The critical probability for random Voronoi percolation in the plane is 1/2

Béla Bollobás[*†‡§] Oliver Riordan[†§¶]

October 14, 2004

Abstract

We study percolation in the following random environment: let Z be a Poisson process of constant intensity on $\mathbb{R}^2$, and form the Voronoi tessellation of $\mathbb{R}^2$ with respect to Z. Colour each Voronoi cell black with probability p, independently of the other cells. We show that the critical probability is $1/2$. More precisely, if $p > 1/2$ then the union of the black cells contains an infinite component with probability 1, while if $p < 1/2$ then the distribution of the size of the component of black cells containing a given point decays exponentially. These results are analogous to Kesten's results for bond percolation in $\mathbb{Z}^2$.

The result corresponding to Harris' Theorem for bond percolation in $\mathbb{Z}^2$ is known: Zvavitch noted that one of the many proofs of this result can easily be adapted to the random Voronoi setting. For Kesten's results, none of the existing proofs seems to adapt. The methods used here also give a new and very simple proof of Kesten's Theorem for $\mathbb{Z}^2$; we hope they will be applicable in other contexts as well.

Contents

[*]Department of Mathematical Sciences, University of Memphis, Memphis TN 38152, USA

[†]Trinity College, Cambridge CB2 1TQ, UK

[‡]Research supported in part by NSF grant ITR 0225610 and DARPA grant F33615-01-C-1900

[§]Research partially undertaken during a visit to the Forschungsinstitut für Mathematik, ETH Zürich

[¶]Royal Society Research Fellow, Department of Pure Mathematics and Mathematical Statistics, University of Cambridge, UK

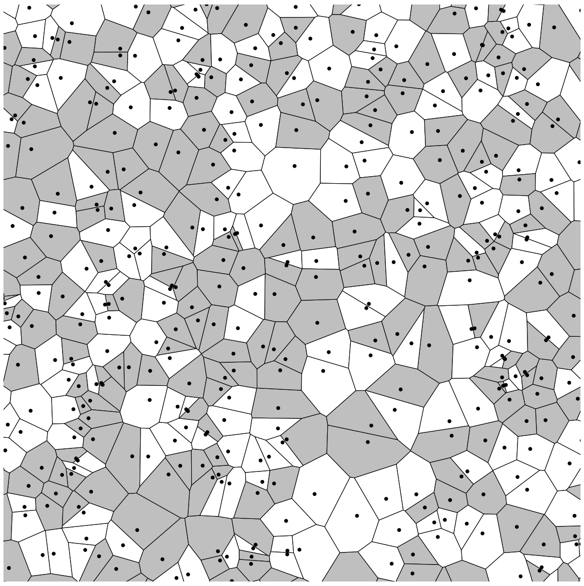

Figure 1: Critical random Voronoi percolation.

1 Introduction and results

Percolation theory is a very active area, with many difficult basic questions still unanswered. Perhaps the most basic question is to establish the critical probability above which percolation occurs; this question was posed by Broadbent and Hammersley [10] in 1957 in a wide variety of contexts. In general, it seems impossible to answer: the exact value is known only in a small number of cases in which something very special happens, involving duality. The best known case is Kesten's celebrated result that for bond percolation in $\mathbb{Z}^2$, which is self-dual, the critical probability is $1/2$ (more about this later). For other exact results, see Grimmett [14]. Another very natural percolation process for which self-duality occurs is *random Voronoi percolation*, introduced in the context of first-passage percolation by Vahidi-Asl and Wierman [27]; we come to the formal description in a moment. Random Voronoi percolation has been studied by many people; see, for example [27, 28, 29, 30, 1, 6, 12, 5]. Unlike the classical lattice examples, here the environment in which percolation occurs is itself random; as we shall see, this means that the techniques needed to establish the critical probability are rather different.

The formal set-up is as follows. A certain parameter p is given. We construct a Poisson process Z with intensity 1 on $\mathbb{R}^2$, and, given Z, assign to each point z of Z independently a colour $\mathrm{col}(z) \in \{\text{black, white}\}$, with $\mathbb{P}(\mathrm{col}(z) = \text{black}) = p$. Equivalently, we may generate Z as the union of two Poisson processes Z_b, Z_w with intensities p, $1-p$, corresponding to the black and white points, i.e., to $\{z \in Z : \mathrm{col}(z) = \text{black}\}$ and $\{z \in Z : \mathrm{col}(z) = \text{white}\}$. Throughout, we shall write $\mathbb{P}_p^{\mathbb{R}^2}$, or simply $\mathbb{P}_p$, for the associated probability measure. We use colouring terminology, as in [26], for example, rather than the more common 'open or closed' terminology; this emphasizes the symmetry and is more natural for figures.

Given Z, we construct the associated Voronoi tiling: the *Voronoi cell* of $z \in Z$ with respect to Z is the set

$$V(z) = V_Z(z) = \{x \in \mathbb{R}^2 : d(x,z) = \inf_{z' \in Z} d(x,z')\},$$

where $d(x,y)$ is the Euclidean metric.

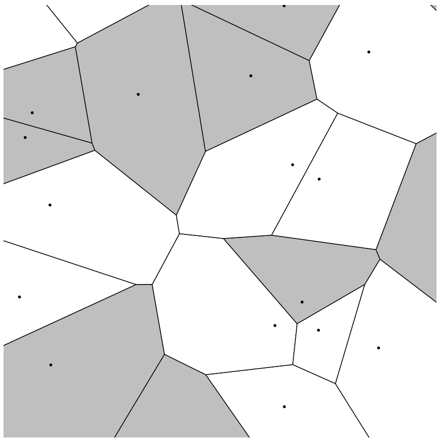

Figure 2: Part of the black (grey) and white-coloured random Voronoi tiling in $\mathbb{R}^2$. The dots are the points of Z.

As we shall note below, with probability 1, every cell is a closed convex k-gon for some k, every vertex of every cell lies in exactly three cells, and if two cells meet, they share an entire edge. Note, however, that cells $V(z_1)$ and $V(z_2)$ may meet without their common edge containing the midpoint of $z_1 z_2$. (There are several examples in Figure 2.) The *graph* of the Voronoi tiling is the countably infinite graph G with Z as vertex set, in which two points z_1, z_2 are adjacent if their Voronoi cells meet. The notions of percolation on the Voronoi tiling will be the usual notions of site percolation on the (random) graph G.

We take the colour of a Voronoi cell $V_Z(z)$ to be the colour of the point z, and say that each $x \in V_Z(z)$ also has this colour. (Thus, a point $x \in \mathbb{R}^2$ may be both black and white if it is in the boundary of two (or three) Voronoi cells.) A *path of cells* corresponds to a path in G. Let z_0 be the (with probability 1 unique) point $z_0 \in Z$ such that the origin lies in $V_Z(z_0)$. In the graph G, the *black vertex cluster containing* v_0 is the set C_0^G of all black vertices of G connected to v_0 by a path consisting of black vertices. Correspondingly, the *black component of the origin* C_0 is the maximal connected set of black points in $\mathbb{R}^2$ containing the origin, i.e., the union of the cells $V_Z(z)$, $z \in C_0^G$.

Writing area(.) and diam(.) for the usual geometric area and diameter of a set in $\mathbb{R}^2$, consider the following three conditions: $|C_0^G| = \infty$, $\mathrm{area}(C_0) = \infty$ and $\mathrm{diam}(C_0) = \infty$. It is easy to see that, except on a set of measure zero, either all three of these conditions hold or none. Let

$$\theta(p) = \mathbb{P}_p(|C_0^G| = \infty).$$

Equivalently, $\theta(p)$ may be defined in terms of C_0. We say that the random Voronoi percolation with parameter p *percolates* if $\theta(p) > 0$. As translationally invariant events have probability 0 or 1, it is easy to show that if $\theta(p) > 0$ then with probability 1 there is an infinite black component somewhere in $\mathbb{R}^2$, while if $\theta(p) = 0$ then with probability 1 there is no such component.

As $\theta(p)$ is increasing, there is a critical probability

$$p_H = \inf\{p : \theta(p) > 0\}.$$

Here, following Welsh (see [25]), the H is in honour of Hammersley.

There is another natural definition of critical probability. Let

$$\chi(p) = \mathbb{E}_p|C_0^G|,$$

and let

$$\chi^{\mathrm{area}}(p) = \mathbb{E}_p \mathrm{area}(C_0).$$

Of course, we cannot determine the exact relationship between $\chi(p)$ and $\chi^{\mathrm{area}}(p)$. However, it will follow from our results that, unsurprisingly, one is finite if and only if the other is. As $\chi(p)$ is increasing in p, there is another critical probability,

$$p_T = \inf\{p : \chi(p) = \infty\},$$

where the T is in honour of Temperley.

The quintessential percolation problem is the study of *bond percolation in* $\mathbb{Z}^2$. Here $\mathbb{Z}^2$ is the planar square lattice, i.e., the graph with vertex set $\mathbb{Z}^2$ in which two vertices are adjacent if they are at Euclidean distance 1. Each edge of this graph is taken to be *open* with probability p independently of all other edges. Edges that are not open are *closed*. The *open cluster* containing a vertex

v is the maximal connected subgraph of $\mathbb{Z}^2$ that contains v, all of whose edges are open. In this context, one can define $p_H(\mathbb{Z}^2)$ and $p_T(\mathbb{Z}^2)$ as above.

In 1957, Broadbent and Hammersley [10] posed the problem of determining p_H in a wide variety of contexts, including bond percolation in $\mathbb{Z}^2$. Hammersley [15, 16, 17] proved general results implying in particular that $0.35 < p_H(\mathbb{Z}^2) < 0.65$.

As the square lattice $\mathbb{Z}^2$ is self-dual, if $p = 1/2$ and R is an $n+1$ by n rectangle then the probability that there is a path of open edges joining some vertex on the left-hand side of R to some vertex on the right-hand side is *exactly* $1/2$. This observation suggests very strongly that $p_H(\mathbb{Z}^2) = 1/2$. Nevertheless, proving this turned out to be highly non-trivial. The first major progress was made by Harris [18] in 1960, when he proved that $p_H(\mathbb{Z}^2) \geq 1/2$. Harris' proof used self-duality as a starting point, but this trivial observation is *just* a starting point; several substantial ideas were needed to deduce the result. From this point it became a well known problem to show that $p_H(\mathbb{Z}^2) = p_T(\mathbb{Z}^2) = 1/2$. Indeed, by the time Kesten [20] finally proved this result 20 years later, the problem was sufficiently well known that Kesten quite rightly felt it needed no introduction: the abstract of his paper, whose title is 'The critical probability of bond percolation on the square lattice equals 1/2', reads 'We prove the statement in the title of the paper'!

The random Voronoi percolation we consider here also has a 'self-duality' property which implies that for $p = 1/2$ the appropriately defined crossing probability for a square is exactly $1/2$. In the light of this property and the Harris-Kesten results, it is extremely natural to conjecture that for the Voronoi set-up, $p_T = p_H = 1/2$ holds also. By adapting one of the many known proofs of Harris' Theorem, Zvavitch [30] has shown that $\theta(1/2) = 0$. The proof Zvavitch adapted is due to Burton and Keane [11] and Zhang; see [14, pages 198 and 289]. Other known proofs of Harris' Theorem do not seem to adapt in this way.

Here we shall prove the analogue of Kesten's results for bond percolation in $\mathbb{Z}^2$, showing that $p_T = p_H = 1/2$ holds for random Voronoi percolation. Although no known proof of Kesten's results seems to translate to the random Voronoi context, the methods used in this paper do give a new and very simple proof of Kesten's Theorem, which we shall describe in [7].

Theorem 1. *For random Voronoi percolation in the plane, $p_H \leq 1/2$. In other words, if $p > 1/2$ then $\theta(p) > 0$.*

In the subcritical case, we establish exponential decay of the 'volume' of the black component containing a given point. Here we may take $|C_0|$ to be $\mathrm{area}(C_0)$, $\mathrm{diam}(C_0)$, or the number $|C_0^G|$ of Voronoi cells making up C_0.

Theorem 2. *For any $p < 1/2$ there is a constant $c(p) > 0$ such that*

$$\mathbb{P}_p(|C_0| \geq n) \leq \exp(-c(p)n)$$

for every $n \geq 1$. In particular, $p_T \geq 1/2$.

As $p_T \leq p_H$, it follows that $p_T = p_H = 1/2$.

One of our intermediate results also gives an alternative proof of Zvavitch's result that $\theta(1/2) = 0$; we state this as a theorem for ease of reference. In principle, our method gives a bound on the rate of decay of $\mathbb{P}_{1/2}(|C_0^G| \geq n)$ as $n \to \infty$, but this bound would be rather weak and we have not calculated it explicitly.

Theorem 3 (Zvavitch [30]). *For random Voronoi percolation in the plane,* $\theta(1/2) = 0$.

The rest of the paper is organized as follows. In Section 2 we collect the results of general probabilistic combinatorics we shall use, and in Section 3 we present some basic facts about random Voronoi percolation. Our first intermediate result, Theorem 12, which will play the role of the Russo-Seymour-Welsh (RSW) Theorem of [24] and [25], is stated and proved in Section 4.

In the remaining sections, our aim is to apply a sharp-threshold result of Friedgut and Kalai [13]; see Section 3. In order to do this, we work with random Voronoi percolation not in $\mathbb{R}^2$, but in a torus, introduced in Section 5. We need to approximate the relevant Poisson process by a discrete process: preliminary lemmas for doing this are given in Section 6. The actual application of the sharp-threshold result is given in Section 7; the proofs of the main results follow in Section 8. In the course of the paper, for technical reasons we need several deterministic lemmas that are not terribly appealing. The proofs of these are postponed to the appendix.

There are two parts of the overall proof that are longer than one might expect. The first is our proof of an analogue of the RSW Theorem, given in Section 4. For bond percolation in $\mathbb{Z}^2$, several proofs of the RSW Theorem are now known, and it is possible to give a very short proof. However, none of the existing proofs carries over to the random Voronoi case, which lacks an important independence property needed. This is discussed further in Section 4.

Secondly, to apply the Friedgut-Kalai sharp-threshold result, we need to approximate a Poisson process with a discrete probability space in some way. We work with a torus T, which has finite area, rather than $\mathbb{R}^2$. Of course, one can 'discretize' a Poisson process on T, essentially by rounding the coordinates of all points to multiples of a small quantity δ. However, there will be a limit as to how small we can take δ in terms of the area of T, so this approximation will introduce 'defects', i.e., places where the discretized Poisson process does not tell us which of the cells in the Voronoi tiling meet. There are not many defects, and, intuitively, it is clear that these defects cannot affect the critical probability. However, proving this turns out to be rather taxing.

Let us remark that this difficultly with defects is not unique to the present paper. Benjamini and Schramm [6] prove a certain 'conformal invariance' property of random Voronoi percolation. This is *not* conformal invariance in the sense of the celebrated conjecture of Aizenman, Langlands, Pouliot and Saint-Aubin [22], which is believed to hold in a very wide variety of contexts, and has been proved for site percolation in the triangular lattice by Smirnov [26]. Rather, Benjamini and Schramm prove (essentially) the following statement specific to Voronoi percolation: fix a region $R \subset \mathbb{R}^2$, and two segments S_1 and S_2 of its boundary. Consider the Voronoi percolation associated with a Poisson process of intensity λ on R, using a certain metric $\mathrm{d}s$ to form the Voronoi cells, rather than the usual Euclidean metric. Then as $\lambda \to \infty$, a fixed conformal change in the metric $\mathrm{d}s$ does not affect the probability that there is a black path from S_1 to S_2 by more than $o(1)$. Benjamini and Schramm also prove a corresponding statement in 3 dimensions. This statement is extremely unsurprising, at least in $\mathbb{R}^2$: because the change in the metric is conformal and, as $\lambda \to \infty$, the density of cells is very high, the graphs associated to the Voronoi tilings for the two metrics coincide almost everywhere. In other words, there are only a few 'defects', where changing the metric causes different cells to meet. In fact, in 2 dimensions, Benjamini and Schramm note that there are (in expectation) only a *bounded* number of defects. In 3 dimensions there are more, but still very few. Despite this fact, the result of Benjamini and Schramm is not at all easy: even dealing with these very few defects requires a lot of work.

The discretization problems described above would not arise in the related setting of *random discrete Voronoi percolation*, in which, instead of considering a Poisson process Z on $\mathbb{R}^2$, we take a random subset L_π of a lattice L (for example, $\mathbb{Z}^2$), where L_π is formed by selecting points of L independently with a certain probability π. Letting $\pi \to 0$ and rescaling suitably, L_π of course converges to Z in a natural sense. If we define the Voronoi tiling with respect to L_π, $\pi > 0$, and then colour the cells black with probability p independently of one another, we can show that percolation does occur for any $p > 1/2$ and $\pi > 0$, with the following proviso. In this setting, more than three Voronoi cells may meet at a vertex (a decreasing proportion of cells do so as $\pi \to 0$), and we must consider all cells meeting at a vertex to be connected. In this setting, while we still need an equivalent of Theorem 12 (which follows from the proof in Section 4), most of the remaining complications in this paper can be avoided. We shall return to this in future work [8].

2 External ingredients

The proofs presented here will be mostly self-contained; we shall make use of two results from probabilistic combinatorics and two observations concerning percolation.

Let X be a fixed ground set with n elements, and let $\mathbf{p} = (p_x)_{x \in X}$ be a vector of probabilities. Let $X_\mathbf{p}$ be a random subset of X obtained by selecting each $x \in X$ independently with probability p_x. For $\mathcal{A} \subset \mathcal{P}(X)$, let $\mathbb{P}_\mathbf{p}(\mathcal{A})$ be the probability that $X_\mathbf{p} \in \mathcal{A}$. As usual, we say that $\mathcal{A}$ is *increasing* if $A \in \mathcal{A}$ and $A \subset B \subset X$ imply $B \in \mathcal{A}$.

The first result we shall need is Harris' lemma, from his 1960 paper [18] in which he proved that $p_c \geq 1/2$ for bond percolation in $\mathbb{Z}^2$.

Lemma 4. *If $\mathcal{A}$ and $\mathcal{B}$ are increasing, then $\mathbb{P}_\mathbf{p}(\mathcal{A} \cap \mathcal{B}) \geq \mathbb{P}_\mathbf{p}(\mathcal{A})\mathbb{P}_\mathbf{p}(\mathcal{B})$.*

In other words, increasing events in the weighted cube (the product of two-element probability spaces) are positively correlated. Harris' original statement is for the case $\mathbf{p}$ constant, but the extension to general $\mathbf{p}$ is essentially equivalent. In fact, once one thinks of the statement, the proof of either form is extremely simple using induction on n. In [18], the set X was a finite set of edges of $\mathbb{Z}^2$, and $X_{\mathbf{p}}$ was the subset of X consisting of the open edges. The extension to infinite X is trivial. Harris' Lemma, which was rediscovered by Kleitman [21] in a different context, led to a series of generalizations culminating in the 'Four-functions Theorem' of Ahlswede and Daykin [3]. However, in the context of percolation, it is often exactly Harris' original lemma that is needed.

The second result we shall need is a modified form of a sharp-threshold result of Friedgut and Kalai [13], which is itself a consequence of a result of Kahn, Kalai and Linial [19] (see also [9]) concerning the influences of coordinates in a product space. Let W_{p_-,p_+} be the (weighted) three-element probability space $\{-1,0,1\}$ where the elements have respective probabilities p_-, $1-p_--p_+$ and p_+. We shall work in the nth power $W^n_{p_-,p_+}$ of this space, writing $\mathbb{P}^n_{p_-,p_+}$ for the corresponding probability measure on $\{-1,0,1\}^n$.

An event $E \subset \{-1,0,1\}^n$ is *increasing* if, whenever $x \in E$ and $x \le x'$ holds pointwise, we have $x' \in E$. This is a natural generalization of the notion of increasing for subsets of $\mathcal{P}(X)$. The event E is *symmetric* if there is a permutation group acting transitively on $[n] = \{1,2,\dots,n\}$ whose action on $\{-1,0,1\}^n$ preserves E.

Theorem 5. *There is an absolute constant c_3 such that if $0 < q_- < p_- < 1/e$, $0 < p_+ < q_+ < 1/e$, $E \subset \{-1,0,1\}^n$ is symmetric and increasing, and $\mathbb{P}^n_{p_-,p_+}(E) > \eta$, then $\mathbb{P}^n_{q_-,q_+}(E) > 1-\eta$ whenever*

$$\min\{q_+ - p_+, p_- - q_-\} \ge c_3 \log(1/\eta) p_{\max} \log(1/p_{\max}) / \log n, \tag{1}$$

where $p_{\max} = \max\{q_+, p_-\}$.

Proof. The case $p_- = q_- = 0$ of this result, a result about the weighted discrete cube, is exactly Theorem 3.2 of Friedgut and Kalai [13]; the proof in [13] extends to Theorem 5 *mutatis mutandis*. Indeed, the key bound $w(f) \le cp\log(1/p)$ in the proof of Theorem 3.1 in [13] holds in our case also, with a different constant. The corresponding formula for our three-element space is

$$w(f) \le cp_+ \log(1/p_+) + cp_- \log(1/p_-),$$

and the right hand side is at most $2cp_{\max}\log(1/p_{\max})$. □

Finally, we shall need two observations concerning k-dependent percolation in $\mathbb{Z}^2$. By a *bond (site) percolation measure on* $\mathbb{Z}^2$ we shall mean a probability measure on the space of assignments of a *state*, i.e., *open* or *closed*, to each edge (vertex) of $\mathbb{Z}^2$, with the usual σ-field of measurable events. For bond (site) percolation on $\mathbb{Z}^2$, the *open cluster C_0 containing the origin* is the set of all vertices of $\mathbb{Z}^2$ that may be reached from the origin by an *open path*, i.e., a path all of whose edges (vertices) are open. In the case of site percolation, $C_0 = \emptyset$ if the state of the origin is closed.

A bond (site) percolation measure on a graph G is *k-dependent* if for every pair S, T of sets of edges (vertices) of G at graph distance at least k, the states (being open or closed) of the edges (vertices) in S are independent of the states of the edges (vertices) in T. For bond percolation, when $k = 1$ the separation condition is exactly that no edge of S shares a vertex with an edge of T.

