VORONOI PERCOLATION: TOPOLOGICAL STABILITY AND GIANT CYCLES

BENJAMIN SCHWEINHART AND MORGAN SHUMAN

Abstract. We study the topological stability of Voronoi percolation in higher dimensions. We show that slightly increasing p allows a discretization that preserves increasing topological properties with high probability. This strengthens a theorem of Bollobás and Riordan and generalizes it to higher dimensions. As a consequence, we prove a sharp phase transition for the emergence of i-dimensional giant cycles in Voronoi percolation on the $2i$-dimensional torus.

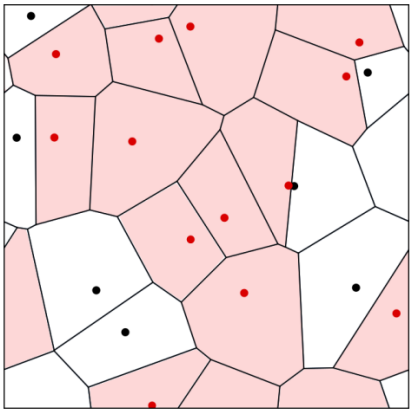

Figure 1. Voronoi percolation at $p = 0.5$.

1. Introduction

The Poisson–Voronoi mosaic on $\mathbb{R}^d$ is the random polyhedral complex obtained by taking the Voronoi diagram of a Poisson process on $\mathbb{R}^d$ with constant intensity. Voronoi percolation is the random subcomplex of the mosaic where each Voronoi cell is included independently with probability p. A seminal paper by Bollobás and Riordan shows that the percolation threshold of Voronoi percolation on $\mathbb{R}^2$ is $1/2$ [4]. We prove a higher-dimensional, finite-volume analogue of this statement for Voronoi percolation on a torus: the i-dimensional homological percolation threshold on the $2i$-dimensional torus is $1/2$. Roughly speaking, homological percolation occurs if there is an i-dimensional giant

cycle spanning the torus. When $i = 1$, a giant cycle is a periodic path. Homological percolation was introduced by Bobrowski and Skraba [1, 2] and further studied by Duncan, Kahle, and Schweinhart [6].

Our proof strategy combines the ideas of [4] and [6]: we construct a discretization of Voronoi percolation at a carefully chosen scale, prove a sharp phase transition for a "stable" homological percolation event in the discrete model, and show it coincides with the transition for the original model. A key technical lemma of [4] posits (roughly) that any paths lost in the discretization process can be made up by slightly increasing p. We strengthen this result by considering a broader class of topological conditions and generalizing to arbitrary dimensions. As a bonus, we obtain a shorter, more readable proof by using results from the literature on Delaunay triangulations [3, 5].

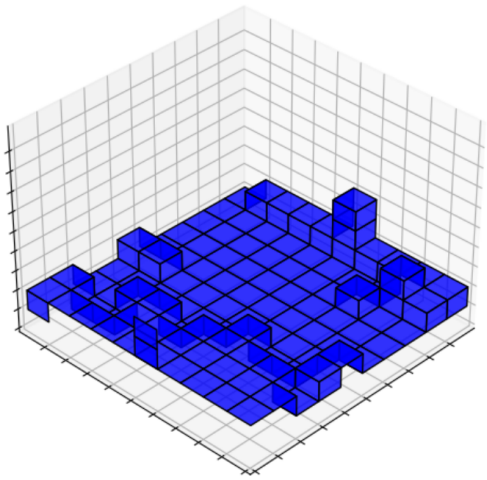

FIGURE 2. A giant cycle in 2-dimensional plaquette percolation. Taken from [6].

We examine two homological events: the existence of an i-dimensional spanning surface, called a **giant i-cycle**; and the presence of a basis of (non-homologous) giant i-cycles. We denote these events A and S, respectively. To illustrate the significance of S, consider A and S for giant 1-cycles on a 2-dimensional torus. If there is only a single periodic path, say from bottom-to-top, then A occurs but S does not. If there are two non-homologous periodic paths, say one from left-to-right and one from top-to-bottom, then both A and S occur. For discrete percolation processes, it isn't difficult to show that the events A and S have sharp threshold functions. However, it is challenging to prove that they coincide for the two events and that they do not depend on the size of the torus. Duncan, Kahle, and Schweinhart [6] investigated homological percolation in models of plaquette percolation and permutohedral site

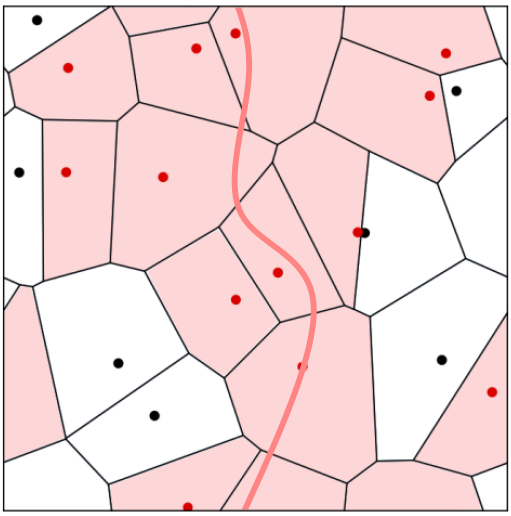

FIGURE 3. A percolation where A occurs, but not S. Both the red and white Voronoi cells contains a giant 1-cycle. The giant 1-cycle in the red cells is highlighted.

percolation. They demonstrated the existence of sharp phase transitions for A and S for i-dimensional homological percolation on a $2i$-dimensional torus (for which the threshold is $1/2$) and $(d-1)$-dimensional homological percolation on the d-dimensional torus (where the threshold is dual to that of classical percolation). We show the former statement for i-dimensional Voronoi percolation on the $2i$-dimensional torus. Our proof strategy combines the ideas of [4] and [6]: we construct a discretization of Voronoi percolation at a carefully chosen scale, prove a sharp phase transition for a "stable" homological percolation event in the discrete model, and show it coincides with the transition for the original model. A key technical lemma of [4] posits (roughly) that any paths lost in the discretization process can be made up by slightly increasing p. We strengthen this result by considering a broader class of topological conditions and generalizing to arbitrary dimensions. As a bonus, we obtain a shorter, more readable proof by using results from the literature on Delaunay triangulations [3, 5].

We quickly introduce the necessary notation to state our results and revisit them later. Let $\mathbb{T}_N^d$ be the cube $[-N, N]^d$ with opposite faces identified. Define Z to be a Poisson point process of intensity 1 on $\mathbb{T}_N^d$; color the points of Z red independently with probability p, and color the remaining points white. Denote by $R = R(p)$ the set of red points of Z and denote the Voronoi percolation as the pair (Z, R). For $W \subset Z$, let $U(W, Z)$ be the union of of the Voronoi cells generated by the points of W in Z.

Roughly speaking, a point set $W \subset Z$ is **δ-good** if the topology of the embedding of $U(W, Z)$ into the torus is invariant to δ-perturbations, where a δ-perturbation is a bijection so that $\|z - \zeta(z)\| < \delta$ for all $z \in Z$. More precisely, we require that every δ-perturbation $\zeta : Z \to Z'$ induces a homeomorphism of the torus $\bar{\zeta} : \mathbb{T}^d \to \mathbb{T}^d$ so that

$$\bar{\zeta}(U(W, Z)) = U(\zeta(W), \zeta(Z)).$$

This amounts to saying that every topological property of $U(W, Z)$ and its embedding in the torus is preserved under δ-perturbations. For example, $U(W, Z)$ contains a periodic path if and only if $U(\zeta(W), Z')$ does. Given this vocabulary, we can now state our main technical result.

Theorem 1. *Let $\epsilon_0 > 0$ and set $\delta_0 = N^{-\epsilon_0}$. Then there exists a coupling of $(Z_1, R_1(p))$ with $(Z_2, R_2(p + \epsilon))$ so that there is a δ_0-good subset $W \subset Z_2$ so that $W \subset R_2$ and*

$$U(R_1, Z_1) \subset U(W, Z_2) \subset U(R_2, Z_2)$$

with high probability.

This allows us to approximate Voronoi percolation with a discrete process that retains the increasing topological properties of R with high probability. In tandem with the sharp threshold theorem of Friedgut and Kalai generalized by Bollobás and Riordan, Theorem 18, we demonstrate and determine the sharp threshold for i-dimensional homological percolation on a $2i$-dimensional torus.

Quickly recall that homology depends on an abelian group of coefficients, which we take to be a field $\mathbb{F}$. We will review homology in Section 2.4 below. For technical reasons involving representation theory we only consider coefficients from fields with characteristic not equal to 2. A full explanation is provided in [6], but the idea is that if, say, $d = 2$, $\mathbb{F}$ has characteristic 2, and A' is the event that there is a periodic path which crosses the torus once left-to-right and once bottom-to-top, then event S cannot be "spun up" from A' by taking the intersection events symmetric to A'.

Theorem 2. *Suppose* $\mathrm{char}(\mathbb{F}) \neq 2$. *If $d = 2i$ then*

$$\begin{cases} \mathbb{P}_p(A) \to 0 & p < \frac{1}{2} \\ \mathbb{P}_p(S) \to 1 & p > \frac{1}{2} \end{cases}$$

as $N \to \infty$.

2. Background

2.1. **Delaunay and Voronoi.** Let Z be a **Poisson point process**. For $z \in Z$, the **Voronoi cell** of z, $V(z)$, is

$$V(z) = \{x \in \mathbb{T}_N^d : d(x,z) \leq d(x,z'), \forall z' \in Z\}.$$

The collection of Voronoi cells forms a **polyhedral complex** called a **Poisson-Voronoi mosaic** that we denote $V(Z)$. Note that a Voronoi tesselation on the torus may fail to be a polyhedral complex if the cells are large relative to the diameter of the ambient space. This happens with probability $o(1)$ in the asymptotic regimes we consider. We give a definition of a polyhedral complex below, borrowed from Definition 2.38 in [10].

Definition 3 (Polytope and Polyhedral complex)**.** *A (convex) n-**polytope** is an n-dimensional convex hull of a finite set of points. A **polyhedral complex** is a collection X of convex polytopes in $\mathbb{R}^n$ such that*

(1) If $\sigma \in X$, then all faces of σ are in X.

(2) If $\sigma_1, \sigma_2 \in X$ and $\sigma_1 \cap \sigma_2 \neq \varnothing$, then $\sigma_1 \cap \sigma_2 \in X$.

A set of points in a d-dimensional Euclidean space is in **general position** if no $(d+2)$-points lie on on the same $(d-1)$-sphere. We will assume this hypothesis as it occurs with probability one for Poisson point processes. It implies that the intersection of k Voronoi cells is either empty or a $(d-k+1)$-face of the Voronoi tesselation, and all $(d-k+1)$-faces occur as such an intersection (on the torus, we also need to assume that the Voronoi cells are small relative to the diameter of the ambient space). In particular, the intersection of $(d+2)$ Voronoi cells is necessarily empty. In three dimensions, two Voronoi cells can intersect in a 2-dimensional face, three Voronoi cells can intersect in a 1-dimensional edge, and four Voronoi cells can intersect in a 0-dimensional vertex. Note that the points of the integer lattice $\mathbb{Z}^d$ are not in general position. The Voronoi diagram of $\mathbb{Z}^2$ consists of unit squares, some of which meet only in a single vertex. This causes topological complications: the complement of a union of these squares is not necessarily topologically equivalent to the union of complementary squares. On the other hand, the complement of a set of Voronoi cells in a Poisson point process is homeomorphic to the union of complementary cells with probability one.

The general position hypothesis also allows the definition of a dual cell complex, called the Delaunay triangulation. Before discussing this in more detail, we define a simplicial complex and its terminology.

