Annales de l'Institut Henri Poincaré - Probabilités et Statistiques
2012, Vol. 48, No. 2, 327–342
DOI: 10.1214/10-AIHP399

ANNALES DE L'INSTITUT HENRI POINCARÉ
PROBABILITÉS ET STATISTIQUES
www.imstat.org/aihp

Stationary map coloring

Omer Angel[a,1], Itai Benjamini[b], Ori Gurel-Gurevich[a], Tom Meyerovitch[a] and Ron Peled[c,2]

[a] *Department of Mathematics, University of British Columbia, Vancouver BC, V6T 1Z2, Canada*
[b] *Faculty of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot 76100, Israel*
[c] *School of Mathematical Sciences, Tel Aviv University, Tel Aviv, 69978, Israel*

Received 23 October 2010; revised 24 October 2010; accepted 1 November 2010

Abstract. We consider a planar Poisson process and its associated Voronoi map. We show that there is a proper coloring with 6 colors of the map which is a deterministic isometry-equivariant function of the Poisson process. As part of the proof we show that the 6-core of the corresponding Delaunay triangulation is empty.

Generalizations, extensions and some open questions are discussed.

Résumé. Nous étudions un processus de Poisson planaire et sa partition de Voronoi associée. Nous montrons qu'il existe une coloration à six couleurs de cette partition satisfaisant les deux propriétés suivantes : la coloration est une fonction déterministe du processus de Poisson. Par ailleurs, elle commute avec les isométries du plan. Une partie de la preuve consiste à montrer que le "6-core" de la triangulation de Delaunay associée est vide.

Des généralisations, extensions et problèmes ouverts sont discutés.

MSC: 60C05; 60D05

Keywords: Poisson process; Graph coloring; Planar graphs; Voronoi tessellation; Delaunay triangulation; Percolation

1. Introduction

The Poisson–Voronoi map is a natural random planar map. Being planar, a specific instance can always be colored with 4 colors with adjacent cells having distinct colors (see Figure 1). The question we consider here is whether such a coloring can be realized in a way that would be isometry-equivariant, that is, that if we apply an isometry to the underlying Poisson process, the colored Poisson–Voronoi map is affected in the same way. In other words, can a Poisson process be equivariantly extended to a colored Poisson–Voronoi map process? How many colors are needed? Can such an extension be deterministic?

Extension of spatial processes, particularly of the Poisson process, have enjoyed a surge of interest in recent years. The general problem is to construct in the probability space of the given process, a richer process that (generally) contains the original process. Notable examples include allocating equal areas to the points of the Poisson process [6,7,11,12,16,18]; matching points in pairs or other groups [1,9,14,15,24]; thinning and splitting of a Poisson process [2,3,13]. Coloring extensions of i.i.d. processes on $\mathbb{Z}^d$ are considered in [4].

[1] Research supported by NSERC.
[2] Partially completed during stay at the Institut Henri Poincare – Centre Emile Borel. Research supported by NSF Grant OISE 0730136.

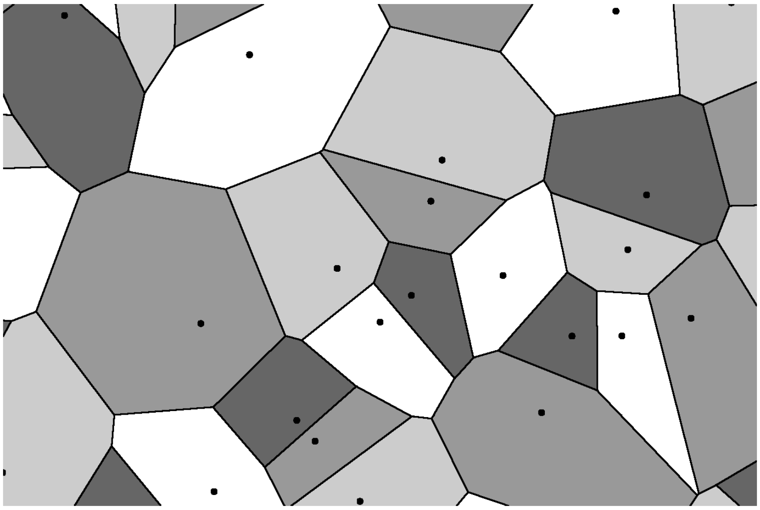

Fig. 1. A proper 4-coloring of a portion of the Poisson–Voronoi map.

We now proceed with formal definitions and statement of the main results. A non-empty, locally finite subset $S \subset \mathbb{R}^d$ defines a partition of $\mathbb{R}^d$, called the *Voronoi tessellation*, as follows: The Voronoi cell $C(x)$ of a point $x \in S$ contains the points of $\mathbb{R}^d$ whose distance to S is realized at x:

$$C(x) = \{z \in \mathbb{R}^d\colon\ d(z, x) = d(z, S)\}.$$

Points in the intersection $C(x) \cap C(y)$ have equal distance to x and y. It follows that the cells cover $\mathbb{R}^d$ and have disjoint interiors.

For the purposes of coloring, we consider the adjacency graph G of these cells, with vertices S and edge (x, y) if $C(x) \cap C(y) \neq \varnothing$. In the case $d = 2$, this graph is called the *Delaunay triangulation*, and is a triangulation of the plane for generic point sets. (In general, this graph is the 1-skeleton of a simplicial cover of $\mathbb{R}^d$.) A k-coloring of the Voronoi tessellation is a proper k-coloring of the Delaunay triangulation, i.e. an assignment of one of k colors to each cell so that adjacent cells have distinct colors. Note that if S does not contain four or more co-cyclic points, then no more than three cells meet at a single point. This is a.s. the case for the Poisson process. However, for greater generality one needs the more careful definition, where (x, y) is an edge if $|C(x) \cap C(y)| > 1$. This ensures that the graph is planar.

Given a standard (unit intensity) Poisson process $\mathcal{P} \subset \mathbb{R}^2$, the Poisson–Voronoi map is the Voronoi map of its support. By the 4 color theorem, the Poisson–Voronoi map can always be properly colored with 4 colors. Our main question is whether it is possible to color the Poisson–Voronoi map in an isometry equivariant way and if so, how many colors are needed.

To make this precise, let $\mathbb{M}$ be the space of locally finite sets in $\mathbb{R}^d$, endowed with the local topology and Borel σ-algebra.[3] Let $\mathbb{P}$ be the probability on $\mathbb{M}$ which is the law of the Poisson process. Each realization $\mathcal{P} \in \mathbb{M}$ has the Delaunay graph associated with it. A (proper) k-coloring of $\mathcal{P}$ is a disjoint partition $\mathcal{P} = \bigcup_{i=1}^k \mathcal{P}_i$ such that if $x \sim y$ in the Delaunay graph of $\mathcal{P}$, then x, y are not in the same $\mathcal{P}_i$. Thus the space of k-colored maps is a subset of $\mathbb{M}^k$.

As in other works on extensions of spatial processes, we consider both deterministic and randomized versions of our main object of study, as follows:

A *deterministic* k-coloring scheme of the Voronoi map is a measurable function $F : \mathbb{M} \to \mathbb{M}^k$ such that $F(\mathcal{P})$ is $\mathbb{P}$-a.s. a k-coloring of $\mathcal{P}$. Informally, given the point process, F assigns a color to each point so that the result is a proper coloring.

A *randomized* k-coloring scheme of the Voronoi map is a probability measure μ on $\mathbb{M}^k$, supported on proper k-colorings, such that the law under μ of $\mathcal{P} := \bigcup_{i=1}^k \mathcal{P}_i$ is $\mathbb{P}$. Given such a measure μ, one may consider μ conditioned

[3] It is also common to let $\mathbb{M}$ be the set of non-negative integer valued measures on $\mathbb{R}^d$ with $\mu(\{x\}) \in \{0, 1\}$. The distinction will not be important to us.

on $\mathcal{P}$. This conditional distribution is defined $\mathbb{P}$-a.s., and is supported on k-colorings of $\mathcal{P}$. Thus a randomized k-coloring can be interpreted as assigning to each $\mathcal{P} \in \mathbb{M}$ a probability measure on k colorings of $\mathcal{P}$. Note that any deterministic coloring scheme is also a randomized one, with μ being the push-forward of $\mathbb{P}$ by F.