Liggett, Schonmann and Stacey [23] proved in a more general context that for any k and any $p_1 < 1$ there is a $p_2 < 1$ such that any k-dependent probability measure in which each edge is open with probability at least p_2 stochastically dominates the product measure in which edges are open with probability p_1. An immediate consequence of this result is the following.

Lemma 6. *There is a constant $p_0 < 1$ such that for any 1-dependent bond percolation measure on $\mathbb{Z}^2$ satisfying the additional condition that each edge is open with probability at least p_0, the probability that the origin is in an infinite open cluster is positive.*

The best value of p_0 in this lemma that is currently known is due to Balister, Bollobás and Walters [4], who showed that one may take $p_0 = 0.8639$. For us, the value of p_0 is not important. As stated, Lemma 6 is essentially trivial from first principles.

The second observation is a variant of a converse of Lemma 6, giving exponential decay rather than percolation. This time, it is easier to work with site percolation. Recall that in this context the open cluster C_0 of the origin is the set of vertices of $\mathbb{Z}^2$ joined to the origin by a path in $\mathbb{Z}^2$ every one of whose vertices is open. Although the lemma below is also essentially trivial, we give a proof.

Lemma 7. *Let k be a fixed positive integer, and let $\widetilde{\mathbb{P}}$ be a k-dependent site percolation measure on $\mathbb{Z}^2$ in which every vertex $v \in \mathbb{Z}^2$ is open with probability at most p. There is a constant $p_1 = p_1(k) > 0$ such that for every $p \le p_1$ there is a $c(p,k) > 0$ for which*

$$\widetilde{\mathbb{P}}(|C_0| \ge n) \le \exp(-c(p,k)n)$$

for all $n \ge 1$.

Proof. If $|C_0| \ge n$, then the subgraph of $\mathbb{Z}^2$ induced by the open vertices contains a tree T with n vertices, one of which is the origin. It is well known and easy to check that the number of such trees in $\mathbb{Z}^2$ grows exponentially, and is at most $(4e)^n$. Fix any such tree T. Then there is a subset S of at least $n/(2k^2 - 2k + 1)$ vertices of T such that any $a, b \in S$ are at distance at least k; indeed, one can find such a set by a greedy algorithm: whenever a vertex a is chosen, the number of other vertices it rules out is at most the number of other vertices of $\mathbb{Z}^2$ within graph distance $k-1$ of a, namely $4\binom{k}{2} = 2k^2 - 2k$. The vertices of S are open independently, so the probability that every vertex of T is open is at most $p^{|S|}$. Hence,

$$\widetilde{\mathbb{P}}(|C_0| \ge n) \le (4e)^n p^{n/(2k^2-2k+1)}.$$

Provided p is small enough that $r = 4ep^{1/(2k^2-2k+1)} < 1$, the conclusion follows, taking $c(p,k) = -\log r$. □

3 Basic preliminaries

Before we get down to the real work, let us eliminate some degenerate cases. Recall that Z is a Poisson process of intensity 1 on $\mathbb{R}^2$. It is easy to check that with probability 1 every Voronoi cell $V_Z(z)$ is bounded. The probability that Z contains 4 points lying on a circle is zero. Thus, with probability 1, no $x \in \mathbb{R}^2$ lies in more than three Voronoi cells. Finally, any bounded region contains finitely many points of Z. It follows that, with probability 1, every Voronoi cell $V_Z(z)$ is a closed convex k-gon for some k, and that when two Voronoi cells meet, they share an edge. We shall assume throughout that these non-degeneracy conditions hold *always*.

Here is a simple observation about Voronoi tilings, which we state as a lemma for ease of reference.

Lemma 8. *Suppose that there is a point $x \in \mathbb{R}^2$ such that the closest two points of Z to x are z_1 and z_2. Then the Voronoi cells of z_1 and z_2 meet.*

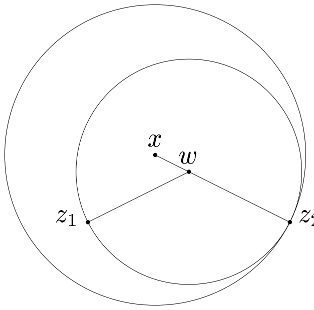

Figure 3: A point $x \in \mathbb{R}^2$ and the two closest points z_1, z_2 of Z. There are no points of Z inside the inner circle shown.

Proof. Writing $\overline{uv}$ for the Euclidean distance between two points u, $v \in \mathbb{R}^2$, we may assume that $\overline{xz_1} \leq \overline{xz_2}$. But then there is a (unique) point w on the segment xz_2 with $\overline{wz_1} = \overline{wz_2}$, and for this point w we have $\overline{wz_i} \geq \overline{wz_2}$ for every $z_i \in Z$; see Figure 3. Hence $V(z_1)$ and $V(z_2)$ do meet – they both contain w. □

Throughout this paper, there will be a parameter s controlling the *scale* of the region we are considering. We say that an event $E = E(s)$ holds *with high probability*, or **whp**, if it holds with probability tending to 1 as $s \to \infty$, with all other parameters (for example, p) fixed. All $o(.)$ notation will refer to the same limit, so an event holds **whp** if and only if it holds with probability $1 - o(1)$.

Most of our results concern the case $s \to \infty$. In the statement and proof of such results we may assume that s is larger than some fixed constant; often we shall do so without comment. Note that in a region of area s^2, we expect the largest Voronoi cell to have diameter $\Theta(\sqrt{\log s})$.

For a rectangle $R \subset \mathbb{R}^2$ and a real number $r > 0$, let us define the *r-neighbourhood* $R[r]$ of R in the usual way, i.e., as the set of all points within distance r of some point of R. Given $\rho > 1$ and a ρs by s rectangle R_s, $s > 1$, let $E_{\rm dense}(R_s)$ be the event that for every $x \in R_s[r]$ there is some point $z \in Z$ at distance $\overline{xz} < r$, where $r = 2\sqrt{\log s}$. We collect some simple properties of the event $E_{\rm dense}(R_s)$ in the following easy lemma.

Lemma 9. *Let $\rho > 1$ be constant. Let Z be a Poisson process of intensity 1 on $\mathbb{R}^2$, let $R_s \subset \mathbb{R}^2$ be a ρs by s rectangle, and set $r = 2\sqrt{\log s}$. Then $E_{\rm dense}(R_s)$ holds* **whp**. *Also, $E_{\rm dense}(R_s)$ depends only on the restriction of Z to $R_s[2r]$, and if $E_{\rm dense}(R_s)$ holds, then the colour of every point of R_s is determined by the restriction of $(Z, {\rm col})$ to $R_s[2r]$.*

Proof. We may cover $R_s[r]$ with $O(s^2/\log s) = o(s^2)$ (half-open) squares S_i of side length $r/\sqrt{2}$ so that the S_i are disjoint, and their union lies in $R_s[2r]$. Each S_i has area $r^2/2 = 2\log s$, so the number of points of Z in S_i has a Poisson distribution with mean $2\log s$. Hence, the probability that a particular S_i contains no points of Z is exactly s^{-2}, and **whp** every S_i contains at least one point of Z. It follows that $E_{\rm dense}(R_s)$ holds **whp**: given any $x \in R_s[r]$, any point $z \in Z$ in the same S_i as x satisfies $\overline{xz} < r$.

The remaining two claims are immediate from the definition of $E_{\rm dense}(R_s)$. $\square$

3.1 Increasing events are positively correlated

Recall that we are working with $\mathbb{P}_p = \mathbb{P}_p^{\mathbb{R}^2}$, the probability measure associated to the coloured Poisson process $(Z, {\rm col})$ where each point is coloured black independently with probability p. We may think of this process as assigning a state from $\{-1, 0, 1\}$ to every point of $\mathbb{R}^2$: the black points of Z have state 1, the white points of Z state -1, and every point of $\mathbb{R}^2 \setminus Z$ state 0. An event E is *black-increasing*, or simply *increasing*, if f is increasing in terms of the states of the points of $\mathbb{R}^2$. In other words, E is preserved under the addition of black points to Z and under the deletion of white points from Z. The following result follows easily from Harris' Lemma.

Lemma 10. *Let E_1 and E_2 be increasing events. Then for any $0 \le p \le 1$ we have*

$$\mathbb{P}_p(E_1 \cap E_2) \ge \mathbb{P}_p(E_1)\mathbb{P}_p(E_2).$$

We shall apply Lemma 10 to events such as 'there is a black path across a certain rectangle', to be defined below.

3.2 Crossing a rectangle: two definitions

Given a rectangle $R = [x_1, x_2] \times [y_1, y_2]$, $x_1 < x_2$, $y_1 < y_2$, let $H'(R) = H'_b(R)$ be the event that there is a path $P' = C_1C_2 \dots C_n$ of black cells such that $C_i \cap C_{i+1} \cap R$ is non-empty for $i = 1, 2, \dots, n-1$, with C_1 meeting the left-hand side of R and C_n meeting the right-hand side. We call such a path a *path of black cells across* R. In the notation, the H stands for horizontal, and the b for black.

Recall that a point x of $\mathbb{R}^2$ is black if it lies in a black cell, i.e., if (at least one of) the point(s) $z \in Z$ closest to x is black. We say that a subset $S \subset \mathbb{R}^2$ is *black* if every point of S is black.

Let $H(R) = H_b(R)$ be the event that there is a piecewise linear black path P inside R from some point on the left-hand side to some point on the right-hand side. We call such a path P a *black path across* R, and say that R may be *crossed (horizontally)* by a black path if $H(R)$ holds. Note that $H(R)$ is black-increasing: the blackness of any point of $\mathbb{R}^2$, or of any set, is black-increasing.

It is easy to see that $H(R)$ and $H'(R)$ coincide (at least up to probability zero degenerate cases, which we have ruled out). While the the definition given for $H'(R)$ is perhaps easier to visualize, that given for $H(R)$ is easier to work with, and this is what we shall use most of the time. We shall write $V(R) = V_b(R)$ for the event that there is a *black path crossing* R *vertically*, defined analogously. Also, we shall write $H_w(R)$ and $V_w(R)$ for the corresponding events with black replaced by white.

The basic starting point for the study of critical points for random Voronoi percolation in $\mathbb{R}^2$ is the following well-known fact.

Lemma 11. *If $p = 1/2$ and S is any square $[x, x + a] \times [y, y + a]$, $a > 0$, then the probability of the event $H(S)$ that there is a black path across S is* exactly $1/2$.

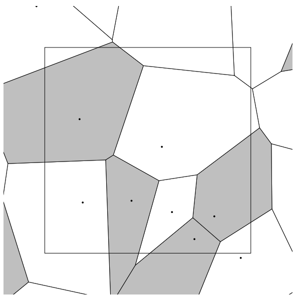

Figure 4: An example in which $H_b(S)$ does not hold and $V_w(S)$ does.

This result is the exact analogue of the well-known result for $\mathbb{Z}^2$ based on duality. It is an immediate consequence of the fact that, with probability 1, exactly one of the events $H(S) = H_b(S)$ and $V_w(S)$ holds; see Figure 4. By symmetry, these events have the same probability. Random Voronoi percolation differs from bond percolation in $\mathbb{Z}^2$ in the following important respect: events such as $H(R)$ for different rectangles R are *not* independent, as $H(R)$ depends on the points of Z outside but near R as well as those inside.

4 Crossing a rectangle at $p = 1/2$

One of our main intermediate results is Theorem 12 below, which will play a role corresponding to the Russo-Seymour-Welsh (RSW) Theorem for bond percolation in $\mathbb{Z}^2$. The RSW Theorem, proved independently by Russo [24] and Seymour and Welsh [25] (see also [14]), states that, if the probability of a horizontal open crossing of an s by s square is at least $c > 0$, then for $\rho > 1$ the probability of a horizontal (i.e., long) open crossing of a ρs by s rectangle is at least $f_\rho(c) > 0$. The key point is that $f_\rho(c)$ (which is, in fact, a polynomial in c) does not depend on s. Together with an analogue of Lemma 11, the RSW Theorem shows that in $p = 1/2$ bond percolation on $\mathbb{Z}^2$ the probability of a horizontal open crossing of a $10s$ by s rectangle, say, is bounded away from zero as $s \to \infty$. It is very easy to deduce that percolation does not occur at $p = 1/2$.

The result we shall give does not directly correspond to the RSW Theorem: starting from the fact that a square of any given size can be crossed by a black path with probability bounded away from zero (namely, probability 1/2 if $p = 1/2$), we deduce that *certain* (not necessarily *all*) very large ρs by s rectangles can be crossed with probability bounded away from zero. The proof is much longer than that of the RSW Theorem: none of the existing proofs of the latter seems to transfer to this context. This observation was made by Zvavitch [30], who pointed out that the simplest proof of the RSW Theorem, based on finding a 'lowest' black left-right crossing of a square, and noting that this is independent of edges above the crossing, breaks down for the Voronoi tiling: the fact that various Voronoi cells form a path depends on the *absence* of points of Z even above the path. Zvavitch notes that the proof of Alexander's result [2] for a different Poisson percolation model does not carry over to the Voronoi setting either. In proving that $p_c \geq 1/2$, Zvavitch uses a totally different method, not going via the RSW Theorem, but rather adapting a proof of Harris' Theorem due to Burton and Keane [11] and Zhang (see [14, pages 198 and 289]), which does carry over to the Voronoi case in a simple way.

In any case, the question of whether a direct analogue of the RSW Theorem holds for random Voronoi percolation is open. The following weaker result will suffice for our purposes. In the results below and for the rest of the paper, we shall make use of the function

$$f_p(\rho, s) = \mathbb{P}_p(H([0, \rho s] \times [0, s])).$$

In other words, $f_p(\rho, s)$ is the $\mathbb{P}_p$-probability that a ρs by s rectangle has a horizontal black crossing.

Theorem 12. *Let* $0 < p < 1$ *and* $\rho > 1$ *be fixed. If* $\liminf_{s\to\infty} f_p(1, s) > 0$, *then* $\limsup_{s\to\infty} f_p(\rho, s) > 0$.

Before we prove this result, let us note that it has the following immediate consequence.

Corollary 13. *Let $\rho > 1$ be fixed. There is a constant $c_0 = c_0(\rho) > 0$ such that for every s_0 there is an $s > s_0$ with $\mathbb{P}_{1/2}(H(R_s)) \geq c_0$, where R_s is a ρs by s rectangle.*

Proof. Lemma 11 states exactly that $f_{1/2}(1, s) = 1/2$ for all $s > 0$. From Theorem 12 it follows that $\limsup f_{1/2}(\rho, s) > 0$, which is precisely the conclusion of the corollary. □

In principle, we could give an explicit value for $c_0(\rho)$, but it would be rather small for the values of ρ we shall consider, for example, $\rho = 3$. It is likely that $f_{1/2}(\rho, s)$ tends to a limit as $s \to \infty$ with ρ fixed, but as far as we are aware this has not been proved. Indeed, this statement would be a weak form of a very special case of the conformal invariance conjecture of [22] mentioned in the introduction. However, even for bond percolation in $\mathbb{Z}^2$, which has been studied much more extensively than random Voronoi percolation, the corresponding question is still open.

Proof of Theorem 12. Assume, for a contradiction, that Theorem 12 does not hold, and fix a particular value of p for which the result fails. Then there is a constant $c_1 > 0$ such that

$$f_p(1, s) \geq c_1 \tag{2}$$

for all large enough s, but, for some fixed $\rho > 1$,

$$f_p(\rho, s) \to 0. \tag{3}$$

Here, as usual, the limit is taken as $s \to \infty$ with all other parameters fixed. We shall make several claims during the proof; all of these will be conditional on our assumptions, which we shall show to be inconsistent.

Note that $f_p(\rho, s)$ is decreasing in ρ, since the corresponding events for different ρ are nested. We claim that for any $a, b \geq 1$ we have

$$f_p((a + b - 1), s) \geq f_p(a, s) f_p(b, s) f_p(1, s). \tag{4}$$

To see this, consider the events E_1, E_2 that there are (piecewise linear) black paths P_1, P_2 across $[0, as] \times [0, s]$ and $[(a-1)s, (a+b-1)s] \times [0, s]$, respectively, and the event E_3 that there is a piecewise linear black path P_3 crossing the square $S = [(a - 1)s, as] \times [0, s]$ from top to bottom. The E_i are increasing events so, by Lemma 10,

$$\mathbb{P}_p(E_1 \cap E_2 \cap E_3) \geq \mathbb{P}_p(E_1)\mathbb{P}_p(E_2)\mathbb{P}_p(E_3) = f_p(a, s) f_p(b, s)\mathbb{P}_p(E_3).$$

But, of course, $\mathbb{P}_p(E_3) = \mathbb{P}_p(V(S)) = \mathbb{P}_p(H(S)) = f_p(1, s)$, using rotational symmetry of the model. Finally, if the E_i all hold then P_3 meets both P_1 and P_2, so the union of the P_i contains a path across $[0, (a + b - 1)s] \times [0, s]$ which is black. This proves (4).

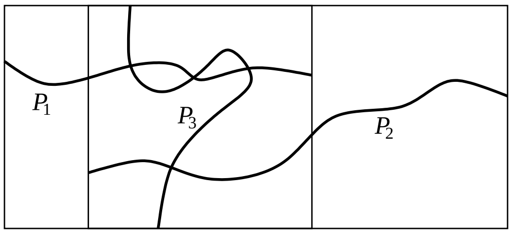

Combining this with our assumptions (2) and (3), for any $\varepsilon > 0$ we have

$$f_p(1+\varepsilon, s) \to 0 \tag{5}$$

as $s \to \infty$; otherwise, as $f_p(1, s)$ is bounded away from zero by (2), using (4) $\lceil(\rho - 1)/\varepsilon\rceil$ times we have $f_p(\rho', s) \not\to 0$ for some $\rho' \geq \rho$, contradicting (3).

Using (5) we shall be able to deduce increasingly implausible properties of black paths crossing certain rectangles. Throughout the argument below we work within the strip $T_s = [0, s] \times \mathbb{R}$. The first step is to show that if we try to cross the infinite strip T_s, then we almost always stay nearly within height $s/2$ of the point we start from.

Claim 12.1. *Let $\varepsilon > 0$ be fixed, and let L be the line-segment $\{0\} \times [-\varepsilon s, \varepsilon s]$. Assuming that (2) and (3) hold, the probability that there is a black path P in T_s starting from L and going outside $S' = [0, s] \times [-(1/2 + 2\varepsilon)s, (1/2 + 2\varepsilon)s]$ tends to zero as $s \to \infty$.*

Proof. Considering the point at which the path P first leaves S', by symmetry it suffices to show that the event E that there is a black path P_1 lying entirely within S' and connecting some point of the segment L to some point on the top side of S' has probability tending to zero.

Let E_1 be the event that there is such a path P_1 lying within the rectangle $R = [0, s] \times [-s/2, s/2 + 2\varepsilon s]$. Note that if E_1 holds and E does not, then there is some black path within T_s joining a point at height $-s/2$ to a point at height $s/2 + 2\varepsilon s$. The shortest such path is a black path crossing a rectangle of width s and height $(1 + 2\varepsilon)s$ from top to bottom. Since such a path exists with probability $f_p((1 + 2\varepsilon), s)$, which tends to zero by (5), we have $\mathbb{P}_p(E \setminus E_1) \leq f_p((1 + 2\varepsilon), s) \to 0$. Consequently, it suffices to show that $\mathbb{P}_p(E_1) \to 0$.

Reflecting vertically in the line $y = \varepsilon s$, a symmetry axis of R, let $L' = \{0\} \times [\varepsilon s, 3\varepsilon s]$ be the image of the segment L, and let E_2 be the event that there is a black path P_2 within R from L' to some point with height $-s/2$; see Figure 5. Then $\mathbb{P}_p(E_2) = \mathbb{P}_p(E_1)$ by symmetry. Now E_1 and E_2 are increasing events, so both hold with probability at least $\mathbb{P}_p(E_1)^2$. But if both hold then the corresponding paths P_1 and P_2 cross, as L lies entirely below L'. Thus $P_1 \cup P_2$ contains a black path crossing R from top to bottom. The probability that such a path exists is $f_p((1 + 2\varepsilon), s)$, which tends to zero by (5). Consequently $\mathbb{P}_p(E_1) \to 0$ and hence $\mathbb{P}_p(E) \to 0$, proving the claim. $\square$

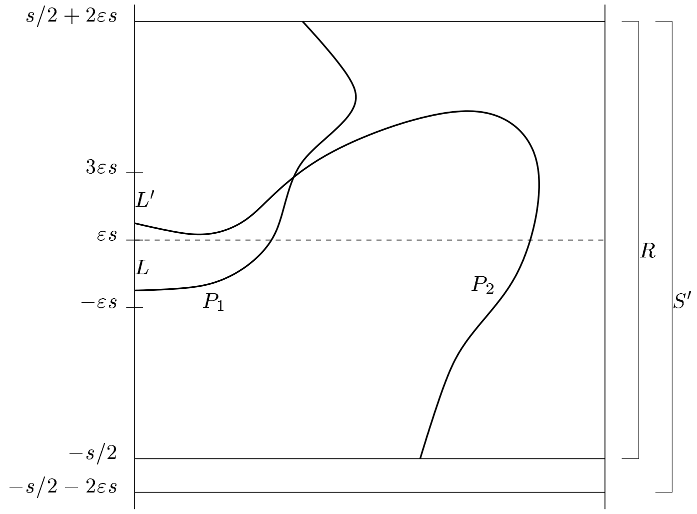

Figure 5: Paths P_1 and P_2 guaranteeing the events E_1 and E_2

Given a path $P \subset \mathbb{R}^2$, let $y_+(P)$ and $y_-(P)$ be the supremum and infimum of the y-coordinates of points on P. In other words, (as P is closed) y_+ and y_- are the maximum and minimum 'heights' attained by P. Recall that an event holds *with high probability*, or **whp**, if it holds with probability $1-o(1)$ as $s \to \infty$ with any other parameters (e.g., p, ε) fixed. Recall also that when we say that a path P crosses a rectangle R horizontally, we mean that P lies entirely within R, starts on the left-hand side of R, and ends on the right-hand side of R. We shall write $y_0(P)$ and $y_1(P)$ for the y-coordinates of the left and right endpoints of such a path P.

