Delaunay graphs of point sets in the plane with respect to axis-parallel rectangles

Xiaomin Chen[*], *János Pach*[†], *Mario Szegedy*[‡] and *Gábor Tardos*[§]

Abstract

Given a point set P in the plane, the *Delaunay graph* with respect to axis-parallel rectangles is a graph defined on the vertex set P, whose two points $p, q \in P$ are connected by an edge if and only if there is a rectangle parallel to the coordinate axes that contains p and q, but no other elements of P. The following question of Even et al. [ELRS03] was motivated by a frequency assignment problem in cellular telephone networks. Does there exist a constant $c > 0$ such that the Delaunay graph of any set of n points in general position in the plane contains an independent set of size at least cn? We answer this question in the negative, by proving that the largest independent set in a randomly and uniformly selected point set in the unit square is $O(n \log^2 \log n / \log n)$, with probability tending to 1. We also show that our bound is not far from optimal, as the Delaunay graph of a uniform random set of n points almost surely has an independent set of size at least $cn/\log n$.

We give two further applications of our methods. 1. We construct 2-dimensional n-element partially ordered sets such that the size of the largest independent sets of vertices in their Hasse diagrams is $o(n)$. This answers a question of Matoušek and Přívětivý [MaP06] and improves a result of Kříž and Nešetřil [KrN91]. 2. For any positive integers c and d, we prove the existence of a planar point set with the property that no matter how we color its elements by c colors, we find an axis-parallel rectangle containing at least d points, all of which have the same color. This solves an old problem from [BrMP05].

1 Delaunay graphs and conflict-free colorings

The *Delaunay graph* associated with a set of points P in the plane is a graph $D'(P)$ whose vertex set is P and whose edge set consists of those pairs $\{p, q\} \subset P$ for which there exists a closed disk that contains p and q, but does not contain any other element of P. The Delaunay graph of P is a planar graph and its dual is the *Dirichlet–Voronoi diagram* of P (see, e.g., [BKOS00]). As any other planar graph, $D'(P)$ contains an independent set of size at least $|P|/4$. It was discovered by Even, Lotker, Ron, and Smorodinsky [ELRS03] that this fact easily implies that any set P of n points in the plane has a *conflict-free coloring* with respect to disks, which uses at most $O(\log n)$ colors, that is, a coloring with the property that any closed disk C with $C \cap P \neq \emptyset$ has an element whose color is not assigned to any other element of $C \cap P$. Here, the logarithmic bound is tight for every point set [PaT03].

[*]Google, 76 Ninth Avenue, New York, NY, 10011.

[†]City College, CUNY and Courant Institute, NYU, 251 Mercer Street, New York, NY 10012. Supported by NSF grant CCF-05-14079 and grants from NSA, PSC-CUNY, Hungarian Research Foundation, and BSF.

[‡]Rutgers University, 110 Frelinghuysen Road Piscataway, NJ 08854-8019. Supported by NSF grant 0105692.

[§]Simon Fraser University, 8888 University Drive, Burnaby, B.C., Canada V5A 1S6 and Rényi Institute, H-1055 Budapest, Reáltanoda utca 13-15, Hungary. Supported by NSERC grant 329527.

The question was motivated by a frequency assignment problem in cellular telephone networks. The points correspond to *base stations* interconnected by a fixed backbone network. Each *client* continuously scans frequencies in search of a base station within its (circular) range with good reception. Once such a base station is found, the client establishes a radio link with it, using a frequency not shared by any other station within its range. Therefore, a conflict-free coloring of the points corresponds to an assignment of frequencies to the base stations, which enables every client to connect to a base station without interfering with the others. For many results on conflict-free colorings, consult [AlS06], [FiLM05], [HaS05].

The same scheme can be used to construct conflict-free colorings of point sets with respect to various other families of geometric figures. In general, let P be a set of points in $\mathbf{R}^d$, and let $\mathcal{C}$ be a family of d-dimensional convex bodies. Define the *Delaunay graph* $D_{\mathcal{C}}(P)$ *of* P *with respect to* $\mathcal{C}$ on the vertex set P by connecting two elements $p, q \in P$ with an edge if and only if there is a member of $\mathcal{C}$ that contains p and q, but no other element of P. The existence of large independent sets in such graphs implies that P has a *conflict-free coloring with respect to* $\mathcal{C}$, which uses a small number of colors. That is, a coloring with the property that any member $C \in \mathcal{C}$ with $C \cap P \neq \emptyset$ has an element whose color is not assigned to any other element of $C \cap P$.

In this paper, we consider this problem in the special case when $\mathcal{C}$ is the family of *axis-parallel rectangles* in the plane. The Delaunay graph $D(P)$ of a point set P with respect to axis-parallel rectangles is also called the *rectangular visibility graph* of P. Computing all rectangularly visible pairs of an n-element point set, that is, all edges of $D(P)$, is a classical problem, solved in $O(n \log n + |D(P)|)$ time by Güting, Nurmi, and Ottman [GuNO85], in a paper presented at the *First ACM Symposium on Computational Geometry* in 1985. See also [GuNO89], [OvW88], [DeH91], [JaT92].