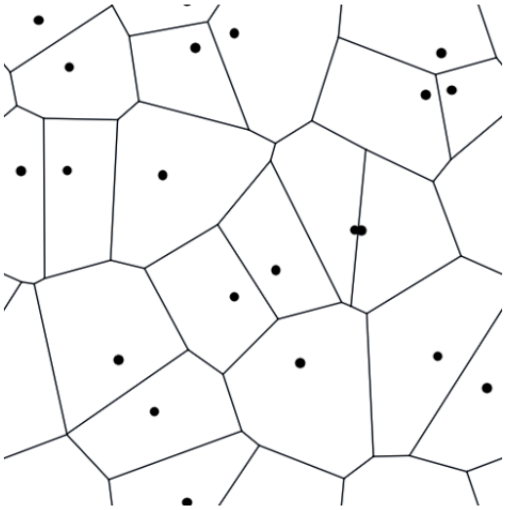

FIGURE 4. A Poisson-Voronoi mosaic

Definition 4 (Simplex and simplicial complex)**.** *An* n***-simplex*** *is the convex hull of* $(n+1)$*-vertices that are affinely independent. A* ***simplicial complex*** $\mathcal{K}$ *is a collection of simplices such that*

(1) For all $\sigma \in \mathcal{K}$*, every face of* σ *is in* $\mathcal{K}$*.*

(2) For any $\sigma_1, \sigma_2 \in \mathcal{K}$*,* $\sigma_1 \cap \sigma_2$ *is a face of* σ_1, σ_2*.*

The Delaunay complex is a natural simplicial complex generated by a point set $Z \subset X$ and a metric on X, dist. (In our case, we take $X = \mathbb{T}_N^d$ and dist as the Euclidean metric.)

Definition 5 (Delaunay ball)**.** *A* ***Delaunay ball****, or a* ***maximal empty ball****, is a ball* $B \subset X$ *such that* $|B \cap Z| = 0$*,* $|\partial B \cap Z| \geq d+1$*.*

If the points are in general position then the Delaunay complex is a simplicial complex called the **Delaunay triangulation**.

Definition 6 (Delaunay triangulation)**.** *If* Z *is in general position, then its* ***Delaunay triangulation****,* $\mathrm{Del}(Z)$*, is*

$$\mathrm{Del}(Z) := \{\mathrm{conv}\{z_1, \ldots, z_{d+1}\} \,|\, z_i \in \partial B, B \text{ is a Delaunay ball}\},$$

where $\mathrm{conv}\{z_1, \ldots, z_{d+1}\}$ *denotes the convex hull. Moreover, we define* $\mathrm{Del}(W; Z)\}$ *as the subcomplex of* $\mathrm{Del}(Z)$ *consisting of simplices whose vertices are in* W*.*

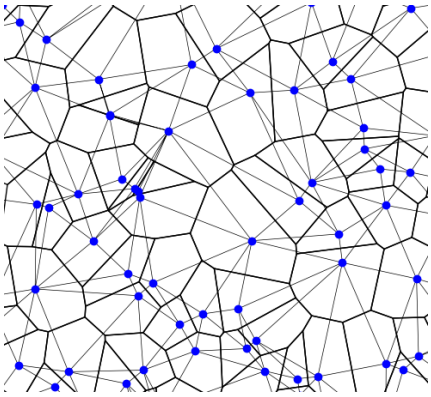

FIGURE 5. Voronoi diagram with Delaunay complex superimposed.

The duality relationship between Delaunay and Voronoi is given by

$$\mathrm{Del}(Z) = \{\varnothing \neq \sigma \subset Z \mid \bigcap_{u\in\sigma} V(u) \neq \varnothing\}.$$

That is, the Voronoi cells of $z_1, \ldots, z_k$ intersect non-trivially if and only if $(z_1, \ldots, z_k)$ forms a $(k-1)$ simplex of the Delaunay triangulation. For example, in Figure 5, the vertices of the Delaunay triangulation are the points Z, two vertices form a Delaunay edge if and only if their Voronoi cells intersect in a Voronoi edge, and three vertices form a Delaunay face if and only if their Voronoi cells intersect in a point.

This is the higher dimensional generalization of graph duality, where i-polytopes are dual to $(d-i)$-polytopes. More generally, the duality mapping of a polyhedral complex is given by the **Nerve** (definition borrowed from 10.3 of [7]),

$$\mathrm{Nerve}(X) := \{\varnothing \neq \sigma \subseteq X \mid \bigcap \sigma \neq \varnothing\}.$$

One can observe that $\mathrm{Nerve}(V(Z)) = \mathrm{Del}(Z)$ and $\mathrm{Nerve}(\mathrm{Del}(Z)) = V(Z)$.

Moreover, we introduced the duality relationship to show the following topological relation.

Lemma 7. *Let Z be in general position and let $P \subset Z$ finite. Then,* $\mathrm{Del}(P; Z)$ *and* $U(P; Z)$ *have the same homotopy type.*

Proof. This is an immediate corollary of the Nerve Theorem in [7] □

This lemma allows us to translate topological results on the Delaunay triangulation to the Voronoi setting. Lastly, we introduce the **star** of a simplex.

Definition 8 (Star)**.** *The star of a vertex $w \in Z$ is the union of all simplices of* $\mathrm{Del}(X)$ *containing* w*. More generally, for* $W \subset Z$, $\mathrm{Star}(W, X)$ *is the union of the stars of the vertices of* W*.*

Notation note: in this paper, we work solely with polyhedral complexes. More general than polyhedral complexes are cell complexes, which are built of i-cells which are homeomorphic to i-dimensional balls but are not required to be polytopes. We use terminology from this more general setting: we refer to the faces of the Voronoi tesselation as **cells**, and the set of i-cells as the i**-skeleton**.

2.2. **A crash course in homology.** Roughly speaking, homology measures the i-dimensional holes of a space. The homology of a topological space can be defined using singular homology, and the topology of a polyhedral complex can be defined using cellular homology [9], but the definition is simpler and less technical for simplicial complexes. As such, we define homology for simplicial complexes $\mathcal{K}$ and use Lemma 7 to define it for unions of Voronoi cells.

Definition 9 (n-chain)**.** *Let $C_n(\mathcal{K})$ be the space of formal $\mathbb{F}$-linear combinations of n-simplices of $\mathcal{K}$, that is*

$$C_n(\mathcal{K}) = C_n(\mathcal{K}, \mathbb{F}) := \{\sum c_\alpha \sigma_\alpha^n \,:\, c_n \in \mathbb{F}\}.$$

$C_n(\mathcal{K})$ forms a vector space over $\mathbb{F}$.

The meaning of coefficients. The simplest choice of coefficients is $\mathbb{F} = \mathbb{Z}_2$. The coefficients, $0, 1$, can be viewed as excluding or including a face, so $C_n(\mathcal{K}; \mathbb{Z}_q)$ can be formally identified with the power set of the collection of n-faces. A different choice of coefficients is $\mathbb{F} = \mathbb{Z}_3$. The coefficients, $0, 1, 2$, can be viewed as excluding $c_\alpha = 0$ a face or including $c_\alpha = 1, 2$ face with orientation (say up or down). As we discuss briefly below, the ranks of the homology groups depend on the coefficients in general.

We will discuss the meaning of the coefficients after defining homology. First, we introduce the boundary homomorphism.

Definition 10 (Boundary homomorphism)**.** *The* ***boundary homomorphism****,* $\partial_n : C_n(\mathcal{K}) \to C_{n-1}(\mathcal{K})$ *is given by*

$$\partial \sigma^k = \sum_{i=0}^{n} (-1)^i \sigma^k | [v_0, \ldots, \hat{v_i}, \ldots, v_n].$$

The term $[v_0, \dots, \hat{v_i}, \dots, v_n]$ *indicates the simplex generated by* $v_0, \dots, v_n$ *excluding* v_i*. Moreover,* $\operatorname{Im} \partial_{n+1}$ *are the* n***-boundaries*** *and* $\operatorname{Ker} \partial_n$ *are the* n***-cycles****. We define the set of all* n*-boundaries as* $B_n(X)$ *and the set of all* n*-cycles as* $Z_n(X)$*.* $B_n(X)$ *and* $Z_n(X)$ *are vector subspaces of* $C_n(X)$*.*

The boundary satisfies the fundamental relation $\partial_n \circ \partial_{n-1} = 0$. In other words, $B_n(X) \subset Z_n(X)$. We often write $\partial_n = \partial$ when its clear and $\partial^2 = 0$ to express the previous relation. If there exists an n-cycle that is not the boundary of an $(n+1)$-chain, e.g, $Z_n(X) \subsetneq B_n(X)$, then we say there exists an n-dimensional hole. This is formalized with homology.

Definition 11 (Homology)**.** *The* n*th homology is defined as*

$$H_n(\mathcal{K}) = H_n(\mathcal{K}; \mathbb{F}) := \operatorname{Ker} \partial_n / \operatorname{Im} \partial_{n+1}.$$

Moreover, the n*th homology is a vector space over* $\mathbb{F}$*.*

Example: Consider a triangle $\sigma = [v_0, v_1, v_2]$. As $\partial_0[v_1] = 0$, $\partial_1[v_0, v_1] = v_1 - v_0$, and $\partial_2[v_0, v_1, v_2] = [v_1, v_2] - [v_0, v_2] + [v_0, v_1]$, we have

$$C_n \cong \begin{cases} 0 & n \geq 3 \\ \mathbb{F} & n = 2 \\ \mathbb{F}^3 & n = 1 \\ 0 & n = 0, \end{cases} \quad \operatorname{Im} \partial_n \cong \begin{cases} 0 & n \geq 3 \\ \mathbb{F} & n = 2 \\ \mathbb{F}^2 & n = 1 \\ 0 & n = 0, \end{cases} \quad \operatorname{Ker} \partial_n \cong \begin{cases} 0 & n \geq 3 \\ 0 & n = 2 \\ \mathbb{F} & n = 1 \\ \mathbb{F}^3 & n = 0. \end{cases}$$

Therefore,

$$H_n \cong \begin{cases} 0 & n \geq 1 \\ \mathbb{F} & n = 0. \end{cases}$$

As a triangle has no holes, the homology is trivial for all n (except $n = 0$, as H_0 counts the number of connected components, 1 in this case). The above example, although plain and straightforward, shows how laborious computations of homology can be.

More on coefficients: When $i = 0$ or when $\mathcal{K}$ is embeddable $\mathbb{R}^2$ or $\mathbb{R}^3$ then the rank of the homology does not depend on the choice of coefficients and counts the number of connected components. The higher homology groups depend on the choice of coefficients in general. For example, if S is a non-orientable surface such as the Klein bottle or real projective plane then $H_2(S; \mathbb{Z}_2) \cong \mathbb{Z}_2$ but $H_2(S; \mathbb{Z}_q) = 0$ for odd primes q: the 2-simplices cannot be oriented so that their boundaries cancel out.

A continuous function $f : X \to Y$ induces a linear transformation on homology $f_* : H_n(X) \to H_n(Y)$. This is particularly easy to define for an inclusion map

$\iota : X \hookrightarrow Y$ when $X \subset Y$. A cycle in X is automatically a cycle in Y and $B_n(X)$ is a subspace of $B_n(Y)$, so ι_* sends a coset $z + B_n(X)$ to the potentially larger coset $z + B_n(Y)$.

A key result is that homology is a homotopy invariant.

Proposition 12 (Corollary 2.11 of [9])**.** *The maps $f_* : H_n(X) \to H_n(Y)$ induced by a homotopy equivalence $f : X \to Y$ are isomorphisms for all n.*

Corollary 13. *Let $P = U(Q, Z)$ and let $I = \mathrm{Del}(Q; Z)$. Then, $H_n(P) \cong H_n(I)$.*

Proof. By Lemma 7, I is homotopy equivalent to P. By the previous corollary, the rest follows. □

2.3. **Voronoi Percolation. Voronoi percolation** is defined by including the Voronoi cells of a Poisson point process Z independently with probability p. More formally, the underlying probability space it is a pair (Z, R) where we include each point of Z in R independently with probability p. See Figure 1.

In [4], Bollobás and Riordan proved for $d = 2$, the critical probability is $1/2$. A key technical lemma is the following coupling result, adapted to our terminology

Theorem 14 (Bollobás and Riordan, 2005)**.** *Let (Z_1, R_1) be a Voronoi percolation on $\mathbb{T}^2_N$. For $\epsilon > 0$ and $\delta \leq N^{-\epsilon}$, there exists a coupling of $(Z, R(p))$ with $(Z_2, R_2(p + \epsilon))$ such that, for every path γ in $U(R_1, Z_1)$ there exists a path γ' in $U(R_2, Z_2)$ and every point of γ' is within distance $(\log N)^2$ of some point of γ and vice-versa so that every point 4δ-close to γ' is red.*

Theorem 1 generalizes this statement to d-dimensions and strengthens it to control all increasing topological properties.