Claim 12.2. *Let $\varepsilon > 0$ and $C > 0$ be fixed, and let R_s be an s by $2Cs$ rectangle. Assuming that* (2) *and* (3) *hold,* **whp** *every black path P crossing R_s horizontally satisfies*

$$\big||y_\pm(P) - y_i(P)| - s/2\big| \leq \varepsilon s, \tag{6}$$

for $y_\pm(P) = y_+(P), y_-(P)$ and $i = 0, 1$. In particular,

$$|y_0(P) - y_1(P)| \leq 2\varepsilon s. \tag{7}$$

In other words, the maximum and minimum heights attained by P are almost exactly $s/2$ larger and smaller respectively than its starting and ending heights.

Proof. Without loss of generality we may take $R_s = [0, s] \times [-Cs, Cs]$. It suffices to prove that **whp** any black path P crossing R_s horizontally satisfies

$$y_0 + s/2 - \varepsilon s \leq y_+ \leq y_0 + s/2 + \varepsilon s. \tag{8}$$

Here, and in what follows, we suppress the dependence of $y_+(P)$ on P when the path being considered is clear from context, and use other self-explanatory abbreviations. Inequality (8) gives one of the four cases included in (6). The others follow by reflecting horizontally and vertically. The bound (7) then follows by applying (6) twice, with $y_\pm(P) = y_+(P)$, say, and $y = 0, 1$.

Let us cover the left-hand side of R_s by $\lceil 2C/(\varepsilon/2) \rceil = O(1)$ line-segments L_i of length $\varepsilon s/2$. Fixing i, let us call a black path P *eligible* if P crosses R_s horizontally, starting from a point $(0, y_0) \in L_i$. As any path P crossing R_s horizontally must contain a point $(0, y_0) \in L_j$ for some j, it suffices to show that **whp** (8) holds for every eligible path P.

Let $(0, y)$ be the midpoint of L. Applying Claim 12.1 with $\varepsilon/4$ in place of ε, and using translational invariance of the model, **whp** every eligible path P satisfies

$$y_+ \le y + s/2 + \varepsilon s/2 \tag{9}$$

and

$$y_- \ge y - s/2 - \varepsilon s/2. \tag{10}$$

As $|y_0 - y| \le \varepsilon s/4$ from the definition of eligibility, (9) implies the upper bound in (8). It remains to prove the lower bound. Suppose that it is not true that the lower bound in (8) holds **whp**. Then with probability bounded away from 0 there is an eligible path with

$$y_+ \le y_0 + s/2 - \varepsilon s \le y + s/2 - 3\varepsilon s/4. \tag{11}$$

As **whp** every eligible path satisfies (10), we see that with probability bounded away from zero there is some black path P crossing R_s horizontally for which both (10) and (11) hold. As such a path crosses a fixed rectangle of width s and height $(1 - \varepsilon/4)s$ horizontally, this conclusion contradicts (5), completing the proof of (8) and hence of the claim. □

The conclusions of Claims 12.1 and 12.2 may seem fairly plausible. Recall, however, that we are aiming for a contradiction. In the next two claims, we show that our assumptions imply properties of paths crossing certain rectangles that are far from plausible.

Claim 12.3. *Let $C > 0$ be fixed, and let $R = R_s$ be the s by $2Cs$ rectangle $[0, s] \times [-Cs, Cs]$. For $i = 0, 1$, set $R_i = [is/100, (i + 99)s/100] \times [-Cs, Cs]$. Assuming that* (2) *and* (3) *hold,* **whp** *every black path P crossing R contains disjoint black paths P_0 and P_1 such that P_i crosses R_i horizontally.*

Proof. Let E be the event that every black path P crossing R horizontally satisfies

$$\big| |y_\pm(P) - y_j(P)| - s/2 \big| \le s/1000 \tag{12}$$

for $j = 0, 1$. Note that $\mathbb{P}_p(E) = 1 - o(1)$ by Claim 12.2, applied with $\varepsilon = 1/1000$ to R. Similarly, for $i = 0, 1$, let E_i be the event that every black path P_i crossing R_i horizontally satisfies

$$\big| |y_\pm(P_i) - y_j(P_i)| - 0.495s \big| \le s/1000 \tag{13}$$

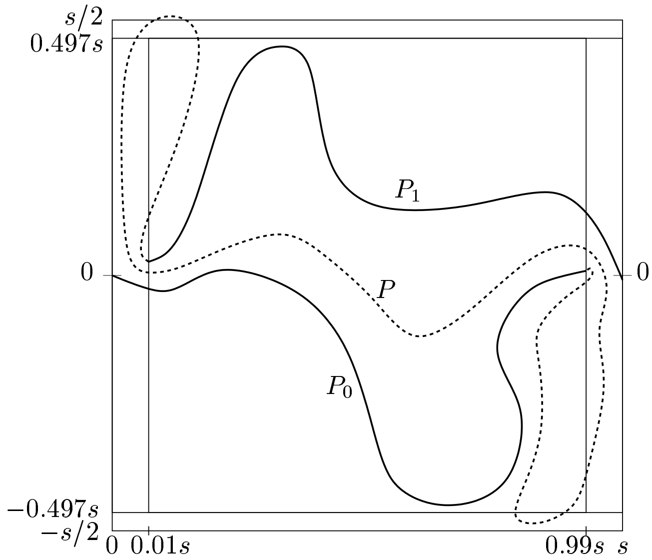

Figure 6: Two sub-paths P_0, P_1 (solid lines) of a path P crossing horizontally from $x = 0$ to $x = s$. The numbers outside the rectangle are x- or y-coordinates; to avoid clutter, y-coordinates are shown relative to $y_0(P)$.

for $j = 0, 1$. By Claim 12.2 applied to R_i with $0.99s$ in place of s, $C/0.99$ in place of C, and $\varepsilon = 1/1000$, each event E_i holds **whp**.

We shall assume, as we may, that $E \cap E_0 \cap E_1$ holds. Let P be any black path crossing R horizontally, starting at $(0, y_0)$ and ending at (s, y_1), say. Consider the first time that P crosses the line $x = 0.99s$. Let P_0 be the initial segment of P from $(0, y_0)$ up to this point. Similarly, let P_1 be the final segment of P, defined by going backwards along P from (s, y_1) until the line $x = 0.01s$ is first reached; see Figure 6. Then each P_i crosses R_i horizontally. Applying (12) twice, as in the proof of (7), we have $|y_0 - y_1| \leq 2s/1000$. Hence, from (13), each of P_0 and P_1 remains between the heights $y_0 - 0.497s$ and $y_0 + 0.497s$. It follows that P_0 and P_1 are disjoint: if they meet, their union contains a path P' crossing R from left to right with $y_+(P') - y_-(P') \leq 0.994s$. Such a path P' cannot satisfy (12), contradicting our assumption that E holds. $\square$

Note that the properties of the path P and its sub-paths P_0, P_1 given by (12) and (13) are even less plausible than Figure 6 might suggest: P_0 and P_1 both have maximum and minimum heights very close to $y_0 \pm 0.495s$. Taking P as a shortest path, the rest of P must somehow thread its way between P_0 and P_1 to join one to the other, without crossing either.

Claim 12.4. *Let $C > 0$ be fixed, and let $R = R_s$ be the s by $2Cs$ rectangle $[0, s] \times [-Cs, Cs]$. For $0 \leq j \leq 4$, set $R_j = [js/100, (j + 96)s/100] \times [-Cs, Cs]$. Assuming that* (2) *and* (3) *hold,* **whp** *every black path P crossing R horizontally* *contains 16 disjoint black paths P_i, $1 \le i \le 16$, where each P_i crosses some R_j horizontally.*

Proof. Arguing as in the proof of Claim 12.3, **whp** any path P_0 obtained as above has disjoint initial and final segments P_{00} and P_{01} crossing rectangles of width $0.98s$, and so on. Alternatively, we could apply a slightly modified Claim 12.3 (with 0.99 replaced by 0.98/0.99, etc.) multiple times. $\square$

At this point it seems that it should be easy to reach a contradiction and hence complete the proof of Theorem 12. Note that all paths Q we consider are piecewise linear. Hence the geometric length $|Q|$ is well-defined. Let $R = R_s = [0, s] \times [-s, s]$. We have shown that **whp** any black path P crossing R horizontally contains 16 paths P_i crossing strips of width $0.96s$, with $|P| \ge \sum_{i=1}^{16} |P_i|$.

Digressing from the formal proof for a moment, suppose for simplicity that $p = 1/2$ (which is the important case). Under our assumption that the conclusion of Theorem 12 does not hold, from (5) we have $f_{1/2}(1+\varepsilon, s) \to 0$ for any $\varepsilon > 0$. Recalling that $V_w(R)$ is the event that there is a vertical crossing of R by a white path, it follows that $\mathbb{P}_{1/2}(V_w(R)) = o(1)$, so $\mathbb{P}_{1/2}(H(R)) = 1 - o(1)$. Hence, **whp** R has a horizontal black crossing; let P be a shortest such crossing. Considering the sub-paths P_i of P given by Claim 12.4, one might hope that the expected length of P is at least $16 + o(1)$ times the expected length of the shortest crossing of a slightly smaller rectangle, which would lead to a contradiction. Unfortunately, we cannot work with expected lengths, as we have no control of the path lengths in the $o(1)$ probability exceptional cases. In general, there is a problem with these cases, as their probabilities tend to accumulate. To deal with this, we shall subdivide the paths P_i once more, this time in such a way as to introduce independence.

Returning to the formal proof of Theorem 12, given any rectangle R, let $L(R)$ be the minimum length of a (piecewise-linear) black path P crossing R horizontally. We take $L(R) = \infty$ if there is no such path. Of course, we could define $L(R)$ as an infimum, but it is easy to see that this infimum is attained.

Let R_s be an s by $2s$ rectangle, where $s > 1$. Setting $r = 2\sqrt{\log s}$, let $E_{\mathrm{dense}}(R_s)$ be the event considered in Lemma 9. We define a modification $\widetilde{L}(R_s)$ of the random variable $L(R_s)$ as follows: if $E_{\mathrm{dense}}(R_s)$ does not hold, set $\widetilde{L}(R_s) = 0$; otherwise, set $\widetilde{L}(R_s) = L(R_s)$. By Lemma 9, $E_{\mathrm{dense}}(R_s)$ holds **whp**, so **whp**

$$\widetilde{L}(R_s) = L(R_s). \tag{14}$$

As any path across R_s has length at least s, it follows that the inequality

$$\widetilde{L}(R_s) \ge s \tag{15}$$

holds **whp**. As R_s has width s and height $2s > s$, with probability at least $f_p(1, s)$ there is some black path P crossing R horizontally. From (2) there is a constant $c_1 > 0$ such that $f_p(1, s) \ge c_1$ for all (large enough) s. Using (14), it follows that if s is large enough, then

$$\mathbb{P}_p(\widetilde{L}(R_s) < \infty) \geq c_1/2 \tag{16}$$

holds for any s by $2s$ rectangle R_s.

Let $0 < \eta < \min\{c_1/3, 10^{-4}\}$ be fixed, and define a function $t(s)$ by

$$t(s) = \sup\{x : \mathbb{P}_p(\widetilde{L}(R_s) < x) \leq \eta\},$$

where R_s is any s by $2s$ rectangle. This makes sense as the distribution of $\widetilde{L}(R_s)$ depends on s only and not on the location of the rectangle. From (16), if s is large enough, which we shall assume from now on, then $t(s) < \infty$, as $c_1/2 > \eta$. It follows that the supremum in the definition of $t(s)$ is attained, so

$$\mathbb{P}_p(\widetilde{L}(R_s) < t(s)) \leq \eta. \tag{17}$$

The key property of $\widetilde{L}(R_s)$ is described in the following claim.

Claim 12.5. *Let R_1 and R_2 be two s by $2s$ rectangles separated by a distance of at least $s/100$. If s is large enough, then the random variables $\widetilde{L}(R_1)$ and $\widetilde{L}(R_2)$ are independent.*

Proof. Let us write $r = r(s)$ for $2\sqrt{\log s}$ as before. We take s large enough that $4r(s) \leq s/100$. By Lemma 9, for $R = R_1$ or R_2, whether $E_{\text{dense}}(R)$ holds depends only on the intersection of the Poisson process Z with the $2r$-neighbourhood $R[2r]$ of R. Also, given that $E_{\text{dense}}(R)$ does not hold, the colour of every point of R is determined by $Z \cap R[2r]$ and the colours of these points. Hence $\widetilde{L}(R)$ depends only on $Z \cap R[2r]$ and the black/white colouring of these points. As $R_1[2r]$ and $R_2[2r]$ are disjoint, the claim follows. $\square$

After all this work the proof of Theorem 12 is not far away. However, in order to complete it, we need one more technical assertion.

Claim 12.6. *Let R_s be a fixed $0.96s$ by $2s$ rectangle. Assuming that (2) and (3) hold, if s is large enough then*

$$\mathbb{P}_p\left(L(R_s) < t(0.47s)\right) \leq 200\eta^2.$$

In other words, with probability at least $1 - 200\eta^2$, every black path P crossing R horizontally has length $|P| \geq t(0.47s)$.

Proof. We take $R_s = [0, 0.96s] \times [-s, s]$. Let us cover the left-hand side of R_s by 100 line-segments L_i of length 0.02. Fixing i, let us say that a black path P crossing R horizontally is *eligible* if P starts from a point of L_i. Let B $(= B_i)$ be the event that there is an eligible path P with $|P| < t(0.47s)$. As every black P crossing R starts from a point of some L_j, it suffices to prove that $\mathbb{P}_p(B) \leq 2\eta^2$.

Let $(0, y)$ be the midpoint of L_i. Consider an eligible path P, and let $(0, y_0)$, $(0.96s, y_1)$ be its endpoints. Let $P_0 = P_0(P)$ be the initial segment of P from $(0, y_0)$, stopping the first time P reaches the line $x = 0.47$, a little less than half way across R_s. Similarly, let $P_1 = P_1(P)$ be the final segment defined by going backwards along P from $(0.96s, y_1)$ to the line $x = 0.49s$. Let $R_0 = [0, 0.47s] \times [y - 0.47s, y + 0.47s]$ and $R_1 = [0.49s, 0.96s] \times [y - 0.47s, y + 0.47s]$ be rectangles of width $0.47s$ and height twice this; see Figure 7.

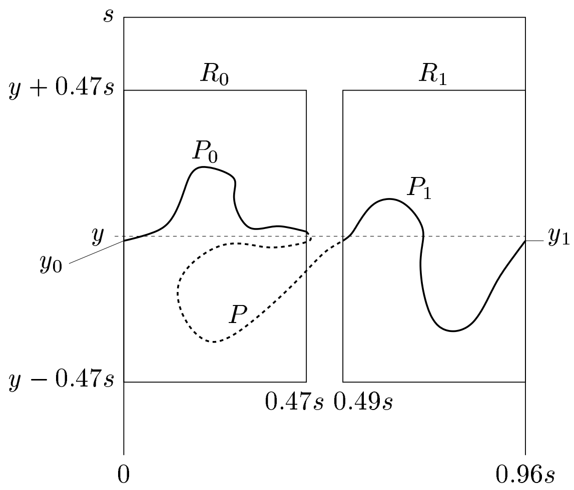

Figure 7: The upper portion of R_s, showing a path P and its initial and final segments P_i (drawn solid).

From our previous claims, **whp** every eligible path P has the following properties. Firstly, $|y_0 - y| \le 0.01s$ (by definition of eligibility). Secondly, $|y_1 - y_0| \le 0.01s$, by Claim 12.2 applied to the rectangle R_s, with $0.96s$ in place of s. Thirdly, P_0 is contained in R_0, by Claim 12.2 applied to the rectangle $[0, 0.47s] \times [-s, s]$ with $0.47s$ in place of s, using $|y_0 - y| \le 0.01s$. Fourthly, P_1 is contained in R_1, using $|y_1 - y| \le 0.02s$ and Claim 12.2.

It follows that **whp** every eligible path P has length at least $L(R_0) + L(R_1)$. As the R_i are $s' = 0.47s$ by $2s'$ rectangles, using (14) it follows that **whp** every eligible path satisfies $|P| \ge \widetilde{L}(R_0) + \widetilde{L}(R_1)$. But $\widetilde{L}(R_0)$ and $\widetilde{L}(R_1)$ are independent by Claim 12.5, applied with $0.47s$ in place of s. Hence, using (17),

$$\mathbb{P}_p\big(\widetilde{L}(R_0) + \widetilde{L}(R_1) < t(0.47s)\big) \le \eta^2.$$

Thus, the probability that some eligible path P has $|P| < t(0.47s)$ is at most $\eta^2 + o(1) \le 2\eta^2$. □

We can now complete the proof of Theorem 12. Indeed, the results above imply that, if s is large enough, then

$$t(s) \ge 16t(0.47s). \tag{18}$$

To see this, consider the rectangle $R = R_s = [0, s] \times [-s, s]$. Consider the 5 sub-rectangles $R_j = [js/100, (j + 96)s/100] \times [-s, s]$, $0 \le j \le 4$, of R_s. Note that each R_j has width $0.96s$ and height $2s$. By Claim 12.6, with probability at least $1 - 1000\eta^2$ we have $L(R_j) \geq t(0.47s)$ for each j. But by Claim 12.4, **whp** every black path P crossing R horizontally has length at least $16 \min_j L(R_j)$. Hence, with probability at least $1 - 1001\eta^2$ we have $L(R) \geq 16t(0.47s)$. From (14) we have $\widetilde{L}(R) = L(R)$ **whp**, so it follows that if s is large enough then

$$\mathbb{P}_p(\widetilde{L}(R) < 16t(0.47s)) \leq 1002\eta^2 \leq \eta.$$

But then $t(s) \geq 16t(0.47s)$ by the definition of $t(s)$, proving (18).

Finally, taking logarithms, we have $\log t(s) \geq \log(16) + \log t(0.47s)$ for all large enough s. Combined with (15), which gives $t(s) \geq s$ for large enough s, this leads to the absurd conclusion that $t(s)$ grows at least as fast as a constant times $s^{\log 16/\log(1/0.47)}$, and is in particular at least s^3 for s large enough. This conclusion is absurd because, considering an s by $2s$ rectangle R_s, from (14), (16) and (17), for large s we have

$$\mathbb{P}_p(t(s) \leq L(R_s) < \infty) \geq c_1/2 - \eta - o(1) \geq c_1/7 > 0.$$

Hence, with probability bounded away from zero, there is a black path P crossing R_s horizontally, but the shortest such P has length at least $t(s) \geq s^3$. As the rectangle R_s has area $2s^2$, **whp** there are at most $100s^2$ Voronoi cells meeting R_s, and **whp** each has diameter at most $\log s$; any shortest crossing P meets each cell at most once, and then in a line segment, and so has length at most $100(\log s)s^2$, which is smaller than s^3 if s is large enough.

This contradiction completes the proof of Theorem 12. □

Let us note for possible future reference that, while our proof of Theorem 12 is rather indirect, it is also rather general. The only properties of the measure $\mathbb{P}_p$ and notion of black crossing needed for the proof are the following. Firstly, crossings can be defined by the 'blackness' of geometric paths (so, for example, horizontal and vertical crossings of the same rectangle meet, enabling the combination of certain crossings to form longer crossings). Secondly, the existence of certain crossings must be 'increasing' events, so that positive correlation of such events holds. Thirdly, the crossing probabilities must be invariant under certain symmetries of $\mathbb{R}^2$. Note that invariance under the symmetries of $\mathbb{Z}^2$ certainly suffices, as we need only consider rectangles with integer coordinates. Fourthly, some kind of asymptotic independence is needed: certainly independence of regions separated by more than a certain distance (which does not hold for random Voronoi percolation) is more than enough.

From the remarks above, the proof of Theorem 12 applies, for example, to bond or site percolation in the square lattice; of course, the existing proofs of the stronger RSW Theorem are much simpler in these contexts. Our proof also applies to the Poisson model considered by Alexander [2], whose stronger result for this simpler model has a rather long proof, and to the random discrete Voronoi setting considered in the introduction.

In what follows, all we shall use from this section is Corollary 13. The claims used in the proof, which in any case assumed a false assumption, will not be needed later.

5 The torus

From now on we shall work mostly in the torus $\mathbb{T}(s)$, $s > 0$, obtained by identifying opposite sides of the square $[0, s]^2 \subset \mathbb{R}^2$. To reduce unnecessary clutter, most of the time we shall suppress the dependence on s. Throughout this paper, we are interested only in the $s \to \infty$ limit: in the end, we shall show that for any s_0, certain events hold with certain probabilities in some $\mathbb{T}(s)$, $s \geq s_0$. Thus we may assume that s is larger than any given constant; often we shall do so without comment, assuming inequalities such as $\log s \geq 2\sqrt{\log s}$ without comment.

We consider the (metric induced by) the Euclidean metric on $\mathbb{T} = \mathbb{T}(s)$. We may define a Poisson process $Z = Z^{\mathbb{T}}$ on $\mathbb{T}$ with intensity 1, and the corresponding random Voronoi tiling, which we may colour as before: each point z of $Z^{\mathbb{T}}$ is black with probability p, independently of all other points, and white otherwise. Every point x in the Voronoi cell associated to z inherits the colour of z. (So some boundary points are both black and white.) We write $\mathbb{P}_p^{\mathbb{T}} = \mathbb{P}_p^{\mathbb{T}(s)}$ for the corresponding probability measure.

Of course, we may couple $Z^{\mathbb{T}}$ and $Z^{\mathbb{R}^2}$ so that they agree on $S = [0, s]^2 \subset \mathbb{R}^2$. We shall see below that **whp** every disc of radius $\sqrt{\log s}$ contained in S contains some point of $Z^{\mathbb{T}}$. It follows that **whp** the Voronoi cells associated with $Z^{\mathbb{T}}$ meeting $S' = [\log s, s - \log s]^2$ and the cells associated with $Z^{\mathbb{R}^2}$ meeting S' coincide. In particular, we shall be able to apply Corollary 13 to a rectangle in $\mathbb{T}$ as long as the longer side has length at most $s - \log s$.

