The simple random walk on a random Voronoi tiling

L. Addario-Berry*, A. Sarkar†

October 18, 2005

Abstract

Let $\mathcal{P}$ be a Poisson point process in $\mathbb{R}^d$ with intensity 1. We show that the simple random walk on the cells of the Voronoi diagram of $\mathcal{P}$ is almost surely recurrent in dimensions $d = 1$ and $d = 2$ and is almost surely transient in dimension $d \geq 3$.

1 Introduction

The aim of this work is to extend classical results for random walks on the d-dimensional lattice $\mathbb{Z}^d$ to various random graphs. Our main results concern the *Voronoi tilings* of Poisson point processes in $\mathbb{R}^d$: however, we expect that the method will work for many other random geometric graphs, such as the k-nearest neighbour graph, and the Gilbert disc model. Our results constitute a first step towards understanding the large scale behaviour of random walks on these random graphs.

Given a Poisson point process $\mathcal{P}$ in $\mathbb{R}^d$, the *Voronoi cell* of a point $p \in \mathcal{P}$ is the set of all points of $\mathbb{R}^d$ closer to p than to any other point in $\mathcal{P}$; the Voronoi diagram is the set of Voronoi cells of $\mathcal{P}$.

The *Delaunay triangulation* $DT(\mathcal{P})$ is the facial dual of the Voronoi diagram - thus a simple random walk on the points of a Delaunay triangulation is equivalent to a simple random walk on the cells of the Voronoi diagram. At

*Department of Computer Science, McGill University, Montreal, Canada

†Department of Mathematical Sciences, University of Memphis, Memphis, USA

each step of the walk, we move from a point $p \in \mathcal{P}$ to one of the points $p' \in \mathcal{P}$ whose Voronoi cell shares a $d-1$-dimensional face with the Voronoi cell of p - equivalently, the edge pp' must belong to $DT(\mathcal{P})$. We choose p' among the neighbours of p uniformly at random. We hereafter focus on the Delaunay triangulation; all results immediately follow for the Voronoi diagram.

A useful equivalent definition of $DT(\mathcal{P})$, stated for $d = 2$ but easily generalized to higher dimensions, is the following. Given points x, y, z of $\mathcal{P}$, $\triangle_{xyz}$ is a triangle of $DT(\mathcal{P})$ if and only if the unique disc with x, y and z on its boundary contains no points aside from x, y and z. At first sight this seems somewhat unwieldy - however, it turns out to be very useful in practice, since it shows that a long edge in $DT(\mathcal{P})$ corresponds to a large empty region in $\mathbb{R}^d$, that is, a region containing no points of $\mathcal{P}$.

In this paper, we solve the problem of whether the simple random walk on $DT(\mathcal{P})$ is recurrent or transient. As with many other questions in this area, it is not hard to guess the correct answer - recurrence for $d \leq 2$ and transience for $d \geq 3$ - the difficulty lies in providing a rigorous proof. The intuition is that the Delaunay triangulation is "essentially" a lattice, and so the results on $DT(\mathcal{P})$ should follow from those for $\mathbb{Z}^d$. However, the stochastic nature of $DT(\mathcal{P})$ makes direct comparison with $\mathbb{Z}^d$ difficult. (In [4], for example, it is proved that the critical site percolation probability for $DT(\mathcal{P})$ is $1/2$; much of the difficulty of that proof lies in finding ways to rederive tools developed for deterministic lattices in this random setting.) Furthermore, the result does not follow easily from the deep theorems of Thomassen [10].

Instead, we imitate the *proof* of the result for $\mathbb{Z}^d$, as presented in [6]. For $d \leq 2$ we utilize a bound on the *stabbing number* of $DT(\mathcal{P})$ obtained by Addario-Berry, Broutin and Devroye [1] in carrying out our plan. For $d \geq 3$, the picture is still more complicated: in addition to the result from [1] we require both a result of Grimmett, Kesten and Zhang [7] on random walks in percolation clusters and the use of a certain dependent percolation model, whose study was initiated by Liggett, Schonmann and Stacey [9] and continued by Balister, Bollobás and Walters [2]. (For our purposes, we only need the simple bound from [9].)

Both the proof of the recurrence of the simple random walk in $\mathbb{Z}^2$ and that of the transience of the walk in $\mathbb{Z}^3$ presented in [6] exploit the relationship between random walks and electrical networks. Briefly, given any graph G, we replace each edge by a resistor of unit resistance, and maintain a fixed potential difference between two vertices of G (or, more generally, disjoint sets of vertices of G). The currents and potential differences in the resulting electrical network have probabilistic interpretations in terms of the simple random walk on G, and, reciprocally, the walk on G can be studied in terms of the electrical network. In particular, fix a vertex v of an infinite graph G, short together all vertices at graph distance n from v to form a single vertex v_n, remove all vertices at distance greater than n and write $R_n(v)$ for the effective resistance between v and v_n in the new network. By Rayleigh's monotonicity law, the *effective resistance from v to infinity* $R_\infty(v) = \lim_{n\to\infty} R_n(v)$ exists for all v, and it is known that the simple random walk on G, starting at v, is transient if and only if $R_\infty(v) < \infty$. This, together with Rayleigh's monotonicity law that the effective resistance of a network increases upon the removal of a resistor, will be our main tool in what follows.