The maximum size of an independent set of vertices in a graph G is called the *independence number* of G, and is usually denoted by $\alpha(G)$ in the literature. Smorodinsky et al. [ELRS03], [HaS05] asked whether the Delaunay graph of every set of n points in the plane with respect to axis-parallel rectangles has independence number at least cn, for an absolute constant $c > 0$. In Section 3, we give a negative answer to this question. More precisely, we establish

Theorem 1. *There are n-element point sets in the plane such that the independence numbers of their Delaunay graphs with respect to axis-parallel rectangles are at most* $O\left(n\frac{\log^2 \log n}{\log n}\right)$.

In fact, a randomly and uniformly selected set of n points in the unit square will meet the requirements with probability tending to 1.

For randomly selected point sets, this result is not far from being best possible. In Section 2, we prove

Theorem 2. *The independence number of a randomly and uniformly selected n-element point set in the unit square is almost surely* $\Omega\left(\frac{n}{\log n}\right)$.

For arbitrary point sets, Ajwani, Elbassioni, Govindarajan, and Ray [AjEG07] proved that the independence number of the Delaunay graph of any set of n points in the plane with respect to axis-parallel rectangles is at least $\Omega\left(n^{.617}\right)$. This implies that any set of n points in the plane admits a conflict-free coloring using $O\left(n^{.383}\right)$, with respect to the family of all axis-parallel rectangles. For weaker results, consult [PaT03],[MaP06], [ElM06].

Matoušek and Přívětivý raised another closely related problem. Given a finite partially ordered set $(X, <)$, we say that $p \in X$ is an *immediate predecessor* of $q \in X$ if $p < q$ and there is no $r \in X$ with $p < r < q$. The *Hasse diagram* $H(X, <)$ of $(X, <)$ is an undirected graph on the vertex set X, in which two vertices are connected if and only if one is an immediate predecessor of the other. The (Dushnik-Miller) *dimension* of a partial ordering $<$ is the smallest number of linear (that is, total) orderings whose intersection is $<$. Matoušek and Přívětivý [MaP06] asked whether the Hasse diagram of every two-dimensional partially ordered set of n elements contains an independent set whose size is linear in n. The next theorem provides a negative answer to this question.

Theorem 3. *There are two-dimensional partially ordered sets with n elements such that the independence numbers of their Hasse diagrams are at most $O\left(n\frac{\log^2 \log n}{\log n}\right)$.*

Given a finite point set P in the plane and a fixed (x, y) coordinate system, we can define a partial ordering on P by letting $p \leq q$ if the x-coordinate of p does not exceed the x-coordinate of q and the y-coordinate of p does not exceed the y-coordinate of q. This ordering is called the *domination order* of P with respect to the coordinate system. Reversing the direction of the x-axis (that is, replacing x by $-x$), we obtain another domination order of P. Denoting the Hasse diagrams of these two domination orders by $H(P)$ and $H'(P)$, we have that their union $H(P) \cup H'(P)$ is equal to $D(P)$, the Delaunay graph $D(P)$ with respect to axis-parallel rectangles. Therefore, the independence number of $H(P)$ satisfies $\alpha(H(P)) \geq \alpha(D(P))$. Theorem 3 can be established by a slight modification of the proof of Theorem 1; the details are left to the reader. The argument also gives that if $<$ is the intersection of two randomly and uniformly selected linear orderings of an n-element set X, then $\alpha(H(X, <)) = O\left(n\frac{\log^2 \log n}{\log n}\right)$, with probability tending to 1.

Kříž and Nešetřil [KrN91] gave an explicit construction proving that the chromatic numbers of the Hasse diagrams of planar point sets are not bounded. A little calculation based on their construction shows that there exist n-element point sets P such that the chromatic numbers of their Hasse diagrams grow as fast as $\log^* n$, the iterated logarithm of n. Theorem 3 implies a better bound.

Corollary 4. *There are two-dimensional partially ordered sets with n elements such that the chromatic numbers of their Hasse diagrams are at least $\Omega\left(\frac{\log n}{\log^2 \log n}\right)$.*

Note that the construction of Kříž and Nešetřil contains linear sized independent sets.

In geometric discrepancy theory [BeCh87], [Ch00], [Ma99], there are plenty of results that indicate some unavoidable irregularities in geometric configurations. In Section 4, we generalize Theorem 1. As a corollary of our results, we obtain a solution to Problem 5, Chapter 2.1 in [BrMP05].

Theorem 5. *For any positive integers c and d, there is a finite point set in the plane with the property that no matter how we color its elements with c colors, there always exists an axis-parallel rectangle containing at least d points, all of which have the same color.*

Pach and Tardos proved the "dual" statement: For any positive integers c and d, there is a d-fold covering of the unit square $[0, 1]^2$ with finitely many axis-parallel rectangles such that no matter how we color them with c colors, there always exists a point in $[0, 1]^2$ with the property that all rectangles containing it are of the same color. (A collection of sets is said to form a *d-fold covering* of the unit square if every point of $[0, 1]^2$ is contained in at least d of its members.) It was shown in [Pa86] that no such d-fold covering exists with *translates* of a fixed rectangle if $c > 1$ and d is large enough. See also [PaTT07].

