



The Raymond and Beverly Sackler Faculty of Exact Sciences
The Blavatnik School of Computer Science

Geometric Incidences and Repeated Configurations

Thesis submitted for the degree of “Doctor of Philosophy”
by

Roel Apfelbaum

This work was carried out under the supervision of
PROF. MICHA SHARIR

Submitted to the Senate of Tel-Aviv University
November 2012

To my mother, father, wife, sisters, son, and all my family. Many thanks for your support and encouragement.

Abstract

In this thesis we present several results in the area of geometric incidences and repeated configurations. These include improved upper and lower bounds on various incidence-related quantities.

In Chapter 3, we derive a fundamental property of point-hyperplane incidences: In point-hyperplane configurations with many incidences, there always exists a large complete subconfiguration, i.e., a relatively large subset of the points together with a relatively large subset of the hyperplanes, such that, all the points and hyperplanes from both subsets are incident to one another. This property is implicit in practically all previous studies on point-hyperplane incidences, but was never explicitly stated and quantified before. We provide thresholds for the number of incidences above which a large complete subconfiguration must exist, and derive upper and lower bounds on the size of such a subconfiguration in terms of the numbers of points, hyperplanes, and incidences between them. In \mathbb{R}^3 , we show that if a set of m points together with a family of n planes have I incidences between them, where $I = \Omega(mn^{1/2} + nm^{1/2})$, then there exist a subset of r points, and a subfamily of s planes, all incident to one another, such that $rs \geq \frac{I^2}{mn} - O(m+n)$. This bound is tight, up to the additive term, because, as we show, there exist point-plane configurations with m points, n planes, and $I = \Omega(mn^{1/2} + nm^{1/2})$ incidences, such that any complete subconfiguration of r points and s planes has $rs \leq \frac{I^2}{mn} - \frac{1}{16}(m+n)$. We also provide tight bounds for the order of magnitude of the largest complete point-plane subconfiguration in \mathbb{R}^3 in case $I = O(mn^{1/2} + nm^{1/2})$. In \mathbb{R}^d , we show that if a set of m points together with a set of n hyperplanes have $I = \Omega(mn^{1-1/(d-1)} + nm^{1-1/(d-1)})$ incidences between them, then there exists a complete subconfiguration of r points and s hyperplanes such that $rs = \Omega\left(\left(\frac{I}{mn}\right)^{d-1} mn\right)$. On the other hand, we construct a set of m points and a family of n hyperplanes, for general m and n , with $I = \Omega((mn)^{1-1/(d-1)})$ incidences between them, such that any complete subconfiguration of r points and s hyperplanes has $rs = O\left(\left(\frac{I}{mn}\right)^{(d+1)/2} mn\right)$.

In Chapter 4, we study the problem of upper bounding the number of triangles of the same area spanned by a set of n points in the plane. We improve upon a previous recent bound of $O(n^{44/19})$ due to Dumitrescu et al. [33], and show that n points in the plane span at most $O(n^{9/4})$ same-area triangles^{1 2}.

In Chapter 5, we study incidences between points and *nondegenerate* spheres in \mathbb{R}^3 . Given a set of points P in \mathbb{R}^3 , some parameter $0 < \eta < 1$, and a sphere $\sigma \subset \mathbb{R}^3$, we say that σ is η -degenerate with respect to P , if there exists a circle $\gamma \subset \sigma$ incident to at least an η -fraction of the incidences of σ with P . That is, $|\gamma \cap P| \geq \eta|\sigma \cap P|$. If no such circle γ exists on σ , then we say that σ is η -nondegenerate. In [4] (Agarwal et al.), it was shown that for any parameter $0 < \eta < \frac{1}{2}$, for any set P of n points in \mathbb{R}^3 , and for any integer $k \leq n$, the number of η -nondegenerate spheres with respect to P that are k -rich (i.e., contain at least k points of P) is at most $O\left(\frac{n^4}{k^5} + \frac{n^3}{k^3}\right)$, with the

¹Originally, we have published a paper [13] (coauthored with Sharir) showing a bound of $O^*(n^{9/4})$. A very recent work of Zahl [84] has enabled us to get rid of the subpolynomial factor.

²The notation $O^*(\cdot)$ specifies an upper bound, while neglecting subpolynomial factors. That is, the statement $I = O^*(f(n))$ is equivalent to saying that the bound $I = O(f(n)n^\varepsilon)$ holds for any arbitrarily small $\varepsilon > 0$, where the constant of proportionality may depend on ε .

constant of proportionality depending on η . We improve this bound in Chapter 5 to $O^*(\frac{n^4}{k^{11/2}} + \frac{n^2}{k^2})$, and show it to apply for all $0 < \eta < 1$. This improved bound is later used in Chapter 6 to improve the upper bound on the number of mutually similar triangles spanned by n points in \mathbb{R}^3 .

In Chapter 6 we study mutually similar k -simplices spanned by n points in \mathbb{R}^d , for any $d \geq 3$, and $k \leq d$. We provide several upper bounds on the maximum possible number of such simplices for various values of k and d . A preliminary version of this study, coauthored with Agarwal, Purdy, and Sharir, was presented in [4], and Chapter 6 elaborates and improves upon the results therein. Concerning similar triangles, we show that n points in \mathbb{R}^3 span at most $O^*(n^{58/27})$ mutually similar triangles, that the number of mutually similar triangles in \mathbb{R}^4 is $O(n^{12/5})$, and that the number of mutually similar triangles in \mathbb{R}^5 is $O(n^{8/3})$. For general $d \geq 4$, and $k = d - 1$ or $k = d - 2$, we show that n points in \mathbb{R}^d span at most $O(n^{d-2+o(1)})$ mutually similar k -simplices, where the $o(1)$ term is exponentially decaying in d .

Acknowledgement

I am deeply grateful and indebted to my supervisor, Professor Micha Sharir, for introducing me to the field of computational geometry, and for his endless help and guidance.

I also wish to thank an anonymous reviewer for his extremely thorough review. His insightful comments have greatly improved this thesis.

Contents

I	Background and Preliminaries	3
1	Overview	5
1.1	Point-line incidences	5
1.2	Point-circle incidences	7
1.3	Nonoverlapping lenses in arrangements of curves	8
1.4	Other point-curve incidence problems in the plane	9
1.5	Point-hyperplane incidences	10
1.5.1	Incidences between hyperplanes and vertices of their arrangement.	10
1.5.2	No $K_{r,s}$ incidence subgraph.	10
1.5.3	Incidences with nondegenerate hyperplanes.	11
1.5.4	Representation complexity of point-hyperplane incidence graphs.	12
1.6	Other point-surface incidence results	13
1.7	Distinct distances	13
1.8	Repeated distances	14
1.9	Repeated triangles and angles	15
1.9.1	Repeated angles	15
1.9.2	Repeated areas	16
1.10	Repeated simplices	16
1.11	Summary of the results of the thesis	17
2	Techniques	19
2.1	Space decomposition	19
2.1.1	Cuttings	19
2.1.2	Partition trees	24
2.1.3	Partition by polynomials	25
2.2	Crossing numbers in graph drawings	27
2.2.1	Cutting curves into pseudo-segments	29
II	Results	31
3	Point-Hyperplane Incidences	33

3.1	Introduction	33
3.1.1	Our results	34
3.2	Large complete bipartite incidence subgraphs in \mathbb{R}^3	35
3.3	Large complete bipartite incidence subgraphs in higher dimensions	41
3.4	Conclusion	49
4	Unit Area Triangles	51
4.1	Introduction	51
4.2	Unit-area triangles in the plane	51
4.3	Appendix: Asymptotes of cubic curves	58
5	Nondegenerate Spheres	63
5.1	Introduction	63
5.2	Our results	64
5.3	Proof of Theorem 5.2.2	65
5.4	Extending the results to partial incidence graphs	70
5.5	Conclusion	71
6	Repeated Simplices in d Dimensions	73
6.1	Introduction	73
6.2	Similar triangles in three dimensions	75
6.2.1	A simpler and weaker bound	75
6.2.2	An improved bound	76
6.3	General bounds in higher dimensions	80
6.3.1	Completing a partial simplex	80
6.3.2	Bounding the number of congruent anchored simplices	88
6.4	Similar simplices	97
6.5	Similar triangles in five dimensions	98
6.6	Discussion	99

Part I

Background and Preliminaries

Chapter 1

Overview

Combinatorial geometry is the study of the combinatorial configurations between geometric objects such as points, lines, circles, etc. Incidence problems are a subcategory that focuses on upper bounds on the maximum number of incidences between geometric objects, and lower bound constructions for those values. For example, given a set of points P , and a set of lines L , an *incidence* is a point-line pair $(p, \ell) \in P \times L$, such that $p \in \ell$. The problem at hand is to bound the number $I(P, L)$ of incidences between the points of P and the lines of L , and to show that the bound is tight by some construction of points and lines that achieves the claimed upper bound.

More generally, if we think of the points P and lines L (or any other family of geometric objects) as vertices of a graph, and of the incidence pairs as edges in that graph, we get a bipartite graph, called the *incidence graph*, and denoted by $G(P, L)$. Incidence problems are concerned with studying the properties of the incidence graph, most notably, but not exclusively, the number of edges in that graph as a function of the number of vertices on both sides, and the class of underlying geometric objects.

Incidence problems play an important role in the design and analysis of geometric algorithms. The complexity of many such algorithms depends on some related incidence bound. On the other hand, the study of incidences has resulted in better understanding of the algorithmic aspects, and has yielded several useful computational techniques, most notably, the divide-and-conquer technique known as *cutting* (see Subsection 2.1.1 for details).

In this chapter we review the known results on various problems of geometric incidences, as well as the major open problems, and then, in Chapter 2, we elaborate about the theory and tools that are employed in the analysis and derivation of these results.

1.1 Point-line incidences

This is the most basic problem: What is the maximum possible number of incidences between n points and l lines in the plane? Szemerédi and Trotter [80] showed that this number is $\Theta(n^{2/3}l^{2/3} + n + l)$, thus settling in the affirmative a conjecture of Erdős [44] (for the case $n = l$). Erdős showed that an $n^{1/2} \times n^{1/2}$ lattice section together with the n richest lines (in terms of the number of incident points) achieve $\Omega(n^{4/3})$ incidences, and conjectured that this order of magnitude is best

possible (If we take the $n^{1/2} < l < \binom{n}{2}$ richest lines in the lattice section, we get $\Omega(n^{2/3}l^{2/3})$ incidences. Also, to construct $n + l - 1$ incidences is trivial). The conjecture has been settled in [80], and in several subsequent papers (e.g. [31, 79]) by simpler methods.

Elekes [37] later gave a simpler construction of $\Omega(n^{2/3}l^{2/3})$ incidences (simpler to analyze, that is) as follows. Suppose $2l^{1/2} < n < 2l^2$. Put $k = 4^{-1/3}n^{2/3}l^{-1/3}$ and $m = 2^{1/3}l^{2/3}n^{-1/3}$ (both $k, m \geq 1$), and let

$$P = \{(x, y) \in \mathbb{Z}^2 \mid 0 \leq x < k, 0 \leq y < 2mk\}$$

be a $k \times 2mk$ lattice section. The set of lines L is defined to be all lines of the form $y = ax + b$, where $0 \leq a < m$ and $0 \leq b < mk$ are integers. We have $|P| = 2mk^2 = n$ points and $|L| = m^2k = l$ lines. Each line is incident to k points so the number of incidences is $I = m^2k^2 = 2^{-2/3}n^{2/3}l^{2/3}$.

Since the exact order of magnitude of the point-line incidence bound has been found, one remaining challenge is to find the exact constant of proportionality of the $n^{2/3}l^{2/3}$ term. The Elekes construction shows that this constant is at least $2^{-2/3} \approx 0.63$. The currently best upper bound is ≈ 2.5 [67]. (These values are very small compared with the ‘‘astronomical’’ constant of 10^{60} in the original proof [80].)

One can ask the same question for points and lines in space, namely, how many incidences can there be between n points and l lines in \mathbb{R}^3 (or in \mathbb{R}^d for any $d > 2$). The answer to this question, however, is trivial, because one can project any point-line configuration in space onto some generic plane, and apply the planar bound, which implies that n points and l lines in space of any dimension can have at most $O(n^{2/3}l^{2/3} + n + l)$ incidences.

To make the spatial problem more interesting, one has to eliminate ‘‘planar’’ configurations by some additional requirements. For example, one can require that the n points must all be *joints*, where a joint is defined as a meeting point of three or more *non-coplanar* lines. In this context there are two questions to be asked: (a) How many joints can there be (as a function of the number of lines l)? and (b) How many incidences can there be between l lines and n of their joints? Consider k planes in general position in \mathbb{R}^3 . In the arrangement of these planes, the vertices are joints of the intersection lines between the pairs of planes. This simple construction shows that the number of joints can be $\Omega(l^{3/2})$, as the number of lines is $l = \binom{k}{2}$, whereas the number of joints is $\binom{k}{3} = \Theta(l^{3/2})$. It was conjectured that this construction is best possible up to a multiplicative constant [28]. Recently, Guth and Katz [50] settled this conjecture in the affirmative, showing that the number of joints is indeed at most $O(l^{3/2})$. In doing so, they introduced a new and exciting technique based on algebra and differential geometry. Consequently, their proof was greatly simplified by Kaplan et al. [55], and Quilodran [71], and their results extended by Elekes et al. [39], who showed that the number of incidences between l lines and n of their joints, if $n = \Omega(l)$, is $O(ln^{1/3})$. On the other hand, if $n = O(l)$, the number of incidences can be as high as the Szemerédi-Trotter bound $\Theta(n^{2/3}l^{2/3} + l)$ (construct a planar configuration of n points and $l - n$ lines that achieves these many incidences, and pierce the plane at the n points by n additional lines).

Alternatively, one can require that no more than a small number of the lines or of the points are coplanar. For these variants, the following tight bounds are known

Theorem 1.1.1 (Elekes et al. [39]). *Let L be a set of (at most) l lines in \mathbb{R}^3 and let P be a set of n arbitrary points in \mathbb{R}^3 , such that (i) no plane contains more than bl points of P , for some absolute constant $b \geq 1$, and (ii) each point of P is incident to at least three lines of L . Then*

$$I(P, L) = O\left(\min\{n^{1/3}l, n^{2/3}l^{2/3} + l\}\right).$$

Theorem 1.1.2 (Guth and Katz [50]). *Let L be a set of l lines in \mathbb{R}^3 and let P be a set of points in \mathbb{R}^3 . Suppose no more than $l^{1/2}$ lines of L lie in the same plane and suppose that each line of L contains at least $l^{1/2}$ points of P . Then $|P| = \Omega(l^{3/2})$.*

Theorem 1.1.3 (Guth and Katz [51]). *Let $k \geq 3$. Let L be a set of l lines in \mathbb{R}^3 with at most m lines in any plane. Let P be the set of points in \mathbb{R}^3 incident to at least k lines of L . Then the following inequality holds:*

$$|P| = O\left(\frac{l^{3/2}}{k^2} + \frac{lm}{k^3} + \frac{l}{k}\right).$$

Theorem 1.1.4 (Guth and Katz [51]). *Let L be a set of l lines in \mathbb{R}^3 with at most $O(l^{1/2})$ lines in any plane or regulus. Let P be the set of points in \mathbb{R}^3 incident to at least two lines of L . Then*

$$|P| = O(l^{3/2}).$$

All these bounds are derived using the algebraic method that Guth and Katz have introduced. Theorems 1.1.1 and 1.1.2 are related to the joints problem, while Theorems 1.1.3 and 1.1.4 are related to the distinct distances problem in the plane, which we review later on in Section 1.7.

1.2 Point-circle incidences

The best known upper bound on the number of incidences I between n points and c circles in the plane is

$$I = O(n^{6/11}c^{9/11} \log^{2/11}(n^3/c) + n^{2/3}c^{2/3} + n + c). \quad (1.2.1)$$

The polynomial factor $n^{6/11}c^{9/11}$ in the leading term was first derived by Aronov and Sharir [17], but with a worse subpolynomial factor. This factor was subsequently improved by Agarwal et al. [6], and then by Marcus and Tardos [63] to $\log^{2/11}(n^3/c)$. See Section 1.3 for more details about this bound. A consequence of (1.2.1) is the following useful variant (similar to a corresponding variant of the classical Szemerédi-Trotter theorem [80]), which will be exploited throughout the thesis: Given a set P of n points in the plane, the number c of distinct circles that contain at least k points of P (we refer to such circles as k -rich) is

$$c = O\left(\frac{n^3 \log k}{k^{11/2}} + \frac{n^2}{k^3} + \frac{n}{k}\right), \quad (1.2.2)$$

and the number I of incidences between the points of P and these circles is

$$I = O\left(\frac{n^3 \log k}{k^{9/2}} + \frac{n^2}{k^2} + n\right). \quad (1.2.3)$$

For a lower bound construction, consider n points and c lines with $\Omega(n^{2/3}c^{2/3} + n + c)$ incidences between them, and invert the plane about another point that is not incident to any of the lines. This gives us n points and c circles with $\Omega(n^{2/3}c^{2/3} + n + c)$ incidences. This is the best known lower bound construction for general n and c . However, this bound can not be tight, as for $c = \Theta^*(n^2)$, there exists a slightly better construction, as follows. Consider n points arranged in a $\sqrt{n} \times \sqrt{n}$ lattice section. As shown by Erdős [42], the number of distinct distances between these n points is $\Theta(n/\log^{1/2} n)$. Now consider all circles with center at one of the points, and containing some other point of this lattice section. By the aforementioned bound of Erdős, the number of such circles is $c = \Theta(n^2/\log^{1/2} n)$. Every pair of points generates two point-circle incidences, so the total number of point-circle incidences is $n(n-1) = \Theta(n^2)$. This number is strictly of a higher order than $n^{2/3}c^{2/3} + n + c = \Theta(n^2/\log^{1/3} n)$. Finding the right order of the maximum number of point-circle incidences in the plane is at present one of the major open incidence problems.

The number of point-circle incidences in space is also of interest. Aronov, Koltun, and Sharir [15] showed that the same upper bounds (1.2.1), (1.2.2), and (1.2.3) apply also to incidences between n points and c circles in \mathbb{R}^d for any $d \geq 3$ (resp., for the number of k -rich circles with respect to some n points in space, and for the number of incidences between these points and circles). This result should be compared to point-line incidences in space. In this case, as mentioned in Section 1.1, the number of point-line incidences in space can not exceed the number of point-line incidences in the plane. This statement is not known to hold for circles, that is, for all we know, it might be that the maximum number of point-circle incidences in \mathbb{R}^2 is strictly smaller than that of point-circle incidences in \mathbb{R}^d for $d \geq 3$. The fact that, at present, point-circle incidences in space have the same upper bound as in the planar case is a result of the proof technique in [15] rather than an inherent property of point-circle incidences (which may or may not exist — we do not know yet).

Another problem of interest is incidences between points and unit circles. This problem is strongly related to the unit distances problem of Erdős [42], which seeks a sharp bound on the maximum number of unit distances determined by pairs of points in any set of n points in the plane. At present, the best known upper bound, for n points and c circles, is $O(n^{2/3}c^{2/3} + n + c)$, but it is conjectured that the correct bound is almost linear; see Section 1.8 for more details.

1.3 Nonoverlapping lenses in arrangements of curves

The study of incidences between points and circles brings up a related topological problem of independent interest. This problem asks for the maximum possible number of nonoverlapping lenses in an arrangement of circles. A *lens* is the union of two circular arcs from two intersecting circles that have the same pair of endpoints. In an arrangement of circles, two lenses are said to be overlapping, if there exists a circular arc that is contained in two arcs from both lenses.

Tamaki and Tokuyama [81] observed that this number is of the same order as the number of cuts needed to turn a family of circles into *pseudo-segments*. A family of curves in the plane is called a family of *pseudo-segments* if any two members intersect at most once. By cutting the circles into sufficiently many arcs, one can ensure that the resulting arcs are pseudo-segments, but

the worst case minimal number of necessary cuts is an open problem.

It is known that the number of nonoverlapping lenses in an arrangement of c circles can be $\Omega(c^{4/3})$ (Tamaki and Tokuyama [81]). A family of circles with these many lenses can be constructed using a simple transformation of a configuration of c points and c lines that achieve the maximal possible number of $\Theta(c^{4/3})$ incidences; see [81]. The current upper bound stands at $O(c^{3/2} \log c)$ (Marcus and Tardos [63]), and is the latest in a series of improvements [6, 10, 17] on the initial bound of $O(c^{5/3})$ by Tamaki and Tokuyama [81].

The above mentioned $O(c^{3/2} \log c)$ bound implies the currently best known bound on point-circle incidences in the plane (given in (1.2.1)). If the logarithmic factor can be removed from the bound on the number of lenses, it will also lead to the removal of the logarithmic factor from the point-circle incidences bound. If it turns out that the number of nonoverlapping lenses is only $O(c^{4/3})$, it will imply that n points and c circles have at most $O(n^{1/2}c^{5/6} + n^{2/3}c^{2/3} + n + c)$ incidences. Since it is generally believed that the number of point-circle incidences is significantly smaller than that (in the range of n and c where the $n^{1/2}c^{5/6}$ term dominates), we probably need some new idea in order to find the right order of magnitude of point-circle incidences.

Generalizations of the problem of lens elimination in arrangements of general curves have been studied by Chan [26]. He showed that for a family of c curves in which any two members intersect at most $2s$ times (we call such a family $(2s)$ -intersecting), $O(c^{2-1/(2s+1)})$ cuts suffice to transform the curves into a family of $(2s - 1)$ -intersecting curve segments [26, remark after Lemma 7.1].

1.4 Other point-curve incidence problems in the plane

The same upper bound of $O(n^{6/11}c^{9/11} \log^{2/11}(n^3/c) + n^{2/3}c^{2/3} + n + c)$ on point-circle incidences [6, 17, 63], applies also for incidences between points and parabolas (of the form $y = ax^2 + bx + c$) [6]. The best lower bound construction of point-parabola incidences gives $\Omega(n^{1/2}c^{5/6} + n^{2/3}c^{2/3} + n + c)$ incidences between n points and c parabolas [37]. This bound is conjectured to be tight. A supporting evidence for this conjecture comes from the fact that if the upper bound on the number of nonoverlapping lenses in an arrangement of c parabolas in the plane turns out to be $O(c^{4/3})$, then the Aronov and Sharir analysis [17] (on point-circle incidences, but it can be adapted to parabolas) would give the matching upper bound $O(n^{1/2}c^{5/6} + n^{2/3}c^{2/3} + n + c)$.

Pach and Sharir [69] studied the following generalization of point-curve incidences. A family of curves \mathcal{C} is said to have k degrees of freedom and multiplicity type s , if for any tuple T of k points in the plane, there exist at most s curves from \mathcal{C} that pass through all the points of T , and any pair of curves from \mathcal{C} intersects in at most s points. Pach and Sharir showed:

Theorem 1.4.1 (Pach and Sharir [69]). *The number of incidences between n points and c simple curves with k degrees of freedom and multiplicity type s in the plane is*

$$O(n^{k/(2k-1)}c^{(2k-2)/(2k-1)} + n + c),$$

with the constant of proportionality depending on k and s .

On the lower bound side, Elekes [37] gave a construction for incidences between points and graphs of $(k - 1)$ -degree polynomials in the plane. His construction gives $\Theta(n^{2/(k+1)}c^{1-2/(k(k+1))})$ incidences between n points and c graphs of $(k - 1)$ -degree polynomials. Note that $(k - 1)$ -degree polynomials are a special case of curves with k degrees of freedom. This leaves a big gap between the lower and upper bounds for point-curve incidences with k degrees of freedom.

1.5 Point-hyperplane incidences

In higher dimensions, the simplest problem involves incidences between points and hyperplanes. A new technical issue that has to be faced is that, without imposing any restrictions on the points or hyperplanes, the number of incidences between n points and m hyperplanes can be mn or close to that. This happens when the intersection of all (or many of) the hyperplanes is a nonzero-dimensional affine subspace, and all (or many of) the points lie in that subspace. Several attempts can be (and have been) made to study this problem in more restricted settings, which we will shortly review. We will also review works that study the nonrestricted setting. In this case, the incidence graph $G(P, \Pi)$ can contain large complete bipartite subgraphs (as noted, in the worst case, it can be a complete bipartite graph). Brass and Knauer [22] use this fact to derive a succinct representation of the incidence graph as the disjoint union of complete bipartite subgraphs, whereas we [12] (coauthored with Sharir, see also Chapter 3) derive lower bounds on the size of the largest such complete bipartite subgraph.

1.5.1 Incidences between hyperplanes and vertices of their arrangement.

One way to avoid the situation mentioned above, where all the points are incident to all the hyperplanes, is to require that all the points be vertices of the arrangement of the hyperplanes. In this setting, for each pair of points p_1 and p_2 , there is at least one hyperplane incident to p_1 but not to p_2 , and vice versa.

Lower bound constructions are given in [34]. For n points and m hyperplanes in \mathbb{R}^d , if $n = O(m^{d-2})$, then the number of incidences can be as high as $\Omega(mn)$. If, however, $\Omega(m^{d-2}) \leq n \leq \binom{m}{d}$, then one can construct point-hyperplane configurations with $\Omega(n^{2/3}m^{d/3} + m^{d-1})$ incidences. Note that n can not exceed $\binom{m}{d}$. These bounds were later shown to be tight by Edelsbrunner et al. [35], and by Agarwal and Aronov [5]. Elekes and Tóth gave a simpler proof of the upper bound [41].

1.5.2 No $K_{r,s}$ incidence subgraph.

Another way to prevent situations where many points are incident to many hyperplanes (giving rise to the trivial bound of $I = \Omega(mn)$ incidences), is to require that in any such complete subconfiguration, there must not be too many (i.e., not more than a constant number of) points or too many hyperplanes.

Edelsbrunner et al. [35] showed that if the incidence graph of n points and m planes in \mathbb{R}^3 does not contain a $K_{3,2}$ subgraph (i.e. three points and two planes all incident to one another), then the

number of incidences is upper bounded by¹ $I = O(n^{3/5}m^{4/5} + n + m)$. By point-plane duality in \mathbb{R}^3 , the symmetric bound $O(n^{4/5}m^{3/5} + n + m)$ holds when the incidence graph does not contain a $K_{2,3}$ subgraph. Brass and Knauer [22] gave a matching lower bound construction showing that these bounds are worst case tight.

We note that Edelsbrunner et al. [35] were actually trying to study a more restricted case. They asked what is the maximum number of incidences if no three points are collinear. Since in that case, the incidence graph does not contain $K_{3,2}$, the upper bound $I = O(n^{3/5}m^{4/5} + n + m)$ follows. It is not known, however, whether this bound is tight for the case of no three collinear points. We do not know of any published lower bound constructions, and the best one we could think of gives $\Omega(n^{5/8}m^{3/4})$ incidences. This lower bound is constructed as follows. Let p be some prime number such that $p \equiv -1 \pmod{4}$, put $n = p^2$, and consider all points of the integer lattice that lie on the standard paraboloid modulo p , i.e., all $(x, y, z) \in \mathbb{Z}^3 \cap [0, p)^3$ such that $x^2 + y^2 \equiv z \pmod{p}$. There are exactly n such points. Call this set P . It is easy to see that no three points in P are collinear (indeed, three collinear points would give rise to three collinear points in \mathbb{Z}_p^3 that simultaneously satisfy a nontrivial quadratic relation — a contradiction; the requirement $p \equiv -1 \pmod{4}$ ensures that the quadratic relation in \mathbb{Z}_p^3 is indeed nontrivial). Next, let V be the set of all integer vectors (a, b, c) all of whose coordinates are in $[0, q]^3$, for some $q > 0$, such that $\gcd(a, b, c) = 1$. There are $\Theta(q^3)$ such vectors. For each point $r \in P$, construct all the planes through r that are orthogonal to the vectors of V . Let Π denote the resulting set of planes. Each such plane has the equation $ax + by + cz = e$, where $(a, b, c) \in V$. From the choices of P and V it follows that $e \leq 3pq$. Thus, the number of planes is at most $m \leq 3pq|V| = O(pq^4)$. The number of incidences is $I = n|V| = \Theta(p^2q^3)$. It is now easy to verify that $I = \Omega(n^{5/8}m^{3/4})$.

More generally, for n points and m hyperplanes in \mathbb{R}^d , if the incidence graph does not contain any $K_{r,s}$ subgraph, for some constants r and s , then the number of incidences is upper-bounded by $I = O((nm)^{d/(d+1)} + n + m)$, with the constant of proportionality depending on r and s . This was shown by Brass and Knauer [22] (with an extra logarithmic factor that can be removed). It is not known whether this bound is tight. Brass and Knauer [22] also provide some lower bound constructions. They showed that the number of incidences can be $I = \Omega^*((nm)^{1-2/(d+3)})$ if d is odd, and $I = \Omega^*((nm)^{1-2(d+1)/(d+2)^2})$ if d is even.

1.5.3 Incidences with nondegenerate hyperplanes.

The concept of *degeneracy* of a hyperplane was introduced by Elekes and Tóth in [41]. Given a finite point set $P \subset \mathbb{R}^d$ and a constant $0 < \eta < 1$, a hyperplane π in \mathbb{R}^d is said to be η -*degenerate* (with respect to P), if there exists some lower-dimensional affine subspace $\pi' \subset \pi$ such that

$$|\pi' \cap P| \geq \eta|\pi \cap P|.$$

If no such affine subspace π' exists, then π is said to be η -*nondegenerate*. As in an earlier definition, a hyperplane π is called k -*rich* (with respect to P) if $|\pi \cap P| \geq k$. Elekes and Tóth showed the following bound on the number of k -rich η -nondegenerate hyperplanes spanned by n points in \mathbb{R}^d .

¹In the original paper, this bound is multiplied by a subpolynomial factor of the form $n^\delta m^\delta$, for any $\delta > 0$. This factor, however, can be eliminated using a more refined analysis.

Theorem 1.5.1 (Elekes and Tóth [41]). *For any dimension $d \geq 3$ there exists a constant $\eta_d \leq 1$ which depends on d , so that, for any $\eta < \eta_d$ and for any set of n points in \mathbb{R}^d , the number h of η -nondegenerate k -rich hyperplanes is*

$$h = O\left(\frac{n^d}{k^{d+1}} + \frac{n^{d-1}}{k^{d-1}}\right), \quad (1.5.4)$$

with the multiplicative constant depending on η and d , and this bound is best possible. Equivalently, the number I of incidences between n points and h η -nondegenerate hyperplanes in \mathbb{R}^d is

$$I = O\left(n^{\frac{d}{d+1}} h^{\frac{d}{d+1}} + nh^{\frac{d-2}{d-1}}\right), \quad (1.5.5)$$

with the multiplicative constant depending on η and d , and this bound is best possible.

The equivalence between (1.5.4) and (1.5.5) can be established by a standard argument, going all the way back to the Szemerédi-Trotter incidence bound [80].

The currently known thresholds are $\eta_3 = 1$, $\eta_4 = 1/2$, and as d increases, η_d gets smaller and smaller. It is not known what the true values of these thresholds are, or even whether there is a need for a threshold at all. For what we know, it may be that the bound applies to all $0 < \eta < 1$ in any dimension. This theorem can be thought of as a generalization of the Szemerédi-Trotter bound for point-line incidences [80] (see Section 1.1). Indeed, if we substitute $d = 2$ in (1.5.4), we get exactly the Szemerédi-Trotter bound. Of course, for lines in the plane, there is no such thing as degeneracy, so the Elekes-Tóth bound is trivial (that is, equivalent to the Szemerédi-Trotter bound) for this case.

In Chapter 5 we extend the notion of degeneracy to spheres, and derive an upper bound on incidences between points and nondegenerate spheres in \mathbb{R}^3 (see also [14] coauthored with Sharir).

1.5.4 Representation complexity of point-hyperplane incidence graphs.

Brass and Knauer [22] considered the general case, where the incidence graph $G(P, \Pi)$ can contain large complete bipartite subgraphs. Rather than bounding $I(P, \Pi)$ itself, they have obtained an upper bound for the overall minimum possible complexity of a representation of $G(P, \Pi)$ as the disjoint union of complete bipartite graphs, that is, $G(P, \Pi) = \bigcup_{i=1}^s A_i \times B_i$, where $A_i \subseteq P$ and $B_i \subseteq \Pi$, for all $i = 1, \dots, s$, and each incidence is recorded exactly once in this union. The *complexity* of such a representation of $G(P, \Pi)$ is defined to be $\sum_{i=1}^s (|A_i| + |B_i|)$, and the smallest complexity of such a representation, or the *representation complexity* of $G(P, \Pi)$, is denoted by $J(P, \Pi)$. We let $J_d(m, n)$ denote the maximum of $J(P, \Pi)$ over all sets P of m points, and Π of n hyperplanes in \mathbb{R}^d . Brass and Knauer [22] have shown that

$$J_d(m, n) = O\left(\left((mn)^{1-\frac{1}{d+1}} + m + n\right) \log(m + n)\right). \quad (1.5.6)$$

One way to interpret (1.5.6) is that if the number of incidences $I(P, \Pi)$ is much larger than $J_d(m, n)$, then $G(P, \Pi)$ should contain large complete bipartite subgraphs (or else the succinct representation would not be possible). This has been one of our main motivations to study how large must these complete bipartite subgraphs be. In Chapter 3 we present the results of that study (see also [12] coauthored with Sharir).

1.6 Other point-surface incidence results

The new algebraic methods, as introduced by Guth and Katz [50, 51], have, very recently, started to yield new results on point-surface incidences in three and higher dimensions. These results are still in the making, so our review may not be up to date. However, we would like to note some of them, as they are strongly related to the results achieved and discussed in this thesis.

Point-sphere incidences have been studied earlier, mostly within works on repeated and distinct distances [16, 31], and repeated simplices [4, 7]. Clarkson et al. [31] have shown that m points and n unit spheres in \mathbb{R}^3 have $O(m^{3/4}n^{3/4}\beta(m, n) + m + n)$ incidences, where $\beta(m, n)$ is an extremely slowly growing function related to the inverse Ackermann function. Very recently, and independently, Kaplan et al. [53] and Zahl [84] have improved this bound to $O(m^{3/4}n^{3/4} + m + n)$ using the polynomial partitioning scheme discussed in Section 2.1.3 below. This result implies the upper bound in (1.8.8) on repeated distances in \mathbb{R}^3 , which in itself is also an improvement of a previous result of Clarkson et al. [31].

Zahl's work [84] has actually studied the more general context of incidences between points and algebraic surfaces of bounded degree in \mathbb{R}^3 . We quote his result in detail in Chapter 4; see Theorem 4.2.4. This result also implies an improved upper bound on the number of unit area triangles in the plane; see Section 1.9.2, and Chapter 4.

Solymosi and Tao have derived the following generalization of the Szemerédi-Trotter theorem to general dimensions.

Theorem 1.6.1 (Solymosi and Tao [73]). *Let $d \geq 2k$, and let P and L be respectively, a set of n points and a family of l k -flats in \mathbb{R}^d , such that any two k -flats in L intersect in at most one point. Then, for any $\varepsilon > 0$, we have*

$$I(P, L) = O(n^{2/3+\varepsilon}l^{2/3} + n + l),$$

with the constant of proportionality depending on k and ε .

1.7 Distinct distances

One of the basic open incidence problems has been the planar distinct distances problem, posed by Erdős in 1946 [42]: Let P be a set of n points in the Euclidean plane. Denote by $g(P)$ the number of distinct distances determined by pairs of points of P , and let $g_2(n)$ denote the minimum of $g(P)$ taken over all sets P of n points in \mathbb{R}^2 . The problem is to obtain sharp upper and lower bounds on $g_2(n)$. By considering points in a $n^{1/2} \times n^{1/2}$ lattice section, one can obtain only $O(n/\log^{1/2} n)$ distinct distances (a bound already noted by Erdős [42]), and so $g_2(n) = O(n/\log^{1/2} n)$.

On the other hand, lower bounds are more difficult to obtain. As observed by Erdős [42], an easy argument, ultimately based on the fact that any two circles intersect in at most two points, gives the lower bound $g_2(n) = \Omega(n^{1/2})$. The exponent $1/2$ has been slowly increasing over the years by a series of increasingly intricate arguments. However, these methods seemed to fall quite short of getting to the optimal exponent of 1. Indeed, until quite recently, the best lower bound known was approximately $\Omega(n^{0.8641})$, due to Katz and Tardos [56].

Very recently, though, Guth and Katz [51] have obtained a near-optimal result:

$$g_2(n) = \Omega(n/\log n).$$

The proof neatly combines several powerful and modern tools in a new way: a recent geometric reformulation of the problem due to Elekes and Sharir [40]; the polynomial method as already been used previously by Guth and Katz to solve the joints problem [50]; and the somewhat older method of cell decomposition (discussed in Section 2.1). A key new insight is that the polynomial method (and more specifically, the polynomial Ham Sandwich theorem) can be used to efficiently create cell decompositions (see Section 2.1.3).

In higher dimensions there is still a gap between the known upper and lower bounds. We denote by $g_d(n)$ the minimal number of distinct distances determined by any set of n points in \mathbb{R}^d . The d -dimensional $n^{1/d} \times \dots \times n^{1/d}$ lattice section contains n points and determines $O(n^{2/d})$ distinct distances, for $d \geq 3$. This is the best upper bound construction known, and Erdős conjectured it to be tight [42]. On the other hand, the current lower bound, based on a recurrence relation by Solymosi and Vu [74] gives, in three dimensions, $g_3(n) = \Omega^*(n^{3/5})$, and more generally, in d dimensions, $g_d(n) = \Omega(n^{2/d-2/(d(d+2))})$.

1.8 Repeated distances

Let $u_d(n)$ denote the maximum number of unit distances between pairs of points of some set of n points in the Euclidean d -dimensional space. For $d = 2$ and $d = 3$, a sharp bound for the value of $u_d(n)$ is not known. We know the bounds:

For $d = 2$, we have

$$n^{1+\frac{c}{\log \log n}} \leq u_2(n) \leq Cn^{4/3}, \quad (1.8.7)$$

for some constants c and C , with the lower bound shown by Erdős [42] and the upper bound by Spencer et al. [75].

For $d = 3$, we have

$$cn^{4/3} \log \log n \leq u_3(n) \leq Cn^{3/2}, \quad (1.8.8)$$

where c and C are some constants. The lower bound is due to Erdős [43] and the upper bound is due to Kaplan et al. [53], and Zahl [84].²

As for $d \geq 4$, put $p = \lfloor d/2 \rfloor$, and consider p mutually orthogonal circles of radius $1/\sqrt{2}$ centered at the origin, each containing about n/p points. This construction, attributed to Lenz [62], gives about $\frac{(p-1)n^2}{2p}$ unit distances. Swanepoel [77] showed that this construction, up to some minor tweaks, is best possible, and thus determined the exact value of $u_d(n)$ for d even, and almost exactly

²The upper bound is a very recent improvement over the earlier bound $n^{3/2}2^{O(\alpha(n)^2)}$ due to Clarkson et al. [31], where $\alpha(\cdot)$ is the extremely slowly growing inverse Ackermann function.

for odd values of d . This value, for even d , is

$$u_d(n) = \frac{(p-1)n^2}{2p} + n - O(1),^3$$

and for odd d is

$$u_d(n) = \frac{(p-1)n^2}{2p} + \Theta(n^{4/3}).$$

Other variants of the problem, mostly two-dimensional ones, have also been studied. Valtr [83] studied the planar case under non-Euclidean norms, and constructed a simple strictly convex norm such that the number of unit distances in the plane can be $\Theta(n^{4/3})$. We note that for norms that are not strictly convex, like L_1 and L_∞ , the number of unit distances can be $\Theta(n^2)$. On the other hand, the upper bound on $u_2(n)$ in (1.8.7) applies to all strictly convex norms, as the known proofs carry over almost verbatim, so Valtr's construction is worst-case optimal. An important conclusion of this result is that if the conjecture $u_2(n) = o(n^{4/3})$ is true, then one cannot derive the exact order of $u_2(n)$ just from the topology of the problem, and one has to use the special properties of the Euclidean norm, and of circles in particular.