In $\mathbb{Z}^d$, if vertices v and w satisfy $\|v - w\|_\infty = k$, then there is a path of length at most kd between v and w. We will need similar information about $DT(\mathcal{P})$ in $\mathbb{Z}^3$ in analyzing its effective resistance to infinity. We obtain this by way of the *stabbing number* of $DT(\mathcal{P}) \cap [0,n]^d$. This number, denoted $st_n(DT(\mathcal{P}))$, is the maximum number of Delaunay cells that intersect a single line in $DT(\mathcal{P}) \cap [0,n]^d$, where the maximum is taken over all lines in $\mathbb{Z}^d$. It is easy to see that if $st_n(DT(\mathcal{P})) \leq K$, then there is a path between any two points of $DT(\mathcal{P}) \cap [0,n]^d$ of length at most K. Thus, bounds on the stabbing number will yield information on the graph distance between points of $DT(\mathcal{P})$. In the 2-dimensional case, bounds on the stabbing number also provide bounds on the number of edges of $DT(\mathcal{P})$ *leaving* $[0,n]^2$, which will prove useful in Section 2.

2 Recurrence in $\mathbb{R}^2$

The k^{th} annulus A_k is the set of points of $\mathbb{R}^2$ with L_∞-norm r, for $(k-1) \leq r < k$. Following Doyle and Snell [6], our strategy will be to short together all points contained within each annulus. This yields a network with resistance to infinity strictly less than in the original network. Ideally, the new network would simply be a path of resistors with resistances $r_1, r_2, \ldots$, such that $r_k \geq c/k$ for some fixed constant c. This network trivially has infinite resistance to infinity, establishing the theorem. In fact, due to the existence of long edges in $DT(\mathcal{P})$, reducing the original network to a path is slightly more involved than in the lattice.

Write A'_k for the set of points with L_∞-norm k, i.e., the outer boundary of

A_k. We will need bounds on the *maximum length* of an edge that crosses such a set A'_k, and on the *number* of such edges. Lemmas 1 and 3, respectively, give us the required bounds on these two quantities.

Lemma 1. *There exists a fixed constant c such that for all $k > 1$, $r > 1$,*

$$\mathbb{P}\left(A'_k \text{ is crossed by an edge of length } \geq cr\sqrt{\log k}\right) \leq e^{-r^2}.$$

To prove this lemma, we use the following fact:

Fact 2. *If e is an edge of $\mathrm{DT}(\mathcal{P})$, then one of the half-circles with diameter e contains no points of $\mathcal{P}$.*

Proof. If e is an edge of $DT(\mathcal{P})$ then, by the definition of the Delaunay triangulation given in the introduction, there is an empty circle that has e as a chord. Such a circle necessarily contains one of the two half-circles with diameter e. □

Proof of Lemma 1. Denote by B_i the box $A_1 \cup \ldots \cup A_i$. Fix $k > 1$ and $r > 1$ arbitrarily.

Let $E(k,r)$ be the event that A'_k is crossed by an edge of length $\geq cr\sqrt{\log k}$. Let D_t be the event that there exists an edge e with one endpoint in B_k of length between $ct\sqrt{\log k}$ and $2ct\sqrt{\log k}$. Trivially, such an e has its second endpoint in $B_{k'}$, where $k' = \lceil k + 2ct\sqrt{\log k}\rceil$. Note that $E(k,r) \subset \bigcup_{i=0}^{\infty} D_{2^i r}$. We will prove that

$$\mathbb{P}(D_t) \leq e^{-t^2}/2. \tag{1}$$

From this it follows that

$$\begin{aligned} \mathbb{P}(E(k,r)) &\leq \sum_{i=0}^{\infty} \mathbb{P}(D_{2^i r}) \\ &\leq \sum_{i=0}^{\infty} e^{-2^{2i} r^2}/2 \\ &\leq e^{-r^2}. \end{aligned}$$

It thus remains to prove (1).

Let k' be defined as above and let $m < k'$ be as large as possible such that $\lceil ct\sqrt{\log k}/4\rceil$ divides m. Let $l = m/\lceil ct\sqrt{\log k}/4\rceil$ and partition B_m into boxes $Q_1, \ldots, Q_{l^2}$ of side length $\lceil ct\sqrt{\log k}/4\rceil$. It is an easy exercise (based on Fact 2) to see that if D_t holds then one of $Q_1, \ldots, Q_{l^2}$ must be empty. In what follows we use the very crude bound

$$l^2 \leq (k')^2 \leq 5 \max\{k^2, (2ct\sqrt{\log k})^2\} \leq 20k^2c^2t^2 \log k =_{\text{def}} R$$

which holds for sufficiently large c and which follows trivially by expanding $(k')^2$. By linearity of expectation (of the number empty boxes) we thus have

$$\begin{aligned}
\mathbb{P}(D_t) &\leq l^2 \mathbb{P}(Q_1 \text{ is empty}) \\
&\leq R \cdot \exp(-(ct)^2 \log k/16) \\
&= e^{-t^2} \left(\frac{20k^2c^2t^2 \log k}{\exp(t^2(c^2 \log k/16 - 1))} \right) \\
&\leq e^{-t^2}/2,
\end{aligned}$$

provided $\exp(t^2(c^2 \log k/16 - 1)) \geq 40k^2c^2t^2 \log k$, which certainly holds for sufficiently large c. □

The following lemma is a weakening of Theorem 1 from [1], and will also be used in the 3-dimensional case.

