8.3.6 Let n be a large integer. Using Theorem 1.6, show that all but at most $o(n^2)$ elements of $[1,n]\cdot[1,n]$ have $(2+o(1))\log\log n$ prime divisors. (Note that the convergence of the sum $\sum_{m=1}^{\infty}\frac{1}{m^2}$ shows that one can neglect those elements which have a large square factor.) Conclude that $|[1,n]\cdot[1,n]| = o(n^2)$. For much more precise estimates, see [106].

8.4 Cell decompositions and the distinct distances problem

Given a finite point set $P \subset \mathbf{R}^d$, let $g(P) := \{|x-y| : x, y \in P\}$ denote the number of distinct distances between the elements of P. Define $g_d(n) = \min_{P\subset\mathbf{R}^d, |P|=n} g(P)$. The well-known *distinct distances problem* of Erdős, posed in 1946 [83], asks to determine the correct rate of growth of $g_d(n)$ in n for each fixed d; this question remains open even when $d = 2$. (Clearly we have $g_1(n) = n-1$.)

By considering the progression $P = [1, n^{1/d}]^d$ it is easy to see that $g_d(n) = O(dn^{2/d})$. Erdős and many other researchers conjecture that $g_d(n)$ is close to this upper bound; in particular it is conjectured that $g_d(n) = \Omega_{\varepsilon,d}(n^{2/d-\varepsilon})$ for any $\varepsilon > 0$.

It is quite easy to establish the lower bound $g_d(n) = \Omega_d(n^{1/d})$: see Exercise 8.4.2. There is a series of improvements for the case $d = 2$, due to Moser [252], Chung [57], Chung–Szemerédi–Trotter [60], Székely [342], Solymosi–Tóth [328], Tardos [353], Katz and Tardos [197]. The most current bound is $g_2(n) = \Omega(n^{0.8635})$ [197], using the approach in [328] combined with clever entropy arguments. Here we will present a slightly weaker bound due Székely; this argument forms the base for all the subsequent bounds mentioned above.

Theorem 8.17 *[342] We have $g_2(n) = \Omega(n^{4/5})$.*

Proof Let P be a set of n points in $\mathbf{R}^2$. Define an *isosceles triangle* to be a triple (p,q,q') of distinct points in P such that $|p-q| = |p-q'|$. We say that the isosceles triangle is *narrow* if the circular arc from q to q' with center at p contains no other points in P. We refer to the pair (q,q') as the *base* of the isosceles triangle, and p as the *apex*. For any $k \geq 1$, we say that a pair (q,q') is *k-rich* if it is the base of at least k narrow isosceles triangles, and *k-poor* otherwise.

Let N be the number of narrow isosceles triangles (p,q,q'). We shall apply a double counting argument to N. We begin with the lower bound. There are $|P|$ choices for p. Given p, the remaining $|P|-1$ points in p are contained in at most $g_2(|P|)$ circles centered at p. Let $\mathcal{C}$ be the collection of such circles, then we easily verify that the number of isosceles triangles with apex p is

$$\sum_{C\in\mathcal{C}:|C\cap P|\geq 2} 2|C\cap P|.$$

Since $\sum_{C \in \mathcal{C}} |C \cap P| = |P| - 1$, we can write the above quantity as $2|P| - O(g_2(P))$. Summing over all p we conclude that

$$N \geq 2|P|^2 - O(|P| g_2(|P|)).$$

Now we obtain the upper bound. We let $k \geq 1$ be a parameter to be chosen later, and split $N = N_{\text{rich}} + N_{\text{poor}}$, where N_{rich} (resp. N_{poor}) is the number of narrow isosceles triangles with a k-rich (resp. k-poor) base. Observe that if (p, q, q') is an isosceles triangle with a k-rich base, then the perpendicular bisector l of q, q' contains p and also contains at least k points from P. Conversely, for fixed l and p there are at most $4g_2(|P|)$ pairs (q, q') with perpendicular bisector l for which (p, q, q') is a narrow isosceles triangle; this can be seen by covering the points in $P \backslash \{p\}$ into at most $g_2(|P|)$ circles and observing that each circle contributes at most four such triangles. Applying Exercise 8.2.5 we conclude

$$N_{\text{rich}} = O(g_2(|P|) \left(\frac{|P|^2}{k^2} + |P| \log |P| \right)).$$