If the n points are in convex position in the plane, then the number of unit distances is at most $O(n \log n)$ [24, 49], and can be $2n - 7$ [36].

Ábrego and Fernández-Merchant [3] studied the case where the points are in convex position and symmetric about the origin. They showed that the maximal number of unit distances in this case is $2n - e$, where e is a lower order term for which they provide the bounds $3 \leq e \leq O(\sqrt{n})$.

Erdős et al. [45] studied unit distances on a sphere. The planar upper bound $O(n^{4/3})$ applies for this variant just as well (where the same proof carries over verbatim). They showed that a simple transformation on a configuration of n points and n lines with $\Theta(n^{4/3})$ incidences gives a configuration of n points on a sphere of radius $1/\sqrt{2}$ with $\Theta(n^{4/3})$ unit distances. For spheres of other radii they constructed a slightly superlinear number of unit distances. This construction was later improved by Swanepoel and Valtr [78] who constructed $\Theta(n\sqrt{\log n})$ unit distances on a sphere of any radius greater than $\frac{1}{2}$.

1.9 Repeated triangles and angles

1.9.1 Repeated angles

Let P be a set of n points in the plane, and consider the angles spanned by triples of points of P . The number of times the same angle α can occur (i.e., the number of distinct triples that span the same angle) is $O(n^2 \log n)$, and this bound is tight for all angles α such that $\tan \alpha = \sqrt{a/b}$ for some integers a and b (Pach and Sharir [68]). It is not known what is the tight bound for other angles, but any angle α can be repeated $(n-1)^2/4$ times simply by evenly distributing the points on two rays that form an angle α , and placing one point at the apex.

³The $O(1)$ term is also known exactly; see [77] for details.

In three dimensions, Apfelbaum and Sharir [11] showed that an angle α can be repeated at most $O(n^{7/3})$ times, and that this bound is tight for the right angle $\alpha = \pi/2$. It is not known what is the tight bound for other angles.

In four dimensions, the right angle can be repeated $\Theta(n^3)$ times. For other angles, Apfelbaum and Sharir [11] showed the bound $O(n^{5/2}C^{\alpha(n)^2})$, where $C > 0$ is some constant.⁴

In five dimensions, we do not know of any nontrivial bounds, and in six dimensions and above, any acute angle can be repeated $\Theta(n^3)$ times using the Lenz construction (construct three mutually orthogonal circles of appropriate radii and place $n/3$ points on each of them).

1.9.2 Repeated areas

Let P be a set of n points in the plane, and consider the areas of triangles spanned by triples of points of P . Erdős and Purdy [46] showed that a lattice section of size $n/\log^{1/2} n \times \log^{1/2} n$ spans $\Omega(n^2 \log \log n)$ triangles of the same area. No better construction is known. The current upper bound is $O(n^{9/4})$ (Apfelbaum and Sharir [13]; we derive this bound in detail in Chapter 4).

In three dimensions, Dumitrescu et al. [33] showed an $O^*(n^{17/7}) \approx O^*(n^{2.4286})$ upper bound.

1.10 Repeated simplices

Let P be a set of n points in \mathbb{R}^d , and let Δ be a prescribed k -simplex, for some $0 \leq k \leq d$ (we consider $k = 0$ and $k = 1$ as degenerate cases; a 0-simplex is a single point and a 1-simplex is a line segment). Let $G(P, \Delta)$ denote the number of k -simplices spanned by P that are congruent to Δ . Set

$$G_{k,d}(n) = \max G(P, \Delta),$$

where the maximum is taken over all sets P of n points in \mathbb{R}^d and over all k -simplices Δ in \mathbb{R}^d . It is an open problem to determine the order of $G_{k,d}(n)$, for general k , d , and n , and in this subsection we review the known results.

First, we note that $G_{1,d}(n) = u_d(n)$, where $u_d(n)$ is defined in Section 1.8 as the maximum number of repeated distances in \mathbb{R}^d . That is, repeated congruent one-dimensional simplices, are simply repeated distances; see Section 1.8 for an overview of the known results on that subject. Another simple observation is that $G_{d,d}(n) \leq 2G_{d-1,d}(n)$. To see that, consider a $(d-1)$ -face Δ' of the repeated d -simplex Δ , and observe that there can be only two points p in \mathbb{R}^d , that together with Δ' make a d -simplex congruent to Δ .⁵ Thus, for example, the number of repeated congruent triangles spanned by some n points in the plane is upper bounded by the number of repeated distances in the plane, with bounds given in (1.8.7).

⁴The new works of Kaplan et al. [53], and Zahl [84] (see Section 1.6) imply an improved bound of $O(n^{5/2})$.

⁵This statement is true if Δ has no symmetries, but if it has symmetries, then there may be more than two points completing the $(d-1)$ -simplex. To handle symmetries, we assume that the vertices of Δ are ordered, and only consider similar simplices with the same vertex order. This assumption is implicit throughout the discussion on simplices here and in Chapter 6, and we keep it hidden for better readability.

Agarwal and Sharir [7] studied this problem for general k and d and showed the following bounds: $\Omega(n^{4/3}) \leq G_{2,3}(n) \leq O^*(n^{5/3})$, $\Omega(n^2) \leq G_{2,4}(n) \leq O^*(n^2)$, $G_{2,5}(n) = \Theta(n^{7/3})$, and $G_{3,4}(n) = O^*(n^{9/4})$. A simple construction, attributed to Lenz [62] (a variant of the one described in Section 1.8), shows that $G_{k,d}(n) = \Omega(n^{d/2})$ for even values of $d \geq 4$ and $k \geq d/2 - 1$, and $G_{k,d}(n) = \Omega(n^{d/2-1/6})$ for odd values of $d \geq 3$ and $k \geq (d-1)/2$. Erdős and Purdy [47] conjectured that this construction is asymptotically best possible for d even, and that, in general, $G_{k,d}(n) = O(n^{d/2})$. Agarwal and Sharir [7] derived a recurrence for $G_{k,d}(n)$, for general values of k and d . The solution of this recurrence is $O(n^{\zeta(d,k)+\epsilon})$, where $\zeta(d,k)$ is a rather complicated function of d and k . They show, though, that $\zeta(d,k) \leq d/2$ for $d \leq 7$ and $k \leq d-2$, and conjecture that $\zeta(d,k) \leq d/2$ for all d and $k \leq d-2$, in accordance with the Erdős-Purdy conjecture just mentioned.

The number of repeated *similar* simplices is also of interest. We study this problem in Chapter 6. Similarly to the above notation $G_{k,d}(n)$, we denote by $F_{k,d}(n)$ the maximal possible number of mutually similar simplices spanned by n points in \mathbb{R}^d . A $n^{1/2} \times n^{1/2}$ square lattice section generates $\Theta(n^2)$ isosceles right triangles, which shows that $F_{2,2}(n) = \Theta(n^2)$. Akutsu et al. [9] showed that $F_{2,3}(n) = O(n^{2.2})$. We improve this bound in Chapter 6 to $F_{2,3}(n) = O^*(n^{58/27})$ (see also [4] coauthored with Agarwal, Purdy, and Sharir). Other than that, there are practically no known upper bounds in $d \geq 3$ dimensions. Brass [21] conjectures, though, that $F_{3,3}(n) = o(n^2)$; the best known lower bound is $F_{3,3}(n) = \Omega(n^{4/3})$.

1.11 Summary of the results of the thesis

In this section we summarize the results obtained in this thesis.

Point-hyperplane incidences. In Chapter 3 we show that if the number I of incidences between m points and n planes in \mathbb{R}^3 is sufficiently large, then the incidence graph contains a large complete bipartite subgraph involving r points and s planes, so that $rs \geq \frac{I^2}{mn} - a(m+n)$, for some constant $a > 0$. This is shown to be almost tight in the worst case because there are examples of arbitrarily large sets of points and planes where the largest complete bipartite incidence subgraph records only $\frac{I^2}{mn} - \frac{m+n}{16}$ incidences. We also make some steps towards generalizing this result to higher dimensions.

The results in Chapter 3 were published, coauthored with Sharir, in *SIAM Journal on Discrete Mathematics* [12].

Unit area triangles. In Chapter 4 we show that the number of unit-area triangles determined by a set of n points in the plane is $O(n^{9/4})$, improving upon the bound $O(n^{44/19})$ of Dumitrescu et al. [33].

The results of Chapter 4 were published, coauthored with Sharir, in *Discrete and Computational Geometry* [13], but with a slightly weaker bound of $O^*(n^{9/4})$. A very recent result of Zahl [84] has enabled us to get rid of the subpolynomial factor.

Nondegenerate spheres. In Chapter 5 we show that the number of k -rich η -nondegenerate spheres spanned by a set of n points in \mathbb{R}^3 is at most $O^*\left(\frac{n^4}{k^{11/2}} + \frac{n^2}{k^2}\right)$, with the constant of proportionality depending on η .

The results of Chapter 5 were published, coauthored with Sharir, in *Combinatorics, Probability and Computing* [14].

Repeated simplices in d dimensions. In Chapter 6, we consider the problem of bounding the maximum possible number $F_{k,d}(n)$ of k -simplices that are spanned by a set of n points in \mathbb{R}^d and are similar to a given simplex. We first show that $F_{2,3}(n) = O^*(n^{58/27})$, and then tackle the general case, and show that $F_{d-2,d}(n) = O(n^{d-2+o(1)})$ and $F_{d-1,d}(n) = O(n^{d-2+o(1)})$, for any d , where the $o(1)$ terms both diminish exponentially in d . Our technique extends to derive bounds for other values of k and d , and we illustrate this by showing that $F_{2,5}(n) = O(n^{8/3})$.

A preliminary version of these results, with weaker upper bounds, appeared, coauthored with Agarwal, Purdy, and Sharir, at the *23rd Annual Symposium on Computational Geometry (2007)* [4]. A new version is being completed and will be submitted for publication.

Chapter 2

Techniques

2.1 Space decomposition

Geometric divide and conquer methods are a very useful tool with many algorithmic, as well as combinatorial applications. Among other things, they are very effective in analyzing incidence problems. The analysis proceeds by decomposing space into cells, estimating the number of incidences in each cell separately, and summing up the results over all cells. Because of their importance, those methods were studied extensively. There have been numerous works aimed at constructing various space decompositions, studying their properties, and establishing their usefulness. We shall give a superficial description of some of the available constructions, and show an example or two of how they can be used to derive incidence bounds.

2.1.1 Cuttings

We start with the following construction, known as a *cutting*, described first for a simple instance involving a family L of n lines in the plane. Let R be a random sample of L , constructed by taking each line of L into the sample independently at random with probability r/n (where $1 < r < n$ is some specified parameter). The expected number of lines in R is r . In the arrangement $\mathcal{A}(R)$ of the lines of R , every edge is intersected by an expected number of n/r lines of L . Now construct the vertical decomposition of $\mathcal{A}(R)$, and get a decomposition of the plane into an expected number of $O(r^2)$ trapezoidal cells. It will usually be convenient for us to work with pairwise disjoint cells, so we define the 2-dimensional cells of the decomposition as the relative interiors of the trapezoids, and add the relatively open edges and vertices of the cells as cells of the decomposition, of dimensions 1 and 0, respectively. The additional vertical edges are also intersected by an expected number of $O(n/r)$ lines of L . In fact, any straight line segment that does not intersect the lines of R , intersects an expected number of $O(n/r)$ lines of L , and with high probability, no such segment intersects more than $O(n \log r/r)$ lines of L . This follows from the theory of ε -nets in range spaces of bounded VC-dimension [52]. Now because every cell has at most four edges, we have, with high probability, that every cell is crossed by at most $O(n \log r/r)$ lines of L , and on average, by $O(n/r)$ lines. One should keep in mind that this is a probabilistic construction. Thus, it is not

necessarily the case that we get a decomposition with at most $O(n \log r/r)$ crossings for every cell, or with no more than $O(r^2)$ cells. However, because we have a nonzero probability of getting this result, this serves as a proof of existence, and we may assume to be given a good sample that results in a good decomposition. For more details, see [27, 30, 31].

This construction can be generalized to arrangements of hyperplanes in \mathbb{R}^d as follows. We use the same random sampling process, and then construct the arrangement of the sample hyperplanes. This gives us an expected number of $O(r^d)$ cells, where, with high probability, each straight line segment entirely contained in a cell is *crossed* by (i.e., intersected by but not contained in) an expected number of n/r hyperplanes, and at most $O(n \log r/r)$ hyperplanes. Next, we need to cut the cells into subcells of constant description complexity (the complexity of a polyhedral cell is defined as the total number of its faces of all dimensions). In this case, however, we shall not, in general, use vertical decomposition, as it does not guarantee that the (expected) number of cells remains $O(r^d)$. Instead, one can use *bottom-vertex decomposition* [29], which does keep the number of cells $O(r^d)$, while guaranteeing both constant complexity of the cells (which are simplices) and, consequently, at most $O(n/r)$ expected number of crossings per cell. We thus get a decomposition of space into an expected number of $O(r^d)$ cells, each of which is crossed by an expected number of $O(n/r)$ hyperplanes, and, with high probability, by at most $O(n \log r/r)$ hyperplanes.

By using bottom-vertex instead of vertical decomposition, we achieve a decomposition with the desired properties, but bottom-vertex decomposition has its drawbacks. The advantage of vertical over bottom-vertex decomposition, is that vertical decomposition can be perturbed by a change of the coordinate system, while bottom-vertex decomposition cannot. As a result, one can make certain general-position assumptions (which greatly simplify the analysis) on vertical decomposition, that cannot be made under bottom-vertex decomposition. For example, we can assume that the additional vertical walls do not pass through any point of some given finite set. Because of that, it is preferable to use vertical decomposition whenever possible. Vertical decomposition can be used in \mathbb{R}^d for $d = 3, 4$, where it is known to produce $O(r^d)$ cells as desired. The result for \mathbb{R}^4 is due to Koltun [59].

Some notations. We typically denote space decompositions in \mathbb{R}^d by $\Xi = \{\tau_1, \dots, \tau_t\}$, where τ_i , for $i = 1, \dots, t$ are the cells. These cells are relatively open connected subsets, of constant description complexity, and of dimensions $0, \dots, d$. The cells are pairwise disjoint, and their union equals the whole of \mathbb{R}^d . A hyperplane π is said to *cross* a cell τ if π intersects τ but does not contain it. Let Π be a family of n hyperplanes in \mathbb{R}^d . We say that a space decomposition Ξ is a $\frac{1}{r}$ -*cutting* of Π , if every cell $\tau \in \Xi$ is crossed by at most n/r hyperplanes of Π . The above construction (of a vertical or bottom-vertex decomposition of a random sample of the hyperplanes with individual selection probability r/n) is, with high probability, a $O(\frac{\log r}{r})$ -cutting with $O(r^d)$ cells. To achieve a $\frac{1}{r}$ -cutting, one has to sample with larger probability $O(r \log r/n)$, but then the number of cells will be larger, $O((r \log r)^d)$. We can summarize this result in the following theorem.

Proposition 2.1.1. *Let Π be a family of n hyperplanes in \mathbb{R}^d , and let $1 < r < n$ be a given parameter. Then there exists a $\frac{1}{r}$ -cutting of Π with $O((r \log r)^d)$ cells.*

This construction is already quite useful. For example, Clarkson et al. [31] used it to derive many incidence bounds, including a simpler proof of the Szemerédi-Trotter bound than the one originally given by Szemerédi and Trotter [80]. However, the $\log r$ factor makes it sub-optimal. One can, in fact, achieve a $\frac{1}{r}$ -cutting with only $O(r^d)$ cells, as was subsequently shown by Chazelle and Friedman [29]. This is done using the random sampling process with individual selection probability r/n as an initial decomposition, and then splitting any cell with too many crossings by a second level of cutting with an appropriately chosen sampling probability. By a careful choice of parameters, this will, with high probability, yield an expected number of $O(r^d)$ cells, each of which is crossed by at most n/r hyperplanes. Since a cutting with this number of cells and crossings occurs with nonzero probability, it follows that a $\frac{1}{r}$ -cutting with $O(r^d)$ cells exists. We thus have,

Theorem 2.1.2. *Let Π be a family of n hyperplanes in \mathbb{R}^d , and let $1 < r < n$ be a given parameter. Then there exists a $\frac{1}{r}$ -cutting of Π with $O(r^d)$ cells.*

We next show how to use cuttings to prove the Szemerédi-Trotter incidence bound. For an initial weak bound, we use the bound of Kővari, Sós, and Turán in extremal graph theory:

Lemma 2.1.3 (Kővari, Sós, and Turán [57]). *Let $G = (M \cup N, E \subseteq M \times N)$ be a bipartite graph with $|M| = m$ vertices on one side, $|N| = n$ vertices on the other side, and $|E| = e$ edges. Assume further that G does not contain a $K_{r,s}$, i.e., a complete bipartite subgraph with r vertices from M and s vertices from N . Then the number of edges is*

$$e \leq (s-1)^{1/r} mn^{1-1/r} + (r-1)n,$$

and, symmetrically,

$$e \leq (r-1)^{1/s} nm^{1-1/s} + (s-1)m.$$

Theorem 2.1.4 (Szemerédi and Trotter [80], this proof – Clarkson et al. [31]; see also Székely [79] and Section 2.2 for yet another, simpler proof). *Let P be a set of n points and let L be a family of l lines in the plane. Then the number I of incidences between P and L satisfies*

$$I = O(n^{2/3}l^{2/3} + n + l). \quad (2.1.1)$$

Equivalently, given a set P of n points in the plane, and a number $k \leq n$, the number l of lines that contain at least k points of P (k -rich lines) is

$$l = O\left(\frac{n^2}{k^3} + \frac{n}{k}\right), \quad (2.1.2)$$

and the number of incidences between these lines and the points of P is

$$I = O\left(\frac{n^2}{k^2} + n\right). \quad (2.1.3)$$

Proof. We shall prove only the first inequality (2.1.1). The equivalent formulations (2.1.2) and (2.1.3) easily follow, and are noted here only for completeness and cross reference.

First, note that the incidence graph does not contain a $K_{2,2}$ as a subgraph. Thus, by Lemma 2.1.3, the number of edges in the incidence graph, which is also the number of incidences, is bounded by

$$I \leq nl^{1/2} + l, \quad (2.1.4)$$

and, symmetrically,

$$I \leq ln^{1/2} + n. \quad (2.1.5)$$

Thus, if $n > l^2$, then $I = O(n)$, and if $l > n^2$, then $I = O(l)$. It remains to handle the case $l^{1/2} \leq n \leq l^2$.

Let $r = n^{2/3}/l^{1/3}$, and note that $1 \leq r \leq l$ in the assumed range of n . Let $\Xi = \{\tau_1, \dots, \tau_t\}$ be a $\frac{1}{r}$ -cutting of L , with $t = O(r^2)$. For each $i = 1, \dots, t$, we denote by $P_i = \tau_i \cap P$ the subset of points of P in τ_i , and by L_i — the subset of lines of L that cross τ_i . By construction we have $|L_i| \leq \frac{l}{r}$ for all $i = 1, \dots, t$. Summing the number of incidences, over all cells of the cutting, between the points and crossing lines of each cell, and denoting this sum by I_1 , we get

$$I_1 = \sum_{i=1}^t I(P_i, L_i) \leq \sum_{i=1}^t (|P_i||L_i|^{1/2} + |L_i|) \leq \left(\frac{l}{r}\right)^{1/2} \sum_{i=1}^t |P_i| + \sum_{i=1}^t |L_i| = O\left(\frac{nl^{1/2}}{r^{1/2}} + lr\right).$$

From the choice of r , it follows that

$$I_1 = O(n^{2/3}l^{2/3}).$$

It remains to bound the incidences $p \in \ell$, where $p \in P \cap \tau$ for some cell $\tau \in \Xi$, and $\ell \in L$ is a line containing τ . In this case, τ is either one-dimensional, i.e., a line segment, or zero dimensional, i.e., a point. In the former case, any point in a one-dimensional cell can generate just one such incidence, so the number of incidences is at most n . In the latter case, we use (2.1.5) with l lines and $O(r^2)$ vertices of the cutting to conclude that the number of incidences of this kind is $O(lr + r^2) = O(lr) = O(n^{2/3}l^{2/3})$. Summing the bounds for all cases, we get the asserted bound

$$I = O(n^{2/3}l^{2/3} + n + l).$$

□

Here is another important example. It proves the point-circle incidence bound (1.2.1), and also demonstrates how sometimes, it is useful to do the cutting in dual space.

Theorem 2.1.5 (See [6, 17, 63]). *The number I of incidences between n points and c circles in the plane is*

$$I = O(n^{6/11}c^{9/11} \log^{2/11}(n^3/c) + n^{2/3}c^{2/3} + n + c).$$

Proof (based on Aronov and Sharir [17]). As an initial weak bound, we shall use the bound of Lemma 2.2.3 given below,

$$I = O(n^{2/3}c^{2/3} + n + c^{3/2} \log c).$$

Note that if $n > c^{5/4} \log^{3/2} c$, then the $n^{2/3} c^{2/3} + n$ terms dominate, and this bound is optimal (because one can always construct $\Omega(n^{2/3} c^{2/3} + n + c)$ incidences). We thus need to consider only the case $n < c^{5/4} \log^{3/2} c$. Also, since the incidence graph does not contain a $K_{3,2}$ (three points and two circles all incident to one another), we have, by Lemma 2.1.3, that $I = O(nc^{2/3} + c)$, from which it follows that if $n < c^{1/3}$, then $I = O(c)$, which is optimal. Thus, it remains to consider only the case $c^{1/3} < n < c^{5/4} \log^{3/2} c$.

Let P and Γ be, respectively, a set of n points and a family of c circles in the plane, such that $c^{1/3} < n < c^{5/4} \log^{3/2} c$, and let $I = I(P, \Gamma)$. We now map the setting into dual 3-space, where we first lift each point to the standard paraboloid, so that each circle becomes a plane, and then apply duality in \mathbb{R}^3 . The resulting transformation maps the circles to points in \mathbb{R}^3 and the points of P to planes, and it preserves incidences, so that a point $p \in P$ lies on a circle $\gamma \in \Gamma$ if and only if the dual point γ^* lies on the dual plane p^* .

It is at this point that we apply a cutting. Let r be some parameter to be specified later, and let Ξ be a $\frac{1}{r}$ -cutting of the arrangement of the planes of P^* . Ξ consists of $O(r^3)$ relatively open polyhedral cells of dimensions 3,2,1,0, each of constant complexity, such that each cell is crossed by at most n/r planes of P^* . We may further assume that each cell contains at most $O(c/r^3)$ points of Γ^* ; otherwise we split a cell into subcells with this property, and note that the number of cells remains $O(r^3)$. By Lemma 2.2.3, each cell has at most

$$O\left(\left(\frac{n}{r}\right)^{2/3} \left(\frac{c}{r^3}\right)^{2/3} + \frac{n}{r} + \left(\frac{c}{r^3}\right)^{3/2} \log(c/r^3)\right) = O\left(\frac{n^{2/3} c^{2/3}}{r^{8/3}} + \frac{n}{r} + \frac{c^{3/2}}{r^{9/2}} \log(c/r^3)\right)$$

incidences of the points that it contains with the planes that cross it. We multiply by $O(r^3)$ to get the total number I_1 of such incidences over all the cells,

$$I_1 = O\left(n^{2/3} c^{2/3} r^{1/3} + nr^2 + \frac{c^{3/2}}{r^{3/2}} \log(c/r^3)\right).$$

Now we substitute $r = c^{5/11} n^{-4/11} \log^{6/11}(n^3/c)$. Note that by our assumption that $c^{1/3} < n < c^{5/4} \log^{3/2} c$, it follows that $\Omega(1) \leq r \leq O(n)$, so, for appropriate constants of proportionality, the choice of r is acceptable. This gives us

$$I_1 = O\left(n^{6/11} c^{9/11} \log^{2/11}(n^3/c) + n^{3/11} c^{10/11} \log^{12/11}(n^3/c)\right) = O\left(n^{6/11} c^{9/11} \log^{2/11}(n^3/c)\right),$$

where the $O(n^{6/11} c^{9/11})$ term is dominant by the assumption that $n > c^{1/3}$.

It remains to bound the number of incidences, for each cell τ of dimension 0, 1, or 2, between the points of Γ^* in τ , and the planes of P^* that contain τ . Each two-dimensional cell is contained in at most one plane, so the number of incidences is at most the number of points in that cell, which adds up to a total of $O(c)$ over all two-dimensional cells. For one-dimensional cells, by a similar argument, we have a total of $O(n + c)$ incidences. Indeed, a one-dimensional cell either contains at most two points (which map back to at most two circles of Γ), or else, contains three or more collinear points (which map back to a pencil of circles of Γ , see Figure 2.1), but then, this cell is contained by at most two planes (which map back to two point of P), and no other point can be

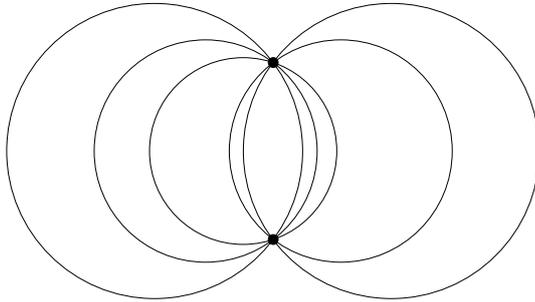


Figure 2.1: A pencil of circles through a pair of points.

incident to all these circles. As for the zero-dimensional cells, which are simply the vertices of the cutting, we have at most $O(r^3)$ such vertices, which map back to $O(r^3)$ circles in the plane. The number of incidences between these circles and the points of P can be bounded, using Lemma 2.1.3, by $O(n(r^3)^{2/3} + r^3) = O(nr^2 + r^3)$. One can easily verify that, for the assumed range $c^{1/3} < n < c^{5/4} \log^{3/2} c$, this bound is dominated by $O(n^{6/11} c^{9/11} \log^{2/11}(n^3/c) + c)$. Summing all bounds together, we get

$$I = O\left(n^{6/11} c^{9/11} \log^{2/11}(n^3/c) + n^{2/3} c^{2/3} + n + c\right),$$

as claimed. □

As a final note about cuttings, we remark that the same construction (of an arrangement of a random sample, further partitioned into cells of constant complexity) is applicable also to algebraic curves in the plane, and surfaces in higher dimensions. The main new technical problem is the need to decompose the random arrangement into cells of constant complexity (in order to apply the ε -net theory of [52], or the alternative technique of Clarkson [30]). In general, we do not have good bounds on the number of such cells for an arrangement of r algebraic surfaces in \mathbb{R}^d (of constant bounded degree).

Vertical decomposition does work in two dimensions, and yields $O(r^2)$ cells. In three dimensions, it yields $O^*(r^3)$ cells, with a sub polynomial factor that is expressed in terms of the inverse Ackermann function and depends on the degree of the surfaces. For example, if the surfaces are spheres, then the number of cells is $O(r^3 c^{\alpha(r)^2})$, for some constant c [31]. In four dimensions, the result of Koltun [58] yields a bound of $O^*(r^4)$, with a worse sub-polynomial factor, for the size of a vertical decomposition. In higher dimensions, however, the best known upper bound on the size of a vertical decomposition is $O^*(n^{2d-4})$ [58]. No other general efficient techniques are known for decomposing arrangements of algebraic surfaces in higher dimensions into cells of constant complexity.

2.1.2 Partition trees

Cutting allows us to use a divide-and-conquer approach on a set of hyperplanes (or other surfaces), but sometimes it is desirable to apply the divide-and-conquer paradigm to a set of points rather

than hyperplanes. That is, we want to partition the underlying space into cells, where each cell contains roughly the same number of points, in some useful manner. Matoušek [64] proposed the following construction.

Theorem 2.1.6. *Let P be a set of n points in general position in \mathbb{R}^d , let $1 \leq s \leq n$ be a given parameter, and put $t = \lfloor n/s \rfloor$. Then there exists a set of t pairs $(\Delta_1, P_1), \dots, (\Delta_t, P_t)$, such that*

1. P_1, \dots, P_t is a partition of P into t subsets of size s each, except, possibly, for the last one, which is of size $s \leq |P_t| < 2s$.
2. For each $1 \leq i \leq t$, Δ_i is a simplex containing P_i .
3. Any hyperplane crosses at most $O(t^{1-1/d})$ simplices Δ_i .

This construction is useful mainly for the design and implementation of efficient data structures supporting *range searching* queries. These are computational problems of the form: Given a finite set of points P , preprocess it into a data structure, so that, given any query simplex Δ (or any other kind of a simply shaped range), one can efficiently report/count the points of P that are contained in Δ . But it can also be used for incidence problems in a similar manner to cuttings. Among other things, one can prove the Szemerédi-Trotter bound using this partition instead of cutting.

One problem with this construction, though, is that the simplices Δ_i are not necessarily disjoint, which sometimes complicates the analysis. Recently, Chan [25] came up with a probabilistic algorithm, that constructs a simplicial partition which satisfies roughly the same conditions as those of Theorem 2.1.6 (there is a small difference, that the sizes $|P_i|$ of the subsets are bounded only from above by $O(s)$ and not from below), but with the additional property that all simplices are pairwise openly disjoint.

2.1.3 Partition by polynomials

In all the above mentioned partitions, the "walls" of the cells are flat (supported by hyperplanes). Recently, Guth and Katz [51] devised another partition scheme which uses algebraic surfaces (zero sets of polynomials) rather than hyperplanes. We say that a polynomial p *bisects* a finite set of points P , if at least half of the points of P satisfy $p \geq 0$, and at least half of the points of P satisfy $p \leq 0$. The partition is based on the following theorem which is a discrete version of the polynomial ham sandwich theorem of Stone and Tukey [76].

Theorem 2.1.7 (Guth and Katz [51]). *Let P_1, \dots, P_M be disjoint finite sets of points in \mathbb{R}^d , with $M = \binom{d+D}{d} - 1$. Then there exists a polynomial $p(x_1, \dots, x_d)$ of degree at most D that bisects each of the sets P_1, \dots, P_M .*

Now let P be a set of n points in \mathbb{R}^d and consider the following process. First, bisect P with a polynomial p_1 . Put $P_1^1 = \{x \in P \mid p_1(x) < 0\}$, and $P_2^1 = \{x \in P \mid p_1(x) > 0\}$. Each of the sets P_1^1 and P_2^1 contains at most $n/2$ points. Next, bisect P_1^1 and P_2^1 with a polynomial p_2 , and denote the resulting sets by $P_1^2, P_2^2, P_3^2, P_4^2$. After i iterations, we have $t = 2^i$ sets P_1^i, \dots, P_t^i , each containing at most n/t points, which are separated from one another by the zero sets of the

polynomials p_1, \dots, p_i . By theorem 2.1.7, one can ensure that each p_j , for $j = 1, \dots, i$, is of degree at most D_j , where D_j is the smallest integer that satisfies $\binom{d+D_j}{d} \geq 2^{j-1}$, hence $D_j = O(2^{j/d})$. By multiplying all polynomials, we can consider them as one polynomial $p = \prod_{j=1}^i p_j$ of degree at most $D = \sum_{j=1}^i D_j = O(2^{i/d}) = O(t^{1/d})$. The complement of the zero set $p = 0$ is then a subdivision of \mathbb{R}^d into t cells, such that each cell contains some subset P_k^i of P of at most n/t points (note that these cells need not be connected, but may be the union of several connected subcells. The number of connected components, however, is also $O(t)$ — see [65, 66, 82]). We thus have

Theorem 2.1.8 (Guth and Katz [51]). *Let P be a set of n points in \mathbb{R}^d , and let $t < n$ be some integer. Then there exists a polynomial p of degree $O(t^{1/d})$ that partitions space into $\Theta(t)$ cells, such that each cell contains at most n/t points of P .*

There may also be some points of P that do not belong to any cell. These points lie on the surface $p = 0$. When applying a divide-and-conquer analysis using Theorem 2.1.8, one has to analyze, as usual, the points in each cell separately. In addition, however, one also needs to consider the points on the surface $p = 0$, which require a separate analysis. For these points, one can exploit the fact that they lie on a low degree algebraic surface (in comparison to the general case, where a polynomial that vanishes on a set of n points in \mathbb{R}^d has degree $O(n^{1/d})$). We shall demonstrate this method by proving the Szemerédi-Trotter theorem for the special case of equal number of points and lines. The following proof was given by Kaplan et al. [54].

Theorem 2.1.9. *Let P be a set of n points and let L be a set of n lines in \mathbb{R}^2 . Then the number of their incidences is $I(P, L) = O(n^{4/3})$.*

Proof. Put $t = n^{2/3}$, and let p be a polynomial satisfying Theorem 2.1.8 with respect to P and t . p has degree $O(t^{1/2}) = O(n^{1/3})$, and it partitions P into $n^{2/3}$ subsets P_1, \dots, P_t , each of size at most $n^{1/3}$. In addition, there may be points of P that lie on the curve $p = 0$, and therefore, do not belong to any of the above subsets. Denote the set of these points by P_0 .

For each $i = 1, \dots, t$, let $L_i \subseteq L$ denote the subset of lines incident to the points of P_i . In addition, let L_0 denote the lines contained in the curve $p = 0$. We have

$$I = I(P, L) = \sum_{i=1}^t I(P_i, L_i) + I(P_0, L).$$

By Lemma 2.1.3, we may write

$$\begin{aligned} I &\leq \sum_{i=1}^t (|P_i| \cdot |L_i|^{1/2} + |L_i|) + I(P_0, L) \\ &= O(n^{1/3}) \sum_{i=1}^t |L_i|^{1/2} + \sum_{i=1}^t |L_i| + I(P_0, L). \end{aligned}$$

Applying the Cauchy-Schwarz inequality, we get

$$\begin{aligned}
I &= O(n^{1/3}) \left(\sum_{i=1}^t |L_i| \right)^{1/2} t^{1/2} + \sum_{i=1}^t |L_i| + I(P_0, L) \\
&= O(n^{2/3}) \left(\sum_{i=1}^t |L_i| \right)^{1/2} + \sum_{i=1}^t |L_i| + I(P_0, L).
\end{aligned} \tag{2.1.6}$$

Each line of L is either in L_0 , or intersects the curve $p = 0$ at most $O(n^{1/3})$ times. In the former case, the line is not incident to any point from the sets P_1, \dots, P_t , and may only be incident to points of P_0 . In the latter case, the line passes through at most $O(n^{1/3})$ cells, and is thus incident to points from at most $O(n^{1/3})$ of the subsets P_1, \dots, P_t . It thus follows that $\sum_{i=1}^t |L_i| = O(n^{4/3})$, which we substitute in (2.1.6) to get

$$I = O(n^{4/3}) + I(P_0, L). \tag{2.1.7}$$

It remains to bound $I(P_0, L)$. We have $I(P_0, L) = I(P_0, L_0) + I(P_0, L \setminus L_0)$. The number of lines in L_0 is at most the degree of p , which is $O(n^{1/3})$, so we have $I(P_0, L_0) \leq |P_0| \cdot |L_0| = O(n^{4/3})$. As for the incidences of P_0 with $L \setminus L_0$, each of these lines intersects $p = 0$ in at most $O(n^{1/3})$ points and is thus incident to at most $O(n^{1/3})$ points of P_0 , so we have $I(P_0, L \setminus L_0) = O(n^{4/3})$, and in total,

$$I(P_0, L) = O(n^{4/3}).$$

By substituting in (2.1.7), we get

$$I = O(n^{4/3}),$$

as claimed. □

This method was used by Guth and Katz [51] to solve almost completely the celebrated planar distinct distances problem of Erdős, and show that the number of distinct distances determined by n points in the plane is $\Omega(n/\log n)$, almost closing the gap from the upper bound construction of $O(n/\log^{1/2} n)$ distinct distances. Another application is due to Kaplan et al. [53], and independently to Zahl [84], who showed that the number of unit distances between n points in \mathbb{R}^3 is $O(n^{3/2})$. As mentioned in Section 1.8, the previously best known bound, which was proved using cuttings, was $O(n^{3/2} C^{\alpha(n)^2})$ (Clarkson et al. [31]). This is the first discovered example where polynomial decomposition gives strictly better results than cuttings.

2.2 Crossing numbers in graph drawings

In 1997, Székely [79] presented a simple and elegant proof of the Szemerédi-Trotter theorem. The proof is based on the following result on graph drawings in the plane.

Theorem 2.2.1 (The Crossing Lemma, Ajtai et al. [8], Leighton [61]). *Let G be a simple graph with v vertices and e edges. Assume further that $e \geq 4v$. Then, in any drawing of G in the plane, the number of edge crossings X satisfies*

$$X \geq \frac{e^3}{64v^2},$$

and this bound is tight, up to a constant factor.

Székely observed that any point-line configuration can be viewed as a graph drawing, and used this observation to analyze point-line incidences as follows.

Székely's proof of the Szemerédi-Trotter theorem (Theorem 2.1.4). Let P be a set of n points and L a set of l lines in the plane realizing I incidences. We construct from P and L a certain topological graph $G = (V, E)$, that is, a graph together with its drawing in the plane. Each point $p \in P$ becomes a vertex of G , and two points $p, q \in P$ are connected by an edge if they lie on a common line $\ell \in L$ and are consecutive along ℓ ; the edge is drawn as the portion of ℓ between the two points. This yields a drawing of G where the edges are straight line segments. Clearly, G is simple.

If a line $\ell \in L$ contains $k \geq 1$ points of P , then it contributes $k - 1$ edges to G and hence $I \leq |E| + l$. Assume $|E| \geq 4n = 4|V|$. Since the edges are portions of the l lines, at most $\binom{l}{2}$ pairs may cross, so $X \leq \binom{l}{2}$, where X is the number of edge crossings in the specific drawing of G described above. On the other hand, from Theorem 2.2.1, we get $X \geq \frac{1}{64}|E|^3/n^2$. Hence $\frac{|E|^3}{64n^2} \leq X \leq \binom{l}{2}$, and thus $|E| = O(n^{2/3}l^{2/3})$, or $I = O(n^{2/3}l^{2/3} + l)$. Otherwise, if $|E| < 4n$, we have $I < 4n + l$. Hence, $I = O(n^{2/3}l^{2/3} + n + l)$. \square

The idea that a point-curve configuration can be viewed as a graph drawing can also be applied to other incidence problems such as incidences between points and circles, or between points and unit circles. However, if the resulting graph is not simple, the lower bound of Theorem 2.2.1 on the number of crossings does not necessarily hold. In point-circle incidences, for example, two points can be connected by arbitrarily many circular arcs (a configuration known as a *pencil*; see Figure 2.1).

To use Székely's method in more general settings, we need a more general crossing lemma that handles drawings of multigraphs, i.e., graphs that can have multiple edges between the same pair of vertices.