2 The size of Delaunay graphs of random point sets: The proof of Theorem 2

The aim of this section is to prove Theorem 2. First, we estimate the average number of edges of a Delaunay graph, and then the standard deviation of the number of edges from its expected value.

Let $P = \{(x_i, y_i) : 1 \le i \le n\}$ be a point set in the unit square, whose no two elements share the same x-coordinate or y-coordinate. Clearly, the Delaunay graph $D(P)$ with respect to axis-parallel rectangles depends only on the relative position of the points in P and not on their actual coordinates. That is, there exists a permutation $\pi : \{1, 2, \dots, n\} \to \{1, 2, \dots, n\}$ such that for the set $P' = \{(i, \pi(i)) : 1 \le i \le n\}$ we have $D(P) \approx D(P')$. Moreover, for a random set of points in the square, the corresponding permutation π is uniformly random. With a slight abuse of notation, we write $D(\pi)$ for the Delaunay graph $D(P')$. In our arguments about Delaunay graphs of randomly selected point sets in the square, it will be convenient to consider the graph $D(\pi)$ for a *random permutation* π. The number of edges of $D(\pi)$ will be denoted by $|D(\pi)|$.

Lemma 6. *Let $\pi : \{1, 2, \dots, n\} \to \{1, 2, \dots, n\}$ be a random permutation. The expected value of the number of edges of the Delaunay graph $D(\pi)$ satisfies*

$$\mathsf{E}[|D(\pi)|] = 2n \log n - O(n).$$

Proof. Two points $p_i = (i, \pi(i))$ and $p_j = (j, \pi(j))$ with $i < j$ are connected by an edge in $D(P)$ if and only if $\pi(i)$ and $\pi(j)$ are consecutive elements in the natural ordering of the set $S = \{\pi(k) | i \le k \le j\}$. Among all $\binom{j-i+1}{2}$ pairs of elements in this set, precisely $j - i$ consist of consecutive elements. Clearly, after fixing $\pi(k)$ for $k < i$ or $k > j$, the pair $\{\pi(i), \pi(j)\}$ is equally likely to be any one of the pairs in S. Therefore, the probability that p_i and p_j are connected is equal to

$$\frac{j-i}{\binom{j-i+1}{2}} = \frac{2}{j-i+1}.$$

Thus, the expected number of edges in $D(P)$ is

$$\sum_{l=1}^{n-1} \frac{2(n-l)}{l+1} = (2n+2) \sum_{l=1}^{n} \frac{1}{l} - 4n = 2n \log n - O(n).$$

$\square$

Lemma 6 easily implies that

$$\mathsf{E}[\alpha(D(\pi))] = \Omega\left(\frac{n}{\log n}\right).$$

To complete the proof of Theorem 2, it is sufficient to show that the number of edges of $D(\pi)$, for a random permutation π, is concentrated around its expected value.

Lemma 7. *Let $\pi : \{1, 2, \dots, n\} \to \{1, 2, \dots, n\}$ be a random permutation, and let $\sigma(|D(\pi)|) = \sqrt{\mathsf{var}|D(\pi)|}$ denote the standard deviation of the number of edges of the Delaunay graph $D(\pi)$. We have*

$$\sigma(|D(\pi)|) = O(\sqrt{n} \log n).$$

Proof. Let p_i denote the same as in the proof of Lemma 6, and let ξ_{ij} be the characteristic function of the event that p_i and p_j form an edge in $D(\pi)$. Clearly, we have $|D(\pi)| = \sum_{i,j} \xi_{ij}$.

We have to estimate the variance

$$\mathsf{var}\left[\sum_{i,j} \xi_{ij}\right] = \sum_{i,j} \mathsf{var}[\xi_{ij}] + 2 \sum_{\{i,j\} \ne \{i',j'\}} \mathsf{cov}[\xi_{ij}, \xi_{i'j'}]. \tag{1}$$

In the proof of Lemma 6, we have shown that $E(\xi_{ij}) = \frac{2}{j-i+1}$, for any $j > i$. Using the fact that ξ_{ij} is a 0-1 valued function, we obtain that

$$\mathsf{var}[\xi_{ij}] = \mathsf{E}[\xi_{ij}](1 - \mathsf{E}[\xi_{ij}]) < \frac{2}{j-i+1}.$$

Therefore, we have $\sum_{i,j} \mathsf{var}[\xi_{ij}] = O(n \log n)$, and in (1) it remains to estimate the pairwise covariances of the random variables ξ_{ij}.

Let ξ_{ij} and $\xi_{i'j'}$ be two characteristic functions, as above. We distinguish two cases and several subcases.

Case 1: The indices i, j, i', j' are all distinct.

Subcase 1a: The intervals $[i, j]$ and $[i'j']$ are disjoint. In this case, obviously, the random variables ξ_{ij} and $\xi_{i'j'}$ are independent, so that we have $\mathsf{cov}[\xi_{ij}, \xi_{i'j'}] = 0$.