Lemma 3. *Fix $d \geq 1$. Then there are constants $\kappa = \kappa(d)$, $K = k(d)$ such that*

$$\mathbb{E}(\text{st}_n(\text{DT}(\mathcal{P}))) \leq \kappa n, \tag{2}$$

and, for any $\alpha > 0$,

$$\mathbb{P}(\text{st}_n(\text{DT}(\mathcal{P})) \geq (\kappa + \alpha)n) \leq e^{-\alpha n/K \log n}. \tag{3}$$

Using these two lemmas, we can now prove:

Theorem 4. *The simple random walk on $\text{DT}(\mathcal{P})$ is recurrent in dimensions one and two, with probability one.*

Proof. In one dimension, the simple random walk on $DT(\mathcal{P})$ is just the simple random walk on $\mathbb{Z}$, which is recurrent.

Given an edge $e = uv$ with $u \in A_i$ and $v \in A_j$ for $j > i$, divide e into $j - i$ resistors in series, each of resistance $1/(j-i)$ and such that the k^{th} resistor has endpoints in the annuli A_{i+k-1} and A_{i+k}. This network is clearly equivalent to the original.

We now define a new network by shorting together all the points in each annulus A_k, yielding a path of resistors $R_1, R_2, \ldots$ in series, having resistances

$r_1, r_2, \ldots$. The resistance of this new network is at most the resistance of the original network. It thus remains to prove that with probability one,

$$\sum_{i=1}^{\infty} r_i = \infty.$$

To calculate r_i, we split the edges crossing A_i' into groups based on their length. Let

$$E_j(i) = \{e \in E | e = uv, u \in A_{k_1}, v \in A_{k_2}, k_1 \leq i < k_2 \text{ and } k_2 - k_1 = j\}.$$

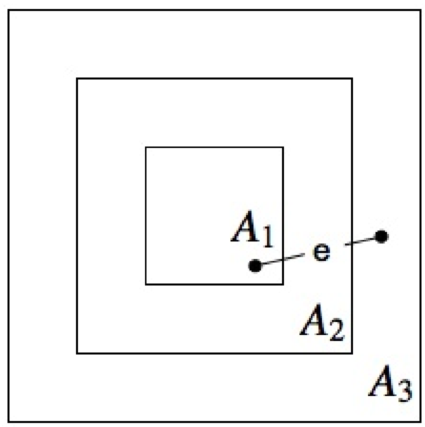

Figure 1: Annuli A_1 through A_3. Edge e is in $E_2(1)$ and $E_2(2)$ but not $E_2(3)$ as it does not leave $B_3 = A_1 \cup A_2 \cup A_3$.

In other words, $E_j(i)$ is the set of edges with one endpoint in B_i, one endpoint outside of B_i, and crossing or touching $j+1$ annuli in total. This is pictured in Figure 1. In the new network, such an edge will have been subdivided into j edges, each of resistance $1/j$. By the resistance rule for resistors in parallel, this allows us to write

$$\frac{1}{r_i} = \sum_{j=1}^{\infty} \sum_{e \in E_j(i)} \frac{1}{1/j} = \sum_{j=1}^{\infty} j|E_j(i)|.$$

Fix $\epsilon > 0$. We show that there exists a large fixed constant L such that for all i,

$$\mathbb{P}\left(\sum_{j=1}^{\infty} j|E_j(i)| \geq Li \log i\right) \leq \frac{1}{i^2}. \tag{4}$$

Assuming (4) holds, it follows that for some i_0,

$$\mathbb{P}\left(\exists i \geq i_0 \quad \text{with } \sum_{j=1}^{\infty} j|E_j(i)| \geq Li \log i\right) \quad \leq \quad \sum_{i=i_0}^{\infty} \frac{1}{i^2} < \epsilon.$$

Therefore, with probability greater than $1-\epsilon$,

$$\sum_{i=1}^{\infty} r_i \geq \sum_{i=1}^{i_0-1} r_i + \sum_{i=i_0}^{\infty} \frac{1}{Li \log i} = \infty.$$

It thus remains to prove (4), for which we use Lemmas 1 and 3. Let $E(k,r)$ be defined as in Lemma 1, and set $F_i = E(i, \sqrt{2\log i + 1})$ - in other words, F_i is the event that A_i' is crossed by an edge of length more than $r_i = c(2\log i + 1)^{1/2}(\log i)^{1/2}$. By Lemma 1,

$$\mathbb{P}(F_i) \leq e^{-(\sqrt{2\log i+1})^2} \leq \frac{1}{ei^2}.$$

Let G_i be the event that $\sum_{i=1}^{\infty} |E_j(i)| \geq 8(\kappa + K)i$, where κ and K are as defined in Lemma 3. Notice that $\sum_{i=1}^{\infty} |E_j(i)|$ is the total number of edges crossing A_i', the outer boundary of A_i.