2.4. **Homological Percolation and Discretization.** Let $P \subset \mathbb{T}^d_N$. The inclusion map $P \hookrightarrow \mathbb{T}^d_N$ induces a map on the homology $\phi_i : H_i(P, \mathbb{F}) \to H_i(\mathbb{T}^d_N, \mathbb{F})$. We define our homological percolation events below.

Definition 15. *Let* $A := \{\operatorname{rank} \phi_i > 0\}, \quad S := \{\operatorname{nullity} \phi_i = 0\}$.

Our coupling result, Theorem 1, states that slightly increasing the probability and discretizing the Voronoi percolation preserves increasing events invariant under homeomorphisms of $\mathbb{T}^d$, including A and S. With the coupling in mind,

we define our formal discretization here, which is the same as the crude state discretization constructed in [4].

Definition 16. *Let $\delta = N^{-\epsilon}$. We partition $\mathbb{T}^d$ into $\lceil (N/\delta)^d \rceil$ cubes of length δ. Each cube is assigned a value in $\{-1, 0, 1\}$,*

(1) -1 if it contains a white point

(2) 0 if it contains no points

(3) 1 if it contains only red points.

*The assignment of (Z, R) is called the **coarse state**, denoted* $\mathrm{Coarse}_\delta(Z, R)$.

Later, we study the probability space of the coarse state to apply Theorem 18. To make events well-defined on the coarse state, we introduce the notion of a **stable event**.

Definition 17 (Stable event)**.** *For any event X define*

$X_{stable} := \{V \mid X \text{ holds for all Voronoi percolations with the same coarse state}\}$.

Events on percolation processes are usually defined in terms of the geometry of the percolation process itself. Such examples include: there exists an infinite open component containing the origin, there exists k disjoint open paths from the origin to the boundary of a box of length n (called a k-arm event), and so on. A stable event is different, instead it is an invariance under the placement of points into their respective cubes. As $X_{\text{stable}} \subset X$, we obtain the useful inequality

$$\mathbb{P}_p(X_{\text{stable}}) \leq \mathbb{P}_p(X).$$

2.5. **Probabilistic tools.** Foundational to our proof is the sharp threshold theorem of Friedgut and Kalai [8], stated in a form generalized by Bollobás and Riordan [4] to a probability space over $\{-1, 0, 1\}^n$ (rather than $\{0, 1\}^n$).

Theorem 18 (Bollobás and Riordan)**.** *Let X be an increasing event on $\{-1, 0, 1\}^n$ with some probability measure $\mathbb{P}^n_{p_1,p_2}$ that is invariant under a transitive group action and*

$$0 < q_1 < p_1 < 1/e, \quad 0 < p_2 < q_2 < 1/e\,. \tag{1}$$

If $\mathbb{P}^n_{p_1,p_2}(X) > \eta$, then $\mathbb{P}^n_{q_1,q_2}(X) > 1 - \eta$ whenever

$$\min\{q_2 - p_2, p_1 - q_1\} \geq c \log(1/\eta) p_{max} \log(1/p_{max}) / \log n \tag{2}$$

where $p_{max} = \max\{p_1, q_2\}$.

2.6. **Paper outline.** The majority of the paper is focused on proving Theorem 1. We develop key results, Lemma 23 and Theorem 36 in sections 3 and 4 respectively. Then, in section 5, we demonstrate Theorem 1 by carefully constructing a coupling that yields an appropriate discretization of a Voronoi percolation with high probability. Lastly, in section 6, with our discretization in hand and results from [6], we prove Theorem 2 on homological percolation.

As mentioned, many of these arguments, especially in Sections 3, 5, are modifications of Bollobás and Riordan's seminal work on Voronoi percolation.

3. Asymptotics of bad points

3.1. **Delaunay Triangulations.** Crucially, [3] established several theorems on the stability of Delaunay triangulations under perturbations of both the point locations and the metric. These results imply corresponding stability theorems for Voronoi tessellations. We review the the terminology and results of [3] and our main result of the section.

Our first definition controls the sparseness of a point sample.

Definition 19 (R-sample set)**.** *A point set $Z \subset \mathbb{T}_N^d$ is an R-**sample set** if the distance from any point $x \in \mathbb{T}_N^d$ to Z is at most R.*

A simplex formed by $d+1$ points is δ-unstable if there is another point that is close to its Delaunay ball.

Definition 20 (δ-unstable)**.** *Let σ be a simplex with $(d+1)$-points and a Delaunay ball B. We say σ is δ-unstable if*

$$\operatorname{dist}_{\mathbb{T}_N^d}(q, \partial B) > \delta$$

for some $q \in P \setminus \sigma$.

If a simplex is δ-unstable, then it is fragile under point perturbation. That is, if there exists a $q \in P \setminus \sigma$ that is δ-close to ∂B, then one may perturb q by at most δ into B. As a result, B is no longer a maximal empty ball, and the simplex shatters! In [4], they study the asymptotics of "δ-close pairs" and "δ-bad quadruples", the latter of which is related to δ-unstable simplices. We generalize their analysis, and rename δ-bad quadruple to δ-unstable tuple for notational consistency.

Definition 21 (δ-close pair and δ-unstable $(d+2)$-tuple)**.** *A pair of points $\{z_1, z_2\}$ is a δ-close pair if the distance between them is less than δ. Moreover,*

we call a $(d+2)$-tuple δ-unstable if there exists a spherical shell of inner radius $r < \log N$ and outer radius $r + \delta$ that contains $(d+2)$-points with no δ-close pairs. Additionally, we say a point is δ-unstable if it belongs to a δ-close pair or a δ-unstable $(d+2)$-tuple. Lastly, we denote the set of all δ-close pairs in Z as Z_δ and the set of all δ-unstable $(d+2)$-tuple as Z_T.

A simplex can be δ-unstable only if is in the star of a δ-close pair or if its $(d+1)$ vertices form a δ-unstable $(d+2)$-tuple with some other vertex. Moreover, a point belongs to a δ-unstable $(d+2)$-tuple of Z if and only if x belongs to a δ-unstable simplex with circumradius at most $\log N + \delta$ and no δ-close pairs for some $Z' \subset Z$. We will also account for neighbors of unstable points.

Definition 22 (δ-bad point, and δ-bad cluster)**.** *Let $P \subset Z$. We say that a point $z \in Z$ is δ-**bad** if it is contained in the star of a δ-unstable point. A δ-**bad cluster** is a maximal component of δ-bad points.*

For the rest of this section, just as in [4], we do not explicitly discuss the underlying Delaunay triangulation. Instead, we only discuss unstable and bad points.

Lastly, we introduce the following notation to streamline the statements of some of our technical results. When we write that a random variable X is $o(f(N))$ (resp. $O(f(N))$) with probability $1 - o(1)$, we mean that for any $\eta > 0$ (resp. there exists $\eta > 0$ so that) , $X < \eta f(N)$ with high probability as $N \to \infty$. Our goal in this section is to prove the following lemma, which corresponds to to Lemma 6.4 of [4].

Lemma 23. *Every δ-bad cluster C has diameter at most $\log N$ with probability $1-o(1)$. Moreover, every δ-bad cluster C has size $o(\log N)$ and contains $O(1)$ unstable points with probability $1 - o(1)$.*

By "diameter" we mean the longest distance between any two points of the cluster. In [4], they enlist a technical lemma (Lemma 19) to obtain a bound on the maximal degree of δ-bad points with high probability. In explaining the technical complications, they give a quick argument for the existence of a δ-unstable point with degree $\Theta(\frac{\log N}{\log \log N})$. It turns out that this is asymptotically sharp.

Theorem 24 (Bonnet, Chenavier [5])**.** *Let Δ_N be the maximal degree in a Poisson-Delaunay triangulation over all vertices in $N^{1/d}[0,1]^d$, $d \geq 2$. Then there exists a deterministic function $N \mapsto J_N$, $N > 0$, with values in $\mathbb{N}$, such that*

(i) $\mathbb{P}(\Delta_N \in \{J_N, J_N+1, \dots, J_N+l_d\}) \underset{N\to\infty}{\longrightarrow} 1$, *where* $l_d = \lfloor \frac{d+3}{2} \rfloor$;

(ii) $J_N \underset{N\to\infty}{\sim} \frac{d-1}{2}\frac{\log N}{\log\log N}$.

We don't need the full strength of this theorem so we write the following corollary for convenience.

Corollary 25. *The maximal degree of the Poisson point process on* $\mathbb{T}_N^d$ *is* $o(\log N)$ *with probability* $1-o(1)$.

Proof. Let Δ_{N^d} be the maximal degree in a Poisson-Delaunay graph over all nodes in $N[0,1]^d = (N^d)^{1/d}[0,1]^d$. By the previous theorem,

$$\mathbb{P}\left(\Delta_{N^d} = O\left(\frac{d-1}{2}\frac{\log(N^d)}{\log\log(N^d)}\right)\right) \to 1\,.$$

Lastly, $\frac{d-1}{2}\frac{\log(N^d)}{\log\log(N^d)} = o(\log N)$. □

3.2. **Poisson-Delaunay preliminaries.** We first review some elementary results on Poisson point processes. Throughout this section, recall that $\delta = N^{-\epsilon}$. Moreover, for a Borel set $\Omega \subset \mathbb{T}_N^d$, we write $\|\Omega\|$ to denote its volume.

Lemma 26. *Every simplex* σ *has circumradius at most* $\rho\sqrt[d]{\log N}$ *for some* $\rho > 0$ *with probability* $1-o(1)$. *In particular,* Z *is a* $(\log N)$*-sample set with probability* $1-o(1)$.

Proof. For any Borel set Ω and number of points $\xi := |\Omega \cap Z|$, we have

$$\mathbb{P}(\xi = 0) = e^{-\|\Omega\|}.$$

If Ω is a ball of radius $\rho\sqrt[d]{\log N}$ and $\rho > \sqrt[d]{d}$, then $\|\Omega\| > d\log N$. Hence,

$$\mathbb{P}(\xi = 0) = e^{-\rho^d \log N} = N^{-\rho^d} = o(N^{-d}).$$

Moreover, the torus may be covered with $O(N^d)$ overlapping balls Ω. The probability that at least one ball is empty is $o(1)$. Hence, the probability that none are empty is $1-o(1)$. □

Lemma 27. *Let* Z *be a point set in general position. For every* $z_1, \dots, z_{d+1} \in Z$, *define* $\mathcal{S}$ *to be the set all of points* x *so that* $\{z_1, \dots, z_{d+1}, x\}$ *is a* δ*-unstable tuple. Then,* $\mathcal{S}$ *has volume* $o(\delta^{1/2})$.

Proof. Equivalently, the set $\mathcal{S}$ is the union of all such spherical shells S that contain $z_1, \ldots, z_{d+1}$ with inner radius $r' < \log N$, outer radius $r' + \delta$. Let B be a ball of radius r whose boundary contains $z_1, \ldots, z_{d+1}$. We may assume that $z_1, \ldots, z_{d+1}$ does not contain a δ-close pair, otherwise they cannot belong to an unstable $(d+2)$-tuple by definition. We will show that each spherical shell S must contain either B, or choose from a finite selection of reflections of B.

Case 1: Assume there does not exist a hyperplane H so that $\operatorname{dist}(z_i, H) < \delta$ for all i. Then, S must contain ∂B. Otherwise, $z_1, \ldots, z_{d+1}$ must belong to the cross-section $S \cap \partial B$, a subset of a δ-wide hyperplane, thus contradicting our assumption. As S contains ∂B, S is contained in the spherical shell S_0 of inner radius $r - \delta$ and outer radius $r + 2\delta$ centered at the center of B. In other words, $\mathcal{S} \subset S_0$, so $\|\mathcal{S}\| \leq \|S_0\|$.

Let ν_d denote the volume of the d-dimensional ball, then

$$\begin{aligned} \|S_0\| &= \nu_d (r + 2\delta)^d - \nu_d (r - \delta)^d \\ &= \nu_d d(c)^{d-1} (3\delta) \end{aligned}$$

for some $c \in (r, r + \delta)$ by the Mean Value Theorem.

$$\begin{aligned} &\leq 3\nu_d d(r + \delta)^{d-1} \delta \\ &\leq 3\nu_d d(\rho \sqrt[d]{\log N} + \delta)^{d-1} \delta \\ &= o(\delta^{1/2}) \,. \end{aligned}$$

Note that the choice of exponent $1/2$ is arbitrary. That is, we just as well write $o(\delta^\alpha)$ for any $\alpha < 1$ in the last line.

