Cambridge University Press
978-1-108-83249-6 — Polynomial Methods and Incidence Theory
Adam Sheffer
Excerpt
More Information

1

Incidences and Classical Discrete Geometry

My most striking contribution to geometry is, no doubt, my problem on the number of distinct distances. This can be found in many of my papers on combinatorial and geometric problems.

Paul Erdős, in a survey of his favorite contributions to mathematics, compiled for the celebration of his 80th birthday (Erdős, 1993).

1.1 Introduction to Incidences

We begin our study of geometric incidences by surveying the field and deriving a few first bounds. In this chapter we only discuss classical discrete geometry from before the discovery of the new polynomial methods. This makes the current chapter rather different from the rest of the book (outrageously, it even includes some graph theory). We also learn basic tricks that are used throughout the book, such as double counting, applying the Cauchy–Schwarz inequality, and dyadic decomposition. These techniques are presented in full detail in this chapter, while some details are omitted in the following chapters.

Consider a set $\mathcal{P}$ of points and a set $\mathcal{L}$ of lines, both in $\mathbb{R}^2$. An *incidence* is a pair $(p, \ell) \in \mathcal{P} \times \mathcal{L}$ such that the point p is contained in the line ℓ. We denote the number of incidences in $\mathcal{P} \times \mathcal{L}$ as $I(\mathcal{P}, \mathcal{L})$. For example, Figure 1 (in the Introduction) depicts a configuration with nine incidences. For any m and n, Erdős constructed a set $\mathcal{P}$ of m points and a set $\mathcal{L}$ of n lines with $\Theta(m^{2/3}n^{2/3} + m + n)$ incidences. Erdős (1985) conjectured that no point-line configuration has an asymptotically larger number of incidences. This conjecture was proved by Szemerédi and Trotter in 1983.

Theorem 1.1 (The Szemerédi–Trotter theorem) *Let $\mathcal{P}$ be a set of m points and let $\mathcal{L}$ be a set of n lines, both in $\mathbb{R}^2$. Then $I(\mathcal{P}, \mathcal{L}) = O(m^{2/3}n^{2/3} + m + n)$.*

The original proof of the Szemeredi–Trotter theorem is rather involved. In this chapter we present a later elegant proof by Székely (1997). A more general algebraic proof is presented in Chapter 3.

Finding the maximum number of point-line incidences in $\mathbb{R}^2$ is one of the simplest incidence problems. It is also one of very few incidence problems that are solved asymptotically. Other problems involve incidences with circles or other types of curves, incidences with varieties in $\mathbb{R}^d$, with semi-algebraic objects in $\mathbb{R}^d$, in complex spaces $\mathbb{C}^d$, in spaces over finite fields, and much more. In each of these problems, we wish to find the maximum number of incidences between a set of points and a set of geometric objects. If you ever need to snub a discrete geometer, try pointing out how they can barely solve any of these problems after decades of work.

One reason for studying incidence problems is that they are natural combinatorial problems. Throughout this chapter, we start to see additional reasons for studying incidence problems, including:

- *Incidence problems are not purely combinatorial, but also require an understanding of the underlying geometry.* One example of this appears in Section 1.5, where we introduce the unit distances problem. This problem involves studying properties that distinguish the Euclidean metric from almost all other distance metrics.
- *Incidence results are also useful for problems that may not seem related to geometry.* In Section 1.8, we use incidences to study the sum-product problem. This problem started as a number-theoretic problem that does not involve any geometry.

1.2 First Proofs

We now develop some initial intuition about incidences. We begin by deriving our first bound for an incidence problem. This is a weak bound, but it is still useful in some cases.

Lemma 1.2 *Let $\mathcal{P}$ be a set of m points and let $\mathcal{L}$ be a set of n lines, both in $\mathbb{R}^2$. Then $I(\mathcal{P},\mathcal{L}) = O(m\sqrt{n}+n)$ and $I(\mathcal{P},\mathcal{L}) = O(n\sqrt{m}+m)$.*

Why do we say that Lemma 1.2 is weaker than Theorem 1.1? For some intuition, consider the case where $m = n$. In this case, Theorem 1.1 leads to the bound $O(n^{4/3})$, while Lemma 1.2 only gives $O(n^{3/2})$.