We note two basic facts about the maximum and minimum number of points in discs of certain sizes in the torus. From now on we write Z for $Z^{\mathbb{T}}$.

Lemma 14. *With probability $1 - o(s^{-10})$ every disc of radius $\sqrt{\log s}$ in $\mathbb{T}(s)$ contains at least one point of Z. Furthermore, for every constant $a > 0$ there is a $\rho = \rho(a) > 0$, and for every constant $\rho > 0$ there is an $a = a(\rho) > 0$, such that with probability $1 - o(s^{-10})$ no disc of radius $\rho\sqrt{\log s}$ in $\mathbb{T}(s)$ contains more than $a \log s$ points of Z.*

We omit the standard proof based on the first-moment method.

6 Approximation preliminaries

We are going to approximate the continuous Poisson process on $\mathbb{T}$ using a fine grid. There is a problem with the order of limits, as our application of Theorem 5 will restrict how fine the grid can be as a function of s, the scale of the torus. This gives rise to 'defects' or potential ambiguities, i.e., places where the discrete approximation of the location of the points does not tell us which Voronoi cells actually meet. If we could take a fine enough grid (in terms of s), we could say that with high probability there are no defects in $\mathbb{T}(s)$; unfortunately we cannot do this. While the density of defects will be low, some will occur, and we must argue that they do not make much difference. More precisely, we shall compare their effect to the effect of changing p slightly, and show that the latter dominates in some precise sense. This is not very surprising, but seems to be a little fiddly to prove rigorously.

An analogous difficulty was encountered by Benjamini and Schramm in [6], and also required considerable work to overcome. Essentially, the result of [6] is that at any p, a very small density of 'defects' (a bounded number in two dimensions) has negligible effect on the probability that a large region may be crossed. (See the introduction for further details.) Here, the density of defects is much higher (an arbitrarily small negative power of the area of the region). On the other hand, the fact that we can vary p gives us much more elbow room.

Let us say that a point $x \in \mathbb{T}$ is *δ-robustly black* if the closest black point of Z is at least a distance δ closer to x than the closest white point. A piecewise linear path P in $\mathbb{T}$ is *δ-robustly black* if every point $x \in P$ is. When we come to make our discrete approximation, δ-robustly black paths will be very useful; such a path clearly remains black if we adjust the position of each point of Z by a small amount (less than $\delta/2$). In the following result, for $p_1 < p_2$ we couple the two probability measures (on black/white-coloured Poisson processes) $\mathbb{P}_{p_1}^{\mathbb{T}(s)}$ and $\mathbb{P}_{p_2}^{\mathbb{T}(s)}$. In other words, we construct simultaneously Poisson processes Z_1 and Z_2 of intensity 1 on $\mathbb{T}(s)$, with associated colourings col_1 and col_2, so that given Z_i each point of Z_i is coloured black with probability p_i independently of all other points of Z_i, and white otherwise.

Theorem 15. *Let $0 < p_1 < p_2 < 1$ and $\varepsilon > 0$ be given. Let $\delta = \delta(s)$ be any function with $\delta(s) \leq s^{-\varepsilon}$. We may construct in the same probability space coloured Poisson processes (Z_1, col_1) and (Z_2, col_2) on $\mathbb{T}(s)$ as described above, in such a way that the following global event holds* **whp** *as $s \to \infty$: for every piecewise linear path P_1 which is black with respect to (Z_1, col_1) there is a piecewise linear path P_2 which is 4δ-robustly black with respect to (Z_2, col_2), such that every point of P_2 is within distance $(\log s)^2$ of some point of P_1 and vice-versa.*

The only aim of this section is to prove the result above. This will involve overcoming several bothersome technicalities. Nothing used in the (rather long) proof of Theorem 15 will be used in other parts of the paper.

6.1 Good points are useful

Let us say that a quadruple of points $\{z_1, z_2, z_3, z_4\}$ of the Poisson process Z on $\mathbb{T} = \mathbb{T}(s)$ is *δ-bad* if there is a point x of $\mathbb{T}$ and a distance $r < 2\sqrt{\log s}$ such that (a) each z_i is within distance $r + \delta$ of x, and (b) no point of Z is within distance r of x; see Figure 8. In other words, the four points are close to a circle which is empty of points of Z. Let us say that a pair of points z_1, z_2 of Z is *δ-close* if they are within distance δ. Finally, a single point $z \in Z$ is *δ-bad* if it is in a δ-bad quadruple or a δ-close pair. Note that if $\delta < \delta'$, then every δ-bad point is also δ'-bad. A point $z \in Z$ is *δ-good* if it is not δ-bad. The reason for considering δ-good points is that paths of cells associated to δ-good points give rise to piecewise linear paths that are robust. We write $\overline{xy}$ for the Euclidean distance between two points x and y of $\mathbb{T}$.

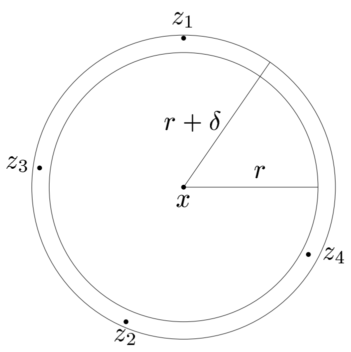

Figure 8: A δ-bad quadruple $\{z_1, z_2, z_3, z_4\} \subset Z$. There are no points of Z in the inner circle.

Lemma 16. *Let $s \geq e^{100}$ and $\delta < 1/\log s$ be given, and suppose that no disc of radius $2\sqrt{\log s}$ in $\mathbb{T}$ is free of points of Z. Let z_1, $z_2 \in Z$ be points whose Voronoi cells meet, and let M be the midpoint of their common edge. If some $z \in Z \setminus \{z_1, z_2\}$ is at distance less than $\overline{z_1 M} + \delta^6$ from M, then either $\{z_1, z_2\}$ is a δ-close pair, or there is a δ-bad quadruple $\{z_1, z_2, z, z'\}$, $z' \in Z$.*

The proof of this deterministic statement is given in Section 9.1 of the appendix. Lemma 16 has the following immediate corollary.

Corollary 17. *Suppose that no disc of radius $2\sqrt{\log s}$ in $\mathbb{T}$ is free of points of Z. Let z_1, z_2 be δ-good black points of Z whose Voronoi cells share an edge, and let M be the midpoint of that edge. Then the piecewise-linear path $z_1 M z_2$ is δ^6-robustly black.*

Proof. By assumption, z_1 and z_2 are not δ-close, and there is no δ-bad quadruple containing z_1 and z_2, so by Lemma 16 every point $z \neq z_1, z_2$ of Z is at least δ^6 further away from M than z_1 is. Hence M is δ^6-robustly black, and so is the line-segment $z_1 M$: as we move towards z_1 along this line segment at rate 1, the distance from z_1 decreases at rate 1, while the distances from all other points decrease at most this fast. Similarly $z_2 M$ is δ^6-robustly black. $\square$

We have shown that a 'good' black path gives rise to a robustly-black path. It remains to show that we can find good black paths. This will require a slight increase in the probability p with which we colour points of Z black.

6.2 Clusters of bad points

To avoid bad points, we have to show that they are not very common. In fact, we shall prove (and need) a much stronger statement, namely that 'large clusters' of bad points essentially never occur. (See below for a precise statement.) Given Z, a *δ-bad component* is a component in the hypergraph on the δ-bad points of Z with edges given by the δ-close pairs and δ-bad quadruples. In other words, we partition the bad points into minimal components such that every bad quadruple or close pair involves points from a single component.

Lemma 18. *Let $\varepsilon > 0$ and $\eta > 0$ be fixed, let Z be a Poisson process of intensity 1 on $\mathbb{T}(s)$, and suppose that $\delta = \delta(s) \leq s^{-\varepsilon}$. Then* **whp** *no δ-bad component contains more than $\eta \log s$ vertices.*

The proof will be a little long; it makes use of the purely deterministic statement below, which will be proved in Section 9.2 of the appendix.

Lemma 19. *Let $0 < \eta < 1/100$ be fixed. If t is large enough, then it is impossible to arrange $M \geq \eta^{-6}t$ points $P_1, \ldots, P_M$ and M associated circles $C_1, \ldots, C_M$ in $\mathbb{R}^2$ so that no two P_i are within distance η of each other, the distance of every P_i from the origin is between ηt and t, each C_i has radius at most t, no C_i contains any P_j, and C_i passes within distance η^3 of both P_i and the origin.*

Proof of Lemma 18. Throughout *bad* and *close* will mean δ-bad and δ-close respectively. We shall write $f = O^*(g)$ to mean that there is an absolute constant C such that $f = O((\log s)^C g)$ as $s \to \infty$. Suppose the lemma fails for some particular values ε, $\eta > 0$, which we may assume to be smaller than $1/10$.

Any bad quadruple or close pair lies in a disc of radius $4\sqrt{\log s} + 2\delta < \log s/2$, say. (As noted earlier, throughout we use without comment the assumption that s is larger than some constant, depending on ε and η.) Hence, if z is a bad point in a bad component C of more than $\eta \log s$ vertices, the intersection of this component with a square S of side $2(\log s)^2$ centred on z must contain at least $\eta \log s$ points of C: one can not get further than distance $(\log s)^2$ from z in fewer than $\log s$ steps, each step passing from one point of a bad quadruple/close pair to another.

Fix a square S of side $3(\log s)^2$. As we may cover $\mathbb{T}$ with $O(s^2)$ such squares S_i, in such a way that any square of side $2(\log s)^2$ lies within one S_i, it suffices to prove the following statement.

$$\mathbb{P}(S \text{ contains more than } \eta \log s \text{ bad points}) = o(s^{-2}). \tag{19}$$

The number N of points of Z in S has a Poisson distribution with mean $\text{area}(S) = 9(\log s)^4$. Let B_1 be the event that $N > 10(\log s)^4$. Then

$$\mathbb{P}(B_1) = \mathbb{P}(N > 10(\log s)^4) = \exp(-\Theta((\log s)^4)) = o(s^{-2}). \tag{20}$$

Our first step will be to count close pairs in S. In fact, we shall count δ'-close pairs, where $\delta' = \delta^{1/10}$, so δ' is much larger than δ. Any upper bound on the number of δ'-close pairs applies to δ-close pairs as well. Let B_2 be the event that S contains more than $900/\varepsilon^2$ pairs of δ'-close points of Z. Our first aim is to show that $\mathbb{P}(B_2) = o(s^{-2})$. In doing this we shall condition on N; in the light of (20), we may and shall assume that $N \leq 10(\log s)^4$.

Given N, we may realize the positions of the N points of $Z \cap S$ as a random sequence $z_1, \ldots, z_N$, where the z_i are independent and each is uniformly distributed in S. Consider placing the z_i one by one: let C_i be the event that z_i is

δ'-close to one of $\{z_1, \ldots, z_{i-1}\}$. As $i \le N \le 10(\log s)^4 = O^*(1)$, there are $O^*(1)$ previously placed points z_i could be close to. As the area δ'-close to a given point is at most $\pi\delta'^2$, whatever the positions of $z_1, \ldots, z_{i-1}$, the conditional probability of C_i, given N and the positions of $z_1, \ldots, z_{i-1}$, is $O^*(\delta'^2)$. As $N = O^*(1)$ by assumption, it follows that having placed some points $z_1, \ldots, z_j$, the probability that there is an $i > j$ for which C_i holds is $O^*(N\delta'^2) = O^*(\delta'^2) \le \delta'$. Thus the probability that more than $30/\varepsilon$ points z_i are close to earlier points is at most $\delta'^{30/\varepsilon} = \delta^{3/\varepsilon} \le s^{-3} = o(s^{-2})$. A point z_i may be close to many earlier points, but is very unlikely to be close to more than $30/\varepsilon$, say – the probability that S contains any disc of radius δ' containing $30/\varepsilon$ points of Z is $o(s^{-2})$. It follows that with probability $1 - o(s^{-2})$ there are at most $900/\varepsilon^2$ pairs of δ'-close points in S. In other words, $\mathbb{P}(B_2) = o(s^{-2})$.

We should like to count points in bad quadruples in a similar way. The problem is that it is not in general true that, given three points, the probability that there is a fourth point forming a bad quadruple with them is $O(\delta)$. If the first three points are very close together, the probability may in fact be rather large. Some quadruples are easy to deal with, however; we start with those.

Let us say that a δ-bad quadruple is *separated* if no two of its points are δ'-close. Let x_1, x_2, x_3 be any three points of S, and let x be chosen from S uniformly at random. It will be straightforward to check that the probability that the set $\{x_1, x_2, x_3, x\}$ forms a separated bad quadruple is $o(\delta')$. More formally, let us say that a quadruple $\{x_1, x_2, x_3, x_4\}$ is *weakly δ-bad* if there is an $x \in \mathbb{T}$ and an $r < 2\sqrt{\log s}$ such that $r < \overline{xx_i} \le r + \delta$ for every i. This is the same as the definition of δ-badness, except that we have omitted the condition that no other point of Z is within distance r of x. In particular, δ-badness implies weak δ-badness.

Claim 18.1. *Suppose that x_1, x_2, and x_3 are points of S, no two within distance δ' of each other. Let B be the set of $x \in S$ for which the quadruple $\{x_1, x_2, x_3, x\}$ is weakly δ-bad. Then* $\mathrm{area}(B) = o(\delta')$.

The straightforward proof of this purely deterministic statement is given in Section 9.3 of the appendix.

Let B_3 be the event that there are at least $M = 200/\varepsilon$ points in separated bad quadruples in S. Recall that $\varepsilon < 1/10$ is constant, so M is a constant with $M \ge 2000$. We claim that $\mathbb{P}(B_3) = o(s^{-2})$. We condition on N, the number of points of Z in S. We shall assume, as we may, that B_1 does not hold, so $N \le 10(\log s)^4$. Given such an N, construct $Z \cap S$ as $\{z_1, \ldots, z_N\}$ with the z_i iid as above. If B_3 holds, then we can find subsets $\emptyset = Y_0, Y_1, \ldots, Y_r$ of $[N] = \{1, 2, \ldots, N\}$ with $Y_i \subset Y_{i+1}$, $1 \le |Y_{i+1} \setminus Y_i| \le 4$ and $M \le |Y_r| \le M + 3$, where Y_{i+1} is the union of Y_i with a quadruple $\{a, b, c, d\}$ such that $\{z_a, z_b, z_c, z_d\}$ is a separated bad quadruple. (We can continue the sequence from Y_i to Y_{i+1} using the indices $\{a, b, c, d\}$ of any separated bad quadruple such that $\{a, b, c, d\} \not\subset Y_i$. By assumption such a quadruple must exist unless $|Y_i| \ge M$, when we stop.)

It follows that if B_3 holds, then we can find a subset I of $[N]$ of size at most $M + 3$, an order on I, and a subset J of I of size $r \ge M/4$ so that if we add the points $\{z_i : i \in I\}$ in this particular order, then whenever $j \in J$ the point z_j completes a specified separated bad quadruple. Fixing I, the order, and J, the probability that one z_j, $j \in J$, completes a separated bad quadruple is $o(\delta') = o(\delta^{1/10})$, using uniform distribution of z_j and Claim 18.1. Hence, the probability that each z_j, $j \in J$, completes a separated bad quadruple is $o(\delta^{M/40})$. Crudely, the number of choices for I is at most $4N^{M+3} \leq (\log s)^{5M}$. The number of choices for the order is $|I|! \leq |I|^{|I|} \leq M^{2M}$, while the number of choices for J is at most 2^{2M}. Hence,

$$\mathbb{P}(B_3 \mid B_1^c) \leq \left(4(\log s)^5 M^2 \delta^{1/40}\right)^M \leq \delta^{M/50} = o(s^{-2}).$$

Since $\mathbb{P}(B_1) = o(s^{-2})$, it follows that $\mathbb{P}(B_3) = o(s^{-2})$ as claimed.

We have now shown that there is a constant C (depending on the parameters ε and η appearing in the statement of Lemma 18) such that with probability $1-o(s^{-2})$ the square S contains (a) at most C pairs that are δ'-close, and (b) at most $200/\varepsilon$ points in separated δ-bad quadruples. To complete the proof of (19) it suffices to show that with high enough probability there are at most $\eta \log s/2$ points $z \in Z \cap S$ with the property that z is not in a δ'-close pair, but is in a δ-bad quadruple of points of $Z \cap S$ containing a δ'-close pair $\{x_1, x_2\}$.

Let B_4 be the event that there are more than $\eta \log s/2$ such points z. If B_4 holds but B_2 does not, then there is a δ'-close pair $\{x_1, x_2\} \subset Z \cap S$ and a set D of points of $Z \cap S$ so that every $z \in D$ is in a δ-bad quadruple with $\{x_1, x_2\}$, where D contains at least $\eta' \log s$ points, $\eta' = \eta'/(2C)$. From the definition of δ-badness, each point of D is within distance $2\sqrt{\log s} + 2\delta \leq 3\sqrt{\log s}$ of x_1. By Lemma 14, there is a constant $\eta_2 > 0$ such that the event B_5 that some disc in $\mathbb{T}$ of radius $\eta_2\sqrt{\log s}$ contains more than $\eta' \log s/2$ points of Z has probability $\mathbb{P}(B_5) = o(s^{-2})$. Assuming that B_5 does not hold, we see that D contains a set D' of at least $\eta' \log s/2$ points whose distance from x_1 is between $\eta_2\sqrt{\log s}$ and $3\sqrt{\log s}$.

We shall next show that for some constant η_3, with very high probability not many points in D' are η_3-close to other points in D'. The argument is as at the beginning of the lemma for δ'-close points, but we must consider a total area of the right order of magnitude.

Fix any disc X of radius $4\sqrt{\log s}$. From Lemma 14, there is a constant C' such that the event $B_1'(X)$ that X contains more than $C' \log s$ points of Z has probability $o(s^{-4})$. Let $\eta_3 > 0$ be constant. Arguing as at the start of the proof of this lemma, as we place points one by one in X, assuming that $B_1'(X)$ does not hold, at each stage the probability that the new point z_i is within distance η_3 of an earlier point is $O(\eta_3^2)$: the area η_3-close to some earlier point is at most $\pi\eta_3^2 C' \log s$, while z_i is chosen uniformly from X, a domain of area $\text{area}(X) = 16\pi \log s$. Hence the number $Y(X)$ of points z_i η_3-close to an earlier point is dominated by a binomial distribution $\text{Bi}(n,p)$ with parameters $n = C' \log s$ and $p = O(\eta_3^2)$. Choosing η_3 small enough, and using the weak bound $\mathbb{P}(\text{Bi}(n,p) \geq t) \leq (enp/t)^t$, the probability that $Y(X)$ exceeds $\eta' \log s/4$ is $o(s^{-4})$.

Let B_6 be the event that some disc X' of radius $3\sqrt{\log s}$ meeting S contains at least $\eta' \log s/4$ points $z_i \in Z$ each closer than η_3 to some earlier point of Z in X' (taking a random order on the points of $Z \cap X'$ as above). Then (placing $O(s^2)$ discs of radius $4\sqrt{\log s}$ so that any disc of radius $3\sqrt{\log s}$ is contained in one), we have $\mathbb{P}(B_6) = o(s^{-2})$. Assuming that B_6 does not hold, and deleting from D' each point z_i closer than η_3 to an earlier point, we find a set $D'' \subset D'$ in which any two points are at least η_3 apart, with $|D''| \geq \eta' \log s/4$.

Let us write $D'' = \{P_1, \ldots, P_M\}$ where $M \geq \eta' \log s/4$, and set $t = 3\sqrt{\log s}$. Then any two P_i are at least η_3 apart, and the distance from any P_i to x_1 is between $\eta_2\sqrt{\log s} = \eta_2 t/3$ and t. Furthermore, as each P_i is in a δ-bad quadruple with x_1 and x_2, for each P_i there is a circle C_i with radius at most $2\sqrt{\log s} < t$ passing within distance $\delta = o(1)$ of P_i and x_1, and containing no points of Z, and hence no P_j. For s large enough, the existence of points P_i and circles C_i with the above properties contradicts the deterministic Lemma 19, applied with $t = 3\sqrt{\log s}$ and $\eta = \min\{\eta_2/3, \eta_3, 1/100\} = \Theta(1)$.

In summary, if S contains many bad points, then one of the 'bad' events B_i, $1 \leq i \leq 6$, holds. Each has probability $o(s^{-2})$, so with probability $1 - o(s^{-2})$, S contains fewer than $\eta \log s$ bad points, proving (19). As noted at the start of the proof, this suffices to prove Lemma 18. $\square$

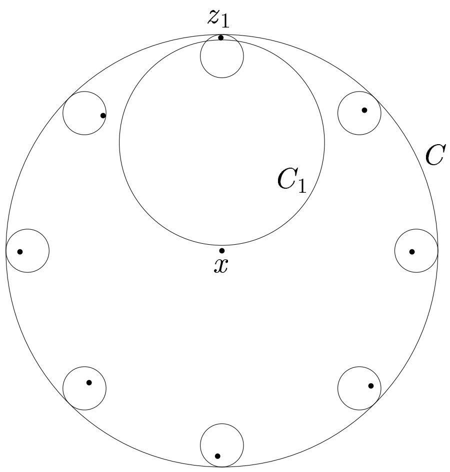

Figure 9: Three points $\{x_1, x_2, x_3\}$ of Z are very close to x. If one places one more point z_i of Z in each small circle in an arbitrary way, and C is otherwise free of points of Z, then each z_i forms a bad quadruple with $\{x_1, x_2, x_3\}$; the circle C_1 illustrates this for z_1.