Theorem 2.2.2 (Generalized Crossing Lemma, Székely [79]). *Let G be a multigraph with v vertices, e edges, and maximum edge multiplicity m . Assume further that $e \geq 5mv$. Then, in any drawing of G in the plane, the number of edge crossings X satisfies*

$$X = \Omega\left(\frac{e^3}{v^2m}\right),$$

and this bound is tight.

We can now easily apply Székely's method to bound the number of incidences between points and unit circles, since two points can be connected by at most two short and two long unit circular arcs. Hence, the edge multiplicity of a graph drawing with edges drawn as portions of unit circles is at most 4, and the same proof for lines applies to unit circles, giving an upper bound of the same order of magnitude.

For incidences between points and other classes of curves, such as arbitrary circles, the edge multiplicity need not be bounded by a constant. However, using a more involved argument, it is possible to apply Székely's method to general curves and derive nontrivial incidence bounds. For example, Pach and Sharir [69] used this method to derive Theorem 1.4.1.

2.2.1 Cutting curves into pseudo-segments

Another way to use the Crossing Lemma where edge multiplicity may be unbounded, is by cutting the curves into pseudo-segments (see Section 1.3). We shall demonstrate this trick on point-circle incidences.

Lemma 2.2.3 (Aronov and Sharir [17], Marcus and Tardos [63]). *The number of incidences between n points and c circles in the plane is $O(n^{2/3}c^{2/3} + n + c^{3/2} \log c)$.*

Remark: The contribution of Marcus and Tardos [63] was to replace a slightly weaker term $O^*(c^{3/2})$ in the original bound of [17] by $O(c^{3/2} \log c)$; see below.

Proof. Let P and Γ be, respectively, a set of n points and a set of c circles in the plane, and let $I = I(P, \Gamma)$. Let $\Gamma' \subseteq \Gamma$ be all the circles of Γ incident to at least three points of P . We have $I \leq I' + 2c$, where $I' = I(P, \Gamma')$. We define the topological graph $G = (V, E)$ by letting $V = P$, and by connecting each pair of consecutive points of P along a circle $\gamma \in \Gamma$ by an edge, drawn as the connecting circular arc. We have $|E| = I'$.

In general, G is not simple. To transform it into a simple graph, we cut some of its edges by cutting the circles of Γ' into a family Γ'' of pseudo-segments. By the result of Marcus and Tardos [63], mentioned above, we can do it with $O(c^{3/2} \log c)$ cuts. Any such cut is, in fact, a removal of a point from a circle of Γ'' . Without loss of generality, we may assume that none of these cuts coincides with any point of P , and so, each cut occurs at the interior of some edge drawing (i.e., an arc). Each such edge is removed from G , and we denote the resulting edge set by E' . We have $|E'| \geq I' - O(c^{3/2} \log c)$, or $I' \leq |E'| + O(c^{3/2} \log c)$, from which we have

$$I \leq |E'| + 2n + O(c^{3/2} \log c).$$

We now apply the crossing lemma (Theorem 2.2.1) to the resulting simple graph $G' = (V, E')$. On one hand, the number of crossings X is at most the number of circle intersections, i.e., $X \leq 2 \binom{c}{2}$. On the other hand, by the crossing lemma, we have that either $|E'| \leq 4n$, or $X \geq \frac{|E'|^3}{64n^2}$. In the former case, we get

$$I = O(n + c^{3/2} \log c),$$

which satisfies the assertion of the lemma. In the latter case, we get $|E'| = O(n^{2/3}c^{2/3})$, or

$$I = O(n^{2/3}c^{2/3} + n + c^{3/2} \log c),$$

and the lemma holds in this case too. This completes the proof. \square

This lemma is used in [17] as an initial weak bound, which is later on strengthened by cuttings to derive the bound (1.2.1). See Theorem 2.1.5

Part II

Results

Chapter 3

Point-Hyperplane Incidences

In this chapter we show that if the number I of incidences between m points and n planes in \mathbb{R}^3 is sufficiently large, then the incidence graph contains a large complete bipartite subgraph involving r points and s planes, so that $rs \geq \frac{I^2}{mn} - a(m+n)$, for some constant $a > 0$. This is shown to be almost tight in the worst case because there are examples of arbitrarily large sets of points and planes (where m and n are arbitrarily large and independent, and I is sufficiently large for those m and n), where the largest complete bipartite incidence subgraph records only $\frac{I^2}{mn} - \frac{m+n}{16}$ incidences. We also make some steps towards generalizing this result to higher dimensions.

The results in this chapter were published, coauthored with Sharir, in *SIAM Journal on Discrete Mathematics* [12].

3.1 Introduction

Let P be a set of m points, and let Π be a set of n hyperplanes in \mathbb{R}^d . As we show, an interesting property of point-hyperplane incidence graphs is that, if the number of incidences is large (close to mn), then the incidence graph contains large complete bipartite subgraphs. Such a subgraph is in fact a configuration consisting of many hyperplanes of Π intersecting at a common lower-dimensional affine subspace H , together with many points of P , all incident to H . This property arises, in one way or another, in almost all previous works on point-hyperplane incidences; see Section 1.5 for details. In this chapter we continue to study this property and ask: Given a point-hyperplane configuration with many incidences, what is the size of the largest complete bipartite incidence subgraph? To state the question more precisely, we define

$$\text{rs}(P, \Pi) = \max \{rs \mid K_{r,s} \subseteq G(P, \Pi)\},$$

where $K_{r,s}$ denotes the complete bipartite subgraph with r vertices on one side and s vertices on the other, and the notation $K_{r,s} \subseteq G(P, \Pi)$ means that $K_{r,s}$ is a subgraph of $G(P, \Pi)$, such that there are some r points of P and s hyperplanes of Π all incident to one another. We let

$$\text{rs}_d(m, n, I) = \min_{\substack{|P|=m \\ |\Pi|=n \\ I(P, \Pi) \geq I}} \text{rs}(P, \Pi)$$

denote the minimum of $\text{rs}(P, \Pi)$ over all choices of a set $P \subset \mathbb{R}^d$ of m points and a set Π of n hyperplanes in \mathbb{R}^d , such that $I(P, \Pi) \geq I$. Note that $\text{rs}_d(m, n, I) \geq \max\{\frac{I}{m}, \frac{I}{n}\}$, since there always exists a point incident to at least I/m hyperplanes, and a hyperplane incident to at least I/n points, which give rise to both subgraphs $K_{I/n, 1}$ and $K_{1, I/m}$. We thus have $\text{rs}_d(m, n, I) \geq \frac{I}{\min\{m, n\}} = \Omega\left(\frac{I}{m} + \frac{I}{n}\right)$, and any non-trivial estimate must exceed this lower bound.

3.1.1 Our results

For the case $d = 3$, we can estimate $\text{rs}_d(m, n, I)$ almost exactly:

Theorem 3.1.1. (i) *If $I = \Omega(m\sqrt{n} + n\sqrt{m})$, with a sufficiently large multiplicative constant, then*

$$\text{rs}_3(m, n, I) = \frac{I^2}{mn} - \Theta(m + n).$$

(ii) *If $m \leq n$, $I = O(n\sqrt{m})$, and $I = \Omega((mn)^{3/4})$, for appropriate multiplicative constants, then*

$$\text{rs}_3(m, n, I) = \Theta\left(\frac{I^4}{m^2n^3} + \frac{I}{m}\right).$$

(iii) *Symmetrically, if $m \geq n$, $I = O(m\sqrt{n})$, and $I = \Omega((mn)^{3/4})$, then*

$$\text{rs}_3(m, n, I) = \Theta\left(\frac{I^4}{m^3n^2} + \frac{I}{n}\right).$$

(iv) *If $I = O(m^{3/4}n^{3/4} + m + n)$, then*

$$\text{rs}_3(m, n, I) = \Theta\left(\frac{I}{m} + \frac{I}{n}\right).$$

The interesting case is (i), where the number of incidences is largest. The upper bound construction for this case consists of almost disjoint complete bipartite subconfigurations.

As the dimension d increases beyond 3, there are progressively more different ranges of I (as a function of m, n) where the lower or upper bounds for rs_d might change qualitatively. A complete analysis of all these cases seems hard at this point, already in four dimensions, and even the simpler question of focusing on the range where I is largest (such as in Theorem 3.1.1 (i)) seems quite hard. There are actually two subproblems here: (a) Determine the range itself (i.e., how large must I be to ensure the best available bounds on rs_d), and (b) determine these best bounds for rs_d in this range.

We have obtained some non-trivial results for the higher dimensional case, but not tight ones:

Theorem 3.1.2 (Lower bound). *For any dimension $d \geq 3$, there exist some constants b_d and c_d such that for any m, n , and I such that $I \geq b_d(mn^{1-\frac{1}{d-1}} + m^{1-\frac{1}{d-1}}n)$, we have*

$$\text{rs}_d(m, n, I) \geq c_d \left(\frac{I}{mn}\right)^{d-1} mn.$$

Theorem 3.1.3 (Upper bound). *For any dimension $d \geq 3$, there exist some constants B_d and C_d such that for any m, n , and I such that $I \geq B_d(mn)^{1-\frac{1}{d-1}}$, we have*

$$rs_d(m, n, I) \leq C_d \left(\frac{I}{mn} \right)^{\frac{d+1}{2}} mn.$$

Note that for $d = 3$, both theorems yield the same bound on rs_3 , which is also identical (up to multiplicative constants) to that in Theorem 3.1.1(i). Moreover, all three theorems apply within the same (up to multiplicative constants) range $I = \Omega(m\sqrt{n} + n\sqrt{m})$ (Theorem 3.1.3 applies within a wider range).

It is interesting to compare these bounds to the equivalent bounds for general graphs. There are (m, n) -bipartite graphs with $\frac{1}{2}mn$ edges, such that the largest complete bipartite subgraph has fewer than $2(m + n)$ edges. In fact, a random graph satisfies this property with very high probability. In contrast, Theorems 3.1.1 and 3.1.2 assert that a point-hyperplane incidence graph with these many edges has a complete bipartite subgraph with $\Omega(mn)$ edges.

3.2 Large complete bipartite incidence subgraphs in \mathbb{R}^3

In this section we prove Theorem 3.1.1. The main technical effort is in proving the lower bound in part (i) of the theorem. Before getting to this, we first dispose of the much simpler upper bound construction.

Lemma 3.2.1 (Theorem 3.1.1(i) — upper bound). *There exist configurations of m points and n planes in \mathbb{R}^3 , with arbitrarily large m , arbitrarily large n , and with $I \geq \max\{m\sqrt{n}, n\sqrt{m}\}$ incidences, such that every $K_{r,s}$ incidence subgraph satisfies*

$$rs \leq \frac{I^2}{mn} - \frac{1}{16}(m + n).$$

Remark: In these constructions, there exist numbers between $\max\{m\sqrt{n}, n\sqrt{m}\}$ and mn that I cannot attain. For example, we cannot have $I = mn - 1$ unless $m \leq 2$ or $n \leq 2$. However, in terms of order of magnitude, I can assume values proportional to any value between $\Theta(m\sqrt{n} + n\sqrt{m})$ and $\Theta(mn)$.

Proof. Without loss of generality, we present the construction for $m \geq n$. Fix three arbitrarily large numbers, $r \geq s \geq k \geq 2$. Take a set L of k parallel lines such that no three lines are coplanar. Then each pair of lines in L determine a distinct plane. We include all these $\binom{k}{2}$ planes in the set Π of n planes. We include in Π additional planes, each of which contains just one of the lines of L , so that each line is incident to exactly s planes. The set P of points consists of $m = rk$ elements, so that each line of L contains r points. The set Π consists of $sk - \binom{k}{2}$ planes, and $I(P, \Pi) = krs$. Put $n_0 = sk = n + \binom{k}{2}$. Note that $rs(P, \Pi) = rs$, because the corresponding subgraph $K_{r,s}$ cannot contain points from three lines — no plane passes through three lines of L ,

and if it contains points from two lines then there is only one plane that passes through both lines. This gives

$$rs(P, \Pi) = rs = \frac{I^2}{mn_0} = \frac{I^2}{m} \cdot \frac{1}{n + \binom{k}{2}} \leq \frac{I^2}{m} \cdot \frac{1}{n + k^2/4}.$$

We now use the inequality

$$\frac{1}{x+h} \leq \frac{1}{x} - \frac{h}{2x^2},$$

which holds for all $x \geq h > 0$, with $x = n$ and $h = k^2/4 \leq n$, to get

$$rs \leq \frac{I^2}{m} \cdot \left(\frac{1}{n} - \frac{k^2/4}{2n^2} \right) = \frac{I^2}{mn} - \frac{(kI)^2}{8mn^2},$$

and since $k = \frac{mn_0}{I}$, we get

$$rs \leq \frac{I^2}{mn} - \frac{(mn_0)^2}{8mn^2} \leq \frac{I^2}{mn} - \frac{m}{8} \leq \frac{I^2}{mn} - \frac{1}{16}(m+n),$$

as claimed. Note that $I = \frac{mn_0}{k} = ms \geq m\sqrt{n_0} \geq m\sqrt{n}$ is in the required range.

The case $m \leq n$ is handled in a symmetric manner, using duality between points and planes. This completes the proof. \square

Remark: A simpler construction, consisting of *disjoint* copies of $K_{r,s}$, yields the lower bound $I^2/(mn)$. Our construction shows that a lower order term proportional to $m+n$ is unavoidable.

Lemma 3.2.2 (Theorem 3.1.1(iv) — upper bound). *There exist sufficiently large constants $b, c > 0$ and a sufficiently small constant $C > 0$, such that for any $m, n > b$, and for any I such that $c(m+n) < I < C(mn)^{3/4}$, there exist configurations of at most m points and at most n planes in \mathbb{R}^3 with at least I incidences, such that every $K_{r,s}$ incidence subgraph satisfies*

$$rs \leq \frac{6I}{\min\{m, n\}}.$$

Remark: The construction of [22], in which there are no $K_{3,2}$ or $K_{2,3}$ incidence subgraphs, provides us with an example where $rs = O(I/\min\{m, n\})$. This construction, however, is good only for the range $I = O(m^{3/5}n^{4/5} + m^{4/5}n^{3/5})$, and cannot be used for larger values of I within the assumed larger range $O((mn)^{3/4})$. In contrast, our construction is good for the entire range specified in Lemma 3.2.2, but may have complete bipartite incidence subgraphs with an arbitrarily large number of elements on both sides.

Proof. This construction resembles similar constructions of Elekes [37]. Put $k = \lfloor \sqrt{2I/n} \rfloor$, $l = \lfloor \sqrt{6I/m} \rfloor$, and $t = \lfloor (mn)^{3/2}/(12\sqrt{3}I^2) \rfloor$. With an appropriate choice of the constants, we may assume that $k, l, t \geq 100$, say, and so $k^3l^3t \geq I$. Define

$$P = \{(x_1, x_2, x_3) \mid x_1, x_2 \in \{1, \dots, k\}, \text{ and } x_3 \in \{1, \dots, 3klt\}\},$$

and

$$\Pi = \{x_3 = a_1tx_1 + a_2tx_2 + b \mid a_1, a_2 \in \{1, \dots, l\}, \text{ and } b \in \{1, \dots, klt\}\}.$$

The set P consists of $3k^3lt \leq m$ points, and the set Π consists of $kl^3t \leq n$ planes. Each plane is incident to k^2 points, so the number of incidences is $k^3l^3t \geq I$. In addition, each point is incident to at most l^2 planes.

Now there are three types of complete bipartite incidence subgraphs $K_{r,s}$; in each case we consider the maximal such subgraph.

1. Between a point and all its incident planes. Then $r = 1$, $s \leq l^2$, and $rs \leq l^2 \leq 6I/m$.
2. Between a plane and all its incident points. Then $r = k^2$, $s = 1$, and $rs = k^2 \leq 2I/n$.
3. Between some r collinear points and s collinear planes, all incident to the same line. Then $r \leq k$, $s \leq l$, and $rs \leq kl = 2\sqrt{3}I/\sqrt{mn} \leq 2\sqrt{3}I/\min\{m, n\}$.

In either case, we have $rs \leq \frac{6I}{\min\{m, n\}}$. This completes the proof. \square

We now turn to the lower bound in Theorem 3.1.1(i), which will be a consequence of the following result.

Theorem 3.2.3 (Cf. Theorem 3.1.1(i) — lower bound). *Let P be a set of m points and Π a set of n planes in \mathbb{R}^3 , with I incidences between them. Then there exists a line ℓ containing r points of P and contained in s planes of Π , such that*

$$\sqrt{rs} \geq \frac{I}{\sqrt{mn}} - \frac{a(m+n)\sqrt{mn}}{I},$$

where $a > 0$ is some sufficiently large constant.

This inequality, when squared, implies that $rs \geq \frac{I^2}{mn} - 2a(m+n)$. This establishes the lower bound of Theorem 3.1.1(i). Note that here there is no lower bound requirement on I , as opposed to Theorem 3.1.1(i), where it is required that $I = \Omega(m\sqrt{n} + n\sqrt{m})$. However, if $I < \sqrt{amn(m+n)}$, then the right hand side is negative. Thus the theorem is interesting only for point-plane configurations with $I > \sqrt{amn(m+n)} = \Omega(m\sqrt{n} + n\sqrt{m})$.

To prove Theorem 3.2.3, we use Theorem 2.1.4 (Szemerédi and Trotter [80]), and Theorem 1.5.1 (Elekes and Tóth [41]). Using these bounds, we can prove the following result, which is slightly weaker than Theorem 3.2.3; see below for a more detailed comparison.

Lemma 3.2.4. *Let P be a set of m points and Π a set of n planes in \mathbb{R}^3 , such that $I = I(P, \Pi) = \Omega((mn)^{3/4} + m\sqrt{n})$, with a sufficiently large multiplicative constant. Then there exists a line ℓ containing r points of P and contained in s planes of Π , such that*

$$rs = \Omega\left(\min\left\{\frac{I^4}{m^2n^3}, \frac{I^2}{mn}\right\}\right).$$

This bound is worst case tight for all combinations of m, n, I that satisfy the above conditions.

Proof. Applying Theorem 1.5.1 with $d = 3$, we see that, when the constant of proportionality is chosen sufficiently large, and for some appropriately chosen $\eta < \eta_3$, most incidences are with planes of Π that are η -degenerate, i.e., for each such plane, at least an η -fraction of its incident points are contained in a single line.

Let $\Pi' \subseteq \Pi$ be the subset of those planes in Π that contain at least $I/(2n)$ points each, and are η -degenerate. By the preceding argument, if the constant of proportionality in the assumed lower bound on I is sufficiently large, then $I(P, \Pi') \geq I/3$, say. We replace each plane of Π' with a line that lies on it and contains an η -fraction of its incident points. Thus, each such line contains at least $\eta I/(2n)$ points of P . By projecting these lines and the points of P onto some generic plane, and applying (2.1.3) from Theorem 2.1.4, the number of incidences between the points of P and these lines is

$$I' = O\left(\frac{m^2}{(\eta I/(2n))^2} + m\right) = O\left(\frac{m^2 n^2}{I^2} + m\right).$$

Note that $I(P, \Pi')$ differs from I' , because we count in $I(P, \Pi')$ each line with its *multiplicity*, equal to the number of planes of Π' that contain it. The average multiplicity of a line is thus

$$s = \frac{I(P, \Pi')}{I'} = \Omega\left(\frac{I}{I'}\right) = \Omega\left(\min\left\{\frac{I^3}{m^2 n^2}, \frac{I}{m}\right\}\right).$$

By the pigeonhole principle, some line ℓ does have at least this multiplicity, i.e., it is contained in at least s planes. By construction, it also contains $r = \Omega(I/n)$ points. Altogether, we get

$$rs = \Omega\left(\min\left\{\frac{I^4}{m^2 n^3}, \frac{I^2}{mn}\right\}\right).$$

We have thus found a line ℓ with the asserted properties.

We can obtain a point-plane configuration that has a matching upper bound on $\text{rs}(P, \Pi)$ of the same order of magnitude as the lower bound we have just proved (that is, unless the trivial bound $\text{rs}(P, \Pi) \geq \max\{I/m, I/n\}$ dominates). This is done as follows. We take m coplanar points spanning the maximal number of lines incident to $r = \Theta(I/n)$ of the points, which, by Theorem 2.1.4, is

$$\Theta\left(\frac{m^2}{r^3} + \frac{m}{r}\right) = \Theta\left(\frac{m^2 n^3}{I^3} + \frac{mn}{I}\right).$$

We then let each such line occur on $s = \Theta(\min\{I^3/(m^2 n^2), I/m\})$ planes. The constants of proportionality are chosen so that the total number of planes is n , and the number of incidences is I . We have thus shown that the bound asserted in the lemma is worst case tight up to multiplication by a constant. \square

From this lemma, we have:

Proof of Theorem 3.1.1(ii,iii). If $I = O(n\sqrt{m})$, for an appropriate multiplicative constant, then, in the bound of Lemma 3.2.4, the first term $I^4/(m^2 n^3)$ is smaller than the second term $I^2/(mn)$. Moreover, the lower bound on I that the lemma requires, holds under the assumptions $I = \Omega((mn)^{3/4})$ and $m \leq n$. Hence, under the assumptions of part (ii) of the theorem, $\text{rs}_3(m, n) =$

$\Omega(I^4/(m^2n^3))$, which clearly implies the lower bound of Theorem 3.1.1(ii). The upper bound also follows easily from Lemma 3.2.4. Finally, Theorem 3.1.1(iii) follows by point-plane duality. \square

On the other hand, if $I = \Omega(n\sqrt{m})$, then the second term in Lemma 3.2.4, namely $I^2/(mn)$, dominates. We thus get:

Corollary 3.2.5. *Let P be a set of m points and Π a set of n planes in \mathbb{R}^3 , such that $I = I(P, \Pi) = \Omega(m\sqrt{n} + n\sqrt{m})$, with a sufficiently large multiplicative constant. Then there exists a line ℓ containing r points of P and contained in s planes of Π , such that*

$$rs = \Omega\left(\frac{I^2}{mn}\right).$$

Note that Lemma 3.2.4 is applicable in Corollary 3.2.5, since $(mn)^{3/4}$ is always dominated by $m\sqrt{n} + n\sqrt{m}$.

This is already very close to the bound we are trying to prove, except for the multiplicative constant. We shall now get rid of this constant and finish the proof of Theorem 3.2.3. Recall that the theorem states that

$$\sqrt{rs} \geq \frac{I}{\sqrt{mn}} - \frac{a(m+n)\sqrt{mn}}{I},$$

for some r points and s planes all incident to one another, and for some constant $a > 0$.

Proof of Theorem 3.2.3. Let P be a set of m points and Π a set of n planes in \mathbb{R}^3 , with $I = I(P, \Pi)$ incidences. By Corollary 3.2.5, there exist positive absolute constants A, k , and η , such that for all $m > k$ and $n > k$, if $I > A(m\sqrt{n} + n\sqrt{m})$, then

$$\sqrt{rs(P, \Pi)} \geq \frac{\eta I}{\sqrt{mn}}.$$

We choose the constant a so that it satisfies $a \geq \max\left\{4, 2A^2, k, \frac{2}{\eta}\right\}$.

The proof proceeds by induction on m and n . It is easy to see that the theorem holds for sufficiently small values of m , n , or I . More precisely, if $I < \sqrt{amn(m+n)}$, then $\frac{I}{\sqrt{mn}} - \frac{a(m+n)\sqrt{mn}}{I} < 0$, and the theorem holds trivially. Moreover, if $m \leq a$ or $n \leq a$, then $I \leq mn < \sqrt{amn(m+n)}$. Hence, the theorem holds for all m and n such that $m \leq a$ or $n \leq a$.

Suppose then that $m > a$ and $n > a$ are arbitrary, and that the claim holds for all (m', n') satisfying $m' < m$ and $n' < n$. Let P be a set of m points and Π a set of n planes in \mathbb{R}^3 with $I > \sqrt{amn(m+n)}$ incidences between them. Let ℓ be a line that maximizes rs , where $r = |\ell \cap P|$, and $s = |\{\pi \in \Pi \mid \pi \supset \ell\}|$.

Remove from the setting all the points and planes incident to ℓ . We are left with $m - r$ points and $n - s$ planes, and denote by I_1 the number of incidences among them. We note that

$$I_1 \geq I - rs - (m+n) + (r+s). \quad (3.2.1)$$

Indeed, by removing the elements incident to ℓ , we lose the rs incidences between these elements. We may also lose incidences between the removed points and the surviving planes, and between

the removed planes and the surviving points. However, each surviving plane can be incident to at most one removed point, and each surviving point can be incident to at most one removed plane. This implies the asserted inequality (3.2.1).

We next choose a line ℓ_1 incident to r_1 of the remaining points, and to s_1 of the remaining planes, such that $r_1 s_1$ is maximized. If $\sqrt{r_1 s_1} \geq \frac{I}{\sqrt{mn}} - \frac{a(m+n)\sqrt{mn}}{I}$, then we are done, since, by construction, $rs \geq r_1 s_1$. Otherwise, we may write

$$\frac{I}{\sqrt{mn}} - \frac{a(m+n)\sqrt{mn}}{I} > \sqrt{r_1 s_1} \geq \frac{I_1}{\sqrt{(m-r)(n-s)}} - \frac{a((m+n)-(r+s))\sqrt{(m-r)(n-s)}}{I_1},$$

where the right inequality follows from the induction hypothesis. Thus,

$$\frac{I}{\sqrt{mn}} > \frac{I_1}{\sqrt{(m-r)(n-s)}} + a \left(\frac{(m+n)\sqrt{mn}}{I} - \frac{((m+n)-(r+s))\sqrt{(m-r)(n-s)}}{I_1} \right).$$

Put

$$h = \frac{(m+n)\sqrt{mn}}{I} - \frac{((m+n)-(r+s))\sqrt{(m-r)(n-s)}}{I_1},$$

so we have

$$\frac{I}{\sqrt{mn}} > \frac{I_1}{\sqrt{(m-r)(n-s)}} + ah.$$

We now distinguish between the two cases $h \geq 0$ and $h < 0$. If $h \geq 0$, then we have

$$\frac{I}{\sqrt{mn}} > \frac{I_1}{\sqrt{(m-r)(n-s)}}, \quad (3.2.2)$$

or, using the inequality $(m-r)(n-s) \leq (\sqrt{mn} - \sqrt{rs})^2$, and applying (3.2.1),

$$\begin{aligned} \frac{I}{\sqrt{mn}} &> \frac{I - rs - (m+n)}{\sqrt{mn} - \sqrt{rs}} \\ \implies I\sqrt{mn} - I\sqrt{rs} &> I\sqrt{mn} - \sqrt{mnr}s - (m+n)\sqrt{mn} \\ \implies rs - \frac{I}{\sqrt{mn}}\sqrt{rs} + (m+n) &> 0. \end{aligned}$$

This quadratic inequality in the variable \sqrt{rs} solves to

$$\begin{aligned} \sqrt{rs} &> \frac{\frac{I}{\sqrt{mn}} + \sqrt{\frac{I^2}{mn} - 4(m+n)}}{2}, \quad \text{or} \\ \sqrt{rs} &< \frac{\frac{I}{\sqrt{mn}} - \sqrt{\frac{I^2}{mn} - 4(m+n)}}{2}. \end{aligned}$$

Note that, since $a \geq 4$, it follows that $\frac{I^2}{mn} - 4(m+n) \geq 0$ for the assumed range of I . We can then use the inequality $\sqrt{x - \Delta x} \geq \sqrt{x} - \frac{\Delta x}{\sqrt{x}}$, which holds for $0 \leq \Delta x \leq x$, to obtain

$$\begin{aligned}\sqrt{rs} &> \frac{I}{\sqrt{mn}} - \frac{2(m+n)\sqrt{mn}}{I}, \quad \text{or} \\ \sqrt{rs} &< \frac{2(m+n)\sqrt{mn}}{I}.\end{aligned}$$

Since $a \geq 2A^2$, it is easily checked that Corollary 3.2.5 is applicable for the assumed range of I , and implies that $\sqrt{rs} \geq \frac{\eta I}{\sqrt{mn}}$. Hence, if the second case were possible, we would have $\frac{\eta I}{\sqrt{mn}} < \frac{2(m+n)\sqrt{mn}}{I}$, or $I < \sqrt{\frac{2}{\eta}mn(m+n)}$, which, having chosen $a \geq \frac{2}{\eta}$, would contradict our assumption on I . Hence, only the first inequality is possible, and the theorem holds in this case.

Consider now the case $h < 0$. We have

$$\begin{aligned}\frac{(m+n)\sqrt{mn}}{I} &< \frac{((m+n) - (r+s))\sqrt{(m-r)(n-s)}}{I_1} \\ &< \frac{(m+n)\sqrt{(m-r)(n-s)}}{I_1} \\ \implies \frac{I}{\sqrt{mn}} &> \frac{I_1}{\sqrt{(m-r)(n-s)}}.\end{aligned}$$

But this is exactly inequality (3.2.2), which, as we have already seen, implies $\sqrt{rs} > \frac{I}{\sqrt{mn}} - \frac{a(m+n)\sqrt{mn}}{I}$, so the theorem holds in this case too.

This completes the induction step, and thus the proof of the theorem. \square

Figure 3.1 summarizes our findings. Each differently-shaded region represents certain values of m, n and I , and has a different lower bound for rs .

3.3 Large complete bipartite incidence subgraphs in higher dimensions

In Lemma 3.2.4 we require that $I = \Omega((mn)^{3/4} + m\sqrt{n})$, because we want to ensure that most planes are ‘degenerate’ in the sense that they can be replaced by lines that they contain, and the number of incidences will stay roughly the same. However, the lemma holds in a considerably more general setting, involving any family of ‘degenerate’ subsets of points in any dimension. Specifically, we call a finite set of points $S \subset \mathbb{R}^d$ (η, j) -degenerate, if some j -flat contains at least an η -fraction of the points of S . In other words, if

$$|F \cap S| \geq \eta|S|$$

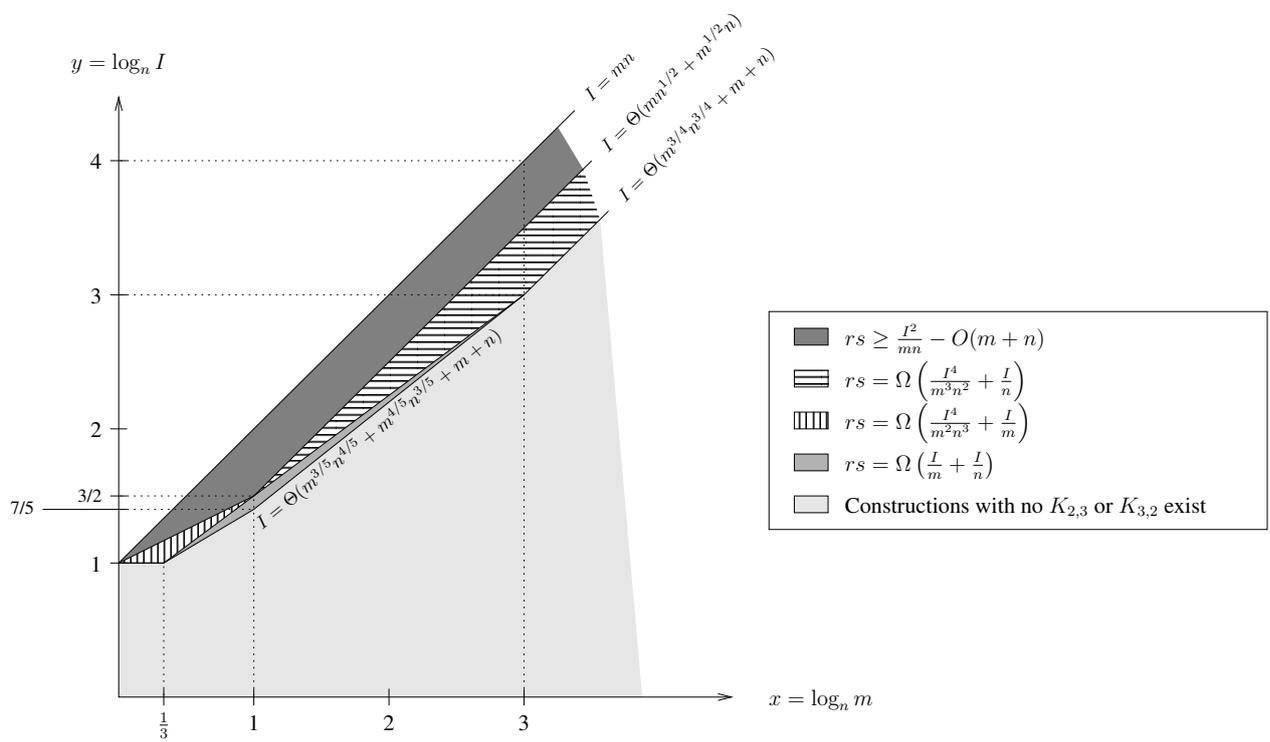


Figure 3.1: The known lower bounds for the maximum number of edges in a complete bipartite incidence subgraph in \mathbb{R}^3 .

for some j -flat F of \mathbb{R}^d . If no such j -flat exists, we call S (η, j) -non-degenerate.

In addition to degeneracy, we will also require that the subsets be *affinely closed* with respect to the ground set P . For a subset $S \subset P$, we say that S is an *affinely closed subset* of P , if for any j -flat $F \subset \mathbb{R}^d$ of any dimension j , if F is spanned by S , then $S \cap F = P \cap F$.

With this notion of degeneracy and affine closedness, Lemma 3.2.4 becomes a special case of the following lemma (with each plane $\pi \in \Pi$ being mapped to the set $\pi \cap P$, and the entire set of planes Π being mapped to a multiset of subsets of P).

Lemma 3.3.1. *Let $P \subset \mathbb{R}^d$ be a set of m points, let $\mathcal{T} \subseteq 2^P$ be a multiset of n affinely closed subsets of P , and let $0 < \eta < 1$ be some constant, such that all the members of \mathcal{T} are $(\eta, 1)$ -degenerate. Then there exist a subset $R \subseteq P$ of $|R| = r$ points and a subfamily $\mathcal{S} \subset \mathcal{T}$ of $|\mathcal{S}| = s$ subsets (counted with multiplicity), such that $R \subseteq S$ for each $S \in \mathcal{S}$, and*

$$rs = \Omega \left(\min \left\{ \frac{I^4}{m^2 n^3}, \frac{I^2}{mn} \right\} \right),$$

where $I = \sum_{T \in \mathcal{T}} |T|$.

In particular, the multiset \mathcal{T} need not be induced by planes, as in Lemma 3.2.4, but can be induced by k -flats of any dimension k . The proof is omitted, but essentially, it is identical to that of Lemma 3.2.4: We replace each subset $S \in \mathcal{S}$ by a line that contains a fraction of its points, and estimate the average multiplicity of the lines using the Szemerédi-Trotter bound within a generic 2-plane onto which we project the points and lines.

Before we proceed, we would like to recall the Elekes-Tóth bound (Theorem 1.5.1), and to reformulate it slightly:

Theorem 3.3.2. *Let $P \subset \mathbb{R}^d$ be a set of m points, and let Π be a family of n j -flats in \mathbb{R}^d , for $j < d$. Let η be some parameter such that $\eta \leq \eta_{j+1}$, where η_{j+1} is the constant defined in Theorem 1.5.1, and assume further that the j -flats of Π are η -nondegenerate with respect to P . Then there exist constants $E_{\eta,j}$ and $\hat{E}_{\eta,j}$, that depend on η and j , but not on m and n , such that*

1. *If all the j -flats of Π are k -rich with respect to P , for any $k \leq m$, then*

$$n \leq E_{\eta,j} \cdot \max \left\{ \frac{m^{j+1}}{k^{j+2}}, \frac{m^j}{k^j} \right\}. \quad (3.3.3)$$

2. *The number of incidences is upper bounded by*

$$I(P, \Pi) \leq \hat{E}_{\eta,j} \cdot \max \left\{ m^{(j+1)/(j+2)} n^{(j+1)/(j+2)}, mn^{j/(j+1)} \right\}. \quad (3.3.4)$$

The reader can verify that this theorem is indeed equivalent to Theorem 1.5.1. Note that in Theorem 1.5.1 we always have $d = j + 1$, but the theorem also applies to larger values of d , because we can always project the setup onto some generic $(j + 1)$ -flat. We put

$$E_j = E_{\eta_{j+1},j}, \quad \text{and} \\ \hat{E}_j = \hat{E}_{\eta_{j+1},j}.$$

Next, we obtain the following generalization of Lemma 3.3.1.

Lemma 3.3.3. *Let $P \subset \mathbb{R}^d$ be a set of m points, let $\mathcal{T} \subseteq 2^P$ be a multiset of n affinely closed subsets of P , and let $\eta > 0$ and $j \geq 1$ be some constants, such that all the members of \mathcal{T} are (η, j) -degenerate. Then there exist a subset $R \subseteq P$ of $|R| = r$ points and a subfamily $\mathcal{S} \subset \mathcal{T}$ of $|\mathcal{S}| = s$ subsets (again, counted with multiplicity), such that $R \subseteq S$ for each $S \in \mathcal{S}$, and*

$$rs \geq C_{\eta,j} \cdot \min \left\{ \frac{I^{j+3}}{m^{j+1}n^{j+2}}, \frac{I^{j+1}}{m^j n^j} \right\},$$

where $I = \sum_{T \in \mathcal{T}} |T|$, and $C_{\eta,j}$ is a sufficiently small constant that depends only on η and j (but not on m , n , or I).

Proof of Lemma 3.3.3. The proof proceeds by double induction on j and n . The base case $j = 1$ is given by Lemma 3.3.1 (for any n), with $C_{\eta,1}$ being the implicit constant in the $\Omega(\cdot)$ expression in that lemma. Suppose now that the lemma holds for $j - 1 \geq 1$, and also for j and for $n' < n$, and we shall see that it also holds for j and for n . (The base case for n , at any fixed j , is trivial, with an appropriate choice of $C_{\eta,j} \leq C_{\eta,j}^{(0)}$, where $C_{\eta,j}^{(0)}$ is sufficiently small.)

Delete from \mathcal{T} all the members containing fewer than $I/(2n)$ points, and let \mathcal{T}' denote the multiset of the remaining sets. We have $I' = \sum_{T \in \mathcal{T}'} |T| \geq I/2$. If $|\mathcal{T}'| < n/4$, then, by induction on n , we have subsets $R \subseteq P$ and $\mathcal{S} \subseteq \mathcal{T}'$, such that $R \subseteq S$ for each $S \in \mathcal{S}$, and

$$\begin{aligned} |R| \cdot |\mathcal{S}| &\geq C_{\eta,j} \cdot \min \left\{ \frac{(I/2)^{j+3}}{m^{j+1}(n/4)^{j+2}}, \frac{(I/2)^{j+1}}{m^j (n/4)^j} \right\} \\ &= C_{\eta,j} \cdot \min \left\{ 2^{j+1} \frac{I^{j+3}}{m^{j+1}n^{j+2}}, 2^{j-1} \frac{I^{j+1}}{m^j n^j} \right\}. \end{aligned}$$

Since $j \geq 2$, we obtain R and \mathcal{S} that satisfy the asserted lower bound. We can therefore assume that there are at least $n/4$ remaining sets in \mathcal{T}' .