Proof of Lemma 1.2 We only derive $I(\mathcal{P},\mathcal{L}) = O(m\sqrt{n}+n)$. The other bound is obtained in a symmetric manner. Consider the set of triples

$$T = \left\{(a,b,\ell) \in \mathcal{P}^2 \times \mathcal{L} \ : \ a \text{ and } b \text{ are both incident to } \ell\right\}.$$

Note that T also contains triples (a,b,ℓ) where $a = b$.

Let m_j be the number of points of $\mathcal{P}$ that are incident to the jth line of $\mathcal{L}$. Then the number of triples of T that include the jth line of $\mathcal{L}$ is m_j^2. This implies that $|T| = \sum_{j=1}^n m_j^2$. Also, note that $I(\mathcal{P},\mathcal{L}) = \sum_{j=1}^n m_j$. We apply the Cauchy–Schwarz inequality (Theorem A.1). We present this first application of the inequality in full detail. Throughout the rest of the book, we skip the intermediary steps. For $1 \le j \le n$, we set $a_j = m_j$ and $b_j = 1$. The Cauchy–Schwarz inequality implies that

$$\sum_{j=1}^n m_j \le \left(\sum_{j=1}^n m_j^2\right)^{1/2} \left(\sum_{j=1}^n 1\right)^{1/2} = \left(\sum_{j=1}^n m_j^2\right)^{1/2} \cdot n^{1/2}.$$

Squaring both sides and rearranging leads to

$$|T| = \sum_{j=1}^n m_j^2 \ge \frac{\left(\sum_{j=1}^n m_j\right)^2}{n} = \frac{I(\mathcal{P},\mathcal{L})^2}{n}. \tag{1.1}$$

The number of triples $(a,b,\ell) \in T$ with $a = b$ is $I(\mathcal{P},\mathcal{L})$. The number of triples $(a,b,\ell) \in T$ with $a \ne b$ is at most $\binom{m}{2}$, since each pair of distinct a, $b \in \mathcal{P}$ is contained in at most one line of $\mathcal{L}$. Thus, $|T| \le \binom{m}{2} + I(\mathcal{P},\mathcal{L})$. Combining this with Equation (1.1) gives

$$\frac{I(\mathcal{P},\mathcal{L})^2}{n} \le \binom{m}{2} + I(\mathcal{P},\mathcal{L}). \tag{1.2}$$

When $\binom{m}{2} \ge I(\mathcal{P},\mathcal{L})$, rearranging Equation (1.2) leads to $I(\mathcal{P},\mathcal{L}) = O(mn^{1/2})$. Otherwise, rearranging Equation (1.2) leads to $I(\mathcal{P},\mathcal{L}) = O(n)$. □

To prove Lemma 1.2, we used a common combinatorial method called *double counting*. In this method, we bound some quantity X in two different ways and then compare the two bounds. This leads to new information that does not involve X. In the proof of Lemma 1.2, we derived upper and lower bounds for the size of T. By comparing these two bounds, we obtained a bound for the number of incidences. Double counting is ubiquitous in this book.

In the proof of Lemma 1.2, we did not use any geometry beyond observing that two points are contained in one line. This implies that the proof still holds after removing all the other geometric properties of the problem. That is, when replacing the lines with abstract sets of points, such that every two sets have at most one common element. For example, instead of the lines in Figure 1 (in the Introduction), we can consider the sets

$$A = \{a, d\}, \quad B = \{a, c\}, \quad C = \{a, d\}, \quad D = \{b, c, d\}.$$

In this abstract setting, the bounds of Lemma 1.2 are asymptotically tight. There exist n subsets of m elements with the above property and $\Theta(mn^{1/2})$ incidences (or $\Theta(nm^{1/2})$). Thus, to derive a stronger upper bound for point-line incidences, we must rely on additional geometric properties of lines.

We now consider an asymptotically tight lower bound for Theorem 1.1. Instead of Erdős's original construction, we present a simpler construction due to Elekes (2001).

Claim 1.3 *For every m and n there exist a set $\mathcal{P}$ of m points and a set $\mathcal{L}$ of n lines, both in $\mathbb{R}^2$, such that $I(\mathcal{P}, \mathcal{L}) = \Theta(m^{2/3}n^{2/3} + m + n)$.*

Proof The term m dominates the bound $\Theta(m^{2/3}n^{2/3}+m+n)$ when $m = \Omega(n^2)$. In this case we can simply take m points on a single line to obtain m incidences. Similarly, the term n dominates the bound when $n = \Omega(m^2)$. In this case we take n lines that pass through a single point to obtain n incidences. It remains to construct a configuration with $\Theta(m^{2/3}n^{2/3})$ incidences when $m = O(n^2)$ and $n = O(m^2)$.

