

Chapter 3: Polynomial Partitioning

Adam Sheffer

April 19, 2015

In this chapter we finally start to talk about the polynomial method. We introduce the polynomial partitioning theorem, use this theorem to derive a bound for incidences with general algebraic curves in \mathbb{R}^2 , and then prove the theorem.

1 The polynomial partitioning theorem

Consider a set \mathcal{P} of m points in \mathbb{R}^d . Given a polynomial $f \in \mathbb{R}[x_1, \dots, x_d]$, we define the *zero set* of f to be $Z(f) = \{p \in \mathbb{R}^d \mid f(p) = 0\}$. For any $r > 1$, we say that $f \in \mathbb{R}[x_1, \dots, x_d]$ is an *r -partitioning polynomial* for \mathcal{P} if every connected component of $\mathbb{R}^d \setminus Z(f)$ contains at most m/r^d points of \mathcal{P} .¹ Notice that there is no restriction on the number of points of \mathcal{P} that lie in $Z(f)$. Figure 1 depicts a 2-partitioning polynomial for a set of 12 points in \mathbb{R}^2 .




Figure 1: A 2-partitioning polynomial for a set of 12 points in \mathbb{R}^2 .

The following result is due to Guth and Katz [4] (the hypersurface condition is by Zahl [11]).

Theorem 1.1 (Polynomial partitioning [4]). *Let \mathcal{P} be a set of m points in \mathbb{R}^d . Then for every $1 < r \leq m$, there exists an r -partitioning polynomial $f \in \mathbb{R}[x_1, \dots, x_d]$ of degree $O(r)$. Moreover, we may assume that $Z(f)$ is a hypersurface.*

To estimate the number of cells in such a partition, we rely on the following theorem.

Theorem 1.2 (Warren [10]). *For a polynomial $f \in \mathbb{R}[x_1, \dots, x_d]$ of degree k , the number of connected components of $\mathbb{R}^d \setminus Z(f)$ is $O_d(k^d)$.*

¹Currently there is no standard definition for an r -partitioning polynomial. Some authors define it to be a polynomial with every cell of $\mathbb{R}^d \setminus Z(f)$ containing at most m/r points, some authors use the notation $1/r$ -partitioning polynomial, etc. We chose the definition that in our opinion is the easiest one to work with.

By Theorem 1.2, a polynomial of degree $O(r)$ yields $O(r^d)$ cells, which is asymptotically the minimum required to have at most m/r^d points in each cell. Thus, Theorem 1.1 is asymptotically tight (usually most of the points will indeed be in the cells and not on the partition itself).

In some degenerate cases, it might happen that all of the cells of the partition are empty. For example, consider a set \mathcal{P} of m points in \mathbb{R}^2 , all on the x -axis, and let f be an r -partitioning polynomial \mathcal{P} (for r asymptotically smaller than m). If $Z(f)$ does not contain the x -axis, then by Bézout's theorem $Z(f)$ intersects the x -axis in $O(r)$ points. This in turn implies that most of the points of \mathcal{P} are in $O(r)$ cells, contradicting the property that every cell contains at most m/r^2 points of \mathcal{P} . Therefore $Z(f)$ contains the x -axis and all of the cells are empty.

2 Incidences with algebraic curves in \mathbb{R}^2

Consider a point set \mathcal{P} and a set of curves Γ , both in \mathbb{R}^2 . The *incidence graph* of $\mathcal{P} \times \Gamma$ is a bipartite graph $G = (V_1 \cup V_2, E)$, where the vertices of V_1 correspond to the points of \mathcal{P} , the vertices of V_2 correspond to the curves of Γ , and an edge $(v_i, v_j) \in E$ implies that the point that corresponds to v_i is incident to the curve that corresponds to v_j ; that is, E can be thought as the subset of incidences in $\mathcal{P} \times \Gamma$. An example is depicted in Figure 2.




Figure 2: A point-line configuration and its incidence graph.

Recall that $K_{s,t}$ is a complete bipartite graph with s vertices on one side and t vertices on the other. Our goal in this section is to prove the following theorem (variants of this result originally appeared in [3, 8]).

