pi r^2} \frac{(\pi r^2)^{k-1}}{(k-1)!} (\log n)^{2\ell - (k-1) + o(\ell) - 2.2\ell} \\
&\leq \mathbb{P}(\text{Po}(\pi r^2) < k) (\log n)^{-\omega(1)}
\end{aligned}$$

since $\ell \rightarrow \infty$ as $n \rightarrow \infty$. Since (by the proof of Lemma 2) $\mathbb{E}(X_k) = \mathbb{E}(X)\mathbb{P}(\text{Po}(\pi r^2) < k)$, there are on average $(\log n)^{-\omega(1)}\mathbb{E}(X_k)$ such pairs of ℓ -thinly covered intersection or marked points.

Lemma 6. *Assume $\pi r^2 = (1 + o(1)) \log n$, $\ell = 3k \log \log n$, and ϵ is defined as in (1). Then, with probability $1 - (\log n)^{-\omega(1)}\mathbb{E}(X_k)$, there are discs D_1, \dots, D_m in S_n , each of radius ϵ , such that all ℓ -thinly covered regions of S_n lie in $\bigcup_{i=1}^m D_i$, and each D_i lies at distance at least $4r$ from the other D_j s and from ∂S_n .*

Proof. By Lemma 5, with probability $1 - (\log n)^{-\omega(1)}\mathbb{E}(X_k)$ there are no thinly covered intersection or marked points u and v of S_n with $\epsilon \leq \|u - v\| \leq 4r + 2\epsilon$, and no thinly covered intersection or marked points within distance $4r + \epsilon$ of ∂S_n . Assume this holds, and choose inductively a sequence u_1, u_2, \dots, u_m of thinly covered intersection or marked points in S_n with $u_i \notin \bigcup_{j < i} D_\epsilon(u_j)$. This process will terminate (almost surely) with all thinly covered intersection and marked points in S_n lying in some $D_i = D_\epsilon(u_i)$. By assumption, $\|u_i - u_j\| \geq \epsilon$, so $\|u_i - u_j\| \geq 4r + 2\epsilon$ for all $i \neq j$. Hence each disc D_i is at distance at least $4r$ from any other disc D_j or ∂S_n . Now consider any thinly covered region R in S_n . The intersection and marked points on the boundary of R are thinly covered. But there are no thinly covered intersection points on ∂S_n . Thus the components of ∂R must each lie either entirely within ∂S_n or entirely within the interior of S_n . Consider a component of ∂R that lies inside S_n . This boundary is formed from arcs of discs. Since any arc between marked points on a disc has length less than $r - \epsilon$ (for sufficiently large n), no point on this component of ∂R can be at distance more than $r - \epsilon$ from a thinly covered marked or intersection point. Hence no point on this component of ∂R can be at distance more than r from some u_i . Thus $\partial R \subseteq D \cup \partial S_n$, where $D = \bigcup_{i=1}^m D_r(u_i)$. Hence either $R \subseteq D$ or $S_n \setminus D \subseteq R$. If we fix two points z_1 and z_2 arbitrarily, but deterministically, in S_n with $\|z_1 - z_2\| = 2r$, then at most one of these can lie in D , as they cannot both lie in the same $D_r(u_i)$, and $D_r(u_i)$ and $D_r(u_j)$ are more than $2r$ apart when $i \neq j$. Hence if $S_n \setminus D \subseteq R$ then at least one of z_1 and z_2 is thinly covered. But the probability that at least one of these (deterministically chosen) points is thinly covered is at most $2\mathbb{E}(X_k)/\mathbb{E}(X) = (\log n)^{-\omega(1)}\mathbb{E}(X_k)$. Thus we may assume $R \subseteq D$. As R is connected, this implies $R \subseteq D_r(u_i)$ for some i . The external boundary of R is formed from arcs of discs that curve inwards, since the exterior of R is covered more times than the interior. Thus R is contained within the convex hull of the intersection points on its boundary, all of which lie in $D_i = D_\epsilon(u_i)$. Thus $R \subseteq D_i$.

Next we shall show that usually, inside each D_i , the $3k$ -thinly covered regions are covered by $k - 1$ common discs from $C_r(S_n)$. In fact we shall prove a slightly stronger result. Define r' by the formula

$$r'^2 = r^2 - \epsilon^2 = r^2 - \frac{36k^2(\log \log n)^4}{\log n}, \quad (2)$$

where ϵ is defined by (1).

Lemma 7. *Fix k and assume $\pi r^2 = (1 + o(1)) \log n$. Then with probability $1 - O(\mathbb{E}(X_{k-1}))$, any set of $3k$ -thinly covered points of S_n that lie within a single disc of radius ϵ are covered by a common set of $k - 1$ discs of $C_{r'}(S_n)$.*

Here, the two families of discs $C_r(S_n)$ and $C_{r'}(S_n)$ are constructed from the same underlying instance of \mathcal{P} .

Proof. Assume otherwise, but that the conclusion of Lemma 6 holds. As $3k \leq \ell$ (for sufficiently large n), all $3k$ -thinly covered points lie in one of the discs D_i of Lemma 6. Since these discs are far apart from each other, we may without loss of generality assume that the single disc of radius ϵ is one of the D_i and the set of $3k$ -thinly covered points is the set R of all $3k$ -thinly covered points in D_i . Note that R might not be connected, but R is bounded by arcs of circles $\partial D_r(x)$, $x \in \mathcal{P}$. Since $D_i \setminus R$ is more highly covered than R , these boundary arcs curve inwards towards R . Thus R is contained within the convex hull of the $3k$ -thinly covered intersection points in D_i . We may therefore assume there are two thinly covered intersection points u and v in R with $\alpha = \|u - v\|$ equal to the diameter of R . Thus there are at most $6k - 2$ points of \mathcal{P} in $D_r(u) \cup D_r(v)$, and hence at most $6k - 2$ points in the slightly smaller region $A_1 = D_{r'}(u) \cup D_{r'}(v)$. If R does not have a common $(k - 1)$ -cover from $C_{r'}(S_n)$, then there are less than $k - 1$ points of \mathcal{P} in $A_2 = D_{r'-\alpha}(u) \cup D_{r'-\alpha}(v)$. For fixed u and v the probability of this occurring is at most

$$p = e^{-|A_2| - |A_1 \setminus A_2|} \sum_{i=0}^{k-2} \sum_{i+j \leq 6k-2} \frac{|A_2|^i |A_1 \setminus A_2|^j}{i! j!}$$

(even conditioned on the points of \mathcal{P} giving rise to u and v). Now $A_2 \subseteq D_r(u)$, so $|A_2| \leq \pi r^2$. Also $|A_1 \setminus A_2| \leq 2|D_{r'}(u) \setminus D_{r'-\alpha}(u)| \leq 4\pi r' \alpha$. Finally, for all $\alpha \leq 2r'$, $|A_2| + |A_1 \setminus A_2| = |A_1| \geq \pi r'^2 + \pi r' \alpha / 2$ (by concavity of $|A_1|$ as a function of $\alpha \in [0, 2r']$). Thus, setting $t = \pi r' \alpha / 2$ and recalling that $r'^2 = r^2 - o(1)$, we have

$$\begin{aligned}
p &\leq e^{-\pi r^2 + o(1) - t} \sum_{i=0}^{k-2} \sum_{i+j \leq 6k-2} \frac{(\pi r^2)^i}{i!} \frac{(8t)^j}{j!} \\
&\leq \left(e^{-\pi r^2} \sum_{i=0}^{k-2} \frac{(\pi r^2)^i}{i!} \right) \left(e^{o(1) - t} \sum_{j=0}^{6k-2} \frac{(8t)^j}{j!} \right).
\end{aligned}$$

Fixing u and letting v vary, we have on average $(4\pi r^2)(\pi \alpha^2) = (16r^2/r'^2)t^2$ choices for the intersection v with t -value less than t , and so $(32 + o(1))t dt$ choices of v with t -value between t and $t + dt$. Thus the expected number of intersections v satisfying the above conditions with a fixed u is at most

$$\left(e^{-\pi r^2} \sum_{i=0}^{k-2} \frac{(\pi r^2)^i}{i!} \right) \int_0^\infty e^{o(1) - t} \sum_{j=0}^{6k-2} \frac{(8t)^j}{j!} (32 + o(1))t dt. \quad (3)$$

But $\int_0^\infty e^{-t} t^{j+1} dt = (j+1)!$ for all j , so the integral above becomes

$$C_k = (32 + o(1)) \sum_{j=0}^{6k-2} 8^j (j+1) = O(k8^{6k}), \quad (4)$$

which is bounded for fixed k , independently of n for large enough n . Since the first factor in (3) is just the probability $\mathbb{P}(\text{Po}(\pi r^2) < k-1)$, the expected number of intersections v satisfying the above conditions with a fixed u is at most $C_k \mathbb{P}(\text{Po}(\pi r^2) < k-1)$.

Now letting u vary, we note that the expected number of (ordered) pairs (u, v) satisfying the conditions above is at most C_k times the expected number of intersections u that are not $(k-1)$ -covered, in other words, at most $C_k \mathbb{E}(X_{k-1})$. In particular, the probability of such a pair existing is $O(\mathbb{E}(X_{k-1}))$. Adding the probability of failure $(\log n)^{-\omega(1)} \mathbb{E}(X_k) = (\log n)^{-\omega(1)} \Theta(\log n) \mathbb{E}(X_{k-1}) = o(\mathbb{E}(X_{k-1}))$ from Lemma 6 gives the result.