Let $r = (m^2/4n)^{1/3}$ and $s = (2n^2/m)^{1/3}$ (for simplicity, instead of taking the ceiling function of s and r, we assume that these are integers). We set

$$\mathcal{P} = \{\ (i, j)\ :\ 1 \le i \le r \quad \text{and} \quad 1 \le j \le 2rs\ \},$$

and

$$\mathcal{L} = \{\ y = ax + b\ :\ 1 \le a \le s \quad \text{and} \quad 1 \le b \le rs\ \}.$$

Note that $\mathcal{P}$ is a rectangular section of the integer lattice. The slopes and y-intercepts of the lines of $\mathcal{L}$ also form such a lattice. Figure 1.1 depicts an example configuration rotated by 90°. We also have that

$$|\mathcal{P}| = 2r^2 s = 2 \cdot \frac{m^{4/3}}{(4n)^{2/3}} \cdot \frac{(2n^2)^{1/3}}{m^{1/3}} = m,$$

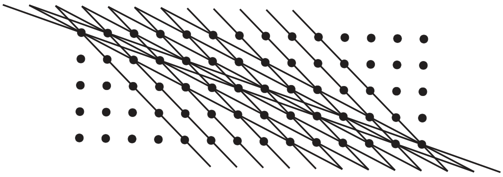

Figure 1.1 Elekes's construction, rotated by 90°.

and

$$|\mathcal{L}| = rs^2 = \frac{m^{2/3}}{(4n)^{1/3}} \cdot \frac{(2n^2)^{2/3}}{m^{2/3}} = n.$$

Consider a line $\ell \in \mathcal{L}$ that is defined by the equation $y = ax + b$. For any $x \in \{1, \dots, r\}$, there exists $y \in \{1, \dots, 2rs\}$ such that the point (x, y) is incident to ℓ. That is, every line of $\mathcal{L}$ is incident to exactly r points of $\mathcal{P}$, which in turn implies that

$$I(\mathcal{P}, \mathcal{L}) = r \cdot |\mathcal{L}| = \frac{m^{2/3}}{(4n)^{1/3}} \cdot n = 2^{-2/3} m^{2/3} n^{2/3}. \qquad \square$$

1.3 The Crossing Lemma

One elegant proof of Theorem 1.1 is based on the *crossing lemma*. We study this proof in Section 1.4. Here, we first go over some required preliminaries. For a brief review of graph theory notation, see Section A.2.

The *crossing number* of a graph $G = (V, E)$, denoted $\mathrm{cr}(G)$, is the smallest integer k such that we can draw G in the plane with k edge crossings. Figure 1.2(a) depicts a drawing of K_5 with a single crossing. Since K_5 cannot be drawn without crossings, we have that $\mathrm{cr}(K_5) = 1$. Intuitively, we expect a graph with a lot more edges than vertices to have a large crossing number. Given a graph $G = (V, E)$, we are interested in a lower bound for $\mathrm{cr}(G)$ with respect to $|V|$ and $|E|$.

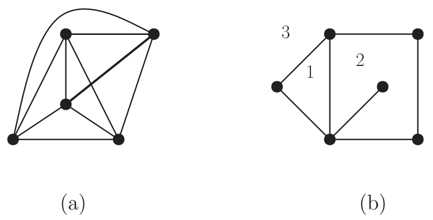

Figure 1.2 (a) A drawing of K_5 with a single crossing. (b) A graph with two bounded faces and one unbounded face.

A graph G is *planar* if $\mathrm{cr}(G) = 0$. We consider a connected planar graph $G = (V, E)$ with v vertices and e edges. More specifically, we consider a drawing of G in the plane with no crossings. The *faces* of this drawing are the maximal two-dimensional connected regions that are bounded by the edges. This includes one outer, infinitely large region. For an example, see Figure 1.2(b). Denote by f the number of faces in the drawing of G. Then *Euler's formula* states that

$$v + f = e + 2. \tag{1.3}$$

For planar graphs that are not connected, we instead have that $v + f > e + 2$.