Case 2: Assume there does exist a hyperplane H so that $\operatorname{dist}(z_i, H) < \delta$ for all i. Consider the set $\mathcal{H}$ of such hyperplanes. Consider a maximal set of orthogonal hyperplanes $H_1, \ldots, H_n \in \mathcal{H}$. If $n > d + 1$, then the cross-section of the hyperplanes with ∂B implies $z_1, \ldots, z_{d+1}$ contains a δ-close pair, a contradiction. Consider the balls $B_1, \ldots, B_n$ obtained by reflecting B across the hyperplanes $H_1, \ldots, H_n$ respectively. Then, arguing as in case 1, S must also contain B, or B_1, or $\ldots$, or B_n. Otherwise, it would imply the existence of another orthogonal hyperplane, a contradiction. Define S_k to be the spherical shell with inner radius $r - \delta$, outer radius $r + 2\delta$, and centered at B_k. Then, S is contained in S_k for some $k \in \{0, \ldots, n\}$. As $\|S_k\| = \|S_0\|$, $\|\mathcal{S}\| \leq (n+1)\|S_k\| = o(\delta^{1/2})$. □

Let Λ be a cube of length $\log N$ and let $|\Lambda|$ denote the number of vertices of Z in the cube, so $|\Lambda|$ is a Poisson random variable with mean $(\log N)^d$.

Lemma 28. *Let $B_1(\Lambda)$ be the event that Λ contains more than $3(\log N)^d$ points. Then, $\mathbb{P}(B_1) = o(N^{-d})$.*

Proof. By using a Chernoff bound, we have

$$\mathbb{P}(|\Lambda| \geq a) \leq \left(\frac{e\|\Lambda\|}{a}\right)^a e^{-\|\Lambda\|}.$$

Setting $a = 3(\log N)^d$ and $\|\Lambda\| = (\log N)^d$, we obtain

$$\begin{aligned}\mathbb{P}(|\Lambda| \geq 3) &\leq \left(\frac{e(\log N)^d}{3(\log N)^d}\right)^{3(\log N)^d} e^{-(\log N)^d} \\ &= \left(\frac{e}{3}\right)^{3(\log N)^d} e^{-(\log N)^d} \\ &= o(N^{-d}).\end{aligned}$$

□

3.3. **Unstable points and bad clusters.** The results of this section are a *mutatis muntandis* adaptation of Lemma 6.4 of [4]. That is, we follow the same line of arguments, but extend them to higher dimensions and employ recent results such as the Maximum Degree Theorem of [5] for greater brevity.

Lemma 29. *Let Ω be a Borel set satisfying $\|\Omega\| \geq 1$. Place $\xi = O((\log N)^d)$ points i.i.d. uniformly in Ω and let $M > 0$. The probability that there are more than $M(d+1)/d\epsilon$ δ-close pairs is $O(\delta^{M/2})$.*

Proof. Let $z_1, \ldots, z_\xi$ be an ordering of the points and let $B_\delta(z)$ denote a δ-ball centered at a point z. The conditional probability that z_i is δ-close to an earlier point $z_1, \ldots, z_{i-1}$ is at most

$$\frac{(i-1)\|B_\delta\|}{\|\Omega\|} \leq \xi\|B_\delta\| = O(\xi\delta) = O(\delta(\log N)^d) \leq O(\delta^{1/2}),$$

where we have used $\|\Omega\| \geq 1$, $\xi = O((\log N)^d)$, and $\|B_\delta\| = O(\delta)$. So, the probability that M points z_i are δ-close to a previously placed point is $O(\delta^{M/2})$.

Each z_i could be δ-close to many earlier points, however, with probability $1 - o(N^{-d})$, there are at most $(d+1)/d\epsilon$ possibly earlier points. Indeed, the

probability of placing $a = (d+1)/d\epsilon$ points in $B_\delta(z_i)$ is:

$$\begin{aligned}
\mathbb{P}(B_\delta(z_i) > a \,|\, \xi) &= \sum_{k=a+1}^{\xi} \binom{\xi}{k} \left(\frac{\|B_\delta(z_i)\|}{\|\Omega\|}\right)^k \left(1 - \frac{\|B_\delta(z_i)\|}{\|\Omega\|}\right)^{\xi-k} \\
&\leq \left(\frac{\|B_\delta(z_i)\|}{\|\Omega\|}\right)^a \underbrace{\sum_{k=a+1}^{\xi} \binom{\xi}{k} \left(\frac{\|B_\delta(z_i)\|}{\|\Omega\|}\right)^{k-a} \left(1 - \frac{\|B_\delta(z_i)\|}{\|\Omega\|}\right)^{\xi-k}}_{\leq 1} \\
&\leq \|B_\delta(z_i)\|^a \\
&= O(\delta^a) \\
&= o(N^{-d}).
\end{aligned}$$

Thus, the conditional probability that Ω has more than $M(d+1)/d\epsilon$ points given ξ and an ordering of points is $O(\delta^{M/2})$. As the bound on the conditional probability is independent of ξ and the ordering of points, we infer that the probability that Ω contains more than $M(d+1)/d\epsilon$ points that are δ-close to an earlier point is $O((\delta^{1/2})^M)$. □

In particular, we are concerned about the case where Ω is a cube of length $\log N$, which we wrote as Λ in the last subsection.

Corollary 30. *Let $B_2(\Lambda)$ be the event that Λ has more than $\frac{2(d+1)^2}{d\epsilon^2}$ points in δ-close pairs. Then, $\mathbb{P}(B_2) = o(N^{-d})$.*

Proof. We apply the previous lemma using $M = 2(d+1)/\epsilon$. Firstly, the conditional probability of B_2 given that Λ contains fewer than $3(\log N)^d$ points is:

$$\mathbb{P}(B_2 \,|\, \xi \leq 3(\log N)^d) = O(\delta^{(d+1)/\epsilon}) = o(N^{-d}).$$

By Lemma 28, $\mathbb{P}(B_1) = o(N^{-d})$. Thus,

$$\mathbb{P}(B_2) = \mathbb{P}(B_2 \,|B_1)\mathbb{P}(B_1) + \mathbb{P}(B_2 \,|B_1^c)\mathbb{P}(B_1^c) = o(N^{-d}).$$

□

Lemma 31. *Let Ω be a Borel set satisfying $\|\Omega\| \geq 1$. Place $\xi := O((\log N)^d)$ points i.i.d. uniformly in Ω. Let T_M be the event that there are more than M points of Ω belonging to δ-unstable $(d+2)$-tuples. Then, $\mathbb{P}(T_M) = o(\delta^{M/2(d+3)})$.*

Proof. If T_M occurs, that is, there are more than M points in δ-unstable $(d+2)$-tuples, then there exists a subset $I \subset [\xi] = \{1, 2, \ldots, \xi\}$ where $|I| \leq M+(d+1)$, an order on I, and a subset $J \subset I$ with $|J| \leq \frac{M}{d+2}$ so that the following

occurs: for all $j \in J$, z_j belongs to a δ-unstable $(d+2)$-tuple with previously placed points. Let C_J denote this event. Our argument, borrowed from [4], is to compute a crude upper-bound on the number of possible I, orders on I, and choices of J, and to multiply this with the conditional probability $\mathbb{P}(C_J \,|\, I, \text{order on } I)$.

First, we compute an upper-bound for the number of choices of I. There are precisely $\sum_{k=0}^{M+d+1} \binom{P}{k}$ subsets I where $|I| \leq M+d+1$. By using the upper-bound $\binom{n}{k} \leq (\frac{en}{k})^k$, we obtain

$$\sum_{k=0}^{M+d+1} \binom{\xi}{k} \leq \sum_{k=0}^{M+d+1} \left(\frac{e\xi}{k}\right)^k \leq \xi^{M+d+1} \sum_{k=0}^{M+d+1} \left(\frac{e\xi}{k\xi^{M+d+1}}\right)^k \leq \xi^{M+d+2}.$$

Secondly, the number of choices of orders on I is $|I|! \leq |I|^{|I|} = (M+d+1)^{M+d+1}$ and we have at most 2^{M+d+1} choices for J.

Lastly, we compute a bound on the conditional probability $\mathbb{P}(C_J \,|\, I, \text{order on } I)$. Consider placing points of I down one-by-one with respect to the order. By Lemma 26, with probability $1 - o(1)$, every For each $j \in J$, the conditional probability that z_j forms a δ-unstable tuple with previously placed points in I is $o(\binom{|I|}{d+1}\delta^{1/2})/\|\Omega\| \leq o(\delta^{1/2})$ by the uniform distribution of z_j, Lemma 27, and the hypothesis $\|\Omega\| \geq 1$. So, $\mathbb{P}(C_J \,|\, I, \text{order on } I) = o(\delta^{|J|/2}) \leq o(\delta^{M/2(d+2)})$.

Putting it all together,

$$\begin{aligned}
\mathbb{P}(T_M) &= \overbrace{\left(O((\log N^d))^{(M+d+2)}\right)}^{\#(\text{Choices of } I)} \overbrace{\left((M+d+2)^{M+d+1}\right)}^{\#(\text{Orders on } I)} \overbrace{\left(2^{M+d+1}\right)}^{\#(\text{Choices of } J)} \overbrace{\left(\mathbb{P}(C_J \,|\, I, \text{order on } I)\right)}^{\text{Conditional Probability}} \\
&\leq O\left(\delta^{M/2(d+2)} (\log N)^{d(M+d+2)}\right) \\
&\leq o(\delta^{M/2(d+3)}).
\end{aligned}$$

□

Corollary 32. *Let $B_3(\Lambda)$ be the event that Λ contains more than $M = \frac{2(d+3)(d+1)}{\epsilon}$ δ-unstable $(d+2)$ tuples. Then, $\mathbb{P}(B_3) = o(N^{-d})$.*

Proof. By the previous lemma and Lemma 28, $\mathbb{P}(B_3 | B_1^c) = o(\delta^{\frac{d+1}{\epsilon}}) = o(N^{-d})$. □

Corollary 33. *Let $1 > \eta > 0$ be fixed. Define $G = G(\eta)$ to be the event that $|Z_\delta \cup Z_T \cap \Lambda| = O(1)$ and $|\text{Star}((Z_\delta \cup Z_T) \cap \Lambda)| = o(\log N)$ for every cube Λ of length $\eta \log N$. Then, $\mathbb{P}(G) = 1 - o(1)$.*

Proof. Let Λ' be a cube of length $\log N$. If $B_2(\Lambda')$ and $B_3(\Lambda')$ do not occur, then $|Z_\delta \cap \Lambda'|$ and $|Z_T \cap \Lambda'|$ are $O(1)$. Moreover, we may cover $\mathbb{T}_N^d$ with $O(N^d)$ cubes of length $\log N$ so that every cube Λ of length $\eta \log N$ sits inside a cube Λ' from the covering. By the union bound, the probability that $B_k(\Lambda)^c$ occurs for all Λ in the covering is $O(N^d) \cdot o(N^{-d}) = o(1)$.

Since each cube Λ is contained in at least one cube Λ' we may conclude that $|Z_\delta \cap \Lambda| = O(1)$ and $|Z_T \cap \Lambda| = O(1)$. Lastly, by Corollary 25, with probability $1 - o(1)$, every vertex of Z has degree $o(\log N)$. Thus,

$$|\mathrm{Star}((Z_\delta \cap \Lambda) \cup (Z_T \cap \Lambda))| = o(\log N).$$

□

We are now ready to prove 23.

Proof of Lemma 23. Let $1 < \eta < 0$. We may assume that $G(\eta)$ occurs by the by the previous corollary. Fix a cube Λ of length $\eta \log N$ and center it at some $z \in C$. By $G(\eta)$, $|(Z_\delta \cup Z_T) \cap \Lambda| = O(1)$. Moreover, $\mathrm{Star}(z)$ has diameter at most $2\rho\sqrt[d]{\log N}$ with probability $1 - o(1)$ by Lemma 26. Recalling that C is a maximal component of $\mathrm{Star}(Z_\delta \cup Z_T)$, we infer that the diameter of a maximal component of $\mathrm{Star}((Z_\delta \cap \Lambda) \cup (Z_T \cap \Lambda))$ is $O(1) \cdot 2\rho\sqrt[d]{\log N}$. Therefore, $C \subset \Lambda$ when N is sufficiently large, as Λ has length $\eta \log N >> 2\rho\sqrt[d]{\log N} \cdot O(1)$. Thus, if $G(\eta)$ occurs, then $|C| \leq |\mathrm{Star}(Z_\delta \cup Z_T) \cap \Lambda| = o(\log N)$. □

4. Topology of Voronoi Cells

We investigate the consequences of the Delaunay stability theorems of [5] for Voronoi tessellations. As such, the results of this section are stated deterministically under certain hypotheses that will, in later arguments, hold with high probability.