Subcase 1b: One of the intervals $[i, j]$ and $[i'j']$ contains the other. In this case, we can still argue that ξ_{ij} and $\xi_{i'j'}$ are independent. Indeed, assume without loss of generality assume that $[i, j]$ contains $[i'j']$. Generate the permutation π by first fixing its values outside the interval $[i', j']$, and then, at the second stage, by fixing the values of the elements in $[i', j']$. Observe that after the first stage we know whether $p_i p_j$ is an edge of $D(\pi)$. It is decided at the second stage whether $p_{i'} p_{j'}$ is an edge, but the probability of this event is exactly $\frac{2}{j'-i'+1}$, independently of the outcome at the first stage. Thus, again we have $\mathsf{cov}[\xi_{ij}, \xi_{i'j'}] = 0$.

Subcase 1c: The intervals $[i, j]$ and $[i'j']$ are intertwined. We may assume without loss of generality that $i < i' < j < j'$. Generate π by the following process. First we fix the values of π outside of the interval $[i, j']$. Next we determine the values of π for i and j'. In the third step, we temporarily fix π for the remaining elements in the open interval (i, j'). Finally, in the fourth step we swap the image, $\pi(x)$, of a random element $x \in [i', j') \setminus \{j\}$ with $\pi(i')$. Clearly, this way we obtain a random permutation. We need the fourth step for technical reasons. The probability that after the third step the rectangle induced by p_i and p_j is either empty or contains the point $p_{i'}$ (but no other point) is exactly $\frac{2}{j-i}$. Let us denote this event by W. If after the last step $p_i p_j$ is an edge, then W holds. Compute the probability, conditioned on W, that $p_{i'} p_{j'}$ is an edge after the fourth step. Let $x_m = \min_{k \in [i',j'] \setminus \{j\}} \pi(k)$ and $x_M = \max_{k \in [i',j'] \setminus \{j\}} \pi(k)$. If $\pi(j') \notin \{x_m, x_M\}$, then before the final (fourth) step exactly two elements of $[i', j') \setminus \{j\}$ have the property that swapping their π values with $\pi(i')$, the rectangle induced by $p_{i'}$ and $p_{j'}$ becomes empty or it only contains p_j. (We think of p_j as invisible.) If $\pi(j') \in \{x_m, x_M\}$, there is exactly one such element. Hence, the probability that $p_{i'} p_{j'}$ becomes an edge after the fourth step is at most $\frac{2}{j'-i'}$, regardless of how π is fixed on $[i, j] \setminus \{i'\}$. Therefore, we have

$$\begin{aligned} \mathsf{cov}[\xi, \xi'] = \mathsf{E}[\xi\xi'] - \mathsf{E}[\xi]\mathsf{E}[\xi'] &\leq \frac{2}{j-i} \frac{2}{j'-i'} - \frac{2}{j-i+1} \frac{2}{j'-i'+1} \\ &< \frac{8}{(j-i)(j'-i') \min\{(j-i), (j'-i')\}}. \end{aligned}$$

Remark. It is easy to see that if $j - i = j' - i' = 2$, and the intervals $[i, j]$ and $[i', j']$ are intertwined, then the covariance is $5/12 - 4/9 = -1/36$. If $i = 1, i' = 2, j = 3, j' = 5$, the covariance is $38/120 - 1/3 = -1/60$.

Case 2: The indices i, j, i', j' are not all distinct.

Subcase 2a: $i = i' < j' < j$. We obtain by direct computation that

$$\begin{aligned}
\mathsf{cov}[\xi, \xi'] &= \mathsf{E}[\xi\xi'] - \mathsf{E}[\xi]\mathsf{E}[\xi'] \\
&= \frac{4(j-i-1)}{(j-i+1)(j-i)(j'-i')} - \frac{4}{(j-i+1)(j'-i'+1)} \\
&= O\left(\frac{1}{(j-i)(j'-i')^2}\right).
\end{aligned}$$

Subcase 2b: $i' < j' = i < j$. An argument similar to the one applied in Subcase 1c yields that if $j - i \geq j' - i'$, then $\mathsf{cov}(\xi, \xi') = O\left(\frac{1}{(j-i)(j'-i')^2}\right)$, and if $j - i \leq j' - i'$, then $\mathsf{cov}(\xi, \xi') = O\left(\frac{1}{(j'-i')(j-i)^2}\right)$.

Summarizing: the last term of (1), and therefore the variance of $|D(\pi)| = \sum_{i,j} \xi_{ij}$, can be estimated by

$$\sum_{\{i,j\} \neq \{i',j'\}} \mathsf{cov}[\xi_{i,j}, \xi_{i',j'}] = O(n \log^2 n),$$

completing the proof of Lemma 7. $\square$

Proof of Theorem 2. By Lemma 6, the expected number of edges in the Delaunay graph of a random permutation π on n elements satisfies

$$\mathsf{E}[|D(\pi)|] = \Theta\left(n \log n\right).$$

By Chebyshev's Inequality, as long as

$$\sigma(|D(\pi)|) = \sqrt{\mathsf{var}|D(\pi)|} = o(\mathsf{E}[|D(\pi)|])$$

holds, the number of edges is almost surely within a factor of $1 + \varepsilon$ of its expectation, for any $\varepsilon > 0$. Lemma 7 shows that this is the case, therefore almost surely we have $|D(\pi)| = \Theta(n \log n)$.