As A_i' is composed of four line segments, if G_i occurs then necessarily one side of A_i' crosses at least $2(\kappa + K)i$ edges. Since each edge is in precisely two Delaunay cells, this implies $st_n(DT(\mathcal{P})) \geq (\kappa + K)i$. By Lemma 3, we thus have that

$$\mathbb{P}(G_i) \leq \mathbb{P}(st_n(DT(\mathcal{P})) \geq (\kappa + K)i) \leq e^{-i/\log i} \leq \frac{1}{2i^2}$$

for all integers $i \geq 1$. Thus, $\mathbb{P}(F_i \cup G_i) \leq 1/ei^2 + 1/2i^2 < 1/i^2$, and assuming $F_i \cup G_i$ does not occur, we have

$$\begin{aligned}
\sum_{j=1}^{\infty} j|E_j(i)| &= \sum_{j=1}^{r_i} j|E_j(i)| \\
&\leq r_i \sum_{i=1}^{\infty} |E_j(i)| \\
&< r_i \cdot 8(\kappa + K)i \\
&= c(2\log i + 1)^{1/2}(\log i)^{1/2} \cdot 8(\kappa + K)i \\
&\leq 16c(\kappa + K)i \log i,
\end{aligned}$$

so (4) holds with $L = 16c(\kappa + K)$. □

3 Transience in $\mathbb{R}^3$

Fix a large integer M to be determined later. We superimpose a lattice Λ with spacing M on $\mathbb{R}^3$; this defines a natural partition of the points of $\mathcal{P}$. Let $V(\Lambda)$ be the set of cells of Λ.

Our approach is to find a set of conditions which make a cell S of Λ "good"; whether or not S is good will be determined by the intersection of $DT(\mathcal{P})$ with S. We will show that the good sites define a process which dominates a supercritical independent site percolation process. The definition of "good" will then imply the existence of a subgraph H of $DT(\mathcal{P})$, contained entirely within the good sites, for which the resistance to infinity is finite.

Let $M' = M/25$ and superimpose a second lattice Λ' with spacing M' such that each cell of Λ is partitioned by cells of Λ' into 25^3 parts. By $DT_S(\mathcal{P})$ we denote the *restriction* of the drawing $DT(\mathcal{P})$ to cell S. Note that by considering $DT(\mathcal{P})$ as a drawing, we are defining $DT_S(\mathcal{P})$ to include portions of edges and faces not fully contained within S.

We say that cell $S' \in \Lambda'$ is *empty* if it contains no points of $\mathcal{P}$, and cell $S \in \Lambda$ is empty if one of its 25^3 subcells is empty. $S \in \Lambda$ is *dangerous* if it is at L_∞-distance 1 from an empty cell, and is *bad* if $st(DT_S(\mathcal{P})) > 1.1\kappa M$, where κ is the same constant as in the statement of Lemma 3. Finally, S is *good* if it is neither empty, dangerous or bad.

Let $E(\mathcal{P}), D(\mathcal{P})$ and $B(\mathcal{P})$ denote the sets of empty, dangerous and bad sites for the process $\mathcal{P}$, respectively. Where there is no danger of confusion, we will abbreviate these sets to E, D and B. The key lemma, which we prove in Section 3.1, is the following.

Lemma 5. *The random set of good sites dominates a supercritical independent site percolation process.*

The following result of Grimmett, Kesten and Zhang from [7] thus implies that there exists an infinite connected component G of good sites such that the resistance of G to infinity is finite (where adjacent sites are viewed as being joined by a resistor of resistance 1). In the two subsequent lemmas, we show how the definition of good allows us to find a subgraph H of $DT(\mathcal{P})$ contained within G with the same property, by "simulating" each resistor in G by a bounded number of edges of H.

Theorem 6. *Let G be the (unique) infinite component of a supercritical site percolation process in $\mathbb{Z}^3$. Then the resistance of G to infinity is finite, with probability one.*

Grimmett, Kesten and Zhang proved their result for bond percolation, but the result also holds for site percolation.

We also require the following result, often referred to as the fundamental theorem of electrical networks [3], page 318. To set the scene, let M_1 and M_2 be electrical networks, each containing a set U of vertices called the vertices of *attachment*. We say that (M_1, U) is *equivalent* to (M_2, U) if whenever N is a network sharing with each M_i the set U and nothing else, and we set some vertices of N at certain potentials, then in $N \cup M_1$ and $N \cup M_2$ we obtain precisely the same currents in the edges of N aside from those in M_1,M_2. The intuition behind this definition of "equivalence" is that the network $N \cup M_1$ and $N \cup M_2$ "look the same" with respect to currents everywhere in N except possibly within M_1 or M_2. The fundamental theorem of electrical networks states that no matter how complicated the graph M_1 may be, from the perspective of the rest of N it might as well just be the set of vertices U and some resistors between these vertices. To be precise:

Lemma 7 (Fundamental Theorem of Electrical Networks)**.** *Every network with attachement set U is equivalent to a network with vertex set U.*

Say that H is a (m, n)-*blowup* of G if it can be obtained by replacing each vertex v of G by a connected graph G_v with at most n and at least d_v vertices, and then replacing each edge $e = uv$ of G by a path $P(u, v)$ of length at most l with endpoints in G_u and G_v and disjoint from all other such paths. Then we have:

Lemma 8. *Let G be a transient connected graph with maximum degree Δ and let n, m be positive integers. If H is a (m, n)-blowup of G then H is transient.*

Proof. We define the graph G_f as follows. Let $f(m, n)$ be some positive integer function of m and n, and replace each edge uv of G by a path of length $f(m, n)$. If v is a vertex of G, clearly the resistance to infinity from v to infinity in G_f is precisely $f(m, n)$ times the resistance to infinity from v in G, which is finite. Thus G_f is transient.