The last part of the proof above may seem unnecessarily complicated, but the situation is a little delicate. One might expect that there is some argument that the size of all bad components will be at most a constant, as (locally) each individual 'badness' has probability $n^{-\Theta(1)}$, but this is not true. In fact, with high probability, $\mathbb{T}(s)$ will contain a bad component with $\Theta(\log s/\log\log s)$ vertices: consider three points very close together, and k points spaced almost evenly around a circle of radius 1, say. Each of these k points can each be placed anywhere in a small disc of radius $\Theta(k^{-2})$ so that all k points form bad quadruples with the three central points; this is illustrated for $k = 8$ in Figure 9. Such a configuration will occur somewhere if $k = c\log s/\log\log s$ for a small enough constant c. Thus the bound $o(\log s)$ given by Lemma 18 is not far from best possible.

6.3 Bad points can be avoided

In the next lemma, for $p_1 < p_2$ we couple the probability measures (on black/white-coloured Poisson processes) $\mathbb{P}_{p_1}^{\mathbb{T}(s)}$ and $\mathbb{P}_{p_2}^{\mathbb{T}(s)}$. In other words, we shall construct simultaneously Poisson processes Z_1 and Z_2 of intensity 1 on $\mathbb{T}(s)$, with associated colourings col_1, col_2, so that given Z_i each point of Z_i is coloured black with probability p_i independently of all other points of Z_i, and white otherwise.

Let us write E_{good} for the global event that the following three conditions hold:

(a) there is a bijection ϕ from the points of Z_1 to those of Z_2 mapping black points into black points,

(b) each point z of Z_1 is within distance 1 of the corresponding point $\phi(z)$ of Z_2, and

(c) if z, z' are black points of Z_1 whose Voronoi cells share an edge, then there is a sequence $\phi(z) = z_1, z_2, \ldots, z_t = \phi(z')$, $t = O(\log s)$, of black, δ-good points of Z_2, such that the Voronoi cells of z_i and z_{i+1} with respect to Z_2 share an edge.

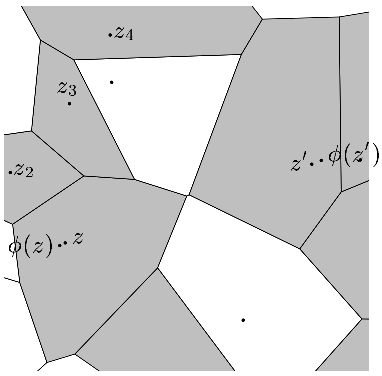

Figure 10: Part of the Voronoi tiling with respect to Z_2, together with two points z, z' of Z_1 whose Voronoi cells with respect to Z_1 meet; these cells are not drawn.

The event E_{good} is illustrated in Figure 10: part of the Voronoi tiling with respect to Z_2 is shown, together with the points z, z' of Z_1. For simplicity we assume that $\phi(z'') = z''$ for all $z'' \in Z_1 \setminus \{z, z'\}$, so z and z' are the only points that move when we pass from Z_1 to Z_2. The Voronoi cells of z and z' with respect to Z_1, which are not shown, meet. The points z_i, with $z_1 = \phi(z)$ and $z_5 = \phi(z')$, exhibit the path required for condition (c).

As usual, in the following statement $o(1)$ notation refers to $s \to \infty$ with all other parameters (here p_1, p_2 and ε) fixed. The statement is not very appealing, but will allow us to deduce Theorem 15, at which point the lemma can be forgotten.

Lemma 20. *Let $0 < p_1 < p_2 < 1$ and $\varepsilon > 0$ be given. Let $\delta = \delta(s)$ be any function with $\delta(s) \leq s^{-\varepsilon}$. We may construct in the same probability space coloured Poisson processes (Z_1, col_1) and (Z_2, col_2) on $\mathbb{T}(s)$ as described above, in such a way that the global event E_{good} defined above holds* **whp**.

Proof. An outline of the idea is as follows: we first decide roughly where the points z_i of Z are. Then in most cases we can already see that a particular z_i is δ-good. There will be some clusters of points that may be δ-bad; however, these will have size $o(\log s)$. Different clusters will behave independently when we decide the precise location of the z_i, and each will contain some bad point with only small probability. We can allow for this by using the difference in the two colouring probabilities.

We now turn to the details. As usual, throughout the proof we assume that s is larger than some very large constant. We may assume that $\delta = s^{-\varepsilon}$.

Pick δ_1 with $\delta_1 \sim \delta^{1/5} = n^{-\varepsilon/5}$ so that s/δ_1 is an integer, and partition $\mathbb{T}(s)$ into $K = (s/\delta_1)^2$ squares S_k of side δ_1 in the natural way. Set $\delta_2 = \delta_1^{1/10}$. Note that $\delta << \delta_1 << \delta_2 \sim n^{-\varepsilon/50}$.

Let N be a Poisson random variable with mean s^2. For $i = 1, 2$ we shall realize Z_i as $\{z_{i,1}, \ldots, z_{i,N}\}$, where the $z_{i,j}$ are uniformly distributed on $\mathbb{T}$, and for each i, the $z_{i,j}$, $1 \leq j \leq N$, are independent. For $1 \leq j \leq N$, choose s_j independently and uniformly at random from $[K] = \{1, 2, \ldots, K\}$. We shall place both $z_{1,j}$ and $z_{2,j}$ into the square S_{s_j}. As long as the locations of each $z_{i,j}$ within S_{s_j} are chosen uniformly at random, with independence as j varies and i is fixed, we shall obtain coupled Poisson processes Z_1, Z_2 so that there is a bijection ϕ satisfying condition (b) of E_{good}.

Having chosen the s_j, we say that a pair $\{a, b\} \subset [N]$ is *potentially close* if it is possible, given that $z_{2,j} \in S_{s_j}$ for all j, that $z_{2,a}$ and $z_{2,b}$ are δ_2-close, i.e., if some two points of the squares S_{s_a} and S_{s_b} are within distance δ_2. Note that in this case, wherever $z_{2,a}$ and $z_{2,b}$ are placed in their squares, they are certainly $(\delta_2 + 2\sqrt{2}\delta_1)$-close and hence $(2\delta_2)$-close.

A quadruple $\{a, b, c, d\} \subset [N]$ is *potentially bad* if it is possible, given that $z_{2,j} \in S_{s_j}$ for all j, that $Q = \{z_{2,a}, z_{2,b}, z_{2,c}, z_{2,d}\}$ forms a δ_2-bad quadruple. This time it follows that wherever the points $z_{2,j}$ are within their squares, Q is a $(2\delta_2)$-bad quadruple. Indeed, by definition there is a placement of the $z_{2,j}$ so that Q is a δ_2-bad quadruple. Let $x \in \mathbb{T}$ and r be the point and radius witnessing this, so no $z_{2,j}$ is within distance r of x and all four points of Q are within distance $r + \delta_2$. Moving each $z_{2,j}$ arbitrarily within S_{s_j}, each $z_{2,j}$ moves by at most a distance $\sqrt{2}\delta_1$. Reducing r by $\sqrt{2}\delta_1$ (stopping at 0 if r is smaller than this to start with!), we find a smaller r' so that none of the relocated $z_{2,j}$ is within distance r' of x, but all four relocated points of Q are within distance $r + \delta_2 + \sqrt{2}\delta_1 \leq r' + \delta_2 + 2\sqrt{2}\delta_1 \leq r' + 2\delta_2$ of x.

In analogy with previous definitions we say that a point $z_{2,a}$, or the corresponding index $a \in [N]$, is *potentially bad* if it is in a potentially close pair or potentially bad quadruple, and define *potentially bad components* as components of the hypergraph induced on the set of potentially bad points by the potentially close pairs/bad quadruples. Note that these definitions all depend only on N and the s_j, i.e., on which squares each $z_{2,j}$ lies in. As we place $z_{1,j}$ and $z_{2,j}$ into the same square S_{s_j}, we could have written $z_{1,j}$ instead of $z_{2,j}$ throughout all definitions of potential badness.

We claim that for any fixed $\eta > 0$, with probability $1 - o(1)$ no potentially bad component contains more than $\eta \log s$ points. This follows from Lemma 18, applied with $2\delta_2 \leq s^{-\varepsilon/51}$ in place of δ: we just realize the exact locations of the $z_{2,j}$, and apply Lemma 18 to see that **whp** there is no component of more than $\eta \log s$ $(2\delta_2)$-bad points. But every potentially bad pair/quadruple gives rise to a $(2\delta_2)$-bad pair/quadruple, so the largest $(2\delta_2)$-bad component is certainly at least as large as the largest potentially bad component.

As outlined above, we construct the coupling as follows. First, we decide N, and the s_j. Later we shall position $z_{1,j}$ and $z_{2,j}$ in S_{s_j} appropriately. Also, we shall define the colourings so that whenever $\mathrm{col}_1(z_{1,j})$ is black, so is $\mathrm{col}_2(z_{2,j})$.

Let B_1 be the 'bad' event that there is some disc X of radius $2\sqrt{\log s}$ in $\mathbb{T}$ so that every square S_k contained entirely in X is free of points of Z_1 (and hence of Z_2). Then $\mathbb{P}(B_1) = o(1)$ by Lemma 14: if B_1 holds, then, when we place the points $z_{1,j}$ randomly in the corresponding S_{s_j}, there is always a smaller disc within X, of radius at least $2\sqrt{\log s} - \sqrt{2}\delta_1 \geq \sqrt{\log s}$, containing no points of Z_1, a Poisson process of intensity 1 on $\mathbb{T}$. The existence of such a disc has probability $o(1)$ by Lemma 14, so $\mathbb{P}(B_1) = o(1)$. Note that if B_1 does not hold, then every disc of radius $2\sqrt{\log s}$ will contain points of both Z_1 and Z_2, however we complete the coupling. Let B_2 be the event that there is a potentially bad component of more than $\eta \log s$ points, where $\eta > 0$ is a very small constant to be chosen later. As shown above, $\mathbb{P}(B_2) = o(1)$. Note that B_1 and B_2 depend only on N and the s_j.

If B_1 or B_2 holds, we abandon the construction and complete the coupling any old way; E_{good} will not hold in this case, but this does not matter, as $\mathbb{P}(B_1 \cup B_2) = o(1)$. Suppose from now on that N and the s_j are fixed and that $B_1 \cup B_2$ does not hold.

Let $\{a, b\} \subset [N]$ be a pair not in the same potentially bad component. Thus a and b are not potentially close, and are not two of the four points in a potentially bad quadruple. We claim that, no matter how we position the $z_{i,j}$ in their squares S_{s_j}, if the Voronoi cells of $z_{1,a}$ and $z_{1,b}$ with respect to Z_1 are adjacent, then the cells of $z_{2,a}$ and $z_{2,b}$ with respect to Z_2 will also be adjacent. Indeed, consider any placements Z_1, Z_2 consistent with our choice of squares. By definition of potential badness, $z_{1,a}$ and $z_{1,b}$ are not δ_2-close, and do not belong to a δ_2-bad quadruple. Now consider the midpoint M of the common edge of the relevant cells of Z_1. By Lemma 16, no point $z_{1,c}$ of $Z_1 \setminus \{z_{1,a}, z_{1,b}\}$ is within distance $\overline{z_1 M} + \delta_2^6$ of M. Moving from Z_1 to Z_2, all points move a distance of at most $\sqrt{2}\delta_1 < \delta_2^6/2$, so $z_{2,a}$ and $z_{2,b}$ are the two points of Z_2 closest to M. It follows by Lemma 8 that the Voronoi cells of $z_{2,a}$ and $z_{2,b}$ with respect to Z_2 meet.

We are now ready to construct the coupling; the observation above suggests that we consider each potentially bad component separately. Let us start with the trivial components: suppose that $j \in [N]$ is not potentially bad. Then we couple the $z_{i,j}$ and their colourings in the following natural way: take $z_{1,j} = z_{2,j}$ uniformly random on S_{s_j}, colour $z_{2,j}$ black with probability p_2 independently of everything else, and colour $z_{1,j}$ black only if $z_{2,j}$ is black, and then only with probability p_1/p_2.

Next, let C be a non-trivial potentially bad component. We start by coupling the $z_{i,j}$, $j \in C$, in the natural way as above; we shall have to adjust the coupling slightly. Consider the event $B(C)$ that there is a pair $\{a, b\} \subset C$ so that $z_{2,a}$ and $z_{2,b}$ are δ-close, or a quadruple $\{a, b, c, d\} \subset C$ so that $\{z_{2,a}, z_{2,b}, z_{2,c}, z_{2,d}\}$ forms a δ-bad quadruple. Note that here we consider δ-badness, a much stronger condition than δ_2-badness.

We claim that $\mathbb{P}(B(C)) \leq s^{-\varepsilon/6}$. Indeed, considering quadruples first, there are at most $|C|^4 = O^*(1)$ quadruples $\{a, b, c, d\}$ to consider, using the assumption that B_2 does not hold. For each, the joint distribution of $(z_{2,a}, z_{2,b}, z_{2,c}, z_{2,d})$ is that these points are chosen independently and uniformly at random from certain squares S_{s_a}, etc. Recall that each square S_i has side-length δ_1, and that $\delta_1 \sim \delta^{1/5}$. In particular, δ_1 is much larger than δ. It is easy to check that no matter how the squares are located, the probability that the quadruple is δ-bad is $O^*(\delta/\delta_1^4) = O^*(\delta_1)$. Indeed, one argument goes as follows. If the quadruple is δ-bad, then there is a point x and a radius r such that r and both coordinates of x are multiples of δ, so that all four points are in the annulus centred at x with inner radius r and outer radius $r + 5\delta$. (We find x and r close to the x' and r' witnessing δ-badness.) For any such annulus, its intersection with each S_j has area $O(\delta\delta_1)$, so the probability that all four points lie in one given annulus is $O((\delta\delta_1/\delta_1^2)^4) = O(\delta^4/\delta_1^4)$. But the number of choices for the annulus is $O^*(\delta^{-3})$; the range of the coordinates of x and of r is at most $O(\sqrt{\log s})$, as the definition of δ-badness limits r to $O(\sqrt{\log s})$. Hence any particular quadruple is δ-bad with probability $O^*(\delta/\delta_1^4) = O^*(\delta_1)$, and the expected number of δ-bad quadruples is $O^*(\delta_1)$. The argument for pairs is much simpler: each of the $O^*(1)$ pairs is δ-close with probability at most $O(\delta\delta_1/\delta_1^2) = O(\delta_1^4)$. Thus $\mathbb{P}(B(C)) = O^*(\delta_1) \leq s^{-\varepsilon/6}$, as claimed.

Let $G(C)$ be the event that every point in the component is white with respect to (Z_1, col_1) and black with respect to (Z_2, col_2). From the way the coupling is defined, $\mathbb{P}(G(C)) = (p_2 - p_1)^{|C|}$. As ε, p_1 and p_2 are constants, while $|C| \leq \eta \log s$, we have $\mathbb{P}(G(C)) \geq 2s^{-\varepsilon/6}$ provided we chose η small enough. Hence $\mathbb{P}(G(C) \setminus B(C)) \geq \mathbb{P}(B(C))$.

Let $G'(C) \subset G(C) \setminus B(C)$ be an event with probability exactly $\mathbb{P}(B(C))$. Our final coupling is defined by 'crossing over' from the natural coupling using

$B(C)$ and $G'(C)$: deleting the portion of the probability space in which $B(C)$ or $G'(C)$ holds, we add instead the distribution of (Z_1, col_1) on $B(C)$ coupled in an arbitrary way with that of (Z_2, col_2) on $G'(C)$, and vice versa. Thus, exactly one of the following three events holds in the final crossed over coupling:

1. $z_{1,j} = z_{2,j}$ for all $j \in C$, there is no δ-close pair/bad quadruple in C, and if $\mathrm{col}(z_{1,j})$ is black, so is $\mathrm{col}(z_{2,j})$,

2. Z_1 contains a δ-close pair or δ-bad quadruple in C while Z_2 does not, and $\mathrm{col}_2(z_{2,j})$ is black for every $j \in C$, or

3. Z_2 contains a δ-close pair or δ-bad quadruple in C while Z_1 does not, and $\mathrm{col}_1(z_{1,j})$ is white for every $j \in C$.

We extend the coupling from individual components C to the whole of Z_1, Z_2 by taking the product of the relevant measures, i.e., treating different components independently.

Note that we have indeed coupled the correct distributions for the (Z_i, col_i): the 'natural' coupling clearly does this, and the marginal distributions are preserved by the crossing over. From the way we constructed the coupling, placing points always in the same squares S_j, conditions (a) and (b) of the event E_{good} hold: note that in each case above $z_{2,j}$ is black whenever $z_{1,j}$ is. It remains to check condition (c). For this we shall use the following claim.

Claim 20.1. *Suppose that $a, b \in C$ for some potentially bad component C, and that $z_{1,a}$ and $z_{1,b}$ are black points whose Voronoi cells with respect to Z_1 meet. Then there is a piecewise linear path P joining $z_{2,a}$ to $z_{2,b}$ so that, with respect to Z_2, every point of P lies in the Voronoi cell of some $z_{2,c}$, $c \in C$, where $z_{2,c}$ is black and δ-good.*

The straightforward proof of this deterministic statement is given in Section 9.4 of the appendix.

Claim 20.1 contains everything we need to finish the proof of Lemma 20: it remains only to establish the existence and properties of the path (of cells) required by condition (c) in the definition of E_{good}. Note first that from the properties of the three cases 1,2,3 above, whenever a point is black in Z_1 it is black in Z_2 and δ-good in Z_2. (Recall that if $j \in C$ is in a δ-close pair/bad quadruple in Z_2, then this pair/quadruple is contained in C, as C is a component in the hypergraph of potentially close pairs/bad quadruples.)

Suppose that $z_{1,a}$ and $z_{1,b}$ are black points of Z_1 whose Voronoi cells meet. Then $z_{2,a}$ and $z_{2,b}$ are black and δ-good. We have already shown above that if a and b do not lie in the same potentially bad component, then the Voronoi cells of $z_{2,a}$ and $z_{2,b}$ must also meet, and (c) holds. The claim gives (c) in the case that a and b do lie in the same potentially bad component, using $|C| = O(\log s)$ and noting that if there is an x-y path in the union of certain Voronoi cells, then there is such a path visiting each of the cells at most once. $\square$

We now have all the tools we need to prove Theorem 15.

Proof of Theorem 15. We use the coupling guaranteed by Lemma 20, applied with $(4\delta)^{1/6}$ in place of δ. We shall assume that every disc of radius $\sqrt{\log s}$

contains points of both Z_1 and Z_2; this holds **whp** by Lemma 14. Given P_1, there is a sequence of (not necessarily distinct) black points $z_{1,i}$, $1 \le i \le t$, in Z_1 so that the Voronoi cells of $z_{1,i}$ and $z_{1,i+1}$ share an edge for each i, and so that P_1 meets the Voronoi cells corresponding to each $z_{1,i}$ and lies in the union of these cells. From Lemma 20, for each i there is a sequence S_i of $(4\delta)^{1/6}$-good black points $z_{2,j}^{(i)}$ of Z_2 so that for each j the Voronoi cells of $z_{2,j}^{(i)}$ and $z_{2,j+1}^{(i)}$ with respect to Z_2 share an edge, where S_i starts with the unique point $\phi(z_{1,i}) \in Z_2$ corresponding to $z_{1,i}$, ends with $\phi(z_{1,i+1})$, and has length $O(\log s)$. Piecing together the sequences S_i, we obtain a single sequence $z_{2,j}$, $1 \le j \le t'$. By Corollary 17, the piecewise-linear path P_2 formed by the $z_{2,j}$ and the midpoints of the edges where the cells of consecutive $z_{2,j}$ meet is 4δ-robustly black with respect to Z_2. Finally, as all Voronoi cells have diameter $O(\sqrt{\log s})$, each S_i has length at most $O(\log s)$, and $z_{2,1}^{(i)} = \phi(z_{i,1})$ is within distance 1 of $z_{1,i}$, the paths P_1 and P_2 are geometrically 'close' to each other as claimed. $\square$

The entire purpose of the present section was to develop a tool (Theorem 15) that will enable us to deal with problems arising when we discretize the continuous Poisson process Z. As noted in the introduction, these problems would not arise in, for example, the random discrete Voronoi setting.

7 Crossing a rectangle when $p > 1/2$

Using Corollary 13, Theorem 5 and Theorem 15, we shall prove that for any $p > 1/2$, some large enough long thin rectangle can be crossed by a black path with probability at least 99%. Once we have this, it will be easy to prove our main results, Theorems 1 and 2. Recall that $f_p(\rho, s)$ is the probability that $H(R)$ holds, i.e., that there is a (piecewise linear) black path crossing R from left to right, when R is a ρs by s rectangle, in the probability measure $\mathbb{P}_p = \mathbb{P}_p^{\mathbb{R}^2}$.

Theorem 21. *Let $p > 1/2$, $c_1 < 1$, $\rho > 1$ and s_1 be given. There is an $s > s_1$ such that $f_p(\rho, s) > c_1$.*

Proof. It suffices to prove the result for any fixed $\rho > 1$, for example, $\rho = 3$. The general result follows by using positive correlation (Lemma 10) as in the proof of equation (4).

Throughout we fix $p > 1/2$, $c_1 < 1$, $\rho = 3$ and s_1. All implicit constants in our notation ($O(.)$, etc) will depend on p and c_1 only. We shall choose constants ε and s_0 as follows. We first fix $\varepsilon = \varepsilon(p, c_1) > 0$ small enough that a certain condition we shall encounter in the course of the proof holds; this condition does not involve s_0 (or s). Then we choose an s_0 depending on all parameters chosen so far, large enough that the statements '... provided s is sufficiently large' in what follows hold for all $s \ge s_0$. (In principle, we could give an explicit value.) We assume throughout that $s_0 \ge 6s_1$.

By Corollary 13, applied with $\rho = 10$, there is an absolute constant $c_0 > 0$ such that for some $s \ge s_0$ we have with probability at least c_0 a horizontal (i.e., long) black crossing of $R_1 = [0, 10s/13] \times [0, s/13]$ in the $p = 1/2$ random Voronoi tiling in $\mathbb{R}^2$. Let us fix such an s throughout the proof.