We will also require a lemma of a somewhat more technical nature. Define a *bad lune* \mathcal{L} to be the intersection of two discs from $C_r(S_n)$ with the following properties: i) the diameter of \mathcal{L} is at most 2ϵ ; ii) \mathcal{L} lies at distance at least 2ϵ from ∂S_n ; and iii) some point in the interior of \mathcal{L} is covered less than k times by the discs of $C_r(S_n)$.

Lemma 8. *Assume $\pi r^2 = (1 + o(1)) \log n$. Then the expected number of k -thinly covered intersection points that lie within distance 2ϵ of a bad lune is at most $(\log n)^{-2+o(1)} \mathbb{E}(X_k)$. Also, with probability at least $1 - (\log n)^{-2+o(1)} \mathbb{E}(X_k)$ there are no bad lunes.*

Proof. A lune \mathcal{L} satisfying i) and ii) above is defined by two points of \mathcal{P} at distance between $2r$ and $2\sqrt{r^2 - \epsilon^2} \geq 2r - 2\epsilon^2/r$. There are on average $O(\epsilon^2 n)$ choices for two such points, and conditioning on their locations does not change the distribution of the remaining points. There are no boundary

intersection points in $D_{2\epsilon}(\mathcal{L}) = \{x : d(x, \mathcal{L}) < 2\epsilon\}$ as \mathcal{L} is at distance at least 2ϵ from ∂S_n . The area of $D_{2\epsilon}(\mathcal{L})$ is $O(\epsilon^2)$, and so the expected number of interior intersections in $D_{2\epsilon}(\mathcal{L})$ not involving the discs forming \mathcal{L} is $O(\epsilon^2 r^2)$ as in the proof of Lemma 2. There are also on average $O(\epsilon r)$ intersections involving the boundary of the lune. (A full boundary of $D_r(x)$ contains on average $2\pi(2r)^2$ intersection points and the lune consists of two segments of total length at most about 4ϵ , so contains on average at most about $(8\pi r^2)(4\epsilon/2\pi r)$ intersection points.) Finally there are 2 intersections forming the lune so there are in total on average $O(\epsilon^2 r^2 + \epsilon r + 2) = O(\epsilon^2 r^2)$ intersections in $D_{2\epsilon}(\mathcal{L})$.

Suppose we now fix a lune \mathcal{L} and an intersection point in $D_{2\epsilon}(\mathcal{L})$. Conditioning on this information, the distribution of the remaining points of \mathcal{P} not involved in \mathcal{L} or the intersection point is unaffected. Thus if the intersection point is not k -covered, there can be at most $k - 1$ of these remaining points of \mathcal{P} within distance r of the intersection. The probability of this is just $\mathbb{E}(X_k^\circ)/\mathbb{E}(X^\circ) = \mathbb{E}(X_k^\circ)/(4\pi r^2 n)$ by Lemma 2. Consequently, the expected number of k -thinly covered intersections in $D_{2\epsilon}(\mathcal{L})$ is $O((\epsilon^2 n)(\epsilon^2 r^2)\mathbb{E}(X_k^\circ)/(4\pi r^2 n)) = O(\epsilon^4 \mathbb{E}(X_k^\circ))$. But $\epsilon = (\log n)^{-1/2+o(1)}$, so the expected number of such intersections is at most $(\log n)^{-2+o(1)}\mathbb{E}(X_k)$.

To estimate the expected number of bad lunes we replace r by r' and S_n by \mathcal{L} and consider coverage and intersection points defined by $C_{r'}(\mathcal{L})$. Following the above argument we see that the expected number of intersection points in, or on the boundary of \mathcal{L} is $O(\epsilon^2 r'^2)$, even including the two endpoints of \mathcal{L} and some intersections outside of \mathcal{L} (which, strictly speaking, are no longer intersection points). If \mathcal{L} is not k -covered then one of these points will be thinly covered by Lemma 1. (If there are no intersections then \mathcal{L} is k -covered iff one of its endpoints is.) The probability that any one of these point is thinly covered is $\mathbb{P}(\text{Po}(\pi r'^2) < k) = (1 + o(1))\mathbb{P}(\text{Po}(\pi r^2) < k)$ as $r^2 = r'^2 + o(1)$ by (2). Thus the expected number of bad lunes is at most $O((\epsilon^2 n)(\epsilon^2 r'^2)\mathbb{E}(X_k^\circ)/(4\pi r^2 n))$ which is at most $(\log n)^{-2+o(1)}\mathbb{E}(X_k)$ as above. Hence with probability at least $1 - (\log n)^{-2+o(1)}\mathbb{E}(X_k)$ there are no bad lunes.

4 Probability of k -coverage

Proof (Proof of Theorem 3).

Write $\pi r^2 = \log n + k \log \log n + t$ and assume for the moment that $t = o(\log n)$. By Lemma 2, the average number of k -thinly covered intersection points is $\mathbb{E}(X_k) = (4 + o(1))e^{-t}/(k-1)!$. However, these intersection points are clearly not independent of one another. Indeed, by considering the boundary of a k -thinly covered region of S_n , if one thinly covered intersection exists, then there must be at least one other one, and so $\mathbb{P}(X_k = 1) = 0$. Define Y_k to be the number of k -thinly covered regions of S_n other than S_n itself. (If the

whole of S_n is thinly covered, we define Y_k to be zero as, for technical reasons, we wish to have at least one intersection point on every thinly covered region counted by Y_k .) We will show that $\mathbb{E}(Y_k) = (1 + o(1))e^{-t}/(k-1)!$ and Y_k is given by an approximately Poisson distribution. The result will then follow, at least for small t .

To show that Y_k has an approximately Poisson distribution, we must show that the k -thinly covered regions do not occur in clusters (for r in the range we are considering). This will follow from the fact that there is usually at most one k -thinly covered region in any one D_i , where the D_i are as in Lemma 6. This in turn follows from the observation that, within a D_i , the disc boundaries $\partial D_r(x)$ are almost straight lines, so that the discs $D_r(x)$ behave almost like half-planes. If they actually were half-planes, Lemma 7 would show that (usually) each D_i contains at most one k -thinly covered region. This follows since, having removed the $k-1$ half-planes from Lemma 7, the intersection of the complements of the remaining half-planes bordering the now uncovered region would actually be a single convex polygonal region. It turns out that for the discs to behave differently, two of them must form a bad lune, which we have already shown is extremely unlikely.

Let Y'_k be the number of k -thinly covered regions of S_n that i) do not border ∂S_n , ii) have diameter at most 2ϵ , and iii) do not contain an intersection point that is at one of the ends of a bad lune. Let X'_k be the number of intersection points on the boundary of such a region. Thus any intersection point counted by $X_k - X'_k$ is either within 2ϵ of ∂S_n , within 2ϵ of a bad lune, lies in a thinly covered region of diameter more than 2ϵ , or lies in the interior of a thinly covered region. By (the proof of) Lemma 2, the average number of thinly covered intersection points within 2ϵ of ∂S_n is $O(\epsilon n^{-\delta} \mathbb{E}(X_k^\circ) + \mathbb{E}(X_k^\partial)) = n^{-\Omega(1)} \mathbb{E}(X_k)$. By Lemma 8, the average number of thinly covered intersection points within 2ϵ of a bad lune is at most $O((\log n)^{-2+o(1)} \mathbb{E}(X_k))$. By Lemma 5, the average number of intersection points that are on the boundary of a thinly covered region with diameter more than 2ϵ , but not within 2ϵ of ∂S_n , is $(\log n)^{-\omega(1)} \mathbb{E}(X_k)$. (Any such intersection point is at distance between ϵ and $2r + \epsilon$ of at least one other intersection point on the boundary of the same region.) Finally, any intersection point in the interior of a k -thinly covered region must be $(k-2)$ -thinly covered, so the expected number of these is at most $\mathbb{E}(X_{k-2}) = \Theta((\log n)^{-2} \mathbb{E}(X_k))$. Hence

$$\mathbb{E}(X_k - X'_k) = O((\log n)^{-2+o(1)} \mathbb{E}(X_k)). \quad (5)$$

Since each thinly covered region has at least one intersection on its boundary, $Y_k - Y'_k \leq X_k - X'_k$, so $\mathbb{E}(Y_k - Y'_k) = O((\log n)^{-2+o(1)} \mathbb{E}(X_k))$. Finally, $X'_k \leq X_k^\circ \leq X_k$, so $\mathbb{E}(X_k^\circ - X'_k) = O((\log n)^{-2+o(1)} \mathbb{E}(X_k))$ as well.

Pick any internal intersection point u defined by points $x, y \in \mathcal{P}$. The event that it is thinly covered is independent of the choice of x and y . Moreover, the angle θ_u between the circles $\partial D_r(x)$ and $\partial D_r(y)$ at u depends only on x and y . As x and y vary, the average value is $\mathbb{E}(\theta_u \mid \|x - y\| < 2r) = \pi/2$.