According to Turán's theorem, any graph with n vertices, e edges, and average degree $d = \frac{2e}{n}$ has an independent set of size at least $\frac{n}{d+1} = \frac{n^2}{2e+n}$. In particular, we have

$$\alpha(D(\pi)) \geq \frac{n^2}{2|D(\pi)| + 1},$$

so that almost surely $\alpha(D(\pi)) = \Omega(n/\log n)$, as required. $\square$

3 The independence number of Delaunay graphs of random point sets: The proof of Theorem 1

We reformulate and prove Theorem 1 in a more precise form.

Theorem 8. *Let P be a set of n randomly and uniformly selected points in the square $[0,1]^2$. Then there exists a constant c such that*

$$\mathrm{Prob}_{n\to\infty}\left(\alpha(D(P)) < c\frac{n \log^2 \log n}{\log n}\right) \to 1.$$

Proof. The points $p_i \in P$ will be defined in two steps. First we select the x-coordinates from the interval $[0,1]$ uniformly at random. With probability 1, all the x coordinates are distinct. Let us relabel the points so that

$$0 \le x_1 < x_2 < \cdots < x_n \le 1.$$

In the second step, we select the y-coordinates of $p_i = (x_i, y_i)$ uniformly and independently from $[0,1]$. Note that, after the x_i's have been fixed, the edge set of the Delaunay graph $D(P)$ depends only on the relative order of the y_i's.

The coordinates y_i are generated as follows. Fix an integer $L \ge 2$ to be specified later. We write the numbers $y_i \in [0,1]$ in base L:

$$y_i = (0.d_i^{(1)} d_i^{(2)} \ldots)_L.$$

The digits $d_i^{(t)}$ of y_i are chosen independently and uniformly from the set $\{0, \ldots, L-1\}$. For $t \ge 1$, denote by $y_i^{(t)}$ the truncated L-ary fraction of y_i, consisting of $t-1$ digits after 0:

$$y_i^{(t)} = (0.d_i^{(1)} \ldots d_i^{(t-1)})_L.$$

The digits of y_i will be chosen one by one. At *stage* t, we determine $d_i^{(t)}$ (and, hence, $y_i^{(t+1)}$), for all i. Note that *before* stage t, the truncated fractions $y_i^{(t)}$ have already been fixed. As soon as we complete stage t, we know the y-coordinates of the points p_i up to an error of at most L^{-t}. If $y_i^{(t+1)} = y_j^{(t+1)}$, then the relative order of y_i and y_j has not yet been decided. Otherwise, if we have $y_i^{(t+1)} < y_j^{(t+1)}$, say, then $y_i < y_j$ holds in the final configuration.

Let $1 \le i < j \le n$ be fixed. Suppose that for some t, the following two conditions are satisfied:

1. $y_i^{(t+1)} = y_j^{(t+1)}$,
2. $y_k^{(t+1)} \ne y_i^{(t+1)}$ holds for all k satisfying $i < k < j$.

Then the rectangle $[x_i, x_j] \times [y_i^{(t+1)}, y_i^{(t+1)} + L^{-t})$ contains p_i and p_j, but no other element of P. Thus, in this case, p_i and p_j are connected in $D(P)$, and we say that this edge is *forced at stage* t. Although $D(P)$ may contain many edges that are not forced at any stage, we are going to use only forced edges in proving our upper bound on the independence number of $D(P)$.

Let us fix a subset $I \subset \{1, \ldots, n\}$, and let $Q = Q(I) = \{p_i : i \in I\}$. We want to estimate from above the probability that Q is an *independent* set in $D(P)$.

Let $t \ge 1$, and consider stage t of our selection process. Before this stage, $y_i^{(t)}$ has been fixed for every i. For any L-ary fraction y of the form $y = (0.d^{(1)} d^{(2)} \cdots d^{(t-1)})_L$, define a subset $H_y \subseteq \{1, \ldots, n\}$ by

$$H_y = \{1 \le i \le n : y_i^{(t)} = y\}.$$

Obviously, these sets partition $\{1, \ldots, n\}$, and hence I, into at most L^{t-1} nonempty parts. If two indices $i, j \in I$ are consecutive elements of the same part $H_y \cap I$, then we call them *neighbors*. That is, $i < j$ are neighbors if

1. $y_i^{(t)} = y_j^{(t)} = y$ holds for some y, and
2. $H_y \cap \{k \in I : i < k < j\} = \emptyset$.

For any two neighbors $i, j \in H_y$ $(i < j)$, define

$$S_{i,j} = \{k \in H_y : i < k < j\}.$$

Two neighbors $i, j \in I$ $(i < j)$ are called *close neighbors* if $|S_{i,j}| \leq L$.