4.1. **Background.**

Definition 34. *Let $Z \subset \mathbb{T}_N^d$ be a finite point set. Let $Q \subset Z$ and $z \in Z$. We define*

(1) the Voronoi cell of z,

$$U(z, Z) = \{x \in \mathbb{T}_N^d : d(x, z) \leq d(x, z') \ \forall z' \in Z\}.$$

(2) the cell complex of Voronoi cells generated by Q in Z as $V(Q, Z)$.

(3) the union of Voronoi cells generated by Q in Z

$$U(Q, Z) := \bigcup_{q \in Q} U(q, Z)$$

(4) the interior site boundary of Q

$$\partial Q := \{q \in Q : \mathrm{Star}(q) \not\subset Q\}$$

(5) the interior of Q

$$\mathrm{Int}\, Q = Q \setminus \partial Q$$

In the previous section, we discussed the asymptotics of bad geometries, e.g δ-bad points and clusters. Here, we introduce our notions of good geometry.

Definition 35 (δ-good and δ-bubble wrapped)**.** *If $z \in Z$ is not δ-bad, then we call it δ-**good**. Given a set $Q \subset Z$, we say Q is δ-**bubble wrapped** if $\mathrm{Star}(Q) \setminus Q$ is δ-good.*

The utility of bubble wrap, as we will see, is that its topology is stable to reasonably small perturbations. And, just like real-world bubble wrap, while the fragile interior may break, the bubble wrap remains! The goal of this section is to prove the following Theorem.

Theorem 36. *Let $Z \subset \mathbb{T}_N^d$ in general position and let $Q \subset Z$ be δ-bubble wrapped. Let ζ be a δ^5-perturbation and let C be the union of all δ-bad points of Z. If $Z \setminus C$ is a $\log N$-sample set, then $\zeta|_{Q \setminus \mathrm{Int}\, C}$ can be extended to a homeomorphism $\bar{\zeta} : \mathbb{T}^d \to \mathbb{T}^d$ so that $\bar{\zeta}\left(U\left(Q, Z\right)\right) = U\left(\zeta\left(Q\right), \zeta\left(Z\right)\right)$.*

To prove this, we apply a special case of Theorem 4.14 from [3], which we have reformatted using our notation.

Definition 37 (Combinatorial isomorphism)**.** *Given two (regular) cell complexes Q_1, Q_2, we say that a bijection ζ between the vertices of Q_1 and the vertices of Q_2 induces **combinatorial isomorphism** from Q_1 to Q_2 if $v_1, \ldots, v_k$ are vertices of an i-dimensional face of Q_1 if and only if $\zeta(v_1), \ldots, \zeta(v_k)$ are vertices of an i-dimensional face in Q_2.*

Proposition 38. *Let $Z \subset \mathbb{T}_N^d$ be in general position, $Q \subset Z$, and C the set of all δ-bad points of Z. Further suppose $Z \setminus C$ is a $\log N$-sample set. If Q is δ-good and $\zeta : Z \to Z'$ is a δ^5 perturbation, then ζ induces a combinatorial isomorphism between $\mathrm{Star}(Q; \mathrm{Del}(Z))$ and $\mathrm{Star}(\zeta(Q); \mathrm{Del}(Z'))$.*

Proof. This follows from combining Theorem 4.14 with Lemma 4.10 of [3], using that every simplex of $\mathrm{Del}(Z \setminus Q)$ is not δ-unstable. □

We review our notation here. We call δ-close pairs and δ-unstable tuples δ-unstable points. Points adjacent to δ-unstable points in the Delaunay triangulation are called δ-bad; all other points are δ-good.

In a nutshell, good points are good, bad points are bad, and unstable points are worse. The previous section (Lemma 23) revealed that unstable points are asymptotically few and far between but bad points are non-negligible. As unstable points induce bad points, deleting unstable points reveals a set of entirely δ-good points.

4.2. **Stability of Voronoi cells.**

Lemma 39. *Let $q \in \operatorname{Int} Q$. Then,*

$$U(Q, Z) = U(Q \setminus \{q\}, Z \setminus \{q\}).$$

Moreover, the boundary does not change. That is, $\partial Q = \partial(Q \setminus \{q\})$.

Proof. $U(Q \setminus \{q\}, Z) \subset U(Q \setminus \{q\}, Z \setminus \{q\})$. Let $x \in \operatorname{Int} V(q, Z)$ and suppose $x \in V(y, Z \setminus \{q\})$. Since $\operatorname{dist}(x, q) < \operatorname{dist}(x, y)$, there exists a point v on the line segment L between x and y that is equidistant to y and q. $V(q, Z \setminus \{q\})$ is convex so $L \subset \operatorname{Int} V(z, Z \setminus \{q\})$ and $\operatorname{dist}(v, q) = \operatorname{dist}(v, y) < \operatorname{dist}(v, Z \setminus \{y, q\})$. Thus, $v \in V(y, Z) \cap V(q, Z)$ so $y \in \operatorname{Star}(q)$ which implies $y \in Q$ and $x \in U(Q \setminus \{q\}, Z \setminus \{q\})$.

The second statement directly follows from the observations that for $q' \in Q \setminus \{q\}$ $U(q', Z \setminus q) \subset U(q', Z) \cup U(q, Z)$ and that $y \in \operatorname{Int} Z' \iff U(y, Z') \subset \operatorname{Int} U(Z', Z)$ for any $Z' \subset Z$ and $z \in Z'$, where we are using two different notions of interior. □

If we iteratively throw away every interior point, the union of Voronoi cells does not change!

Corollary 40.

$$U(Q, Z) = U(\partial Q, Z \setminus \operatorname{Int} Q).$$

Lemma 41. *Let $\zeta : Z \to Z'$ be a δ^5-perturbation where both Z and Z' are in general position, let C be the set of δ-bad points of Z, and assume that $Z \setminus C$ is a $\log N$-sample set. Then the restriction $\zeta|_{Z \setminus \operatorname{Int} C}$ induces a combinatorial isomorphism between the Voronoi diagrams $V(Z \setminus \operatorname{Int} C; Z \setminus \operatorname{Int} C)$ and $V(Z' \setminus \operatorname{Int} \zeta(C); Z' \setminus \operatorname{Int} \zeta(C))$.*

Proof. By Proposition 38, every point in $Z \setminus \operatorname{Int} C$ is δ-good relative to the point set $Z \setminus \operatorname{Int} C$. This induces a combinatorial isomorphisms on the level of Delaunay triangulations,

$$\operatorname{Del}(Z \setminus \operatorname{Int} C; Z \setminus \operatorname{Int} C) \cong \operatorname{Del}(\zeta(Z \setminus \operatorname{Int} C); \zeta(Z \setminus \operatorname{Int} \mathrm{C})),$$

where we used that $\operatorname{Star}(\hat{Z}; \hat{Z}) = \operatorname{Del}(\hat{Z}; \hat{Z})$ for any point set $\hat{Z}$. Furthermore, a combinatorial isomorphism of Delaunay triangulations induces a combinatorial isomorphism of the corresponding Voronoi mosaics by their duality relationship. That is,

$$V(Z \setminus \operatorname{Int} C; Z \setminus \operatorname{Int} C) \cong V(\zeta(Z \setminus \operatorname{Int} C); \zeta(Z \setminus \operatorname{Int} C)).$$

□

Lemma 42. *Assume the hypotheses of the previous lemma, and let $Q \subset Z$ be δ-bubble wrapped. Then the boundary of Q is invariant under ζ:*

$$\partial\zeta(Q) = \zeta\partial(Q)\,.$$

Proof. Let $W = \operatorname{Star}(Q) \setminus Q$. As Q is δ-bubble wrapped, W is δ-good. By Proposition 38, ζ induces a combinatorial isomorphism between $\operatorname{Star}(W; Z)$ and $\operatorname{Star}(\zeta(W); \zeta(Z))$. Thus

$$\zeta\left(\partial Q\right) = \zeta\left(\operatorname{Star}(W; Z) \cap Q\right) = \operatorname{Star}(\zeta(W); \zeta(Z)) \cap \zeta\left(Q\right) = \partial\zeta\left(Q\right)\,.$$

□

The last ingredient in our proof of Theorem 36 is to demonstrate the stability of the topology of Voronoi cells of $Z \setminus C$ under ζ. That is, a combinatorial isomorphism should induce a homeomorphism of the torus that maps between the Voronoi tessellations. In particular, topological features like the presence or absence of giant cycles are preserved. We demonstrate this over the next two lemmas.

Lemma 43. *Suppose B_1, B_2 are homeomorphic to the d-ball and $f : \partial B_1 \to \partial B_2$ is a homeomorphism. Then, f may be extended to a homeomorphism $\bar{f}$ from B_1 to B_2.*

Proof. We proceed in cases.

Case 1: B_1 and B_2 are unit d-balls in $\mathbb{R}^d$ centered at the origin. We may define $\bar{f} : B_1 \to B_2$ by

$$\bar{f}(x) := \begin{cases} \|x\| f(x/\|x\|) & \text{if } x \neq 0 \\ 0 & \text{else.} \end{cases}$$

It is a homeomorphism as we may define a continuous inverse, $\bar{f}^{-1}$ by

$$\bar{f}^{-1}(x) := \begin{cases} \|x\| f^{-1}(x/\|x\|) & \text{if } x \neq 0 \\ 0 & \text{else.} \end{cases}$$

As $\bar{f} \circ \bar{f}^{-1} = id$ and $\bar{f}^{-1} \circ \bar{f} = id$, we conclude that $\bar{f}$ is a homeomorphism that extends f.

Case 2: general case. Let D be the unit d-ball in $\mathbb{R}^d$ and $g_1 : B_1 \to D, g_2 : B_2 \to D$ be homeomorphisms. Then, $h := g_2 \circ f \circ g_1^{-1}|_{\partial D} : \partial D \to \partial D$ is a homeomorphism. Thus, by the previous case, h extends to a homeomorphism $\bar{h} : D \to D$ such that $\bar{h}|_{\partial D} = g_2 \circ f \circ g^{-1}|_{\partial D}$. Therefore

$$\bar{f} := g_2^{-1} \circ \bar{h} \circ g_1 : B_1 \to B_2$$

is a homeomorphism that also satisfies the condition

$$\bar{f}|_{\partial B_1} = g_2^{-1} \circ \bar{h} \circ g_1|_{\partial B_1} = f.$$

□

Recall that a 0-cell is a vertex, a 1-cell is an edge, a 2-cell is a face, and so on.

Definition 44 (i-skeleton)**.** *Given a cell complex, the* i***-skeleton*** *is the collection of all* i*-cells.*

As each vertex is a 0-cell, a bijective map on the vertices is trivially a homeomorphism from the 0-skeleton to the 0-skeleton. We would hope that a combinatorial isomorphism between two Voronoi diagrams implies that they are homeomorphic in the ambient space. However, to rule out any topological degeneracies, such as a cell wrapping around the torus of length N, we enforce an upper bound on the diameter on every Voronoi cell.

Lemma 45. *If* ζ *induces a combinatorial isomorphism of* $V(Z; Z)$ *with* $V(f(Z); f(Z))$ *and the diameter of every Voronoi cell is less than* N *then* ζ *can be extended to a homeomorphism* $\bar{\zeta} : \mathbb{T}^d \to \mathbb{T}^d$ *so that* $\bar{\zeta}(V(z; Z)) = V(\zeta(z), \zeta(Z))$ *for all* $z \in Z$.

Proof. Let $V^i(Z; Z)$ denote the i-skeleton of $V(Z; Z)$. We induct on i. The base case $i = 0$ is trivial.