(Fix x and average over $y \in D_{2r}(x)$ to obtain $\mathbb{E}(\theta_u \mid \|x - y\| < 2r) = (4\pi r^2)^{-1} \int_0^{2r} 4\pi x \sin^{-1}(x/2r) dx = \pi/2$.) The sum of the external angles of the intersection points on the boundary of an internal thinly covered region R is $2\pi + |\partial R|/r$, since the boundary arcs all have radius of curvature r and curve inwards toward R . However, if the diameter of R is less than 2ϵ , then $|\partial R| = O(\epsilon)$ since R is “almost” convex. Now if an intersection u on ∂R is covered exactly $k - 1$ times then the exterior angle at u is on average $\pi/2$, but if u is covered exactly $k - 2$ times the average exterior angle at u is $-\pi/2$. Nonetheless, since $\mathbb{E}(X'_{k-1}) = O(\mathbb{E}(X'_k)/\log n)$, we may write

$$(2\pi + O(\epsilon/r))\mathbb{E}(Y'_k) = (\frac{\pi}{2} + O(1/\log n))\mathbb{E}(X'_k)$$

and thus

$$\mathbb{E} \sum_{u \text{ thin}} \theta_u = \frac{\pi}{2} \mathbb{E}(X_k^o) = \frac{\pi}{2} \mathbb{E}(X_k^o - X'_k) + (2\pi + o(1))\mathbb{E}(Y'_k).$$

Hence by Lemma 2 and (5), $\mathbb{E}(Y'_k) = (1 + o(1))e^{-t}/(k - 1)!$.

To prove that Y'_k is approximately Poisson we use the Stein-Chen method. A simple form of it is given in Theorem 1 of [1], which immediately implies the following.

Theorem 4. *Let ξ_1, ξ_2, \dots be a countable collection of independent random variables and let Z_1, Z_2, \dots be a countable collection of Bernoulli random variables where Z_i is a function of the values of ξ_j , $j \in S_i$. Suppose that $\sum_i \mathbb{E}(Z_i) = \lambda$ and let $b_1 = \sum_{i,j: S_i \cap S_j \neq \emptyset} \mathbb{E}(Z_i)\mathbb{E}(Z_j)$, $b_2 = \sum_{i,j: S_i \cap S_j \neq \emptyset, i \neq j} \mathbb{E}(Z_i Z_j)$ and $Z = \sum_i Z_i$. Then for all r ,*

$$|\mathbb{P}(Z = r) - e^{-\lambda} \lambda^r / r!| \leq \frac{1 - e^{-\lambda}}{\lambda} (b_1 + b_2).$$

(Theorem 1 in [1] includes another term b_3 that bounds dependency when $S_i \cap S_j = \emptyset$, but in our case $b_3 = 0$.)

To make the collection of events countable, divide S_n up into a very fine grid and move all points of \mathcal{P} to their nearest grid point. It is clear that for fixed r, n, k , we can make the number of thinly covered regions in this discretized version equal to the number of thinly covered regions in the original with probability arbitrarily close to 1. The random variables ξ_i record whether or not the i^{th} grid point is occupied, and for every conceivable thinly covered region R_i satisfying i)–iii) above, we introduce a variable Z_i indicating whether this thinly covered region exists. The set S_i can be taken to be the set of grid points within distance $r + 2\epsilon$ of R_i . Then $\mathbb{E}(|Y'_k - Z|)$ can be made arbitrarily small.

We now bound b_1 and b_2 . First consider b_1 . Fix a potential thinly covered region R_i of diameter at most 2ϵ . Then $\sum_{j: S_i \cap S_j \neq \emptyset} \mathbb{E}(Z_j)$ is bounded by the expected number of thinly covered intersection points within a disc of radius

$2r + 4\epsilon$, which is $z = O(\mathbb{E}(X_k)r^2/n) = (\log n)^{-\omega(1)}\mathbb{E}(Z)$, uniformly in the choice of R_i . Thus $b_1 \leq \sum_i \mathbb{E}(Z_i)z = \mathbb{E}(Z)z = (\log n)^{-\omega(1)}\mathbb{E}(Z)^2$. Now consider b_2 . This counts the expected number of pairs of thinly covered regions satisfying i)–iii) that lie within distance $2r + 4\epsilon$ of each other. By Lemma 5, there is a contribution of $(\log n)^{-\omega(1)}\mathbb{E}(X_k)$ to b_2 from pairs of regions that are not contained within a single disc of radius ϵ . Then by the proof of Lemma 7, there is a contribution of at most $O(\mathbb{E}(X_{k-1})) = O((\log n)^{-1}\mathbb{E}(X_k))$ from pairs of regions that do not share a common $(k-1)$ -cover from $C_{r'}(S_n)$, where r' is defined by (2). (Any such pair gives rise to a pair of thinly covered intersection points u and v with no common $(k-1)$ -cover, and these are (over-) counted in the proof of Lemma 7.) Finally, assume two k -thinly covered regions R_i and R_j have a common $(k-1)$ -cover by elements of $C_{r'}(S_n)$ (and hence also by elements of $C_r(S_n)$). Thus the set of discs of $C_r(S_n)$ covering R_i is the same as the set of discs covering R_j . We shall now show that if R_i and R_j lie in a disc of radius ϵ , then at least one of them contains an endpoint of a bad lune, so this pair does not contribute to b_2 .

Let $u \in R_i$ and $v \in R_j$ be chosen with $\|u - v\|$ minimal. Then at least one of u or v must be an intersection point, say u . Let $D_r(x_1)$ and $D_r(x_2)$ be two discs whose boundaries intersect at u . Since these discs do not cover R_i , they also do not cover R_j . The regions R_j and R_i must be separated by the lune formed by $D_r(x_1)$ and $D_r(x_2)$, since otherwise u would not be the closest point in R_i to R_j . (Recall that R_i and R_j are contained in a region of diameter much smaller than r .) In particular, this lune is of diameter at most 2ϵ . The point u is k -thinly covered by $C_r(S_n)$ and hence any point within distance $r - r' > 0$ of u is k -thinly covered by $C_{r'}(S_n)$. Thus the lune is bad and R_i contains the endpoint u .

Now $b_1 = (\log n)^{-\omega(1)}\mathbb{E}(Z)^2$ and $b_2 = O((\log n)^{-1}\mathbb{E}(Z))$. Taking $r = 0$ in Theorem 4 we have

$$\mathbb{P}(Z = 0) = e^{-\lambda} + O((\log n)^{-1}) + (\log n)^{-\omega(1)}\lambda,$$

where $\lambda = \mathbb{E}(Z) = (1 + o(1))e^{-t}/(k-1)!$. By (5) and the fact that $\mathbb{P}(X'_k = 0) = \mathbb{P}(Y'_k = 0)$ is arbitrarily close to $\mathbb{P}(Z = 0)$, we have

$$\mathbb{P}(X_k = 0) = e^{-\lambda} + O((\log n)^{-1}) + O((\log n)^{-2+o(1)})\lambda.$$

Now by Lemma 2 and Lemma 1,

$$\mathbb{P}(k\text{-coverage succeeds}) = e^{-\lambda} + O((\log n)^{-1}) + O((\log n)^{-2+o(1)})\lambda + n^{-\omega(1)}. \quad (6)$$

For $1/\log n \leq \lambda \leq \log \log n$ we have $t = o(\log n)$ and so (6) is valid. But for this range of λ the last two error terms are dominated by the first and $\lambda = (1 + o(1))e^{-t}/(k-1)! = e^{-t+o(1)}/(k-1)!$, so

$$\mathbb{P}(k\text{-coverage succeeds}) = e^{-e^{-t+o(1)}/(k-1)!} + O((\log n)^{-1}). \quad (7)$$

This probability is monotonic in t , and for $\lambda = \log \log n$ it is $O((\log n)^{-1})$. Thus (7) holds for all $\lambda \geq \log \log n$. For $\lambda = 1/\log n$, the probability is $1 - O((\log n)^{-1})$, so (7) holds for all $\lambda \leq 1/(\log n)$. Thus (7) holds for all t .

5 k -Partitionability

We turn to our main goal, to prove that if $C_r(S_n)$ forms a k -cover of S_n , then it can almost always be partitioned into k single covers of S_n . We will suppose that $C_r(S_n)$ forms a k -cover, and colour the discs in $C_r(S_n)$ with k colours in two steps, the first deterministic and the second random. Our aim is for every region to be covered by discs of every colour; the colour classes will then form our desired partition.

The discs in $C_r(S_n)$ decompose S_n into regions R_i . A naive random partitioning runs into trouble with lightly covered regions, which have a good chance of not being fully covered. Fortunately these regions occur in well-separated groups by Lemma 6.

For the second, random, step of our colouring, we will need the following form of the Lovász local lemma.

Lemma 9. *Let A_1, \dots, A_n be events in a probability space. Suppose that for each i , the event A_i is independent of (the σ -algebra generated by) all but at most d other events. Suppose further that $\mathbb{P}(A_i) \leq p$ for all i , and that $pe(d+1) < 1$. Then there is a positive probability that no event A_i occurs.*

Proof. See [2].

The following theorem bounds the probability that $C_r(S_n)$ is a non- k -partitionable k -cover in terms of the probability that $C_r(S_n)$ is not a $(k-1)$ -cover.

Let $n, r \in \mathbb{R}$. Recall that, for $k \in \mathbb{N}$, E_r^k is the event that $C_r(S_n)$ is a k -cover of S_n , and F_r^k is the event that $C_r(S_n)$ may be partitioned into k single covers of S_n .