Let us consider the fixed rectangle R_1 as a rectangle in the torus $\mathbb{T} = \mathbb{T}(s)$. Let E_1 be the event that R_1 has a black crossing P_1 in the Voronoi tiling defined on the torus. As noted in Section 5, since R_1 does not come close to 'wrapping round' the torus, if s is large enough we have $\mathbb{P}^{\mathbb{T}}_{1/2}(E_1) = \mathbb{P}^{\mathbb{R}^2}_{1/2}(H(R_1)) + o(1) \geq c_0/2$.

Let δ be chosen so that $s^{-\varepsilon} \leq \delta \leq (1 + o(1))s^{-\varepsilon}$, and so that s/δ is an integer. Set $p' = (1/2 + p)/2$, so $1/2 < p' < p$. Consider the coupling given by Theorem 15, applied with $p_1 = 1/2$, $p_2 = p'$, and 2ε in place of ε. Suppose that E_1 holds with respect to the process (Z_1, col_1); as noted above, this event has probability at least $c_0/2$. Suppose also that the global event described in Theorem 15 holds, as it does with probability $1 - o(1)$. Then with respect to (Z_2, col_2) there is a 4δ-robustly black path P_2 geometrically close to P_1. In particular, if s is large enough, P_2 remains entirely within the rectangle R_1' with the same centre as R_1 but with corresponding sides $2(\log s)^2 < s/1000$ longer. Furthermore, P_2 contains points within distance $s/1000$ of the left and right sides of R_1. Hence, a sub-path P_2' of P_2 crosses R_2 from left to right, where R_2 is the $3s/4$ by $s/12$ rectangle with the same centre as R_1; see Figure 11.

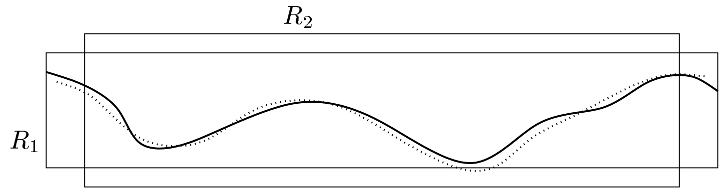

Figure 11: The rectangles R_1 and R_2, together with a black path P_1 (shown solid) crossing R_1 horizontally, and the nearby path P_2 (shown dotted), part of which crosses R_2 horizontally.

To summarize, let R_2 be the fixed $3s/4$ by $s/12$ rectangle in $\mathbb{T}$ described above, and let E_2 be the event that there is a 4δ-robustly black path P crossing R_2 horizontally. As the assumptions made above hold simultaneously with probability at least $c_0/2 - o(1)$, and these assumptions imply that E_2 holds for (Z_2, col_2), we have

$$\mathbb{P}^{\mathbb{T}}_{p'}(E_2) \geq c_0/3,$$

if s is large enough.

Our main aim is to apply Theorem 5 to an appropriate symmetric event in a suitable discrete probability space. Let E_3 be the event that there is a 4δ-robustly black path crossing *some* $3s/4$ by $s/12$ rectangle in $\mathbb{T}$ horizontally. Clearly E_3 is 'symmetric', and

$$\mathbb{P}^{\mathbb{T}}_{p'}(E_3) \geq \mathbb{P}^{\mathbb{T}}_{p'}(E_2) \geq c_0/3.$$

The next step is to define a corresponding symmetric event in a discrete probability space. Divide $\mathbb{T} = \mathbb{T}(s)$ up into $(s/\delta)^2$ small squares S_i of side-length δ in the natural way. The *crude state* of each S_i will be *bad* if S_i contains at least one white point of Z (and perhaps one or more black points), *neutral* if S_i contains no points of Z, and *good* if S_i contains at least one black point of Z and no white points. If we generate Z as a Poisson process with intensity 1 and colour each point independently black with probability p'' bounded away from 0 and 1 and white with probability $1-p''$, then writing $\gamma = \delta^2 = o(1)$ for the area of S_i, each S_i is bad/neutral/good with respective probabilities

$$\begin{aligned} p_{\text{bad}} &= 1-\exp(-\gamma(1-p'')) \sim \gamma(1-p''), \\ p_{\text{neutral}} &= \exp(-\gamma), \\ p_{\text{good}} &= \exp(-\gamma(1-p''))(1-\exp(-\gamma p'')) \sim \gamma p'', \end{aligned} \tag{21}$$

independently of all other S_j. The *crude state* CS of $\mathbb{T}$ is given by the crude states of the S_i.

Define E_3^{crude} to be the event that the crude state CS of $\mathbb{T}$ is such that E_3 is possible given CS. As $\mathbb{P}(E_3) = \mathbb{E}(\mathbb{P}(E_3 \mid CS)) \le \mathbb{P}(E_3^{\text{crude}})$, we have

$$\mathbb{P}_{p'}^{\mathbb{T}}(E_3^{\text{crude}}) \ge c_0/3 \tag{22}$$

if s is large enough.

If we write $n = (s/\delta)^2$ for the number of squares S_i, and take $p_- = p_{\text{bad}}$ and $p_+ = p_{\text{good}}$, then E_3^{crude} may be regarded as an event E_3^{disc} in the finite probability space $W^n_{p_-,p_+}$ defined in Section 3. Here the ith coordinate in the state space $\{-1,0,1\}^n$ of $W^n_{p_{\text{bad}},p_{\text{good}}}$ is -1, 0 or 1 if the crude state of S_i is bad, neutral or good respectively.

Now E_3 is 'symmetric' in the sense that it is preserved under translations of the torus. It follows that E_3^{crude} is preserved under horizontal and vertical translations of the torus by integer multiples of δ. Such translations can map any S_i into any other S_j. Thus E_3^{disc} is symmetric with respect to the coordinates $[n] = \{1,2,\ldots,n\}$ of the product space $W^n_{p_{\text{bad}},p_{\text{good}}}$; there is a permutation group acting transitively on $[n]$ whose action on $\{-1,0,1\}$, the state space of $W^n_{p_{\text{bad}},p_{\text{good}}}$, preserves E_3^{disc}. This is precisely the symmetry condition required for Theorem 5.

We claim that E_3^{disc} is increasing in $\{-1,0,1\}^n$. In other words, we claim that E_3^{crude} is preserved if we change the crude state of one S_i from bad to neutral or from neutral to good. To see this, suppose that we are given CS (i.e., the crude states of all the S_j) such that E_3^{crude} holds. Fix one S_i and suppose that this S_i is neutral. By the definition of E_3^{crude}, the event E_3 is possible given this overall crude state. Changing the state of S_i to good clearly preserves the possibility of E_3, as it corresponds to adding one or more black points, and E_3 is black-increasing. Suppose now that the state of S_i is bad, and pick a particular realization of (Z, col) consistent with CS for which E_3 holds. In Z, there is (at least one) white point w in S_i; there may also be one or more black points b. To show that E_3 is still possible after changing S_i to neutral, it suffices to check that E_3 still holds after deleting all points of Z in S_i. As E_3 is black-increasing, the only possible problem is that deleting the black points destroys E_3. But this cannot happen, because of the point w, and the fact that a black point b within distance 2δ of some white point cannot help E_3 to hold – such a black point cannot be the closest point of Z to a 4δ-robustly black point of $\mathbb{T}$.

We are ready to apply Theorem 5. Let $p_{\mathrm{bad},0}$ and $p_{\mathrm{good},0}$ be the values defined by (21) with $p'' = p'$, and let $p_{\mathrm{bad},1}$ and $p_{\mathrm{good},1}$ be the values defined by (21) with $p'' = p$. From the correspondence between $W^n_{p_{\mathrm{bad}},p_{\mathrm{good}}}$ and the crude state of the torus, we have

$$\mathbb{P}^n_{p_{\mathrm{bad},0},p_{\mathrm{good},0}}(E_3^{\mathrm{disc}}) = \mathbb{P}^{\mathbb{T}}_{p'}(E_3^{\mathrm{crude}}) \tag{23}$$

and

$$\mathbb{P}^n_{p_{\mathrm{bad},1},p_{\mathrm{good},1}}(E_3^{\mathrm{disc}}) = \mathbb{P}^{\mathbb{T}}_{p}(E_3^{\mathrm{crude}}). \tag{24}$$

From (22) and (23) it follows that

$$\mathbb{P}^n_{p_{\mathrm{bad},0},p_{\mathrm{good},0}}(E_3^{\mathrm{disc}}) \geq c_0/3. \tag{25}$$

Now

$$p_{\mathrm{bad},0} \sim (1-p')\gamma \quad \text{and} \quad p_{\mathrm{good},0} \sim p'\gamma, \tag{26}$$

while

$$p_{\mathrm{bad},1} \sim (1-p)\gamma \quad \text{and} \quad p_{\mathrm{good},1} \sim p\gamma. \tag{27}$$

We shall apply Theorem 5 with $n = (s/\delta)^2$, and with $p_- = p_{\mathrm{bad},0}$, $p_+ = p_{\mathrm{good},0}$, $q_- = p_{\mathrm{bad},1}$ and $q_+ = p_{\mathrm{good},1}$. Recall that p' and p are constants with $p' < p$, and note that

$$q_+ - p_+,\; p_- - q_- \sim (p-p')\gamma. \tag{28}$$

Let $c_2 < 1$ be any absolute constant, and let η be the absolute constant $\min\{c_0/4, 1-c_2\}$. From (26) and (27) all four of p_-, p_+, q_-, q_+ are at most γ if s is large enough. Hence the quantity $p_{\max}$ appearing in (1) is at most γ, and the right hand side of (1) is at most

$$\Delta = c_3 \log(1/\eta)\gamma \log(1/\gamma)/\log n,$$

where c_3 is an absolute constant. As $n = (s/\gamma)^2 \geq s^2$ and $\gamma \geq s^{-2\varepsilon}$, we have $\log(1/\gamma)/\log n \leq \varepsilon$. It follows that $\Delta \leq C\varepsilon\gamma$ for some absolute constant C. By our choice of ε we may ensure that $C\varepsilon < (p-p')/2$, so $\Delta < (p-p')\gamma/2$. Hence, (28) implies that if s is large enough, then $q_+ - p_+$ and $p_- - q_-$ are both at least Δ. Thus, by (25) and Theorem 5,

$$\mathbb{P}^n_{p_{\mathrm{bad},1},p_{\mathrm{good},1}}(E_3^{\mathrm{disc}}) \geq c_2.$$

Hence, from (24), we have $\mathbb{P}^{\mathbb{T}}_p(E_3^{\mathrm{crude}}) \geq c_2$.

In summary, we have shown that for any absolute constant $c_2 < 1$ there is some $s \geq s_0 \geq 6s_1$ such that $\mathbb{P}^{\mathbb{T}}_p(E_3^{\mathrm{crude}}) \geq c_2$.

Consider any crude state CS_0 such that E_3^{crude} holds, and let (Z_1, col_1) be a realization of (Z, col) consistent with $CS = CS_0$ and such that E_3 holds, i.e., such that there is a 4δ-robustly black path P crossing some $3s/4$ by $s/12$ rectangle in $\mathbb{T}$ horizontally. By the definition of E_3^{crude}, such a (Z_1, col_1) exists. Let (Z_2, col_2) be any other realization of Z consistent with $CS = CS_0$. In Z_1, for every point x of P there is a black point b at some distance r and no white point within distance $r + 4\delta$. In particular, the square S_i containing b contains no white points and is thus good. As Z_2 corresponds to the same crude state, there is (at least one) black point of Z_2 in S_i; this point is within distance $r+2\delta$ of x. Also, there is no white point $w \in Z_2$ within distance $r+2\delta$ of x: otherwise, the square S_j containing w is bad, and there is a white point w' of Z_1 in this square, and hence within distance $r + 4\delta$ of x, contradicting our assumption. Hence every point x of P is black with respect to Z_2.

Let E_4 be the event that there is a black path crossing *some* $3s/4$ by $s/12$ rectangle in $\mathbb{T}$ horizontally. We have shown that whenever CS is such that E_3^{crude} holds, the conditional probability, given CS, that E_4 holds is 1. Hence,

$$\mathbb{P}_p^{\mathbb{T}}(E_4) \geq \mathbb{P}_p^{\mathbb{T}}(E_3^{\text{crude}}) \geq c_2.$$

We can cover $\mathbb{T}$ with a fixed number M of $s/2$ by $s/6$ rectangles R_i so that whenever E_4 holds, there is a black path crossing some R_i horizontally. For example, we may take $M = 48$ and use rectangles whose x-coordinates are all multiples of $s/4$ and whose y-coordinates are all multiples of $s/12$. Recall that $f_p(3, s/6)$ is the probability in the random Voronoi percolation on $\mathbb{R}^2$ that an $s/2$ by $s/6$ rectangle has a horizontal black crossing. As before the corresponding probability in the torus is within $o(1)$. The events H_i that R_i has a horizontal black crossing are black-increasing. Hence, their complements H_i^c are white-increasing. Applying Lemma 10, which holds in the torus, with black and white exchanged, each H_i^c is positively correlated with $\cap_{j<i} H_j^c$, so

$$\mathbb{P}_p^{\mathbb{T}}\left(\bigcap_{i=1}^{M} H_i^c\right) \geq \prod_{i=1}^{M} \mathbb{P}_p^{\mathbb{T}}(H_i^c) = \mathbb{P}_p^{\mathbb{T}}(H_1^c)^M = (1 - f_p(3, s/6) - o(1))^M.$$

Since if no H_i holds then E_4 fails, the probability of the intersection above is at most $1 - c_2$, so

$$1 - f_p(3, s/6) \leq (1 - c_2)^{1/M} + o(1).$$

(This is sometimes known as the 'square-root' trick.) Choosing c_2 sufficiently close to 1, we see that for some $s \geq s_0 \geq 6s_1$ we have

$$f_p(3, s/6) > c_1.$$

This completes the proof of Theorem 21. $\square$

8 Proofs of the main results

Now all the technicalities are behind us. To prove our main results, we only need the ingredients in Section 2, Theorem 21 and Lemma 10.

We first prove Theorem 1, that for any $p > 1/2$ we have $\theta(p) > 0$, i.e., that percolation occurs in the black cells of the random Voronoi tiling in which cells are coloured black independently with probability p. This corresponds to Kesten's result for bond-percolation in $\mathbb{Z}^2$. It turns out that it is easy to deduce Theorem 1 from Theorem 21 using the idea of 1-dependent percolation, and in particular the simple observation given as Lemma 6.

Proof of Theorem 1. Let $p > 1/2$ be fixed. Let $p_0 < 1$ be a value for which Lemma 6 holds, and set $\varepsilon = (1 - p_0)/4$. For $s > 1$, let R_s be a $3s$ by s rectangle. Let us recall the definition of the event $E_{\text{dense}}(R_s)$ considered in Lemma 9: writing $R_s[d]$ for the set of all points within Euclidean distance d of R_s, $E_{\text{dense}}(R_s)$ is the event that for every $x \in R_s[r]$ there is some point $z \in Z$ at distance $\overline{xz} < r$, where $r = r(s) = 2\sqrt{\log s}$. By Lemma 9, there is an s_1 such that $\mathbb{P}_p(E_{\text{dense}}(R_s)) \geq 1 - \varepsilon$ for all $s \geq s_1$. Also, the event $E_{\text{dense}}(R_s)$ depends only on the restriction of Z to $R_s[2r]$. We may and shall assume that s_1 is large enough that $r = r(s) < s/4$ for $s \geq s_1$.

By Theorem 21, there is an $s \geq s_1$ such that $f_p(3, s) \geq 1 - \varepsilon$. From now on we fix such an s and suppress the dependence on s in our notation. Thus, for R a $3s$ by s rectangle with the long side horizontal, the probability of the event $H(R)$ that there is a black path within R crossing R from left to right is at least $1 - \varepsilon$. Let S_1 and S_2 be the two s by s squares making up the left and right thirds of R. As $H(R)$ implies $H(S_1)$, we have $\mathbb{P}_p(V(S_i)) = \mathbb{P}_p(H(S_i)) \geq 1 - \varepsilon$ for $i = 1, 2$.

Let

$$G(R) = H(R) \cap V(S_1) \cap V(S_2) \cap E_{\text{dense}}(R). \tag{29}$$

This event is illustrated in Figure 12.

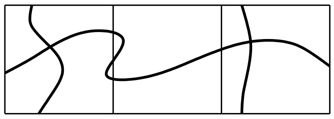

Figure 12: A $3s$ by s rectangle R such that $H(R)$, $V(S_1)$ and $V(S_2)$ hold.

For R an s by $3s$ rectangle we define $G(R)$ similarly: $G(R)$ is the event that $G(R')$ holds after rotating $\mathbb{R}^2$ through 90 degrees mapping R to R'. As each of the four events intersected in (29) has probability at least $1 - \varepsilon$, we have

$$\mathbb{P}_p(G(R)) \geq 1 - 4\varepsilon = p_0$$

for any $3s$ by s or s by $3s$ rectangle R.

We define a 1-dependent bond percolation measure $\widetilde{\mathbb{P}}$ on $\mathbb{Z}^2$ as follows: for $x, y \in \mathbb{Z}$ the edge e from (x, y) to $(x+1, y)$ is open in $\widetilde{\mathbb{P}}$ if and only if $G(R_e)$ holds in $\mathbb{P}_p$ for the $3s$ by s rectangle $R_e = [2sx, 2sx + 3s] \times [2sy, 2sy + s]$. Similarly, the edge e from (x, y) to $(x, y + 1)$ is open in $\widetilde{\mathbb{P}}$ if and only if $G(R_e)$ holds in $\mathbb{P}_p$ for the s by $3s$ rectangle $R_e = [2sx, 2sx + s] \times [2sy, 2sy + 3s]$.

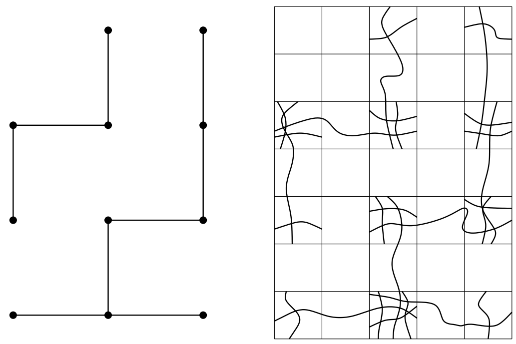

Figure 13: A set of open edges in $\widetilde{\mathbb{P}}$ (left), and corresponding rectangles R drawn with black paths showing that $G(R)$ holds in $\mathbb{P}_p$.

It is easy to check that $\widetilde{\mathbb{P}}$ is 1-dependent. Indeed, recall that the event $E_{\mathrm{dense}}(R)$ depends only on the restriction of (Z, col) to the $2r$-neighbourhood $R[2r]$ of R, and that if $E_{\mathrm{dense}}(R)$ holds, then the colours of all points in R are determined by this restriction. It follows that the event $G(R)$ depends only on this restriction. As $r < s/4$, if S and T are sets of edges of $\mathbb{Z}^2$ at graph distance at least 1, then the regions $D_U = \bigcup_{e\in U} R_e[2r]$, $U = S, T$, are disjoint. Hence the restrictions of (Z, col) to D_S and D_T are independent, implying 1-dependence of $\widetilde{\mathbb{P}}$.

By Lemma 6 we have percolation in $\widetilde{\mathbb{P}}$, i.e., with positive $\widetilde{\mathbb{P}}$-probability the origin is in an infinite open path in $\mathbb{Z}^2$. However, we have defined $G(R)$ in such a way that a $\widetilde{\mathbb{P}}$-open path in $\mathbb{Z}^2$ guarantees a corresponding (much longer) black path in the coloured Voronoi tiling of $\mathbb{R}^2$, using only the fact that horizontal and vertical crossings of a square must meet; see Figure 13. Hence, with positive $\mathbb{P}_p$-probability, some point of $[0, s] \times [0, s]$ is in an infinite path of black Voronoi cells. It follows that $\theta(p) > 0$, completing the proof of Theorem 1. $\square$

We next turn to our second main result, Theorem 2, giving exponential decay of the size (defined in any reasonable way) of C_0, the black component of the origin, when $p < 1/2$. Again the result follows easily from Theorem 21 using the idea of locally dependent percolation. This time, we use the observation given as Lemma 7.

Proof of Theorem 2. The basic idea is to divide $\mathbb{R}^2$ up into s by s squares, and to show that for each square S it is very unlikely that any point of S is connected to any point far outside S. Then we shall apply Lemma 7.

Let S be an s by s square. We can surround S with a ring of four overlapping $3s$ by s (or s by $3s$) rectangles R_i, $1 \le i \le 4$, as shown in Figure 14.

Let $\overline{S} = S \cup \bigcup_{i=1}^4 R_i$ be the L_∞ s-neighbourhood of S. Suppose that each R_i is crossed the long way by a *white* path P_i. (As before, this means that there is a piecewise linear path of white points inside R_i starting and ending

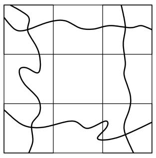

Figure 14: A square S surrounded by four rectangles R_i, containing four white paths P_i (drawn in black) whose union must contain a cycle surrounding S.

on appropriate sides of R_i.) Then the paths P_i meet pairwise, so their union contains a piecewise-linear white cycle C surrounding S as in Figure 14. Given the existence of such a cycle C, it follows with probability 1 that no point $x \in S$ is connected to any point x' outside $\overline{S}$ by a black path B. (Recall that a point may be both black and white, if it is in two Voronoi cells of different colours. But this is not a problem – with probability 1 the Voronoi cells meet three at a vertex. Corresponding to the white cycle C and to the black path B there is a cycle/path of white/black Voronoi cells, and these cannot cross.)