For each set $T \in \mathcal{T}'$, let π_T be a j -flat (which exists by assumption) containing at least $\eta|T| \geq \frac{\eta I}{2n}$ points of P . Project these j -flats and the points of P onto some generic $(j+1)$ -space Q , and partition \mathcal{T}' into two subfamilies:

$$\mathcal{T}_1 = \{T \in \mathcal{T}' \mid \text{The projection of } \pi_T \text{ is } \eta_{j+1}\text{-non-degenerate in } Q\}, \quad \text{and} \quad \mathcal{T}_2 = \mathcal{T}' \setminus \mathcal{T}_1,$$

where η_{j+1} is the constant given in Theorem 1.5.1 in $j+1$ dimensions. Note that all the members of \mathcal{T}_2 are $(\eta_{j+1}, j-1)$ -degenerate, that is, informally, they are ‘more degenerate’ than the other members of \mathcal{T}' , but with respect to a lower-dimensional flat. One of these two families contains at least half of the members of \mathcal{T}' . If $|\mathcal{T}_2| \geq |\mathcal{T}'|/2 \geq n/8$, we have, by induction on j ,

$$\begin{aligned} rs &\geq C_{\eta_{j+1},j-1} \cdot \min \left\{ \frac{I^{j+2}}{m^j n^{j+1}}, \frac{I^j}{m^{j-1} n^{j-1}} \right\} \\ &\geq C_{\eta_{j+1},j-1} \cdot \min \left\{ \frac{I^{j+3}}{m^{j+1} n^{j+2}}, \frac{I^{j+1}}{m^j n^j} \right\}, \end{aligned}$$

and the lemma holds in this case, provided that $C_{\eta,j} \leq C_{\eta_{j+1},j-1}$.

Suppose then that $|\mathcal{T}_1| \geq |\mathcal{T}'|/2 \geq n/8$. Put $\Pi = \{\pi_T \mid T \in \mathcal{T}_1\}$. Since the j -flats $\pi \in \Pi$ are $\frac{\eta I}{2n}$ -rich and η_{j+1} -nondegenerate with respect to P (in the space Q of projection), it follows from (3.3.3) that the number of these j -flats is upper-bounded by

$$\begin{aligned} |\Pi| &\leq E_j \cdot \max \left\{ \frac{m^{j+1}}{(\eta I/(2n))^{j+2}}, \frac{m^j}{(\eta I/(2n))^j} \right\} \\ &\leq A_{\eta,j} \cdot \max \left\{ \frac{m^{j+1}n^{j+2}}{I^{j+2}}, \frac{m^j n^j}{I^j} \right\}, \end{aligned}$$

for $A_{\eta,j} = E_j \left(\frac{2}{\eta}\right)^j \cdot \max \left\{ \frac{4}{\eta^2}, 1 \right\}$. Taking into account that $|\mathcal{T}_1| \geq n/8$, the average multiplicity of an element of Π is

$$\frac{|\mathcal{T}_1|}{|\Pi|} \geq \frac{1}{8A_{\eta,j}} \cdot \min \left\{ \frac{I^{j+2}}{m^{j+1}n^{j+1}}, \frac{I^j}{m^j n^{j-1}} \right\}.$$

Let $\pi \in \Pi$ be a j -flat with at least this multiplicity. Define $R = \pi \cap P$, and $\mathcal{S} = \{T \in \mathcal{T}_1 \mid \pi_T = \pi\}$. We have (i) $r = |R| \geq \frac{\eta I}{2n}$, (ii) $s = |\mathcal{S}| \geq |\mathcal{T}_1|/|\Pi|$, (iii) R is contained in every member of \mathcal{S} , and

$$rs \geq \frac{\eta I}{2n} \cdot \frac{|\mathcal{T}_1|}{|\Pi|} \geq \frac{\eta}{16A_{\eta,j}} \cdot \min \left\{ \frac{I^{j+3}}{m^{j+1}n^{j+2}}, \frac{I^{j+1}}{m^j n^j} \right\},$$

and the lemma holds in this case too for any $C_{\eta,j} \leq \frac{\eta}{16A_{\eta,j}}$.

Altogether, if we choose

$$C_{\eta,j} = \min \left\{ \frac{\eta}{16A_{\eta,j}}, C_{m_{j+1},j-1}, C_{\eta,j}^{(0)} \right\},$$

then the lemma holds in all cases. This completes the proof. \square

As a corollary, we obtain:

Theorem 3.3.4. *If $I > 2\hat{E}_{d-1} \max \left\{ (mn)^{1-\frac{1}{d+1}}, mn^{1-\frac{1}{d-1}} \right\}$, then*

$$rs_d(m, n, I) \geq 2^{-d-1} C_{\eta_d, d-2} \min \left\{ \frac{I^{d+1}}{m^{d-1}n^d}, \frac{I^{d-1}}{m^{d-2}n^{d-2}} \right\}.$$

Proof. Let P be a set of m points in \mathbb{R}^d and Π a set of n hyperplanes in \mathbb{R}^d , with $I = I(P, \Pi)$ in the assumed range. By (3.3.4), most incidences are with hyperplanes of Π that are η_d -degenerate with respect to P . We map each hyperplane $\pi \in \Pi$, which is η_d -degenerate, to the set $T_\pi = P \cap \pi$, and let \mathcal{T} be the multiset of all those T_π 's. This multiset has $n' < n$ elements, all of which are $(\eta_d, d-2)$ -degenerate, and $I' \geq I/2$ incidences. By Lemma 3.3.3, there are subsets $R \subseteq P$ and $\mathcal{S} \subseteq \mathcal{T}$, such that $R \subseteq S$ for each $S \in \mathcal{S}$ and

$$\begin{aligned} rs &\geq C_{\eta_d, d-2} \cdot \min \left\{ \frac{(I')^{d+1}}{m^{d-1}(n')^d}, \frac{(I')^{d-1}}{m^{d-2}(n')^{d-2}} \right\} \\ &\geq 2^{-d-1} C_{\eta_d, d-2} \cdot \min \left\{ \frac{I^{d+1}}{m^{d-1}n^d}, \frac{I^{d-1}}{m^{d-2}n^{d-2}} \right\}, \end{aligned}$$

where $r = |R|$ and $s = |S|$.

We map each member $S \in \mathcal{S}$ back to the hyperplane $\pi \in \Pi$ that satisfies $S = T_\pi$ (by the multiset structure of \mathcal{S} , this inverse mapping can be assumed to be well defined). We denote the resulting set of hyperplanes by Σ . Then $G(R, \Sigma) = K_{r,s}$ and rs has the asserted lower bound. This completes the proof. \square

We can now prove Theorem 3.1.2, which states that in the range $I \geq b_d(mn^{1-\frac{1}{d-1}} + nm^{1-\frac{1}{d-1}})$, for a suitable constant b_d , we have the lower bound

$$rs_d(m, n, I) \geq c_d \left(\frac{I}{mn} \right)^{d-1} mn,$$

for another suitable constant c_d . That is, in this range of m, n , and I , the minimum in the expression provided by Theorem 3.3.4 is attained by the second term.

Proof of Theorem 3.1.2. Put $b_d = 2\hat{E}_{d-1}$. Let P be a set of m points and Π a set of n hyperplanes in \mathbb{R}^d , with $I = I(P, \Pi) \geq b_d(mn^{1-\frac{1}{d-1}} + nm^{1-\frac{1}{d-1}})$ incidences. Then we have $I \geq 2\hat{E}_{d-1} \max \left\{ mn^{1-\frac{1}{d-1}}, nm^{1-\frac{1}{d-1}} \right\}$. This lower bound is larger than the one required in Theorem 3.3.4. Indeed, we have $(mn)^{1-\frac{1}{d+1}} \leq mn^{1-\frac{1}{d-1}}$ when $n \leq m^{(d-1)/2}$, and, symmetrically, $(mn)^{1-\frac{1}{d+1}} \leq nm^{1-\frac{1}{d-1}}$ when $m \leq n^{(d-1)/2}$; since $(d-1)/2 \geq 1$, at least one of the latter inequalities must hold. Therefore, we have in this range

$$rs_d(m, n, I) \geq 2^{-d-1} C_{\eta_d, d-2} \min \left\{ \frac{I^{d+1}}{m^{d-1}n^d}, \frac{I^{d-1}}{m^{d-2}n^{d-2}} \right\}. \quad (3.3.5)$$

Furthermore, we have

$$I \geq 2\hat{E}_{d-1} \max \left\{ mn^{1-\frac{1}{d-1}}, nm^{1-\frac{1}{d-1}} \right\} \geq 2\hat{E}_{d-1} nm^{1-\frac{1}{d-1}} \geq 2\hat{E}_{d-1} nm^{1/2},$$

which implies

$$I^2 \geq 4 \left(\hat{E}_{d-1} \right)^2 n^2 m.$$

When we substitute this inequality into (3.3.5), we get

$$\begin{aligned} rs_d(m, n, I) &\geq 2^{-d-1} C_{\eta_d, d-2} \min \left\{ 4 \left(\hat{E}_{d-1} \right)^2 \frac{I^{d-1}}{m^{d-2}n^{d-2}}, \frac{I^{d-1}}{m^{d-2}n^{d-2}} \right\} \\ &= 2^{-d-1} C_{\eta_d, d-2} \min \left\{ 4 \left(\hat{E}_{d-1} \right)^2, 1 \right\} \left(\frac{I}{mn} \right)^{d-1} mn. \end{aligned}$$

This gives us the asserted lower bound on $rs_d(m, n, I)$, e.g., for

$$c_d = 2^{-d-1} C_{\eta_d, d-2} \min \left\{ 4 \left(\hat{E}_{d-1} \right)^2, 1 \right\}.$$

\square

Next, we give an upper bound construction showing that

$$\text{rs}_d(m, n, I) = O\left(\left(\frac{I}{mn}\right)^{\frac{d+1}{2}} mn\right),$$

as asserted in Theorem 3.1.3.

Proof of Theorem 3.1.3. We start with the following d -dimensional structure, which is similar to constructions of Elekes [37]. For arbitrary integers $k, l > 0$, let $P_{d,k,l}$ and $\Pi_{d,k,l}$ denote the following respective sets of points and hyperplanes in \mathbb{R}^d :

$$P_{d,k,l} = \{(x_1, \dots, x_d) \mid x_1, \dots, x_{d-1} \in \{1, \dots, k\}, \text{ and } x_d \in \{1, \dots, dkl\}\},$$

$$\Pi_{d,k,l} = \left\{ x_d = \sum_{i=1}^{d-1} a_i x_i + b \mid a_1, \dots, a_{d-1} \in \{1, \dots, l\}, \text{ and } b \in \{1, \dots, kl\} \right\}.$$

Note that $|P_{d,k,l}| = dk^d l$ and $|\Pi_{d,k,l}| = kl^d$. For any hyperplane $\pi \in \Pi_{d,k,l}$, and for each choice of $x_1, \dots, x_{d-1} \in \{1, \dots, k\}$, there is a point $(x_1, \dots, x_{d-1}, x_d) \in P_{d,k,l} \cap \pi$. The set $P_{d,k,l} \cap \pi$ is thus a $(d-1)$ -lattice isomorphic to the hypercube $\{1, \dots, k\}^{d-1}$, and contains k^{d-1} points. Hence the number of incidences between $P_{d,k,l}$ and $\Pi_{d,k,l}$ is $I = k^d l^d$.

Each j -flat $F \subset \mathbb{R}^d$, which is the intersection of some $d-j$ or more hyperplanes of $\Pi_{d,k,l}$, is the image of some j -flat of the hypercube, as embedded into any of the hyperplanes $\pi \in \Pi_{d,k,l}$ that contain F . Since any j -flat of the hypercube contains at most k^j points, we have $|F \cap P_{d,k,l}| \leq k^j$. Furthermore, we have:

Observation 3.3.5. Any j -flat $F \subset \mathbb{R}^d$ (for $j < d$) is contained in at most l^{d-j-1} hyperplanes of $\Pi_{d,k,l}$.

Proof. F is the image of some affine mapping $T : \mathbb{R}^j \rightarrow \mathbb{R}^d$, that is, $T(y) = My + v$, for some matrix $M \in \mathbb{R}^{d \times j}$, with rank $\rho(M) = j$, and vector $v \in \mathbb{R}^d$.

Let $\pi \in \Pi_{d,k,l}$ be a hyperplane containing F , given by the linear equation $x_d = \sum_{i=1}^{d-1} a_i x_i + b$, for some $a_1, \dots, a_{d-1} \in \{1, \dots, l\}$ and $b \in \{1, \dots, kl\}$. Put $a_d = -1$, and $a = (a_1, \dots, a_d) \in \mathbb{R}^d$. Thus we can write $\pi = \{x \in \mathbb{R}^d \mid a^T x + b = 0\}$.

Since $\pi \supset F$, we have $a^T(My + v) + b = 0$ for all $y \in \mathbb{R}^j$. In particular, for $y = 0$, we have

$$a^T v + b = 0.$$

This gives $a^T My = 0$, for all $y \in \mathbb{R}^j$, which is equivalent to

$$M^T a = 0.$$

Thus, a is in the kernel of $M^T \in \mathbb{R}^{j \times d}$. We have

$$\dim \text{Ker}(M^T) = d - \dim \text{Im}(M^T) = d - \rho(M) = d - j.$$

Hence a lies in the $(d-j)$ -flat $K = \text{Ker}(M^T)$. In addition, the requirement $a_d = -1$ constrains a to a hyperplane H . Note that $H \not\supset K$, since $0 \in K$, but $0 \notin H$. Hence a lies in the $(d-j-1)$ -flat $K \cap H$. This flat can contain at most l^{d-j-1} points of the $l \times \cdots \times l \times 1$ lattice section. Hence there are at most l^{d-j-1} possible values of a . Once a has been determined, $b = -a^T v$ is also uniquely determined. Thus, there are at most l^{d-j-1} possible hyperplanes $\pi \in \Pi_{d,k,l}$ containing F , and the observation is established. \square

By adding another dimension to the construction, an x_{d+1} -axis, we turn every point of $P_{d,k,l}$ into a line parallel to the x_{d+1} -axis, and every $(d-1)$ -hyperplane of $\Pi_{d,k,l}$ into a d -hyperplane parallel to the x_{d+1} -axis. We denote the resulting set of lines by $P'_{d,k,l}$, and the set of d -hyperplanes by $\Pi'_{d,k,l}$. These sets have the same incidence relations as the original sets of points and $(d-1)$ -hyperplanes. In particular, every j -flat in \mathbb{R}^{d+1} , which is the intersection of some $d-j+1$ or more d -hyperplanes of $\Pi'_{d,k,l}$ contains at most k^{j-1} lines of $P'_{d,k,l}$ (all parallel to the x_{d+1} -axis), and is contained in at most l^{d-j} d -hyperplanes of $\Pi'_{d,k,l}$.

To construct an example that attains the asserted bound $rs = O\left(\left(\frac{I}{mn}\right)^{\frac{d+1}{2}} mn\right)$, we proceed as follows. Let $P' = P'_{d-2,k,k}$ and $\Pi' = \Pi'_{d-2,k,k}$ be sets of $(d-2)k^{d-1}$ lines and k^{d-1} $(d-2)$ -flats in \mathbb{R}^d (these lines and flats are constructed in \mathbb{R}^{d-1} , but we embed them in a natural way in \mathbb{R}^d). For every line $\ell \in P'$, choose μ arbitrary points on ℓ , and let P denote the overall resulting set of points, and put $m = |P| = (d-2)\mu k^{d-1}$. For every $(d-2)$ -flat $\pi' \in \Pi'$, choose ν distinct arbitrary hyperplanes, i.e., $(d-1)$ -flats, containing π' , and let Π denote the overall resulting set of hyperplanes. The hyperplanes are chosen so that no two hyperplanes containing two different flats from Π' coincide. Put $n = |\Pi| = \nu k^{d-1}$.

Now every hyperplane $\pi \in \Pi$ contains one flat $\pi' \in \Pi'$, which contains k^{d-3} lines of P' , yielding a total of μk^{d-3} points of P incident to π . The number of incidences between P and Π is thus $I = \mu \nu k^{2d-4} = \Theta(k^{-2}mn)$, or $\frac{I}{mn} = \Theta(k^{-2})$. Note that the freedom of choice of the parameters k, μ and ν allows I to have any order of magnitude from $\Theta((mn)^{1-\frac{1}{d-1}})$ (choose $\mu = \nu = 1$) up to $\Theta(mn)$ (choose $k = 1$). In particular, we may assume $I = \Omega(mn^{1-\frac{2}{d+1}} + m^{1-\frac{2}{d+1}}n)$. Suppose now that $G(P, \Pi)$ contains a $K_{r,s}$ subgraph, that is, there exists some j -flat F (for some $j = 1, \dots, d-2$) containing r points of P , and contained in s hyperplanes of Π . Without loss of generality, we may take F to be the intersection of these s hyperplanes. Thus, F is parallel to the x_{d+1} -axis, so any line of P' that meets F is fully contained in F . F contains at most k^{j-1} lines of P' , hence, $r \leq \mu k^{j-1}$. Also, F is contained in at most k^{d-j-2} flats of Π' , hence, $s \leq \nu k^{d-j-2}$. Altogether,

$$rs \leq \mu \nu k^{d-3} = \underbrace{\mu k^{d-1}}_{\approx m} \cdot \underbrace{\nu k^{d-1}}_{=n} \cdot \underbrace{k^{-d-1}}_{\approx \left(\frac{I}{mn}\right)^{\frac{d+1}{2}}} = O\left(\left(\frac{I}{mn}\right)^{\frac{d+1}{2}} mn\right),$$

as claimed. \square

We leave it as an open problem to close the gap between the bounds in Theorem 3.1.2 and Theorem 3.1.3, for $d \geq 4$.

3.4 Conclusion

We have studied the structure of point-hyperplane incidence graphs, and have shown that whenever the number of incidences is large, the incidence graph contains large complete bipartite subgraphs. Specifically,

1. We have derived lower bounds on the number of edges in the largest complete bipartite incidence subgraph in three dimensions (Theorems 3.1.1) and in higher dimensions (Theorem 3.1.2).
2. We have obtained matching upper bound constructions for these lower bounds. The three-dimensional constructions (Lemma 3.2.1, Lemma 3.2.2, and Lemma 3.2.4) are worst-case tight, whereas the higher-dimensional one (Theorem 3.1.3) is not known to be tight.
3. For each of these bounds, we have provided an estimate of how many incidences must there be in order to ensure the existence of large complete bipartite incidence subgraphs that attain the asserted lower bounds. The three-dimensional estimates are tight, whereas the higher-dimensional ones are not known to be tight.

We leave as an open problem to close the gap between the higher-dimensional bounds on the number of edges in the largest complete bipartite point-hyperplane incidence subgraph.

Chapter 4

Unit Area Triangles

In this chapter we show that the number of unit-area triangles determined by a set of n points in the plane is $O(n^{9/4})$, improving upon the bound $O(n^{44/19})$ of Dumitrescu et al. [33].

The results of this chapter were published, coauthored with Sharir, in *Discrete and Computational Geometry* [13]. We note that in the original publication [13], we derive a slightly weaker bound of $O^*(n^{9/4})$, but a very recent work of Zahl [84] has enabled us to get rid of the subpolynomial factor.

4.1 Introduction

In 1967, A. Oppenheim (see [48]) asked the following question: Given n points in the plane and $A > 0$, how many triangles spanned by the points can have area A ? By applying a scaling transformation, one may assume $A = 1$ and count the triangles of *unit* area. Erdős and Purdy [46] showed that a $n/\log^{1/2} n \times \log^{1/2} n$ section of the integer lattice determines $\Omega(n^2 \log \log n)$ triangles of the same area. They also showed that the maximum number of such triangles is at most $O(n^{5/2})$. In 1992, Pach and Sharir [68] improved the bound to $O(n^{7/3})$, using the Szemerédi-Trotter theorem [80] on the number of point-line incidences. Recently, Dumitrescu et al. [33] have further improved the upper bound to $O(n^{44/19}) = O(n^{2.3158})$, by estimating the number of incidences between the given points and a 4-parameter family of quadratic curves.

In this chapter we further improve the bound to $O(n^{9/4})$. Our proof borrows some ideas from [33], but works them into a different approach, which reduces the problem to bounding the number of incidences between points and a certain kind of surfaces in three dimensions.

Very recently, Zahl [84] has made new progress on bounding the number of incidences between points and algebraic surfaces in three dimensions. His results have enabled us to improve our bound from $O^*(n^{9/4})$ in the published paper [13] (coauthored with Sharir) to $O(n^{9/4})$ in this thesis.

4.2 Unit-area triangles in the plane

Theorem 4.2.1. *The number of unit-area triangles spanned by n points in the plane is $O(n^{9/4})$.*

Proof. We begin by borrowing some notation and preliminary ideas from [33]. Let S be the given set of n points in the plane. Consider a triangle $\Delta = \Delta abc$ spanned by S . We call the three lines containing the three sides of Δ , *base lines* of Δ , and the three lines parallel to the base lines and incident to the respective third vertices, *top lines* of Δ .

For a parameter k , $1 \leq k \leq \sqrt{n}$, to be optimized later, call a line ℓ *k-rich* (resp., *k-poor*) if ℓ contains at least k (resp., fewer than k) points of S . Call a triangle Δ *k-rich* if each of its three top lines is *k-rich*; otherwise Δ is *k-poor*.

We first observe that the number of *k-poor* unit-area triangles spanned by S is $O(n^2k)$. Indeed, assign a *k-poor* unit-area triangle Δabc whose top line through c is *k-poor* to the opposite base ab . Then all the triangles assigned to a base ab are such that their third vertex lies on one of the two lines parallel to ab at distance $2/|ab|$, where that line contains fewer than k points of S . Hence, a base ab can be assigned at most $2k$ triangles, and the bound follows.

So far, the analysis follows that of [33]. We now focus the analysis on the set of *k-rich* unit-area triangles spanned by S , and use a different approach.

Let L denote the set of *k-rich* lines, and let Q denote the set of all pairs

$$\{(\ell, p) \mid \ell \in L, p \in S \cap \ell\}.$$

By the Szemerédi-Trotter theorem [80] (see Theorem 2.1.4), we have, for any $k \leq \sqrt{n}$, $m := |L| = O(n^2/k^3)$, and $N := |Q| = O(n^2/k^2)$.

A pair $(\ell_1, p_1), (\ell_2, p_2)$ of elements of Q is said to *match* if the triangle with vertices $p_1, p_2, \ell_1 \cap \ell_2$ has area 1; see Figure 4.1.

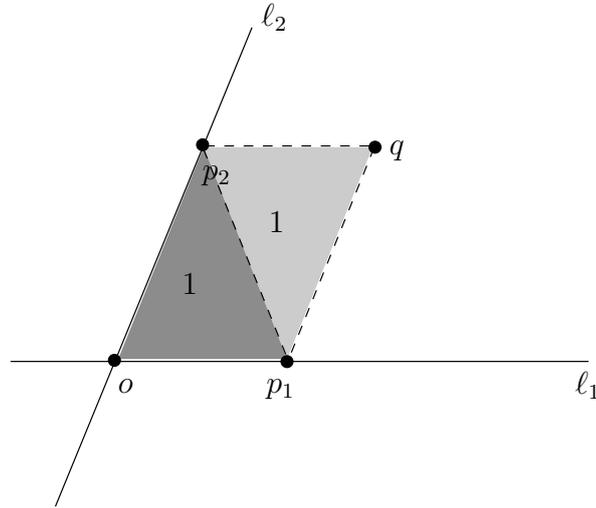


Figure 4.1: The ordered pair $((\ell_1, p_1), (\ell_2, p_2))$ is a matching pair of elements of Q .

To upper bound the number of unit-area triangles, all of whose three top lines are *k-rich*, it suffices to bound the number of matching pairs in Q . Indeed, given such a unit-area triangle $\Delta p_1 p_2 q$, let ℓ_1 (resp., ℓ_2) be the top line of $\Delta p_1 p_2 q$ through p_1 (resp., through p_2). Then (ℓ_1, p_1)

and (ℓ_2, p_2) form a matching pair in Q , by definition (again, see Figure 4.1). Conversely, a matching pair $(\ell_1, p_1), (\ell_2, p_2)$ determines at most one unit-area triangle p_1p_2q , where q is the intersection point of the line through p_1 parallel to ℓ_2 and the line through p_2 parallel to ℓ_1 ; we get an actual triangle if and only if the point q belongs to S .

In other words, our problem is now reduced to that of bounding the number of matching pairs in Q . (Since we do not enforce the condition that the third point q of the corresponding triangle belong to S , we most likely over-estimate the true bound.)

Since elements of Q have three degrees of freedom, we can represent them in an appropriate 3-dimensional parametric space. For example, we can assume that no line in L is vertical, and parametrize an element (ℓ, p) of Q by the triple (a, b, κ) , where (a, b) are the coordinates of p , and κ is the slope of ℓ . For simplicity of notation, we refer to this 3-dimensional parametric space as \mathbb{R}^3 .

So far, the matching relationship is symmetric. To simplify the analysis, and with no loss of generality, we make it asymmetric, by requiring that, in an (ordered) matching pair $(\ell_1, p_1), (\ell_2, p_2)$, \vec{op}_2 lies counterclockwise to \vec{op}_1 , where $o = \ell_1 \cap \ell_2$. See Figure 4.1.

Let us express the matching condition algebraically. Let $(a, b, \kappa) \in \mathbb{R}^3$ be the triple representing a pair (ℓ, p) , and $(x, y, w) \in \mathbb{R}^3$ be the triple representing another pair (ℓ', p') . Clearly, $w \neq \kappa$ in a matching pair. The lines ℓ and ℓ' intersect at a point o , for which there exist real parameters t, s which satisfy

$$o = (a + t, b + \kappa t) = (x + s, y + ws),$$

or

$$\begin{aligned} t &= \frac{y - b - w(x - a)}{\kappa - w} \\ s &= \frac{y - b - \kappa(x - a)}{\kappa - w}. \end{aligned}$$

It is now easy to verify that the condition of matching, with \vec{op}' lying counterclockwise to \vec{op} , is given by

$$\left(y - b - \kappa(x - a) \right) \left(y - b - w(x - a) \right) = 2(w - \kappa) \quad \text{and} \quad w \neq \kappa,$$

or, alternatively,

$$w = \frac{\left(y - b - \kappa(x - a) \right) (y - b) + 2\kappa}{\left(y - b - \kappa(x - a) \right) (x - a) + 2} \quad \text{and} \quad w \neq \kappa. \quad (4.2.1)$$

Similarly, the condition of “reverse” matching, with \vec{op}' lying clockwise to \vec{op} , is given by

$$w = \frac{\left(y - b - \kappa(x - a) \right) (y - b) - 2\kappa}{\left(y - b - \kappa(x - a) \right) (x - a) - 2} \quad \text{and} \quad w \neq \kappa. \quad (4.2.2)$$

Fix an element (ℓ, p) of Q , and associate with it a surface $\sigma_{\ell, p} \subset \mathbb{R}^3$, which is the locus of all pairs (ℓ', p') that match (ℓ, p) (i.e., $(\ell, p), (\ell', p')$ is an ordered matching pair). By the preceding analysis, $\sigma_{\ell, p}$ satisfies (4.2.1), where (a, b, κ) is the parametrization of (ℓ, p) , and is thus a 2-dimensional algebraic surface in \mathbb{R}^3 of degree 3. We thus obtain a system Σ of N 2-dimensional algebraic surfaces in \mathbb{R}^3 , and a set Q of N points in \mathbb{R}^3 , and our goal is to bound the number of incidences between Q and Σ .

The main technical step in the analysis is to rule out the possible existence of *degeneracies* in the incidence structure, where many points are incident to many surfaces; this might happen when many points lie on the common intersection of many surfaces (a situation which might arise, e.g., in the case of planes and points in \mathbb{R}^3). However, for the class of surfaces under consideration, namely, the surfaces $\sigma_{\ell, p}$ generated by some line-point incidence pair (ℓ, p) , such a degeneracy is impossible, because, as we will show, the incidence graph does not contain a $K_{3,10}$ incidence subgraph. We establish this fact using the following lemma.

Lemma 4.2.2. *Let (ℓ_1, p_1) and (ℓ_2, p_2) be two line-point incidence pairs, let $\gamma = \sigma_{\ell_1, p_1} \cap \sigma_{\ell_2, p_2}$ be the intersection of their associated surfaces, and assume that γ is non-empty. Then*

1. *The projection of γ onto the xy -plane is the zero set of an irreducible cubic polynomial.*
2. *Let (ℓ, p) be some incidence pair and assume further that $\sigma_{\ell, p} \supset \gamma$. Then either $(\ell, p) = (\ell_1, p_1)$ or $(\ell, p) = (\ell_2, p_2)$.*

Proof. We establish the equivalent claim to part 2, that, given a set γ , which is the intersection of some unknown pair of surfaces σ_{ℓ_1, p_1} and σ_{ℓ_2, p_2} , one can reconstruct (ℓ_1, p_1) and (ℓ_2, p_2) uniquely (up to a swap between the two incidence pairs) from γ . Moreover, it is enough to know the projection γ^* of γ onto the xy -plane in order to uniquely reconstruct the incidence pairs (ℓ_1, p_1) and (ℓ_2, p_2) that generated γ .

We start by computing the algebraic representation of γ^* . Let (a_1, b_1, κ_1) and (a_2, b_2, κ_2) be the respective parametrizations of (ℓ_1, p_1) and (ℓ_2, p_2) . By (4.2.1), γ^* satisfies the equation

$$\frac{\left(y - b_1 - \kappa_1(x - a_1)\right)(y - b_1) + 2\kappa_1}{\left(y - b_1 - \kappa_1(x - a_1)\right)(x - a_1) + 2} = \frac{\left(y - b_2 - \kappa_2(x - a_2)\right)(y - b_2) + 2\kappa_2}{\left(y - b_2 - \kappa_2(x - a_2)\right)(x - a_2) + 2}. \quad (4.2.3)$$

Recall the additional requirement in (4.2.1), namely that $w \neq \kappa_1$ and $w \neq \kappa_2$. This requirement is implicit in (4.2.1) and in (4.2.3), meaning that equation (4.2.3) is defined only for values of x and y for which the value of w is not any of κ_1 or κ_2 . Consulting (4.2.1), this implies that (x, y) cannot satisfy $y - b_1 = \kappa_1(x - a_1)$ or $y - b_2 = \kappa_2(x - a_2)$. Put

$$\begin{aligned} L_1 &= y - b_1 - \kappa_1(x - a_1), & \text{and} \\ L_2 &= y - b_2 - \kappa_2(x - a_2), \end{aligned}$$

and write (4.2.3) as

$$\frac{L_1(y - b_1) + 2\kappa_1}{L_1(x - a_1) + 2} = \frac{L_2(y - b_2) + 2\kappa_2}{L_2(x - a_2) + 2},$$

or

$$\begin{aligned} (L_1(y - b_1) + 2\kappa_1)(L_2(x - a_2) + 2) = \\ (L_2(y - b_2) + 2\kappa_2)(L_1(x - a_1) + 2), \end{aligned}$$

which we can rewrite as

$$L_1L_2L_3 + 2L_1L_4 - 2L_2L_5 + 4C = 0,$$

where

$$\begin{aligned} L_3 &= (b_2 - b_1)x - (a_2 - a_1)y + (a_2b_1 - a_1b_2), \\ L_4 &= y - b_1 - \kappa_2(x - a_1), \\ L_5 &= y - b_2 - \kappa_1(x - a_2), \\ C &= \kappa_1 - \kappa_2. \end{aligned}$$

We can further simplify the equation by noting that $L_6 = L_1L_4 - L_2L_5$ is a linear expression of x, y . That is,

$$L_6 = Dx + Ey + F,$$

where

$$\begin{aligned} D &= 2\kappa_1\kappa_2(a_2 - a_1) - (\kappa_1 + \kappa_2)(b_2 - b_1), \\ E &= 2(b_2 - b_1) - (\kappa_1 + \kappa_2)(a_2 - a_1), \\ F &= \kappa_1\kappa_2(a_1^2 - a_2^2) + (\kappa_1 + \kappa_2)(a_2b_2 - a_1b_1) + (b_1^2 - b_2^2). \end{aligned}$$

We can thus write (4.2.3) as

$$L_1L_2L_3 + 2L_6 + 4C = 0, \quad \text{and} \quad L_1 \neq 0, L_2 \neq 0. \quad (4.2.4)$$

Figure 4.2 illustrates the different lines defined by the linear equations $L_i = 0$, and their relations with (ℓ_1, p_1) and (ℓ_2, p_2) . The linearity of L_1, L_2, L_3 , and L_6 implies that the equation (4.2.4) of γ^* is cubic. We have the following two special cases to rule out:

1. If $p_1 = p_2$, that is, $a_1 = a_2$ and $b_1 = b_2$, then $L_3 = 0$, $L_4 = L_2$, and $L_5 = L_1$. But then the equation becomes $4C = 0$, so it has no solutions, meaning that γ is empty and the surfaces do not intersect.
2. If $\ell_1 = \ell_2$ but $p_1 \neq p_2$, that is, $\kappa_1 = \kappa_2 = (b_2 - b_1)/(a_2 - a_1)$, then $L_1 = L_2 = L_4 = L_5$, $L_3 = (a_1 - a_2)L_1$, and $C = 0$, resulting in the equation $(L_1)^3 = 0$, which is not allowed in (4.2.4). Hence γ is not defined in this case either.

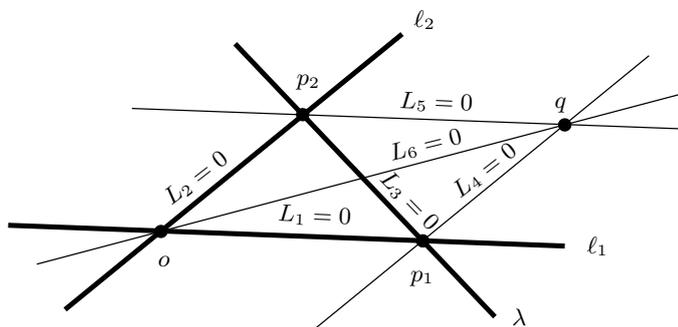


Figure 4.2: The lines $L_i = 0$, for $i = 1, \dots, 6$. The line $L_3 = 0$ connects p_1 and p_2 , $L_4 = 0$ passes through p_1 and is parallel to ℓ_2 , $L_5 = 0$ passes through p_2 and is parallel to ℓ_1 , and $L_6 = 0$ connects $o = \ell_1 \cap \ell_2$ with the intersection point q of $L_4 = 0$ and $L_5 = 0$.

We can therefore restrict our attention to the general case. Consider the cubic part of the equation $L_1L_2L_3$. In this term, each factor can be thought of as a line defined by the equation $L_i = 0$, for $i = 1, 2, 3$. The lines $L_1 = 0$ and $L_2 = 0$ respectively are simply ℓ_1 and ℓ_2 , whereas $L_3 = 0$ represents the line λ passing through p_1 and p_2 (see Figure 4.2). Note that λ may coincide with one of the other two lines. Indeed, if p_1 happens to be incident with ℓ_2 , then λ coincides with ℓ_2 . Similarly, if $p_2 \in \ell_1$ then λ coincides with ℓ_1 (these are the only possible coincidences, since we have ruled out the case $\ell_1 = \ell_2$). These cases will be handled shortly, but for now, we ignore them and consider the general case. In this case, γ^* has three distinct asymptotes given by $L_1 = 0$, $L_2 = 0$, and $L_3 = 0$; the proof of this fact is given in Lemma 4.3.3 in the appendix of this chapter.

Note that in either case, the resulting polynomial in (4.2.4), $f(x, y) = L_1L_2L_3 + 2L_6 + 4C$, is irreducible, as shown in Lemmas 4.3.3 and 4.3.4 in the appendix, which proves part 1 of the lemma.

Using the fact that γ^* has three distinct asymptotes, $L_1 = 0$, $L_2 = 0$, and $L_3 = 0$, one can reconstruct the two line-point pairs that generated γ^* as follows. Suppose we are given a curve γ^* generated by some unknown pair of incidence pairs, (ℓ_1, p_1) and (ℓ_2, p_2) , and we want to reconstruct these pairs. γ^* is given as the zero set of some cubic bivariate polynomial $f(x, y) = 0$, where f can be written as $f(x, y) = c(L_1L_2L_3 + 2L_6 + 4C)$, but the decomposition of f into L_1, L_2, L_3, L_6, C , and c is unknown, and, moreover, is not known a priori to be unique (a fact which we prove in this lemma). First, we find its three asymptotes $\Lambda_1 = 0$, $\Lambda_2 = 0$, and $\Lambda_3 = 0$, where for each $i = 1, 2, 3$, Λ_i is linear in x and y . Since, by Lemma 4.3.3, these asymptotes are $L_1 = 0$, $L_2 = 0$, and $L_3 = 0$, we know that each Λ_i is equal to some L_j multiplied by a constant, but we do not know which is which. To determine the roles of the asymptotes correctly, observe that $\Lambda_1\Lambda_2\Lambda_3 = \mu L_1L_2L_3$ for some constant μ . Thus, there exists some unique constant ν , such that $f(x, y) - \nu\Lambda_1\Lambda_2\Lambda_3 = \Lambda_4$ is linear in x and y . The line $\Lambda_4 = 0$ is parallel to the line $L_6 = 0$, which happens to be the median of the triangle spanned by the three asymptotes, which emanates from the vertex $o = \ell_1 \cap \ell_2$, and bisects the edge p_1p_2 ; see Figure 4.2. We thus have enough information to determine which vertex of the triangle is o , and which are p_1 and p_2 , and which edges of the triangle are supported by ℓ_1 and ℓ_2 . This proves part 2 of the lemma for the general case where all

the points and lines are distinct, and no point coincides with both lines.

Finally, consider the case where $p_2 \in \ell_1$ (a symmetric argument applies when $p_1 \in \ell_2$). In this case, $L_1 = L_5$, and $L_3 = (a_1 - a_2)L_1$, so the equation of the curve γ^* can be rewritten as

$$(a_1 - a_2)L_1^2L_2 + 2L_1(L_4 - L_2) + 4C = 0.$$

Note that $a_1 \neq a_2$ under the preliminary assumption that there are no vertical lines in the system, since both $p_1 = (a_1, b_1)$ and $p_2 = (a_2, b_2)$ are on ℓ_1 . Note also that $C = \kappa_1 - \kappa_2 \neq 0$, for otherwise, ℓ_1 and ℓ_2 would have to coincide, a case which we have ruled out earlier. Finally, note that $s = L_4 - L_2 = b_2 - b_1 - \kappa_2(a_2 - a_1) = C(a_2 - a_1)$ is a nonzero constant. Hence, the equation of γ^* is, up to a constant multiple,

$$(a_1 - a_2)L_1^2L_2 + 2sL_1 + 4C = 0. \quad (4.2.5)$$

This equation defines a cubic curve with two asymptotes given by $L_1 = 0$, and $L_2 = 0$, namely, the lines ℓ_1 and ℓ_2 ; the proof is given in Lemma 4.3.4 in the appendix. Since $C \neq 0$, it follows that γ^* does not intersect $L_1 = 0$, whereas $L_2 = 0$ is intersected at a single point (x, y) for which $L_1 = 2/(a_1 - a_2)$. Using this point, one can compute the values of $(a_1 - a_2)$, C , and s , and hence, reconstruct the line $L_4 = 0$. The point p_1 is then simply the intersection of the lines $L_1 = 0$ and $L_4 = 0$. Thus, one can uniquely reconstruct ℓ_1 , ℓ_2 , p_1 , and p_2 in this case too. This completes the proof of Lemma 4.2.2. \square

As a corollary of Lemma 4.2.2, we get, as claimed earlier, the following result.