Theorem 5. *Let $k \in \mathbb{N}$ and suppose $\pi r^2 = \log n + k \log \log n + o(\log \log n)$. Then*

$$\mathbb{P}(E_r^k \setminus F_r^k) = O(\mathbb{E}(X_{k-1})).$$

Proof. Let B be the event that at least one of the conclusions of Lemma 4, Lemma 6, Lemma 7, or Lemma 8 fails. We shall also include in B some probability zero events, such as the event that $C_r(S_n)$ is infinite. We shall in fact show that $E_r^k \setminus F_r^k \subseteq B$, so that the result follows from the probability bounds given by these lemmas, namely

$$\begin{aligned} \mathbb{P}(B) &\leq o((\log n)^{-2}) + (\log n)^{-\omega(1)} \mathbb{E}(X_k) \\ &\quad + O(\mathbb{E}(X_{k-1})) + (\log n)^{-2+o(1)} \mathbb{E}(X_k), \end{aligned}$$

which is $O(\mathbb{E}(X_{k-1}))$ since by Lemma 2, $\mathbb{E}(X_k) = \Theta(\log n)\mathbb{E}(X_{k-1})$ and $\mathbb{E}(X_k) = (\log n)^{o(1)}$. Fix a configuration \mathcal{P} for which E_r^k holds, but B fails. It is enough to show k -partitionability of the cover $C_r(S_n)$.

Since B fails, we have (by Lemma 6) discs D_1, \dots, D_m of radius ϵ enclosing all the $(3k \log \log n)$ -thinly covered regions. Recall that, outside $\bigcup D_i$, each point of S_n is covered at least $3k \log \log n$ times, and that no D_i is within distance $4r$ of any of the others. We shall examine the D_i one by one. For a fixed D_i , we will colour at most $3k$ discs intersecting D_i with k colours so that all of D_i is covered by a disc of each colour. When we have done this for each D_i , we will complete the colouring randomly and apply Lemma 9 in conjunction with the conclusion of Lemma 4 to show that each colour class of the resulting colouring covers the whole of S_n with positive probability, and hence a good colouring exists. Note that in applying Lemma 9 we are randomly colouring a *fixed* configuration $C_r(S_n)$ (one based on a fixed instance of \mathcal{P}).

Our first task, then, is to k -colour some discs intersecting $D = D_1$ so that D is covered by a disc of each colour. For each disc $D_r(x)$ of $C_r(S_n)$ intersecting D , we replace $D_r(x)$ by a half-plane $H(x)$ as follows: if $\partial D_r(x)$ intersects D then the boundary of $H(x)$ is the straight line L through the intersection points $\partial D_r(x) \cap \partial D$ and we choose $H(x)$ to lie on the same side of L as the majority of $D_r(x)$. If $D_r(x) \supseteq D$, then $H(x)$ is any half-plane that covers $D_r(x)$ (and hence D). We also add half-planes surrounding (but not intersecting) D so that the exterior of D is covered at least as many times as the boundary of D (in particular, at least $3k \log \log n$ times). This will require the addition of at most a finite number of half-planes since $C_r(S_n)$ (and hence the number of intersections in $S_n \setminus D$) is finite. Note that the half-planes $H(x)$ still contain the discs $D_{r'}(x)$ of $C_{r'}(S_n)$.

We now consider the half-planes as a finite configuration of \mathbb{R}^2 , which is in fact a $(3k \log \log n)$ -cover outside of D . The boundaries of these half-planes divide \mathbb{R}^2 into (sometimes infinite) polygonal regions. We classify the polygonal regions in \mathbb{R}^2 into two types: the *red* regions, those covered at most $3k - 1$ times, and the *green* regions, which are covered at least $3k$ times.

Since B fails, we can assume (Lemma 7) that there are half-planes Π_1, \dots, Π_{k-1} each of which covers the entire red region. (Recall that $D_{r'}(x) \subseteq H(x)$ for every $D_r(x)$ intersecting D , and all the red regions are within D . If there are no red regions, then the Π_i can be chosen arbitrarily.) If we remove these half-planes, there are two cases.

Case 1. \mathbb{R}^2 is still covered by the remaining half-planes.

Remove Π_1, \dots, Π_{k-1} . Since the remaining half-planes cover \mathbb{R}^2 , there is a subset consisting of (at most) three half-planes $H(x_1)$, $H(x_2)$ and $H(x_3)$, which cover \mathbb{R}^2 (see Lemma 11 below). Hence $D_r(x_1)$, $D_r(x_2)$ and $D_r(x_3)$ cover D .

Case 2. The remaining half-planes do not cover \mathbb{R}^2 .

In this case the uncovered region must have been covered at most $k - 1$ times originally, so must lie inside D . The uncovered region is convex, so

is a polygonal region $P \subseteq D$. We convert the half-planes back to discs by decreasing their radii of curvature from ∞ to r continuously (while keeping the intersection of their boundaries with D fixed), and stop at the moment D is covered. This will indeed happen, since even if we remove the discs corresponding to the Π_i , which contain the half-planes within D , the remaining discs must cover D as D is k -covered. The last point y of D to be covered lies at the intersection of (almost surely, and if not, then include this event in B) three discs which correspond to three of our original discs $D_r(x_1)$, $D_r(x_2)$ and $D_r(x_3)$, whose intersection contains y . If $D_r(x_i)$, $i = 1, 2, 3$, do not cover ∂D , but have a common intersection in ∂D then y could not be the last uncovered point inside D . If two of the $D_r(x_i)$ cover disjoint arcs of ∂D , then as they intersect in y , these two discs form a lune inside D . This lune is bad as y is not k -covered by the half-planes and hence not k -covered by $C_{r'}(S_n)$. However, no bad lune exists since B fails (Lemma 8). The only other possibility is that the $D_r(x_i)$ cover ∂D . But then they cover D as each $D_r(x_i)$ covers the convex hull of $\partial D \cap D_r(x_i)$ and y .

In both case 1 and case 2 we colour the three discs $D_r(x_1)$, $D_r(x_2)$ and $D_r(x_3)$ with colour k .

Let $\mathcal{A} = \{H(x_1), H(x_2), H(x_3)\}$, $\mathcal{B} = \{\Pi_1, \dots, \Pi_{k-1}\}$, and let \mathcal{C} be the set of all the half-planes not in \mathcal{A} or \mathcal{B} . Now all the red regions are covered by each of the half-planes in \mathcal{B} and the green regions are covered at least $3k - 3 = 3(k - 1)$ times by $\mathcal{B} \cup \mathcal{C}$.

Now assume $k \geq 2$. The green regions are covered by at least $2(k - 1)$ half-planes from \mathcal{C} , and the red regions are covered by Π_{k-1} . Thus $\mathcal{C} \cup \{\Pi_{k-1}\}$ is a cover of \mathbb{R}^2 . By Lemma 11 we can find a subset of three of these half-planes that cover \mathbb{R}^2 . Colour these with colour $k - 1$ and remove them. If Π_{k-1} was not used then move it to the set \mathcal{C} . Now $|\mathcal{B}| = k - 2$ and all the red regions are covered by every half-plane in \mathcal{B} , while every green region is covered at least $3(k - 1) - 3 = 3(k - 2)$ times by $\mathcal{B} \cup \mathcal{C}$. Repeating this process we colour $3(k - 1)$ half-planes in total with $k - 1$ colours so that the whole of \mathbb{R}^2 is covered by half-planes of each colour.

Some of our coloured half-planes correspond to some discs of radius r : we colour these accordingly. (Some of the coloured half-planes may be additional half-planes we added to cover the exterior of D . These may be safely ignored since they do not contribute to the coverage of D .) Together with $D_r(x_1)$, $D_r(x_2)$ and $D_r(x_3)$, which receive colour k , we have covered every point of D with discs of all k colours.

Next we repeat this process for D_2, \dots, D_m . At each stage we are colouring different discs from those coloured in previous stages since no disc can intersect both D_i and D_j for $i \neq j$. Indeed, no disc coloured in one stage can even intersect a disc coloured in a different stage, since the distance between discs D_i is at least $4r$. When we have finished, we will have coloured some of the discs $C' \subseteq C_r(S_n)$, but the level of coverage provided by $C'' = C_r(S_n) \setminus C'$ outside the D_i is still at least $3k \log n - 3k$.

We now colour the discs of C'' randomly with k colours, each used with equal probability, and apply Lemma 9. Our “bad” events A_u will correspond to the intersection points u outside $\bigcup_i D_i$: indeed A_u will be the event that u is not covered by discs of all colours. For all u ,

$$\mathbb{P}(A_u) \leq k \left(1 - \frac{1}{k}\right)^{3k \log \log n - 3k} \leq e^3 k (\log n)^{-3} = p.$$

Also, A_u is independent of A_v whenever $\|u - v\| \geq 2r$, since then no disc can cover both u and v . Since B fails we know (Lemma 4) that, for each u , A_u is independent of all but at most $d = 677(\log n)^2$ other A_v . Since for sufficiently large n

$$pe(d+1) \leq e^4 k (\log n)^{-3} (677(\log n)^2 + 1) < 1$$

there is, by Lemma 9, a positive probability that no A_u occurs. Hence the required colouring of C'' exists, and we obtain our desired colouring of $C_r(S_n)$ on combining the two colourings above.

Proof (Proof of Theorem 1).