Let $L(S)$ be the event that some point in S is connected by a black path to a point in the boundary of $\overline{S}$. As each R_i is crossed the long way by a white path with probability $f_{1-p}(3,s)$, and these events are white-increasing and hence positively correlated by Lemma 10, we find a white cycle C as above with probability at least $f_{1-p}(3,s)^4$, and

$$\mathbb{P}_p(L(S)) \leq 1 - f_{1-p}(3,s)^4. \tag{30}$$

Let $G(S) = L(S) \cup E_{\mathrm{dense}}(\overline{S})^c$, and set $\varepsilon = p_1(7)/2$, where $p_1(7)$ is the quantity appearing in Lemma 7 for $k = 7$, i.e., for 7-dependent site percolation. Applying Lemma 9 to the $3s$ by $3s$ square $\overline{S}$, we have $\mathbb{P}_p(E_{\mathrm{dense}}(\overline{S})) \to 1$ as $s \to \infty$, so if s is large enough, then $\mathbb{P}_p(E_{\mathrm{dense}}(\overline{S})^c) < \varepsilon$. Also, by Theorem 21 and (30), there are arbitrarily large s for which $\mathbb{P}_p(L(S)) < \varepsilon$. It follows that there is some s for which

$$\mathbb{P}_p(G(S)) < 2\varepsilon = p_1(7). \tag{31}$$

Let us fix such an s, chosen large enough that $r = r(3s) = 2\sqrt{\log(3s)} < s/4$.

As before, by Lemma 9, $E_{\mathrm{dense}}(\overline{S})$ depends only on the restriction of (Z, col) to the (usual) $2r$-neighbourhood $\overline{S}[2r]$ of $\overline{S}$. If $E_{\mathrm{dense}}(\overline{S})$ does not hold, then $G(S)$ holds. On the other hand, if $E_{\mathrm{dense}}(\overline{S})$ does hold, then $G(S)$ holds if and only if $L(S)$ holds, which depends only on the colours of points of $\overline{S}[2r]$. Hence, $G(S)$ depends only on the restriction of (Z, col) to $\overline{S}[2r]$. To each $v = (a,b) \in \mathbb{Z}^2$ let us assign an s by s square $S_v = [as, as+s] \times [bs, bs+s]$. We say that $v \in \mathbb{Z}^2$ is *open* if $G(S)$ holds. We claim that this assignment of states to vertices defines a 7-dependent site percolation measure $\widetilde{\mathbb{P}}$ on $\mathbb{Z}^2$. Indeed, if v, $w \in \mathbb{Z}^2$ are at graph distance at least 7, then either their x-coordinates or their y-coordinates differ by at least 4, so the sets $\overline{S}_v[2r]$ and $\overline{S}_w[2r]$ are disjoint. For A, $B \subset \mathbb{Z}^2$ at distance at least 7, which vertices of A are open and which of B depend respectively on the restrictions of (Z, col) to the domains D_A, D_B, where $D_C = \bigcup_{v \in C} \overline{S}_v[2r]$, so D_A and D_B are disjoint. Hence the states of vertices in A are independent of the sates of the vertices in B, proving 7-dependence.

Let C'_0 be the set of vertices $v \in \mathbb{Z}^2$ in the open cluster of the origin in the site percolation $\widetilde{\mathbb{P}}$ on $\mathbb{Z}^2$ we have just defined. From (31), each vertex v is open with probability at most $p_1(7)$, so by Lemma 7 there is some $c > 0$ such that

$$\mathbb{P}_p(|C'_0| \geq m) = \widetilde{\mathbb{P}}(|C'_0| \geq m) \leq \exp(-cm)$$

for all $m \geq 1$.

Now let C_0 be the black component of the origin in the random Voronoi percolation, i.e., the maximal connected union of black Voronoi cells containing the origin (which is empty if the origin is white). Let $N = |C_0^G|$ be the number of Voronoi cells in C_0. We consider two cases. The first is that C_0 is contained in a disc D of radius $10s$ centred at the origin. Note that in this case area(C_0) and diam(C_0) are bounded by absolute constants. It is easy to check that the number of Voronoi cells contained entirely in D decays exponentially (the number is dominated by the number of points of Z in the disc, a Poisson distribution with a certain fixed expectation). Hence, there is a $c' > 0$ such that

$$\mathbb{P}\big(C_0 \subset D \text{ and } N \geq n\big) \leq \exp(-c'n) \tag{32}$$

for all n.

On the other hand, if $C_0 \not\subset D$, then C_0 has geometric diameter at least $10s$, and for every point $x \in C_0$ there is another point $x' \in C_0$ at distance at least $5s$. It follows that if $x \in C_0 \cap S_v$, then $G(S_v)$ holds; indeed, x is a point of S_v connected by a black path in C_0 to a point x' outside $\overline{S}_v$. Stopping the first time the path hits the boundary of $\overline{S}_v$, we see that $L(S_v)$ holds, and hence so does $G(S_v)$. Let $C_1 = \{v \in \mathbb{Z}^2 : C_0 \cap S_v \neq \emptyset\}$. We have shown that if $C_0 \not\subset D$, then $G(S_v)$ holds for every $v \in C_1$, i.e., that every $v \in C_1$ is open. But C_1 is with probability 1 a connected set in $\mathbb{Z}^2$ that, if nonempty, contains the origin, so $C_1 \subset C'_0$, and we have

$$\mathbb{P}(|C_1| \geq m) \leq \exp(-cm) \tag{33}$$

for all m, with the same $c > 0$ as above. As the area of C_0 is at most $s^2|C_1|$, this establishes exponential decay of the area of C_0. As diam$(C_0) \leq \sqrt{2}s + s(|C_1| - 1)$, say, exponential decay of diam(C_0) also follows.

Finally, to show exponential decay of the number $N = |C_0^G|$ of Voronoi cells in C_0 note that, given that $|C_1| = m$, there are are most $(4e)^m$ possibilities for C_1. As every cell in C_0 is contained entirely in the union U of the m squares S_v, $v \in C_1$, the number of cells in C_0 is at most the number of points of Z in U, which has area $s^2 m$. For any possible C_1, the probability that the corresponding U contains more than $100s^2m$ points of Z is certainly at most $\exp(-10s^2m) \leq \exp(-10m)$. (Recall that s is a large constant.) Multiplying by the number of possibilities for C_1, we see that

$$\mathbb{P}\big(C_0 \not\subset D \text{ and } N \geq 100s^2m \text{ and } |C_1| = m\big) \leq \exp(-5m).$$

From (33) it follows that

$$\mathbb{P}\big(C_0 \not\subset D \text{ and } N \geq 100s^2m\big) \leq \exp(-5m) + \exp(-cm).$$

Combined with (32) this establishes exponential decay of $N = |C_0^G|$, completing the proof. □

As noted in Section 1, our proofs of Theorems 1 and 2 also give a new proof of Theorem 3, Zvavitch's result that $\theta(1/2) = 0$. As the deduction of Theorem 3 from our intermediate results is rather short, we give it here.

Proof of Theorem 3. The result follows almost immediately from Corollary 13, applied in a similar way to the Russo-Seymour-Welsh Theorem for $\mathbb{Z}^2$. Let A_s be the square annulus centered on the origin, with inner and outer diameters $8s$ and $10s$. Let $C(A_s)$ be the event that A_s contains a (piecewise-linear) white cycle surrounding the origin. Arguing as in the proof of Theorem 2, from Corollary 13, applied with $\rho = 10$ and with white and black swapped, and positive correlation (Lemma 10), there is an absolute constant $\varepsilon = c_0(10)^4/2 > 0$ and an unbounded set $\mathcal{S}$ of values of s for which $\mathbb{P}_{1/2}(C(A_s)) \geq c_0(10)^4 = 2\varepsilon$. Let $L(A_s)$ be the event that $E_{\mathrm{dense}}(R)$ holds for each of the four rectangles R making up A_s. By Lemma 9, $\mathbb{P}_{1/2}(E_{\mathrm{dense}}(R)) \to 1$ as $s \to \infty$, so $\mathbb{P}_{1/2}(L(A_s)) \geq 1 - \varepsilon$ if s is large enough. Thus, deleting small values from $\mathcal{S}$ if necessary, the event $G(A_s) = C(A_s) \cap L(A_s)$ holds with probability at least ε for every $s \in \mathcal{S}$. Finally, as in the proof of Theorem 2, it follows from Lemma 9 that $G(A_s)$ depends only on the restriction of (Z, col) to the $4\sqrt{\log(10s)}$ neighbourhood N_s of A_s. If $s_2 \geq 2s_1$ and s_1 is large enough, then N_{s_1} and N_{s_2} are disjoint. As $\mathcal{S}$ is unbounded, we can pick an infinite sequence s_i, $i = 1, 2, \ldots$, of values from $\mathcal{S}$ with s_1 large enough and $s_{i+1} \geq 2s_i$ for every i. Then the events $G(A_{s_i})$ are independent. As each has probability at least ε, with probability 1 some $G(A_{s_i})$ holds. But then $C(A_{s_i})$ holds, and the origin is surrounded by a white cycle, so the black component C_0 of the origin is bounded. This proves that $\theta(1/2) = 0$. □

The arguments given in this section were very simple. Indeed, once one thinks of using k-dependent percolation measures, the main results are more or less immediate from Theorem 21. Also, the basic idea needed to derive Theorem 21 from Theorem 12, our analogue of the RSW Theorem, is very simple: construct a suitable symmetric event and apply a sharp-threshold result such as Theorem 5. Together, these ideas give a very simple proof of, for example, Kesten's results on the critical probability for bond percolation in $\mathbb{Z}^2$; we shall return to this in future work [7, 8].

9 Appendix

In this appendix we give the proofs of several deterministic statements used in earlier sections. Some of these stand alone and will be repeated here; others only make sense in the context in which they arise.

9.1 Proof of Lemma 16

We start with a little geometric lemma concerning distances in $\mathbb{R}^2$. As before, for two points A, B of $\mathbb{R}^2$ we write $\overline{AB}$ for the Euclidean distance between them.

Lemma 22. *Let C_1 and C_2 be two circles with centres O_1, O_2 and radii r_1, r_2, with $r_1, r_2 \leq 1/\delta$, where $0 < \delta < 1/100$. Suppose that C_1, C_2 meet in two points P_1, P_2 with $\overline{P_1P_2} \geq \delta$, and that $\overline{O_1O_2} \geq \delta$. Let C_3, C_4 be the circles with centres P_1, P_2 and radii $\delta/4$, and let M be the midpoint of the line-segment O_1O_2. Then every point outside $\cup_{i=1}^4 C_i$ is at distance at least $\overline{P_1M} + \delta^6$ from M.*

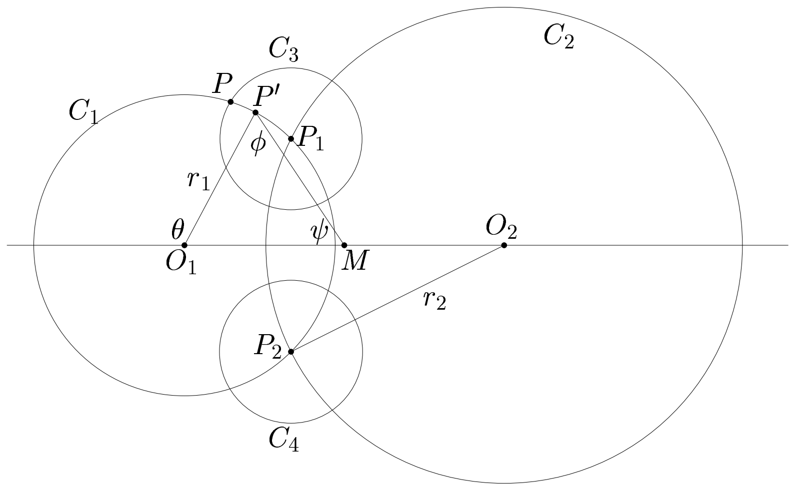

Figure 15: An illustration of the notation in Lemma 22.

Proof. We may assume that the centres of C_1, C_2 are $(-a, 0)$ and $(a, 0)$, $a \geq \delta/2$, so M is the origin, and that $P_1 = (c, d)$ with $d > 0$. In fact, as $\overline{P_1P_2} \geq \delta$, we have $d \geq \delta/2$, so C_3 lies entirely in the upper half plane. Consider a point $P = (x, y)$ not in the interior of any C_i at minimal distance from M. By symmetry we may assume $y \geq 0$.

Note that M lies inside at least one of C_1, C_2, as otherwise these circles cannot meet in two points. Hence $P \neq M$, so P must lie on at least one of the C_i. By minimality of $\overline{PM}$, either P is the closest point of some C_i to M, or P lies on at least two C_i (otherwise, there is a point very close to P on the same circle that is closer to M). The first case is easy to exclude: let Q_i be the closest point of C_i to M. Whether or not M lies inside C_1, Q_1 is the rightmost point of C_1, i.e., the point $(-a + r_1, 0)$. This lies strictly inside C_2, as otherwise C_1 and C_2 do not meet in two points. So we cannot have $P = Q_1$. Similarly, $P \neq Q_2$. Now Q_3 is a point strictly between P_1 and M along the line-segment P_1M. (Recall that C_3 lies entirely in the upper half plane, so M lies outside

C_3.) Now M is inside C_i for $i = 1$ or $i = 2$, and P_1 is on the boundary of this C_i, so Q_3 is inside C_i and we cannot have $P = Q_3$.

It follows that P lies on two of the circles C_i. As C_4 lies in the lower half plane and, clearly, $P \neq P_1, P_2$, we have without loss of generality that P lies on C_1 and C_3. Note that P lies anticlockwise around the upper-half of C_1 from P_1; the upper half of C_1 on the other side of P_1 lies inside C_2. Let us consider moving a point $P' = (x', y')$ along C_1, starting from P_1 and moving anticlockwise towards P. As we move, since $y' \geq 0$, the distance $\overline{P'M}$ from M increases monotonically, and it suffices to show that it increases by at least δ^6. To get to P we must move a distance at least $\delta/4$ along C_1; note that $y' \geq d - \delta/4 \geq \delta/4$ holds throughout.

Now consider the angles $\phi = O_1P'M$, $\psi = O_1MP'$ and θ shown in Figure 15. Note that if $a = r_1$, i.e., if M is on C_1, then $\phi = \psi = \theta/2$. Hence if $a > r_1$, then $\phi \geq \theta/2$. On the other hand, if $a < r_1$, $\psi \geq \theta/2$, while by the sine rule, $\sin\phi = (a/r_1)\sin\psi \geq (a/r_1)\sin(\theta/2)$. Using $a \geq \delta/2$, $r_1 \leq 1/\delta$, it follows that in either case

$$\sin\phi \geq \delta^2 \sin(\theta/2)/2.$$

Now $y' \geq \delta/4$ while $r_1 \leq 1/\delta$, so $\sin(\theta/2) \geq \sin\theta/2 \geq \delta^2/8$, and $\sin\phi \geq \delta^4/16$.

Finally, the rate at which the distance $\overline{P'M}$ increases as we move P' at rate 1 around C_1 is exactly $\sin\phi$, so after moving distance $\delta/4$ starting at $P' = P_1$ we have increased $\overline{P'M}$ by at least $\delta^5/64 \geq \delta^6$. Thus $\overline{PM} \geq \overline{P_1M} + \delta^6$, completing the proof. $\square$

It is now easy to deduce Lemma 16. We recall the statement.

Lemma 16. *Let $s \geq e^{100}$ and $\delta < 1/\log s$ be given, and suppose that no disc of radius $2\sqrt{\log s}$ in $\mathbb{T}$ is free of points of Z. Let z_1, $z_2 \in Z$ be points whose Voronoi cells meet, and let M be the midpoint of their common edge. If some $z \in Z \setminus \{z_1, z_2\}$ is at distance less than $\overline{z_1M} + \delta^6$ from M, then either $\{z_1, z_2\}$ is a δ-close pair, or there is a δ-bad quadruple $\{z_1, z_2, z, z'\}$, $z' \in Z$.*

Proof of Lemma 16. Note first that the assumption concerning point-free discs ensures that all Voronoi cells have diameter much less than s. In particular, they are all contractible in the torus, and are hence convex polygons. Thus the cells of z_1 and z_2 do meet in a line-segment O_1O_2, say. Furthermore, there are distinct points x_1, $x_2 \in Z \setminus \{z_1, z_2\}$ so that O_1 is a vertex of the cell of x_1 and O_2 of the cell of x_2. Thus, $\overline{O_1z_1} = \overline{O_1z_2} = \overline{O_1x_1} = r_1$, say, and $\overline{O_2z_1} = \overline{O_2z_2} = \overline{O_2x_2} = r_2$, say, and no point of Z is closer to O_i than distance r_i. The notation is shown in Figure 16. As $x_1 \neq x_2$, we may assume without loss of generality that $z \neq x_1$.

Suppose first that $\overline{O_1O_2} < \delta$. Let $r = \overline{z_1M} = \overline{z_2M}$. Using the triangle inequality twice, $\left|\overline{x_1M} - \overline{z_1M}\right| \leq \left|\overline{x_1O_1} - \overline{z_1O_1}\right| + 2\overline{O_1M} = 0 + \overline{O_1O_2} < \delta$. By assumption, $\overline{zM} < r + \delta^6$. As M is in the Voronoi cells of both z_1 and z_2, these are the two closest points of Z to M, so no point of Z is at distance less than r from M. As $\overline{z_1M}, \overline{z_2M}, \overline{x_1M}, \overline{zM} < r + \delta$, it follows that $\{z_1, z_2, x_1, z\}$ is a δ-bad quadruple. (The assumption on point free discs ensures that $r < 2\sqrt{\log s}$.) We may thus suppose that $\overline{O_1O_2} \geq \delta$. We may also assume that $\overline{z_1z_2} \geq \delta$; otherwise there is nothing to prove.

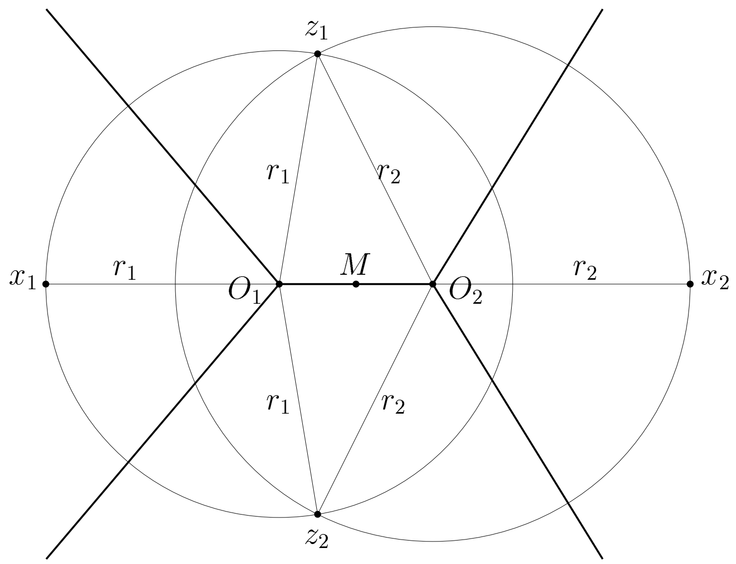

Figure 16: The region where four Voronoi cells $V_Z(z_1)$, $V_Z(z_2)$, $V_Z(x_1)$ and $V_Z(x_2)$ come together. The cells associated to z_1 and z_2 meet in the line segment O_1O_2 with midpoint M. The thick lines are parts of the boundaries of the Voronoi cells.

We shall now apply Lemma 22. There is no problem applying this lemma in the torus: by our assumption on point free discs, $r_1, r_2 \leq 2\sqrt{\log s}$, so all points we consider are within distance $10\sqrt{\log s} < s/3$ of z_1, say, and we need only consider a region of $\mathbb{T}(s)$ isomorphic to the corresponding region of $\mathbb{R}^2$. We apply Lemma 22 to the circles C_i with centres O_i and radii r_i, which meet at the points z_1, z_2. The conditions do apply; we are assuming $\overline{z_1z_2} \geq \delta$ and $\overline{O_1O_2} \geq \delta$, while, as just noted, $r_i \leq 2\sqrt{\log s} \leq 1/\delta$.

As C_1 and C_2 have no points of Z in their interiors, we conclude from Lemma 22 that z must be $\delta/4$-close to z_1 or z_2. But then z_1, z_2 and x_1 are at distance exactly r_1 from O_1, while z, which is distinct from these three points, can be at distance at most $r_1 + \delta/4$. Hence, as C_1 is free of points of Z, the quadruple $\{z_1, z_2, x_1, z\}$ is δ-bad as required. $\square$

9.2 Proof of Lemma 19

Lemma 19. *Let $0 < \eta < 1/100$ be fixed. If t is large enough, then it is impossible to arrange $M \geq \eta^{-6}t$ points $P_1, \ldots, P_M$ and M associated circles $C_1, \ldots, C_M$ in $\mathbb{R}^2$ so that no two P_i are within distance η of each other, the distance of every P_i from the origin is between ηt and t, each C_i has radius at most t, no C_i contains any P_j, and C_i passes within distance η^3 of both P_i and the origin.*

Proof of Lemma 19. Consider any arrangement as described in the statement of the lemma. By assumption, all the P_i lie in an annulus of inner radius ηt and outer radius t centred on the origin O. We may write this annulus as a union of fewer than η^{-2} annuli A_i, each of which has outer radius $(1+\eta)$ times its inner radius. It suffices to show that we cannot arrange at least $\eta^2 M$ points as described in any one of the A_i. Suppose we can. Then there are two P_i, say P_1 and P_2, such that the angle P_1OP_2 is at most $\theta_0 = 2\pi/(\eta^2 M) \le \eta^3/t$. Suppose that $\overline{OP_1} \ge \overline{OP_2}$. Let us rotate the coordinates so that P_1 is the point $(0, y)$, $y > 0$, and P_2 the point (x, y'), $y' < y$. As $\overline{OP_2}$ is at most t, we see that $|x| \le t\theta_0 \le \eta^3$. Using $\overline{P_1P_2} \ge \eta$, it follows that $|y - y'| \ge \eta/2$.

Now it is easy to see that C_1 must contain P_2: otherwise, translating through a distance of at most η^3 so that the image C_1' of C_1 passes exactly through the origin, the image P_2' of P_2 lies outside C_1' and satisfies $\overline{P_2P_2'} \le \eta^3$. Also, there is a point P_1' (close to the image of P_1) which lies exactly on C_1', with $\overline{P_1P_1'} \le 2\eta^3$. Rotating again, we may write $P_1' = (0, y_1)$, $P_2' = (x_2, y_2)$ with, crudely, $\eta t/2 < y_2 < y_1$, $y_1 - y_2 \ge \eta/3$ and $|x_2| \le 5\eta^3$.