It remains to show that there is a function $f(m, n)$ such that for any (m, n)-blowup H of G, if G_f is transient then H is transient. We do this by a sequence of transformations to H that yield such a graph G_f and increase the resistance to infinity at every step.

First, in H we may replace each G_u by an equivalent graph R_u on d_u vertices (these are the vertices of attachment), so that all external currents and potential differences remain the same, by Lemma 7. Further, there are only a finite number of possibilities for G_u and the attachment vertices as Δ is bounded. Since in addition G_u is connected, the resistances along the edges of R_u are all bounded above by some constant value $r(n)$. By Rayleigh's monotonicity law, we may replace R_u by a spanning tree T_u of R_u, and the overall resistance to infinity will increase.

In T_u the resistance between any pair of vertices is at most $(n-1)r(n)$, as all vertices are joined by a path of length at most $(n-1)$. It seems intuitive that replacing T_u by a star S_u with hub u and leaves the vertices of T_u, each edge of S_u having resistance at least $(n-1)r(n)$, the resistance to infinity from any vertex should increase. We prove something similar to this by a sequence of transformations of T_u, each of which can not decrease the resistance to infinity.

First select a root $r \in V(T_u)$. Select a vertex v_2 at graph distance 2 from r, and let rv_1v_2 be a path of length 2 from r to v_2. Suppose the resistances along rv_1 and v_1v_2 are a and b respectively. Create a new vertex v_1', and edges rv_1', v_1v_1' and $v_1'v_2$, and set the resistances along rv_1, rv_1', v_1v_1', $v_1'v_2$ and v_1v_2 to be $2a$, $2a$, 0, $2b$ and $2b$ respectively: this network is clearly equivalent to the old one.

Cut edges v_1v_1' and v_1v_2 and replace the path $rv_1'v_2$ by a single edge with resistance $2(a+b)$. The first step increases all effective resistances in the network, and the second has no effect on them. In this way we have increased the degree of r by one.

Repeating this procedure at most n times, T_u will be replaced by a graph S_u' in which r is adjacent to all other vertices. Each time we apply the procedure at most quadruples the largest resistance so in the end all resistances are bounded above by some fixed value $g(n)$.

We create S_u from S_u' by (a) removing all edges of S_u' of which neither endpoint is r, and (b) replacing the root r by an edge ru with resistance 0, so that r is connected to some vertex not in S_u' and u is incident to the neighbours of r in S_u'. Neither operation decreases the resistance to infinity.

The final step is to replace each edge connecting neighbouring stars S_u and S_v with a path $P(u,v)$ of length m. At this point, the hubs of neighbouring stars are connected by a path consisting of two resistors of resistance at most $g(n)$ and m resistors of resistance 1. This is equivalent to a path of length $2g(n)+m$. Taking $f(m,n) = 2g(n)+m$, this last graph is equivalent to G_f. As the effective resistance to infinity rose at every step of this transformation, if G_f is transient then H is transient, as claimed. □

In the preceding proof we glossed over the technicalities of what is meant by "effective resistance to infinity" in the absence of a chosen vertex for the resistance to infinity to be *from.* However, in a connected graph in which every resistance is bounded, every node has either finite resistance to infinity or infinite resistance to infinity. As we are only interested in the finitude of this quantity and not a precise bound, we feel this imprecision is justified.

We now show that we can carry out the above replacement process in the infinite cluster of good sites of Λ. As this cluster has finite resistance to infinity by Lemma 5, this establishes:

Theorem 9. *The simple random walk on* $\mathrm{DT}(\mathcal{P})$ *is transient in dimension* $d \geq 3$*, with probability one.*

Proof of Theorem 9. All the essential ingredients in the proof occur in the case $d = 3$; we thus restrict our attention to this case. The proof can be exactly reproduced for $d > 3$ - only certain constants will change.

Given a cell $S \in \Lambda$, let S_m be the cell of Λ' at the center of S, and let S_c be the box with side length $13M/25 = 13M'$ and with the same center as S and S_m. Note that by Theorem 6, the set of good sites, viewed as a graph with edges between adjacent cells, has finite resistance to infinity. (Here and for the remainder of the proof, adjacent is used to mean that the cells share a face.) We exhibit a subgraph $G = (V, E)$ of $DT(\mathcal{P})$ contained within the set of good sites of $V(\Lambda)$ so that the following conditions hold:

(1) For any cell S, $|V \cap S_c| \leq 6.6\kappa M + 6$

(2) Given adjacent cells S, S', there is a path from $|V \cap S_c|$ to $|V \cap S'_c|$ disjoint from $|V \cap S_c|$ and $|V \cap S'_c|$ except at its endpoints and disjoint from all other such paths. Furthermore, this path contains at most $2.2\kappa M + 1$ vertices of V.

It follows from Lemma 5 and Lemma 8 that G, and thus $DT(\mathcal{P})$, has finite resistance to infinity. It thus remains to prove (1) and (2).

As S is good, S_m contains a point; call it v_S. Given adjacent cells S, S', let $\ell_{S,S'}$ be the line segment connecting $v_S, v_{S'}$. As S and S' are both good, $st(DT_S(\mathcal{P})) \leq 1.1\kappa M$ and $st(DT_{S'}(\mathcal{P})) \leq 1.1\kappa M$. Thus $\ell_{S,S'}$ crosses at most $2.2\kappa M$ cells of $DT(\mathcal{P})$. It is left as an exercise to the reader to show that this implies there exists a path $P(S, S')$ in G from v_S to $v_{S'}$ with at most $2.2\kappa M$ edges, all of whose edges border Delaunay cells that intersect $\ell_{S,S'}$. (A picture is given in Figure 2.)