A deterministic coloring scheme is said to be *isometry equivariant* if every isometry γ of $\mathbb{R}^d$, acting naturally on $\mathbb{M}$ and $\mathbb{M}^k$, has $\gamma F(\mathcal{P}) = F(\gamma\mathcal{P})$. For randomized schemes equivariance is defined by $\mu \circ \gamma = \mu$. These definitions coincide for deterministic schemes.

Theorem 1.1. *There exists a deterministic isometry equivariant* 6*-coloring scheme of the Poisson–Voronoi diagram in* $\mathbb{R}^2$.

The requirement of determinism complicates things significantly. In contrast, we have the much simpler result

Proposition 1.2. *There exists a randomized isometry equivariant* 4*-coloring scheme of the Poisson–Voronoi diagram in* $\mathbb{R}^2$.

In dimensions other than 2 the problem is not as interesting.

Proposition 1.3. *In* $\mathbb{R}$, *there is a randomized isometry equivariant coloring of the Poisson–Voronoi map with* 2 *colors and a deterministic one with* 3 *colors. In both cases this is the best possible.*

In $\mathbb{R}^d$ *for* $d > 2$, *the chromatic number of the Poisson–Voronoi map is a.s.* ∞.

The rest of the paper is organized as follows: In Section 2 we outline the proof of Theorem 1.1, and present our deterministic coloring algorithm and the two main propositions needed to prove its correctness. In Section 3 we discuss related questions: randomized colorings, dimensions other than 2, and mention some open problems. Section 4 contains the proof of our main theorem.

2. Proof outline

We outline the proof of Theorem 1.1. The idea is to find an isometry equivariant adaptation to the Voronoi map of a 6 coloring algorithm for finite planar graphs, originating in Kempe's attempted proof of the four color theorem. By Euler's formula it is known that any finite planar graph G has a vertex of degree at most 5. The algorithm proceeds by iteratively removing such a vertex until the graph is empty, then putting back the vertices one by one in reverse order. As each vertex is put back into the graph, it is assigned a color distinct from those already assigned to any of its neighbors. Since a vertex has at most 5 neighbors when it is put back, this produces a proper 6 coloring.

To adapt this algorithm to the Poisson–Voronoi isometry equivariant setting, one must deal with several issues. First, there exist infinitely many vertices of degree at most 5 and there is no way to pick just one of them in an isometry-equivariant way. Second, even if we iteratively remove *all* vertices of degree at most 5, the graph will not become empty after any finite number of steps. Finally, when returning the vertices, it is not clear in what order to do so (which may be important if some of them are neighbors). We need a way to order them which is isometry-equivariant.

We overcome these issues by proving that for a Poisson–Voronoi map, the following two properties hold almost surely. Let $G = (V, E)$ be the Delaunay graph formed by the Poisson–Voronoi map. For a cell $v \in V$ write $A(v)$ for its area as a planar region. Inductively, define $G_0 = G$ and G_{i+1} as the graph formed from G_i by removing all vertices of degree (in G_i) at most 5.

Proposition 2.1. *There exists an integer* $M > 0$ *such that, almost surely,* G_M *contains only finite connected components.*

Proposition 2.2. *Almost surely, all cells have different areas and there is no infinite path in* G *with decreasing areas.*

We now exhibit a deterministic algorithm which takes as input a graph $G = (V, E)$ with chromatic number at most 6 and an area function $A : V \to \mathbb{R}^+$ satisfying the two propositions above and returns a proper 6-coloring of the graph. Since the algorithm only depends on the graph structure G and areas A which are preserved by isometries, it is clear that when applying it to the Delaunay graph of a Poisson–Voronoi map we will get a deterministic isometry-equivariant 6-coloring.

The algorithm starts with all vertices uncolored. Once a vertex is colored, its color never changes. Consider first G_M. Each of its components is finite and hence may be colored with 4 colors in an isometry equivariant way (e.g. take the minimal coloring in lexicographic order, when the vertices of the component are ordered by their area).

Next, having colored G_k, we color G_{k-1} inductively. Once $G = G_0$ is colored, we are done. Consider the vertices of $G_{k-1} \setminus G_k$. Each has at most 5 neighbors in G_{k-1}. We order these vertices by increasing areas and wish to color them in order, i.e., coloring a vertex v only after its neighbors of smaller area have been colored. The color of these neighbors is determined using the same method in an iterative manner. Proposition 2.2 implies that there are just finitely many vertices that need to be considered before v (see also Lemma 4.19). Hence, going over these finitely many vertices in the order of their areas, we color each one by a color which is unused by its neighbors (say, the minimal such color) until we finally color v.

Proposition 2.1 is more difficult than Proposition 2.2 and the main lemma required for its proof (Lemma 4.9) says that if we consider a square of side length $6R$ and iteratively remove vertices *inside* this square having degree at most 5, then the square of side length $2R$ with the same center will eventually become empty with probability tending to 1 as $R \to \infty$. This is shown using several probabilistic estimates and uses of Euler's formula. We then show that for well separated squares of side length $6R$, the events just described, applied to these squares, are nearly independent. A small variation on the above event (requiring that the boxes are also *sealed*; see below) makes separated boxes completely independent. Proposition 2.1 then follows by standard k-dependent percolation arguments. Proposition 2.2 is proved using a similar but easier k-dependent percolation argument.

As a corollary of the proofs of the above propositions we obtain that our coloring is finitary with exponential tails. That is, for any given point $p \in \mathbb{R}^2$, the probability that the color of the cell containing p is not determined by the points of the Poisson process within a ball of radius R around p is at most $C\mathrm{e}^{-cR}$ for some $C, c > 0$. See remark after the proof of the main theorem for more details.

Note that instead of the area A, we could use any other parameter of the cell (e.g. diameter) which satisfies Proposition 2.2 (in fact, one can relax the requirement that all cells have different areas to the requirement that adjacent cells have different areas). The sole purpose of A is to induce a well founded order on cells which would "break ties" when putting back vertices. We chose to use the area because it is a very natural parameter to consider, but it is as easy to prove the required properties for other parameters (see Section 4.2). A related result is that there is no infinite path where each Poisson point is the closest to the previous one in the path [8].

3. Generalizations, extensions and questions

In this section we explain some variants and extensions of the question and settings discussed in our paper.

3.1. *Randomized colorings*

The fact that there is a randomized 4-coloring scheme of the Poisson–Voronoi map follows from the four color theorem by a soft argument. This involves an averaging consideration of ergodic theory and works for any amenable transitive space.

Proof of Proposition 1.2. We first claim that the 4-color theorem implies existence of a measurable function F (not necessarily equivariant) which assigns each Voronoi diagram a 4-coloring. One such map is the lexicographically minimal proper coloring (when the centers are ordered by their distance to the origin). It is measurable since it is the limit of the lexicographically minimal coloring of the n nearest centers.

To get a randomized equivariant coloring, let τ_x be a translation by $x \in \mathbb{R}^2$, ρ_θ a rotation by θ, and ε the reflection about the x axis. Let $\sigma = \tau_x \circ \rho_\theta \circ \varepsilon^u$ be a random isometry, where $u \in \{0, 1\}$, $\theta \in [0, 2\pi]$ and $x \in B(0, R)$ are uniform and independent.

This defines a probability measure F_R on 4-colored maps by conjugating F by σ. It is clear (due to compactness of the space of distributions over 4-colorings) that $\{F_R\}$ has a subsequential weak limit as $R \to \infty$, and any such limit is an isometry equivariant 4-coloring. $\square$

Explicit randomized colorings
While the previous argument is clearly optimal with respect to the number of colors used, it is not constructive. It is instructive to consider an explicit construction with 7 colors. The construction below will be algorithmic, i.e. there is an algorithm, that determines the color of each cell by accessing a finite (but unbounded) number of cells along with a random independent bit for each cell.

As a first stage, we explain how to get an 8-coloring. Start by assigning a fair coin toss to each cell independently. Consider the subgraph of $H \subset G$ where an edge is present if its endpoints have the same coin result. The connected components in this graph are components of site percolation on G with $p = 1/2$. By a result of Zvavitch [25], almost surely all connected components of both the heads and tails will be finite (in fact, Bollobás and Riordan [5] proved that the critical percolation threshold is indeed $p = \frac{1}{2}$).