Theorem 2.1. *Let \mathcal{P} be a set of m points and let Γ be a set of n distinct irreducible algebraic curves of degree at most k , both in \mathbb{R}^2 . If the incidence graph of $\mathcal{P} \times \Gamma$ contains no copy of $K_{s,t}$, then*

$$I(\mathcal{P}, \Gamma) = O_{s,t,k} \left(m^{\frac{s}{2s-1}} n^{\frac{2s-2}{2s-1}} + m + n \right).$$

It is straightforward to generalize the theorem to sets of curves that are neither distinct nor irreducible. We include these restrictions only to simplify the analysis. To emphasize the strength of Theorem 2.1, we consider some common types of curves:

- If Γ is a set of lines, since two lines intersect in at most one point, the incidence graph contains no copy of $K_{2,2}$. That is, Theorem 2.1 generalizes the Szemerédi-Trotter theorem.
- If Γ is a set of unit circles, the incidence graph contains no copy of $K_{2,3}$. That is, Theorem 2.1 generalizes the best bound for the unit distances problem.
- If Γ is a set of arbitrary circles, the incidence graph contains no copy of $K_{3,2}$. In this case we obtain the bound $I(\mathcal{P}, \Gamma) = O(m^{3/5}n^{4/5} + m + n)$.

It is known that Theorem 2.1 is not tight in various cases, such as for parabolas and for arbitrary circles. The following common conjecture suggests that a significantly stronger result might hold.

Conjecture 2.2. *Let \mathcal{P} be a set of n points and let Γ be a set of n algebraic curves of degree at most k , both in \mathbb{R}^2 . If the incidence graph of $\mathcal{P} \times \Gamma$ contains no copy of $K_{s,t}$, then*

$$I(\mathcal{P}, \Gamma) = O_{s,t,k}(n^{4/3}).$$

We begin by proving a weaker incidence bound, which is purely combinatorial. This bound can be seen as a special case of the Kővari–Sós–Túran theorem from extremal graph theory (e.g., see [7, Section 4.5]). We recall *Hölder's inequality*: Given $a_1, \dots, a_n, b_1, \dots, b_n \in \mathbb{R}$, and $1 < p, q$ such that $1/p + 1/q = 1$, we have

$$\sum_{i=1}^n |a_i b_i| \leq \left(\sum_{i=1}^n |a_i|^p \right)^{1/p} \left(\sum_{i=1}^n |b_i|^q \right)^{1/q}. \quad (1)$$

Lemma 2.3. *Let \mathcal{P} be a set of m points and let Γ be a set of n algebraic curves of degree at most k , both in \mathbb{R}^2 . If the incidence graph of $\mathcal{P} \times \Gamma$ contains no copy of $K_{s,t}$, then*

$$I(\mathcal{P}, \Gamma) = O_{s,t,k}(mn^{1-\frac{1}{s}} + n).$$

Proof. Let T be the set of $(s+1)$ -tuples $(a_1, \dots, a_s, \gamma)$ such that $a_1, \dots, a_s \in \mathcal{P}$, $\gamma \in \Gamma$, and $a_1, \dots, a_s \in \gamma$. We prove the lemma by double counting $|T|$.

On one hand, there are $\binom{m}{s}$ subsets of s points of \mathcal{P} , and any such subset is contained in at most $t-1$ curves of Γ . That is,

$$|T| \leq \binom{m}{s}(t-1) = O_{s,t}(m^s). \quad (2)$$

Let $\Gamma = \{\gamma_1, \dots, \gamma_n\}$. For each $\gamma_i \in \Gamma$ put $d_i = |\mathcal{P} \cap \gamma_i|$, so that $I(\mathcal{P}, \Gamma) = \sum_{i=1}^n d_i$. We have

$$|T| = \sum_{i=1}^n \binom{d_i}{s} = \Omega_s \left(\sum_{i=1}^n (d_i - s)^s \right).$$

By applying Hölder's inequality with $a_i = d_i - s$ and $b_i = 1$ for every $1 \leq i \leq n$, and with $p = s$, we get

$$\sum_{i=1}^n (d_i - s) \leq \left(\sum_{i=1}^n (d_i - s)^s \right)^{1/s} \left(\sum_{i=1}^n 1 \right)^{(s-1)/s} = \left(\sum_{i=1}^n (d_i - s)^s \right)^{1/s} n^{(s-1)/s}.$$