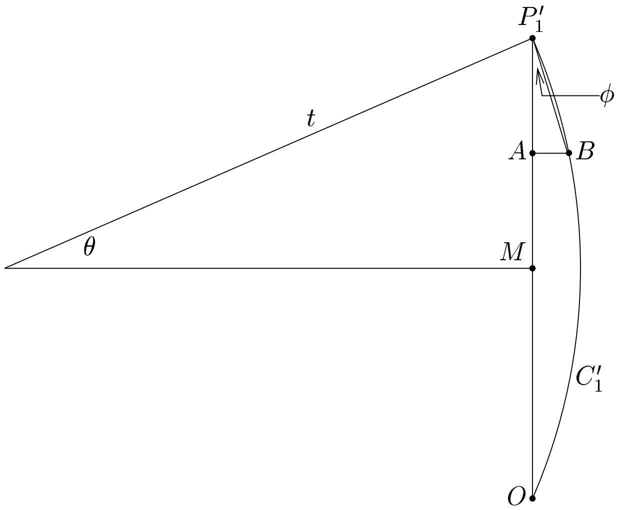

Figure 17: The notation in the proof of Lemma 19.

As the circle C_1' passes through the origin and $P_1' = (0, y_1)$, it contains $A = (0, y_2)$ and certain other points on the line $y = y_2$. The extreme case is when the circle has maximum radius, t, by assumption. This case is shown in Figure 17. Let B be the unique point (x_3, y_2) on both C_1' and the line $y = y_2$ for which $x_3 > 0$. As $\overline{OP_1'} = y_1 \ge \eta t/2$, writing M for the midpoint of OP_1 we have $\overline{MP_1'} \ge \eta t/4$, so the angle θ shown is at least $\arcsin(\eta/4)$. By assumption $\overline{OP_1} \le (1+\eta)\overline{OP_2}$, as P_1 and P_2 lie in the same annulus A_i. It follows that, crudely, $y_2 > y_1/2$, so the angle $\phi = BP_1'A$ shown in Figure 17 is at least $\theta/2$, the value it takes when $y_2 = y_1/2$, i.e., when $A = M$. Thus

$$x_3 = \overline{AB} = (y_1 - y_2)\tan\phi \ge (\eta/3)\sin\phi \ge (\eta/3)\sin(\arcsin(\eta/4)/2) \ge \eta^2/24.$$

As the case $r = t$ was extreme, in all cases C_1' contains all points (x, y_2) with $|x| \leq \eta^2/24$. In particular, C_1' contains P_2'. But then C_1 contains P_2, contradicting our assumptions. $\square$

9.3 Proof of Claim 18.1

Claim 18.1, repeated below, is part of the proof of Lemma 18. The context is as follows: $\varepsilon > 0$ is fixed, and we have $\delta \leq s^{-\varepsilon}$ and $\delta' = \delta^{1/10}$. The square S is a square of side length $\ell = 3(\log s)^2 = O(1/\delta')$ in the torus $\mathbb{T}(s)$. Recall that a quadruple $\{x_1, x_2, x_3, x_4\}$ is *weakly δ-bad* if there is a point x and a radius $r < 2\sqrt{\log s}$ such that $r < \overline{xx_i} \leq r + \delta$ holds for $1 \leq i \leq 4$. Note that this definition makes no reference to the Poisson process Z. Finally, all limits are as $s \to \infty$.

Claim 18.1. *Suppose that x_1, x_2, and x_3 are points of S, no two within distance δ' of each other. Let B be the set of $x \in S$ for which the quadruple $\{x_1, x_2, x_3, x\}$ is weakly δ-bad. Then* $\mathrm{area}(B) = o(\delta')$.

Proof of Claim 18.1. Although in the proof of Lemma 18 we are working in the torus $\mathbb{T}(s)$, here we shall work in $\mathbb{R}^2$, treating S as a square in $\mathbb{R}^2$ of side-length $\ell = O(1/\delta')$. We may do this because of the restriction $r < 2\sqrt{\log s}$ in the definition of weak δ-badness, which ensures that no circle we need consider comes close to 'wrapping round' the torus.

To avoid degeneracy we shall assume that x_1, x_2 and x_3 are not collinear. (This is valid as we may shift one slightly and increase δ slightly.)

We say that four points of $\mathbb{R}^2$ are *concyclic* if there is a circle passing through them. Suppose that $x \in B$, i.e., that the quadruple $\{x_1, x_2, x_3, x\}$ is weakly δ-bad. Then there are points $y_i \in \mathbb{R}^2$ such that y_1, y_2, y_3 and x are concyclic, with $\overline{x_i y_i} \leq \delta$. (Fixing the centre, adjust the radius of the circle witnessing weak δ-badness by at most δ so that this circle passes through x.) As the statement of the claim depends only on the set $\{x_1, x_2, x_3\}$, we may renumber so that y_1, y_2, y_3 and x appear in this order around the circle on which they lie.

Consider the linear transformation $T : \mathbb{R}^2 \to \mathbb{R}^2$ obtained as follows: first translate $\mathbb{R}^2$ so that the point y_1 is mapped to x_1. Then rotate and rescale around x_1 so that (the image of) y_3 is mapped to x_3. As $\overline{x_i y_i} \leq \delta$, while $\overline{x_1 x_3} \geq \delta'$, the angle of rotation is $O(\delta/\delta')$, and the scale factor is $1 + O(\delta/\delta')$. It follows that T, which maps x_1 to y_1 and x_3 to y_3, moves each point inside S, or indeed within distance ℓ of S, through a distance of at most $O(\ell\delta/\delta') = O({\delta'}^8)$. Hence,

$$\overline{T(y_2)x_2} \leq \overline{T(y_2)y_2} + \overline{y_2 x_2} = O({\delta'}^8) + O(\delta) = O({\delta'}^8). \tag{34}$$

Applying T to $\{y_1, y_2, y_3, x\}$ we see that x_1, $T(y_2)$, x_3 and $T(x)$ are concyclic, and appear in this order around the circle on which they lie. Let P be the intersection of the line-segments $x_1 x_3$ and $T(y_2)T(x)$, so P lies in the interior of these two line-segments; see Figure 18. Then

$$\overline{x_1 P}\ \overline{x_3 P} = \overline{T(y_2)P}\ \overline{T(x)P}. \tag{35}$$

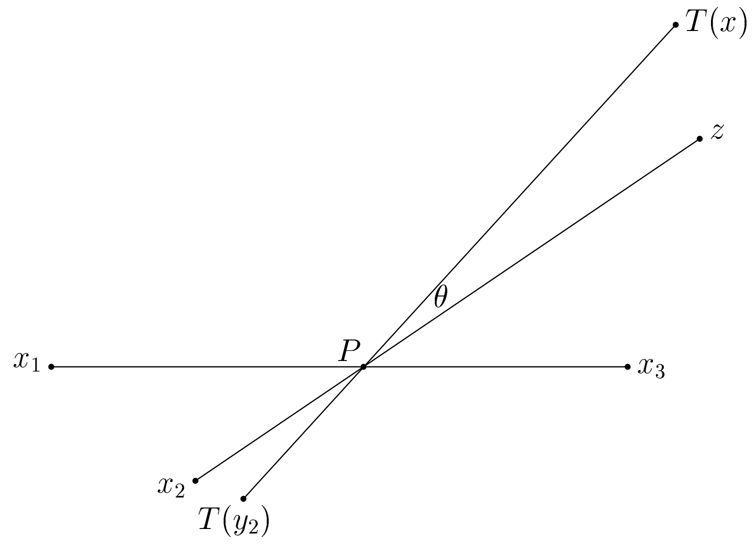

Figure 18: The points $x_1, T(y_2), x_3, T(x)$ are concyclic, as are the points x_1, x_2, x_3, z.

We claim that $\overline{T(y_2)P} \geq \delta'^3$: otherwise, as x and $T(y_2)$ lie in S, $\overline{T(x)P} \leq \overline{T(x)T(y_2)} \leq 2\ell$, so from (35), $\overline{x_1P}\ \overline{x_3P} \leq 2\ell\delta'^3$. Assuming without loss of generality that $\overline{x_1P} \leq \overline{x_3P}$, as $\overline{x_1x_3} \geq \delta'$ by assumption, we have $\overline{x_3P} \geq \delta'/2$. It follows that $\overline{x_1P} \leq 4\ell\delta'^2$. But then using (34), if s is large enough we have

$$\overline{x_1x_2} \leq \overline{x_1P} + \overline{PT(y_2)} + \overline{T(y_2)x_2} \leq 4\ell\delta'^2 + \delta'^3 + O(\delta'^8) < \delta',$$

contradicting one of our assumptions.

Let C be the circle through x_1, x_2 and x_3, and let z be the point on C obtained by extending the line-segment x_2P. Let θ be the angle $T(y_2)Px_2 = T(x)Pz$ as shown in Figure 18. Now from (34) and the fact that $\overline{T(y_2)P} \geq \delta'^3$, we see that $\theta = O(\delta'^5)$. Also, $\overline{x_2P}/\overline{T(y_2)P} = 1 + O(\delta'^5)$. As x_1, x_2, x_3 and z are concyclic,

$$\overline{zP} = \frac{\overline{x_1P}\ \overline{x_3P}}{\overline{x_2P}} = \overline{T(x)P}\frac{\overline{T(y_2)P}}{\overline{x_2P}} = \overline{T(x)P}(1 + O(\delta'^5)) = \overline{T(x)P} + O(\ell\delta'^5).$$

As $\overline{T(x)P} = O(\ell)$ and $\theta = O(\delta'^5)$, it follows that $\overline{T(x)z} = O(\ell\delta'^5)$. Thus, as $\overline{T(x)x} = O(\delta'^8)$ from the properties of T, we have $\overline{xz} = O(\ell\delta'^5)$.

In summary, we have shown that any point x for which $\{x_1, x_2, x_3, x\}$ is weakly δ-bad is within distance $O(\ell\delta'^5)$ of some point z on the circle through $\{x_1, x_2, x_3\}$. Thus, such an x lies in a certain fixed annulus whose inner and outer radii differ by $O(\ell\delta'^5)$. Any such annulus meets S in a total area of $O(\ell^2\delta'^5)$, completing the proof. $\square$

9.4 Proof of Claim 20.1

Although purely deterministic, Claim 20.1 is part of the proof of Lemma 20, and the proof below should be read in this context. We refer the reader to page 34 for the statement of Claim 20.1, and to page 34 for the three cases describing the properties of the coupling constructed for a particular potentially bad component C.

Proof of Claim 20.1. In the context of the proof of Lemma 20, recall that we have constructed the coupling of (Z_1, col_1) and (Z_2, col_2) in such a way that one of the three cases 1,2,3 defined on page 34 holds. As $a, b \in C$ and $z_{1,a}$ and $z_{1,b}$ are black, case 3 does not hold, and we must consider only cases 1 and 2. Unfortunately, within each of these cases we shall need to divide further into the two cases (i) that $z_{1,a}$ and $z_{1,b}$ are not δ_2-close, and (ii) that they are. Let us write M for the midpoint of the common edge of the Voronoi cells of $z_{1,a}$ and $z_{1,b}$ (with respect to Z_1).

Suppose first that case 1(i) holds. Then as $z_{1,a}$ and $z_{1,b}$ are not δ_2-close, for any $c \notin C$ we have by Lemma 16 that $\overline{z_{1,c}M} \geq \overline{z_{1,a}M} + \delta_2^6$; otherwise $z_{1,c}$ would be in a δ_2-bad quadruple with $z_{1,a}$ and $z_{1,b}$. (The assumption of Lemma 16 on point free discs is guaranteed by our present assumption that B_1 does not hold; see page 32.) Hence, by definition of potential badness, c would be in a potentially bad quadruple with a and b, contradicting $a \in C$, $c \notin C$. Passing from Z_1 to Z_2, as we are in case 1 only points outside C may move (and then by at most $\sqrt{2}\delta_1 < \delta_2^6$), so $z_{2,a}$ and $z_{2,b}$ are the closest points of Z_2 to M, and by Lemma 8 their Voronoi cells meet (indeed, meet at M). As $z_{2,a}$ and $z_{2,b}$ are black and δ-good with respect to Z_2, this proves the claim in this case.

Now suppose that case 2(i) holds. Here all points of C are black and δ-good with respect to Z_2, so we need only show that there is a path joining the cells of $z_{2,a}$ and $z_{2,b}$ staying within cells associated to the component. We claim that the path $P = z_{2,a}Mz_{2,b}$ will do: as above, all points $z_{1,c}$, $c \notin C$, are at least $\delta_2^6 \geq 2\sqrt{2}\delta_1$ further from M than $z_{1,a}$ and $z_{1,b}$ are. Passing to Z_2, it follows that $z_{2,a}$ and $z_{2,b}$ are both closer to M than any point $z_{2,c}$, $c \notin C$. Hence, as any point $x \in P$ can be reached from M by moving in a straight line towards $z_{2,a}$ or $z_{2,b}$, for any $x \in P$ one of $z_{2,a}$, $z_{2,b}$ is closer to x than any $z_{2,c}$, $c \notin C$. Thus the closest point of Z_2 to x is some $z_{2,d}$, $d \in C$, as required.

The case 2(ii) is similar but simpler: now $z_{2,a}$ and $z_{2,b}$ are at distance at most $\delta_2 + 2\sqrt{2}\delta_1$, and we can use the path $P = z_{2,a}z_{2,b}$: letting M' be the midpoint of this line-segment, it suffices to show that no $z_{2,c}$, $c \notin C$, is closer to M' than $z_{2,a}$, $z_{2,b}$ are. But if some $z_{2,c}$ were, it would be δ_2-close to $z_{2,a}$ or $z_{2,b}$, contradicting $a, b \in C$, $c \notin C$.

Finally, we have the case 1(ii). As in 1(i), it suffices to show that the Voronoi cells of $z_{2,a}$ and $z_{2,b}$ meet. They certainly do unless some $z_{2,c}$ is closer to M', the midpoint of $z_{1,a}z_{1,b} = z_{2,a}z_{2,b}$, than $z_{2,a}$ is, so we may suppose that there is such a c. As in case 2(ii) above, for any such c we must have $c \in C$.

Consider the common edge e of the Voronoi cells of $z_{1,a}$ and $z_{1,b}$ with respect to Z_1. The location of $z_{1,c} = z_{2,c}$ ensures that M' is not a point of this edge,

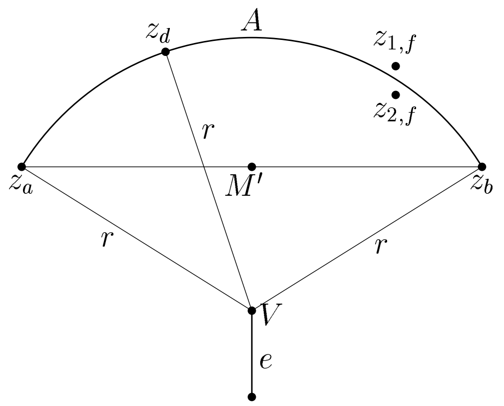

Figure 19: The points z_i represent $z_{1,i} = z_{2,i}$.

so the edge lies entirely to one side of M'. Let V be the endpoint of e closest to M', and let $z_{1,d}$ be the third point of Z_1 at the same distance r from V as $z_{1,a}$ and $z_{1,b}$. Then (as V is the closer end of e to M'), $z_{1,d}$ lies on an arc-segment A joining $z_{1,a}$ and $z_{1,b}$ every point of which is within distance $\overline{z_{1,a}M'}$ of M'; see Figure 19. Hence $d \in C$, as above. Passing to Z_2, the Voronoi cells of $z_{2,a} = z_{1,a}$ and $z_{2,b} = z_{1,b}$ with respect to Z_2 meet at V unless for some f we have $\overline{z_{2,f}V} < \overline{z_{2,a}V} = \overline{z_{2,b}V} = r$. Since only points outside C move, and then only by a distance at most $\sqrt{2}\delta_1$, we have $f \notin C$, so a, b, d and f are all distinct. Furthermore, $\overline{z_{2,f}V} \geq r - \sqrt{2}\delta_1$ for any such f. The quadruple $\{z_{2,a}, z_{2,b}, z_{2,d}, z_{2,f}\}$ is thus δ_2-bad, as witnessed by the annulus with centre V, inner radius $r' = r - \sqrt{2}\delta_1$, and outer radius $r' + \delta_2 > r$, in which all four points lie. Note that the condition $r' < 2\sqrt{\log s}$ required for δ_2-badness holds as our assumption that B_1 holds guarantees that $r < 2\sqrt{\log s}$. As $\{z_{2,a}, z_{2,b}, z_{2,d}, z_{2,f}\}$ is δ_2-bad, the quadruple $\{a, b, d, f\}$ is potentially bad, contradicting $a \in C$, $f \notin C$ and completing the proof. $\square$

References

[1] M. Aizenman, Scaling limit for the incipient spanning clusters, in Mathematics of multiscale materials (Minneapolis, MN, 1995–1996), *IMA Vol. Math. Appl.* **99** (1998), 1–24.

[2] K.S. Alexander, The RSW theorem for continuum percolation and the CLT for Euclidean minimal spanning trees, *Ann. Appl. Probab.*, **6** (1996), 466-494.

[3] R. Ahlswede and D.E. Daykin, An inequality for the weights of two families of sets, their unions and intersections, *Z. Wahrsch. Verw. Gebiete* **43** (1978), 183–185.

[4] P. Balister, B. Bollobás and M. Walters, Continuum percolation with steps in the square or the disc, to appear in *Random Structures and Algorithms*.

[5] P. Balister, B. Bollobás and A. Quas, Percolation in Voronoi tilings, to appear in *Random Structures and Algorithms*. Preprint available from `www.msci.memphis.edu/preprint/bbq2.pdf`.

[6] I. Benjamini and O. Schramm, Conformal invariance of Voronoi percolation, *Commun. Math. Phys.* **197** (1998), 75–107.

[7] B. Bollobás and O.M. Riordan, A short proof of the Harris-Kesten Theorem. Preprint available from `http://www.dpmms.cam.ac.uk/~omr10/`.

[8] B. Bollobás and O.M. Riordan, Sharp thresholds and percolation in the plane, in preparation.

[9] J. Bourgain, J. Kahn, G. Kalai, Y. Katznelson and N. Linial, The influence of variables in product spaces, *Israel J. Math.* **77** (1992), 55–64.

[10] S.R. Broadbent and J.M. Hammersley, Percolation processes. I. Crystals and mazes, *Proc. Cambridge Philos. Soc.* **53** (1957), 629–641.

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

[12] M.H. Freedman, Percolation on the projective plane, *Math. Res. Lett.* **4** (1997), 889–894.

[13] E. Friedgut and G. Kalai, Every monotone graph property has a sharp threshold, *Proc. Amer. Math. Soc.* **124** (1996), 2993–3002.

[14] G. Grimmett, *Percolation*, Second edition. Springer-Verlag, Berlin, 1999. xiv+444 pp. ISBN 3-540-64902-6

[15] J.M. Hammersley, Percolation processes. II. The connective constant, *Proc. Cambridge Philos. Soc.* **53** (1957), 642–645.

[16] J.M. Hammersley, Percolation processes: Lower bounds for the critical probability, *Ann. Math. Statist.* **28** (1957), 790–795.

[17] J.M. Hammersley, Bornes supérieures de la probabilité critique dans un processus de filtration, Le calcul des probabilités et ses applications. Paris, 15-20 juillet 1958, Colloques Internationaux du Centre National de la Recherche Scientifique, LXXXVII (1959), pp. 17–37.

[18] T.E. Harris, A lower bound for the critical probability in a certain percolation process, *Proc. Cam. Philos. Soc.* **56** (1960), 13–20.

[19] J. Kahn, G. Kalai and N. Linial, The influence of variables on boolean functions, Proc. 29-th Annual Symposium on Foundations of Computer Science, 68-80, Computer Society Press, 1988.

[20] H. Kesten, The critical probability of bond percolation on the square lattice equals 1/2, *Comm. Math. Phys.* **74** (1980), 41–59.

[21] D. J. Kleitman, Families of non-disjoint subsets, *J. Combinatorial Theory* **1** (1966), 153–155.

[22] R. Langlands, P. Pouliot and Y. Saint-Aubin, Conformal invariance in two-dimensional percolation, *Bull. Amer. Math. Soc.* **30** (1994), 1–61.

[23] T.M. Liggett, R.H. Schonmann and A.M. Stacey, Domination by product measures, *Annals of Probability* **25** (1997), 71–95.

[24] L. Russo, A note on percolation, *Z. Wahrsch. Verw. Gebiete* **43** (1978), 39–48.

[25] P.D. Seymour and D.J.A. Welsh, Percolation probabilities on the square lattice, in Advances in graph theory (Cambridge Combinatorial Conf., Trinity College, Cambridge, 1977). *Ann. Discrete Math.* **3** (1978), pp 227–245.

[26] S. Smirnov, Critical percolation in the plane: conformal invariance, Cardy's formula, scaling limits, *Comptes Rendus de l'Académie des Sciences. Série I. Mathématique*, **333** (2001), 239–244. Expanded version available at `www.math.kth.se/~stas/papers`.

[27] M.Q. Vahidi-Asl and J.C. Wierman, First-passage percolation on the Voronoi tessellation and Delaunay triangulation, in Random graphs '87 (Poznań, 1987), Wiley, Chichester (1990), pp 341–359.

[28] M.Q. Vahidi-Asl and J.C. Wierman, A shape result for first-passage percolation on the Voronoi tessellation and Delaunay triangulation, in *Random graphs, Vol. 2 (Poznań, 1989)*, Wiley-Intersci. Publ., Wiley, New York (1992), pp 247–262.

[29] M.Q. Vahidi-Asl and J.C. Wierman, Upper and lower bounds for the route length of first-passage percolation in Voronoi tessellations, *Bull. Iranian Math. Soc.* **19** (1993), 15–28.

[30] A. Zvavitch, The critical probability for Voronoi percolation, MSc. thesis, Weizmann Institute of Science (1996).