Throughout the proof, k will be fixed and all asymptotic notation will be as $n \rightarrow \infty$. We shall for simplicity assume S_n is a disc centred at the origin. Write

$$\pi r^2 = \log n + k \log \log n - \log(k-1)! - \log t.$$

Then by Theorem 3, $\mathbb{P}(E_r^k) = e^{-(1+o(1))t} + O(1/\log n)$. If $t > 2 \log \log n$ then $\mathbb{P}(E_r^k) = O(1/\log n)$ and so the conclusion holds automatically. If $t \leq 1/\log n$ then $\mathbb{P}(E_r^k) = 1 - O(1/\log n)$. If the conclusion holds for $t = 1/\log n$ then $\mathbb{P}(F_r^k) = 1 - O(1/\log n)$. Monotonicity of $\mathbb{P}(F_r^k)$ then implies that the conclusion of the theorem holds for all $t \leq 1/\log n$. Hence it is enough to prove the result when $1/\log n \leq t \leq 2 \log \log n$.

Write

$$\begin{aligned} S'_n &= \{(x, y) \in S_n : y > -\epsilon\}, \\ T'_n &= \{(x, y) \in S_n : y < -\epsilon - 2r\}, \\ S''_n &= \{(x, y) \in S_n : y < \epsilon\}, \\ T''_n &= \{(x, y) \in S_n : y > \epsilon + 2r\}, \end{aligned}$$

where ϵ is defined by (1). Let B be the event that the conclusions of either Lemma 4, Lemma 6, or Lemma 8 fail in S_n . Coverage events in S'_n are independent from those in T'_n , and similarly for S''_n and T''_n . Write I' for the event that S'_n is k -covered, but that there exist a set of $3k$ -thinly covered points in S'_n within a disc of radius ϵ that do not have a common $(k-1)$ -cover by elements of $C_{r'}(S_n)$. Let E' be the event that T'_n is k -covered. Define I'' and E'' analogously. If E_r^k occurs but F_r^k does not, then by the proof of Theorem 5, either B occurs, or there is a set of $3k$ -thinly covered points in S_n within a disc of radius ϵ that do not have a common $(k-1)$ -cover by elements of $C_{r'}(S_n)$. These points must lie entirely within S'_n or entirely within S''_n .

Thus if E_r^k holds, F_r^k fails and B also fails, one of I' or I'' must occur, and both E' and E'' must occur. Consequently,

$$\mathbb{P}(E_r^k \setminus F_r^k) \leq \mathbb{P}(I')\mathbb{P}(E') + \mathbb{P}(I'')\mathbb{P}(E'') + \mathbb{P}(B) = 2\mathbb{P}(I')\mathbb{P}(E') + \mathbb{P}(B). \quad (8)$$

Now

$$\begin{aligned} \mathbb{P}(B) &\leq o((\log n)^{-2}) + (\log n)^{-\omega(1)}\mathbb{E}(X_k) + (\log n)^{-2+o(1)}\mathbb{E}(X_k) \\ &\leq (\log n)^{-2+o(1)}, \end{aligned}$$

as $\mathbb{E}(X_k) = (4 + o(1))t = O(\log \log n)$.

Write X'_k and X'_{k-1} for the numbers of intersections in T'_n and S'_n respectively that are not covered k and $k-1$ times respectively. By Theorem 3 (noting that $|T'_n| = (\frac{1}{2} + o(1))n$)

$$\mathbb{P}(E') = \mathbb{P}(X'_k = 0) + n^{-\omega(1)} = e^{-(1/2+o(1))t} + O(1/\log n).$$

Also, by Lemma 2 we have

$$\mathbb{E}(X'_{k-1}) = (2 + o(1))t(k-1)/\log n.$$

Hence by Lemma 7

$$\mathbb{P}(I') = O\left(\frac{t}{\log n}\right). \quad (9)$$

It follows that

$$\mathbb{P}(E_r^k \setminus F_r^k) \leq \frac{c_k t e^{-(1/2+o(1))t}}{\log n} + O(1/\log n) + \mathbb{P}(B) \leq \frac{c'_k}{\log n}$$

which completes the proof.

6 Hitting time results

We consider two hitting time variants of Theorem 1. In the first we fix the points of the Poisson process in \mathbb{R}^2 , but increase r until k -coverage of S_n occurs. We show that with high probability, the instant we have k -coverage we also have k -partitionability. In the second variant we fix r and n and consider placing points uniformly at random in some large region containing S_n (thus effectively increasing the intensity of the Poisson process). We show that with high probability, the first point added that results in S_n being k -covered by the discs of radius r around these points, also results in this cover being k -partitionable. Here “with high probability” means with probability tending to 1 as $n \rightarrow \infty$ in the first case and as $n/r^2 \rightarrow \infty$ in the second. We start with a useful lemma.

Lemma 10. *Suppose $\pi r^2 = \log n + k \log \log n + o(\log \log n)$. Then with probability $1 - o(1)$ (as $n \rightarrow \infty$), $C_r(S_n)$ can be partitioned into $k - 1$ single covers $\mathcal{A}_1, \dots, \mathcal{A}_{k-1}$ of S_n , together with a collection \mathcal{A}_k of disks that cover all the points of S_n that are k -covered by $C_r(S_n)$. Moreover, we may assume \mathcal{A}_k covers ∂S_n .*

Proof. By Lemma 2, $\mathbb{E}(X_k) = o((\log n)^{1/2})$ and $\mathbb{E}(X_{k-1}) = o((\log n)^{-1/2})$. Following the proof of Lemma 7, we see that with probability at least $1 - o((\log n)^{1/2})\mathbb{E}(X_{k-1}) = 1 - o(1)$ all $(0.1 \log \log n)$ -thinly covered points within any single disk of radius ϵ are covered by a common set of $k - 1$ discs of $C_r(S_n)$. Indeed, the only change to the proof of Lemma 7 is in the expression (4) for C_k which is now bounded by $O((\log \log n)8^{0.2 \log \log n}) = o((\log n)^{1/2})$. Now we follow the proof of Theorem 5. We may assume that the conclusions of Lemma 4, Lemma 6, and Lemma 8 hold, since the probability that any of these fail is $o(1)$. We may also assume from the above that the conclusion of Lemma 7 holds with $(0.1 \log \log n)$ -thinly covered regions in place of $3k$ -thinly covered regions. Finally we may assume that $X_k \leq 0.1 \log \log n - 3k$, since once again this fails with probability $o(1)$.

Fix such a configuration \mathcal{P} . Then (by Lemma 6) we have discs D_1, \dots, D_m of radius ϵ enclosing all the $(3k \log \log n)$ -thinly covered regions. Recall that, outside $\bigcup D_i$, each point of S_n is covered at least $3k \log \log n$ times, and that no D_i is within distance $4r$ of any of the others or of ∂S_n . We shall examine the D_i one by one. For a fixed D_i , we will colour at most $0.1 \log \log n$ discs intersecting D_i with k colours so that all of D_i is covered by a disc of each of the first $k - 1$ colours, and the discs of colour k cover all the k -covered regions in D_i . When we have done this for each D_i , we will complete the colouring randomly and apply Lemma 9 in conjunction with the conclusion of Lemma 4 to give the required partition.

Our first task, then, is to k -colour some discs intersecting $D = D_1$. We replace the discs $D_r(x)$ of $C_r(S_n)$ intersecting D with half-planes $H(x)$ and add surrounding half-planes as in the proof of Theorem 5. These half-planes divide \mathbb{R}^2 into polygonal regions. We classify the polygonal regions in \mathbb{R}^2 into two types: the *red* regions, those covered less than $0.1 \log \log n$ times, and the *green* regions, which are covered at least this many times. As in the proof of Theorem 5, there are half-planes Π_1, \dots, Π_{k-1} , each of which covers the entire red region. If we remove these half-planes, there are two cases.

Case 1. \mathbb{R}^2 is still covered by the remaining half-planes.

Remove Π_1, \dots, Π_{k-1} . Since the remaining half-planes cover \mathbb{R}^2 , there is a subset consisting of (at most) three half-planes $H(x_1)$, $H(x_2)$ and $H(x_3)$, which cover \mathbb{R}^2 . Hence $D_r(x_1)$, $D_r(x_2)$ and $D_r(x_3)$ cover D , and we colour these discs with colour k .

Case 2. The remaining half-planes do not cover \mathbb{R}^2 .

In this case the uncovered region must have been covered at most $k - 1$ times originally, so must lie inside D . The uncovered region is convex, so is a polygonal region $P \subseteq D$. We convert the half-planes back to discs by decreasing their radii of curvature from ∞ to r continuously (while keeping

the intersection of their boundaries with D fixed), and stop at the moment when either D is covered or when the radius of curvature is r . In the first case, the last point y of D to be covered lies at the intersection of (almost surely) three discs which correspond to three of our original discs $D_r(x_1)$, $D_r(x_2)$ and $D_r(x_3)$, whose intersection contains y . Together, $D_r(x_1)$, $D_r(x_2)$ and $D_r(x_3)$ cover D since there are no bad lunes. We colour these discs with colour k . In the second case, colour with colour k all discs which bound the remaining uncovered region (which is exactly the region that was not k -covered originally in D). Once again, the absence of bad lunes implies that these discs cover all of D that was originally k -covered. Moreover, since the number of intersections that are not k -covered is at most $0.1 \log \log n - 3k$ (in the whole of S_n), there can be at most $0.1 \log \log n - 3k$ discs bordering regions that are not k -covered (each intersection that is not k -covered meets two of these discs, but each such disc has at least two such intersections on its boundary). Thus we have coloured at most $0.1 \log \log n - 3k$ discs with colour k .