Since $I(\mathcal{P}, \Gamma) = \sum_{i=1}^n d_i$, we get

$$|T| = \Omega \left(\sum_{i=1}^n (d_i - s)^s \right) = \Omega \left(\frac{(I(\mathcal{P}, \Gamma) - sn)^s}{n^{s-1}} \right). \quad (3)$$

By combining (2) and (3), we obtain

$$\frac{(I(\mathcal{P}, \Gamma) - sn)^s}{n^{s-1}} = O_{s,t}(m^s).$$

Hence $I(\mathcal{P}, \Gamma) = O_{s,t}(mn^{(s-1)/s} + n)$, as asserted. \square

We will prove Theorem 2.1 by using Theorem 1.1 to partition \mathbb{R}^2 into cells and applying the bound of Lemma 2.3 separately in each cell. That is, we amplify the weak incidence bound by combining it with polynomial partitioning. We first present some intuition for why this technique works. One way to think of the bound of Lemma 2.3 is as if each point of \mathcal{P} contributes $O(n^{(s-1)/s})$ incidences (where n is the number of curves). Once we apply the bound of the lemma separately in each cell, it is as if every point p contributes $O(n_p^{(s-1)/s})$ incidences, where n_p is the number of curves that intersect the specific cell that contains p .

Proof of Theorem 2.1. By Theorem 1.1, there exists an r -partitioning polynomial $f \in \mathbb{R}[x, y]$ for \mathcal{P} of degree $O(r)$. We may assume that f is square-free, since removing repeated factors has no effect on $Z(f)$.

Let c denote the number of cells in (i.e., connected components of) $\mathbb{R}^2 \setminus Z(f)$. We denote by $\mathcal{P}_0 = Z(f) \cap \mathcal{P}$ the set of points of \mathcal{P} that are contained in $Z(f)$. Similarly, we denote by Γ_0 the set of curves of Γ that are fully contained in $Z(f)$. For $1 \leq i \leq c$, let \mathcal{P}_i denote the set of points that are contained in the i -th cell and let Γ_i denote the set of curves of Γ that intersect the i -th cell. Notice that

$$I(\mathcal{P}, \Gamma) = I(\mathcal{P}_0, \Gamma_0) + I(\mathcal{P}_0, \Gamma \setminus \Gamma_0) + \sum_{i=1}^c I(\mathcal{P}_i, \Gamma_i).$$

We bound each of these three expressions separately.

We begin with $\sum_{i=1}^c I(\mathcal{P}_i, \Gamma_i)$, for which we require an upper bound on the number of cells c . By Theorem 1.2, we have $c = O(r^2)$. We set $m_i = |\mathcal{P}_i| \leq m/r^2$ and $n_i = |\Gamma_i|$. By applying Lemma 2.3 separately in each cell, we have

$$\sum_{i=1}^c I(\mathcal{P}_i, \Gamma_i) = O \left(\sum_{i=1}^c (m_i n_i^{\frac{s-1}{s}} + n_i) \right) = O \left(\frac{m}{r^2} \sum_{i=1}^c n_i^{\frac{s-1}{s}} + \sum_{i=1}^c n_i \right).$$

We claim that any curve $\gamma \in \Gamma$ intersects $O_k(r)$ cells of the partition. Recall that by Bézout's theorem, the number of intersection points between a curve $\gamma \in \Gamma$ and $Z(f)$ is $O_k(r)$. By Harnack's curve theorem, γ has $O_k(1)$ connected components. These two bounds do not suffice to claim that γ intersects $O_k(r)$ cells, since in any intersection point between γ and the partition, γ may "split" into several cells (e.g.,




Figure 3: The dashed lines represent the partition. In an intersection with the partition, the red curve splits into six cells.

See Figure 3). Consider such an intersection point p and let C_p be a circle that is centered at p and of a small generic radius. By Bézout theorem, γ and C_p intersect in at most $2k$ points. Thus, γ may split into $O_k(1)$ cells in an intersection point with the partition.