Every edge of G is either on the boundary of two faces or has both of its sides on the boundary of the same face. Moreover, the boundary of every face of G consists of at least three edges. Thus, we have $2e \geq 3f$. Plugging this into Equation (1.3) yields

$$e \leq v + f - 2 \leq v + \frac{2e}{3} - 2.$$

That is, for any planar graph $G = (V, E)$, we have that

$$|E| \leq 3|V| - 6. \tag{1.4}$$

The above leads to our first lower bound on $\text{cr}(G)$.

Lemma 1.4 *For any graph $G = (V, E)$, we have* $\text{cr}(G) \geq |E| - 3|V| + 6$.

Proof Consider a drawing of G in the plane that minimizes the number of crossings. Let $E' \subset E$ be a maximum subset of the edges such that no two edges of E' intersect in the drawing. By Equation (1.4), we have that $|E'| \leq 3|V| - 6$. Since every edge of $E \backslash E'$ intersects at least one edge of E', and since $|E \backslash E'| \geq |E| - 3|V| + 6$, there are at least $|E| - 3|V| + 6$ crossings in the drawing. □

Since K_5 has 5 vertices and 10 edges, Lemma 1.4 gives the correct value $\text{cr}(K_5) = 1$. However, in general the bound of this lemma is rather weak. For example, it is known that $\text{cr}(K_n) = \Theta(n^4)$, while Lemma 1.4 only implies that $\text{cr}(K_n) = \Omega(n^2)$. We can amplify the lower bound of Lemma 1.4 by combining it with a probabilistic argument. The following lemma was originally derived in Ajtai et al. (1982); Leighton (1983), with different proofs.

Lemma 1.5 (The crossing lemma) *Let $G = (V, E)$ be a graph with $|E| \geq 4|V|$. Then* $\text{cr}(G) = \Omega(|E|^3/|V|^2)$.

Proof Consider a drawing of G with $\text{cr}(G)$ crossings. Set $p = \frac{4|V|}{|E|}$. The assumption of the lemma implies that $0 < p \leq 1$. We remove every vertex of V from the drawing with probability $1 - p$ (together with the edges adjacent to the vertex). Let $G' = (V', E')$ denote the resulting subgraph. Let c' denote the number of crossings in the drawing of G that have both of their edges in E'.

To avoid confusion with the edge set E, we denote expectation of a random variable as $\mathbb{E}[\cdot]$. Since every vertex remains with probability p, we have that $\mathbb{E}[|V'|] = p|V|$. Since every edge remains if and only if its two endpoints remain, we have that $\mathbb{E}[|E'|] = p^2|E|$. Finally, since each crossing remains if and only if the two corresponding edges remain, we have that $\mathbb{E}[c'] = p^4\text{cr}(G)$. By linearity of expectation,

$$\begin{aligned}\mathbb{E}[c' - |E'| + 3|V'|] &= p^4\mathrm{cr}(G) - p^2|E| + 3p|V| \\ &= \frac{4^4|V|^4}{|E|^4}\mathrm{cr}(G) - \frac{4^2|V|^2}{|E|^2}\cdot|E| + 3\cdot\frac{4|V|}{|E|}\cdot|V| \\ &= \frac{4^4|V|^4}{|E|^4}\mathrm{cr}(G) - \frac{4|V|^2}{|E|}.\end{aligned}$$

Since this is the expected value, there exists a subgraph $G^* = (V^*, E^*)$ with c^* crossings remaining from the drawing of G, such that

$$c^* - |E^*| + 3|V^*| \leq \frac{4^4|V|^4}{|E|^4}\mathrm{cr}(G) - \frac{4|V|^2}{|E|}. \tag{1.5}$$

By Lemma 1.4, we have $c^* \geq |E^*| - 3|V^*| + 6$. Combining this with Inequality (1.5) implies

$$0 < 6 \leq c^* - |E^*| + 3|V^*| \leq \frac{4^4|V|^4}{|E|^4}\mathrm{cr}(G) - \frac{4|V|^2}{|E|}.$$

That is, $\frac{4|V|^2}{|E|} < \frac{4^4|V|^4}{|E|^4}\mathrm{cr}(G)$. Tidying up this inequality leads to the required bound. □

Lemma 1.5 implies the asymptotically tight bound $\mathrm{cr}(K_n) = \Omega(n^4)$.

1.4 Szemerédi–Trotter via the Crossing Lemma

We are now ready to prove Theorem 1.1. We first restate this theorem.