Corollary 4.2.3. *In the incidence graph $G(Q, \Sigma)$ there is no $K_{3,10}$ subgraph.*

Proof. The intersection points of three distinct surfaces σ_{ℓ_i, p_i} , for $i = 1, 2, 3$, are the intersection points of the two curves $\gamma_{1,2} = \sigma_{\ell_1, p_1} \cap \sigma_{\ell_2, p_2}$, and $\gamma_{1,3} = \sigma_{\ell_1, p_1} \cap \sigma_{\ell_3, p_3}$. These intersection points project to the intersection points of the projections $\gamma_{1,2}^*$ and $\gamma_{1,3}^*$ of $\gamma_{1,2}$ and $\gamma_{1,3}$, respectively, onto the xy -plane. By Lemma 4.2.2, these two curves are distinct (or empty). Since each of them is cubic, and since, by Lemma 4.2.2, they are the zero sets of irreducible polynomials, Bézout's theorem [72] implies that they intersect in at most $3^2 = 9$ points. \square

Bounding the number of incidences. Recall that we need to bound the number of incidences between the set Σ of surfaces $\sigma_{\ell, p}$, for $(\ell, p) \in Q$, and the set Q of points. This is done using the following result of Zahl [84]:

Theorem 4.2.4 (Zahl [84]). *Let $t \geq 3$ and C be some constants, let Q be a set of m points in \mathbb{R}^3 , and let Σ be a set of n algebraic surfaces in \mathbb{R}^3 of bounded degree. Assume further, that, for any three distinct surfaces $\sigma, \sigma', \sigma'' \in \Sigma$, we have $|\sigma \cap \sigma' \cap \sigma''| \leq C$, and for any set of t points in \mathbb{R}^3 , there are at most C surfaces of Σ that contain all t points. Then the number of incidences between the points in Q and the surfaces in Σ is*

$$O\left(m^{\frac{2t}{3t-1}}n^{\frac{3t-3}{3t-1}} + m + n\right),$$

where the implicit constant depends only on t , C , and the degree of the algebraic surfaces.

In our case, the points of Q and the surfaces of Σ satisfy the assumptions of Theorem 4.2.4 with $t = 3$ and $C = 9$. Indeed, by Corollary 4.2.3, the intersection of any three surfaces contains at most $C = 9$ points, and, Symmetrically, for any $t = 3$ points, there are at most nine surfaces of Σ that contain all the three. Thus, by Theorem 4.2.4, the number of incidences between the N points of Q and the N surfaces of Σ is $O(N^{3/4} \cdot N^{3/4}) = O(N^{3/2})$.

Recall now that $N = O(n^2/k^2)$, and that we also have the bound $O(n^2k)$ for the number of unit-area triangles with at least one k -poor top line. Thus, the overall bound on the number of unit-area triangles is

$$O\left(\frac{n^3}{k^3} + n^2k\right),$$

which, if we choose $k = n^{1/4}$, becomes $O(n^{9/4})$, as asserted. \square

Discussion. In the published paper [13] (coauthored with Sharir), we bounded $I(Q, \Sigma)$ using cuttings in \mathbb{R}^3 , which in turn were based on vertical decomposition of the arrangement of (a sample of) the surfaces in Σ . Since the complexity of the vertical decomposition of such surfaces involves an extra subpolynomial factor, our analysis gave the slightly weaker bound of $O^*(n^{9/4})$. Zahl's new result has allowed us to get rid of this subpolynomial factor.

Theorem 4.2.1 constitutes a major improvement over previous bounds, but it still leaves a substantial gap from the near-quadratic lower bound. One major weakness of our proof is that, in bounding the number of matching pairs, it ignores the constraint that a matching pair is relevant only when the (uniquely defined) third vertex q of the resulting triangle belongs to S , and that the (uniquely defined) top line of this triangle through q is k -rich. Moreover, even without the requirement that the third point-line pair exist in the system, we do not know whether the bound $O(n^{9/4})$ on the number of matching pairs is tight. It is therefore natural to conjecture that our bound is not tight, and that the true bound is nearly quadratic, perhaps coinciding with the lower bound of [46].

4.3 Appendix: Asymptotes of cubic curves

In this appendix, we analyze the class of cubic curves defined by equation (4.2.4) and its special case (4.2.5) of Section 4.2, derive their asymptotes, and show them to be the zero sets of irreducible bivariate polynomials. We start by analyzing a normalized version of these equations, in which two of the generating lines (and, as we show henceforth, the asymptotes) are the x and y -axes. We then reduce equation (4.2.4) to the normalized case. Finally, we handle equation (4.2.5) in a different and simpler way.

Lemma 4.3.1. *Let λ_1 and λ_2 be two distinct lines in \mathbb{R}^2 , given by the equations $\Lambda_i = 0$, where $\Lambda_i = \alpha_i x + \beta_i y + \gamma_i$, and α_i and β_i are both nonzero, for $i = 1, 2$. Let $f(x, y)$ be the bivariate cubic polynomial*

$$f(x, y) = xy\Lambda_1 + \Lambda_2.$$

Then $f(x, y)$ is irreducible.

Proof. Assume, to the contrary, that f is reducible. Then it has a linear factor $L = ax + by + c$. Without loss of generality, $b \neq 0$ (a symmetric argument follows for the case $a \neq 0$), so we can assume $b = 1$. Then f , as a polynomial in y with coefficients from $\mathbb{R}[x]$, has $y = -ax - c$ as root, i.e., if we put

$$p(x) := f(x, -ax - c) = -x(ax + c)L_1 + L_2,$$

where $L_i = \alpha_i x - \beta_i(ax + c) + \gamma_i$, for $i = 1, 2$, then $p(x) \equiv 0$. But then, the term $x(ax + c)L_1$ can not be properly cubic, nor quadratic, so, $a = c = 0$, or L_1 is a constant, possibly zero. In the former case, $L_2 \equiv 0$, but $L_2 = \alpha_2 x + \gamma_2$ and $\alpha_2 \neq 0$ by assumption, a contradiction. If $L_1 = 0$, then we must also have $L_2 = 0$, and so both lines λ_1 and λ_2 coincide (with the line $L = 0$), in contradiction. If L_1 is a nonzero constant, then the term $(ax + c)$ must also be constant, or else $p(x)$ is a proper quadratic polynomial, hence $a = 0$. But then, for $L_1 = \alpha_1 x - \beta_1 c + \gamma_1$ to be constant, we must have $\alpha_1 = 0$, in contradiction. Either way, f cannot be reducible. \square

Lemma 4.3.2. *Let λ_1 and λ_2 be two lines in \mathbb{R}^2 , given by the equations $\Lambda_i = 0$, where $\Lambda_i = \alpha_i x + \beta_i y + \gamma_i$, for $i = 1, 2$, such that α_1 and β_1 are both nonzero. Let Γ be the algebraic cubic curve defined by the equation*

$$xy\Lambda_1 + \Lambda_2 = 0. \tag{4.3.6}$$

Then Γ is asymptotic to the x -axis, and, symmetrically, to the y -axis.

Proof. We only prove in detail that the x -axis is an asymptote. Note that, for any fixed $x \neq 0$, (4.3.6) is a quadratic equation in y , which we rewrite as

$$xy(\alpha_1 x + \beta_1 y + \gamma_1) + \alpha_2 x + \beta_2 y + \gamma_2 = 0,$$

or

$$\beta_1 x y^2 + (\alpha_1 x^2 + \gamma_1 x + \beta_2)y + (\alpha_2 x + \gamma_2) = 0.$$

Hence

$$y = -\frac{\alpha_1 x^2 + \gamma_1 x + \beta_2}{2\beta_1 x} \pm \frac{\sqrt{(\alpha_1 x^2 + \gamma_1 x + \beta_2)^2 - 4\beta_1 x(\alpha_2 x + \gamma_2)}}{2\beta_1 x}$$

We only consider the solution with positive square root, which is

$$y = \frac{\alpha_1 x^2 + \gamma_1 x + \beta_2}{2\beta_1 x} \left[\sqrt{1 - \frac{4\beta_1 x(\alpha_2 x + \gamma_2)}{(\alpha_1 x^2 + \gamma_1 x + \beta_2)^2}} - 1 \right].$$

The expression in the square brackets is of the form $\sqrt{1+t} - 1$. Since $\alpha_1 \neq 0$, t tends to 0 as $x \rightarrow \pm\infty$. Using the inequalities $1 - |t| \leq \sqrt{1+t} \leq 1 + \frac{|t|}{2}$, for $|t| < 1$, we obtain, for $|x|$ sufficiently large,

$$|y| \leq \frac{|\alpha_1 x^2 + \gamma_1 x + \beta_2|}{|2\beta_1 x|} |t| = \frac{2|\alpha_2 x + \gamma_2|}{|\alpha_1 x^2 + \gamma_1 x + \beta_2|}, \tag{4.3.7}$$

which tends to 0 as $x \rightarrow \pm\infty$. This shows that the x -axis is indeed an asymptote of Γ (on both sides). A symmetric argument shows that the y -axis is also an asymptote. \square

Note that if $\alpha_2 = 0$ and $\gamma_2 = 0$, then, by (4.3.7), $y = 0$, i.e., the curve Γ coincides with the x -axis. But then λ_2 coincides with the x -axis as well. In all other cases, Γ is asymptotic to the x -axis, but does not coincide with it.

We are now ready to prove the more general cases discussed in Section 4.2.

Lemma 4.3.3. *Let ℓ_1, \dots, ℓ_4 be four distinct lines in \mathbb{R}^2 , given by the equations $L_i = 0$, where $L_i = A_i x + B_i y + C_i$, for $i = 1, \dots, 4$. Assume that no pair of ℓ_1, ℓ_2, ℓ_3 are parallel. Put*

$$f(x, y) = L_1 L_2 L_3 + L_4,$$

and let Γ be the algebraic cubic curve defined by the equation

$$f(x, y) = 0.$$

Then, f is irreducible, and Γ is asymptotic to the lines ℓ_1, ℓ_2, ℓ_3 .

Proof. We may assume, by an appropriate change of variables, that one of ℓ_1, ℓ_2 , and ℓ_3 is the x -axis and another one is the y -axis. For example, put $u = L_1$, and $v = L_2$, and write $L_3 = \alpha_1 u + \beta_1 v + \gamma_1$, and $L_4 = \alpha_2 u + \beta_2 v + \gamma_2$, for some appropriate coefficients $\alpha_1, \beta_1, \gamma_1, \alpha_2, \beta_2, \gamma_2$. Note that, by the preliminary assumptions on the lines, α_1 and β_1 are both nonzero. Γ can then be written as

$$g(u, v) = uvL_3 + L_4 = 0$$

in the (u, v) coordinate system. It then follows, by Lemma 4.3.2, that ℓ_1 and ℓ_2 are asymptotes of Γ . Note that the choice of ℓ_1 and ℓ_2 as axes is arbitrary, and we could just as well choose any other pair of lines among $\{\ell_1, \ell_2, \ell_3\}$ in any order. Hence, ℓ_3 is also an asymptote of Γ . To prove irreducibility, we assume, without loss of generality, that ℓ_4 is not parallel to any of ℓ_1 and ℓ_2 , i.e., to the u - and v -axes (otherwise, we exchange the roles of the parallel axis and ℓ_3). Then, by Lemma 4.3.1, g is irreducible. But then f is irreducible too, for otherwise, any factorization of f could be transformed into a factorization of g , in contradiction. \square

Lemma 4.3.4. *Let ℓ_1 and ℓ_2 be two distinct intersecting lines in \mathbb{R}^2 , given by the equations $L_i = 0$, where $L_i = A_i x + B_i y + C_i$, for $i = 1, 2$. Put $f(x, y) = L_1^2 L_2 + L_1 + C$, for some constant C , and let Γ be the algebraic curve defined by the equation*

$$f(x, y) = 0.$$

Then Γ is asymptotic to the lines ℓ_1 and ℓ_2 . Furthermore, if $C \neq 0$, then f is an irreducible bivariate polynomial.

Proof. If $C = 0$, then the claim is easy. Indeed, in this case we have $L_1(L_1 L_2 + 1) = 0$, so Γ is the union of the line $L_1 = 0$ and the hyperbola $L_1 L_2 = -1$, which is asymptotic to the lines $L_1 = 0$, and $L_2 = 0$.

If $C \neq 0$, put $u = L_1$, and $v = L_2$. Then, in the (u, v) coordinate system, Γ is defined by the equation

$$g(u, v) := u^2 v + u + C = 0.$$

Note that g is clearly irreducible, and so is f . This equation can be rewritten as

$$v = -\frac{u + C}{u^2}.$$

Clearly, this function tends to 0 as u tends to ∞ , which means it is asymptotic to the u -axis, i.e., to ℓ_2 . Furthermore, the function has a pole at $u = 0$, meaning it is asymptotic to the v -axis, i.e., to ℓ_1 . □

Chapter 5

Nondegenerate Spheres

In this chapter we study incidences between points and spheres in three dimensions and obtain an improved bound for their number. The results in this chapter were published, coauthored with Sharir, in *Combinatorics, Probability and Computing* [14].

5.1 Introduction

Degeneracy and nondegeneracy in spheres is analogous to degeneracy and nondegeneracy in hyperplanes, as described in Section 1.5.3. Given a finite point set $P \subset \mathbb{R}^d$ and a constant $0 < \eta < 1$, a $(d - 1)$ -sphere $\sigma \subset \mathbb{R}^d$ is called η -degenerate (with respect to P) if there exists some $(d - 2)$ -subsphere $\sigma' \subset \sigma$ such that

$$|\sigma' \cap P| \geq \eta |\sigma \cap P|.$$

Otherwise, σ is called η -nondegenerate.

The issue of degeneracy and nondegeneracy in spheres arises in several problems involving distances between points. For example, the study of Aronov et al. [16] on the number of distinct distances determined by n points in \mathbb{R}^3 involves incidences between the points and a certain collection of spheres, and faces the issue of degeneracy of these spheres. Another study by Agarwal, Apfelbaum, Purdy, and Sharir [4] considers the problem of bounding the number of similar triangles (or simplices) determined by n points in d dimensions. This study is described in detail in Chapter 6. Here too the problem is reduced to incidences between points and spheres, and handling degenerate and nondegenerate spheres is a major step in the analysis.

By lifting \mathbb{R}^d to the standard paraboloid in \mathbb{R}^{d+1} (see, e.g., [34]), every sphere is transformed into a hyperplane, and the incidence relation, as well as degeneracy and nondegeneracy, are preserved. It thus follows, by Theorem 1.5.1 that for any $\eta < \eta_{d+1}$, where $\{\eta_d\}_{d=3}^\infty$ are the same constants specified in Theorem 1.5.1, the number of k -rich η -nondegenerate $(d - 1)$ -spheres, with respect to a set of n points in \mathbb{R}^d is

$$O\left(\frac{n^{d+1}}{k^{d+2}} + \frac{n^d}{k^d}\right). \quad (5.1.1)$$

See [4] for more details.

It has been conjectured (see [4]) that this bound is not tight. A supporting evidence comes from the fact that in the plane, where the spheres are circles, which are clearly nondegenerate (if they are k -rich and k is sufficiently large), we have an upper bound of $O^*(n^3/k^{11/2} + n^2/k^3 + n/k)$ on the number of k -rich circles [6, 17, 63] (see (1.2.2) in Chapter 1), which is significantly better than the $O(n^3/k^4 + n^2/k^2)$ bound implied by lifting the circles into nondegenerate planes in \mathbb{R}^3 . We recall that this improved bound holds for circles in any dimension; see [15].

At any rate, the bound for spheres in \mathbb{R}^d should lie in between those for hyperplanes in \mathbb{R}^d and for hyperplanes in \mathbb{R}^{d+1} . The second threshold follows from the lifting transform just discussed, and the first threshold follows from noting that an *inversion* of \mathbb{R}^d takes a collection of hyperplanes in \mathbb{R}^d into a collection of spheres (all passing through the point of inversion), while preserving incidences, richness, degeneracy, and nondegeneracy. In particular, we obtain the lower bound (the best known, as far as we are aware)

$$\Omega\left(\frac{n^d}{k^{d+1}} + \frac{n^{d-1}}{k^{d-1}}\right)$$

on the number of k -rich η -nondegenerate spheres in \mathbb{R}^d , by taking the Elekes-Tóth lower bound construction for nondegenerate hyperplanes in \mathbb{R}^d [41], and by inverting space.

Narrowing this gap between the upper and lower bounds is an interesting problem in its own right. Moreover, any improvement of the upper bound will immediately improve the bounds on the problems studied in the papers [4, 16] mentioned above.

5.2 Our results

In this chapter we improve the upper bound on the number of rich nondegenerate spheres in 3-space. We show:

Theorem 5.2.1. *For any constant $\eta < 1$, any arbitrarily small $\varepsilon > 0$, any set P of n points in \mathbb{R}^3 , and any $k \leq n$, the number m of k -rich η -nondegenerate spheres is*

$$m = O\left(\frac{n^4}{k^{11/2-\varepsilon}} + \frac{n^2}{k^2}\right),$$

with the multiplicative constant depending on η and ε .

Clearly, this is an improvement over the previous bound of $O(n^4/k^5 + n^3/k^3)$. Using this bound, we get an improved upper bound on the number of mutually similar triangles spanned by n points in \mathbb{R}^3 . Previously, it was known that the number of such triangles is at most $O(n^{13/6}) = O(n^{2.1667})$ [4], but using Theorem 5.2.1, the bound becomes $O^*(n^{58/27}) = O(n^{2.1482})$. This bound is established in Chapter 6 of the thesis.

Discussion. There are three conventional ways to present incidence bounds:

- (a) An upper bound on the number of incidences between n points and m objects. For example, see Equation (2.1.1) of the Szemerédi-Trotter theorem (Theorem 2.1.4). As another example, the Elekes-Tóth bound [41] (Theorem 1.5.1), when formulated this way (1.5.5), asserts that the number of incidences between n points in \mathbb{R}^d and m η_d -nondegenerate hyperplanes is $O(n^{d/(d+1)}m^{d/(d+1)} + nm^{(d-2)/(d-1)})$.
- (b) An upper bound on the number of objects incident to at least k out of n points. This, for example, is the formulation of (2.1.2) in Theorem 2.1.4, as well as (1.5.5) in Theorem 1.5.1, which says that the number of k -rich η -nondegenerate hyperplanes spanned by a set of n points in \mathbb{R}^d is $O(n^d/k^{d+1} + n^{d-1}/k^{d-1})$.
- (c) An upper bound on the number of incidences between n points and objects incident to at least k of these points. For example, see Equation (2.1.3) in Theorem 2.1.4. Theorem 1.5.1 can be reformulated this way to say that the number of incidences between n points in \mathbb{R}^d and any family of k -rich η_d -nondegenerate hyperplanes is $O(n^d/k^d + n^{d-1}/k^{d-2})$.

In practically all known instances, all three alternatives are equivalent and any one of them can be derived from any other. Theorem 5.2.1, when stated in the first form (a), reads as follows.

Theorem 5.2.2. *For any $\eta < 1$, and any arbitrarily small $\varepsilon > 0$, the number of incidences I between n points in \mathbb{R}^3 and m η -nondegenerate spheres is*

$$I = O\left(m\left(\frac{n^4}{m}\right)^{2/11+\varepsilon} + nm^{1/2}\right) = O\left(n^{8/11+4\varepsilon}m^{9/11-\varepsilon} + nm^{1/2}\right),$$

with the multiplicative constant depending on η and ε .

We will prove Theorem 5.2.2 rather than Theorem 5.2.1. Theorem 5.2.1 follows by noting that if each of the m spheres is also k -rich then $I \geq mk$. Solving the resulting inequality $mk = O(n^{8/11+4\varepsilon}m^{9/11-\varepsilon} + nm^{1/2})$ for m yields the bound of Theorem 5.2.1 (with a modified value of ε).

We use in the proof the previous bound for the number of incidences between n point in \mathbb{R}^3 and m nondegenerate spheres, and strengthen it by a *cutting* (See Section 2.1.1 for more details about this technique). The bound is

$$I = O\left(m\left(\frac{n^4}{m}\right)^{1/5} + nm^{2/3}\right) = O\left(n^{4/5}m^{4/5} + nm^{2/3}\right), \quad (5.2.2)$$

which is the alternative form (a) of the Elekes-Tóth bound (Theorem 1.5.1) in four dimensions (where the given spheres appear as hyperplanes).

5.3 Proof of Theorem 5.2.2

Let P be a set of n points in \mathbb{R}^3 , and let \mathcal{S} be a set of m η -nondegenerate spheres in \mathbb{R}^3 , for some fixed positive $\eta < 1$. Let $j > 0$ be some integer. We say that a sphere σ is j -*bad* with respect to P

if σ contains some circle $\gamma \subset \sigma$, which contains at least j points of P , i.e., $|\gamma \cap P| \geq j$. If no such circle exists, then σ is said to be j -good. Note that an η -nondegenerate sphere incident to exactly k points is, by definition, ηk -good.

Let $I(P, \mathcal{S})$ denote the number of incidences between P and \mathcal{S} . Denote by $I_\eta(n, m)$, or $I(n, m)$ for short, the maximum of $I(P, \mathcal{S})$ over all sets P of n points and \mathcal{S} of m η -nondegenerate spheres in \mathbb{R}^3 . Note that each sphere $\sigma \in \mathcal{S}$ contains some noncoplanar quadruple of points, and each such quadruple uniquely determines the sphere containing it, so we have $m = |\mathcal{S}| \leq \binom{n}{4}$. In fact, this bound decreases as η decreases, but it remains $m \leq C_\eta^1 n^4$, with the constant of proportionality C_η^1 decreasing with η . Therefore, $I(n, m)$ is defined only for values of n and m satisfying (the appropriate variant of) this relationship.

The next lemma establishes a recurrence relation on $I(n, m)$.

Lemma 5.3.1. *For any $0 < \eta < 1$, for any two positive integers n, m satisfying $m \leq C_\eta^1 n^4$ as above, for any integer $1 < j < n$, and for any number $1 < r < \min\{m, n^{1/3}\}$, we have*

$$I(n, m) \leq C_\eta^2 \left(r^3 \beta(r) \log^3 r \cdot I\left(\frac{n}{r^3}, \frac{m}{r}\right) + m j r^2 \beta(r) \log^3 r + \frac{n^4 \log j}{j^{11/2}} + \frac{n^3}{j^3} + \frac{n^2}{j} \right), \quad (5.3.3)$$

where $C_\eta^2 > 0$ depends only on η , and where $\beta(r) = 2^{C\alpha^2(r)}$, $C > 0$ is some constant, and $\alpha(r)$ is the extremely slowly growing inverse Ackermann function.

Proof. Let P be a set of n points in \mathbb{R}^3 , and let \mathcal{S} be a set of m η -nondegenerate spheres with respect to P . We partition \mathcal{S} into two subsets consisting respectively of the j -good and the j -bad spheres, and bound the incidences of P with each of these subsets separately.

We bound the number of incidences with the j -bad spheres as follows. The number of j -rich circles with respect to P is at most $C_{\text{circ}}((n^3 \log j)/j^{11/2} + n^2/j^3 + n/j)$ for some constant $C_{\text{circ}} > 0$ (see (1.2.2)). Each circle γ can be incident to several spheres of \mathcal{S} . Since each such sphere σ is η -nondegenerate, we have $|\sigma \cap P| > \frac{1}{\eta} |\gamma \cap P|$, or $|\sigma \cap P| < \frac{1}{1-\eta} |(\sigma \setminus \gamma) \cap P|$. Any point of P not on γ can lie in at most one of these spheres, so

$$\sum_{\sigma \in \mathcal{S} | \sigma \supset \gamma} |\sigma \cap P| < \sum_{\sigma \in \mathcal{S} | \sigma \supset \gamma} \frac{1}{1-\eta} |(\sigma \setminus \gamma) \cap P| < \frac{n}{1-\eta}. \quad (5.3.4)$$

Hence, the total number of incidences between the spheres containing γ and P is $\frac{n}{1-\eta}$. Multiplying by the number of circles, we get that the number of incidences between P and the j -bad spheres is

$$I_{\text{bad}} = \frac{C_{\text{circ}}}{1-\eta} \cdot \left(\frac{n^4 \log j}{j^{11/2}} + \frac{n^3}{j^3} + \frac{n^2}{j} \right). \quad (5.3.5)$$

To bound the number of incidences with the j -good spheres, we construct a $\frac{1}{r}$ -cutting of \mathcal{S} (see Section 2.1.1). We will use the simpler version of cutting, which is similar to the one used in Theorem 2.1.1, but with spheres instead of hyperplanes. In this type of cutting, we sample spheres of \mathcal{S} by choosing each sphere in the sample independently at random with probability $\frac{C_{\text{cut}} r}{n} \log r$, for an appropriate sufficiently large constant C_{cut} , and then construct the vertical decomposition of

the arrangement of the random sample (see, e.g., [27,31]). As follows from the analysis of [31], the result is a set of an expected number of $C_{\text{cell}}r^3\beta(r)\log^3 r$ relatively open cells, for some constant $C_{\text{cell}} > 0$ which depends on C_{cut} , each cell having constant description complexity, such that, with high probability, each cell is *crossed* by (i.e., is intersected by but not contained in) at most m/r spheres of \mathcal{S} . We will assume that our cutting does satisfy its expected properties up to a constant factor, i.e., it consists of at most $2C_{\text{cut}}r\log r$ sample spheres, at most $2C_{\text{cell}}r^3\beta(r)\log^3 r$ cells, and each cell is crossed by at most m/r spheres of \mathcal{S} . We may also assume that each cell contains at most n/r^3 points of P . To enforce this, we take each cell that contains tn/r^3 points, for any $t > 1$, and partition it into $\lceil t \rceil$ subcells in some generic way, such that each subcell contains at most n/r^3 points. The number of cells in this refined cutting is at most $2C'_{\text{cell}}r^3\beta(r)\log^3 r$, for some other constant $C'_{\text{cell}} > C_{\text{cell}}$.

For each cell τ separately, we bound the number of incidences between the points of $P \cap \tau$ and the spheres that cross τ , and then sum up the bounds over all cells. A cell τ has at most n/r^3 points and is crossed by at most m/r spheres. If a sphere contains more than j/η points of $P \cap \tau$, then, since it is a j -good sphere, it is locally η -nondegenerate in τ (i.e., with respect to the points of $P \cap \tau$). The number of incidences between such spheres and the points of $P \cap \tau$ is at most $I(\frac{n}{r^3}, \frac{m}{r})$. The number of incidences with the other spheres in τ is at most $mj/(\eta r)$. The total, summed over all cells, is

$$\sum_{\tau} \left(I\left(\frac{n}{r^3}, \frac{m}{r}\right) + \frac{mj}{\eta r} \right) \leq 2C'_{\text{cell}}r^3\beta(r)\log^3 r \cdot I\left(\frac{n}{r^3}, \frac{m}{r}\right) + \frac{2C'_{\text{cell}}}{\eta}mj r^2\beta(r)\log^3 r. \quad (5.3.6)$$

It remains to bound, for each cell τ , the number of incidences between the points of $P \cap \tau$ and the spheres of \mathcal{S} that contain τ , which we do separately for cells of dimensions 2,1 and 0. (Cells of dimension 3 are “exempt” from this analysis.)

Any two-dimensional cell can be contained by at most one sphere, so these cells contribute a total of at most n incidences of this kind.

For one-dimensional cells, we can restrict our attention only to these cells which are (portions of) edges in the arrangement of the sample spheres. (By choosing a generic coordinate frame, the other edges of the vertical decomposition will not contain points of P .) Each of the m spheres intersects the at most $2C_{\text{cut}}r\log r$ spheres of the sample along at most $2C_{\text{cut}}r\log r$ circles, and each such circle contains at most j points of P . Thus, the number of incidences contributed by the one-dimensional cells is at most $2C_{\text{cut}}mjr\log r$.

As for zero-dimensional cells, which are simply the vertices of the arrangement of the sample spheres (as above, we may ignore the additional vertices created by the vertical decomposition), each of the m spheres σ contains at most $(2C_{\text{cut}}r\log r)^2$ such vertices, which are intersection points of the at most $2C_{\text{cut}}r\log r$ circles of intersection of σ with the sample spheres. Hence, the total contribution of these cells is at most $(2C_{\text{cut}})^2mr^2\log^2 r$ incidences.

Thus, the total number of incidences, over all cells τ , between the points of $P \cap \tau$ and the spheres of \mathcal{S} that contain τ is at most $2C_{\text{cut}}mjr\log r + (2C_{\text{cut}})^2mr^2\log^2 r + n$, which we combine

with (5.3.6) to get the total number of incidences with the good spheres:

$$\begin{aligned}
I_{\text{good}} &\leq 2C'_{\text{cell}}r^3\beta(r)\log^3 r \cdot I\left(\frac{n}{r^3}, \frac{m}{r}\right) + \frac{2C'_{\text{cell}}}{\eta}mjr^2\beta(r)\log^3 r \\
&\quad + 2C_{\text{cut}}mjr\log r + (2C_{\text{cut}})^2mr^2\log^2 r + n \\
&\leq 2C'_{\text{cell}}r^3\beta(r)\log^3 r \cdot I\left(\frac{n}{r^3}, \frac{m}{r}\right) + C_{\eta}^3mjr^2\beta(r)\log^3 r + n, \tag{5.3.7}
\end{aligned}$$

where $C_{\eta}^3 = \frac{2C'_{\text{cell}}}{\eta} + 2C_{\text{cut}} + (2C_{\text{cut}})^2$. Summing both inequalities (5.3.5) and (5.3.7), we obtain the bound asserted in the lemma, e.g., with

$$C_{\eta}^2 = \max\left\{2C'_{\text{cell}}, C_{\eta}^3, \frac{2C_{\text{circ}}}{1-\eta}\right\}.$$

□

Next, we simplify the recurrence by eliminating j . We have, on one hand, the term $mjr^2\beta(r)\log^3 r$, which increases with j , and, on the other hand, the terms

$$\frac{n^4\log j}{j^{11/2}} + \frac{n^3}{j^3} + \frac{n^2}{j},$$

which decrease with j , so we choose j so as to balance between them. By comparing the increasing term with each of the decreasing terms (but ignoring the logarithmic factors), this leads to the choice

$$j = \max\left\{\frac{n^{8/13}}{m^{2/13}r^{4/13}}, \frac{n^{3/4}}{m^{1/4}r^{1/2}}, \frac{n}{m^{1/2}r}\right\},$$

which implies

$$I(n, m) \leq 2C_{\eta}^2l(r)\left(r^3I\left(\frac{n}{r^3}, \frac{m}{r}\right) + n^{8/13}m^{11/13}r^{22/13}\log\left(\frac{n^4}{m}\right) + n^{3/4}m^{3/4}r^{3/2} + nm^{1/2}r\right), \tag{5.3.8}$$

where $l(r) = \beta(r)\log^3 r$. Note that we may assume $C_{\eta}^1 < 1$, and then, in the range $m < C_{\eta}^1n^4$ we have

$$n^{8/13}m^{11/13}\log\left(\frac{n^4}{m}\right) = m\left(\frac{n^4}{m}\right)^{2/13}\log\left(\frac{n^4}{m}\right) < \tilde{c}m\left(\frac{n^4}{m}\right)^{2/11} = \tilde{c}n^{8/11}m^{9/11},$$

and

$$n^{3/4}m^{3/4} < \tilde{c}(n^{8/11}m^{9/11} + nm^{1/2})$$

for a suitable constant \tilde{c} . Substituting that in (5.3.8), and also substituting $n^{8/11}m^{9/11} = m(n^4/m)^{2/11}$, we get

$$I(n, m) \leq 2C_{\eta}^2l(r)\left(r^3I\left(\frac{n}{r^3}, \frac{m}{r}\right) + 2m\left(\frac{n^4}{m}\right)^{2/11}r^2 + 2nm^{1/2}r^2\right), \tag{5.3.9}$$

Our next and final step is to solve this recurrence, and derive the bound of Theorem 5.2.2.

Theorem 5.3.2 (Cf. Theorem 5.2.2). *For any $0 < \eta < 1$, and for any arbitrarily small $\varepsilon > 0$, there exist a constant $A = A_{\eta, \varepsilon}$, which depends on ε and η , such that*

$$I_{\eta}(n, m) \leq A \left(m \left(\frac{n^4}{m} \right)^{2/11+\varepsilon} + nm^{1/2} \right). \quad (5.3.10)$$

Proof. We use induction on m . Our end conditions are as follows

1. If m is smaller than some constant, say c , then $I \leq nm \leq cn$. Clearly, in this case (5.3.10) is satisfied, if we choose A sufficiently large.
2. If $m > \frac{1}{2}C_{\eta}^1 n^4$, then $I = O(m) = O(n^4)$ (recall that only the case $m < C_{\eta}^1 n^4$ can arise, so in the present situation $m = \Theta(n^4)$). This can be shown, e.g., using the previous bound of $I = O(n^{4/5}m^{4/5} + nm^{2/3})$. In this case (5.3.10) is satisfied, if we choose A sufficiently large.

We now deal with the case where n and m are large, and $c < m < \frac{1}{2}C_{\eta}^1 n^4$, in which case, we apply the induction hypothesis and rewrite (5.3.9) as

$$\begin{aligned} I(n, m) &\leq 2C_{\eta}^2 l(r) \left(r^3 A \left(\frac{m}{r} \left(\frac{n^4 r}{r^{12} m} \right)^{2/11+\varepsilon} + \frac{n}{r^3} \left(\frac{m}{r} \right)^{1/2} \right) + 2mr^2 \left(\frac{n^4}{m} \right)^{2/11} \right. \\ &\quad \left. + 2nm^{1/2}r^2 \right) \\ &= 2C_{\eta}^2 l(r) \left(Am \left(\frac{n^4}{m} \right)^{2/11+\varepsilon} r^{-11\varepsilon} + Anm^{\frac{1}{2}}r^{-1/2} + 2mr^2 \left(\frac{n^4}{m} \right)^{2/11} \right. \\ &\quad \left. + 2nm^{1/2}r^2 \right) \\ &\leq 2C_{\eta}^2 l(r) (Ar^{-11\varepsilon} + 2r^2) m \left(\frac{n^4}{m} \right)^{2/11+\varepsilon} + 2C_{\eta}^2 l(r) (Ar^{-1/2} + 2r^2) nm^{1/2}. \end{aligned}$$

We put $r = r_{\eta, \varepsilon}$ to be a sufficiently large constant so as to satisfy both inequalities $r^{11\varepsilon} > 4C_{\eta}^2 l(r)$ and $r^{1/2} > 4C_{\eta}^2 l(r)$. With this choice of r , we have

$$I(n, m) \leq \left(\frac{A}{2} + 4r^2 C_{\eta}^2 l(r) \right) m \left(\frac{n^4}{m} \right)^{2/11+\varepsilon} + \left(\frac{A}{2} + 4r^2 C_{\eta}^2 l(r) \right) nm^{1/2}.$$

Next, we choose $A = 8r^2 C_{\eta}^2 l(r)$, and this gives us

$$I(n, m) \leq A \left(m \left(\frac{n^4}{m} \right)^{2/11+\varepsilon} + nm^{1/2} \right).$$

This establishes the induction step, thereby completing the proof of the theorem, and consequently also the proof of Theorem 5.2.2. \square

Remark: The reader may be left wondering where did the exponents $8/11$ and $9/11$ “pop-up” from, since they do not appear explicitly in (5.3.8). The answer is that a solution of (5.3.8) with a leading term $n^\alpha m^\beta$ must satisfy the two inequalities $3\alpha + \beta > 3$, and $\alpha + 4\beta > 4$. The first inequality is needed to control the homogeneous part of the recurrence, i.e., the term $r^3 l(r) \cdot I\left(\frac{n}{r^3}, \frac{m}{r}\right)$, and the second inequality is needed to ensure that the term $n^{8/13} m^{11/13}$ is dominated by $n^\alpha m^\beta$ when $m = O(n^4)$, as can be easily verified. The exponents $8/11 + \varepsilon$ and $9/11 + \varepsilon$, for arbitrarily small ε , are, in a sense, the best solution of these inequalities.

5.4 Extending the results to partial incidence graphs

Theorems 5.2.1 and 5.2.2 can be applied to a more general setting where not necessarily all incidences are considered. In this setting, we have, as before, a set of points P and a set of spheres \mathcal{S} . In addition, we have a subgraph of their incidence graph $E \subseteq G(P, \mathcal{S})$, such that each sphere $\sigma \in \mathcal{S}$ is η -nondegenerate with respect to its neighbors in E . That is, there does not exist a circle $\gamma \subset \sigma$ such that the number of E -incidences on γ (that is, the number of points $p \in P \cap \gamma$ such that $(p, \sigma) \in E$) is larger than η times the number of E -incidences on σ . In that case, the upper bound of Theorem 5.2.2 applies to the number of edges in E , and the upper bound of Theorem 5.2.1 applies to the number of spheres such that each sphere has at least k neighbors in E . We shall state the partial incidences variant of Theorem 5.2.1 explicitly, as we are going to use it in Chapter 6.

Theorem 5.4.1. *Let $0 < \eta < 1$ be an arbitrary constant, and let $\varepsilon > 0$ be arbitrarily small. Let P be a set of n points in \mathbb{R}^3 , let \mathcal{S} be a set of spheres in \mathbb{R}^3 , and let $E \subseteq G(P, \mathcal{S})$ be an arbitrary subset of their incidences such that each sphere $\sigma \in \mathcal{S}$ is η -nondegenerate with respect to its neighbors in E , i.e., with respect to the set*

$$N_E(\sigma) = \{p \in P \mid (p, \sigma) \in E\}.$$

Assume further that $|N_E(\sigma)| \geq k$ for each $\sigma \in \mathcal{S}$, where k is an arbitrary integer. Then

$$|\mathcal{S}| = O\left(\frac{n^4}{k^{11/2-\varepsilon}} + \frac{n^2}{k^2}\right),$$

with the multiplicative constant depending on η and ε .

The proof is essentially the same as the above proof for the simpler case where $E = G(P, \mathcal{S})$. That is, the proof given in Section 5.3 applies almost verbatim to the case of partial incidences by noting that every statement about incidences should be taken as a statement about the incidences in E . For example, a j -bad sphere is a sphere containing a circle of j points of P that are also neighbors of that sphere in E . Most of the bounds derived in the previous analysis are worst-case upper bounds, and they therefore also apply to E -incidences. The only subtlety is in handling j -bad spheres, where we need to observe that (5.3.4) also holds for E -incidences.

5.5 Conclusion

In this chapter we have obtained an improved bound on the number of incidences between points and nondegenerate spheres in \mathbb{R}^3 . The usefulness of this bound is demonstrated in Chapter 6 where it is used for upper-bounding the number of mutually similar triangles in \mathbb{R}^3 . We believe that these results (Theorems 5.2.1 and 5.2.2) can be improved still further. We also believe that they can be extended to higher-dimensional nondegenerate spheres, and that they shall find applications in various other problems of geometric incidences and repeated subconfigurations.