If there are two close neighbors $i, j \in I$ such that the $\{p_i, p_j\}$ is an edge of $D(P)$ forced at stage t, then Q is not an independent set in $D(P)$ and we say that Q *fails at stage* t. Otherwise, Q is said to *survive* stage t, and we indicate this fact by writing $Q \curvearrowright t$.

Let $i < j$ be a pair of close neighbors. Note that $\{p_i, p_j\}$ is an edge of $D(P)$ forced in stage t if and only if $d_i^{(t)} = d_j^{(t)}$, but $d_i^{(t)} \neq d_k^{(t)}$ holds for all $k \in S_{i,j}$. The probability of this event is

$$\mathsf{Prob}(\{p_i, p_j\} \text{ is forced at stage } t) = \frac{1}{L}\left(1 - \frac{1}{L}\right)^{|S_{i,j}|}.$$

Taking into account that $|S_{i,j}| \leq L$, we obtain

$$\mathsf{Prob}(\{p_i, p_j\} \text{ is forced at stage } t) \geq \frac{1}{4L}.$$

Notice that, assuming a fix outcome of previous stages (i.e., $p_k^{(t)}$ is fixed for all k), the presence of edges $\{p_i, p_j\}$ forced at stage t are independent for all neighbors. Thus,

$$\mathsf{Prob}(Q \curvearrowright t | \text{outcome of stages } t' < t) \leq \left(1 - \frac{1}{4L}\right)^m \leq e^{-\frac{m}{4L}},$$

where m stands for the number of pairs $i, j \in I$ that are close neighbors before stage t.

Obviously, every $i \in I$, except the last element in each set H_y, has exactly *one* neighbor $j > i$. As the sets $S_{i,j}$ are pairwise disjoint for different pairs of neighbors $i < j$, there are fewer than $\frac{n}{L}$ pairs that are neighbors but not close neighbors. Thus, we have

$$m > |I| - \frac{n}{L} - L^{t-1}.$$

If $t \leq \log n / \log L$ and $|I| \geq 3n/L$, we have $m \geq n/L$, and thus

$$\mathsf{Prob}(Q \curvearrowright t | \text{outcome of stages } t' < t) \leq e^{-\frac{n}{4L^2}}.$$

As the above bound applies assuming any set of choices made at previous stages, so in particular, it applies to the conditional probability that Q survives stage t, given that it has survived all previous stages:

$$\mathsf{Prob}(Q \curvearrowright t | Q \curvearrowright t' \text{ for all } t' < t) \leq \left(1 - \frac{1}{4L}\right)^m \leq e^{-\frac{n}{4L^2}}.$$

Taking the product of these estimates for all $t \leq \log n / \log L$, we obtain

$$\mathsf{Prob}(Q \text{ survives the first } \lfloor \log n / \log L \rfloor \text{ stages}) \leq \exp\left(-\frac{n}{4L^2}\left(\frac{\log n}{\log L} - 1\right)\right).$$

The last bound is valid for any set $Q = Q(I) \subseteq P$, where $I \subset \{1, \ldots, n\}$ satisfies $|I| \geq 3n/L$. Letting

$$L = \left\lfloor \frac{\log n}{100 \log^2 \log n} \right\rfloor \quad \text{and} \quad a = \left\lceil \frac{3n}{L} \right\rceil,$$

we can conclude that

$$\begin{aligned}\mathsf{Prob}\left(\alpha(D(P)) \geq a\right) &\leq \sum_{Q \subset P, |Q|=a} \mathsf{Prob}\left(Q \text{ survives all stages}\right)\\ &\leq \binom{n}{a} \exp\left(-\frac{n}{4L^2}\left(\frac{\log n}{\log L} - 1\right)\right)\\ &\to 0,\end{aligned}$$

as required. □

4 Discrepancy in colored random point sets

In this section, we strengthen Theorem 1.

Definition. Given an integer $d > 1$ and a finite point set P in the plane, a subset $Q \subseteq P$ is called d-*independent* if there is no axis-parallel rectangle R such that $|R \cap P| = d$ and $R \cap P \subseteq Q$. Let $\alpha_d(P)$ denote the size of the largest d-independent subset of P.

According to this definition, a subset of P is 2-independent if and only if it is an independent set in the Delaunay graph $D(P)$ associated with P. In particular, we have $\alpha_2(P) = \alpha(D(P))$.

Obviously, if a set is d-independent for some $d > 1$, then it is also d'-independent for any $d' > d$. Therefore, $\alpha_d(P)$ is increasing in d.

Theorem 5 is a direct corollary to

Theorem 9. *A randomly and uniformly selected set P of n points in the unit square almost surely satisfies*

$$\alpha_d(P) = O\left(\frac{dn \log^2 \log n}{\log^{1/(d-1)} n}\right).$$

Proof. We modify the proof of Theorem 1. Pick the random points $p_i = (x_i, y_i) \in P$ according to the same multi-stage model as in the previous section, and define the truncated fractions $y_i^{(t)}$ that approximate y_i in exactly the same way as before.