Since any curve $\gamma \in \Gamma$ intersects $O_k(r)$ cells of the partition, we have $\sum_{i=1}^c n_i = O(nr)$. According to Hölder's inequality (1), we have

$$\sum_{i=1}^c n_i^{\frac{s-1}{s}} \leq \left(\sum_{i=1}^c n_i \right)^{\frac{s-1}{s}} \left(\sum_{i=1}^c 1 \right)^{\frac{1}{s}} = O \left((nr)^{\frac{s-1}{s}} (r^2)^{\frac{1}{s}} \right) = O \left(n^{\frac{s-1}{s}} r^{\frac{s+1}{s}} \right).$$

Combining the above implies

$$\sum_{i=1}^c I(\mathcal{P}_i, \Gamma_i) = O \left(\frac{m}{r^2} \sum_{i=1}^c n_i^{\frac{s-1}{s}} + \sum_{i=1}^c n_i \right) = O \left(\frac{mn^{\frac{s-1}{s}}}{r^{\frac{s-1}{s}}} + nr \right). \quad (4)$$

Next, consider a curve $\gamma \in \Gamma \setminus \Gamma_0$. By Bézout's theorem, the number of intersection points between γ and $Z(f)$ is $O_k(r)$. This implies

$$I(\mathcal{P}_0, \Gamma \setminus \Gamma_0) = O(nr). \quad (5)$$

It remains to bound $I(\mathcal{P}_0, \Gamma_0)$. Notice that $Z(f)$ consists of $O(r)$ one-dimensional irreducible components. Since the curves of Γ are irreducible and distinct, each component of $Z(f)$ corresponds to at most one curve of Γ . Recall that a point that is contained in more than one component of $Z(f)$ is a singular point of $Z(f)$. Thus, every regular point of $Z(f)$ is incident to at most one curve of Γ_0 . That is, there are $O(m)$ incidences between curves of Γ_0 and points of \mathcal{P}_0 that are regular points of $Z(f)$.

Since we are in \mathbb{R}^2 , a curve is also a hypersurface so its ideal is generated by a single polynomial. More specifically, we may assume that $\langle f \rangle = \mathbf{I}(\mathbf{V}(f))$. This implies that $f_x = \frac{\partial f}{\partial x}$ vanishes on every singular point of $Z(f)$, and has no common components with f . By Bézout's theorem, $Z(\gamma) \cap Z(f_x)$ consists of $O_k(r)$ points. That is, γ is incident to $O_k(r)$ singular points of $Z(f)$. By summing this up over every $\gamma \in \Gamma_0$, we have

$$I(\mathcal{P}_0, \Gamma_0) = O(nr + m). \quad (6)$$

By combining (4), (5), and (6), we obtain

$$I(\mathcal{P}, \Gamma) = O \left(\frac{mn^{\frac{s-1}{s}}}{r^{\frac{s-1}{s}}} + nr + m \right).$$

It can be easily verified that the optimal value for r is $\Theta\left(m^{\frac{s}{2s-1}}/n^{\frac{1}{2s-1}}\right)$. Setting this value immediately implies the assertion of the theorem.

There remains a minor issue: When $m = O(n^{1/s})$, we might have $m^{\frac{s}{2s-1}}/n^{\frac{1}{2s-1}} < 1$, which may lead to an invalid value of r . Fortunately, in this case Lemma 2.3 implies the bound $I(\mathcal{P}, \Gamma) = O(n)$. \square

3 Proving the polynomial partitioning theorem

In this section we prove the polynomial partitioning theorem. To help the reader, we first repeat it.

Theorem 1.1. *Let \mathcal{P} be a set of m points in \mathbb{R}^d . Then for every $1 < r \leq m$, there exists an r -partitioning polynomial $f \in \mathbb{R}[x_1, \dots, x_d]$ of degree $O(r)$. Moreover, we may assume that $Z(f)$ is a hypersurface.*

Intuitively, to prove the theorem we iteratively partition \mathcal{P} . That is, we first partition \mathcal{P} into two disjoint sets $\mathcal{P}_1, \mathcal{P}_2 \subset \mathcal{P}$ that are not “too large” (and may not contain some of the points of \mathcal{P}). We then partition both \mathcal{P}_1 and \mathcal{P}_2 to obtain four sets $\mathcal{P}'_1, \mathcal{P}'_2, \mathcal{P}'_3, \mathcal{P}'_4$, etc. An example of this process is depicted in Figure 4.




Figure 4: Iteratively partitioning a set of 15 points in the plane.

A hyperplane h in \mathbb{R}^d bisects a finite point set $\mathcal{P} \subset \mathbb{R}^d$ if each of the two open halfspaces bounded by h contains at most $|\mathcal{P}|/2$ points of \mathcal{P} . The bisecting hyperplane may contain any number of points of \mathcal{P} . The following is a discrete version of the *ham sandwich theorem* (e.g., see [6]).