Color each "head" component independently with colors $\{0, 1, 2, 3\}$ in some deterministic isometry-equivariant manner which is a function only of the cells of this component (e.g., again, a lexicographically minimal coloring with vertex order based on cell areas). Color the "tail" components with $\{4, 5, 6, 7\}$. The result is a.s. a proper 8-coloring of G. The randomness comes exclusively from the coin tosses. The color of a cell is determined by its connected component in H (and the size of the corresponding cells).

A trick suggested by Gady Kozma [17] reduces the number of colors required to 7 as follows. A finite planar graph embedded in the plane has a unique unbounded face, called the *external* face. Attaching an additional vertex to the vertices of the external face preserves planarity. Thus a finite planar graph can be 4-colored so vertices of the external face do not use one specified color. Now color the "heads" components using $\{0, 1, 2, 3\}$ so that color 0 does not appear at vertices of the external face of any component. Color the "tails" components using $\{0, 4, 5, 6\}$ with the same constraint. Whenever two connected planar graphs are jointly embedded in the plane, one is contained in the external face of the other. Thus when a "tails" component is adjacent to a "heads" component, it is impossible for them to have adjacent vertices colored 0, and the coloring is proper.

As noted above, in order to determine the color of any cell, it is sufficient to know the map structure and the coin-tosses within a ball of a certain random radius around this cell. In addition, if one modifies the above algorithm by initially performing fair-independent rolls of a 3-sided dice, instead of coin tosses (thus obtaining a proper 10-coloring in the final outcome, after applying Kozma's trick) then the distribution of the aforementioned radius will have exponential tails (see [5]). The radius for our deterministic 6-coloring also has exponential tails, as noted in the proof outline.

3.2. 1-*dimensional Poisson–Voronoi map*

The deterministic isometry equivariant chromatic number of a graph may well be different from its usual chromatic number. For example, consider Z^d translated by a uniform random variable in $[0, 1]^d$ and rotated by a uniform random angle in $[0, 2\pi]$. Clearly, the distribution of this random graph is isometry invariant and it is almost surely 2-colorable. Yet any deterministic isometry equivariant coloring must assign the same color to all vertices and hence cannot be proper.

A different example is furnished by the 1-dimensional Poisson–Voronoi diagram, i.e., the "Voronoi" map composed of line segments around the points of a one-dimensional standard Poisson process. This map is 2-colorable, but we claim that its deterministic isometry equivariant coloring number is 3. First, it is seen to be at most 3 by considering the following algorithm: First color green all cells which are shorter than both their neighbors. Now, from each green cell, proceed to alternately color its neighbors to the right by red and blue, until the next green cell is reached. This produces a deterministic translation equivariant proper 3-coloring. To get an isometry equivariant coloring, instead of coloring red and blue from left to right, start from the shorter of the two green cells bounding the current stretch of uncolored cells.

The following lemma states that at least 3 colors are needed. The argument is standard, see e.g. [14] for a similar context.

Lemma 3.1. *There is no deterministic translation equivariant proper* 2*-coloring of the* 1 *dimensional Poisson–Voronoi map*.

Proof. In order to reach a contradiction, suppose $\mathcal{A}$ is such a coloring scheme. Since $\mathcal{A}$ is measurable there exists an integer L and a function $\mathcal{B}: \mathbb{M} \to \{0, 1\}$ such that $\mathcal{B}$ depends only on the point process in the interval $[-L, L]$ and the probability that $\mathcal{B}$ is the color given by $\mathcal{A}$ to the cell of the origin is at least $\frac{7}{8}$. Consider also another point $x > 2L$. By translation equivariance, with probability at least $\frac{7}{8}$ the $\mathcal{A}$-color of the cell of x is given by $\mathcal{B}$ applied to the map shifted by x, which is determined by the centers in $[x - L, x + L]$.

Hence, with probability at least $\frac{3}{4}$ the $\mathcal{A}$-color of both these cells agrees with a function of the centers in $[-L, L] \cup [x - L, x + L]$. However, The $\mathcal{A}$-colors of these cells determine the parity of the number of cells (i.e. centers) between them. But for the Poisson process, the parity of the number of centers in $[L, x - L]$ is independent of the process in $[-L, L] \cup [x - L, x + L]$, and tends to uniform on $\{0, 1\}$ as $x \to \infty$. When x is large enough this gives a contradiction. □

We remark that a variant of the 3-coloring above can be used to color any invariant point process on $\mathbb{R}$ that is not an arithmetic progression (so that not all points are isomorphic). Furthermore, the proof of impossibility with 2 colors also applies to more general processes as we only use the fact that the parity of the number of points in $[L, x - L]$ is not (nearly) determined by the process in $[-L, L]$ and $[x - L, x + L]$ for x large enough.

3.3. *Higher dimensional Poisson–Voronoi maps*

A natural generalization of our setting is to consider the 3-dimensional Poisson–Voronoi diagram. In this case it is not obvious whether one can properly color the diagram with finitely many colors even without the isometry equivariant condition. Dewdney and Vranch [10], and Preparata [22] discovered that n Voronoi cells in $\mathbb{R}^3$ may be all pairwise adjacent. Indeed, [10] shows that in $\mathbb{R}^3$, the Voronoi cells of $(x_i, x_i^2, x_i^3)_{i=1}^n$ satisfy this for any $\{x_1, \ldots, x_n\}$. Since pairwise adjacency is preserved by sufficiently small perturbations, and since such configurations a.s. appear in the Poisson process, this implies that the chromatic number of the 3-dimensional Poisson–Voronoi diagram is almost surely infinite. Higher dimensional analogues also exist.

Following Proposition 2.1, one can still ask, as a weaker result than having an isometry equivariant coloring, what is the minimal k such that if we iteratively remove all cells having degree at most k we remain with finite components only? Such a k necessarily exists by arguments similar to those of Proposition 2.1. (Simulations indicate that $k = 12$ may suffice in $\mathbb{R}^3$.)

3.4. *Ramblings and open questions*

Fewer colors

Is there a deterministic 4-coloring of the Poisson–Voronoi map? Theorem 1.1 shows that 6 colors suffice, while obviously at least 4 are needed. Recent work by Ádám Timár [23] shows the existence of deterministic, equivariant 5-colorings using different methods. In order to get a 5-coloring using our methods it suffices to prove an analogue of Proposition 2.1 – that after finitely many stages of repeatedly removing vertices of degree at most 4 from the Delaunay graph the remaining graph does not contain an infinite connected component (this implies that the 5-core, that is, what remains after infinitely many such stages, has only finite components).

To show this, it is enough to prove a statement similar to Lemma 4.9 (roughly put, that the probability that a large component of the 5-core intersects the boundary of a box of size R is at most ε for some value of R). Here, R and ε can be given explicit numerical values, hence our methods reduce this question to a finite computation. Simulations suggest that this is indeed the case. Back of an envelope computations suggest $R = 25$ and $\varepsilon = 0.01$ will do.

A small difficulty involved in the case of 5 colors is that not every vertex is removed at some finite stage. Indeed, the 5-core of the Delaunay triangulation will not be empty, since it contains finite sub-graphs with minimal degree 5. The smallest such sub-graph is the dodecahedron, involving 12 vertices.

Applying the same proof for 4 colors cannot work, since the 4-core of the Delaunay triangulation has an infinite component. Indeed, a vertex of degree 3 is necessarily in the interior of the triangle formed by its neighbors. It is straightforward to check that there are no infinite chains of triangles each one inside the next (since the probability of long edges decays exponentially; see also Lemma 4.16 below). Therefore, one can consider all the *maximal* triangles in the Delaunay triangulation. This is also a triangulation of the plane, since every triangle is contained in a maximal one, and these are all disjoint. All the vertices of this triangulation also belong to G_∞ (since none of them are in the interior of another triangle), and they are all in the same connected component, which is therefore infinite.