As for the poor triangles, consider the multi-graph drawing G whose vertices are P and whose edges are the circular arcs corresponding to narrow isosceles triangles (p, q, q') with a k-poor base. This graph has $|P|$ vertices and N_{poor} edges, and has edge multiplicity at most k. Thus by Exercise 8.1.5, we have

$$N_{\text{poor}} = O\big(k|P| + k^{1/3}|P|^{2/3} \text{cross}(G)^{1/3}\big).$$

On the other hand, since the drawing of G is contained in at most $|P| g_2(|P|)$ circles (each center $p \in P$ contributing at most $g_2(|P|)$ circles), and any two circles cross in at most two points, we see that $\text{cross}(G) \leq 2(|P| g_2(|P|))^2$; thus

$$N_{\text{poor}} = O\big(k|P| + k^{1/3}|P|^{4/3} g_2(|P|)^{2/3}\big).$$

Combining our upper bounds for N_{poor} and N_{rich} with the lower bound for N, we obtain

$$|P|^2 \leq O(|P| g_2(|P|)) + O\Big(g_2(|P|)\Big(\frac{|P|^2}{k^2} + |P| \log |P|\Big)\Big) + O\big(k|P| + k^{\frac{1}{3}}|P|^{\frac{4}{3}} g_2(|P|)^{\frac{2}{3}}\big).$$

We optimize this by setting $k := c|P|^{2/5}$ for some small constant $c > 0$, and some elementary algebra then gives $g_2(|P|) = \Omega(|P|^{4/5})$ as desired. □

The above argument generalizes to many other metrics than the Euclidean metric; see [130]. To go beyond $n^{4/5}$, however, it seems that one needs to use the finer arithmetic structure of Euclidean geometry. Very roughly, the results of [328], [353], [197] proceed by analyzing the perpendicular bisectors of all of the narrow isosceles triangles (p, q, q') with a given apex p and a k-rich base; note these bisectors are k-rich in the sense that they contain at least k points in P. Using

polar coordinates around p, one can parameterize these bisectors using the sum of the angles of q and q'. One can then use some bounds on partial sum sets to obtain non-trivial lower bounds on the number of k-rich lines through p, which can then be combined with Exercise 8.2.5 to obtain an improvement to Theorem 8.17; see [328]. The further refinements in [353], [197] proceed similarly, but with a slightly weaker notion of narrow isosceles triangle, allowing the circular arc connecting q with q' to contain $O(1)$ other points from P. This provides several further partial sum sets to yield slightly better lower bounds on the number of k-rich lines through p.

In the higher-dimensional case $d > 2$ much less is known. However there are some reasonable results if one imposes some uniform distribution on the points. Let Q^d be the standard unit cube in $\mathbf{R}^d$, centered at the origin. Let us call a finite set $P \subset \mathbf{R}^d$ *homogeneous* if $P \subset |P|^{1/d} \cdot Q$, and $|P \cap (x + Q)| = O_d(1)$ for all $x \in \mathbf{R}^d$. A good example of a homogeneous set can be obtained by starting with the progression $[1, |P|^{1/d}]^d$ and perturbing each element of this progression by an arbitrary bounded displacement.

A weakened version of Erdős' original problem asks for the number of distinct distances in a homogeneous set. Homogeneous sets are interesting for at least two reasons. First, the best known upper bounds for the distance problem are homogeneous. Second, homogeneous sets play an important role in analysis (see e.g. [189]). In this section we prove

Theorem 8.18 *[326] Let $P \subset \mathbf{R}^d$ be a homogeneous set. Then $g_d(P) = \Omega_d(|P|^{\frac{2}{d} - \frac{1}{d^2}})$.*

This should be compared with the (homogeneous) lattice example $P = [1, |P|^{1/d}]^d$, which gives $g_d(P) = O_d(|P|^{\frac{2}{d}})$. As in Theorem 8.17, the proof starts by a double counting argument applied to narrow isosceles triangles. However, crossing number and Szemerédi–Trotter type results are not available in higher dimensions, and one instead uses the more flexible technique of *cell decomposition*. Given a large, complex incidence system S, we try to break it into many pieces, each of which has only a small number of incidences. After the decomposition is achieved, an (often tricky) double counting argument concerning the number of a properly defined object yields fairly efficient bounds.

Proof We may of course assume $|P| \geq 2$. By hypothesis, P is contained in the cube $|P|^{1/d} \cdot Q$. Let $1 \leq r < |P|^{1/d}$ be an integer to be chosen later. By using hyperplanes parallel to the coordinate axes, we can partition $|P|^{1/d} \cdot Q = C_1 \cup \cdots \cup C_{r^d}$, where each C_i is a cube of side-length $|P|^{1/d}/r$; we assign the boundary points of these cubes arbitrarily to one of the cubes of the partition. We refer to the cubes C_i as *cells*.