Suppose that ζ can be extended to a homeomorphism $\zeta_i : V^i(Z; Z) \to V^i(\zeta(Z); \zeta(Z))$ mapping i-cells to i-cells. Since the diameter of each Voronoi cell is less than N, the $(i+1)$-faces are homeomorphic to $(i+1)$-balls. So, ζ_i extends to a homeomorphism $\bar{\zeta}_i$ mapping an $(i+1)$-cell of $V^{i+1}(Z; Z)$ to an $(i+1)$-cell of

$V^{i+1}(\zeta(Z),\zeta(Z))$ by the previous lemma. As two $(i+1)$-cells may only intersect in an i-cell by the definition of a cell complex, we may glue over intersecting $(i+1)$-cells to extends further to a homeomorphism ζ_{i+1} from $V^{i+1}(Z;Z)$ to $V^{i+1}(\zeta(Z);\zeta(Z))$. □

Now we prove Theorem 36.

Proof of Theorem 36. By Lemma 42, $\zeta(\operatorname{Int} C) = \operatorname{Int}\zeta(C)$. We infer from Corollary 40 that

$$U\left(\zeta\left(Q\right)\setminus\zeta\left(\operatorname{Int}C\right),\zeta\left(Z\right)\setminus\zeta\left(\operatorname{Int}C\right)\right)=U\left(\zeta\left(Q\right),\zeta\left(Z\right)\right).$$

Every vertex of $Z\setminus C$ is δ-good so by Lemma 41 ζ induces a combinatorial isomorphism between $V(Q\setminus\operatorname{Int}C, Z\setminus\operatorname{Int}C)$ and $V(\zeta(Q\setminus\operatorname{Int}C),\zeta(Z)\setminus\zeta(\operatorname{Int}C))$ which in turn is the same as $V(\zeta(Q)\setminus\operatorname{Int}\zeta(C),\zeta(Z)\setminus\operatorname{Int}\zeta(C))$ by Lemma 42. It follows from the previous result that $\zeta|_{Q\setminus\operatorname{Int}C}$ extends to a homeomorphism $\bar{\zeta}:\mathbb{T}^d\to\mathbb{T}^d$ so that

$$\begin{aligned}\bar{\zeta}\left(U\left(Q,Z\right)\right)&=\bar{\zeta}\left(U\left(Q\setminus\operatorname{Int}C,Z\setminus\operatorname{Int}C\right)\right)\\&=U\left(\zeta\left(Q\setminus\operatorname{Int}C\right),\zeta\left(Z\setminus\operatorname{Int}C\right)\right)\\&=U(\zeta(Q),\zeta(Z)).\end{aligned}$$

□

5. Coupling construction

In this section, we prove Theorem 1. The coupling we construct is a generalization of that featured in Lemma 20 of [4] to higher dimensions. The arguments of this section are owed directly to Bollobás and Riordan, modified for our setting. Central to their arguments is **potential instability**, which can be interpreted for now as "nearly δ-unstable".

The ingredients of this coupling between (Z_1,R_1) and (Z_2,R_2) are laid out in the following scheme:

(1) A natural coupling where $Z_1=Z_2$, points are included in R_1 independently with probability p_1, and additional points are colored red in R_2 with probability p_2/p_1.

(2) A coupling for a region of potential instability that "crosses over" unstable regions in Z_1 with red and fully stable regions in Z_2.

(3) Independence between each region of potential instability.

We apply (2) to each region of potential instability and (1) elsewhere, and define our final coupling as an independent product of probability spaces by

(3). The nuance of our coupling is due to (2) as it requires moving the points. We begin by defining the natural coupling.

Definition 46 (Natural coupling)**.** *Let $0 < p_1 < p_2 < 1$. We couple (Z, R_1) with (Z, R_2) by coloring each white point of (Z, R_1) red in (Z, R_2) independently with probability p_2/p_1.*

This coupling satisfies the condition $R_1 \subset R_2$. While most points of R_1 will be δ-good, a small portion will be δ-bad and great care must be taken. We must also specify the conditions under which we can make up for these δ-bad regions by increasing p by ϵ.

Throughout this section, we use the following convention: we write that a random variable X is $o(f(N))$ (resp. $O(f(N))$) with probability $1 - o(1)$ to mean that, for any $\eta > 0$ (resp. there exists $\eta > 0$ so that) , $X < \eta f(N)$ with high probability as $N \to \infty$.

Definition 47 (Allocation scheme)**.** *Let l be the closest number to $\delta^{1/5}$ such that $n := (N/l)^d$ is an integer. We may tile the torus with n cubes Λ_i of length l. We couple two Poisson point processes, Z_1 and Z_2, by setting the number of points in Z_2 to be equal to the number of points in Z_1 in each cube Λ_i and placing them independently. We call the assignment of the number of points in each cube an **allocation scheme** and a placement of the correct number of points in each cube a **realization** of the allocation scheme.*

For every coupling of Poisson point processes with the same allocation scheme, there exists an **induced perturbation** $\zeta : Z_1 \to Z_2$ that matches the points in each cube of Z_1 to the corresponding points in Z_2. If a cube has more than one point assigned we match those points arbitrarily.

Lemma 48. *With probability $1 - o(1)$, each Λ_i has $O(1)$ points.*

Proof. Let Λ be a cube of length l. The probability that Λ contains more than a points is bounded above by:

$$\begin{aligned}\mathbb{P}(|\Lambda| \geq a) &\leq \left(\frac{e\|\Lambda\|}{a}\right)^a e^{-\|\Lambda\|} \\ &\leq \frac{e^a l^{ad}}{a^a} \\ &= \Theta(\frac{e^a}{a^a}\delta^{ad/5}).\end{aligned}$$

Choosing $a = 5(d+1)/\epsilon$ yields $\Theta(N^{-(d+1)}) \leq o(N^{-d})$ in the last expression. As there are $O(N^d)$ cubes, the probability that one of them contains more than a points is $O(N^d) \cdot o(N^{-d}) = o(1)$. □

A choice of allocation scheme determines some instabilities inherited by all of its realizations. For example, if two points are to be placed into a cube, then no matter how they are placed, they will always form an $l\sqrt{d}$-close pair. Despite this, the probability that they form a δ-close pair is still quite small as δ is exponentially smaller than $l = \Theta(\delta^{1/5})$. On the other hand, if only one point is placed into a cube and every adjacent cube is empty, then that point cannot form a δ-close pair in any realization. We would like to measure these instabilities on a scale greater than l, we pick $\delta' = \delta^{1/30}$. A crucial step in our argument is to compare δ-instabilities in realizations to δ'-instabilities in allocation schemes.

Definition 49 (Potential instability and potential badness)**.** *A cube in an allocation scheme is* ***potentially*** δ'***-unstable*** *(resp.* ***potentially*** δ'***-bad****) if there is a realization in which a δ'-unstable (resp. δ'-bad) point is contained in the cube.*

Consider the graph on the potentially δ'-bad cubes of an allocation scheme where two cubes are connected by an edge if there exists a realization with a δ'-bad cluster that intersects both cubes. A ***potentially*** δ'***-bad cluster*** *of an allocation scheme is a maximal component of this graph.*

We now introduce the event that makes the coupling of Theorem 1 asymptotically feasible.

Definition 50 (Coupling feasibility)**.** *For $\eta > 0$, we define $E_{\text{good},\eta}$ to be the event that every potentially δ'-bad cluster is $\eta \log N$ in size and diameter, and contains $O(1)$ potentially unstable points. As η is arbitrary, by abuse of notation, we write E_{good} and choose η later.*

We will adapt our arguments from the previous section to show that E_{good} occurs with high probability.

A potentially δ'-unstable cube Λ is potentially δ'-unstable in only three possible cases:

(1) A ball of diameter δ' intersects Λ and at least one other nonempty cube.

(2) A spherical shell of inner radius r and outer radius $r + \delta'$ intersects Λ and at least $(d+1)$ other nonempty cubes

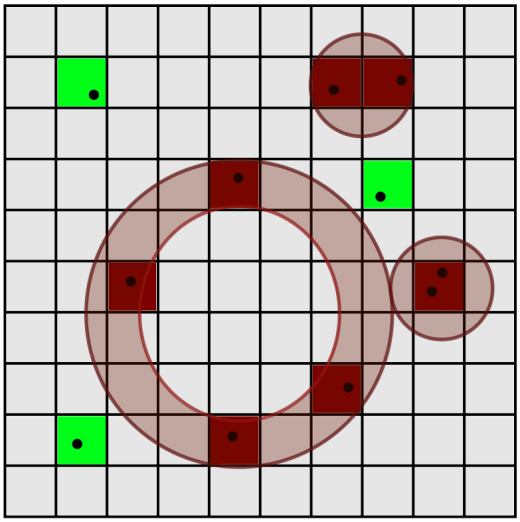

FIGURE 6. An allocation scheme and a realization with a δ'-unstable $(d+2)$-tuple and two δ'-close pairs. The gray cubes have no points, the red cubes are potentially δ'-unstable, and the green cubes are non-empty and not potentially δ'-unstable.

(3) Λ contains more than one point.

Again, we are reiterating the ideas of Section 3.

Lemma 51. *When N is sufficiently large, the points in a potentially δ'-bad cube are $2\delta'$-bad in any realization.*

Proof. First, we argue that the points of a potentially δ'-unstable cube are necessarily $2\delta'$-unstable. We will utilize the facts that $\delta' >> l\sqrt{d} >> \delta$ where $\delta' = \delta^{1/30}$ and $l\sqrt{d} = \Theta(\delta^{1/5})$, as $N \to \infty$.

Case 1: If a ball of diameter δ' intersects two nonempty cubes Λ_1 and Λ_2 then the distance between a point in Λ_1 and a point in Λ_2 is at most

$$\delta' + 2l\sqrt{d} \approx \delta' + 2\delta^{1/5}\sqrt{d} << 2\delta'.$$

Case 2: If a spherical shell of inner radius r and outer radius $r+\delta'$ intersects at least $(d+2)$ nonempty cubes then the points of any realization which are contained in those cubes sit in a spherical shell of inner radius $\max\{r - l\sqrt{d}, \delta'\}$ and outer radius $r + \delta' + l\sqrt{d}$. Thus, they contain as a subset either a $2\delta'$-unstable $(d+2)$-tuple if $r > l\sqrt{d}$ or otherwise a $2\delta'$-close pair.

Case 3: If a cube contains more than one point then it contains a $2\delta'$-close pair.

□

Corollary 52.

$$\mathbb{P}(E_{good}) = 1 - o(1).$$

Proof. By the previous lemma, the points of every potentially bad cluster are contained in a $2\delta'$-bad cluster in every realization. By Lemma 23, every $2\delta'$-bad cluster has $O(1)$ $2\delta'$-unstable points and is $\eta \log N$ in size and diameter with probability $1 - o(1)$. □

Let C be a potentially δ'-bad cluster. For a realization, we say that the points of C are the points placed in cubes contained in C. Such a point is a boundary point if its star contains points in cubes outside of C.

Lemma 53. *If E_{good} occurs, then the boundary points of C are well-defined in the sense that they belong to the same cubes in any realization.*

Proof. Fix an allocation scheme and realization Z. Let C_Z be the set of points in Z that belong to C. Given another realization of the same allocation scheme, Z', there is an induced perturbation $\zeta : Z \to Z'$ where ζ is a $l\sqrt{d} << (\delta')^5$-perturbation. Let C be the set of all δ'-bad points of Z, by E_{good}, $Z \setminus C$ is a $\log N$-sample set. Moreover, as $\mathcal{C}_Z$ is a δ'-bubble wrapped set, we may apply Lemma 42 and obtain: $\zeta(\partial C_Z) = \partial\zeta(C_Z) = \partial C_{Z'}$.

□

Definition 54. *Let $x_1, \ldots, x_j$ be the interior points of C and $y_1, \ldots, y_k$ be the boundary points of C. Define $B(C)$ to be the event that there exist points $y_1', \ldots, y_k'$ in the same l-boxes as $y_1, \ldots, y_k$ so that $C' = \{x_1, ..., x_j, y_1', ..., y_k'\}$ has a δ-bad point.*

By Lemma 53, $B(C)$ only depends on the δ-unstable points of C, which are interior points of C and $O(1)$ in number if E_{good} occurs.