Finally, while we only prove that some G_M (again, deleting vertices of degree < 6) has only finite connected components, simulations suggest that $M = 2$ suffices while $M = 1$ does not. In fact, it appears sufficient to delete in the second iteration roughly half the vertices of degree at most 5. Can one prove any of these assertions?

Other properties of colorings
If there is no deterministic 4-coloring, one could consider intermediate properties between deterministic and unrestricted randomized colorings. For example, one may seek colorings that are ergodic, mixing, finitary, etc. Such properties were first brought to our attention by Russ Lyons [21].

Other planar processes
It might be more interesting to consider other translation or isometry equivariant graph processes in the plane. These could be the Voronoi tessellation of some point process or more general planar graph processes. Except for some obvious counterexamples (see remarks before and after Lemma 3.1), is it true that every such process can be colored deterministically with 4 colors? The aforementioned work of Timár [23] shows the existence of deterministic 5-colorings.

Hyperbolic geometry
What can be done in the hyperbolic plane? Our argument can be adapted to give a deterministic coloring. However, the number of colors diverges as the density of the Poisson process tends to 0, since the average degree diverges. For high enough density we can get a deterministic 6-coloring. Is there a (deterministic or randomized) k-coloring with k independent of the density? While the Poisson–Voronoi map is 4-colorable by Proposition 1.2, our randomized constructions use amenability and fail for the hyperbolic plane.

Prescribed color distribution
What color distributions are achievable (with deterministic or randomized colorings)? We only show that coloring schemes exist such that the color of (say) the cell of 0 is supported on a finite set. If one asks for a particular distribution the question is interesting also in $\mathbb{R}^d$ for $d > 2$. For example, in $\mathbb{R}^d$, it is possible to get a coloring so that color i appears with exponentially (in i) small probability. What is the minimal possible entropy of the color of a cell?

4. Proof of the main result

In this section we prove Theorem 1.1. As explained in the proof outline, the proof is based on Propositions 2.1 and 2.2. These in turn will be proved by reduction to k-dependent percolation. Section 4.1 below gives the basic fact about k dependent percolation we shall need and introduces sealed squares, the tool which allows us to deduce that events taking place in distant locations are almost independent. In Section 4.2 we prove the simpler Proposition 2.2 and in Section 4.3 the more difficult Proposition 2.1. Section 4.4 shows how to deduce the main result from the two propositions.

Notation. *Throughout we shall denote by $G = (V, E)$ the Delaunay graph embedded in the plane where V is the set of points of the Poisson process and the edges are straight lines connecting these points* (*this can be seen to be a planar representation of G*). *We will sometimes call the vertices centers and say that a Voronoi cell is centered at its vertex. We also let $A : V \to \mathbb{R}_+$ be the function which assigns to each vertex the area of the corresponding Voronoi cell. For $x \in \mathbb{R}^2$ we denote $Q(x, R) := x + [-R, R]^2$, i.e., a square centered at x of side length $2R$. We let $B_R(x)$ or $B(x, R)$ stand for a closed ball of radius R around x* (*in the Euclidean metric*). *We write $d(x, y)$ for the Euclidean distance between $x, y \in \mathbb{R}^2$. Similarly $d(x, U) := \inf\{d(x, y) | y \in U\}$ for sets $U \subseteq \mathbb{R}^2$.*

4.1. *Dependent percolation and sealed squares*

A process $\{A_x\}_{x \in \mathbb{Z}^2}$ is said to be *k-dependent* if for any sets $S, T \subset \mathbb{Z}^2$ at ℓ^∞-distance at least k, the restrictions of A to S and to T are independent. Our processes will always take values in $\{0, 1\}$. Vertices $x \in \mathbb{Z}^2$ with $A_x = 1$ are called *open* (and others are *closed*). An *open component* is a connected component in $\mathbb{Z}^2$ of open vertices.

A well known result of Liggett, Schonmann and Stacey [20] states that k-dependent percolation with sufficiently small marginals ($\mathbb{E}A_x$) is dominated by sub-critical Bernoulli percolation. The following simple lemma is weaker, and is a standard argument in percolation theory. We include a proof for completeness:

Lemma 4.1. *For any k there is some $p_0 = p_0(k) < 1$ such that if $\{A_x\}_{x\in\mathbb{Z}^2}$ is k-dependent and for all x, $P(A_x = 1) \le p_0$, then*

$$\mathbb{P}(\exists \textit{ an infinite open component}) = 0.$$

Proof. The number of simple paths of length L starting at a given $x \in \mathbb{Z}^2$ is bounded by 4^L. Any simple path of length L contains at least $\frac{L}{k^2}$ coordinates which are pairwise k-separated. Thus, the probability that any given path of length L is open is at most p_0^{L/k^2}. The expected number of open paths originating at x is bounded by

$$4^L \cdot p_0^{L/k^2} = \left(4p_0^{1/k^2}\right)^L.$$

If $p_0 < 4^{-k^2}$ this quantity tends to 0 as L tends to infinity. However, an infinite open component must contain an open path of any length. □

Definition 4.2. *A set $S \subset \mathbb{R}^2$ is called* α-sealed *w.r.t. the point configuration V if $d(x, V) \le \alpha$ for every point $x \in \partial S$.*

Thus a set is sealed if the point process is not far from any point on the boundary of S. This implies that the Voronoi cells of V which intersect the boundary of S are centered near the boundary. The purpose of this term is that it bounds the dependency between the Voronoi map inside and outside the set. For a set S we denote

$$S^\alpha = \left\{x \in \mathbb{R}^2\colon\ d(x, S) \le \alpha\right\}$$

i.e. the closed (Euclidean) α-neighborhood of S (so that α-sealed is equivalent to $\partial S \subset V^\alpha$). Note that being α-sealed is determined by $V \cap (\partial S)^\alpha$ (the points in the α-neighborhood of the boundary). We denote by $S^{-\alpha}$ the points at distance at least α from the complement S^c (the idea is that if $S = B_R(x)$ then $S^\alpha = B_{R+\alpha}(x)$ for any $\alpha \ge -R$).

Lemma 4.3. *Condition on the points of $V \cap (\partial S)^\alpha$. On the event that S is α-sealed, the Voronoi map in $S^{-\alpha}$ is determined by the process $V \cap S^\alpha$. Moreover, the cell as well as all neighbors of $x \in V \cap S^{-\alpha}$ are contained in S^α.*

Proof. The lemma follows from the following simple geometrical fact: If $V \cap (\partial S)^\alpha$ is such that S is α-sealed, then the center of the cell of any $z \in \partial S$ is in $(\partial S)^\alpha$. Thus the cells of centers in $(\partial S)^\alpha$ separate $S^{-\alpha}$ from $\mathbb{R}^2 \setminus S^\alpha$. It follows that the cell of $x \in S^{-\alpha}$ is contained in S, and is adjacent only to cells centered in S^α. □

Next we argue that squares are likely to be α-sealed

Lemma 4.4. *The probability that $Q(0, R)$ is not α-sealed is at most*

$$\lceil 8R/\alpha \rceil \mathrm{e}^{-\pi\alpha^2/4}.$$

Proof. Take an $\alpha/2$ net in $\partial Q(0, R)$, of size $\lceil 8R/\alpha \rceil$. Each of these points fails to have a center within distance $\alpha/2$ from it with probability $\mathrm{e}^{-\pi\alpha^2/4}$. If none fail to have such a nearby center then the square is α-sealed. A union bound gives the claim. □

4.2. *Areas behave – Proposition 2.2*

Our present goal is to prove Proposition 2.2. To this end we need two properties of the areas of Poisson–Voronoi cells.

Lemma 4.5. *Let μ be the law of the area of the cell containing the origin, then μ is absolutely continuous w.r.t. the Lebesgue measure.*

A *partition* of $\mathbb{R}_+$ is a finite union $\mathbb{R}_+ = \bigcup_{i<M}[x_i, x_{i+1})$, given by a sequence $0 = x_0 < x_1 < \cdots < X_M = \infty$. The following is an immediate corollary of Lemma 4.5.