Fix a subset $I \subseteq \{1, \ldots, n\}$, and let $Q = Q(I) = \{p_i : i \in I\}$. Just like in the proof of Theorem 1, analyze a fixed stage t of the selection process, by introducing the sets H_i.

Instead of using the notion of *neighbors*, we need a new definition. For any two elements $i, j \in I$ $(i < j)$ such that $y_i^{(t)} = y_j^{(t)} = y$ for some y, introduce the sets

$$T_{i,j} = \{k \in H_y \cap I : i \leq k \leq j\} \quad \text{and} \quad S_{i,j} = \{k \in H_y \setminus I : i < k < j\}.$$

The numbers i and j are called d-*neighbors* if $|T_{i,j}| = d$. The pair $\{i, j\}$ of d-neighbors is called a pair of *close* d-*neighbors* if $|S_{i,j}| \leq L$.

We say that the pair of close d-neighbors $\{p_i, p_j\}$ *fails* at stage t if at this stage the y-coordinates of *all* points p_k with $k \in T_{i,j}$ receive the same new digit $d_k^{(t)} = \delta$, but the y-coordinate of no point p_ℓ with $\ell \in S_{i,j}$ receives this digit. The probability of this event is exactly

$$L^{1-d}\left(1 - \frac{1}{L}\right)^{|S_J|} \geq L^{1-d}\left(1 - \frac{1}{L}\right)^{L} \geq \frac{1}{4L^{d-1}}.$$

Obviously, if any pair $\{p_i, p_j\}$ fails at stage t, then Q cannot be d-independent. In this case, we say that Q *fails at stage* t. Otherwise, Q is said to have *survived* stage t, and we write $Q \curvearrowright t$.

The failures of certain pairs at a given stage are not independent events. However, they are independent for any collection of close d-neighbor pairs (i, j) with the property that the corresponding sets $T_{i,j}$ are pairwise disjoint. To find such a collection consisting of many pairs, select at least $\frac{|H_y \cap I|}{d-1} - 1$ pairs of d-neighbors from each H_y with pairwise disjoint sets $T_{i,j}$, and thus a total of at least $\frac{|I|}{d-1} - L^{t-1}$ pairs. Since the corresponding sets $S_{i,j}$ are pairwise disjoint, all but at most n/L of them are close d-neighbors. Thus, as long as $|I| \geq 3(d-1)n/L$ and $t \leq \log n / \log L$, we obtain collection of

$$m \geq \frac{|I|}{d-1} - L^{t-1} - \frac{n}{L} \geq \frac{n}{L}$$

close d-neighbors with the required property.

If any pair of this collection fails at stage t, then Q fails at this stage. As in the proof of Theorem 1, we have

$$\mathsf{Prob}(Q \curvearrowright t | Q \curvearrowright t' \text{ for all } t' < t) \leq e^{-\frac{n}{4L^d}}$$

and

$$\mathsf{Prob}(Q \text{ survives all stages }) \leq \exp\left(-\frac{n}{4L^d}\left(\frac{\log n}{\log L} - 1\right)\right).$$

Letting

$$L = \left\lfloor \frac{\log^{1/(d-1)} n}{100 \log^2 \log n} \right\rfloor \quad \text{and} \quad a = \left\lceil \frac{3(d-1)n}{L} \right\rceil,$$

we obtain

$$\mathsf{Prob}\left(\alpha(D(P)) \geq a\right) < \binom{n}{a} \exp\left(-\frac{n}{4L^d}\left(\frac{\log n}{\log L} - 1\right)\right) \to 0.$$

□

5 Concluding remarks, open problems

The notion of Delaunay graphs for axis-parallel boxes naturally generalizes to higher dimensions. An easy extension of the proof of Theorem 2 proves that for any fixed d, the Delaunay graph of randomly and uniformly selected points in the d-dimensional unit cube has expected average degree $O((\log n)^d)$. This implies that random Delaunay graphs have independent sets of size $n^{1-o(1)}$ in higher dimensions, too. All upper bounds on the independence number that apply to dimension d also apply to every larger dimension. This can easily be seen by projecting a d-dimensional point sets to a coordinate hyperplane. Delaunay graphs can only lose edges under this operation.

In general, by repeated application of the Erdős-Szekeres lemma it is easy to show that the independence number of the Delaunay graph of any set of n points in d-dimensions, with respect to axis-parallel boxes, is at least $\Omega(n^{1/2^{d-1}})$. As far as we know, no significant improvement on this bound is known, although the truth may well be $\Omega(n^{1-o(1)})$, for any fixed d.

Returning to the plane, it is not hard show that the expected number of d-tuples T in a randomly and uniformly selected set P of n points in the plane, for which there exists an axis-parallel rectangle whose intersection with P is T, is $\Theta(d^2 n \log n)$. By a result of Spencer [Sp72], any d-uniform hypergraph with

n vertices and $\Theta(nk)$ edges has an independent set of size $\Omega(n/k^{1/(d-1)})$. Therefore, P contains a d-independent subset of size $\Omega(n/\log^{1/(d-1)} n)$. This is within $O(\log^2 \log n)$ of our upper bound.