Proposition 55. *On the event E_{good} and N sufficiently large, for every potentially δ-bad cluster C, $\mathbb{P}(G(C)) \geq 2\mathbb{P}(B(C))$. In particular, there exists $G'(C) \subset G(C)$ and a measure-preserving bijection $f_C : B(C) \cup G'(C) \to B(C) \cup G'(C)$ with $f_C(B(C)) = G'(C)$, $f_C(G'(C)) = B(C)$ whose induced perturbation only moves δ-unstable points.*

Proof. Under the event E_{good} occurs, every potentially bad cluster C has $o(\log(N))$ points. In particular, if we set

$$\eta = {}^{-2\epsilon}\!/_{(10(d+3)\log\epsilon)}$$

then C has fewer than $\eta \log(N)$ points for N sufficiently large. Thus, recalling that $\epsilon = p_2 - p_1$, we may compute an upper bound for the probability of $G(C)$ directly:

$$\mathbb{P}(G(C)) = \epsilon^{|C|} \geq \epsilon^{\eta \log N} = N^{\eta \log \epsilon}.$$

Our desired inequality becomes

$$N^{\eta \log \epsilon} \geq 2\mathbb{P}(B(C)).$$

By our choice of η, it will suffice to show that $\mathbb{P}(B(C)) \leq N^{-4\epsilon/10(d+3)}$.

Let C_{unstable} denote the union of all potentially δ'-unstable cubes of C and enumerate the cubes of C_{unstable} as $\Lambda_1, \ldots, \Lambda_T$ arbitrarily. By E_{good} there are $O(1)$ potentially δ'-unstable cubes and they have $O(1)$ points. That is, there exist constants η_1, η_2 so that there are at most η_1 cubes and each has at most η_2 points for all N. Let $1 \leq \lambda_i \leq \eta_2$ be the number of points in the cube Λ_i.

To apply Lemmas 29 and 31, we rescale the torus by by $1/l$ so that the cubes of C_{unstable} become unit cubes. Denote by $\mathcal{C}$ the resulting collection of unit cubes. Then, there is a δ-bad point in the original configuration if and only if there is a δ_1-bad point in the rescaled one where $\delta_1 = \delta/l^d$ where we recall that $l \sim \delta^{1/5}$. Let $t = \sum_{i=1}^{T} \lambda_i \leq \eta_2\eta_1$ be the total number of points. If we place t points into $\mathcal{C}$ uniformly (forgetting about the allocation scheme), the probability of a δ_1-close pair is $O(\delta_1^{1/2})$ by Lemma 29. Moreover, the probability that of a point in a δ_1-unstable $(d+2)$-tuple is $o(\delta_1^{1/2(d+3)})$ by Lemma 31. So, the probability that placing t points i.i.d. uniformly in $\mathcal{C}$ forms a δ_1-bad point is $o(\delta_1^{1/2(d+3)}) = o(\delta^{4/10(d+3)})$. Hence, the probability that placing t points i.i.d uniformly in C_{unstable} creates a δ_1-bad point is $o(\delta^{4/10(d+3)})$. This implies the desired statement, since there are only finitely many ways to place the the points into cubes and the probability of the correct assignment is bounded away from zero by a constant that does not depend on N. □

Now, we are ready to begin constructing our coupling. We do so independently for each allocation scheme. If the event E_{good} fails, we define the coupling arbitrarily. Otherwise, we construct the coupling separately for each potentially δ'-bad cluster. As the assignments of point locations and colors to disjoint potentially bad clusters are independent, we may combine these couplings to define a global coupling for the allocation scheme on the event E_{good}.

More formally, let Z_1, Z_2 Poisson point processes coupled to have the same allocation scheme, and condition on a fixed allocation scheme $\mathcal{A}$. and enumerate

the clusters of potentially δ'-bad cubes as C_j for $j = 1, \ldots, n$. Let D_j denote the union of the cubes in C_j, and D_0 be the union of the remaining cubes. , let $Z_j^k = Z_j \cap D_k$ and $R_j^k = R_j \cap D_k$, $j = 1, 2$. As each Z_j^k are disjoint, we can construct a coupling independently for each Z_j^k, $k = 1, \ldots, n$. This partitions our probability space Ω into a product of probability spaces, $\Omega_0 \times \ldots \times \Omega_n$, where each Ω_j is the assignment of point locations and colors to the points in D_j for $j = 0, \ldots, n$. Lastly, the measure-preserving bijection in the previous proposition, $f_{C_j} : B(C_j) \cup G'(C_j) \to B(C_j) \cup G'(C_j)$, may be equivalently defined on the event space of each Ω_j. We write f_j to be this measure-preserving bijection on Ω_j.

Definition 56 (Coupling on a cluster)**.** *Let Z_1, Z_2 be as in the previous paragraph, let $0 < p_1 < p_2 < 1$, and assume that $\mathcal{A} \in E_{\text{good}}$. Let $(Z_1, R_1^*(p_1))$ and $(Z_2^*, R_2^*(p_2))$ be given by the natural coupling with the allocation scheme $\mathcal{A}$. We construct a coupling of (Z_1^k, R_1^k) with (Z_2^k, R_2^k) over the probability space $(\Omega_j, \mathbb{P}_j)$ by*

$$(Z_2^k, R_2^k) = \begin{cases} (Z_2^*(\omega), R_2^*(\omega) & \omega \notin B(C_k) \cup G'(C_k) \\ (Z_2^*(f_k(\omega)), R_2(f_k(\omega)) & \omega \in B(C_k) \cup G'(C_k). \end{cases}$$

In addition, define Q_k by

$$W_k = \begin{cases} Z_k^2 & \text{if } \omega \in B(C_k) \\ R_k^2 & \text{else.} \end{cases}$$

We discuss the explicit properties of the coupling below.

Let ζ denote an induced perturbation from Z_1^k to Z_2^k. For every $z \in Z_1^k$ and $\omega_k \in \Omega_k$

(1) $(\omega \in B(C_k) \implies f(\omega) \in G'(C_k))$.
If z is δ-bad, then $\zeta(z)$ is red and not δ-bad. Moreover, if z is in ∂C_k, then $\zeta(z) = z$. Lastly, $W_k = Z_2^k$ and is δ-good.

(2) $(\omega \in G'(C_k) \implies f(\omega) \in B(C_k))$.
If $G'(C)$ occurs, then z is white and $\zeta(z)$ is possibly δ-bad. Similarly, if z is in ∂C_k, then $\zeta(z) = z$. Lastly, $W_k = \varnothing$.

(3) $(\omega \notin B(C_k) \cup G'(C_k))$.
We take $\zeta = \text{id}$. $\zeta(z) = z$ and has the same color. If z is white, then $\zeta(z)$ is red with probability p_2/p_1. Lastly, $W_k = R_1^k$ and is δ-good.

Finally, we extend this coupling to the entire product space $\Omega_0 \times \Omega_n$. The coupling is a natural coupling over Ω_0 and a cluster coupling over Ω_k for $k = 1, \ldots, n$.

Definition 57 (Final coupling)**.** *Let Z_1, Z_2 Poisson point processes coupled to have the same allocation scheme $\mathcal{A}$. If $\mathcal{A} \notin E_{\text{good}}$, define the coupling arbitrarily. Otherwise, let $f_0 : \Omega_0 \to \Omega_0$ be the natural coupling and let $f_k : \Omega_k \to \Omega_k$ be the map defined in Definition 56. Define $f : \Omega_0 \times \ldots \times \Omega_n \to \Omega_0 \times \ldots \times \Omega_n$ by $f = \prod_{j=0}^{k} f_j$.*

We now check that this final coupling indeed satisfies the properties of Theorem 1. We denote by ζ the point perturbation induced by f. Throughout the proof, we also explicitly construct the point set W.

Note: The statement of the theorem is for $\epsilon_0, \delta_0 \leq N^{-\epsilon_0}$ and they satisfy the hypotheses of the couplings. This changes nothing and avoids notational ambiguity.

Proof of Theorem 1. We may assume that E_{good} occurs by Corollary 52.

Let $W = \bigcup W_i$, where W_i is defined above in Definition 56. Recall by Lemma 53, the boundary is well-defined, by Proposition 55 the construction of f yields an induced perturbation ζ does not move boundary points, and by Corollary 40, $U(Q, Z) = U(Q \setminus \operatorname{Int} \mathrm{Q}, Z \setminus \operatorname{Int} Q)$. Putting these altogether, we have the following analysis for each C_k with $k > 0$:

(1) If $\omega_k \in B(C_k)$, then $\zeta(R_1) \subset R_2^j =: W_i$ is δ-good. Moreover,

$$U(R_1^k, Z_1) \subset U(Z_1^k, Z_1) = U(Z_1^k \setminus \operatorname{Int} Z_1^k, Z_1 \setminus \operatorname{Int} Z_1^k) = U(W, Z_2) = U(R_2^k, Z_2).$$

The second equality uses $Z_1^k \setminus \operatorname{Int} Z_1^k = Z_2^k \setminus \operatorname{Int} Z_2^k$ and that no point of ∂C_k is moved for all k.

(2) If $\omega_k \in G'(C_k)$ then $R_1^k = \varnothing =: W_k$ and

$$U(R_1^k, Z_1) = U(W_k, Z_1) = \varnothing \subset U(R_2^k, Z_2).$$

(3) If $\omega_k \notin B(C_k) \cup G'(C_k)$ then $Z_1^k = Z_2^k$, $\zeta|_{Z_1^k} = \mathrm{id}$ and $W_k = R_1^k$ is δ-good.

$$U(R_1^k, Z_1) = U(W_k, Z_1) \subset U(R_2^k, Z_2).$$

Finally $Z_1^0 = Z_2^0$ and $R_1^0 \subset R_2^0$ so $U(R_1^0, Z_1) \subset U(W_0, Z_2) = U(R_1^0, Z_2)$. □

6. Homological Percolation at $d = 2i$.

6.1. **Homological percolation ingredients.** Consider Voronoi percolation $P = (Z, R)$ on $\mathbb{T}_N^d$. The dual percolation is the union of white Voronoi cells, $P^\bullet = (Z, Z \setminus R)$; it is a Voronoi percolation with parameter $1 - p$. Since Z is in general position, the proof of Lemma C.1 in [1] yields that $U(Z \setminus R, Z)$ is homotopy equivalent to $\operatorname{Int} U(Z \setminus R, Z) = \mathbb{T}_N^d \setminus P$. We may then infer a powerful relation between P and $P^\bullet$ using Lemmas C.1 and C.2 from [1].

Lemma 58 (Bobrowski and Skraba)**.** *Let $\phi_i : H_i(P, \mathbb{F}) \to H_i(\mathbb{T}_N^d, \mathbb{F})$ and $\psi_i : H_i(P^\bullet, \mathbb{F}) \to H_i(\mathbb{T}_N^d, \mathbb{F})$ be the maps on homology induced by inclusion. Then*

$$\operatorname{rank} \phi_i + \operatorname{rank} \psi_{d-i} = \operatorname{rank} H_i(\mathbb{T}^d)\,.$$

In particular, A occurs either in the dual or the primal. Similarly, if S occurs in the dual, then A cannot occur in the primal and vice-versa.

From this, we obtain the higher dimensional analogue of the result "for any square on a Voronoi percolation, the probability of a horizontal crossing at $p = 1/2$ is at least $1/2$".

Corollary 59.

$$\mathbb{P}_{1/2}(A) \geq 1/2\,.$$

We also borrow two more technical results from [6] that allow us to compare the probability of S to A. We define the coarse state once again for easy reference.

Definition 60 (Coarse state)**.** *Let $\{\Lambda_i\}$ be a disjoint covering of $\mathbb{T}_N^d$ by cubes of length δ. Let (Z, R) be a Voronoi percolation. We define* $\operatorname{Coarse}_\delta(\cdot) : (Z, R) \mapsto \{-1, 0, 1\}^n$ *by, for each cube Λ_i,*

$$\Lambda_i \mapsto \begin{cases} -1 & \text{if } \Lambda_i \text{ contains a white point} \\ 0 & \text{if } \Lambda_i \text{ contains no points} \\ 1 & \text{if } \Lambda_i \text{ contains only red points.} \end{cases}$$

The first technical result involves an application of representation theory.