As in Chapter 4, it would be interesting to tighten the bounds in Theorems 5.2.1 and 5.2.2, getting rid of the subpolynomial factors, using the polynomial partitioning technique (see Section 2.1.3).

Chapter 6

Repeated Simplices in d Dimensions

In this chapter, we consider the problem of bounding the maximum possible number $F_{k,d}(n)$ of k -simplices that are spanned by a set of n points in \mathbb{R}^d and are similar to a given simplex. We first show that $F_{2,3}(n) = O^*(n^{58/27})$, and then tackle the general case, and show that $F_{d-2,d}(n) = O(n^{d-2+\frac{2}{3}\varepsilon_d})$ and $F_{d-1,d}(n) = O(n^{d-2+\varepsilon_d})$, for any d , where $\varepsilon_d = \sqrt{2} \cdot 2^{-d/4}$. Our technique extends to derive bounds for other values of k and d , and we illustrate this by showing that $F_{2,5}(n) = O(n^{8/3})$.

A preliminary version of these results, with coarser upper bounds, appeared in the *Annual Symposium on Computational Geometry* 2007 [4] coauthored with Agarwal, Purdy, and Sharir.

6.1 Introduction

We refer the reader to Section 1.10 for an overview of the known bounds on the number of repeated congruent and similar simplices of various dimensions.

Let P be a set of n points in \mathbb{R}^d , and let Δ be a prescribed k -simplex, for some $0 \leq k \leq d$. Let $F(P, \Delta)$ denote the number of k -simplices spanned by P that are similar to Δ . Set

$$F_{k,d}(n) = \max F(P, \Delta),$$

where the maximum is taken over all sets P of n points in \mathbb{R}^d and over all k -simplices Δ in \mathbb{R}^d . Let $f_{k,d} \in \mathbb{R}$ be the polynomial degree of $F_{k,d}$, i.e.,

$$f_{k,d} = \min \{r \mid F_{k,d}(n) = O^*(n^r)\}.$$

Similarly, we denote by $G(P, \Delta)$ the number of k -simplices spanned by P and *congruent* to Δ , by $G_{k,d}(n)$ the maximum number of such simplices over $|P| = n$ and Δ a k -simplex, and by $g_{k,d}$ the polynomial degree of $G_{k,d}$.

We wish to obtain sharp bounds on $F_{k,d}(n)$ and $G_{k,d}(n)$ (or, rather, on $f_{k,d}$ and $g_{k,d}$). It suffices to consider cases with $2 \leq k \leq d-1$, since $F_{0,d}(n) = n$, $F_{1,d}(n) = \binom{n}{2}$, and $F_{d,d}(n) \leq 2F_{d-1,d}(n)$. (Still, it is plausible, as has indeed been conjectured in [21, 23], that $F_{d,d}(n)$ is much smaller than $F_{d-1,d}(n)$.)

The problem of obtaining sharp bounds on $F_{k,d}(n)$ is motivated by *exact pattern matching*: We are given a set P of n points in \mathbb{R}^d and a “pattern set” Q of $m \leq n$ points (in most applications m is much smaller than n ; let us assume $m \geq d + 1$), and we wish to determine whether P contains a copy of Q , under some allowed class of transformations, or, alternatively, to enumerate all such copies. See [23] for a comprehensive review of this and related problems. A commonly used approach to this problem, say, for the case of similarities, is to take a d -simplex Δ spanned by some points of Q , and find all similar copies of Δ that are spanned by points of P . For each such copy Δ' , take the similarity transformation(s) that map Δ to Δ' , and check whether all the other points of Q map to points of P under that transformation. The efficiency of such an algorithm depends on the number of similar copies of Δ in P . Using this approach for congruences, de Rezende and Lee [32] developed an $O(mn^d)$ -time algorithm to determine whether P contains a congruent copy of Q . For $d = 3$, Brass [20] developed an $O(mn^{7/4}\beta(n)\log n + n^{11/7+\epsilon})$ -time algorithm, which improves an earlier result by Boxer [19]. See also [18, 22] for related work. To recap, for applications of this kind, the main quantity of interest is $F_{d-1,d}(n)$.

In their recent monograph [23, pp. 265–266], Brass et al. review the known bounds on $F_{k,d}(n)$ and state various conjectures and open problems (see also [21]). In the plane, $F_{2,2}(n) = \Theta(n^2)$, where the lower bound can be obtained, e.g., from a section of the triangular grid (see [2] for an improved constant of proportionality). For more involved patterns, sufficient and necessary conditions for a tight quadratic bound are given in [60]; see also [38].

There are practically no known upper bounds in $d \geq 3$ dimensions, with the sole exception of a bound of $O(n^{2.2})$ on $F_{2,3}(n)$ (and $F_{3,3}(n)$), given by Akutsu et al. [9]. Brass [21] conjectures, though, that $F_{3,3}(n) = o(n^2)$; the best known lower bound is $F_{3,3}(n) = \Omega(n^{4/3})$.

The case of *congruent* simplices has also been studied, with somewhat greater success. We have reviewed the known results in Section 1.10; see also Agarwal and Sharir [7] and references therein.

Returning to the case of similar simplices, we note that the only known lower bounds for $F_{k,d}(n)$ are the same bounds for $G_{k,d}(n)$, namely $\Omega(n^{d/2})$ for $d \geq 4$ even and $k \geq d/2 - 1$, and $\Omega(n^{d/2-1/6})$, for $d \geq 3$ odd and $k \geq (d - 1)/2$; see also [1, 2, 7].

Our results

We first obtain the bound $F_{2,3}(n) = O^*(n^{58/27}) \approx O^*(n^{2.148})$, improving upon the bound of Akutsu et al. [9]. Recall that Brass [21] conjectures that $F_{3,3}(n) = o(n^2)$ (see also [23, p. 265]), in contrast to the known lower bound $F_{2,3}(n) = \Omega(n^2)$ (in fact, $F_{2,2}(n)$ is already $\Theta(n^2)$).

We then tackle the general case, and try to obtain sharp bounds on $F_{k,d}(n)$ and $G_{k,d}(n)$, focussing on the former quantities. In studying the structure of similar and congruent simplices spanned by finite point sets, we need to analyze *anchored* simplices, a term which will be defined more rigorously later. Informally, any anchored placement of a simplex Δ can be obtained from another such placement by rotating Δ around a set U of j anchor points, so that each vertex of Δ maintains fixed distances to the elements of U , forcing Δ to move on a sphere of some dimension $d' < d$. We thus introduce a third type of simplices to be counted, and define $G_{k,d,j}(n)$ to be the maximum number of congruent simplices in an n -element point set in \mathbb{R}^d , anchored at some

j fixed points. As it turns out, counting congruent anchored simplices is, in some sense, a more basic problem than the non-anchored versions of the problem, and its solution yields bounds for the other counts too.

Using the tools that we develop to deal with anchored simplices, we obtain the first nontrivial bounds for $F_{k,d}(n)$ (and $G_{k,d,j}(n)$). Specifically, for $F_{d-2,d}(n)$ and $F_{d-1,d}(n)$ (and thus also for $F_{d,d}(n)$) we show:

$$F_{d-2,d}(n) = O^*(n^{d-2+\frac{2\sqrt{2}}{3}\cdot 2^{-d/4}}) \quad \text{and} \quad F_{d-1,d}(n) = O^*(n^{d-2+\sqrt{2}\cdot 2^{-d/4}}),$$

for d sufficiently large. This improves upon the respective trivial bounds of $F_{d-2,d}(n) = O(n^{d-1})$ and $F_{d-1,d}(n) = O(n^d)$ (and, for sufficiently large values of d , upon the bounds in the earlier conference version of this chapter: $F_{d-2,d}(n) = O(n^{d-8/5})$, and $F_{d-1,d}(n) = O^*(n^{d-72/55})$). For the particular value of $d = 4$, we apply the bound for $F_{d-2,d}(n)$ from the conference version to obtain the bound $F_{2,4}(n) = O(n^{12/5})$. Finally, we prove that $F_{2,5}(n) = O(n^{8/3})$. Note that this is the last interesting case for triangles because, by Lenz' construction, $F_{2,6}(n) = G_{2,6}(n) = \Theta(n^3)$. Needless to say, none of these bounds (except for the last cubic bound) is known (nor conjectured) to be tight.

Our techniques are strongly based on bounds on the number of incidences between points and spheres (or circles). In particular, we use the improved bound from Chapter 5 on incidences between points and nondegenerate spheres. Considerable progress has been made on these problems in recent years; see [6, 15, 17, 63] and also [70] for a comprehensive survey of these and related results. We also use the bound, reviewed in Section 1.5.3, of Elekes and Tóth [41] on the number of rich nondegenerate hyperplanes. In this regard, our work can be regarded as an application of the recent developments in incidence problems, and it raises several interesting basic open problems in this area, as discussed at the end of this chapter.

6.2 Similar triangles in three dimensions

In this section we derive an improved bound on $F_{2,3}(n)$. Let Δ_0 be a fixed triangle, and let $\mathcal{S}(\Delta_0)$ denote the set of all triangles spanned by P and similar to Δ_0 . As a warm-up exercise, we first derive a simple, albeit weaker, upper bound on $F_{2,3}(n)$, and then establish a tighter bound whose proof is considerably more involved.

6.2.1 A simpler and weaker bound

For each pair a, b of points of P , any triangle Δabc in $\mathcal{S}(\Delta_0)$, with $c \in P \setminus \{a, b\}$, has the property that c lies on a circle $\gamma_{a,b}$, which is orthogonal to ab and whose center lies at a fixed point on the line containing ab .¹ Moreover, given a circle γ , there exist at most two pairs (a, b) , such that $\gamma = \gamma_{a,b}$

¹Actually, there are up to six such circles, depending on the roles of the vertices a and b in the triangle Δabc , i.e., to which vertices in Δ_0 do a and b correspond. To resolve this ambiguity, we have to introduce ordering among the triangle vertices, which we do explicitly in Subsection 6.2.2. Meanwhile, let us assume we have a well defined ordering on the vertices, and so, $\gamma_{a,b}$ is unique.

(if there are two pairs, one is the reflection of the other through the center of γ). Hence, $|\mathcal{S}(\Delta_0)|$ is at most twice the number of incidences between the points of P and the (at most) $\binom{n}{2}$ distinct circles $\gamma_{a,b}$.

As shown by Aronov et al. [15] (see (1.2.1) in Chapter 1), the number of incidences between n points and c distinct circles in \mathbb{R}^3 (or, for that matter, in any dimension ≥ 2) is

$$O(n^{6/11}c^{9/11}\log^{2/11}(n^3/c) + n^{2/3}c^{2/3} + n + c). \quad (6.2.1)$$

Substituting $c = O(n^2)$, the first term dominates, and we obtain

$$F_{2,3}(n) = O(n^{24/11}\log^{2/11}n) = O^*(n^{24/11}) = O(n^{2.182}).$$

We remark that a similar approach was taken by Akutsu et al. [9], except that they used a weaker bound on point-circle incidences (albeit the best known bound at that time). Finally, taking Δ_0 to be an equilateral triangle and P to be a section of a 2-dimensional triangular lattice, it is easy to verify that $F_{2,2}(n) = \Theta(n^2)$, which implies that $F_{2,d}(n) = \Omega(n^2)$ for any $2 \leq d \leq 5$ (as already noted, $F_{2,d}(n) = \Theta(n^3)$ for $d \geq 6$).

6.2.2 An improved bound

To simplify the presentation, let us assume² that Δ_0 is not isosceles, so its edges have distinct lengths. For each triangle $\Delta abc \in \mathcal{S}(\Delta_0)$, we can order its vertices in a unique order, say a, b, c , so that a is incident to the two longest edges of the triangle, and b is the other endpoint of the longest edge. We call a the *main vertex* and b the *secondary vertex* of the triangle. For any pair of points a, b , denote the sphere centered at a and containing b by $\sigma_{a,b}$. Let $\gamma = \gamma_{a,b}$ denote the circle that is the locus of all the points c' such that $\Delta abc' \sim \Delta_0$, a is the main vertex, and b is the secondary vertex. (γ is the same circle constructed in the preceding proof). Clearly, γ is contained in a sphere centered at a (which is $\sigma_{a,c}$, for any $c \in \gamma$).

For $a \in P$, let Σ_a denote the set of all spheres $\sigma_{a,b}$, for $b \in P \setminus \{a\}$. Define a relation Π_a on $\Sigma_a \times \Sigma_a$, which consists of all ordered pairs $(\sigma_{a,b}, \sigma_{a,c})$ of spheres for which $\Delta abc \in \mathcal{S}(\Delta_0)$ (with a, b, c ordered in the above manner). We denote by Γ_a the set of all circles $\gamma_{a,b}$ (note that the pair (a, b) uniquely determines $\gamma_{a,b}$ within Γ_a). Put $\Sigma = \bigcup_{a \in P} \Sigma_a$, $\Gamma = \bigcup_{a \in P} \Gamma_a$, and $\Pi = \bigcup_{a \in P} \Pi_a$. By construction, each sphere in Σ appears in at most two pairs of Π (either as the bigger sphere $\sigma_{a,b}$, or as the smaller sphere $\sigma_{a,c}$, or in both roles).

For an integer $k \leq n$, let $\Pi_{\leq k} \subseteq \Pi$ denote the set of pairs of spheres (σ, τ) such that either σ or τ contains at most k points of P , and let $\Pi_{>k} = \Pi \setminus \Pi_{\leq k}$ denote the complementary set of those pairs in which each of the spheres contains more than k points.

As already discussed, each (ordered) triangle $\Delta abc \in \mathcal{S}(\Delta_0)$ corresponds to an incidence between the point c and the circle $\gamma_{a,b}$, and the same circle $\gamma_{a,b}$ can arise for at most two pairs (a, b) . Hence we have

$$\frac{1}{2}|\mathcal{S}(\Delta_0)| \leq I(P, \Gamma) \leq |\mathcal{S}(\Delta_0)|,$$

²The proof can be carried out without this assumption, but the presentation is simplified in this way.

where $I(P, \Gamma)$ is the number of incidences between the points of P and the circles of Γ . It therefore suffices to bound $I(P, \Gamma)$.

Fix a threshold parameter k , and a sufficiently small constant $\eta > 0$. We call a pair in $\Pi_{\leq k}$ *light* and a pair in $\Pi_{> k}$ *heavy*. We classify the heavy pairs into nondegenerate and degenerate pairs, as follows: A heavy sphere σ (i.e., a sphere containing more than k points) is *degenerate* (or, more concisely, η -degenerate) if there exists a circle $\gamma \subset \sigma$ (not necessarily from the family Γ) such that $|\gamma \cap P| \geq \eta |\sigma \cap P|$; otherwise it is *nondegenerate*. A pair in $\Pi_{> k}$ is called *nondegenerate* if at least one of the spheres in the pair is nondegenerate, and *degenerate* otherwise.

We bound separately the number of incidences between the points of P and the circles determined by each of these three types of pairs. Let I_L (resp., I_N, I_D) denote the number of incidences between P and the circles induced by light (resp., heavy nondegenerate, heavy degenerate) pairs. Then $I(P, \Gamma) = I_L + I_N + I_D$.

Handling light pairs. Let $a \in P$ and put

$$\Pi'_a = \Pi_a \cap \Pi_{\leq k}.$$

For a sphere pair $(\sigma, \tau) \in \Pi'_a$, put

$$P_\sigma = P \cap \sigma, \quad P_\tau = P \cap \tau, \quad \text{and} \quad \Gamma_\tau = \{\gamma_{a,b} \mid b \in P_\sigma\}.$$

Recall that either $|\Gamma_\tau| = |P_\sigma| \leq k$, or $|P_\tau| \leq k$. All the circles of Γ_τ lie on τ and have the same radius. Hence, as follows, e.g., from [79], the number of similar triangles associated with the pair (σ, τ) is bounded by

$$I(P_\tau, \Gamma_\tau) = O(|P_\sigma|^{2/3} |P_\tau|^{2/3} + |P_\sigma| + |P_\tau|).$$

Summing over all $(\sigma, \tau) \in \Pi'_a$, we get

$$I'_a := \sum_{(\sigma, \tau) \in \Pi'_a} I(P_\tau, \Gamma_\tau) = \sum_{(\sigma, \tau) \in \Pi'_a} O(|P_\sigma|^{2/3} |P_\tau|^{2/3} + |P_\sigma| + |P_\tau|).$$

The sums of the last two terms are clearly $O(n)$. As for the first term, one of $|P_\sigma|, |P_\tau|$ is at most k . We may assume, without loss of generality, that $|P_\tau| \leq k$, and then obtain, using Hölder's inequality,

$$\begin{aligned} \sum_{(\sigma, \tau) \in \Pi'_a} |P_\sigma|^{2/3} |P_\tau|^{2/3} &\leq k^{1/3} \sum_{(\sigma, \tau) \in \Pi'_a} |P_\sigma|^{2/3} |P_\tau|^{1/3} \\ &\leq k^{1/3} \left(\sum |P_\sigma| \right)^{2/3} \left(\sum |P_\tau| \right)^{1/3} \\ &= O(nk^{1/3}). \end{aligned}$$

Altogether, summing over all points $a \in P$, the number of point-circle incidences with circles generated by light pairs is

$$I_L = \sum_{a \in P} I'_a = O(n^2 k^{1/3}). \quad (6.2.2)$$

Handling nondegenerate heavy pairs. Let $\Pi_N \subseteq \Pi_{>k}$ denote the set of nondegenerate heavy pairs. We first fix a point $a \in P$ and bound the number of incidences between the points of P and the circles $\gamma_{a,b}$, for which there exists a pair of spheres $(\sigma, \tau) \in \Pi_N \cap \Pi_a$, such that $b \in \sigma$ and $\gamma_{a,b} \subset \tau$. Recall that at least one of the two spheres in each pair of Π_N is nondegenerate and both spheres contain more than k points of P . We first obtain an upper bound on $|\Pi_N|$ as follows.

We use the bound derived in Chapter 5 on the possible number of nondegenerate spheres incident to k out of n points. By this bound (Theorem 5.2.1), the number of spheres, and so, also the number of pairs of Π_N , is

$$\frac{1}{2}|\Sigma_N| \leq |\Pi_N| \leq |\Sigma_N| = O^* \left(\frac{n^4}{k^{11/2}} + \frac{n^2}{k^2} \right). \quad (6.2.3)$$

The bound on $|\Pi_N|$ follows from the fact that any nondegenerate heavy sphere can participate in a pair of Π_N at most twice. Once as the first sphere σ , and once as the second sphere τ .

We partition Π_N into $O(\log n)$ classes, $\Pi_N^{(1)}, \Pi_N^{(2)}, \dots$, where $\Pi_N^{(i)}$ consists of those pairs $(\sigma, \tau) \in \Pi_N$ such that

$$2^{i-1}k < \max\{|\sigma \cap P|, |\tau \cap P|\} \leq 2^i k,$$

for $i = 1, \dots, \log_2(n/k)$. We sum the incidence bounds within each class separately. Fix a class $\Pi_N^{(i)}$, and put $k_i = 2^i k$. For each pair $(\sigma, \tau) \in \Pi_N^{(i)}$, σ induces on τ a set of at most k_i congruent circles, and the number of points on τ is also at most k_i . Hence the number of incidences between these points and circles is $O(k_i^{4/3})$. By the above discussion and Theorem 5.2.1, the number of such pairs of spheres is $|\Pi_N^{(i)}| = O^*(n^4/k_i^{11/2} + n^2/k_i^2)$. Hence, summing over i , the overall number of incidences involving sphere pairs in Π_N is

$$I_N = O^* \left(\sum_i k_i^{4/3} \cdot \left(\frac{n^4}{k_i^{11/2}} + \frac{n^2}{k_i^2} \right) \right) = O^* \left(\frac{n^4}{k^{25/6}} + \frac{n^2}{k^{2/3}} \right). \quad (6.2.4)$$

Handling degenerate heavy pairs. Let $\Pi_D = \Pi_{>k} \setminus \Pi_N$ be the set of degenerate heavy pairs. We apply the following pruning process on each pair $(\sigma, \tau) \in \Pi_D$. Recall that τ is a degenerate sphere containing more than k points of P . Then there is a circle $\gamma_1 \subset \tau$ containing at least $\eta|P_\tau|$ points of $P_\tau = P \cap \tau$. If we remove the points of $P \cap \gamma_1$ from τ , then one of the following may happen:³

1. τ is incident to at most k of the remaining points of $P \setminus \gamma_1$.
2. τ becomes nondegenerate with respect to the remaining points.
3. τ is still degenerate and contains more than k points.

³The pruning process is applied only within τ ; the removal of points from τ does not remove them from other spheres.

In the third case, we continue with the pruning process, until one of the first two events occurs, which will happen after at most $\log_{\frac{1}{1-\eta}} n$ iterations. At the end of the process, if τ contains k or fewer points, we include it, together with its remaining incidences and with the sphere σ , as one of the pairs of $\Pi_{\leq k}$, meaning that we only consider the remaining points when bounding the number of point-circle incidences on τ . Otherwise, τ is nondegenerate with respect to the set of at least k remaining points. In other words, when considering only the surviving points on τ at the end of the pruning process, (σ, τ) becomes a light or nondegenerate pair. We now bound the number of spheres using Theorem 5.4.1 instead of Theorem 5.2.1, to account for the partial incidences that are being considered here, and conclude that the bound (6.2.3) still holds for these latter pairs.

We still have to count the incidences involving the removed points and/or circles on each of these spheres τ . Such a (yet uncounted) incidence $p \in \gamma$, on a degenerate sphere τ , is uncounted either because

- (a) γ was one of the circles whose points were removed from P_τ , or
- (b) γ was not such a circle, but p was removed because it lies on another circle $\gamma' \subset \tau$ whose points have been removed.

For an incidence of type (a), we have $\gamma \in \Gamma_\tau$. There can be at most two spheres τ which produce the same γ . Thus, the number of uncounted incidences of type (a) is at most twice the number of incidences between the points of P and the removed circles. Using the fact that each of these circles contains at least ηk points, we have, by (1.2.3), that the number of these incidences is

$$O\left(\frac{n^3 \log k}{k^{9/2}} + \frac{n^2}{k^2} + n\right).$$

To bound the number of incidences of type (b), we observe that, for each pair $(\sigma, \tau) \in \Pi_D$, the number of removed circles on τ is $O(\log n)$, and each $\gamma \in \Gamma_\tau$ can lose at most two incidences for each removed circle. Thus, the number of such incidences on τ is $O(|\Gamma_\tau| \log n)$. Summing over all degenerate spheres, the overall number of type (b) incidences is $O(n^2 \log n)$. Hence,

$$\begin{aligned} I_D &= O\left(\frac{n^3 \log k}{k^{9/2}} + \frac{n^2}{k^2} + n + n^2 \log n\right) \\ &= O\left(\frac{n^3 \log k}{k^{9/2}} + n^2 \log n\right). \end{aligned} \tag{6.2.5}$$

Adding (6.2.2), (6.2.4), and (6.2.5), we get

$$\begin{aligned} I(P, \Gamma) &\leq I_L + I_N + I_D \\ &= O^*\left(n^2 k^{1/3} + \frac{n^4}{k^{25/6}} + \frac{n^3}{k^{9/2}} + n^2 \log n\right), \end{aligned}$$

for any k . By choosing $k = n^{4/9}$, we obtain the main result of this section:

Theorem 6.2.1. *Let P be a set of n points in \mathbb{R}^3 , and let Δ_0 be some fixed triangle. Then the number of triangles similar to Δ_0 spanned by P is $O^*(n^{58/27}) = O^*(n^{2.148})$.*

6.3 General bounds in higher dimensions

We next consider the general case of bounding the maximum number $F_{k,d}(n)$ (resp., $G_{k,d}(n)$) of mutually similar (resp., congruent) k -simplices in a set of n points in \mathbb{R}^d . We use the same general strategy in both cases. That is, we first choose a $(k-1)$ -face Δ' of the prescribed k -simplex Δ , and bound the number of $(k-1)$ -simplices similar or congruent to Δ' , by applying recursively the bound for repeated simplices of a lower dimension. Then, for each such $(k-1)$ -simplex, we consider the set of all points in \mathbb{R}^d that complete it into a simplex similar or congruent to Δ (more precisely, congruent to a fixed scaled copy of Δ). This set, as we will see, is a $(d-k)$ -dimensional sphere (provided that $k > 1$, or that we are bounding congruent simplices). The number of repeated k -simplices is then equal to the number of incidences between the points of P and these spheres. Our goal now is to derive an incidence bound between the given points and the spheres constructed in the above manner. A significant technical complication which arises in the analysis is that these spheres need not necessarily be distinct, so we have to take their multiplicity into account.

We start this section by studying the setting, where we have a $(k-1)$ -simplex Δ' congruent to a $(k-1)$ -face of a k -simplex Δ , and we want to add another vertex to Δ' such that the resulting k -simplex is congruent to Δ . Note that this situation is also the one we face in the case of similar simplices (assuming $k > 1$): There Δ' is similar to a $(k-1)$ -face of Δ , but then completing Δ' to a k -simplex similar to Δ is the same as completing Δ' to a k -simplex congruent to a fixed scaled copy of Δ . The analysis of such completions leads us to the notion of *anchored* simplices, a term which will be explained later. We continue by deriving a recurrence on the number of congruent anchored simplices, and then apply this recurrence to derive a bound on the number of similar simplices. Our technique can be used to derive bounds on $F_{k,d}(n)$ for all values of k and d , but we demonstrate it only for $F_{d-1,d}(n)$, and $F_{d-2,d}(n)$, and also for $F_{2,5}(n)$.

6.3.1 Completing a partial simplex

We use the notation $\Delta(T)$ to denote the simplex spanned by a tuple T . We denote *congruence* by $\Delta_1 \cong \Delta_2$, and recall that *similarity* is denoted by $\Delta_1 \sim \Delta_2$. We use the term *facet* to refer to a (proper) face of maximal dimension. More formally, a facet of a k -simplex is a $(k-1)$ -face of it.

Let S be a j -sphere in \mathbb{R}^d . The *axis* of S is the $(d-j-1)$ -flat $\pi = \pi_S$ passing through the center of S and orthogonal to the $(j+1)$ -flat spanned by S . Every point $q \in \pi$ lies within the same distance from all the points on S . Conversely, π is the locus of all the points that have this property.

In this subsection we consider the following setting. We have some prescribed k -simplex Δ , a facet Δ' of Δ , and some k -tuple that spans a $(k-1)$ -simplex similar or congruent to Δ' . We study how this partial tuple can be completed into one that spans a simplex congruent to an appropriately scaled copy of Δ , and the structures that arise in this setting. To simplify the presentation, we only consider the case of congruent simplices, but, as noted above, the analysis carries over, with only minor modifications, to the case of similar simplices.

We will need the following simple propositions about sphere intersections.

Proposition 6.3.1. *Let S be a j -sphere, and let S_1 be a $(d - 1)$ -sphere in \mathbb{R}^d , where $j < d$. Then $S \cap S_1$ is either:*

1. *Empty, or*
2. *A single point (in case the two spheres are tangent), or*
3. *S , i.e., a j -sphere (in case $S \subseteq S_1$), or*
4. *A $(j - 1)$ -sphere (otherwise).*

The next proposition follows by induction using Proposition 6.3.1 for the induction step.

Proposition 6.3.2. *Let S_1, \dots, S_k be k $(d - 1)$ -spheres. Then $\bigcap_{i=1}^k S_i$ is either:*

1. *Empty, or*
2. *A single point, or*
3. *A sphere of dimension at least $d - k$.*

Our first lemma asserts that if we are given the first k vertices of a k -simplex congruent to Δ in \mathbb{R}^d , then the last vertex lies on a $(d - k)$ -sphere.

Lemma 6.3.3. *Let Δ be a k -simplex ($1 \leq k \leq d$), and let Δ' be a facet of Δ . Let $T \subset \mathbb{R}^d$ be a k -tuple such that $\Delta(T) \cong \Delta'$. Put*

$$S = \{s \in \mathbb{R}^d \mid \Delta(T \cup \{s\}) \cong \Delta\}.$$

Then S is a $(d - k)$ -sphere, whose axis is the $(k - 1)$ -flat spanned by T .

Proof. Let s be an arbitrary point in S . For each point $t \in T$, s lies on some $(d - 1)$ -sphere centered at t whose radius is equal to the distance between the two vertices in Δ mapped to s and t , respectively. Thus, S is the intersection of k $(d - 1)$ -spheres. We claim that this intersection is a $(d - k)$ -sphere.

Indeed, by Proposition 6.3.2, the intersection of k $(d - 1)$ -spheres can be either empty, or a single point, or a sphere of dimension at least $d - k$. S is certainly not empty, since there exists some isometry that maps the vertices of Δ' to T , and this isometry maps the remaining vertex of Δ into a point $s \in S$. Furthermore, the point symmetric to s about the affine hull of T is also in S , so S cannot be a single point, and is therefore a proper sphere (the point symmetric to s cannot be s itself, since in that case s would be affinely dependent on T , in contradiction to Δ being k -dimensional). As noted above, the dimension of this sphere is at least $d - k$. Now, every point of T is equidistant from all the points of S , and so, it is in the axis of S . Since T is affinely independent, the axis of S is at least $(k - 1)$ -dimensional, so the dimension of S is at most $d - k$. Hence, S is exactly $(d - k)$ -dimensional, as asserted. \square

For a k -simplex Δ , a facet Δ' of Δ , and a k -tuple T such that $\Delta(T) \cong \Delta'$, we denote the sphere S defined in Lemma 6.3.3 by $\sigma_{\Delta, \Delta'}(T)$, or $\sigma(T)$ if Δ and Δ' are clear from the context. The notation $\sigma_{\Delta, \Delta'}(T)$ implicitly assumes that $\Delta(T) \cong \Delta'$.

If Δ is d -dimensional, then $\sigma(T)$ is a 0-sphere, which is simply a pair of points. Thus, a tuple T such that $\Delta(T) \cong \Delta'$ can be completed into a simplex similar to Δ by two points only. It then follows that $G_{d,d}(n) \leq 2G_{d-1,d}(n)$, and $F_{d,d}(n) \leq 2F_{d-1,d}(n)$ (as already noted).

Anchored simplices. Let Δ be a $(k + j)$ -simplex, let Δ' be a k -face of Δ , and let U be a j -tuple of points in \mathbb{R}^d , such that $\Delta(U)$ is congruent to the $(j - 1)$ -face of Δ complementary to Δ' (i.e, the convex hull of the remaining vertices). A $(k + 1)$ -tuple $T \subset \mathbb{R}^d$ is said to be *anchored* at U with respect to Δ and Δ' , if $\Delta(T) \cong \Delta'$, and $\Delta(T \cup U) \cong \Delta$. That is, we fix the subset U of vertices of (a congruent copy of) Δ , and want to complete it into the full Δ , in which case, the vertices that we add form (a copy of) the k -face Δ' . For example, we can regard the situation in Lemma 6.3.3 as a special case of anchoring, where the simplex is anchored at T , and we want to add s as the complementary 0-dimensional face.

In the following lemma, we consider the converse of Lemma 6.3.3. That is, we have some prescribed $(d - k)$ -sphere S , which is known to consist of all the points that complete some unknown k -tuple into a prescribed k -simplex, and we want to characterize the distinct tuples of points which can generate S .

Basically, the lemma says that all these tuples are anchored at the center u of S and lie inside the axis π of S . Since π is $(k - 1)$ -dimensional, and each of the tuples T in question is a k -tuple, it spans a full-dimensional simplex within π . In that case, it suffices to consider only the first $k - 1$ points of T , because given these points, the k -th point is uniquely determined (a fact which is established in the proof of the lemma). Thus, given a sphere S as above, the lemma asserts that it is generated by some $(k - 1)$ -tuple of points anchored at the center u within the axis π , that span a simplex congruent to some prescribed anchored $(k - 2)$ simplex Δ'' , and then the k -th vertex is in π and is uniquely determined, and the $(k + 1)$ -st vertex lies anywhere on S . See Figure 6.1 for an illustration.

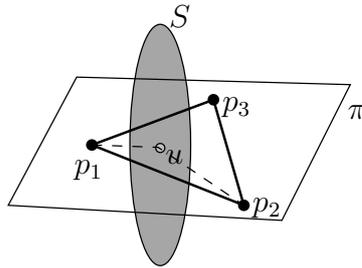


Figure 6.1: Illustrating the claim in Lemma 6.3.4. Here Δ is congruent to $\Delta(p_1, p_2, p_3, s)$ for any $s \in S$, Δ' is congruent to $\Delta(p_1, p_2, p_3)$, Δ'' is congruent to $p_1 p_2 = \Delta(p_1, p_2)$, and Δ^* is congruent to $\Delta(p_1, p_2, u)$.

Lemma 6.3.4. *Let Δ be a k -simplex, and let Δ' be a facet of Δ . Let S be a $(d - k)$ -sphere in \mathbb{R}^d which is equal to $\sigma_{\Delta, \Delta'}(T)$ for at least one tuple T . Let π be the axis of S , and let u be its center.*

Then, for all k -tuples $T = (p_1, p_2, \dots, p_k)$ such that $\sigma(T) = \sigma_{\Delta, \Delta'}(T) = S$, we have $T \subset \pi$. Furthermore, up to a possible permutation of T (which depends only on Δ or Δ'), there exists some $(k - 1)$ -simplex Δ^ in π , a facet Δ'' of Δ^* , and a mapping $f : \pi^{k-1} \rightarrow \pi$, such that $\sigma(T) = S$ if and only if the subtuple $T' = (p_1, p_2, \dots, p_{k-1})$ is anchored at $\{u\}$ with respect to Δ^* and Δ'' , and $p_k = f(p_1, p_2, \dots, p_{k-1})$. Moreover, the mapping f rigidly attaches p_k to $\{u, p_1, \dots, p_{k-1}\}$, in the sense that the sets $\{u, p_1, \dots, p_{k-1}, p_k\}$ are congruent to each other, over all tuples T with $\sigma(T) = S$.*

Remark. See Figure 6.1 for an illustration. The reason for the somewhat involved formulation of the lemma is that the set $T \cup \{u\} \subset \pi$, of size $k + 1$, is not affinely independent, so we need to remove one point from T so that the tuple T' of the remaining $k - 1$ points form with u a $(k - 1)$ -simplex (congruent to Δ^*) anchored at $\{u\}$. The removed point, though, is rigidly attached to that simplex, within π .

Proof. Let $T = (p_1, p_2, \dots, p_k)$ be a k -tuple such that $\sigma(T) = S$. Fix a point $p_i \in T$, and an arbitrary point $s \in S$, and let x_i denote the distance between them. x_i is equal to the distance between the i -th vertex of Δ' (which p_i is mapped to) and the vertex of Δ which is not in Δ' . Thus, x_i does not depend on the choice of s , so all points $s \in S$ have the same distance, x_i , from p_i . Hence, p_i lies in the axis π of S , implying that $T \subset \pi$.

The distance between p_i and u is $\sqrt{x_i^2 - r^2}$, where r is the radius of S , so all vertices of T have fixed distances from u as well. Thus, T is anchored at $\{u\}$ (in the somewhat looser sense discussed in the preceding remark).

Since T is affinely independent, it has a member which can be replaced by u , such that the resulting tuple is also affinely independent. Assume, without loss of generality, that this member of T is p_k . Put $\Delta^* = \Delta(p_1, \dots, p_{k-1}, u)$, and $\Delta'' = \Delta(p_1, \dots, p_{k-1})$.

The only point $p \in \pi$ satisfying both $\Delta(p_1, \dots, p_{k-1}, p) \cong \Delta'$, and $|p - u| = \sqrt{x_k^2 - r^2}$ is $p = p_k$. Indeed, since π is $(k - 1)$ -dimensional, there are only two points in π (one of which is p_k) which complete $\Delta(p_1, \dots, p_{k-1})$ to a simplex congruent to Δ' . If these two points have equal distances from u then u must lie in the $(k - 2)$ -flat spanned by p_1, \dots, p_{k-1} , contradicting the affine independence of p_1, \dots, p_{k-1}, u .

We define the mapping $f : \pi^{k-1} \rightarrow \pi$ as follows. For each $(k - 1)$ -tuple (p_1, \dots, p_{k-1}) of points in π such that $\Delta(p_1, \dots, p_{k-1}) \cong \Delta''$, and $\Delta(p_1, \dots, p_{k-1}, u) \cong \Delta^*$, we put $f(p_1, \dots, p_{k-1}) = p_k$, where p_k is the unique point satisfying $\Delta(p_1, \dots, p_k) \cong \Delta'$, and $|p_k - u| = \sqrt{x_k^2 - r^2}$. For any other $(k - 1)$ -tuple $(p_1, \dots, p_{k-1}) \in \pi^{k-1}$, we arbitrarily put $f(p_1, \dots, p_{k-1}) = u$, say.

From the way Δ^* , Δ'' , and f were constructed, it follows immediately that they satisfy the assertion of the lemma. That is, for each k -tuple $T = (p_1, \dots, p_k) \in \pi^k$, we have $\sigma(T) = S$ if and only if $\Delta(T') \cong \Delta''$, $\Delta(T', u) \cong \Delta^*$, and $f(T') = p_k$. This completes the proof. \square

Next, we generalize this lemma as follows. We are given not the entire $(d - k)$ -sphere, but a subsphere S of dimension $d - k - l$, for some $l > 1$, and we want to characterize all the k -tuples T that generate a $(d - k)$ -sphere $\sigma(T)$ that contains S .

Here the characterization depends on whether or not S is a great subsphere of $\sigma(T)$. To distinguish between the two cases, for a simplex Δ and one of its facets Δ' , we define $\rho = \rho_{\Delta, \Delta'}$ to be the radius of $\sigma_{\Delta, \Delta'}(T)$, where T is any tuple for which $\Delta(T) \cong \Delta'$. Clearly, this radius does not depend on the choice of T . Now $S \subset \sigma(T)$ is a great subsphere if and only if the radius of S is ρ .

As in Lemma 6.3.4, we claim that all these k -tuples T lie in the axis π and are anchored at the center u of S . See Lemma 6.3.6.

Here the situation is simpler than that of Lemma 6.3.4, because T does not span a full dimensional simplex, so we need not consider a $(k - 1)$ -subtuple of T and then claim that this subtuple uniquely determines the last point of T . In fact, this is generally not the case. There is one exception, though, and that is the special case where S happens to be a great subsphere. Then, the last point actually is determined by the first $k - 1$ points of T , and so the subsphere S is determined by some $(k - 1)$ -tuple that spans a simplex similar to some prescribed $(k - 2)$ -simplex and anchored at u . See Lemma 6.3.5.

We first consider the case where S is a great subsphere.

Lemma 6.3.5. *Let Δ be a k -simplex, and let Δ' be a facet of Δ . Let S be a $(d - k - l)$ -sphere of radius $\rho_{\Delta, \Delta'}$ in \mathbb{R}^d , for some $1 \leq l \leq d - k$, let π be its axis, and let u be its center.*

Then, for all k -tuples $T = (p_1, p_2, \dots, p_k)$ such that $\Delta(T) \cong \Delta'$ and $\sigma(T) = \sigma_{\Delta, \Delta'}(T) \supset S$, we have $T \subset \pi$. Furthermore, there exists some $(k - 1)$ -simplex Δ^ in π , a facet Δ'' of Δ^* , and a mapping $f : \pi^{k-1} \rightarrow \pi$, such that $\sigma(T) \supset S$ if and only if $T' = (p_1, p_2, \dots, p_{k-1})$ is anchored at $\{u\}$ with respect to Δ^* and Δ'' , and $p_k = f(p_1, p_2, \dots, p_{k-1})$.*

Proof. Since S is a great subsphere, its center u is also the center of any $(d - k)$ -sphere $\sigma(T)$ containing S . Thus, u lies in the axis of $\sigma(T)$, which is the affine hull of T , and is easily verified to be a subflat of the axis of S . Once we have established that u is affinely dependent on T , the rest of the proof proceeds exactly like that of the previous Lemma 6.3.4. \square

If S is not a great subsphere, then, as previously discussed, the situation is simpler, as we do not need the mapping f .