Theorem 3.1. *Every d finite point sets $\mathcal{P}_1, \dots, \mathcal{P}_d \subset \mathbb{R}^d$ can be simultaneously bisected by a hyperplane.*

A planar example of Theorem 3.1 is depicted in Figure 5. To iteratively partition \mathcal{P} , we can apply Theorem 3.1. However, after about $\log_2 d$ steps we obtain more than d sets of points, and can no longer apply the theorem. Indeed, it is not hard to find $d + 1$ sets of points that cannot be simultaneously bisected by a hyperplane. To overcome this difficulty, we instead use a discrete version of the *polynomial ham sandwich theorem*.

A polynomial f in $\mathbb{R}[x_1, \dots, x_d]$ bisects a finite point set $\mathcal{P} \subset \mathbb{R}^d$ if $f(p) > 0$ for at most $|\mathcal{P}|/2$ points $p \in \mathcal{P}$ and $f(p) < 0$ for at most $|\mathcal{P}|/2$ points $p \in \mathcal{P}$. The zero set $Z(f)$ may contain any number of points of \mathcal{P} .




Figure 5: The line simultaneously bisects the set of blue points and the set of orange points.

Theorem 3.2 (Stone and Tukey [9]). *Let $\mathcal{P}_1, \dots, \mathcal{P}_t \subset \mathbb{R}^d$ be t finite point sets, and let D be an integer such that $\binom{D+d}{d} - 1 \geq t$. Then there exists a nonzero polynomial $f \in \mathbb{R}[x_1, \dots, x_d]$ of degree at most D that simultaneously bisects all of the sets \mathcal{P}_i .*

Proof. The number of monomials that a polynomial of degree at most D in $\mathbb{R}[x_1, \dots, x_d]$ can have is $\binom{D+d}{d}$; this is illustrated in Figure 6. Let $U_D = \{(i_1, \dots, i_d) \mid 1 \leq i_1 + \dots + i_d \leq D\}$ (the set of nonconstant monomials), and notice that $m = |U_D| = \binom{D+d}{d} - 1$. The Veronese map $\nu_D : \mathbb{R}^d \rightarrow \mathbb{R}^m$ is defined as

$$\nu_D(x_1, \dots, x_d) := (x_1^{u_1} x_2^{u_2} \cdots x_d^{u_d})_{u \in U_D}.$$




Figure 6: Every monomial of degree at most D in x_1, \dots, x_d corresponds to a unique choice of the d gray blocks out of a total of $D + d$ blocks.

Every coordinate in \mathbb{R}^m corresponds to a nonconstant monomial of degree at most D in $\mathbb{R}[x_1, \dots, x_d]$, and $\nu_D(\cdot)$ maps a point p in \mathbb{R}^d to the tuple of the values of these monomials at p . For example, the Veronese map $\nu_2 : \mathbb{R}^2 \rightarrow \mathbb{R}^5$ is

$$\nu_2(x, y) = (x^2, xy, y^2, x, y).$$

For $1 \leq i \leq t$, we set $\mathcal{P}'_i = \nu_D(\mathcal{P}_i)$. That is, every \mathcal{P}'_i is a finite point set in \mathbb{R}^m . By the assumption on D , we have $m \geq t$. Thus, by Theorem 3.1 there exists a hyperplane $h \subset \mathbb{R}^m$ that simultaneously bisects all of the sets \mathcal{P}'_i . We denote the coordinates of \mathbb{R}^m as y_u (for each $u \in U$), so that h can be defined as the zero set of a linear equation of the form $h_0 + \sum_{u \in U} y_u h_u$, for a suitable set of constants h_u .