Theorem 1.1 *Let $\mathcal{P}$ be a set of m points and let $\mathcal{L}$ be a set of n lines, both in $\mathbb{R}^2$. Then $I(\mathcal{P}, \mathcal{L}) = O(m^{2/3}n^{2/3} + m + n)$.*

Proof We write $\mathcal{L} = \{\ell_1, \ldots, \ell_n\}$ and denote by m_j the number of points of $\mathcal{P}$ that are on ℓ_j. Notice that $I(\mathcal{P}, \mathcal{L}) = \sum_{j=1}^n m_j$. We may remove any line ℓ_j that satisfies $m_j = 0$, since this would not change the number of incidences.

We build a graph $G = (V, E)$ as follows. Every vertex of V corresponds to a point of $\mathcal{P}$. For $v, u \in V$, we add (v, u) to E if v and u correspond to consecutive points along a line of $\mathcal{L}$. For an example, see Figure 1.3. A line ℓ_j contributes exactly $m_j - 1$ edges of E. Thus, we have $|V| = m$ and $|E| = \sum_{j=1}^n (m_j - 1) = I(\mathcal{P}, \mathcal{L}) - n$.

If $|E| < 4|V|$ then $I(\mathcal{P}, \mathcal{L}) = O(m + n)$, which completes the proof. We may thus assume that $|E| \geq 4|V|$. Then, Lemma 1.5 leads to

$$\mathrm{cr}(G) = \Omega\left(\frac{(I(\mathcal{P}, \mathcal{L}) - n)^3}{m^2}\right). \tag{1.6}$$

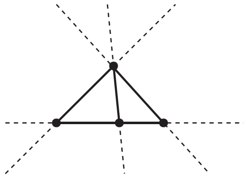

Figure 1.3 (Solid segment) The edges of the graph. (Dashed segment) The portions of the lines ℓ_j that do not form graph edges.

We draw G according to the point-line configuration: Every vertex is at the corresponding point and every edge is the corresponding line segment. Every crossing in this drawing is an intersection of two lines of $\mathcal{L}$. Since every two lines intersect at most once, we have that $\text{cr}(G) \le \binom{n}{2} = O(n^2)$. Combining this with Equation (1.6) implies that

$$\frac{(I(\mathcal{P},\mathcal{L})-n)^3}{m^2} = O(n^2).$$

Rearranging this equation gives $I(\mathcal{P},\mathcal{L}) = O(m^{2/3}n^{2/3}+n)$. □

The proof of Theorem 1.1 is another example of the double counting method. We counted $\text{cr}(G)$ in two different ways. By combining the two resulting bounds, we obtained a bound on the number of incidences.

In the proof of Theorem 1.1, we used the geometric property that two lines intersect at most once. This is similar to the observation that any two points are contained in one line,[1] which was used in the proof of Lemma 1.2. In the proof of Theorem 1.1 we used a second geometric property when stating that the line ℓ_j corresponds to exactly $m_j - 1$ edges of E. This statement relies on the observation that a line consists of a single connected component and does not intersect itself. When replacing the lines with other curves that satisfy the same geometric properties, the proof of Theorem 1.1 remains valid.

1.5 The Unit Distances Problem

The *unit distances problem* is one of the main open problems in discrete geometry. While it is extremely difficult to solve this problem, it easy to state:

In a set of n points in the plane, what is the maximum possible number of pairs of points at distance 1 from each other?

[1] These two geometric properties are equivalent when studying point-line incidences, due to point-line duality. We discuss this concept in Section 1.10.

We denote this maximum number of pairs as $u(n)$. By taking a set of n points equally spaced on a line, we immediately obtain that $u(n) \ge n-1$. Erdős (1946) introduced the problem, while also deriving the bounds $u(n) = O(n^{3/2})$ and $u(n) = \Omega(n^{1+c/\log\log n})$, for some constant c. While many mathematicians have studied this problem, the lower bound for $u(n)$ has not been improved since 1946 and the upper bound was last improved in 1984. That was when Spencer et al. (1984) derived the bound $u(n) = O(n^{4/3})$.

Consider a set $\mathcal{P} \subset \mathbb{R}^2$ of n points such that the number of unit distances between pairs of points of $\mathcal{P}$ is $u(n)$. We draw a unit circle (a circle of radius one) around each point of $\mathcal{P}$, and denote the set of these n circles as C. Every two points $p, q \in \mathcal{P}$ that determine a unit distance correspond to two incidences in $\mathcal{P} \times C$: The circle around p is incident to q and vice versa. See Figure 1.4 for an example. Thus, to bound $u(n)$ it suffices to bound the maximum number of incidences between n points and n unit circles (it is not difficult to show that this maximum number of incidences is asymptotically equivalent to $u(n)$).