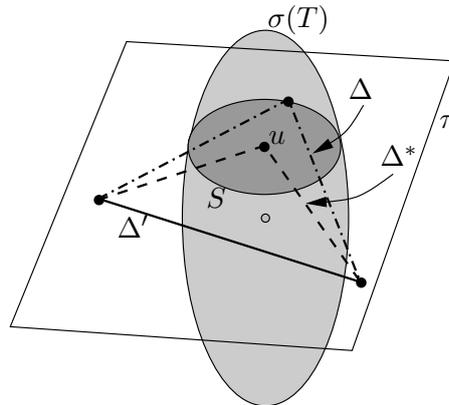


Figure 6.2: Illustrating the claim in Lemma 6.3.6.

Lemma 6.3.6. *Let Δ be a k -simplex, and let Δ' be a facet of Δ . Let S be a $(d - k - l)$ -sphere of radius smaller than $\rho_{\Delta, \Delta'}$ in \mathbb{R}^d , for some $1 \leq l \leq d - k$, let π be its axis, and let u be its center.*

Then, for all k -tuples T such that $\Delta(T) \cong \Delta'$ and $\sigma(T) = \sigma_{\Delta, \Delta'}(T) \supset S$, we have $T \subset \pi$. Furthermore, there exists some k -simplex Δ^ in π , which has a facet congruent to Δ' , such that $\sigma(T) \supset S$ if and only if T is anchored at $\{u\}$ with respect to Δ^* and Δ' .*

Remark. Note that here Δ^* is a k -simplex, rather than a $(k - 1)$ -simplex in the previous lemmas. Δ^* is obtained from Δ by replacing its vertex in S by u . See Figure 6.2.

Proof. Let $T = (p_1, p_2, \dots, p_k)$ be a k -tuple such that $\sigma(T) \supset S$. As in the proof of Lemma 6.3.4, one can easily verify that the distance x_i between p_i and any point $s \in S$ is fixed and does not depend on T and on s . Hence $T \subset \pi$. Furthermore, $|p_i - u| = \sqrt{x_i^2 - r^2}$, where r is the radius of S , so T is anchored at $\{u\}$.

Since S is not a great subsphere of $\sigma(T)$, it follows that u is affinely independent of T . Indeed, the affine hull of T is the axis of $\sigma(T)$, whereas u is not the center of $\sigma(T)$, so it does not have the same distance to all the points on $\sigma(T)$, and thus u is not in the axis of $\sigma(T)$. Thus, the convex hull of $T \cup \{u\}$ is a k -simplex, which we shall denote by $\Delta^* = \Delta(T \cup \{u\})$.

By construction, T is anchored at $\{u\}$ with respect to Δ^* and $\Delta(T) \cong \Delta'$. Furthermore, since we made no special assumptions on T other than $\sigma(T) \supset S$, it follows that for any other T such that $\sigma(T) \supset S$, T is anchored at $\{u\}$ with respect to Δ^* and Δ' . This proves the *only if* part.

For the *if* part, let $T = (p_1, \dots, p_k)$ be a k -tuple in π such that $\Delta(T) \cong \Delta'$, and $\Delta(T \cup \{u\}) \cong \Delta^*$. Then, for each $1 \leq i \leq k$, and for each $s \in S$, we have

$$|p_i - s| = \sqrt{|p_i - u|^2 + |u - s|^2} = \left[\left(\sqrt{x_i^2 - r^2} \right)^2 + r^2 \right]^{1/2} = x_i.$$

Thus, $S \subset \sigma(T)$. This completes the proof. \square

Next, we develop the anchored versions of the above lemmas. In this setting, as in the original lemmas, we study the properties of tuples that are one point short of spanning some prescribed simplex. Here, however, the first j points of any such tuple are assumed to be predetermined, and are the same for all relevant tuples. We assume in what follows that $j \geq 1$.

More formally, we proceed as follows. Let Δ be a $(k + j)$ -simplex, let Δ' be a k -face of Δ , and let Δ'' be a facet of Δ' (a $(k - 1)$ -face of Δ). Denote the vertices of Δ by v_0, \dots, v_{k+j} , the vertices of Δ' by v_j, \dots, v_{k+j} , and the vertices of Δ'' by v_j, \dots, v_{k+j-1} . We define the $(k + j - 1)$ -simplex $\Delta^* = \Delta^*(\Delta, \Delta', \Delta'')$ as the simplex spanned by the vertices v_0, \dots, v_{k+j-1} . Observe that Δ^* is a facet (a $(k + j - 1)$ -face) of Δ , and Δ'' is a $(k - 1)$ -face of Δ^* . See Figure 6.3 for a symbolic illustration.

Now, let U be a j -tuple of anchor points, such that $\Delta(U)$ is congruent to the simplex spanned by v_0, \dots, v_{j-1} . We are looking for k -tuples T such that $\Delta(T) \cong \Delta''$, and $\Delta(T \cup U) \cong \Delta^*$. Any such tuple T can be completed into a $(k + 1)$ -tuple T' , by adding one point, such that $\Delta(T') \cong \Delta'$ and $\Delta(T' \cup U) \cong \Delta$. See Figure 6.4 for an illustration; the anchored vertices of $\Delta(T' \cup U)$ are those corresponding to v_0, \dots, v_{j-1} .

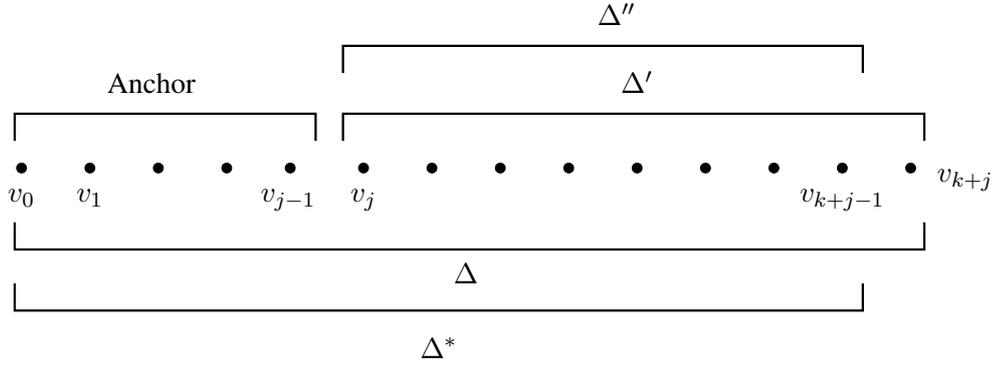


Figure 6.3: A symbolic representation of the simplex Δ and its various subsimplices.

Given that T is one point short of spanning the prescribed anchored simplex Δ , one can ask where can the last point be. This is what the lemmas below answer; they are the “anchored” versions of the four respective preceding lemmas. The first of these, equivalent to Lemma 6.3.3, asserts that given a k -tuple T anchored at U with respect to Δ^* and Δ'' , i.e., $\Delta(T) \cong \Delta''$ and $\Delta(T \cup U) \cong \Delta^*$, any point which completes T into a $(k + 1)$ -tuple anchored at U with respect to Δ and Δ' lies on a $(d - k - j)$ -sphere $\sigma(T)$. Recall that in Lemma 6.3.3 the dimension of the sphere is $d - k$. Thus, each anchor point reduces the dimension of $\sigma(T)$ by one, and the j anchor points together reduce the dimension of $\sigma(T)$ by j .

Lemma 6.3.7. *Let Δ be a $(k + j)$ -simplex, let Δ' be a k -face of Δ , let Δ'' be a facet of Δ' , let U be a j -tuple of points in \mathbb{R}^d so that $\Delta(U)$ is congruent to the $(j - 1)$ -face of Δ complementary to Δ' , and let T be a k -tuple in \mathbb{R}^d which is anchored at U with respect to $\Delta^* = \Delta^*(\Delta, \Delta', \Delta'')$, and Δ'' . Put*

$$S = \{ s \in \mathbb{R}^d \mid T \cup \{s\} \text{ is anchored at } U \text{ with respect to } \Delta \text{ and } \Delta' \}.$$

Then S is a $(d - k - j)$ -sphere.

Proof. Put $V = T \cup U$. We have $|V| = k + j$, $\Delta(V) \cong \Delta^*$, and for each $s \in S$, $\Delta(V \cup \{s\}) \cong \Delta$. Furthermore, S contains all the points that have this property. Thus, by Lemma 6.3.3, S is a $(d - k - j)$ -sphere. \square

We denote this sphere by $\sigma_{\Delta, \Delta', \Delta'', U}(T)$. If $\Delta, \Delta', \Delta''$ and U are clear from the context, we denote it as $\sigma(T)$ for short. Note that both T and U are contained in the axis π of this sphere.

Next, consider the anchored version of Lemma 6.3.4. We have a $(k + j)$ -simplex Δ , a k -face Δ' of Δ , a facet Δ'' of Δ' , a $(d - k - j)$ -sphere S , and a j -tuple U of anchor points, and we want to characterize all k -tuples T such that $\sigma_{\Delta, \Delta', \Delta'', U}(T) = S$. In other words, we want to characterize all $(k + j)$ -tuples V , whose first j points are those of U , such that $\sigma_{\Delta, \Delta^*}(V) = S$. By Lemma 6.3.4, these are exactly all the tuples in the axis π of S that are anchored at the center u of S with respect to some prescribed full-dimensional simplex in π . Here, however, we have the additional constraint that some of the points of V must coincide with the points of U . This implies that the

k -tuples which generate S are anchored at $j + 1$ points of π , namely, the points of $U \cup \{u\}$, with respect to some full-dimensional simplex.

Note that given U , it is not enough to require that S be any $(d - j - k)$ -sphere. First, the axis of S must contain U . In addition, the distance of S from each point $u' \in U$ must be equal to the distance between the respective vertices of Δ . For a tuple U and a sphere S as above, we say that they are *anchored* at each other (with respect to Δ , Δ' , and Δ'').

This is summarized in the following lemma.

Lemma 6.3.8. *Let Δ be a $(k + j)$ -simplex, let Δ' be a k -face of Δ , let Δ'' be a facet of Δ' , and let U be a j -tuple of points in \mathbb{R}^d , spanning a $(j - 1)$ -simplex congruent to the face of Δ complementary to Δ' . Let S be a $(d - k - j)$ -sphere in \mathbb{R}^d anchored at U with respect to Δ , Δ' , and Δ'' , let π be its axis, and let u be its center.*

Then for all k -tuples $T = (p_1, p_2, \dots, p_k)$ such that $\sigma_{\Delta, \Delta', \Delta'', U}(T) = S$, we have $T \subset \pi$. Furthermore, there exists some $(k + j - 1)$ -simplex $\tilde{\Delta}$ in π , a $(k - 2)$ -face Δ''' of $\tilde{\Delta}$, which is also a facet of Δ'' , and a mapping $f : \pi^{k-1} \rightarrow \pi$, such that $\sigma_{\Delta, \Delta', \Delta'', U}(T) = S$ if and only if $T' = (p_1, p_2, \dots, p_{k-1})$ is anchored at $U \cup \{u\}$ with respect to $\tilde{\Delta}$ and Δ''' , and $p_k = f(p_1, p_2, \dots, p_{k-1})$.

The proof, some hints of which were given above, is a straightforward extension of the proof of Lemma 6.3.4.

Next, consider the anchored versions of Lemmas 6.3.5 and 6.3.6. That is, we are given a low-dimensional sphere S , and want to characterize all the tuples, anchored at some given set of points, that generate a sphere which contains S . Here too the proofs extend the respective proofs of Lemmas 6.3.5 and 6.3.6 in a straightforward manner, and show that the tuples satisfying this condition have to be anchored also at one additional point, namely, the center of S .

Lemma 6.3.9. *Let Δ be a $(k + j)$ -simplex, let Δ' be a k -face of Δ , let Δ'' be a facet of Δ' , and let U be a j -tuple of points in \mathbb{R}^d , spanning a $(j - 1)$ -simplex congruent to the face of Δ complementary to Δ' . Let S be a $(d - k - j - l)$ -sphere in \mathbb{R}^d , for some $1 \leq l \leq d - k - j$, let π be its axis, and let u be its center. Assume further that S is anchored at U , and has radius ρ_{Δ, Δ^*} , where $\Delta^* = \Delta^*(\Delta, \Delta', \Delta'')$.*

Then for all k -tuples $T = (p_1, p_2, \dots, p_k)$ such that $\sigma_{\Delta, \Delta', \Delta'', U}(T) \supset S$, we have $T \subset \pi$. Furthermore, there exists some $(k + j - 1)$ -simplex $\tilde{\Delta}$ in π , a $(k - 2)$ -face Δ''' of $\tilde{\Delta}$, which is also a facet of Δ'' , and a mapping $f : \pi^{k-1} \rightarrow \pi$, such that $\sigma_{\Delta, \Delta', \Delta'', U}(T) \supset S$ if and only if $T' = (p_1, p_2, \dots, p_{k-1})$ is anchored at $U \cup \{u\}$ with respect to $\tilde{\Delta}$ and Δ''' , and $p_k = f(p_1, p_2, \dots, p_{k-1})$.

Lemma 6.3.10. *Let Δ be a $(k + j)$ -simplex, let Δ' be a k -face of Δ , let Δ'' be a facet of Δ' , and let U be a j -tuple of points in \mathbb{R}^d , spanning a $(j - 1)$ -simplex congruent to the face of Δ complementary to Δ' . Let S be a $(d - k - j - l)$ -sphere in \mathbb{R}^d , for some $1 \leq l \leq d - k - j$, let π be its axis, and let u be its center. Assume further that S is anchored at U , and has radius smaller than ρ_{Δ, Δ^*} , where $\Delta^* = \Delta^*(\Delta, \Delta', \Delta'')$.*

Then for all k -tuples T such that $\sigma_{\Delta, \Delta', \Delta'', U}(T) \supset S$, we have $T \subset \pi$. Furthermore, there exists some $(k + j)$ -simplex $\tilde{\Delta}$ in π , which has a face congruent to Δ'' , such that $\sigma_{\Delta, \Delta', \Delta'', U}(T) \supset S$ if and only if T is anchored at $U \cup \{u\}$ with respect to $\tilde{\Delta}$ and Δ'' .

6.3.2 Bounding the number of congruent anchored simplices

The setup. Let j, k, d be three nonnegative integers such that $0 < k < d$, and $0 \leq j < d - k$. Here d is the dimension of the ambient space, k is the dimension of the repeated simplex, and j is the number of anchor points. We allow $j = 0$ anchor points, so our analysis applies also for repeated congruent non-anchored simplices.

Let P be a set of n points in \mathbb{R}^d . Let Δ be a $(k + j)$ -simplex in \mathbb{R}^d , let Δ' be a k -face of Δ , and let U be a j -tuple of anchor points, spanning a $(j - 1)$ -simplex congruent to the face of Δ complementary to Δ' . We define

$$G(P, \Delta, \Delta', U) = \# \{T \subseteq P \mid T \text{ is a } (k + 1)\text{-tuple anchored at } U \text{ with respect to } \Delta \text{ and } \Delta'\}.$$

We denote by $G_{k,d,j}(n)$ the maximum of $G(P, \Delta, \Delta', U)$ over all choices of a set P of n points in \mathbb{R}^d , a $(k + j)$ -simplex Δ , a k -face Δ' of Δ , and a j -tuple U spanning a $(j - 1)$ -simplex congruent to the face of Δ complementary to Δ' . We denote the polynomial degree of $G_{k,d,j}(n)$ by $g_{k,d,j}$. In this setting, the previously defined $G(P, \Delta)$ can be regarded as $G(P, \Delta, \Delta, \emptyset)$, and we have $G_{k,d}(n) = G_{k,d,0}(n)$ and $g_{k,d} = g_{k,d,0}$.

In this section we derive recurrences on the quantities $G_{k,d,j}(n)$, and, consequently, on their polynomial degrees.

Let Δ'' be a facet of Δ' , that is, a $(k - 1)$ -simplex, and let $\Delta^* = \Delta^*(\Delta, \Delta', \Delta'')$ be the $(k + j - 1)$ -simplex as defined in the previous section (see Figure 6.3). We denote by \mathcal{T} the following set of k -tuples

$$\mathcal{T} = \{T \subseteq P \text{ a } k\text{-tuple} \mid \Delta(T) \cong \Delta'' \quad \text{and} \quad \Delta(T \cup U) \cong \Delta^*\},$$

and let Σ denote the following *multiset* of $(d - k - j)$ -spheres

$$\Sigma = \{\sigma(T) = \sigma_{\Delta, \Delta', \Delta'', U}(T) \mid T \in \mathcal{T}\}.$$

The fact that the elements of Σ are $(d - k - j)$ -spheres follows from Lemma 6.3.7. We have $|\mathcal{T}| = G(P, \Delta^*, \Delta'', U) \leq G_{k-1,d,j}(n)$.

The total number of simplices spanned by P that are congruent to Δ' and anchored at U is $I(P, \Sigma)$, the number of incidences between the points of P and the spheres of Σ , where each incidence is counted with the multiplicity of its sphere.

We need to estimate the maximal multiplicity of a sphere in Σ . For a $(d - k - j)$ -sphere $S \in \Sigma$, put

$$\mathcal{T}(S) = \{T \in \mathcal{T} \mid \sigma(T) = S\},$$

so $|\mathcal{T}(S)|$ is the multiplicity of S . Let u denote the center of S , and let π denote its axis. By Lemma 6.3.8, $U \subset \pi$, and all the tuples $T \in \mathcal{T}(S)$ are contained in π , and are anchored at $U \cup \{u\}$. Furthermore, we can reduce each k -tuple of $\mathcal{T}(S)$ into an anchored $(k - 1)$ -tuple (at the $j + 1$ points of $U \cup \{u\}$), since the last point (say) in each tuple is uniquely determined by the first $k - 1$ points and by $U \cup \{u\}$. Thus, the number of tuples in $\mathcal{T}(S)$, i.e., the multiplicity of S , can be bounded by $G_{k-2,k+j-1,j+1}(n) \leq 2G_{k-3,k+j-1,j+1}(n)$ (note that the values of the relevant parameters imply that the simplices counted in $G_{k-2,k+j-1,j+1}(n)$ are full-dimensional within their ambient space π).

For any multiplicity μ , we can take the $|\Sigma|/\mu$ spheres of Σ with the largest number of incidences, give them multiplicity μ , and discard the other spheres. In doing so, we have constructed a new multiset of spheres Σ' with the same number of elements, $|\Sigma'| = |\Sigma|$, and at least as many incidences, $I(P, \Sigma') \geq I(P, \Sigma)$. It thus follows, that the upper bound on $I(P, \Sigma)$ is maximized if all the spheres have maximal multiplicity. Hence,

$$I(P, \Sigma) \leq 2G_{k-3, k+j-1, j+1}(n) \cdot I' \left(n, \frac{G_{k-1, d, j}(n)}{2G_{k-3, k+j-1, j+1}(n)} \right), \quad (6.3.6)$$

where $I'(n, m)$ is the maximum possible number of incidences (counted without multiplicity) between n distinct points and m distinct $(d - k - j)$ -spheres in \mathbb{R}^d (for simplicity, we hide the parameters $d, d - k - j$ in this notation, and assume that the second parameter of I' in (6.3.6) is an integer). Note that the right hand side is multiplied by 2. This is done to compensate for possible rounding errors of the quotient $G_{k-1, d, j}(n)/2G_{k-3, k+j-1, j+1}(n)$. We have thus reduced the problem of bounding the number of anchored simplices into a problem involving incidences between points and *distinct* $(d - k - j)$ -spheres (generated by the k -tuples that span a prescribed face of the simplex).

As in Section 6.2, we classify the spheres into η -degenerate and η -nondegenerate spheres, where $\eta = \eta_{k, d}$ is a constant depending on k and d , and bound separately the number of incidences involving degenerate and nondegenerate spheres. We let Σ^D (resp., Σ^N) denote the subset of the degenerate (resp., nondegenerate) spheres of Σ , and put

$$G(P, \Delta, \Delta', U) = G^D(P, \Delta, \Delta', U) + G^N(P, \Delta, \Delta', U),$$

where $G^D(P, \Delta, \Delta', U)$ (resp., $G^N(P, \Delta, \Delta', U)$) is the number of k -tuples T that are counted in $G(P, \Delta, \Delta', U)$ and whose spheres $\sigma(T)$ are degenerate (resp., nondegenerate). Thus

$$\begin{aligned} G^D(P, \Delta, \Delta', U) &= I(P, \Sigma^D) \\ G^N(P, \Delta, \Delta', U) &= I(P, \Sigma^N), \end{aligned}$$

where, as above, incidences are counted with the multiplicities of their spheres.

As above, denote by $G_{k, d, j}^D(n)$ (resp., $G_{k, d, j}^N(n)$) the maximum possible value of $G^D(P, \Delta, \Delta', U)$ (resp., $G^N(P, \Delta, \Delta', U)$), over all choices of a set P of n points in \mathbb{R}^d , a $(k+j)$ -simplex Δ , a k -face Δ' of Δ , and a j -tuple U spanning a $(j-1)$ -simplex congruent to the face of Δ complementary to Δ' . We denote the polynomial degree of $G_{k, d, j}^D(n)$ (resp., $G_{k, d, j}^N(n)$) by $g_{k, d, j}^D$ (resp., $g_{k, d, j}^N$).

We next obtain, in two steps, the following simple but significant extension of the Elekes-Tóth bound, Theorem 1.5.1. In the first step we observe that a similar bound can be derived for j -flats rather than for hyperplanes. That is, if P is a set of n points in \mathbb{R}^d and H is a set of h η_{j+1} -nondegenerate j -flats, for some $1 \leq j \leq d-1$, then we can project the whole scenario onto some generic $(j+1)$ -dimensional subspace, observe that incidences, as well as nondegeneracy, are preserved in the projection, and apply Theorem 1.5.1, to obtain the bound

$$I(P, H) = O \left(n^{(j+1)/(j+2)} h^{(j+1)/(j+2)} + nh^{(j-1)/j} \right). \quad (6.3.7)$$

Next, let P be a set of n points in \mathbb{R}^d and Σ a set of s η_{j+2} -nondegenerate j -spheres in \mathbb{R}^d , for some $1 \leq j \leq d-1$. We lift \mathbb{R}^d , and in particular, the points of P and the spheres of Σ , to the standard paraboloid in \mathbb{R}^{d+1} . A well known property of this lifting transform is that any $(d-1)$ -sphere $\sigma \subset \mathbb{R}^d$ is lifted to a $(d-1)$ -ellipsoid in \mathbb{R}^{d+1} , which is the intersection of some hyperplane $\pi \subset \mathbb{R}^{d+1}$ (a d -flat) with the standard paraboloid in \mathbb{R}^{d+1} (see, e.g, [34]). Furthermore, the mapping $\sigma \mapsto \pi$ is one-to-one. Similarly, any j -sphere, for any $j \leq d-1$, in \mathbb{R}^d is lifted to some $(j+1)$ -flat in \mathbb{R}^{d+1} . Indeed, such a sphere is the intersection of some $d-j$ $(d-1)$ -spheres in general position in \mathbb{R}^d , which, after lifting, form $d-j$ d -hyperplanes in general position in \mathbb{R}^{d+1} , and their intersection is a $(j+1)$ -flat. Note that the lifting of a j -sphere in \mathbb{R}^d to a $(j+1)$ -flat in \mathbb{R}^{d+1} is also a one-to-one mapping. This implies that degeneracy and nondegeneracy of spheres are preserved by lifting. That is, a degenerate (resp. nondegenerate) j -sphere in \mathbb{R}^d with respect to P is lifted to a degenerate (resp. nondegenerate) $(j+1)$ -flat in \mathbb{R}^{d+1} with respect to the lifted points of P . We thus have transformed the above P and Σ into a set P^+ of n points in \mathbb{R}^{d+1} , and a family Σ^+ of s η_{j+2} -nondegenerate $(j+1)$ -flats in \mathbb{R}^{d+1} having the same incidence relations as P and Σ . Applying the previous bound (6.3.7), we obtain

$$I(P, \Sigma) = I(P^+, \Sigma^+) = O\left(n^{(j+2)/(j+3)} s^{(j+2)/(j+3)} + ns^{j/(j+1)}\right). \quad (6.3.8)$$

Incidences with nondegenerate spheres. By (6.3.8), the number of incidences between the n points of P and the s distinct nondegenerate $(d-k-j)$ -spheres in Σ^N is bounded by

$$O\left(n^{(d-k-j+2)/(d-k-j+3)} s^{(d-k-j+2)/(d-k-j+3)} + ns^{(d-k-j)/(d-k-j+1)}\right).$$

Substituting the worst-case estimate (cf. (6.3.6))

$$s = O^*(G_{k-1,d,j}(n)/G_{k-3,k+j-1,j+1}(n)) = O(n^{g_{k-1,d,j} - g_{k-3,k+j-1,j+1}}),$$

we get

$$\begin{aligned} \frac{G_{k,d,j}^N(n)}{G_{k-3,k+j-1,j+1}(n)} &= O\left((n^{(d-k-j+2)/(d-k-j+3)} s^{(d-k-j+2)/(d-k-j+3)} + ns^{(d-k-j)/(d-k-j+1)})\right) \\ &= O^*\left(n^{(1+g_{k-1,d,j} - g_{k-3,k+j-1,j+1})(d-k-j+2)/(d-k-j+3)}\right. \\ &\quad \left.+ n^{1+(g_{k-1,d,j} - g_{k-3,k+j-1,j+1})(d-k-j)/(d-k-j+1)}\right). \end{aligned}$$

Taking logarithms of both sides of the equation, and recalling that,

$$g_{k,d,j} = \max\{g_{k,d,j}^N, g_{k,d,j}^D\},$$

we conclude that one of the following inequalities must hold: Either

$$\begin{aligned} g_{k,d,j} - g_{k-3,k+j-1,j+1} &\leq (1 + g_{k-1,d,j} - g_{k-3,k+j-1,j+1})(d-k-j+2)/(d-k-j+3), \quad \text{or} \\ g_{k,d,j} - g_{k-3,k+j-1,j+1} &\leq 1 + (g_{k-1,d,j} - g_{k-3,k+j-1,j+1})(d-k-j)/(d-k-j+1), \quad \text{or} \\ g_{k,d,j} &\leq g_{k,d,j}^D. \end{aligned}$$

That is,

$$g_{k,d,j} \leq \max \begin{cases} \frac{(d-k-j+2)g_{k-1,d,j} + g_{k-3,k+j-1,j+1} + (d-k-j+2)}{d-k-j+3}, \\ \frac{(d-k-j)g_{k-1,d,j} + g_{k-3,k+j-1,j+1} + (d-k-j+1)}{d-k-j+1}, \\ g_{k,d,j}^D. \end{cases} \quad (6.3.9)$$

The easiest case is when $k = d - j - 1$. Then the s $(d - k - j)$ -spheres of Σ are actually circles (note that the simplices Δ in this case are $(d - 1)$ -simplices, and compare with the situation in Section 6.2), and there are no degeneracies. In this case, as reviewed in Section 1.2, we have a better incidence bound, namely the bound of Aronov et al. [15] (see also [6, 17, 63]) for point-circle incidences in d -space, for any $d \geq 3$, which is $I(n, s) = O^*(n^{6/11}s^{9/11} + n^{2/3}s^{2/3} + n + s)$. The implied bound on the number of simplices then becomes

$$\begin{aligned} G_{k,d,j}(n)/G_{k-3,d-2,j+1}(n) &= O^*(n^{6/11}s^{9/11} + (ns)^{2/3} + n + s) \\ &= O^*(n^{6/11+(9/11)(g_{k-1,d,j}-g_{k-3,d-2,j+1})} + \\ &\quad n^{(2/3)(1+g_{k-1,d,j}-g_{k-3,d-2,j+1})} + \\ &\quad n + n^{g_{k-1,d,j}-g_{k-3,d-2,j+1}}), \end{aligned}$$

or

$$g_{k,d,j} - g_{k-3,d-2,j+1} \leq \max \begin{cases} (6 + 9g_{k-1,d,j} - 9g_{k-3,d-2,j+1})/11, \\ 2(1 + g_{k-1,d,j} - g_{k-3,d-2,j+1})/3, \\ g_{k-1,d,j} - g_{k-3,d-2,j+1}, \\ 1, \end{cases}$$

or

$$g_{k,d,j} \leq \max \begin{cases} (6 + 9g_{k-1,d,j} + 2g_{k-3,d-2,j+1})/11, \\ (2 + 2g_{k-1,d,j} + g_{k-3,d-2,j+1})/3, \\ g_{k-1,d,j}, \\ g_{k-3,d-2,j+1} + 1. \end{cases} \quad (6.3.10)$$

Recall that here $k = d - j - 1$. To simplify this recurrence, let us define

$$h_{k,d} \stackrel{def}{=} \sup_{j \geq 0} g_{k,d+j,j}.$$

We have $g_{k,d,j} \leq h_{k,d-j}$, for any k, d, j . Substituting in (6.3.10), together with the fact that $k = d - j - 1$ gives

$$g_{k,d,j} \leq \max \begin{cases} (6 + 9h_{k-1,k+1} + 2h_{k-3,k-2})/11, \\ (2 + 2h_{k-1,k+1} + h_{k-3,k-2})/3, \\ h_{k-1,k+1}, \\ h_{k-3,k-2} + 1. \end{cases} \quad (6.3.11)$$

Since this system of inequalities applies for any d, j such that $k = d - j - 1$, we get

$$h_{k,k+1} \leq \max \begin{cases} (6 + 9h_{k-1,k+1} + 2h_{k-3,k-2})/11, \\ (2 + 2h_{k-1,k+1} + h_{k-3,k-2})/3, \\ h_{k-1,k+1}, \\ h_{k-3,k-2} + 1. \end{cases} \quad (6.3.12)$$

Put $a_k = h_{k,k+1}$, and use (for now) the naive bound $h_{k-1,k+1} \leq k$ (which follows since the tuples counted in $h_{k-1,k+1}$ are k -tuples of points of P). Then we have

$$a_k \leq \max \left\{ \frac{6 + 9k + 2a_{k-3}}{11}, \frac{2 + 2k + a_{k-3}}{3}, a_{k-3} + 1, k \right\}. \quad (6.3.13)$$

Trivially, we have $a_0 = 1$ (this involves 0-simplices, which can occur at most n times). a_1 involves repeated distances in the plane, or on a 2-sphere, and the best known upper bound for this count is $O(n^{4/3})$ (see [31, 75, 79], and Section 1.8), which is tight for the 2-sphere [45] (but is strongly suspected not to be tight in the plane). Thus $a_1 = 4/3$. a_2 counts congruent triangles in 3-space, or on a 3-sphere. Agarwal and Sharir [7] showed that the number of congruent triangles in 4-space is $G_{2,4}(n) = O^*(n^2)$. Their analysis can be applied verbatim to anchored triangles in \mathbb{R}^d with $d - 3$ anchor points, to yield $h_{2,3} = a_2 \leq 2$. Using the Lenz construction, one can construct $\Theta(n^2)$ congruent triangles in 4-space anchored at $\{0\}$, which shows that $a_2 \geq 2$, so $a_2 = 2$. The values of a_d , for $d \geq 3$ can be bounded recursively using (6.3.13) to yield: $a_3 \leq 3\frac{2}{11}$, $a_4 \leq 4\frac{2}{33}$, $a_5 \leq 5$, $a_6 \leq 6\frac{4}{121}$, $a_7 \leq 7\frac{4}{363}$, $a_8 \leq 8$. In fact, one can show inductively that, for each integer $k \geq 0$, we have

$$a_{3k} \leq 3k + \left(\frac{2}{11}\right)^k, \quad a_{3k+1} \leq 3k + 1 + \frac{1}{3} \left(\frac{2}{11}\right)^k, \quad \text{and} \quad a_{3k+2} \leq 3k + 2.$$

Since $g_{d-1,d} = g_{d-1,d,0} \leq h_{d-1,d}$, we obtain the following result:

Corollary 6.3.11. *The number of $(d - 1)$ -simplices spanned by a set of n points in \mathbb{R}^d and congruent to a given $(d - 1)$ -simplex is $O^*(n^{d-1+\alpha_d})$, where*

$$\alpha_d = \begin{cases} 0, & d \equiv 0 \pmod{3} \\ \left(\frac{2}{11}\right)^{(d-1)/3}, & d \equiv 1 \pmod{3}, \\ \frac{1}{3} \left(\frac{2}{11}\right)^{(d-2)/3}, & d \equiv 2 \pmod{3}. \end{cases}$$

In what follows, we further improve this bound, by further refining the analysis, replacing the naive bound $h_{k-1,k+1} \leq k$ by a sharper bound.

Extending and refining the analysis. To bound $g_{k,d,j}$, for $k < d - j - 1$, we need to analyze incidences with higher-dimensional spheres (i.e., of dimension at least 2). In this case we also need to take $g_{k,d,j}^D$ into account, that is, we need to bound the number of incidences involving degenerate spheres (see (6.3.9)).

Replace each degenerate $(d - k - j)$ -sphere $S \in \Sigma^D$ with a $(d - k - j - 1)$ -subsphere $S' = S'(S) \subset S$ containing a η -fraction of the points of $P \cap S$. Define the multiset

$$\Sigma' = \{S'(S) \mid S \in \Sigma^D\},$$

where each $(d - k - j - 1)$ -sphere appears with appropriate multiplicity.

The next task is to estimate, for a given $(d - k - j - 1)$ -sphere $S' \in \Sigma'$, the multiplicity of S' , i.e., the number of k -tuples $T \in \mathcal{T}$ that satisfy $\sigma(T) \supset S'$. The answer depends on whether or not S' is a great subsphere, i.e., whether the radius of S' is equal to or smaller than the radius of the corresponding sphere of Σ^D . If S' is a great subsphere, then, by Lemma 6.3.9, any k -tuple $T \in \mathcal{T}$ such that $\sigma(T) \supset S'$, is anchored at $j + 1$ points in the $(k + j)$ -dimensional axis of S' , and the last point (say) in T is uniquely determined by the first $k - 1$ points of the tuple. Thus, the number of such tuples is bounded by $G_{k-2, k+j, j+1}(n)$. If, on the other hand, S' is not a great subsphere, then, by Lemma 6.3.10, any k -tuple $T \in \mathcal{T}$ such that $\sigma(T) \supset S'$, is anchored at $j + 1$ points in the $(k + j)$ -dimensional axis of S' (without the above dependence of one point on the other $k - 1$). Thus, the number of such tuples is bounded by $G_{k-1, k+j, j+1}(n) \leq 2G_{k-2, k+j, j+1}(n)$ (the simplices counted in the left-hand side are full-dimensional within the relevant ambient space). Either way, the maximal multiplicity of a $(d - k - j - 1)$ -subsphere is at most $2G_{k-2, k+j, j+1}(n) = O^*(n^{h_{k-2, k-1}})$.

Before continuing, let us recall the setup as defined above. We have

1. A set P of n points in \mathbb{R}^d ;
2. A $(k + j)$ -simplex Δ consisting of j vertices that are anchor points, and another $k + 1$ vertices that can be chosen from the points of P ;
3. A k -face Δ' of Δ which represents the $k + 1$ vertices that can be chosen from the points of P ;
4. A set U of j anchor points spanning a simplex congruent to the face of Δ complementary to Δ' ;
5. A facet Δ'' of Δ' (i.e., a $(k - 1)$ -face), with which respect we look for k -tuples $T \subset P$ such that $\Delta(T) \cong \Delta''$;
6. A set \mathcal{T} of all k -tuples of P that span a simplex congruent to Δ'' and anchored at U ;
7. A multiset Σ of all $(d - k - j)$ -spheres $\sigma(T)$, for all $T \in \mathcal{T}$, which is partitioned to nondegenerate spheres Σ^N and degenerate spheres Σ^D , $\Sigma = \Sigma^N \cup \Sigma^D$; and
8. A multiset Σ' of $(d - k - j - 1)$ -spheres that serve as ‘witnesses’ for the degeneracy of all members of Σ^D , i.e., for each sphere in Σ^D , there is a $(d - k - j - 1)$ -subsphere in Σ' containing a η -fraction of its incidences.

We classify the spheres in Σ' into degenerate and nondegenerate spheres, and bound the number of incidences within each class separately. Denote by $G^{DD}(P, \Delta, \Delta', \Delta'', U)$ the number of incidences with the degenerate spheres of Σ' , and by $G_{k, d, j}^{DD}(n) = O^*(n^{g_{k, d, j}^{DD}})$ the maximum value

of this quantity over all sets P of n points in \mathbb{R}^d , all appropriate simplices $\Delta, \Delta', \Delta''$, and sets U . (Inductively, one can define $g_{k,d,j}^{D^3}, g_{k,d,j}^{D^4}$, etc., by partitioning the spheres further and further into classes of increasing degeneracy, within progressively lower-dimensional subspheres. But we stop at $g_{k,d,j}^{DD}$.)