Returning to \mathbb{R}^d , we consider the polynomial $f(x_1, \dots, x_d) = h_0 + \sum_{u \in U} h_u x_1^{u_1} x_2^{u_2} \cdots x_d^{u_d}$. For any point $a \in \mathbb{R}^d$ and $a' = \nu_D(a)$, we have that $h_0 + (h_u)_{u \in U} \cdot a' = f(a)$. That is, for every point $a \in \mathbb{R}^d$, $f(a) > 0$ (resp., $f(a) < 0$) if and only if $\nu_D(a)$ is in the positive “side” of h (resp., in the negative side of h). Since h bisects every \mathcal{P}'_i , then $Z(f)$ bisects every A_i . This concludes the proof since f is of degree at most D . \square

In other words, Theorem 3.2 implies the existence of a bisecting polynomial of degree at most

$$D \leq \lceil (d! \cdot (t + 1))^{1/d} \rceil \leq 2(d! \cdot t)^{1/d}.$$

Guth and Katz [4] relied on the polynomial ham sandwich theorem to derive the polynomial partitioning theorem.

Proof of Theorem 1.1. Let $c_d = 2(d!)^{1/d}$. We first ignore the hypersurface condition. To prove the theorem without this condition, we show that there exists a sequence of polynomials f_1, f_2, \dots such that the degree of f_i is smaller than $c_d 2^{(i+1)/d} / (2^{1/d} - 1)$ and every connected component of $\mathbb{R}^d \setminus Z(f_i)$ contains at most $m/2^i$ points of \mathcal{P} . An example is depicted in Figure 4. This would complete the proof since we can then choose $f = f_t$, where t is the minimum integer satisfying $2^t > r^d$.

We prove the existence of f_1, f_2, \dots by induction. The existence of f_1 is immediately implied by Theorem 3.1, so we move to the induction step. By the induction hypothesis, there exists a polynomial f_i of degree smaller than $c_d 2^{(i+1)/d} / (2^{1/d} - 1)$ such that every connected component of $\mathbb{R}^d \setminus Z(f_i)$ contains at most $m/2^i$ points of \mathcal{P} . Since $|\mathcal{P}| = m$, the number of these connected components that contain more than $m/2^{i+1}$ points of \mathcal{P} is smaller than 2^{i+1} . Let $S_1, \dots, S_n \subset \mathcal{P}$ be the subsets of \mathcal{P} that are contained in each of these connected components (that is, $|S_i| > m/2^{i+1}$ for each i , and $n < 2^{i+1}$). By Theorem 3.2 (and the remark following it), there is a polynomial g_i of degree smaller than $c_d 2^{(i+1)/d}$ that simultaneously bisects every S_i . We can set $f_{i+1} = f_i \cdot g_i$, since every connected component of $\mathbb{R}^d \setminus Z(f_i \cdot g_i)$ contains at most $m/2^{i+1}$ points of \mathcal{P} and $f_i \cdot g_i$ is a polynomial of degree smaller than

$$\frac{c_d 2^{(i+1)/d}}{2^{1/d} - 1} + c_d 2^{(i+1)/d} = c_d 2^{(i+1)/d} \cdot \left(\frac{1}{2^{1/d} - 1} + 1 \right) = \frac{c_d 2^{(i+2)/d}}{2^{1/d} - 1}.$$

This completes the induction step, and thus also the proof.

The hypersurface condition. It remains to show that we may assume that $Z(f)$ is a hypersurface. First, notice that we cannot simply remove lower dimensional components of $Z(f)$, even though such removal cannot cause cells to merge. The issue is that such components may contain many points of \mathcal{P} , and removing them may cause some cells to contain too many points.

We rely on the following property of gradients. If $g \in \mathbb{R}[x_1, \dots, x_d]$ is an irreducible polynomial with $\mathbf{I}(U) = \langle g \rangle \subset \mathbb{R}^d$, then $\dim U < d - 1$ if and only if ∇g vanishes identically on U (e.g., see [1, Theorem 4.5.1]).²

By the above property, if there exists an irreducible factor f' of f such that $Z(f')$ is not $(d - 1)$ -dimensional, then $\nabla f'$ vanishes identically on $Z(f')$. Thus, we can choose a generic direction v , and replace the factor f' with $v \cdot \nabla f'$ (in f). Notice that $v \cdot \nabla f'$ still vanishes on $Z(f')$, but is not identically zero. This substitution results in a lower-degree polynomial with a zero set that contains $Z(f)$. We then factor this lower-degree polynomial and repeat the process until all of the factors in $Z(f)$ correspond to $(d - 1)$ -dimensional varieties. Since f is of a finite degree, this process must terminate after a finite number of steps. \square

²If you have a simple proof for this property, I would be happy to hear about it.