Acknowledgement. Lemma 6 has been proved independently by Sariel Har-Peled (personal communication). We are indebted to him and to Shakhar Smorodinsky for many interesting discussions on the subject.

References

[AjEG07] D. Ajwani, K. Elbassioni, S. Govindarajan, and S. Ray: Conflict-free coloring for rectangle ranges using $\tilde{O}(n^{.382} + \varepsilon)$ colors, in: em Proc. 19th ACM Symp. on Parallelism in Algorithms and Architectures (SPAA 07), 181–187.

[AlS06] N. Alon and S. Smorodinsky: Conflict-free colorings of shallow discs, in: *Proc. 22nd Ann. ACM Symp. on Computational Geometry (SoCG 2006)*, 41–43.

[BeCh87] J. Beck and W. Chen: *Irregularities of Distribution. Cambridge Tracts in Mathematics* **89**, Cambridge University Press, Cambridge, 1987.

[BKOS00] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf: *Computational Geometry. Algorithms and Applications. 2nd ed.*, Springer–Verlag, Berlin, 2000.

[BrMP05] P. Brass, W. Moser, and J. Pach: *Research Problems in Discrete Geometry*, Springer-Verlag, New York, 2005.

[Ch00] B. Chazelle: *The Discrepancy Method. Randomness and Complexity*, Cambridge University Press, Cambridge, 2000.

[ClS89] K.L. Clarkson and P. W. Shor: Applications of random sampling in computational geometry. II, *Discrete Comput. Geom.* **4** (1989), 387–421.

[DeH91] F. Dehne, A.-L. Hassenklover, J.-R. Sack, and N. Santoro: Computational Geometry algorithms for the systolic screen, *Algorithmica* **6** (1991), 734–761.

[ElM06] K. Elbassioni and N. H. Mustafa: Conflict-free colorings of rectangles ranges, in *23rd Internat. Symp. on Theoretical Aspects of Comp. Science (STACS 2006)*, 254–263.

[ELRS03] G. Even, Z. Lotker, D. Ron, and S. Smorodinsky: Conflict–free colorings of simple geometric regions with applications to frequency assignment in cellular networks, *SIAM J. Comput.* **33** (2003), 94–136.

[FiLM05] A. Fiat, M. Levy, J. Matoušek, E. Mossel, J. Pach, M. Sharir, S. Smorodinsky, U. Wagner, and E. Welzl: Online conflict-free coloring for intervals, in: *Proc. 16th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA 2005)*, 545–554.

[GuNO85] R.H. Güting, O. Nurmi, and T. Ottmann: The direct dominance problem, in: *Proc. First ACM Symp. on Computational Geometry*, Baltimore, 1985, 81–88.

[GuNO89] R.H. Güting, O. Nurmi, and T. Ottmann: Fast algorithms for direct enclosures and direct dominances, *J. of Algorithms* **10** (1989), 170–186.

[HaS05] S. Har-Peled and S. Smorodinsky: Conflict-free coloring of points and simple regions in the plane, *Discrete & Computational Geometry* **34** (2005), 47–70.

[JaT92] J.W. Jaromczyk and G.T. Toussaint: Relative neighborhood graphs and their relatives, *Proc. IEEE* **80** (1992), 1502–1517.

[KeL86] K. Kedem, R. Livne, J. Pach, and M. Sharir: On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete Comput. Geom.* **1** (1986), 59–71.

[KrN91] I. Kříž and J. Nešetřil: Chromatic number of Hasse diagrams, eyebrows and dimension, *Order* **8** (1991), 41–48.

[Ma99] J. Matoušek: *Geometric Discrepancy. An Illustrated Guide. Algorithms and Combinatorics* **18**, Springer-Verlag, Berlin, 1999.

[MaP06] J. Matoušek and A. Přívětivý: The minimum independence number of a Hasse diagram, *Combinatorics, Probability and Computing* **15** (2006), 473–475.

[OvW88] M. Overmars and D. Woods: On rectangular visibility, *J. Algorithms* **9** (1988), 372–390.

[Pa86] J. Pach: Covering the plane with convex polygons, *Discrete Comput. Geom.* **1** (1986), 73–81.

[PaTT07] J. Pach, G. Tardos, and G. Tóth: Indecomposable coverings, in: *Discrete Geometry, Combinatorics and Graph Theory, The China–Japan Joint Conference (CJCDGCGT 2005), Lecture Notes in Computer Science* **4381**, Springer, 2007, 135–148.

[PaT03] J. Pach and G. Tóth: Conflict-free colorings, in: *Discrete and Computational Geometry, Algorithms Combin.*, Vol. 25, Springer, Berlin, 2003, 665–671.

[Sm06] S. Smorodinsky: On the chromatic number of some geometric hypergraphs, in: *Proc. 16th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA 2006)*, 316–323.

[Sp72] J. Spencer: Turán's theorem for k-graphs, *Discrete Math.* **2** (1972), 183–186.