The number $I = I(n, s)$ of incidences between n points and s nondegenerate $(d - k - j - 1)$ -spheres of maximal multiplicity μ can be bounded by the Elekes-Tóth bound for incidences between points and nondegenerate $(d - k - j)$ -flats (see Theorem 1.5.1 and (6.3.8)), so

$$\begin{aligned} I &= O \left(\mu \left(\left(n \cdot \frac{s}{\mu} \right)^{\frac{d-k-j+1}{d-k-j+2}} + n \left(\frac{s}{\mu} \right)^{\frac{d-k-j-1}{d-k-j}} \right) \right) \\ &= O \left(n^{\frac{d-k-j+1}{d-k-j+2}} s^{\frac{d-k-j+1}{d-k-j+2}} \mu^{\frac{1}{d-k-j+2}} + n s^{\frac{d-k-j-1}{d-k-j}} \mu^{\frac{1}{d-k-j}} \right). \end{aligned}$$

Substituting $s \leq G_{k-1,d,j}(n) = O^*(n^{g_{k-1,d,j}})$, and $\mu \leq 2G_{k-2,k+j,j+1}(n) = O^*(n^{g_{k-2,k+j,j+1}})$, we get

$$I = O^* \left(n^{\frac{(d-k-j+1)+(d-k-j+1)g_{k-1,d,j}+g_{k-2,k+j,j+1}}{d-k-j+2}} + n^{\frac{(d-k-j)+(d-k-j-1)g_{k-1,d,j}+g_{k-2,k+j,j+1}}{d-k-j}} \right).$$

Thus we have, for $k > d - j - 2$,

$$g_{k,d,j}^D \leq \max \begin{cases} \frac{(d-k-j+1)g_{k-1,d,j} + g_{k-2,k+j,j+1} + (d-k-j+1)}{d-k-j+2}, \\ \frac{(d-k-j-1)g_{k-1,d,j} + g_{k-2,k+j,j+1} + (d-k-j)}{d-k-j}, \\ g_{k,d,j}^{DD}. \end{cases} \quad (6.3.14)$$

If $k = d - j - 2$, then the spheres of Σ' are actually circles, so instead of the Elekes-Tóth bound, we apply the stronger bound of Aronov et al. [15] for point-circle incidences in space, and obtain

$$\begin{aligned} I &= O^* \left(n^{6/11} s^{9/11} \mu^{2/11} + n^{2/3} s^{2/3} \mu^{1/3} + s + n\mu \right) \\ &= O^* \left(n^{(6+9g_{k-1,d,j}+2g_{k-2,k+j,j+1})/11} + n^{(2+2g_{k-1,d,j}+g_{k-2,k+j,j+1})/3} + n^{g_{k-1,d,j}} + n^{1+g_{k-2,k+j,j+1}} \right), \end{aligned}$$

and then, the bound on $g_{k,d,j}^D$ becomes

$$g_{k,d,j}^D \leq \max \begin{cases} \frac{9g_{k-1,d,j} + 2g_{k-2,k+j,j+1} + 6}{11}, \\ \frac{2g_{k-1,d,j} + g_{k-2,k+j,j+1} + 2}{3}, \\ g_{k-1,d,j}, \\ g_{k-2,k+j,j+1} + 1, \end{cases} \quad (6.3.15)$$

(recalling that there are no degenerate circles). If we substitute (6.3.14) in recurrence (6.3.9), we

get, for $k > d - j - 2$,

$$g_{k,d,j} \leq \max \left\{ \begin{array}{l} \frac{(d-k-j+2)g_{k-1,d,j} + g_{k-3,k+j-1,j+1} + (d-k-j+2)}{d-k-j+3}, \\ \frac{(d-k-j)g_{k-1,d,j} + g_{k-3,k+j-1,j+1} + (d-k-j+1)}{d-k-j+1}, \\ \frac{(d-k-j+1)g_{k-1,d,j} + g_{k-2,k+j,j+1} + (d-k-j+1)}{d-k-j+2}, \\ \frac{(d-k-j-1)g_{k-1,d,j} + g_{k-2,k+j,j+1} + (d-k-j)}{d-k-j}, \\ g_{k,d,j}^{DD} \end{array} \right\},$$

in the general case, and, for $k = d - j - 2$, using (6.3.15),

$$g_{k,d,j} \leq \max \left\{ \begin{array}{l} \frac{4g_{k-1,d,j} + g_{k-3,k+j-1,j+1} + 4}{5}, \\ \frac{2g_{k-1,d,j} + g_{k-3,k+j-1,j+1} + 3}{3}, \\ \frac{9g_{k-1,d,j} + 2g_{k-2,k+j,j+1} + 6}{11}, \\ \frac{2g_{k-1,d,j} + g_{k-2,k+j,j+1} + 2}{3}, \\ g_{k-1,d,j}, \\ g_{k-2,k+j,j+1} + 1. \end{array} \right\}$$

Note that, in this case, even though the spheres of Σ are 2-spheres, we may not use the improved bound from Chapter 5, because it holds only in \mathbb{R}^3 and not in general \mathbb{R}^d . We can rewrite the latter system of inequalities in terms of $h_{k,d} = \sup_{j \geq 0} g_{k,d+j,j}$ as follows

$$h_{k,k+2} \leq \max \left\{ \begin{array}{l} \frac{4h_{k-1,k+2} + h_{k-3,k-2} + 4}{5}, \\ \frac{2h_{k-1,k+2} + h_{k-3,k-2} + 3}{3}, \\ \frac{9h_{k-1,k+2} + 2h_{k-2,k-1} + 6}{11}, \\ \frac{2h_{k-1,k+2} + h_{k-2,k-1} + 2}{3}, \\ h_{k-1,k+2}, \\ h_{k-2,k-1} + 1. \end{array} \right\}$$

If we use the naive bound $h_{k-1,k+2} \leq k$ (simplices counted in $h_{k-1,t}$ for any t have k nondetermined vertices), and put $b_k = h_{k,k+2}$, we get

$$b_k \leq \max \left\{ \frac{4k + a_{k-3} + 4}{5}, \frac{2k + a_{k-3} + 3}{3}, \frac{9k + 2a_{k-2} + 6}{11}, \frac{2k + a_{k-2} + 2}{3}, k, a_{k-2} + 1 \right\}, \quad (6.3.16)$$

and the recurrence (6.3.12) for a_k becomes (compared with (6.3.13), where we have used the naive bound $b_{k-1} \leq k$)

$$a_k \leq \max \left\{ \frac{2a_{k-3} + 9b_{k-1} + 6}{11}, \frac{a_{k-3} + 2b_{k-1} + 2}{3}, a_{k-3} + 1, b_{k-1} \right\}. \quad (6.3.17)$$

Consider the initial values of b_k . Recall that this is the exponent in a bound on the number of k -simplices anchored at any number j of points, in $k + j + 2$ dimensions. We have $b_0 = 1$ (trivial), $b_1 = 2$ (which is clearly an upper bound, and the Lenz construction produces $n^2/4$ unit-length segments in \mathbb{R}^4 anchored at 0), $b_2 = 7/3$ (see Agarwal and Sharir [7], who show that $g_{2,5} = 7/3$, and note that (a) their construction of $\Theta(n^{7/3})$ congruent triangles in \mathbb{R}^5 makes all the triangles anchored at 0, and (b) their upper bound analysis, based on cuttings in \mathbb{R}^5 , carries over to embeddings of the 4-sphere in \mathbb{R}^d , and hence, holds also for b_2). Recall that $a_0 = 1, a_1 = 4/3, a_2 = 2$. By the above recurrences (6.3.16) and (6.3.17), the next bounds on a_k and b_k are: $a_3 \leq 2.636, b_3 \leq 3.4, a_4 \leq 3.570, b_4 \leq 4.267, a_5 \leq 4.4, b_5 \leq 5.2, a_6 \leq 5.279, b_6 \leq 6.127, a_7 \leq 6.208, b_7 \leq 7.114, a_8 \leq 7.166, b_8 \leq 8.08, a_9 \leq 8.116, b_9 \leq 9.056, a_{10} \leq 9.083, b_{10} \leq 10.042$. If we continue to calculate a_k and b_k using these recurrences, we see that a_k is asymptotic to $k - 1$, and b_k is asymptotic to k . More rigorously, put $a_k = k - 1 + a_k^*$ and $b_k = k + b_k^*$. Then the recurrences become

$$b_k^* \leq \max \left\{ \frac{a_{k-3}^*}{5}, \frac{a_{k-3}^* - 1}{3}, \frac{2a_{k-2}^*}{11}, \frac{a_{k-2}^* - 1}{3}, 0, a_{k-2}^* - 2 \right\},$$

and

$$a_k^* \leq \max \left\{ \frac{2a_{k-3}^* + 9b_{k-1}^*}{11}, \frac{a_{k-3}^* + 2b_{k-1}^* - 1}{3}, a_{k-3}^* - 2, b_{k-1}^* \right\}.$$

Lemma 6.3.12. *For all $k \geq 5$, we have $a_k^* \leq 2^{-k/4}$ and $b_k^* \leq \frac{2}{3}2^{-k/4}$. That is,*

$$a_k \leq k - 1 + 2^{-k/4}, \quad \text{and}$$

$$b_k \leq k + \frac{2}{3}2^{-k/4}.$$

Proof. From the above values, one can see that the lemma holds for $k = 5, 6, 7$. For larger k , we have, using induction on k ,

$$\begin{aligned} b_k^* &\leq \frac{a_{k-3}^*}{5} \leq \frac{2^{3/4}}{5}2^{-k/4}, \quad \text{or} \\ b_k^* &\leq \frac{a_{k-3}^* - 1}{3} \leq \frac{2^{3/4}}{3}2^{-k/4} - \frac{1}{3}, \quad \text{or} \\ b_k^* &\leq \frac{2a_{k-2}^*}{11} \leq \frac{2\sqrt{2}}{11}2^{-k/4}, \quad \text{or} \\ b_k^* &\leq \frac{a_{k-2}^* - 1}{3} \leq \frac{\sqrt{2}}{3}2^{-k/4} - \frac{1}{3}, \quad \text{or} \\ b_k^* &\leq 0, \quad \text{or} \\ b_k^* &\leq a_{k-2}^* - 2 \leq \sqrt{2} \cdot 2^{-k/4} - 2. \end{aligned}$$

All the inequalities imply $b_k^* \leq \frac{2}{3}2^{-k/4}$. As for a_k^* , we have

$$\begin{aligned} a_k^* &\leq \frac{2a_{k-3}^* + 9b_{k-1}^*}{11} \leq \frac{2 \cdot 2^{3/4} + 6 \cdot 2^{1/4}}{11} 2^{-k/4}, \quad \text{or} \\ a_k^* &\leq \frac{a_{k-3}^* + 2b_{k-1}^* - 1}{3} \leq \frac{2^{3/4} + \frac{4}{3} \cdot 2^{1/4}}{3} 2^{-k/4} - \frac{1}{3}, \quad \text{or} \\ a_k^* &\leq a_{k-3}^* - 2 \leq 2^{3/4} \cdot 2^{-k/4} - 2, \quad \text{or} \\ a_k^* &\leq b_{k-1}^* \leq \frac{2^{5/4}}{3} \cdot 2^{-k/4}. \end{aligned}$$

The resulting inequalities imply $a_k^* \leq 2^{-k/4}$; the second inequality is the only one that is not upper bounded by $2^{-k/4}$ for every k , but it does satisfy this bound for $k \geq 8$. \square

Corollary 6.3.13. *For any $d \geq 5$,*

1. *The number of $(d-1)$ -simplices spanned by a set of n points in \mathbb{R}^d and congruent to a given $(d-1)$ -simplex is*

$$O^* \left(n^{d-2+2^{-(d-1)/4}} \right).$$

2. *The number of $(d-2)$ -simplices spanned by a set of n points in \mathbb{R}^d and congruent to a given $(d-2)$ -simplex is*

$$O^* \left(n^{d-2+\frac{2}{3}2^{-(d-2)/4}} \right).$$

6.4 Similar simplices

Using the tools developed above, we can derive improved bounds on the number of similar simplices.

Lemma 6.4.1.

$$F_{k,d}(n) \leq O(n^2)G_{k-2,d,2}(n-2), \quad \text{so}$$

$$f_{k,d} \leq g_{k-2,d,2} + 2 \leq h_{k-2,d-2} + 2.$$

Proof. Let P be a set of n points in \mathbb{R}^d , let Δ be a k -simplex, let Δ' be some $(k-2)$ -face of Δ , and let $p_1, p_2 \in P$ be a fixed pair of points from P . Put $P' = P \setminus \{p_1, p_2\}$.

Consider a $(k-1)$ -tuple T such that $\Delta(T) \sim \Delta'$, and $\Delta(T \cup \{p_1, p_2\}) \sim \Delta$. Since the edge length $|p_1 - p_2|$ of the edge p_1p_2 of $\Delta(T \cup \{p_1, p_2\})$ is fixed, all these simplices are mutually congruent. Put $\tilde{\Delta} = \Delta(T \cup \{p_1, p_2\})$ and $\tilde{\Delta}' = \Delta(T)$. Then the number of ways to complete $\{p_1, p_2\}$ into a $(k+1)$ -tuple of points of P' that spans a simplex similar to Δ is

$$G(P', \tilde{\Delta}, \tilde{\Delta}', \{p_1, p_2\}) \leq G_{k-2,d,2}(n-2).$$

The total number of similar simplices can be bounded by multiplying this bound by the number of pairs of points of P , and so we get $F_{k,d}(n) \leq O(n^2)G_{k-2,d,2}(n-2)$, as claimed. \square

For $k = d - 1$ and $k = d - 2$, we can use the bounds on $h_{k,k+1}$, and $h_{k,k+2}$ derived in the previous section. The resulting bounds yield the following result.

Theorem 6.4.2. *The number of mutually similar $(d - 1)$ -simplices spanned by n points in \mathbb{R}^d is $O^*(n^{f_{d-1,d}})$, and the number of such $(d - 2)$ -simplices is $O^*(n^{f_{d-2,d}})$, where*

$$f_{d-1,d} \leq d - 2 + 2^{3/4} \cdot 2^{-d/4}, \quad \text{and}$$

$$f_{d-2,d} \leq d - 2 + \frac{4}{3} 2^{-d/4}.$$

This improves the previous bounds from the conference version of this chapter [4] (coauthored with Agarwal, Purdy, and Sharir) of $f_{d-1,d} \leq d - 72/55$, and $f_{d-2,d} \leq d - 8/5$, for $d \geq 8$.

6.5 Similar triangles in five dimensions

Let P be a set of n points in \mathbb{R}^5 , and let Δ_0 be a triangle. For each pair of points $a, b \in P$, let $\sigma_{a,b}$ denote the 3-sphere orthogonal to ab and containing all the points c for which $\Delta(abc) \sim \Delta_0$. Let Σ be the set of resulting 3-spheres. As in Section 6.2, it is easily checked that no 3-sphere can arise in this way more than twice. Ignoring this constant multiplicity, we face the problem of bounding the number of incidences between P and a set Σ of $O(n^2)$ 3-spheres.

As in earlier sections, we fix a parameter k and a sufficiently small constant $\eta > 0$. We define a 3-sphere to be light, heavy, nondegenerate, or degenerate, with respect to these two parameters, as in Section 6.2. The number of incidences involving light 3-spheres is $O(n^2k)$, so we concentrate on the heavy 3-spheres.

Consider first the *nondegenerate* heavy 3-spheres of Σ . In this case, by lifting the points and 3-spheres into \mathbb{R}^6 , and then projecting them onto some generic 5-space, we can apply, as before, the Elekes-Tóth bound [41], to conclude that the number of such 3-spheres is $O(n^5/k^6 + n^4/k^4)$ (provided that η is small enough), and that the number of incidences between these 3-spheres and the points of P is $O(n^5/k^5 + n^4/k^3)$.

Consider next the *degenerate* heavy 3-spheres in Σ , so each of them contains a 2-sphere that contains more than a η -fraction of the points of P on the 3-sphere. We replace the 3-spheres by the respective 2-spheres, bound the number of incidences with these 2-spheres, counted with the appropriate multiplicities, and lose only a constant factor using this bound. Consider first the case where the 2-spheres themselves are nondegenerate, in the sense that none of them contains a circle that contains more than a η -fraction of the points on the 2-sphere. (For simplicity, we use the same constant η , although the value of η for 2-spheres may be different from its value for 3-spheres.) Since these 2-spheres are (ηk) -heavy, the number of distinct such 2-spheres is, using another variant of the Elekes-Tóth bound, $O(n^4/k^5 + n^3/k^3)$, and the number of incidences with them is $O(n^4/k^4 + n^3/k^2)$.

Here however the 2-spheres may appear with multiplicity, but we claim that the maximum multiplicity of a 2-sphere is at most $O(n)$. Indeed, given a 2-sphere σ' with center o' , if we fix one point a in the defining pair (a, b) of a 3-sphere $\sigma_{a,b}$ containing σ' , the size of the corresponding

triangle is determined, and the center o of $\sigma_{a,b}$ must then lie at a fixed distance from a , within the axis π , which is a 2-plane orthogonal to σ' and passing through its center o' . Also, o lies on another fixed circle in π centered at o' . These two circles intersect at most twice (assuming $a \neq o'$, which can always be guaranteed), so at most two points b can form with a a pair (a, b) for which $\sigma_{a,b} \supset \sigma'$. Hence, the number of triangles similar to Δ_0 that fall into this subcase is $O(n^5/k^4 + n^4/k^2)$.

Finally, consider the subcase where the 2-spheres themselves are degenerate, so we replace each of them by a respective $(\eta^2 k)$ -heavy circle, and bound the number of incidences between these circles and the points of P . The multiplicity of a circle is, trivially, at most $O(n^2)$. Hence, arguing as above, the number of triangles that arise in this subcase is $O((n^5 \log k)/k^{9/2} + n^4/k^2)$. Hence, the overall number of triangles is $O(n^2 k + n^5/k^4 + n^4/k^2)$. Choosing $k = n^{2/3}$, we thus obtain

Theorem 6.5.1. $F_{2,5}(n) = O(n^{8/3})$.

As already noted, $d = 5$ is the last interesting case for triangles, since, already for the congruent case, $G_{2,6}(n) = \Theta(n^3)$ [7].

6.6 Discussion

Examining the proof of the general bound in Section 6.3, we note that there are several sources for potential improvements. First, the proof uses the naive estimate $g_{d-3,d+j,j} \leq d - 2$; one should be able to get a better, nontrivial bound. Indeed, the previous section implies that this is the case when $d = 5$ and $j = 0$; that is, in this case we have $g_{d-3,d,0} \leq f_{d-3,d} \leq d - 7/3$ (for $d = 5$). Another room for improvement is in the estimation of the multiplicities of the spheres that arise in the analysis. We look at the axis π of a sphere, and make the worst case assumption that $|\pi \cap P| = n$. With a more careful analysis (e.g., using the Elekes-Tóth bound), we expect to be able to improve this considerably.

Another observation, already mentioned in Section 6.4, is that we can relate $F_{k,d}(n)$ to $G_{k-2,d,2}(n)$ (the maximum number of $(k - 2)$ -simplices *congruent* to a given simplex and *anchored* at two points), through the following inequality

$$F_{k,d}(n) \leq O(n^2)G_{k-2,d,2}(n).$$

Recall the Erdős-Purdy conjecture that $g_{k-2,d} \leq d/2$ (for even d). If the conjecture were true then we would have $F_{k,d}(n) = O(n^{2+d/2})$, which, for large values of d , is significantly smaller than the general bounds derived in this paper.

Another observation is that the proof technique is essentially a careful analysis of incidences between points and spheres of various dimensions. While the case of circles has already been studied fairly intensively, the case of higher-dimensional spheres has not received much attention. The bounds that we obtain via the Elekes-Tóth bound seem to be weak, because they do not exploit the property that, after lifting, the points lie on the standard (and convex) paraboloid. For example, using this technique for estimating the number of k -rich circles would yield a bound of $O(n^3/k^4 + n^2/k^2)$, whereas the bound (1.2.2) is $O^*(n^3/k^{11/2} + n^2/k^3 + n/k)$. One would hope that

similar improvements could be obtained for incidences with higher-dimensional spheres too, e.g., by sharpening the analysis of [41] for point sets lying in convex position, or on a lower-dimensional surface.

We have recently achieved an improvement in this direction for incidences between points and nondegenerate 2-spheres in \mathbb{R}^3 [14] (coauthored with Sharir, see Chapter 5), which, in turn, helped us in Section 6.2 to improve the bound for similar triangles in \mathbb{R}^3 . This improvement, however, applies only for 2-spheres, and only in \mathbb{R}^3 , and still leaves a gap from the lower bound construction. We leave the problem of closing these gaps for future research.

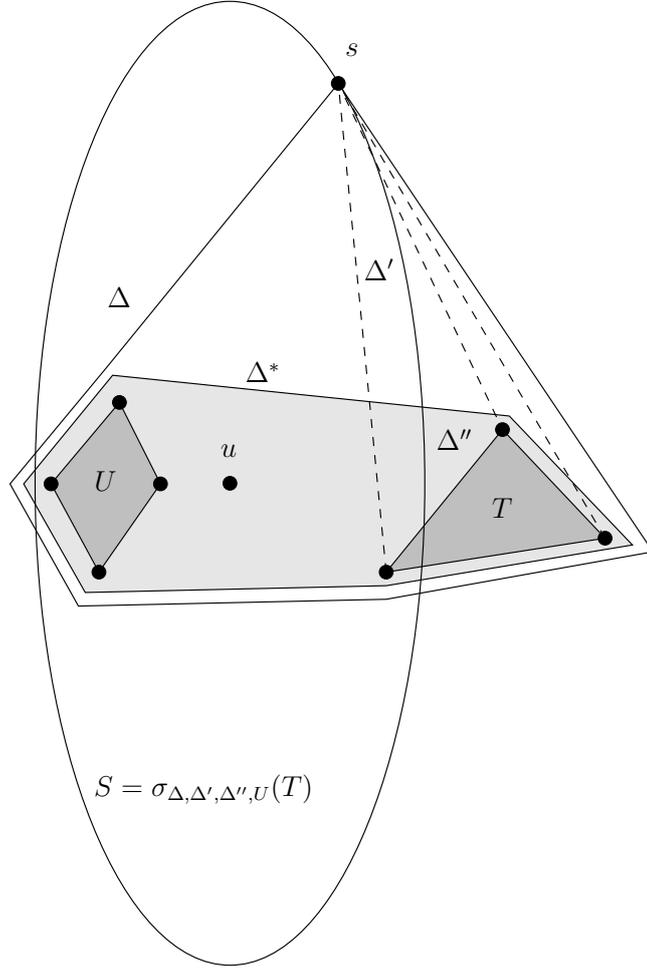


Figure 6.4: Illustrating the claims of Lemmas 6.3.7, 6.3.8, 6.3.9, and 6.3.10. Δ is a $(k + j)$ -simplex in \mathbb{R}^d , Δ' is a k -face of Δ , Δ'' is a $(k - 1)$ -face (a facet) of Δ' , and Δ^* is a $(k + j - 1)$ -face (a facet) of Δ consisting of all the vertices of Δ except for the vertex of Δ' which is missing from Δ'' (s in the figure). U is a j -tuple of anchor points spanning a simplex congruent to the $(j - 1)$ -face of Δ complementary to Δ' , and T is a k -tuple anchored at U with respect to Δ^* and Δ'' , meaning that T spans a simplex congruent to Δ'' and $T \cup U$ span a simplex congruent to Δ^* . S is the set of all points s such that $T \cup \{s\}$ is anchored at U with respect to Δ and Δ' , i.e., $T \cup \{s\}$ spans a simplex congruent to Δ' and $T \cup U \cup \{s\}$ spans a simplex congruent to Δ . We denote S by $\sigma_{\Delta, \Delta', \Delta'', U}(T)$. By Lemma 6.3.7, S is a $(d - k - j)$ -sphere (whose center is labeled here by u). Furthermore, the axis of S is the affine hull of $T \cup U$. Lemma 6.3.8 says that, given $\Delta, \Delta', \Delta'', U$, and S as above, any tuple T such that $S = \sigma_{\Delta, \Delta', \Delta'', U}(T)$ must lie in the axis π of S and be anchored at $U \cup \{u\}$ with respect to some $(k + j - 1)$ -simplex $\tilde{\Delta} \subset \pi$ and a $(k - 2)$ -simplex Δ''' which is a face of both $\tilde{\Delta}$ and Δ'' , and that these are exactly all the tuples that satisfy $S = \sigma_{\Delta, \Delta', \Delta'', U}(T)$. Lemmas 6.3.9 and 6.3.10 are similar to Lemma 6.3.8, but assume that S is not the entire sphere $\sigma_{\Delta, \Delta', \Delta'', U}(T)$, but some subsphere of smaller dimension. They assert that any tuple T such that $\sigma_{\Delta, \Delta', \Delta'', U}(T) \supset S$, must lie in the axis π of S and be anchored at $U \cup \{u\}$ with respect to some $(k + j - 1)$ -simplex $\tilde{\Delta}$ and a $(k - 1)$ -face or $(k - 2)$ -face of $\tilde{\Delta}$ (and of Δ''), and that these are exactly all the tuples with this property.

Bibliography

- [1] ÁBREGO, B. M., ELEKES, G., AND FERNÁNDEZ-MERCHANT, S. Structural results for plane sets with many similar subsets. *Combinatorica* 24 (2004), 541–554.
- [2] ÁBREGO, B. M., AND FERNÁNDEZ-MERCHANT, S. On the maximum number of equilateral triangles, i. *Discrete Comput. Geom.* 23 (2000), 129–136.
- [3] ÁBREGO, B. M., AND FERNÁNDEZ-MERCHANT, S. The unit distance problem for centrally symmetric convex polygons. *Discrete Comput. Geom.* 28 (2002), 467–473.
- [4] AGARWAL, P. K., APFELBAUM, R., PURDY, G., AND SHARIR, M. Similar simplices in a d -dimensional point set. In *Proc. 23rd Annu. ACM Sympos. Comput. Geom.* (2007), pp. 232–238.
- [5] AGARWAL, P. K., AND ARONOV, B. Counting facets and incidences. *Discrete Comput. Geom.* 7 (1992), 359–369.
- [6] AGARWAL, P. K., NEVO, E., PACH, J., PINCHASI, R., SHARIR, M., AND SMORODINSKY, S. Lenses in arrangements of pseudocircles and their applications. *J. ACM* 51 (2004), 139–186.
- [7] AGARWAL, P. K., AND SHARIR, M. On the number of congruent simplices in a point set. *Discrete Comput. Geom.* 28 (2002), 123–150.
- [8] AJTAI, M., CHVÁTAL, V., NEWBORN, M., AND SZEMERÉDI, E. Crossing-free subgraphs. *Ann. Discrete Math.* 12 (1982), 9–12.
- [9] AKUTSU, T., TAMAKI, H., AND TOKUYAMA, T. Distribution of distances and triangles in a point set and algorithms for computing the largest common point set. *Discrete Comput. Geom.* 20 (1998), 307–331.
- [10] ALON, N., LAST, H., PINCHASI, R., AND SHARIR, M. On the complexity of arrangements of circles in the plane. *Discrete Comput. Geom.* 26 (2001), 465–492.
- [11] APFELBAUM, R., AND SHARIR, M. Repeated angles in three and four dimensions. *SIAM J. Discrete Math.* 19 (2005), 294–300.

- [12] APFELBAUM, R., AND SHARIR, M. Large complete bipartite subgraphs in incidence graphs of points and hyperplanes. *SIAM J. Discrete Math.* 21 (2007), 707–725.
- [13] APFELBAUM, R., AND SHARIR, M. An improved bound on the number of unit area triangles. *Discrete Comput. Geom.* 44 (2010), 753–761.
- [14] APFELBAUM, R., AND SHARIR, M. Nondegenerate spheres in three dimensions. *Combinat. Probab. Comput.* 20 (2011), 503–512.
- [15] ARONOV, B., KOLTUN, V., AND SHARIR, M. Incidences between points and circles in three and higher dimensions. *Discrete Comput. Geom.* 33 (2005), 185–206.
- [16] ARONOV, B., PACH, J., SHARIR, M., AND TARDOS, G. Distinct distances in three and higher dimensions. *Combinat. Probab. Comput.* 13 (2004), 283–293.
- [17] ARONOV, B., AND SHARIR, M. Cutting circles into pseudo-segments and improved bounds for incidences. *Discrete Comput. Geom.* 28 (2002), 475–490.
- [18] ATKINSON, M. D. An optimal algorithm for geometric congruence. *J. Algorithms* 8 (1987), 159–172.
- [19] BOXER, L. Point set pattern matching in 3-d. *Pattern Recog. Letts.* 17 (1996), 1293–1297.
- [20] BRASS, P. Exact point pattern matching and the number of congruent triangles in a three-dimensional point set. In *Proc. European Sympos. Algos., Lecture Notes in Computer Science* (2000), vol. 1879, Springer Verlag, pp. 112–119.
- [21] BRASS, P. Combinatorial geometry problems in pattern recognition. *Discrete Comput. Geom.* 28 (2002), 495–510.
- [22] BRASS, P., AND KNAUER, C. On counting point-hyperplane incidences. *Comput. Geom. Theory Appl.* 25 (2003), 13–20.
- [23] BRASS, P., MOSER, W., AND PACH, J. *Research Problems in Discrete Geometry*. Springer Verlag, New York, 2005.
- [24] BRASS, P., AND PACH, J. The maximum number of times the same distance can occur among the vertices of a convex n -gon is $O(n \log n)$. *J. Combinat. Theory, Ser. A* 94 (2001), 178–179.
- [25] CHAN, T. Optimal partition trees. *Discrete Comput. Geom.* 47 (2012), 661–690. Preliminary version in *Proc. 26th Annu. ACM Sympos. Comput. Geom.* (2010) 1–10.
- [26] CHAN, T. M. On levels in arrangements of curves. *Discrete Comput. Geom.* 29 (2003), 375–393.
- [27] CHAZELLE, B. Cuttings. In *Handbook of Data Structures and Applications*, D. Mehta and S. Sahni, Eds. Chapman and Hall/CRC Press, Boca Raton, FL, 2005, ch. 25.

- [28] CHAZELLE, B., EDELSBRUNNER, H., GUIBAS, L. J., POLLACK, R., SEIDEL, R., SHARIR, M., AND SNOEYINK, J. Counting and cutting cycles of lines and rods in space. *Comput. Geom. Theory Appl. 1* (1992), 305–323.
- [29] CHAZELLE, B., AND FRIEDMAN, J. A deterministic view of random sampling and its use in geometry. *Combinatorica 10* (1990), 229–249.
- [30] CLARKSON, K. A randomized algorithm for closed-point queries. *SIAM J. Comput. 17* (1988), 830–847.
- [31] CLARKSON, K., EDELSBRUNNER, H., GUIBAS, L., SHARIR, M., AND WELZL, E. Combinatorial complexity bounds for arrangements of curves and spheres. *Discrete Comput. Geom. 5* (1990), 99–160.
- [32] DE REZENDE, P. J., AND LEE, D. T. Point set pattern matching in d -dimensions. *Algorithmica 13* (1995), 387–404.
- [33] DUMITRESCU, A., SHARIR, M., AND TÓTH, C. Extremal problems on triangle areas in two and three dimensions. *J. Combinat. Theory, Ser. A* (2009), 1177–1198.
- [34] EDELSBRUNNER, H. *Algorithms in Combinatorial Geometry*. Springer Verlag, Heidelberg, 1987.
- [35] EDELSBRUNNER, H., GUIBAS, L., AND SHARIR, M. The complexity of many cells in arrangements of planes and related problems. *Discrete Comput. Geom. 5* (1990), 197–216.
- [36] EDELSBRUNNER, H., AND HAJNAL, P. A lower bound on the number of unit distances between vertices of a convex polygon. *J. Combinat. Theory, Ser. A 56* (1991), 312–316.
- [37] ELEKES, G. Sums versus products in number theory, algebra and Erdős geometry. In *Paul Erdős and his Mathematics II*, G. Halász, Ed. János Bolyai Math. Soc., Budapest, 2002, pp. 241–290.
- [38] ELEKES, G., AND ERDŐS, P. Similar configurations and pseudo-grids. In *Intuitive Geometry* (1994), K. Böröczky and G. F. Tóth, Eds., vol. 63, Colloq. Math. Soc. János Bolyai, pp. 85–104.
- [39] ELEKES, G., KAPLAN, H., AND SHARIR, M. On lines, joints, and incidences in three dimensions. *J. Combinat. Theory, Ser. A 118* (2011), 962–977. Also in [arXiv:0905.1583](https://arxiv.org/abs/0905.1583).
- [40] ELEKES, G., AND SHARIR, M. Incidences in three dimensions and distinct distances in the plane. *Combinat. Probab. Comput. 20* (2011), 571–608. Also in [arXiv:1005.0982](https://arxiv.org/abs/1005.0982).
- [41] ELEKES, G., AND TÓTH, C. D. Incidences of not too degenerate hyperplanes. In *Proc. 21st Annu. ACM Sympos. Comput. Geom.* (2005), pp. 16–21.
- [42] ERDŐS, P. On sets of distances of n points. *Amer. Math. Monthly 53* (1946), 248–250.

- [43] ERDŐS, P. On sets of distances of n points in euclidean space. *Magyar Tud. Akad. Mat. Kutató Int. Kozl.* 5 (1960), 165–169.
- [44] ERDŐS, P. Problems and results in combinatorial geometry. *Annals of the New York Academy of Sciences* 440 (1985), 1–11.
- [45] ERDŐS, P., HICKERSON, D., AND PACH, J. A problem of Leo Moser about repeated distances on the sphere. *Amer. Math. Monthly* 96 (1989), 569–575.
- [46] ERDŐS, P., AND PURDY, G. Some extremal problems in geometry. *J. Combinat. Theory* 10 (1971), 246–252.
- [47] ERDŐS, P., AND PURDY, G. Some extremal problems in geometry iv. In *Proc. 7th South-Eastern Conf. Combinatorics, Graph Theory, and Comput.* (1976), pp. 307–322.
- [48] ERDŐS, P., AND PURDY, G. Extremal problems in combinatorial geometry. In *Handbook of Combinatorics*, R. Graham, M. Grötschel, and L. Lovász, Eds., vol. 1. Elsevier, Amsterdam, 1995, pp. 809–874.
- [49] FÜREDI, Z. The maximum number of unit distances in a convex n -gon. *J. Combinat. Theory, Ser. A* 55 (1990), 316–320.
- [50] GUTH, L., AND KATZ, N. H. Algebraic methods in discrete analogs of the kakeya problem. *Adv. Math.* 225 (2010), 2828–2839. Also in [arXiv:0812.1043](https://arxiv.org/abs/0812.1043).
- [51] GUTH, L., AND KATZ, N. H. On the Erdős distinct distance problem in the plane. [arXiv:1011.4105](https://arxiv.org/abs/1011.4105) (2010).
- [52] HAUSSLER, D., AND WELZL, E. ϵ -nets and simplex range queries. *Discrete Comput. Geom.* 2 (1987), 127–151.
- [53] KAPLAN, H., MATOUŠEK, J., SAFERNOVA, Z., AND SHARIR, M. Unit distances in three dimensions. *Combinatorics, Probability and Computing* 21, 04 (2012), 597–610. Also in [arXiv:1107.5391](https://arxiv.org/abs/1107.5391).
- [54] KAPLAN, H., MATOUŠEK, J., AND SHARIR, M. Simple proofs of classical theorems in discrete geometry via the Guth-Katz polynomial partitioning technique. *Discrete Comput. Geom.* 48 (2012), 499–517. Also in [arXiv:1102.5391](https://arxiv.org/abs/1102.5391).
- [55] KAPLAN, H., SHARIR, M., AND SHUSTIN, E. On lines and joints. *Discrete Comput. Geom.* 44 (2010), 838–843. Also in [arXiv:0906.0558](https://arxiv.org/abs/0906.0558).
- [56] KATZ, N. H., AND TARDOS, G. A new entropy inequality for the Erdős distance problem. In *Towards a Theory of Geometric Graphs*, J. Pach, Ed., vol. 342 of *Contemporary Mathematics*. Amer. Math. Soc., Providence, RI, 2004, pp. 119–126.

- [57] KÓVARI, T., SÓS, V. T., AND TURÁN, P. On a problem of K. Zarankiewicz. *Colloquium Math.* 3 (1954), 50–57.
- [58] KOLTUN, V. Almost tight upper bounds for vertical decompositions in four dimensions. *J. ACM* 51 (2004), 699–730.
- [59] KOLTUN, V. Sharp bounds for vertical decompositions of linear arrangements in four dimensions. *Discrete Comput. Geom.* 31 (2004), 435–460.
- [60] LACZKOVICH, M., AND RUZSA, I. The number of homothetic subsets. In *The Mathematics of Paul Erdős*, R. Graham and J. Nešetřil, Eds., vol. 2. Springer Verlag, 1997, pp. 294–302.
- [61] LEIGHTON, F. T. New lower bound techniques for VLSI. *Math. Systems Theory* 17 (1984), 47–70.
- [62] LENZ, H. Zur zerlegung von punktmengen in solche kleineren durchmessers. *Arch. Math.* 6 (1955), 413–416.
- [63] MARCUS, A., AND TARDOS, G. Intersection reverse sequences and geometric applications. *J. Combinat. Theory Ser. A* 113 (2006), 675–691.
- [64] MATOUŠEK, J. Efficient partition trees. *Discrete Comput. Geom.* 8 (1992), 315–334.
- [65] MILNOR, J. On the Betti numbers of real varieties. *Proc. Amer. Math. Soc.* 15 (1964), 275–280.
- [66] OLEINIK, O. A., AND PETROVSKII, I. G. On the topology of real algebraic surfaces. *Izvestiya Akad. Nauk SSSR. Ser. Mat.* (1949), 389–402.
- [67] PACH, J., RADOIČIĆ, R., TARDOS, G., AND TÓTH, G. Improving the crossing lemma by finding more crossings in sparse graphs. *Discrete Comput. Geom.* 36 (2006), 527–552.
- [68] PACH, J., AND SHARIR, M. Repeated angles in the plane and related problems. *J. Combinat. Theory Ser. A* 59 (1992), 12–22.
- [69] PACH, J., AND SHARIR, M. On the number of incidences between points and curves. *Combinat. Probab. Comput.* 7 (1998), 121–127.
- [70] PACH, J., AND SHARIR, M. Geometric incidences. In *Towards a Theory of Geometric Graphs*, J. Pach, Ed., vol. 342 of *Contemporary Mathematics*. Amer. Math. Soc., Providence, RI, 2004, pp. 185–223.
- [71] QUILODRÁN, R. The joints problem in \mathbb{R}^n . *SIAM J. Discrete Math.* 23 (2010), 2211–2213.
- [72] SHAFAREVICH, I. R. *Basic Algebraic Geometry*. Springer Verlag, 1977.
- [73] SOLYMOSI, J., AND TAO, T. An incidence theorem in higher dimensions. *Discrete Comput. Geom.* 48 (2012), 255–280. Also in arXiv:1103.2926.

- [74] SOLYMOSI, J., AND VU, V. Near optimal bounds for the Erdős distinct distances problem in high dimensions. *Combinatorica* 28 (2006), 113–125.
- [75] SPENCER, J., SZEMERÉDI, E., AND TROTTER, W. Unit distances in the Euclidean plane. In *Graph Theory and Combinatorics (Proc. Cambridge Conf. on Combinatorics)* (1984), B. Bollobas, Ed., Academic Press, pp. 293–308.
- [76] STONE, A. H., AND TUKEY, J. W. Generalized “sandwich” theorems. *Duke Math. J.* 9, 2 (1942), 356–359.
- [77] SWANEPOEL, K. Unit distances and diameters in Euclidean spaces. *Discrete Comput. Geom.* 41 (2009), 1–27.
- [78] SWANEPOEL, K. J., AND VALTR, P. The unit distance problem on spheres. In *Towards a Theory of Geometric Graphs*, J. Pach, Ed., vol. 342 of *Contemporary Mathematics*. Amer. Math. Soc., Providence, RI, 2004, pp. 273–279.
- [79] SZÉKELY, L. A. Crossing numbers and hard Erdős problems in discrete geometry. *Combinat. Probab. Comput.* 6 (1997), 353–358.
- [80] SZEMERÉDI, E., AND TROTTER, W. Extremal problems in discrete geometry. *Combinatorica* 3 (1983), 381–392.
- [81] TAMAKI, H., AND TOKUYAMA, T. How to cut pseudo-parabolas into segments. *Discrete Comput. Geom.* 19 (1998), 265–290.
- [82] THOM, R. Sur l’homologie des variétés algébriques réelles. In *Differential and Combinatorial Topology (A Symposium in Honor of Marston Morse)* (1965), Princeton Univ. Press, Princeton, N.J., pp. 255–265.
- [83] VALTR, P. Strictly convex norms allowing many unit distances and related touching questions. Manuscript, Charles University, Prague, 2005.
- [84] ZAHL, J. An improved bound on the number of point-surface incidences in three dimensions. [arXiv:1104.4987](https://arxiv.org/abs/1104.4987) (2011).