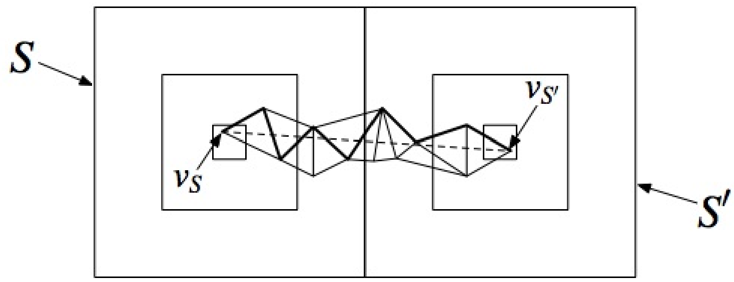

Figure 2: v_S, $v_{S'}$ and a possible path $P(S, S')$ between them (shown in bold). The smallest box containing v_S is S_m; the next smallest is S_c.

We let G consist of the points $\{v_S \mid S \text{ is good}\}$, together with all paths $P(S, S')$ between adjacent good sites. We must check that (1) and (2) hold. For any cell S, $|V \cap S| \leq 6(1.1\kappa M + 1)$ as $V \cap S$ is the union of six paths, each with at most $1.1\kappa M$ edges. It is immedate that $|V \cap S_c| < 6.6\kappa M + 6$, i.e. (1) holds.

We now turn our attention to (2). Let $\ell_{S,S'}$ and $\ell_{S,S''}$ be lines connecting v_S to $v_{S'}$ and $v_{S''}$, and let $\triangle \in DT(\mathcal{P})$ be any Delaunay cell intersecting both lines. It follows that the sphere C circumscribing $\triangle$ intersects both lines; if any intersection is outside of S_c it follows immediately that C has diameter at least $6M'$, as S_c has side length $13M'$ and $\ell_{S,S'}$ and $\ell_{S,S''}$ intersect only at v_S, which has distance at least $6M'$ from all sides.

Recall that as S is good, every cell of $\Lambda' \cap S$ contains a point. Notice that any sphere with diameter of length at least $6M'$ contains a cell of Λ' in its interior. (This can be seen by considering an axis-aligned cube inscribed in such a sphere; in fact the cube must contain a cell of Λ'.) As C has empty interior, it must have diameter less than $6M'$, so all intersections of C with $\ell_{S,S'}$ and $\ell_{S,S''}$ are within S_c. It follows that if either $P(S, S')$ or $P(S, S'')$ contains an edge of $\triangle$, this edge is fully contained within S_c. As $\triangle$ was an arbitrary Delaunay cell intersecting both $\ell_{S,S'}$ and $\ell_{S,S''}$, it follows that $P(S, S')$ and $P(S, S'')$ are disjoint except perhaps within S_c. This proves the first part of (2).

The second part of (2) is trivial as $P(S, S')$ contains at most $2.2\kappa M$ edges, as noted above.

□

3.1 Proof of Lemma 5

It is not hard to see that, for a single site S, $\mathbb{P}(S \in E \cup D \cup B)$ can be made small by taking M large (though we have not yet *proved* even this). However, this fact on its own is insufficient to prove that the good sites dominate a supercritical *independent* site percolation process. To see this, consider the following process: fix an arbitrarily small positive ϵ, and let Y be a $0-1$ Bernoulli random variable with $\mathbb{P}(Y=1) = 1-\epsilon$. Let $\mathcal{Y} = \{Y_s\}_{s \in \mathbb{Z}^d}$ be a family of $(0-1)$ random variables such that $Y_s = 1$ precisely if $Y = 1$, and let $C_{\mathcal{Y}}$ be the event that there exists an infinite open cluster in $\{s | Y_s = 1\}$. Then $\mathbb{P}(C_{\mathcal{Y}}) = 1 - \epsilon$, but for any supercritical independent site percolation process $\mathcal{Y}'$, $\mathbb{P}(C_{\mathcal{Y}'}) = 1$. Thus, $\mathcal{Y}$ does not dominate any such process.

We thus need a result of Liggett, Schonmann and Stacey from [9], which proves that if dependence of a family $\mathcal{X}$ has "limited range" and, for $X \in \mathcal{X}$, $\mathbb{P}(X=1)$ is sufficiently high, then $\mathcal{X}$ dominates a supercritical independent process. To be more precise we make the following definitions.

Let $\mathcal{X} = \{X_s\}_{s \in \mathbb{Z}^d}$ be a lattice-indexed familiy of $0-1$ random variables. We say that $\mathcal{X}$ is a *k-dependent* family if, for each pair $A, B \subset \mathbb{Z}^d$ such that for all $a \in A$ and $b \in B$, $\|a-b\|_\infty > k$, the families of random variables $\{X_a\}_{a \in A}$ and $\{X_b\}_{b \in B}$ are independent of each other. Liggett, Schonmann and Stacey proved the following theorem.