4 Curves containing lattice points

We conclude this chapter with a simple application of Theorem 1.1. Consider a $\sqrt{n} \times \sqrt{n}$ integer lattice \mathcal{P} in \mathbb{R}^2 . It is easy to show that any constant-degree algebraic curve is incident to $O(\sqrt{n})$ points of \mathcal{P} (e.g., by using Bézout's theorem). This bound is tight, since a line can pass through $\Theta(\sqrt{n})$ points of \mathcal{P} . We now show that every other constant-degree algebraic curve passes through an asymptotically smaller number of lattice points (it is based on similar ideas of Iosevich [5]). In the homework we will derive a stronger bound, while an even stronger bound was derived by Bombieri and Pila [2] relying on number-theoretic methods.

Claim 4.1. *Let \mathcal{P} be a $\sqrt{n} \times \sqrt{n}$ integer lattice in \mathbb{R}^2 , and let γ be an irreducible algebraic curve of degree $k \geq 2$. Then γ contains $O\left(n^{k^2/(2k^2+1)}\right)$ points of \mathcal{P} .*

Proof. Let x denote the number of points of \mathcal{P} that are incident to γ , let p be a point of \mathcal{P} that is incident to γ , and let \mathbb{T} denote the set of translations of \mathbb{R}^2 that take p to another point of \mathcal{P} . We apply each of the translations of \mathbb{T} on γ to obtain a set Γ of n copies of γ . An example is depicted in Figure 7(a,b).




Figure 7: (a) A curve γ containing lattice points. (b) Applying the translations of \mathbb{T} on γ . (c) Applying the translations of \mathbb{T} on \mathcal{P} .

Some of the translated copies of γ might contain fewer than x points of \mathcal{P} . To fix this, we also apply each translation of \mathbb{T} on the points of \mathcal{P} (see Figure 7(c)). Notice that this results in a set \mathcal{P}' of less than $4n$ distinct lattice points. To complete the proof, we double count $I(\mathcal{P}', \Gamma)$.

After inserting the additional points, each of the n copies of γ contains at least x points of \mathcal{P} . That is, $I(\mathcal{P}', \Gamma) = \Omega(nx)$.

Notice that two translated copies of an irreducible curve cannot have a common component. Thus, by Bézout's theorem any two curves of Γ have at most k^2 points in common. By Theorem 2.1 with $s = k^2 + 1$, we obtain the bound $I(\mathcal{P}', \Gamma) = O(n^{(3k^2+1)/(2k^2+1)})$. Combining our two bounds for $I(\mathcal{P}', \Gamma)$ immediately implies the assertion of the claim. \square

References

- [1] J. Bochnak, M. Coste, and M. Roy, *Real Algebraic Geometry*, Springer-Verlag, Berlin, 1998.

- [2] E. Bombieri and J. Pila, The number of integral points on arcs and ovals, *Duke Math. J.* **59** (1989), 337–357.
- [3] K. L. Clarkson, H. Edelsbrunner, L. J. Guibas, M. Sharir, and E. Welzl, Combinatorial complexity bounds for arrangements of curves and spheres, *Discrete Comput. Geom.* **5** (1990), 99–160.
- [4] L. Guth and N. H. Katz, On the Erdős distinct distances problem in the plane, *Annals of Mathematics*, **181** (2015), 155–190.
- [5] A. Iosevich, Fourier analysis and geometric combinatorics, In *Topics in mathematical analysis*, 321–335, World Sci. Publ., Hackensack, NJ, 2008.
- [6] C. Y. Lo, J. Matoušek, and W. L. Steiger, Algorithms for ham-sandwich cuts, *Discrete Comput. Geom.* **11** (1994), 433–452.
- [7] J. Matoušek, *Lectures on Discrete Geometry*, Springer Verlag, Heidelberg, 2002.
- [8] J. Pach and M. Sharir, On the number of incidences between points and curves. *Combinat. Probab. Comput.* **7** (1998), 121–127.
- [9] A. H. Stone and J. W. Tukey, Generalized sandwich theorems, *Duke Math. J.* **9** (1942), 356–359.
- [10] H. E. Warren, Lower bound for approximation by nonlinear manifolds, *Trans. Amer. Math. Soc.* **133** (1968), 167–178.
- [11] J. Zahl, An improved bound on the number of point-surface incidences in three dimensions, *Contrib. Discrete Math.*, **8** (2013), 100–121.