For each $p \in P$, the set $P \backslash \{p\}$ is contained in the union of at most $g_d(P)$ spheres centered at p. We denote by $\mathcal{S}_p$ the set of these spheres. For each sphere $S \in \mathcal{S}_p$, let $\mathcal{C}_S$ denote all the cells C_i which intersect S and which contain at least one point of P; from elementary geometry we see that $|\mathcal{C}_S| = O_d(r^{d-1})$.

We now apply a double counting argument to the quantity

$$N := |\{(p, S, C, q, q') : p \in P; S \in \mathcal{S}_p; C \in \mathcal{C}_S; q, q' \in P \cap S \cap C; q \neq q'\}|;$$

informally, N counts the number of isosceles triangles in P where the base points lie in the same cell (cf. the proof of Theorem 8.17). We begin with an upper bound. Observe that there are r^d possible cells C. A cell has side-length $|P|^{1/d}/r > 1$, so by homogeneity it contains $O_d(|P|/r^d)$ points in P. Thus there are $O_d(|P|/r^d)^2$ possible pairs q, q' that can be associated to C. For each such pair, observed that p must lie on the hyperplane bisecting q and q' (since q, q' lie on a sphere centered at p). By homogeneity again, this hyperplane contains at most $O(|P|^{(d-1)/d})$ elements of P. Finally, once p, q, q' are fixed, S is completely determined. Putting this all together we obtain the upper bound

$$N \leq r^d O_d(|P|/r^d)^2 O\big(|P|^{(d-1)/d}\big) = |P|^{3-\frac{1}{d}} r^{-d}. \tag{8.6}$$

Now we obtain a lower bound. Observe the explicit formula

$$N = \sum_{p \in P} \sum_{S \in \mathcal{S}_p} \sum_{C \in \mathcal{C}_S} |P \cap S \cap C|^2 - |P \cap S \cap C|.$$

From Cauchy–Schwarz we have

$$\sum_{C \in \mathcal{C}_S} |P \cap S \cap C|^2 \geq \frac{|P \cap S|^2}{|\mathcal{C}_S|}$$

and

$$\sum_{S \in \mathcal{S}_p} |P \cap S|^2 \geq \frac{(|P| - 1)^2}{g_d(P)}$$

and hence

$$N \geq \sum_{p \in P} \frac{(|P| - 1)^2}{g_d(P)|\mathcal{C}_S|} - (|P| - 1).$$

Since $|\mathcal{C}_S| = O_d(r^{d-1})$, we conclude

$$N \geq \Omega_d\left(\frac{|P|^3}{g_d(P) r^{d-1}}\right) - |P|^2.$$

Combining this with (8.6) and rearranging, we conclude

$$g_d(P) = \Omega_d\left(\frac{|P|^3}{r^{d-1}\big(|P|^{3-\frac{1}{d}} r^{-d} + |P|^2\big)}\right).$$

We optimize this by selecting r to be the nearest integer to $|P|^{\frac{1}{d}-\frac{1}{d^2}}$, and the claim follows. □

For $d \geq 3$ and general (inhomogeneous) sets, little has been known for a long time, as the method based on the Szeméredi–Trotter theorem cannot be generalized to dimension larger than 2. Clarkson, Edelsbrunner, Gubias, Sharir and Welzl [63] proved that $g_3(n) = \Omega(n^{1/2})$. In 2002, Aronov, Pach, Sharir and Tardos [14] proved that $g_3(n) = \Omega_\epsilon(n^{77/141-\epsilon})$ for any $\epsilon > 0$. More generally, they proved that $g_d(n) = \Omega_{d,\epsilon}(n^{1/(d-90/77)-\epsilon})$ for any $d \geq 3$. This result gives a non-trivial improvement for small d, compared to the previous bound $n^{1/d}$. On the other hand, as $d \to \infty$, the exponent $1/(d - 90/77) - \epsilon$ converges to $1/d$, rather than to the conjectured bound $2/d$.

Very recently, Solymosi and Vu [327] managed to show that the exponent $2/d$ is best possible to top order, in the sense that it cannot be replaced by $(2 - \epsilon + o_{d\to\infty}(1))/d$ for any positive constant $\epsilon > 0$. More precisely, they showed that that

$$g_d(n) = \Omega_d\left(n^{\frac{2}{d}-\frac{2}{d(d+2)}}\right)$$

for all $d \geq 4$, and also $g_3(n) = \Omega(n^{.5643})$.