Theorem 10. *Suppose that $\mathbb{P}(X_s = 1) > p$ for all $s \in \mathbb{Z}^d$. Then if p is large enough then $\mathcal{X}$ is dominated from below by the product random field with density ρ, where ρ is a positive constant depending on d, k and p. One can make ρ arbitrarily close to 1 by taking p large enough.*

With this result in our toolkit, we may proceed to the proof of Lemma 5.

We index the sites $V(\Lambda)$ by picking an origin site arbitrarily, and exploiting the self-duality of $\mathbb{Z}^d$ to extend this indexing to the rest of the sites. We will abuse notation by referring to sites, rather than their indices, as elements of $\mathbb{Z}^d$. A few further notational points: we say that S and S' have distance k (written $d(S,S') = k$) from each other if the L_∞-distance between their indices is k. For a set $A \subset \Lambda$, we write $d(S,A) = \min\{d(S,S') | S' \in A\}$, and define the distance $d(A,B)$ between two subsets of Λ similarly. For a set A, we also define $N(A) = \bigcup_{S' | d(S',A)=1} S'$ and $\bar{N}(A) = A \cup N(A)$.

We let X_S be the indicator variable of the event $\{S \notin E \cup D \cup B\}$, i.e., of the event that S is good. We will show that $\mathcal{X} = \{X_S\}_{S \in \mathbb{Z}^d}$ is a 2-dependent family and that $\mathbb{P}(X_S = 1)$ can be made arbitrarily close to 1 by taking

M large enough. Lemma 5 then immediately follows from Theorem 10 on making $\mathbb{P}(X_S = 1)$ large enough so that ρ is greater than the critical site percolation probability in $\mathbb{Z}^3$.

To show that $\mathcal{X}$ is a 2-dependent family it suffices to prove that for all $X_S \in \mathcal{X}$, X_S is determined by the behavior of the Poisson process $\mathcal{P}$ in the region $\bar{N}(S)$. For if this holds then if $d(A, B) > 2$, any event E_1 depending only on $\{X_S | S \in A\}$ is determined by the behaviour of $\mathcal{P}$ in

$$\bigcup_{S \in A} \bar{N}(S) = \bar{N}(A).$$

Similarly, any event E_2 depending only on $\{X_S | S \in B\}$ is determined by the behaviour of $\mathcal{P}$ on $\bar{N}(B)$. E_1 and E_2 are thus independent by the properties of the Poisson process, as $\bar{N}(A)$ and $\bar{N}(B)$ are disjoint.

To see that $\{S \notin E \cup D \cup B\}$ is determined by $\mathcal{P} \cap \bar{N}(S)$, it is useful to write $E \cup D \cup B = E \cup D \cup B'$, where $B' = B - (D \cup E)$. By definition, $\{S \in E\}$ is determined by $\mathcal{P} \cap S$, and $\{S \in D\}$ is determined by $\mathcal{P} \cap N(S)$.

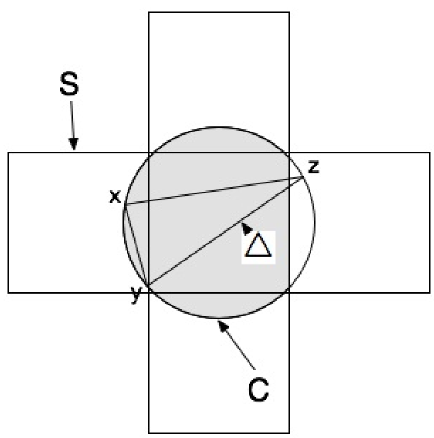

Figure 3: A geometric view of facts (a), (b) and (c) in two dimensions: note that C intersects both S and $\mathbb{R}^3 - \bar{N}(S)$. In this case, if $C \cap \bar{N}(S)$ (the shaded area) contains no points of $\mathcal{P}$ aside from x and y, then $\triangle$ is a cell of $DT(\mathcal{P})$ precisely if $z \in \mathcal{P}$ and $C \cap (\mathbb{R}^3 - \bar{N}(S))$ contains no other points of $\mathcal{P}$.

We now turn our attention to the event $\{S \in B'\}$. We may assume $S \notin D \cup E$, as this is determined by $\mathcal{P} \cap \bar{N}(S)$. Suppose $S \notin D \cup E$, but that the event $\{S \in B'\}$ is *not* determined by $\mathcal{P} \cap \bar{N}(S)$. As $\{S \in B\}$ is the event that $st(DT_S(\mathcal{P})) > 1.1\kappa M$, it is certainly determined by $DT_S(\mathcal{P})$. If $DT_S(\mathcal{P})$ is not determined by $\mathcal{P} \cap \bar{N}(S)$, then there is some tetrahedron $\triangle$ with corners $\{x, y, z, w\}$ that intersects S, such that whether or not $\triangle$ is a cell of $DT(\mathcal{P})$ depends on the behaviour of $\mathcal{P}$ in $\mathbb{R}^3 - \bar{N}(S)$. Considering the sphere C that circumscribes $\triangle$, C must satisfy:

(a) $C \cap S$ is not empty.

(b) $C \cap (\mathbb{R}^3 - \bar{N}(S))$ is not empty.

(c) $C \cap \bar{N}(S) \cap \mathcal{P}$ contains no points aside from one or all of x, y, z, w.