Corollary 4.6. *For any $R, \varepsilon > 0$ there is some sufficiently refined partition $\mathcal{A}$ of $\mathbb{R}_+$ such that for every interval $I \in \mathcal{A}$ the probability that there exists $v \in V \cap B_R(0)$ with $A(v) \in I$ is at most ε.*

Proof. Given R and ε, there is a finite $S \subset B_R(0)$ such that with probability at least $1 - \frac{\varepsilon}{2}$ every cell of every $v \in V \cap B_R(0)$ contains a point in S. Let $\mathcal{A}$ be a partition of $\mathbb{R}_+$ such that $\mu(I) < \frac{\varepsilon}{2|S|}$ for all $I \in \mathcal{A}$. Then for all $I \in \mathcal{A}$ the probability that there exists $v \in V \cap B_R(0)$ with $A(v) \in I$ is bounded by $\frac{\varepsilon}{2} + \frac{\varepsilon}{2|S|}|S| = \varepsilon$. □

However, just knowing that the area distribution is continuous is not enough, since the areas of different cells are not independent. For this reason we also need.

Lemma 4.7. *Almost surely, all cells have different areas.*

These two lemmas are intuitively obvious, though writing a precise proof is delicate. It is possible to get a somewhat simpler proof by replacing the area of a cell by some other quantity. For example, the total distance to the neighbors in the Delaunay graph works. However, the distance to the nearest neighbor does not work since there are pairs of centers that are mutually closest, and so have the same distance to their nearest neighbor.

Proof of Lemma 4.5. The idea of the proof is this: let x be the center of the cell of the origin and let y be the center of an adjacent cell. Conditioned on the location of *all centers other than* y, and on the direction of the vector $y - x$, we get that the area of x is an a.e. differentiable function of $r = \|y - x\|$, the distance between x and y, with positive derivative. Thus, μ conditioned on this σ-algebra is a.s. absolutely continuous w.r.t. Lebesgue and so μ itself is also absolutely continuous.

To make this precise, we partition $\mathbb{R}^d$ into cubes of size ε^d centered around $\varepsilon\mathbb{Z}^d$. We condition on the number of points of the Poisson process in each of these cubes. We then use finer and finer partitions (say, with $\varepsilon_i = 2^{-i}$) until we reach a partition which already reveals in what cube lies the center of the cell of the origin (i.e. x) as well as its nearest neighbor (i.e. y). We then continue according to the previous paragraph: we condition on the exact location of all points of the Poisson process except y and on the direction of $y - x$. After that we get that $A(x)$ is now a monotone function of $r = \|y - x\|$ and its derivative is equal to the length of the intersection of the cells of x and y, which is strictly positive, with probability 1. Since under this conditioning, the distribution of r is absolutely continuous w.r.t. Lebesgue measure on some interval we get that the conditioned μ is also absolutely continuous w.r.t. Lebesgue and so is μ itself. □

Proof of Lemma 4.7. The proof is similar to that of Lemma 4.5. Fixing any two points, a and b we wish to show that the probability that they belong to different cells with equal areas is zero. To that end, we find the two centers of the cells, x and y and find a third cell, centered at z, which is adjacent to one of these cells, say, x, but not to the other. (Such z exists for any x, y in any planar triangulation with no unbounded face.) Now $A(x)$ depends on the exact location of z, as in the proof of Lemma 4.5, but $A(y)$ does not. Of course, all this needs to be done using fine partitions, etc.

The lemma now follows by considering all possible values for a and b with rational coordinates. □

Note that the proof of Lemma 4.7 above does not apply as is to higher dimensions, since in such dimensions, there are configurations with two distinct cells having the same neighbors. Of course, Lemma 4.7 itself remains valid.

We now prove Proposition 2.2. The key idea is that cells with areas in any sufficiently small interval are dominated by sub-critical percolation.

Proof of Proposition 2.2. We show that there is some sufficiently refined partition $\mathcal{A}$ of $\mathbb{R}_+$, such that a.s. for any $I \in \mathcal{A}$ there is no infinite path in G with all areas in I. The proposition will follow since an infinite path with decreasing areas will have all areas in the same interval of $\mathcal{A}$ from some point on.

For some R to be determined later, consider the lattice $L = (2R\mathbb{Z})^2$. For an interval I, if there is an infinite path of cells with areas in I, then there is an infinite path $\{x_i\}$ in L so that every $Q(x_i, R)$ intersects such a cell. The probability that a square intersects a cell with area in I can be made arbitrarily small, but these events are not independent. To overcome this we use sealed boxes.

For an interval I, call a point $x \in L$ open if either $Q(x, R)$ intersects a cell with area in I, or if either one of $Q(x, R + \alpha)$ or $Q(x, R + 3\alpha)$ is not α-sealed. If there is an infinite path in G with areas in I then there is also an infinite open path in L.

The event that the squares are sealed depends only on the Poisson process within $Q(x, R + 4\alpha)$. We claim that on the event that they are sealed, the areas of cells intersecting $Q(x, R)$ also depend only on the process in $Q(x, R+4\alpha)$. Taking $\alpha = R/8$ it follows that the process of open boxes is 2-dependent. To see this claim, note that the center of any cell intersecting $Q(x, R)$ must be within $Q(x, R + 2\alpha)$. The second seal implies that the cell of this center is contained in $Q(x, R + 4\alpha)$ and determined by the process in this box.

To complete the proof, take some $\varepsilon > 0$ so that a 2-dependent percolation with marginal ε is sub-critical (using Lemma 4.1). Using Lemma 4.4, fix R large enough so that with $\alpha = R/8$,

$$\mathbb{P}\big(Q(x, R + i\alpha) \text{ is not } \alpha\text{-sealed}\big) < \varepsilon/3 \quad \text{for } i = 1, 3.$$

Next, using Corollary 4.6 take a partition $\mathcal{A}$ fine enough that for any $I \in \mathcal{A}$, the probability that there exists $v \in V \cap Q(x, R + 2\alpha)$ with area in I is at most $\varepsilon/3$. Then for each $I \in \mathcal{A}$, the probability that any fixed x is open is at most ε and so the process of open points does not contain an infinite open path. $\square$

4.3. *Deleting low degree vertices – Proposition 2.1*

In this section we prove Proposition 2.1. Throughout the section $R > 0$ is a parameter, assumed large enough as needed for the calculations which follow. We also define the square annuli $A(x, r, R) := Q(x, R) \setminus Q(x, r)$.

We now introduce our main object of study in this section:

Definition 4.8. *Inductively, let $G_0^R := G$ and let G_{n+1}^R denote the graph obtained from G_n^R by deleting all vertices in $Q(0, 3R)$ with G_n^R-degree at most 5. Let $G_\infty^R := \bigcap_{n=0}^\infty G_n^R$.*

Thus we iteratively delete vertices of degree at most 5, but only those vertices contained in a fixed large square. We aim to prove the following

Lemma 4.9. *We have $\mathbb{P}(G_\infty^R \cap Q(0, R) \neq \varnothing) \xrightarrow[R\to\infty]{} 0$.*

Corollary 4.10. *For any $\varepsilon > 0$, there are R, M so that $\mathbb{P}(G_M^R \cap Q(0, R) \neq \varnothing) < \varepsilon$.*

Proof. Pick R such that $\mathbb{P}(G_\infty^R \cap Q(0, R) \neq \varnothing) < \varepsilon$. Since

$$\big\{G_\infty^R \cap Q(0, R) \neq \varnothing\big\} = \bigcap_M \big\{G_M^R \cap Q(0, R) \neq \varnothing\big\},$$

the bound will hold for that R and sufficiently large M. $\square$

Before embarking on the proof of Lemma 4.9, let us explain how one can get a similar and simpler result when deleting vertices of degree at most 6 (thus yielding a deterministic 7-coloring). Suppose that G_∞^R contains a vertex in $Q(0, R)$. By Lemma 4.13 G_∞^R is unlikely to contain edges longer then $\log R$ within $Q(0, 3R)$. All vertices of G_∞^R in $Q(0, 3R)$ have degree at least 7. It is an easy consequence of Euler's formula that a planar graph with minimal degree 7 has positive expansion (the boundary of any set is proportional to its size). This implies (in the absence of long edges) that the number of vertices of G_∞^R in $Q(0,3R)$ is exponential in R. Of course, this too is unlikely. When deleting vertices of degree at most 5, the remaining graph has minimal degree 6, which is not as obviously unlikely. However, this can only happen if G contains a large section of the triangular lattice, which we rule out below.