The rest of the proof follows that of Theorem 5, except that we need $0.1 \log \log n$ coverage of the green regions to ensure that we still have $3k - 3$ coverage after removing the discs coloured with k above. Finally, as all the k -thinly covered regions lie inside $\bigcup_i D_i$, ∂S_n is k -covered, and so covered by \mathcal{A}_k .

Theorem 6. *For any fixed k ,*

$$\mathbb{P}(\{r : C_r(S_n) \text{ is a } k\text{-cover}\} = \{r : C_r(S_n) \text{ is } k\text{-partitionable}\}) \rightarrow 1$$

as $n \rightarrow \infty$.

Proof. Choose r_0 so that $\pi r_0^2 = \log n + k \log \log n - (\log \log n)^{1/2}$. Then **whp**, $C_{r_0}(S_n)$ fails to k -cover (Theorem 3), but does have a partition into $k - 1$ single covers $\mathcal{A}_1(r_0), \dots, \mathcal{A}_{k-1}(r_0)$, and a collection $\mathcal{A}_k(r_0)$ which covers all the k -covered regions, including ∂S_n (Lemma 10). Now suppose r is such that $C_r(S_n)$ is a k -cover. Then clearly $r > r_0$. Thus each $\mathcal{A}_1(r), \dots, \mathcal{A}_{k-1}(r)$ covers S_n , where $\mathcal{A}_i(r)$ is the collection of discs of $C_r(S_n)$ corresponding to the discs $\mathcal{A}_i(r_0)$ of $C_{r_0}(S_n)$. Suppose there is a point $x \in S_n$ which is not covered by $\mathcal{A}_k(r)$. Then no point y within distance $r - r_0$ of x is covered by $\mathcal{A}_k(r_0)$. Since $\mathcal{A}_k(r_0)$ covers ∂S_n , all points y within distance $r - r_0$ of x lie in S_n . But then no such point y is k -covered by $C_{r_0}(S_n)$. In particular, all such y are covered by exactly one disc from each $\mathcal{A}_i(r_0)$, $i = 1, \dots, k - 1$, and hence x can be covered by at most one disc from each $\mathcal{A}_i(r)$, $i = 1, \dots, k - 1$. In particular x is not k -covered, a contradiction. Thus $\mathcal{A}_k(r)$ covers S_n , and so $C_r(S_n)$ is k -partitionable.

Theorem 7. *Fix k, r and S_n . Place points x_1, x_2, \dots independently and uniformly at random in the region S' , where S' is some finite region of the plane that contains every point that is within distance r of S_n . Then with probability tending to 1 as $n/r^2 \rightarrow \infty$, the minimum value of N such that $\{D_r(x_i)\}_{i=1}^N$*

is a k -cover of S_n is equal to the minimum value of N such that $\{D_r(x_i)\}_{i=1}^N$ can be partitioned into k single covers of S_n .

Proof. If we choose $N = N_0$ according to a Poisson distribution of mean $\lambda|S'|$, then the points x_1, \dots, x_{N_0} form a Poisson process in S' of intensity λ . Choose λ so that

$$\pi r^2 \lambda = \log(\lambda n) + k \log \log(\lambda n) - (\log \log(\lambda n))^{1/2}.$$

We note that for large n/r^2 , such a λ does exist with $\lambda n \rightarrow \infty$ as $n/r^2 \rightarrow \infty$. By scaling the plane, we see that this is equivalent to taking a Poisson process with intensity 1 and considering coverage of $S_{\lambda n}$ by discs of radius $r\sqrt{\lambda}$. Thus by Theorem 3, **whp**, $\{D_r(x_i)\}_{i=1}^{N_0}$ fails to k -cover S_n , but does have a partition into $k-1$ single covers $\mathcal{A}_1, \dots, \mathcal{A}_{k-1}$, and a collection \mathcal{A}'_k which covers all the k -covered points in S_n (Lemma 10). Now add discs $D_r(x_{N_0+1}), D_r(x_{N_0+2}), \dots, D_r(x_N)$ to the collection \mathcal{A}'_k to obtain a collection \mathcal{A}_k . If $\{D_r(x_i)\}_{i=1}^N$ k -covers any point $x \in S_n$ then either $\{D_r(x_i)\}_{i=1}^{N_0}$ k -covers x , in which case x is covered by \mathcal{A}'_k , or $\{D_r(x_i)\}_{i=1}^{N_0}$ does not k -cover x , in which case x is covered by only one disc from each of $\mathcal{A}_1, \dots, \mathcal{A}_{k-1}$. But in this second case, x must be covered by \mathcal{A}_k since it is k -covered in total. Thus k -coverage implies \mathcal{A}_k covers S_n , so we have a k -partition $\mathcal{A}_1, \dots, \mathcal{A}_k$.

7 Sharpness

Note that it is the failure in Lemma 7 which gives rise to the $\Theta(1/\log n)$ bound in Theorem 1. Moreover, it can be easily seen that if a common $(k-2)$ -cover were sufficient, then the $\Theta(1/\log n)$ bound could be reduced to at least $(\log n)^{-2+o(1)}$. Thus failure of k -partitionability is most likely to occur with small k -covered configurations which have a common $(k-2)$ -cover. Our next result shows that Theorem 1 is essentially sharp by exhibiting such configurations that can occur with probability $\Theta(1/\log n)$.

Theorem 8. *Let $n \in \mathbb{R}$, $k \in \mathbb{N}$ and let $\pi r^2 = \log n + k \log \log n$. Then*

$$\mathbb{P}(E_r^k \setminus F_r^k) \geq \frac{c'_k}{\log n}$$

for sufficiently large n and some $c'_k > 0$ independent of n .

Proof. With n , k and r as in the statement of the theorem, we aim to show that a certain configuration occurs with probability at least $\frac{c'_k}{\log n}$. First we shall describe these configurations. Fix $\varepsilon = 1/r$. Let D_1, \dots, D_6 be six discs of radius $\varepsilon/10$, centred at the vertices of a regular hexagon with radius ε and centre O (see Figure 1, left). A simple calculation shows that there are half-planes which contain D_2, \dots, D_6 , say, but are at strictly positive distance




Fig. 1 Left: C_3 -configuration used in the proof of Theorem 8. Solid discs are precisely 2-covered, but there is no 2-partition. Right: a similar C_5 -configuration.

from D_1 . Thus for n sufficiently large, and hence ε/r sufficiently small, one can construct six discs $D_r(x_i)$, $i = 1, \dots, 6$, of radius r , such that $D_r(x_i)$ contains D_i but does not intersect any D_j , $j \neq i$, when $i = 2, 4, 6$, and $D_r(x_i)$ does not intersect D_i , but contains all the D_j , $j \neq i$, when $i = 1, 3, 5$. One can check (see Figure 2) that the $D_r(x_i)$ form a 2-cover of the disc of radius cr about O for some absolute $c > 0$. However, D_1 , D_3 and D_5 are precisely 2-covered by these discs. Suppose that $x_i \in \mathcal{P}$, but no other point of \mathcal{P} lies within distance $r + 2\varepsilon$ of O . Then the discs are not 2-partitionable, since D_i is covered only by $D_r(x_{i+2})$, and $D_r(x_{i+4})$ (indices mod 6), $i = 1, 3, 5$, and no 2-colouring of the discs $D_r(x_1)$, $D_r(x_3)$, $D_r(x_5)$ results in D_1 , D_3 and D_5 being covered by discs of both colours. It is easy to see that if n is sufficiently large then even removing points inside $D_{r+2\varepsilon}(O)$ from \mathcal{P} leaves $D_{3r}(O) \setminus D_{cr}(O)$ k -covered **whp**. Thus by applying Theorem 3 to the region $S_n \setminus D_{3r}(O)$ we obtain, with probability bounded away from zero, a k -covered but not k -partitionable configuration whenever we condition on the event that $\mathcal{P} \cap D_{r+2\varepsilon}(O)$ consists just of the points x_1, \dots, x_6 plus another $k - 2$ points inside $D_{r/2}(O)$. It remains to estimate the probability of such a configuration. Clearly, if we do not fix O , then x_1 is arbitrary, so we choose x_1 to be any point of $\mathcal{P} \cap S_n$ at distance at least $r + 2\varepsilon$ from ∂S_n (so that our configuration is guaranteed to lie inside S_n). Then x_2 is just confined to be within a certain distance of x_1 within a margin of $O(r)$, and hence within a region of area $O(r^2)$. After that, we may assume that the disks D_1, \dots, D_6 are fixed, so that the remaining x_i are confined to a region of area $O(\varepsilon r)$. Thus the probability of such a configuration occurring is bounded below by

$$p = Cnr^2(\varepsilon r)^4(r^2)^{k-2}e^{-\pi(r+2\varepsilon)^2}.$$

With $\varepsilon = 1/r$, we obtain $p \geq C'n(\log n)^{k-1}e^{-\pi r^2} = \frac{C'}{\log n}$, as claimed.