The fact that (a),(b) and (c) hold follows directly from the definition of the Delaunay triangulation; Figure 3 gives a geometric view of what they state. Conditions (a) and (b) imply that C contains one of the subcells T of $N(S)$ of side length $M' = M/25$ in its interior. Condition (c) then implies that $T \cap \mathcal{P} = \emptyset$, so that one of $S \cup N_\infty(S)$ is in E, and so by definition $S \in E \cup D$, a contradiction. Thus $\{S \in E \cup D \cup B\}$ is determined by $\mathcal{P} \cap \bar{N}(S)$, so $\mathcal{X}$ is a 2-dependent family.

It remains to show that $\mathbb{P}(X_S = 1)$ can be made arbitrarily high. We have

$$\mathbb{P}(X_S = 1) \geq 1 - \mathbb{P}(S \in E) + \mathbb{P}(S \in D) + \mathbb{P}(S \in B),$$

and

$$\begin{aligned} \mathbb{P}(S \in E) + \mathbb{P}(S \in D) &\leq \mathbb{P}(S \in E) + \mathbb{P}(\exists S' \in N(S), S' \in E) \\ &\leq 27\mathbb{P}(S \in E) \\ &\leq 27 \cdot 25^3 \cdot \mathbb{P}\left([0, M/25]^3 \cap \mathcal{P} = \emptyset\right) \\ &\leq 27 \cdot 25^3 \cdot e^{-(M/25)^3}, \end{aligned}$$

which can be made arbitrarily small by choosing M large enough. Turning our attention to $\mathbb{P}(S \in B)$, by Lemma 3 we have

$$\begin{aligned} \mathbb{P}(S \in B) &= \mathbb{P}(st(DT_S(\mathcal{P})) > (1 + 0.1)\kappa M) \\ &\leq e^{-0.1M/K \log M}, \end{aligned}$$

which can also be made arbitrarily small by taking M large enough as κ and K do not depend on M. Thus $\mathbb{P}(X_S = 1)$ can be made arbitrarily large, so Lemma 5 holds.

4 Conclusion

The authors are aware that the proofs presented here can be adapted to other random graphs. In particular, for graphs satisfying "reasonable" conditions on minimum and maximum degree and edge length, it should be possible to reprove our results. On the other hand, the results will not necessarily be optimal in other cases; by way of explanation we briefly discuss the Gilbert disc and k-nearest neighbour models.

Gilbert's disc model G_r connects points x, y of the Poisson process satisfying $\|x - y\|_2 \leq r$. Adaptation of the above proofs should easily yield the recurrence results for the G_r in $\mathbb{Z}^2$, for any r. In $\mathbb{Z}^d$, $d \geq 3$, the proof technique should again apply; however, it will only prove that the simple random walk on G_r is transient *for all r sufficiently large.* It would be nice to show that for *any* r for which G_r contains an infinite connected component, the simple random walk on that component is transient in $\mathbb{Z}^d$, $d \geq 3$. Similarly, the results should extend to the k-nearest neighbour model for all k in the $d = 2$ case, and for k sufficiently large in the $d \geq 3$ case. In both cases it will be necessary to deal with some connectivity complications not present in the Delaunay triangulation. We note that in both cases, a proof would yield an upper bound on the percolation threshold, in terms of r, for the Gilbert disc model (though a better bound follows directly from comparison with bond percolation in this case) and in terms of k for the k-nearest neighbour model.

We also believe it would be interesting to study the mixing or cover time of the walk on a finite portion of such a graph. In fact, such quantities have been studied for finite graphs using the connection with electrical networks; see, e.g., [5].

In [10], Thomassen proves that given an infinite graph G, if $f(k)$ denotes the smallest number of vertices in the boundary of a connected subgraph with k vertices, then the simple random walk on G is transient if $\sum f(k)^{-2}$ converges. This result does not apply to such random graphs as those discussed above, due to the presence of locally "bad" areas of the graph where such a condition will not hold. However, it may be possible to show that if G is a graph for which such a condition holds "almost everywhere with probability 1" then the simple random walk on G is almost surely transient.

References

[1] L. Addario-Berry, N. Broutin and L. Devroye, The slicing number of a Delaunay triangulation, *submitted*.

[2] P.N. Balister, B. Bollobás and M.J. Walters, Continuum percolation with steps in the square or the disc, *Random Structures and Algorithms* **26** (2005), 392-403.

[3] B. Bollobás, *Modern Graph Theory*, Springer, New York (1998).

[4] B. Bollobás and O.M. Riordan, The critical probability for random Voronoi percolation in the plane is 1/2, *Probability Theory and Related Fields*, to appear.

[5] A.K. Chandra, P. Raghavan, W.L. Ruzzo, R. Smolensky, P. Tiwari, The electrical resistance of a graph captures its commute and cover times, *Proc. 21'st Symposium on the Theory of Computing* **21** (1989) 574-586.

[6] P.G. Doyle and L.J. Snell, *Random Walks and Electrical Networks*, Carus Math. Monographs 22, Washington DC (1984).

[7] G.R. Grimmett, H. Kesten and Y. Zhang, Random walk on the infinite cluster of the percolation model, *Probability Theory and Related Fields* **96** (1993), 33-44.

[8] G.R. Grimmett, *Percolation*, second edition, Springer, New York, (1999).

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

[10] C. Thomassen, Isoperimetric inequalities and transient random walks on graphs, *Annals of Probability* **20** (1992), 1592-1600.