This result and the previous bound of Aronov et al. were proved using the decomposition method combined with other arguments. Unlike the homogeneous case, the decomposition used here is more sophisticated and was first developed by Chazelle and Friedman [52] (see also [245]), motivated by problems in geometric searching in computer science. Let us conclude this section by briefly discussing this result.

One of the main techniques for doing a search is divide-and-conquer. In many problems, the situation looks as follows: given a set B of hyperplanes (of co-dimension 1) in $\mathbf{R}^d$, one would like to partition $\mathbf{R}^d$ in not too many parts so that each part intersects only few hyperplanes.

Definition 8.19 A hyperplane H *strongly intersects* a set P if $H \cap P$ is not empty and P has a point on both side of H.

Lemma 8.20 *Let B be a set of k hyperplanes in $\mathbf{R}^d$. For any $1 \leq r \leq k$, one can partition $\mathbf{R}^d$ into r sets $P_1, \ldots, P_r$ such that for each $1 \leq i \leq r$, there are only $O(k/r^{1/d})$ planes which strongly intersect P_i.*

The bound $O(k/r^{1/d})$ is best possible; the hidden constants in O depend on d but not on r. One can also guarantee that the sets P_i are generalized simplices. Strong intersection actually means intersection with the interior (see [245]). Let us now consider a little bit more complex situation when beside B we also have a set A of n points. We can require, in addition, that each part contains not too many points.

Lemma 8.21 *Let A be a set of n points and B be a set of k hyperplanes in $\mathbf{R}^d$. For any $1 \leq r \leq k$, one can partition $\mathbf{R}^d$ into r sets $P_1, \ldots, P_r$ such that for each $1 \leq i \leq r$, $|P_i \cap A| \leq 2n/r$ and P_i strongly intersects $O(k/r^{1/d})$ planes.*

Lemma 8.21 is not restricted to hyperplanes. It still holds if we replace a family of hyperplanes by a family of surfaces satisfying certain topological conditions. In particular, the lemma holds if we replace hyperplanes by (full-dimensional) spheres (see Section 6.5 of [245]). As an analog of Lemma 8.21, we obtain the following lemma, which was actually used in [327].

Definition 8.22 A sphere S strongly intersects a set P if $S \cap P$ is not empty and P has a point on both sides of S.

Lemma 8.23 *Let A be a set of n points and B be a set of k spheres in $\mathbf{R}^d$. For any $1 \leq r \leq k$, one can partition $\mathbf{R}^d$ into r sets $P_1, \ldots, P_r$ such that for each $1 \leq i \leq r$, $|P_i \cap A| = O(n/r)$ and there are only $O(k/r^{1/d})$ spheres which strongly intersect P_i.*

It would be very desirable to have a finite field analog of the above lemmas. Here is the simplest form of the problem: given a set of lines (or simple curves) on a finite plane, we would like to partition the plane into a few parts so that each part intersects only a few lines. The main obstacle here is that one needs to find a proper replacement for the topological condition of strong intersection. This condition was used to rule out extremal cases such as when all the hyperplanes go through the same point.

Exercises

8.4.1 Let A be a finite non-empty set of reals. Show that $|k((A-A)\hat{}2)| \geq g_k(|A|^2)$, where $X\hat{}2 := \{x^2 : x \in X\}$ is the set of squares in X, and $kX := X + \cdots + X$ is the k-fold sum set of X. Thus progress on the Erdős distance problem is linked to progress on questions of sum-product type; see [43] for some further development of this idea.

8.4.2 [83] Let $x_1, \ldots, x_d$ be d points in general position in $\mathbf{R}^d$. Show that if $x \in \mathbf{R}^d$, then the d distances $|x - x_1|, \ldots, |x - x_d|$ determine x up to a multiplicity of $O_d(1)$. Use this to show that $g_d(n) = O_d(n^{1/d})$ for all n. (Note that the degenerate case in which many points lie in a lower-dimensional space can be dealt with by an induction argument.)

8.4.3 [326] (Rich lines in three dimensions) Let P be a homogeneous set in $\mathbf{R}^3$. Show that $\sum_{l:|l\cap P|\geq k} |L \cap P| = O(|P|^2/k^3)$ for all $k \geq 2$.

8.4.4 [326] Let A be a homogeneous set of cardinality n in $\mathbb{R}^3$ and $\mathcal{P}$ be a collection of D pairwise non-parallel planes. Then there is a plane $P \in \mathcal{P}$