Fig. 2 Configuration used in the proof of Theorem 8 with closest point (cross) which is not 2-covered. Since the angle between the two boundaries $\partial D_r(x_i)$ is bounded away from zero where they cross between the small discs D_i , the 2-thinly covered region lies at distance at least cr from the D_i for some $c > 0$.

The configuration in Theorem 8 contains three precisely k -covered regions, with $k - 2$ discs covering all of them and three more discs covering just two of them. We call such a configuration a C_3 -configuration. One can also use a C_5 -configuration (Figure 1, right) which consists of five precisely k -covered regions, with $k - 2$ discs covering all of them, and five more discs each covering two of them (with a different pair of regions covered by each disc). One can follow the proof of Theorem 8 by defining D_1, \dots, D_5 to be discs of radius $\varepsilon/10$ about points on a regular pentagon of radius ε about O , $D_0 = D_{\varepsilon/10}(O)$ and $D_r(x_i)$ covering D_0 , D_i and D_{i+1} for $i = 1, \dots, 5$ (indices taken mod 5).

8 A classification theorem

As stated in the introduction, the aim of this section is to classify the non 2-partitionable 2-covers of \mathbb{R}^2 , where the covers are comprised of open half-planes. Throughout the section, we will assume that no two half-plane boundaries are equal, and no three half-plane boundaries pass through a single point. We do however allow boundaries of half-planes to be parallel. As the half-planes are open, it is clear that one can reduce to this case by, for example, translating some of the half-planes slightly so that they cover a smaller subset of \mathbb{R}^2 while at the same time preserving 2-coverage of \mathbb{R}^2 . We begin with some simple lemmas.

Lemma 11. *If \mathcal{C} is a finite collection of half-planes which cover \mathbb{R}^2 , then there is a subset of \mathcal{C} of size at most three that covers \mathbb{R}^2 .*

Proof. If not, the complements of any three (or fewer) half-planes from \mathcal{C} intersect. By Helly's theorem (see e.g. [3]), the complements of all the half-

planes from \mathcal{C} intersect, contradicting the fact that the original half-planes cover \mathbb{R}^2 .

The boundaries of the half-planes form lines which divide \mathbb{R}^2 into a number of polygonal (possibly infinite) atomic regions.

Definition 1. A *minimal region* is an atomic region that is covered by strictly fewer half-planes than any of its neighbouring regions.

Note that a minimal region may be covered by more half-planes than the least covered region.

Lemma 12. *Suppose \mathcal{C} is a finite collection of half-planes which cover \mathbb{R}^2 . Colour some of the half-planes red. If every minimal region is covered by a red half-plane, then the red half-planes cover \mathbb{R}^2 .*

Proof. Suppose that A is a non-minimal region that is not covered by a red half-plane. One of the regions A' adjacent to A is covered by a strict subset of the half-planes covering A , and so is not covered by a red half-plane either. Repeating, and using the fact that A is covered by a finite number of half-planes, we eventually arrive at a minimal region which is not covered by a red half-plane.

Lemma 13. *Suppose that the boundary $\partial\Pi$ of a half-plane $\Pi \in \mathcal{C}$ is adjacent to t minimal regions. Then each of these minimal regions is $(t - 1)$ -covered.*

Proof. Each of the t minimal regions must lie on the opposite side of $\partial\Pi$ to Π (since they are minimal). The intersections of (the closures) of these minimal regions with $\partial\Pi$ form t disjoint closed intervals I_1, \dots, I_t along $\partial\Pi$. Pick one of these intervals I_i . The endpoint(s) of I_i occur where the boundary lines of other half-planes cross $\partial\Pi$, I_i itself not being covered by these half-planes. Hence every point in $\partial\Pi \setminus I_i$ is covered by one of these half-planes. In particular, all the other intervals I_j , $j \neq i$, are covered by these half-planes. Since two distinct lines can only intersect at a single point, the half-planes corresponding to distinct intervals must be distinct. Thus each interval, and hence its corresponding minimal region, is covered by (at least) $(t - 1)$ half-planes.

Definition 2. A region is *precisely k -covered* if it is covered by exactly k half-planes.

Lemma 14. *Suppose that \mathcal{C} is a 2-cover of \mathbb{R}^2 . Suppose also that the half-planes of \mathcal{C} can be 2-coloured so that the precisely 2-covered regions are each covered by half-planes of each colour. Then there exists a set of (at most) three half-planes $\Pi_1, \Pi_2, \Pi_3 \in \mathcal{C}$ which covers \mathbb{R}^2 and such that each precisely 2-covered region is contained in only one of the Π_i .*

Proof. Let the colours be red and blue, and let the red half-planes be R_1, \dots, R_s and the blue half-planes be B_1, \dots, B_t . Also, write r_i and b_j for the number of precisely 2-covered regions covered by the half-planes R_i and B_j respectively. Suppose first that there are at most two red half-planes, R_1 and R_2 say, with $r_i > 0$. Since removing these half-planes reduces the coverage of any region by at most two, and all the precisely 2-covered regions are covered by blue half-planes, the remaining (red and blue) half-planes must cover \mathbb{R}^2 . Applying Lemma 11 to the remaining half-planes gives the desired three (or fewer) half-planes Π_1, Π_2 and Π_3 , since all the precisely 2-covered regions must be covered by either R_1 or R_2 , which are distinct from the Π_i , and hence can only be covered by at most one Π_i . A similar argument applies if there are at most two blue half-planes with $b_j > 0$.

Consequently we may assume that there are at least three half-planes of each colour, so that $s, t \geq 3$, and indeed that at least three of each of the r_i and the b_j are positive. Note that no precisely 2-covered region can be covered by more than one of the R_i , or by more than one of the B_j , by hypothesis, so that

$$\sum_{i=1}^s r_i = \sum_{j=1}^t b_j = N,$$

where N is the number of precisely 2-covered regions. Suppose that at most two of the r_i are greater than one. Without loss of generality $r_i \leq 1$ unless $i = 1$ or 2 . Now remove R_1 and R_2 . The remaining half-planes still cover \mathbb{R}^2 , so by Lemma 11 we may select three of them which cover \mathbb{R}^2 . Suppose that these half-planes are R_{i_1}, \dots, R_{i_h} and $B_{j_1}, \dots, B_{j_\ell}$ where $h + \ell = 3$. Then, since $r_{i_r} \leq 1$ and there are at least 3 positive b_j ,

$$N \leq \sum_{r=1}^h r_{i_r} + \sum_{b=1}^{\ell} b_{j_b} \leq h + \sum_{b=1}^{\ell} b_{j_b} = (3 - \ell) + \sum_{b=1}^{\ell} b_{j_b} \leq \sum_{j=1}^t b_j = N,$$

so that in fact each precisely 2-covered region is covered exactly once by our three half-planes, as required. Hence we may assume that at least three of the r_i are greater than one, so that in particular there are at least six precisely 2-covered regions.

We now claim that there is at most one minimal region M which is covered more than twice by \mathcal{C} . To see this, fix a precisely 2-covered region A , and let P_1 and P_2 be the half-planes covering A . For each other minimal region B , A must lie in at least one of the half-planes forming the boundary of B (since B is minimal and so $\mathbb{R}^2 \setminus \bar{B}$ is covered by these half-planes). Hence all the other minimal regions lie on the boundary of either P_1 or P_2 . By Lemma 13, at most three of the precisely 2-covered minimal regions can lie on the boundary of each P_i . Since there are at least six precisely 2-covered regions, both P_1 and P_2 are adjacent to at least one precisely 2-covered region. But then by Lemma 13 there can be at most three minimal regions lying on each of the boundaries of P_1 and P_2 . Thus there are at most seven minimal regions in

total (including A and at most three on each of P_1 and P_2). Of these at most $7 - 6 = 1$ can be covered more than twice. Since this region (if it exists) is either covered by a blue half-plane or a red half-plane, either the blue half-planes or the red half-planes cover all the minimal regions, and hence, by Lemma 12, all of \mathbb{R}^2 . Suppose that the red half-planes cover \mathbb{R}^2 , and apply Lemma 11. We find three (or fewer) red half-planes covering \mathbb{R}^2 which satisfy our requirements.

Lemma 15. *Suppose \mathcal{C} is a finite 2-cover of \mathbb{R}^2 that is not 2-partitionable. Then the half-planes of \mathcal{C} cannot be 2-coloured so that every precisely 2-covered region is covered by half-planes of each colour.*

Proof. Suppose the lemma is false, so that the half-planes can be coloured in such a fashion, and apply Lemma 14. Then if we remove the half-planes Π_1, Π_2 and Π_3 , which form a cover of \mathbb{R}^2 , the remaining half-planes from \mathcal{C} do not cover \mathbb{R}^2 by hypothesis. The only way this can happen is that there is a precisely 3-covered minimal region M lying inside the triangle Δ that is 3-covered by the Π_i , since the precisely 2-covered regions are only covered once by the Π_i , and the 3 (or more) covered regions outside Δ are only covered by at most two of the Π_i . (If we just have two half-planes Π_i covering \mathbb{R}^2 , then all 3-covered regions remain covered when we remove them.) We may label the sides of M by the angles which the half-planes forming them make with the horizontal, with the convention that the angle of the half-plane $\{(x, y) : y > 0\}$ is zero. Without loss of generality, suppose that the angles of Π_1, Π_2 and Π_3 are $0, \alpha$ and β respectively, where $0 \leq \alpha \leq \beta$. Since the Π_i cover \mathbb{R}^2 , $0 \leq \beta - \pi \leq \alpha \leq \pi$.