We begin with two combinatorial lemmas on planar maps. For any finite graph H, let $\mathrm{LD} = \mathrm{LD}(H)$ be the number of vertices of low-degree, namely at most 5. For a finite simple planar map H, let $\mathrm{ME} = \mathrm{ME}(H)$ be the number of "missing edges": the number of edges that can be added to the map while keeping it planar and simple. A face of size k can be triangulated using $k-3$ edges, after which no further edges can be added, and so $\mathrm{ME} = \sum_f (\deg[f]-3)$ (where the sum also includes the external face and where we assume $|H| \geq 3$ so that $\deg[f] \geq 3$ for all faces).

Lemma 4.11. *For any finite, connected and simple planar map H with $|V[H]| \geq 3$ we have* $\mathrm{LD} \geq \frac{2}{5}\mathrm{ME} + \frac{12}{5}$.

Proof. Add ME edges to make the map into a triangulation. Let d'_v be the resulting vertex degrees, than we have $\sum_v (6-d'_v) = 12$ (using Euler's formula combined with the triangulation property $3F = 2E = \sum d'_v$), and therefore $\sum_v (6-d_v) = 12 + 2\mathrm{ME}$. The claim follows since low-degree vertices contribute at most 5 to this sum, and high degree vertices at most 0, so that $\sum_v (6-d_v) \leq 5\mathrm{LD}$. □

Lemma 4.12. *Fix $\rho > \ell > 0$. Let H be a simple planar graph embedded in $\mathbb{R}^2$ satisfying the following*:

1. *All vertices in $Q(0,3\rho)$ have degree at least* 6.
2. *All edges of H with an endpoint in $Q(0,3\rho)$ have length at most ℓ.*
3. *There exists a vertex of H in $Q(0,\rho)$.*

Then H has at least $\frac{8\rho^2}{5\ell^2}$ vertices in $Q(0,3\rho)$.

Note that the order of magnitude $(\rho/\ell)^2$ is achieved by a triangular lattice with edge length ℓ.

Proof of Lemma 4.12. We assume that H has only finitely many vertices in $Q(0,3\rho)$ since otherwise the conclusion is trivial. Fix a vertex $v \in Q(0,\rho)$. For $t \in [\rho, 3\rho]$, let H'_t be the sub-graph induced by vertices inside $Q(0,t)$, and let H_t be the connected component of v in H'_t.

Note that the connected component of v in H is not contained in $Q(0,3\rho)$ since otherwise it would be a finite, connected and simple planar map with all degrees at least 6 which is impossible by Lemma 4.11. By our assumptions, all vertices of H_t with neighbors in $H \setminus H_t$ (which includes all vertices of degree at most 5 in H_t) must be in the annulus $A(0, t-\ell, t)$. It follows that the external face of H_t surrounds v and exits $Q(0, t-\ell)$ and so has degree at least $\frac{2(t-\rho-\ell)}{\ell}$. Thus

$$\mathrm{ME}(H_t) \geq \frac{2(t-\rho-\ell)}{\ell} - 3 = \frac{2(t-\rho)}{\ell} - 5.$$

By Lemma 4.11, the number of vertices in $A(0, t-\ell, t)$ is at least $\frac{2}{5}\mathrm{ME}(H_t) + \frac{12}{5} \geq \frac{4}{5}\frac{t-\rho}{\ell} + \frac{2}{5}$. Let $M = \lfloor 2\rho/\ell \rfloor$. Splitting $A(0,\rho,3\rho)$ into annuli $A(0, \rho+(k-1)\ell, \rho+k\ell)$ for $k = 1, \dots, M$ one finds that the number of vertices of H in $Q(0,3\rho)$ is at least

$$1 + \sum_{k=1}^{M} \left(\frac{4}{5}k + \frac{2}{5}\right) = \frac{2(M+1)^2 + 3}{5} \geq \frac{2(2\rho/\ell)^2}{5}.$$ □

Next, a simple lemma showing that long edges in G are unlikely.

Lemma 4.13 (No long edges). *The probability of having an edge of length at least ℓ in $E[G]$ which intersects the square $Q(0,\rho)$ is at most*

$$\left(\frac{\sqrt{32}\rho}{\ell} + 8\right)^2 \mathrm{e}^{-\ell^2/32}.$$

Proof. Suppose (x, y) were such an edge, then some disc with x, y on its boundary has no points in its interior. Consequently at least one of the two semi-circles with diameter (x, y) has no points in its interior. This implies that there is an empty disc $B_{\ell/4}(z)$ for some $z \in Q(0, \rho + \ell)$ (z might be outside $Q(0, \rho)$ since one of x, y may be outside the square).

Cover $Q(0, \rho + \ell)$ by $\lceil \frac{\rho+\ell}{\ell/\sqrt{32}} \rceil^2$ squares of side length $\ell/\sqrt{32}$. It follows that if such a long edge exists than one of the squares (the one containing z) must be empty, and the claim follows. □

We continue by showing that after some low-degree vertices are deleted, many large holes remain in the graph.

Definition 4.14. *Call a square $Q(x, \rho)$ a typical square if there exists some vertex $v \in Q(x, \rho)$ such that*:

1. $\deg[v] < 6$.
2. *v is not in the interior of any triangle in the Delaunay Graph G.*

Otherwise we call the square rare.

To make this clear, the second condition states that there are no $v_1, v_2, v_3 \in V$ which are pairwise adjacent in G such that v is contained in the interior of the triangle (v_1, v_2, v_3).

Lemma 4.15. *There exist points $v \in V$ of degree at most 5 that are not contained in the interior of any triangle.*

Proof. Since the event is invariant, its probability is either 0 or 1, by ergodicity of the Poisson process.

Consider the following event: There are exactly 5 points in $B(0, 3)$, located within 0.1 of each of $(0, 0)$, $(0, 1)$, $(0, -1)$, $(1, 0)$, $(-1, 0)$, and additionally, for every $n > 3$, the ball $B(0, 1.5^n)$ is 1.5^{n-10} sealed. It is easy to check that this event has positive probability.

On this event, the point near the origin has degree 4. If $\{x, y, z\}$ is a triangle in the Voronoi map of diameter R with the point near 0 in its interior, then there is an empty semi-disc with diameter R within distance R of 0. On our event, this does not occur. □

Lemma 4.16 (Rare squares are rare). *For some $\alpha, \beta > 0$ we have $\mathbb{P}(Q(x, \rho)$ is rare$) \le \alpha \exp(-\beta\rho)$.*

Proof. We may assume without loss of generality that $\rho \ge C$ for some large $C > 0$ (otherwise the claim is trivial). Let $\gamma = \sqrt{\rho}$. The square $Q(0, \rho)$ contains at least $c\rho$ disjoint squares $Q(x, 4\gamma)$ for some $c > 0$. Call each of these squares *good* if it satisfies the following:

1. $Q(x, 3\gamma)$ is γ-sealed.
2. $Q(x, \gamma)$ contains a vertex v of degree at most 5 which is not in the interior of any triangle with vertices in $Q(x, 2\gamma)$.

Note that by Lemma 4.3, the event that $Q(x, 4\gamma)$ is good is determined by the Poisson process within it, and so these events are all independent. For large γ, $Q(x, 3\gamma)$ is likely to be γ-sealed. By the previous lemma, $Q(x, \gamma)$ is likely to contain a vertex as needed, hence each square has probability at least $p > 0$ of being good (independent of ρ, for ρ large enough). Thus the probability that no square within $Q(0, \rho)$ is good is at most $e^{-\beta\rho}$ for some $\beta > 0$.