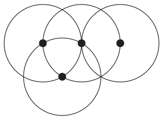

Figure 1.4 Every two points that are at a unit distance correspond to two point-circle incidences.

Theorem 1.6 *Let $\mathcal{P}$ be a set of n points and let C be a set of n unit circles, both in $\mathbb{R}^2$. Then $I(\mathcal{P}, C) = O(n^{4/3})$.*

Theorem 1.6 immediately implies the current best bound $u(n) = O(n^{4/3})$.

Proof of Theorem 1.6 We imitate the proof of Theorem 1.1. Let $C = \{c_1, \ldots, c_n\}$ and let m_j denote the number of points of $\mathcal{P}$ on c_j. Note that $I(\mathcal{P}, C) = \sum_{j=1}^n m_j$. We may remove any circle c_j that satisfies $m_j < 3$, since this reduces the number of incidences by at most $2n$.

We build a graph $G = (V, E)$ as follows. Every vertex of V corresponds to a point of $\mathcal{P}$. For $v, u \in V$, the edge (v, u) is in E if v and u are consecutive points along at least one circle of C. A circle c_j corresponds to exactly m_j edges of E, and every edge originates from at most two unit circles. Note that $|V| = n$ and $|E| \ge \left(\sum_{j=1}^n m_j\right)/2 = I(\mathcal{P}, C)/2$.

If $|E| < 4|V|$ then $I(\mathcal{P}, C) = O(n)$, which completes the proof. We may thus assume that $|E| \ge 4|V|$. By Lemma 1.5, we have that

$$\mathrm{cr}(G) = \Omega\left(\frac{I(\mathcal{P}, C)^3}{n^2}\right). \tag{1.7}$$

We draw G according to the point-circle configuration: Every vertex is at the corresponding point and every edge is one of the corresponding circle arcs. Every crossing in this drawing is the intersection of two circles of C. Since every two circles intersect at most twice, we have that $\text{cr}(G) \le 2\binom{n}{2} = O(n^2)$. Combining this with Equation (1.7) implies that

$$\frac{I(\mathcal{P},C)^3}{n^2} = O(n^2).$$

Rearranging this equation leads to $I(\mathcal{P},C) = O(n^{4/3})$. $\square$

Erdős offered \$250 for proving the following conjecture (Erdős, 1985).

Conjecture 1.7 (Erdős, 1985) $u(n) = O(n^{1+\varepsilon})$ *for any* $\varepsilon > 0$.

This is an example of how little we currently know about incidences. While the problem of point-line incidences in $\mathbb{R}^2$ has been settled for decades, the case of unit circles remains wide open. Hardly any other incidence problems have been solved.

The answer to the unit distances problem significantly depends on the metric:

- For Euclidean distance, this is a long-standing difficult problem.
- For some metrics, there exist sets of n points that span $\Theta(n^2)$ unit distances. See Exercise 1.3.
- Valtr (2005) discovered a well-behaved metric for which $u(n) = \Theta(n^{4/3})$.
- Matoušek (2011) showed that, for most metrics,[2] $u(n) = O(n \log n \log\log n)$.

The bound that is conjectured for the Euclidean distance is different from all other bounds stated above. One may thus say that the unit distances problem is about studying properties of the underlying geometry. A proof of Conjecture 1.7 is likely to require properties that are unique for the Euclidean metric.

1.6 The Distinct Distances Problem

The *distinct distances* problem is a close relative of the unit distances problem. Both problems were introduced in the same 1946 paper of Erdős. For a set $\mathcal{P} \subset \mathbb{R}^2$, let $\Delta(\mathcal{P})$ denote the set of distances spanned by pairs of points of $\mathcal{P}$. Every distance appears in $\Delta(\mathcal{P})$ at most once, no matter how many pairs of points span it. This is why we refer to $\Delta(\mathcal{P})$ as the *set of distinct distances of* $\mathcal{P}$. See Figure 1.5 for an example. The distinct distances problem asks for $\min_{|\mathcal{P}|=n} |D(\mathcal{P})|$. In other words:

[2] The exact meaning of "most metrics" is beyond the scope of this chapter.