Fig. 3 Half-planes Π_1, Π_2, Π_3 and Π'_2 (of type 3) in proof of Lemma 15. Minimal region shown (hexagon) is surrounded by half-planes of types 1–6.

We now introduce a labelling of the sides of M based on their angles, as given by the following table (see Figure 3). The first column lists the angle of the half-plane forming a side of M , column $1 + i$ indicates whether this half-plane completely covers (1), partially covers (p) or doesn't intersect (0) the region $S_i = \Pi_i \setminus (\Pi_{i+1} \cup \Pi_{i+2})$, where the subscripts are taken mod 3, and the final column is the new labelling.

angle	S_1	S_2	S_3	label
$0 \leq \theta < \beta - \pi$	1	p	0	1
$\beta - \pi \leq \theta < \alpha$	p	1	0	2
$\alpha \leq \theta < \pi$	0	1	p	3
$\pi \leq \theta < \beta$	0	p	1	4
$\beta \leq \theta < \alpha + \pi$	p	0	1	5
$\alpha + \pi \leq \theta < 2\pi$	1	0	p	6

We claim that, as we follow the sides of M in a counter-clockwise direction (so that θ increases), the following pairs of consecutive labels are prohibited.

$$11, 22, 33, 44, 55, 66, 12, 13, 24, 34, 35, 46, 56, 51, 62.$$

Let us check this first for the cases 33, 34 and 35. Suppose that there is a half-plane Π'_2 corresponding to a side of M labelled 3 which is followed by a side labelled 3, 4 or 5, corresponding to a half-plane Π_4 . Then we may remove the cover formed by Π_1, Π'_2 and Π_3 , and we are left with a cover of \mathbb{R}^2 . The reason for this is that Π_4 and Π_2 ensure that the region $\Pi_1 \cap \Pi'_2 \cap \Pi_3$ is 4-covered and that the region $(\Pi'_2 \cap \Pi_3) \setminus \Pi_1$ is 3-covered. The region $(\Pi_1 \cap \Pi_3) \setminus \Pi'_2 \subseteq \Pi_1 \cap \Pi_3$ is 3-covered by assumption (since any precisely 2-covered region cannot be covered by both Π_1 and Π_3), as is $(\Pi'_2 \cap \Pi_1) \setminus \Pi_3 \subseteq \Pi_2 \cap \Pi_1$. Thus removing Π_1, Π'_2 , and Π_3 does not result in any uncovered region. Since \mathcal{C} is not 2-partitionable, these cases cannot occur. By symmetry, neither can 11, 12, 13, and 55, 56, 51. By reflection in the vertical axis, reversing the order in which we consider the sides of M , the pairs 24, 34, and 44 are also excluded. Hence by symmetry so are 46, 56, 66, and 62, 12, 22.

Now suppose there is a side labelled 1. The next side must be labelled 4, 5 or 6. The following side must then be labelled 5 or 1 and in both cases there are no possibilities for the fourth side (which might be identical with the first). Similarly there can be no side labelled 3 or 5. But if there is then a side labelled 2, the next side must be labelled 6, and there is no possibility for the following side. Similarly no side can be labelled 4 or 6. Hence no such region M can exist.

If \mathcal{C} is a collection of half-planes, we define a C_3 -configuration to be one which contains three precisely 2-covered regions A_1, A_2 and A_3 and three half-planes P_1, P_2 and P_3 , where A_i is covered by P_j for $j \neq i$. A C_5 -configuration is one which contains five regions A_i , $1 \leq i \leq 5$, and five half-planes P_i ,

$1 \leq i \leq 5$, where A_i is covered by P_i and P_{i+1} , with the subscripts taken mod 5.

Theorem 9. *Suppose \mathcal{C} is a finite 2-cover of \mathbb{R}^2 that is not 2-partitionable. Then \mathcal{C} contains a C_3 or a C_5 -configuration.*

Proof. We define a graph $G = (V, E)$ where the vertices are the half-planes of \mathcal{C} and the edges are the precisely 2-covered regions covered by the half-planes corresponding to their endvertices. This graph is not bipartite, by Lemma 15. Hence it must contain an odd cycle. Suppose first that the vertices of this cycle correspond to half-planes P_1, \dots, P_s , where s is odd and at least seven, and P_i covers the 2-covered regions A_i and A_{i+1} (the subscripts taken mod s). We observe that any three half-planes from \mathcal{C} cannot cover \mathbb{R}^2 since they cannot even cover the A_i . (The three half-planes include at most three of the P_i which between them can cover at most six of the A_i .) Consequently, by Lemma 11, \mathcal{C} cannot cover \mathbb{R}^2 , a contradiction. Thus G contains a cycle on 3 or 5 vertices, corresponding to a C_3 or C_5 -configuration respectively.

9 Deterministic partitionability of k -covers

Theorem 10. *Any finite k -cover of \mathbb{R}^2 by half-planes is $\lceil k/2 \rceil$ -partitionable. Conversely, there exists a k -cover of \mathbb{R}^2 by half-planes which is not $\lceil (2k+2)/3 \rceil$ -partitionable.*

Proof. A 1-cover or 2-cover is clearly 1-partitionable, so for the first part we may assume $k \geq 3$. We know by Lemma 11 that we can choose three or fewer half-planes that cover \mathbb{R}^2 . Suppose first that no two half-planes cover \mathbb{R}^2 . Choose three half-planes Π_1, Π_2, Π_3 covering \mathbb{R}^2 with minimal distance between $\partial\Pi_1$ and the point of intersection of $\partial\Pi_2$ and $\partial\Pi_3$. (As no two half-planes cover \mathbb{R}^2 , $\partial\Pi_2$ and $\partial\Pi_3$ cannot be parallel.) Let Δ be the triangle $\Pi_1 \cap \Pi_2 \cap \Pi_3$. Suppose there is a point in Δ that is precisely k -covered. Then it must lie in some minimal region M . M consists of a polygonal region whose boundary lies in the interior of Δ (since no half-plane covering M can have a boundary in common with M). Let p be a corner of M that is closest to $\partial\Pi_1$. Then the two half-planes Π'_2 and Π'_3 giving rise to the corner p , together with Π_1 cover \mathbb{R}^2 and p is closer to $\partial\Pi_1$ than $\partial\Pi_2 \cap \partial\Pi_3$, contradicting the choice of Π_1, Π_2, Π_3 . Thus Δ is $(k+1)$ -covered, and removing Π_1, Π_2, Π_3 from the k -cover gives rise to a $(k-2)$ -cover. If two half-planes Π_1, Π_2 cover \mathbb{R}^2 , then removing these also gives rise to a $(k-2)$ -cover. In either case, by induction on k , the remaining half-planes are $\lceil (k-2)/2 \rceil$ -partitionable. Adding the cover $\{\Pi_1, \Pi_2, \Pi_3\}$ or $\{\Pi_1, \Pi_2\}$ gives a $\lceil k/2 \rceil$ -partition as required.

Now consider $2k+1$ half planes all containing the origin, and at angles that form all the multiples $2\pi i/(2k+1)$, $i = 0, \dots, 2k$. Every point in \mathbb{R}^2 is

covered by at least k of these half-planes, so these half-planes form a k -cover. However, at least three of these half-planes are needed to form any single cover. Since $3\lceil(2k+2)/3\rceil \geq 2k+2 > 2k+1$, this collection of half-planes is not $\lceil(2k+2)/3\rceil$ -partitionable.

Note that the example of failure of $\lceil(2k+2)/3\rceil$ -partitionability can be extended to covers by discs. Indeed, take an example of this construction and choose $K > 0$ so that the half-planes fail to be partitionable even in $D_K(O)$. By replacing the half-planes by sufficiently large discs and placing many discs far from the origin so that every point outside $D_K(O)$ is covered, we obtain an example with discs.

References

1. R. Arratia, L. Goldstein and L. Gordon, *Two Moments Suffice for Poisson Approximations: The Chen-Stein Method*, Ann. Probab. **17** (1989), 9–25.
2. B. Bollobás, *Random Graphs*, second edition, Cambridge University Press, 2001.
3. B. Bollobás, *The Art of Mathematics: Coffee Time in Memphis*, Cambridge University Press, 2006.
4. E.N. Gilbert, *Random plane networks*, Journal of the Society for Industrial and Applied Mathematics **9** (1961), 533–543.
5. E.N. Gilbert, *The probability of covering a sphere with N circular caps*, Biometrika **56** (1965), 323–330.
6. P. Hall, *On the coverage of k -dimensional space by k -dimensional spheres*, Annals of Probability **13** (1985), 991–1002.
7. S. Janson, *Random coverings in several dimensions*, Acta Mathematica **156** (1986), 83–118.
8. S. Kumar, T.H. Lai and J. Balogh, *On k -coverage in a mostly sleeping sensor network*, Proceedings of the Tenth International Conference on Mobile Computing Networks (MobiCom 2004).
9. P. Wan and C. Yi, *Coverage by randomly deployed wireless sensor networks*, IEEE Transactions on Information Theory **52** (2006), 2658–2669.
10. H. Zhang and J. Hou, *On deriving the upper bound of α -lifetime for large sensor networks*, Proceedings of the Fifth ACM International Symposium on Mobile Ad Hoc Network Computing (MobiHoc 2004).