If the low-degree vertex in a good square is contained in a triangle of G then an edge of that triangle must have length at least γ. Either the triangle intersects $Q(0, \rho)$, which by Lemma 4.13 has probability at most $C_1\rho e^{-\rho/32} \le C_2 e^{-\rho/33}$ for some $C_1, C_2 > 0$. Or the triangle contains $Q(0, \rho)$ in its interior, in which case for some integer $m \ge 1$, its longest edge has length at least $m\rho$ and intersects $Q(0, m\rho)$. By a union bound, this has probability at most $\sum_{m=1}^{\infty} C_2 e^{-m^2\rho^2/33} \le C_3 e^{-\rho^2/33}$ for some $C_3 > 0$. □

In what follows, define

$$L := \log R, \qquad r := R^{1/3}.$$

L will be a bound on lengths of edges that appear in $Q(0, 3R)$ (and can be reduced to $C\sqrt{\log R}$ for some large C). The role of r is more involved, and there is much freedom in its choice. Primarily, we consider a partition of boxes of size of order R into boxes of size r. For simplicity, we assume that $6R/r$ is an odd integer (R can be arbitrarily large under this condition).

Define for each $x \in r\mathbb{Z}^2$ the square $Q_x := Q(x, \frac{r}{2})$. Note that $Q(0, 3R)$ is precisely tiled by the boxes $\{Q_x, x \in r\mathbb{Z}^2 \cap Q(0, 3R)\}$. We now define several events which we will show to be unlikely.

$$\begin{aligned}
\Omega_0 &:= \big\{G^R_\infty \text{ has a vertex in } Q(0, R)\big\},\\
\Omega_1 &:= \big\{\text{There exists } e \in E[G] \text{ of length at least } L \text{ that intersects } Q(0, 3R)\big\},\\
\Omega_2 &:= \big\{Q_x \text{ is rare for some } x \in r\mathbb{Z}^2 \cap Q(0, 3R)\big\},\\
\Omega_3 &:= \big\{\text{There exists } x \in r\mathbb{Z}^2 \cap Q(0, 3R) \text{ and } |V \cap Q_x| \geq 2r^2\big\},\\
\Omega_4 &:= \big\{|V \cap A| > 2\,\mathrm{Area}(A)\big\}, \qquad \text{where } A = A(0, 3R, 3R + L).
\end{aligned}$$

Thus Lemma 4.9 states that $\mathbb{P}(\Omega_0)$ is small.

Lemma 4.17. *With L, r as above, $\mathbb{P}(\Omega_i) \xrightarrow[R\to\infty]{} 0$ for $i = 1, 2, 3, 4$.*

Proof. Lemma 4.13 implies that $\mathbb{P}(\Omega_1) = \mathrm{O}(R^2 \mathrm{e}^{-L^2/32})$ is small. Lemma 4.16 implies $\mathbb{P}(\Omega_2) = \mathrm{O}(R^{4/3} \mathrm{e}^{-\beta r})$ (since there are $(6R/r)^2$ squares to consider).

$\mathbb{P}(\Omega_3)$ and $\mathbb{P}(\Omega_4)$ are bounded by the fact that $\mathbb{P}(\mathrm{Poi}(\lambda) > 2\lambda) \leq \mathrm{e}^{-c\lambda}$ for some constant c. This gives respective bounds $\mathrm{O}(R^2 \mathrm{e}^{-cr^2})$ and $\mathrm{O}(\mathrm{e}^{-cRL})$. □

Define the set

$$S := \big\{x \in r\mathbb{Z}^2 \cap Q(0, 3R)\colon\ Q_x \text{ is typical and } G^R_\infty \cap Q_x \neq \varnothing\big\}.$$

Lemma 4.18. *There exists $C > 0$ such that if Ω_1^c holds then*

$$|S| \leq C\big|V \cap A(0, 3R, 3R + L)\big|.$$

Proof. Let H be the sub-graph of G^R_∞ induced by vertices in $Q(0, 3R + L)$. On the event Ω_1^c, the vertices of $H \cap Q(0, 3R)$ all have degree at least 6. Thus low-degree vertices are all in the annulus $A(0, 3R, 3R + L)$ and by Lemma 4.11,

$$\big|V \cap A(0, 3R, 3R + L)\big| \geq \mathrm{LD}(H) > \frac{2}{5}\mathrm{ME}(H).$$

For each $x \in S$ the square Q_x is typical. Hence there is a vertex $v_x \in Q_x$ of degree at most 5 that is not contained in any triangle in G. The vertex v_x is deleted in the first round and so is not in H. Let f_x be the face of H surrounding v_x. Note that f_x must have an edge e_x that intersects Q_x, since otherwise Q_x is completely in the interior of f_x and there could be no vertex of H in Q_x.

Now, on Ω_1^c, the edge e_x has length at most $L < r$ and therefore can intersect at most 3 different squares Q_x (it can intersect 3 if it passes near a corner of Q_x). Since the face f_x cannot be a triangle by definition of v_x we deduce that

$$\sum_{\substack{f \text{ face of } H\\ \deg[f]>3}} \deg[f] \geq \frac{1}{3}|S|.$$

Hence $\mathrm{ME}(H) = \sum_{f\in H}(\deg[f] - 3) \geq \frac{1}{12}|S|$ proving the claim (with $C = 30$). □

Proof of Lemma 4.9. We show that $\Omega_0 \subset \bigcup_{i=1}^4 \Omega_i$. Assume by negation that Ω_0 and Ω_i^c hold for $i = 1, 2, 3, 4$. Let H be the restriction of G^R_∞ to $Q(0, 3R + L)$, and apply Lemma 4.12 with $\rho = R$, $\ell = L$. Ω_1^c and Ω_0 show that the lemma's hypotheses hold, thus H has at least $\frac{8R^2}{5L^2}$ vertices in $Q(0, 3R)$.

On the other hand we show that H is small. Tile $Q(0,3R)$ by boxes Q_x with $x \in r\mathbb{Z}^2 \cap Q(0,3R)$. On Ω_4^c, Lemma 4.18 implies that $|S| \le CRL$ for some C. On Ω_3^c each of these includes at most $2r^2$ vertices, so the number of vertices of H in $Q(0,3R)$ that are in typical boxes is at most $2r^2|S| \le CRLr^2$. On Ω_2^c there are no vertices in rare boxes.

Thus $\frac{8R^2}{5L^2} \le CRLr^2$ which is a contradiction for R large enough and our choice of r and L. □

Proof of Proposition 2.1. Define $G_M^{R,x}$ similarly to G_M^R, except that low degree vertices are deleted in $Q(x,3R)$ instead of $Q(0,3R)$. Consider the following dependent percolation process on the lattice $\Lambda = R\mathbb{Z}^2$. A point x is open in one of 3 cases:

1. The square $Q(x,4R)$ is not R-sealed.
2. $G_M^{R,x}$ has a vertex in $Q(x,R)$.
3. G has an edge of length at least $R/2$ intersecting $Q(x,R/2)$.

We first argue that the event $\{x \text{ is open}\}$ is determined by the Poisson process in $Q(x,5R)$, so that the process is 11-dependent. Indeed, whether $Q(x,4R)$ is R-sealed depends only on the process in $Q(x,5R)$. If it is R-sealed, the restriction of the Voronoi map to $Q(x,3R)$ is determined by the process in $Q(x,5R)$, which determine the state of x.

By Lemmas 4.4, 4.9 and 4.13, we can choose M, R so that $\mathbb{P}(x \text{ is open})$ is arbitrarily small. In particular, for some M, R, using Lemma 4.1, this percolation is dominated by sub-critical percolation, and has no infinite open component.

Finally, we argue that if there were an infinite component in G_M then there would also be an infinite component in our process on Λ. Consider all squares $Q(x,R/2)$ which intersect the edges of some infinite open component in G_M. For each such x, either there is a vertex of G_M in $Q(x,R)$, or else the edge that passes through $Q(x,R/2)$ has both endpoints outside $Q(x,R)$. Since $G_M \subset G_M^{R,x}$, either case implies x is open. □

4.4. *Equivariant coloring*

We now use Propositions 2.1 and 2.2 to construct a deterministic 6-coloring scheme. Recall that G_n is derived from G_{n-1} by deleting low degree vertices. Define the *level* of a vertex by

$$\ell(v) = \max\{n\colon\ v \in G_n\}.$$

Thus a vertex has level 0 iff its degree is at most 5. For neighboring v, w we direct the edge from v to w, and write $v \to w$, if either $\ell(w) > \ell(v)$ or $(\ell(v) = \ell(w)$ and $A(w) < A(v))$ (recall that by Proposition 2.2 no two areas are equal).