Definition 61 (Irreducible representation)**.** *A vector space V acted on by a group G is an irreducible representation of G if the only subspaces of V are invariant under G or are trivial subspaces.*

Lemma 62 (Duncan, Kahle, Schweinhart)**.** *Let V be a finite dimensional vector space and Y be a set. Let $\mathcal{A}$ be a lattice of subspaces of V. Suppose $f : \{-1, 0, 1\}^{|Y|} \to \mathcal{A}$ is an increasing function, i.e. if $A \subset B$ then $f(A) \subset f(B)$. Let G be a finite group which acts on both Y and V whose action is compatible with f. That is, for each $g \in G$ and $D \in \mathcal{P}(Y)$ $g(f(D)) = f(gD)$. Let X be a $\{-1, 0, 1\}^{|Y|}$-valued random variable with a G-invariant distribution that is positively associated, meaning that increasing events with respect to X are non-negatively correlated. Then if V is an irreducible representation of G, there are positive constants C_0, C_1 so that*

$$\mathbb{P}_p(f(X) = V) \geq C_0 \mathbb{P}_p(f(X) \neq 0)^{C_1},$$

where C_0 only depends on G and C_1 only depends on $\dim V$*.*

We have modified the above lemma slightly, using $\{-1, 0, 1\}^{|Y|}$ to be the domain of f instead of the power set $\mathcal{P}(Y)$. However, the proof is the same, *mutatis mutandis*.

We apply the above lemma to events defined in terms of "stable" giant cycles. Let Ω' be the collection of all possible coarse states of Voronoi percolations on $\mathbb{T}^d_N$. Note that Ω' is of the form $\{-1, 0, 1\}^{|Y|}$.

Definition 63 (Homology of a coarse state)**.** *Let $CS \in \Omega'$ be a coarse state and $P_{CS} = \{P : \text{Coarse}_\delta(P) = CS\}$ the collection of Voronoi percolations with the same coarse state. We define $f : \Omega' \to H_i(\mathbb{T}^d_N, \mathbb{F})$ by $f(CS) = \cap_{P \in P_{CS}} \phi_i(P)$. We call the set, $f(CS)$, the* ***stable giant i-cycles*** *of CS.*

In order to apply the lemma, with $V = H_i(\mathbb{T}^d_N; \mathbb{F})$, Y the cubes in the coarse state, X the coarse state of a Voronoi percolation, f as in the previous definition, and $G = W_d$ the point symmetry group of the torus. It remains to show that V is indeed an irreducible representation of G. Luckily, this result is provided in [6] as well.

Proposition 64. *Let $\mathbb{F}$ be a field, $d > 0$, and $1 \leq i \leq d - 1$. $H_i(\mathbb{T}^d_N; \mathbb{F})$ is an irreducible representation of W_d if and only if* $\text{char}(\mathbb{F}) \neq 2$*.*

Define the events $f(X) = V$ and $f(X) \neq 0$ as S_{discrete} and A_{discrete} respectively. Thus, we obtain the following uniform bound on $\mathbb{P}_p(S_{\text{discrete}})$.

Corollary 65. *There are constants $C_0, C_1 > 0$ not depending on N, i such that*

$$\mathbb{P}_p(S_{\text{discrete}}) \geq C_0 \mathbb{P}_p(A_{\text{discrete}})^{C_1}.$$

6.2. **A sharp transition.** With Theorem 1 in hand, the proof is a straightforward modification of the arguments of [6] which substitutes Bollobás and Riordan's version of Theorem 18 for that of [8].

Our goal now is to discretize our events into events over the coarse state so that we may apply Theorem 18. In the previous subsection, we did this in our "homology of a coarse state" definition, where the coarse state homology is the common homology of all the percolations with that coarse state.

Definition 66 (Stable events and discrete events)**.** *Let X be an event for Voronoi percolation. Define $X_{\text{discrete}} \subset \Omega'$ to be the event that X occurs for every Voronoi percolation with a given coarse state. Similarly, X_{stable} occurs for $P = (Z, R)$ if $\text{Coarse}_\delta(P) \in X_{\text{discrete}}$.*

Note that this definition agrees with the one previously given for A_{discrete} and S_{discrete}.

Aside from sharing the same name, stable events and stable points are related. For a fixed percolation, $\mathcal{P}$, if an event depends only on the topological features of $\mathcal{P}$ in $\mathbb{T}_N^d$ and depends on a set of $\delta^{1/6}$-good points in $\mathcal{P}$, then X will still hold under any δ-perturbation by Theorem 36. Subsequently, we formally introduce an event that depends only on topological features.

Definition 67 (Events invariant under homeomorphisms of the torus)**.** *Let $\mathcal{P}$ be a percolation, $\varphi : \mathbb{T}^d \to \mathbb{T}^d$ a homeomorphism, and $\varphi(\mathcal{P})$ another percolation. We say an event X is invariant under homeomorphisms of the torus if $\mathcal{P} \in X$ implies $\varphi(\mathcal{P}) \in X$.*

We express the previous paragraph now compactly as a Corollary of Theorem 36.

Corollary 68. *Let X be an increasing event that is invariant under homeomorphisms of the torus. If there exists an $R' \subset R$ that is $\delta^{1/6}$-stable and X occurs for (Z, R'), then X_{stable} occurs.*

In this setting, we apply Theorem 1 as follows.

Lemma 69. *Let X be an increasing event that is invariant under homeomorphisms of the torus. Then,*

$$\mathbb{P}_p(X_{\text{stable}}) \leq \mathbb{P}_p(X) \leq \mathbb{P}_{p+\epsilon}(X_{\text{stable}}) + o(1)\,.$$

Proof. The first inequality is trivial as X_{stable} occurs if X occurs. The second inequality follows from Theorem 1 (by using $\delta_0 = \delta^{1/6}$ in the theorem) and the previous Corollary. That is, if X occurs, as $\zeta(R_1)$ is $\delta^{1/6}$-stable, then X_{stable} occurs for (Z_2, R_2). □

The existence of an i-dimensional giant cycle, the event A, is an increasing event that is invariant under homeomorphisms of the torus. So, the above lemma applies for A, which we state as a corollary below.

Corollary 70.

$$\mathbb{P}_p(A_{\text{stable}}) \leq \mathbb{P}_p(A) \leq \mathbb{P}_{p+\epsilon}(A_{\text{stable}}) + o(1)\,.$$

We define a proper probability space on the set of coarse states.

$$\begin{cases} p_{bad} & = 1 - \exp(-\|B\|(1-p)) \approx \|B\|(1-p), \\ p_{neutral} & = \exp(-\|B\|), \\ p_{good} & = \exp(-\|B\|(1-p))(1-\exp(-\|B\|p)) \approx \|B\|p\,. \end{cases} \tag{3}$$

Where $\|B\| = \delta^d$, the size of each cube. Then, as the group of translations on the cubes transitively acts on $\{-1, 0, 1\}^n$, we are ready to apply Theorem 18.

$$\mathbb{P}_p(X_{\text{stable}}) = \mathbb{P}_{p_{bad}, p_{good}}(X_{\text{discrete}}).$$

Lemma 71. *If X is an increasing event invariant under homeomorphisms of the torus so that $\mathbb{P}_p(X_{\text{stable}}) > \eta$ for some $\eta > 0$, then $\mathbb{P}_q(X_{\text{stable}}) \to 1$ for $q > p$.*

Proof. By construction

$$0 < \mathbb{P}_p(X_{\text{stable}}) = \mathbb{P}_p(X_{\text{discrete}}).$$

Let $0 < \eta < \mathbb{P}_p(X_{\text{discrete}})$ and $q > p$. Define q_{bad}, q_{good} by substituting q for p in Equation 3. As a reminder, n is the number of cubes in the coarse state and N is the length of the torus.

We proceed by establishing the inequalities of Theorem 18. Throughout this proof, recall that $\delta = N^{-\epsilon}$ and $\|B\| = \nu_d \delta^d$ Firstly, we can demonstrate inequality 1 — $0 < q_{bad} < p_{bad} < 1/e$ and $0 < p_{good} < q_{good} < 1/e$ — can quickly be observed as

$$\begin{aligned} 0 < \|B\|(1-q) < \|B\|(1-p) < 1/e \\ 0 < \|B\|p < \|B\|q < 1/e. \end{aligned}$$

To show inequality 2, observe that $q_{good} - p_{good}, p_{bad} - q_{good} \sim (q-p)\|B\|$ and $p_{\max} = \max\{\|B\|(1-p), \|B\|q\}$. Moreover, using $n = (\lceil \frac{N}{\delta} \rceil)^d$, the right-hand-side of Equation 2 which we denote Δ has the following upper-bound.

$$\begin{aligned}
\Delta &:= c\log(1/\eta) p_{\max} \log(1/p_{\max})/\log n \\
&\leq \left(c\log(1/\eta)\right)\left(p_{\max}(-\log p_{\max})/\log\left((N/\delta)^d\right)\right) \\
&= \left(c\log(1/\eta)\right)\left((\nu_d N^{-d\epsilon}\max\{1-p,q\})(-\log(\nu_d N^{-d\epsilon}\max\{1-p,q\}))/((d+\epsilon)\log N)\right) \\
&= \left(\nu_d c\log(1/\eta)\max\{1-p,q\})\right)\left((N^{-d\epsilon}(d\epsilon\log N - \log\nu_d \max\{1-p,q\})/((d+\epsilon)\log N)\right) \\
&\leq N^{-dr}\left(C_1\right)\left(\frac{d\epsilon}{d+\epsilon} + \frac{C_2}{\log N}\right) \\
&= C_1\epsilon\|B\| + C_2\frac{\|B\|}{\log N},
\end{aligned}$$

for some constants $C_1, C_2 > 0$. As $\Delta \leq C\epsilon\|B\|$ and $\min\{q_{good} - p_{good}, p_{bad} - q_{bad}\} \sim (q-p)\|B\|$, taking ϵ small enough and N large enough so that $C_1\epsilon + \frac{C_2}{\log N} < q-p$ yields inequality 2 of Theorem 18. Therefore, $\mathbb{P}_{q_{good},q_{bad}}(X_{\text{discrete}}) > 1-\eta$. It follows that $\mathbb{P}_q(X_{\text{stable}}) = \mathbb{P}_{q_{good},q_{bad}}(X_{\text{discrete}}) \xrightarrow{N\to\infty} 1$. □

Corollary 72. *If X is an increasing event invariant under homeomorphisms of the torus and $\mathbb{P}_p(X_{\text{stable}}) > 0$, then $\mathbb{P}_q(X) \to 1$ for $q > p$.*

Corollary 73. *For $p > 1/2$,*

$$\mathbb{P}_p(A) \to 1.$$

Moreover, by Corollary 65, $\mathbb{P}_p(S)$ is uniformly bounded away from 0.

$$\mathbb{P}_p(S) \to 1.$$

When S occurs, the giant cycles block the giant cycles of the dual. That is, by Lemma 58, we infer the following.

Corollary 74. *For $p < 1/2$,*

$$\mathbb{P}_p(A) \to 0$$

7. Acknowledgments

We would like to thank Paul Duncan, Ron Peled, and Anthony Pizzimenti for feedback and helpful comments.

References

[1] Omer Bobrowski and Primoz Skraba. Homological percolation and the Euler characteristic. *Physical Review E*, 101(3):032304, 2020.
[2] Omer Bobrowski and Primoz Skraba. Homological percolation: The formation of giant k-cycles. *International Mathematics Research Notices*, 2022(8):6186–6213, 2022.
[3] Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh. The stability of Delaunay triangulations. *International Journal of Computational Geometry & Applications*, 23(04n05):303–333, 2013.
[4] Béla Bollobás and Oliver Riordan. The critical probability for random Voronoi percolation in the plane is 1/2. *Probability Theory and Related Fields*, 136(3):417–468, December 2005.
[5] Gilles Bonnet and Nicolas Chenavier. The maximal degree in a Poisson–Delaunay graph. *Bernoulli*, 26(2):948–979, 2020.
[6] Paul Duncan, Matthew Kahle, and Benjamin Schweinhart. Homological percolation on a torus: plaquettes and permutohedra. In *Annales de l'Institut Henri Poincare (B) Probabilites et statistiques*, volume 61, pages 2235–2261. Institut Henri Poincaré, 2025.
[7] Herbert Edelsbrunner. *A Short Course in Computational Geometry and Topology*. Springer, 2014.
[8] Ehud Friedgut and Gil Kalai. Every monotone graph property has a sharp threshold. *Proceedings of the American Mathematical Society*, 124(10):2993–3002, 1996.
[9] A. Hatcher. *Algebraic Topology*. Algebraic Topology. Cambridge University Press, 2002.
[10] D. Kozlov. *Combinatorial Algebraic Topology*. Algorithms and Computation in Mathematics. Springer Berlin Heidelberg, 2008.

Email address: bschwei@gmu.edu

Department of Mathematical Sciences, George Mason University, Fairfax, VA 22030, USA

Email address: mshuman4@gmu.edu

Department of Mathematical Sciences, George Mason University, Fairfax, VA 22030, USA