Let $\prec$ to be the transitive closure of $\to$. That is, $w \prec v$ iff there is a finite sequence such that $v = u_0 \to u_1 \to \cdots \to u_n = w$.

Lemma 4.19. *A.s. every $v \in V$ has finitely many $\prec$-predecessors (in particular, $\prec$ is well founded).*

Proof. We first argue that there is no infinite directed path in G. By Proposition 2.2 there are no infinite A-monotone paths, so any infinite directed path must have $\ell(v) \to \infty$. However, by Proposition 2.1 there are no infinite paths with $\ell > M$.

Our conclusion then follows from König's lemma: A locally finite tree with no infinite paths is finite. □

From this we get:

Proposition 4.20. *There exists a unique function $f : V \to \{0, \dots, 5\}$ determined by the recursive formula:*

$$f(u) = \operatorname{mex}\{f(v) : u \to v\},$$

where $\operatorname{mex} S = \min(\mathbb{N} \setminus S)$ *is the minimal excluded integer function.*

Proof. The proof is by induction on $\prec$, which is a well founded order by Lemma 4.19 (see e.g. [19], Chapter 3). Every edge (x, y) of G has $x \prec y$ or $y \prec x$, and any $u \in V$ has at most 5 neighbors v with $\ell(v) \geq \ell(u)$. Thus $|\{v\colon u \to v\}| \leq 5$ and so $f(u) < 6$ and is well defined.

Uniqueness holds since $f(u)$ is determined by $\{f(v)\colon v \prec u\}$. □

Theorem 1.1 now follows, since $\mathcal{P}_i = f^{-1}(i)$ defines a deterministic, isometry equivariant 6-coloring. Note that the resulting coloring is finitary, that is, for every $x \in \mathbb{R}^2$ there exists a finite (but random) $R > 0$ such that the color of the cell containing x is a function of the Poisson process restricted to $B_R(x)$. Indeed, to determine $f(v)$ for $v \in V$, it is sufficient to know the graph G induced on the $\prec$-predecessors of v, which are only finitely many, and only a bounded portion of the plane needs to be explored to do so. Furthermore, there exist $C, c > 0$ such that $\mathbb{P}(R > s) \leq C\mathrm{e}^{-cs}$. This is the case because Proposition 2.2 is proved using domination by sub-critical percolation.

More explicitly, to determine the color of the cell containing the origin, we start by finding out its level. In order to determine whether a given cell belongs to G_{k+1}, we need only know which of its neighbors belong to G_k. Hence, it is enough to observe all the cells of graph distance at most $M + 1$. Consider all squares of side length R, centered at (iR, jR), where $i, j \in \{-M - 1, \dots, M + 1\}$. If all of these are $R/10$-sealed then all cells of graph distance at most $M + 1$ are inside the square of side length $2MR$. This event fails to happen with only exponentially small probability.

Next, if our cell belongs to G_M, we only need to know the location of those cells in its connected component. This component is dominated by subcritical k-dependent percolation (by the proof of Proposition 2.1) and hence its radius also has exponential decay, as is evident from the proof of Lemma 4.1.

Next, we see that by the proof of Proposition 2.2, for small enough ε, and any $n \geq 0$, the components of the graph of cells with areas between $n\varepsilon$ and $(n+1)\varepsilon$ also have exponential decay of their diameter. Hence, the radius needed to determine the color of a cell of G_{M-1} with area less than ε, also decays exponentially, since this color depends only on the other cells in G_{M-1} with areas less than ε and on the cells in G_M. We now proceed by induction to prove that cells in G_{M-1} with areas between $n\varepsilon$ and $(n+1)\varepsilon$ also have this property. A similar argument shows that the subgraph of cells with areas more than $n\varepsilon$, for some large n, is also dominated by a subcritical k-dependent percolation and hence the components of this subgraph also exhibit exponential decay. Thus, after finitely many steps we are able to extend our argument to all the cells in G_{M-1}. Similarly, we continue by backward induction to G_{M-2} etc. until the entire graph $G = G_0$ is shown to have exponential decay for the radius of exploration needed to determine the color of a cell.

References

[1] O. Angel, G. Amir and A. Holroyd. Multi-color matching. Unpublished manuscript.

[2] O. Angel, A. Holroyd and T. Soo. Deterministic thinning of finite Poisson processes. *Proc. Amer. Math. Soc.* **139** (2011) 707–720. MR2736350

[3] K. Ball. Poisson thinning by monotone factors. *Electron. Comm. Probab.* **10** (2005) 60–69 (electronic). MR2133893

[4] I. Benjamini, A. Holroyd, O. Schramm and D. Wilson. Finitary coloring. Unpublished manuscript.

[5] B. Bollobás and O. Riordan. The critical probability for random Voronoi percolation in the plane is 1/2. *Probab. Theory Related Fields* **136** (2006) 417–468. MR2257131

[6] S. Chatterjee, R. Peled, Y. Peres and D. Romik. Phase transitions in gravitational allocation. Preprint. Available at http://arxiv.org/abs/0903.4647. MR2729280

[7] S. Chatterjee, R. Peled, Y. Peres and D. Romik. Gravitational allocation to Poisson points. *Ann. of Math. (2)* **172** (2010) 617–671. MR2680428

[8] D. J. Daley and G. Last. Descending chains, the lilypond model, and mutual-nearest-neighbour matching. *Adv. in Appl. Probab.* **37** (2005) 604–628. MR2156551

[9] M. Deijfen and R. Meester. Generating stationary random graphs on $\mathbb{Z}$ with prescribed independent, identically distributed degrees. *Adv. in Appl. Probab.* **38** (2006) 287–298. MR2264945

[10] A. K. Dewdney and J. K. Vranch. A convex partition of $\mathbf{R}^3$ with applications to Crum's problem and Knuth's post-office problem. *Utilitas Math.* **12** (1977) 193–199. MR0482539

[11] C. Hoffman, A. E. Holroyd and Y. Peres. A stable marriage of Poisson and Lebesgue. *Ann. Probab.* **34** (2006) 1241–1272. MR2257646

[12] C. Hoffman, A. E. Holroyd and Y. Peres. Tail bounds for the stable marriage of Poisson and Lebesgue. *Canad. J. Math.* **61** (2009) 1279–1299. MR2588423

[13] A. Holroyd, R. Lyons and T. Soo. Poisson splitting by factors. *Ann. Probab.* To appear. Available at http://arxiv.org/abs/0908.3409.

[14] A. Holroyd, R. Pemantle, Y. Peres and O. Schramm. Poisson matching. *Ann. Inst. H. Poincaré Probab. Statist.* **45** (2009) 266–287. MR2500239

[15] A. E. Holroyd and Y. Peres. Trees and matchings from point processes. *Electron. Comm. Probab.* **8** (2003) 17–27 (electronic). MR1961286

[16] A. E. Holroyd and Y. Peres. Extra heads and invariant allocations. *Ann. Probab.* **33** (2005) 31–52. MR2118858

[17] G. Kozma. Private communication, 2007.
[18] M. Krikun. Connected allocation to Poisson points in $\mathbb{R}^2$. *Electron. Comm. Probab.* **12** (2007) 140–145 (electronic). MR2318161
[19] K. Kunen. *Set Theory: An Introduction to Independence Proofs*. *Studies in Logic and the Foundations of Mathematics* **102**. North-Holland, Amsterdam, 1980. MR0597342
[20] T. M. Liggett, R. H. Schonmann and A. M. Stacey. Domination by product measures. *Ann. Probab.* **25** (1997) 71–95. MR1428500
[21] R. Lyons. Private communication, 2007.
[22] F. P. Preparata. Steps into computational geometry. Technical report, Coordinated Science Laboratory, Univ. Illinois, 1977.
[23] A. Timár. Invariant colorings of random planar graphs. Preprint. Available at arXiv:0909.1091.
[24] A. Timár. Invariant matchings of exponential tail on coin flips in $\mathbb{Z}^d$. Preprint. Available at arXiv:0909.1090.
[25] A. Zvavitch. The critical probability for Voronoi percolation. Master's thesis, Weizmann Institute of Science, 1996. Available at http://www.math.kent.edu/~zvavitch/.