THE POLYNOMIAL METHOD IN INCIDENCE GEOMETRY

BRYAN CHEN, JULIA WU, EMILY XIE

Abstract. In this paper, we introduce and consider primary use cases of the polynomial method in the context of incidence geometry, expanding on ideas in Dvir's survey paper [1]. By looking at problems that vary from directly being modeled around incidences to those that have less obvious initial ties to incidence theory, we observe the potential results that can be achieved by using the polynomial method. We also consider the significance of these bounds in historical and practical contexts.

Contents

1. Introduction

The polynomial method is a technique that introduces polynomials into problems with no immediately obvious connection to polynomials. The technique relies on a few simple but key observations about polynomials and their zeros. The polynomial method has been a hot topic in math research this last decade for providing simple solutions to longstanding open problems. The problems examined in this paper are from incidence geometry, the study of possible combinatorial configurations between geometric objects such as lines and points.

An *incidence* describes an intersection between two geometric objects. For example, a line can be incident to a plane, or a line can be incident to another line. This paper primarily concerns incidences between different lines and incidences between points and lines. Figure 1 illustrates an example of both types of incidences, with ℓ_1 and ℓ_2 illustrating an example of two incident lines. Although Figure 1 only

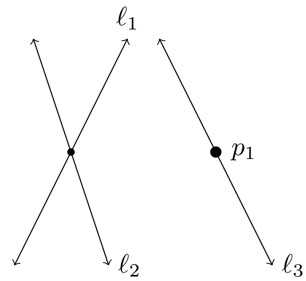

FIGURE 1. The intersection of ℓ_1 and ℓ_2 is an incidence, as is the intersection of p_1 and ℓ_3.

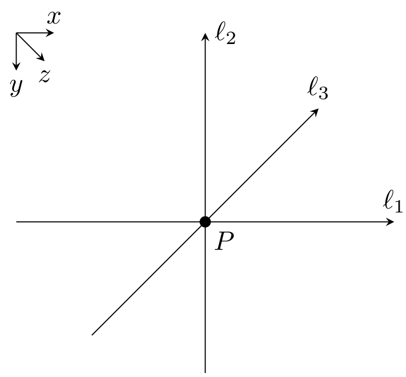

FIGURE 2. The intersection of ℓ_1, ℓ_2, and ℓ_3 at P is a joint.

shows an incidence involving two lines, multiple lines can be incident at the same point.

The first problem we examine in this paper involves incidences between at least three non-coplanar lines in $\mathbb{R}^3$. The point at which these lines are incident is called a joint, as illustrated in Figure 2. A simple question one can ask about joints is what is the maximal number of joints that can arise in N lines. This question is deceptively simply stated, but was actually considered quite hard. In fact, there were a long line of papers that proved incremental improvements to an upper bound with very complicated proof methods [1]. The problem was finally fully answered by Guth and Katz only very recently using the polynomial method. In this paper we show how to introduce polynomials into this problem and how to obtain a tight upper bound on the number of joints in N lines.

The second problem this paper discusses, at first glance, appears unrelated to incidences. It involves distinct distances determined by a set of points. For example, the set of points $P = \{p_1, p_2, p_3, p_4\}$ in Figure 3 form a square. Then it is simple to see that P contains 2 unique distances. One question posed by Erdős, which is still an open problem, asks what is the minimum number of distinct distances determined by a set of points in the real plane? An almost optimal lower bound was obtained by Guth and Katz using the polynomial method[2].

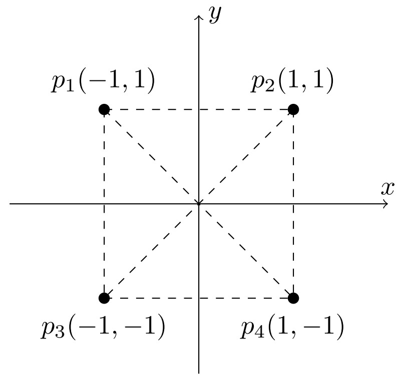

FIGURE 3. The set $P = \{p_1, p_2, p_3, p_4\}$ determines 2 unique distances.

The proof of Guth and Katz's bound involves a clever way of reducing the problem into an incidence counting problem in $\mathbb{R}^3$. They showed that a specific upper bound for the number of incidences between N lines in $\mathbb{R}^3$ implies the desired lower bound for distinct distances in points in $\mathbb{R}^2$. This problem of upper bounding incidences in lines is similar to the joints problem, with the addition that points where only two lines intersect are also counted. Then some ideas in the two proofs will be related. Again, the polynomial method is used to obtain the the upper bound, specifically through a novel technique called "polynomial cell partitioning." Polynomial cell partitioning is used to partition $\mathbb{R}^3$ into "cells" using polynomials, and the number of incidences in each cell can be bounded using old, familiar bounds on incidences. In this paper, we show how polynomial cell partitioning is obtained, how it is used to obtain the desired bound on the number of incidences, and how this implies the desired bound on the Erdős distinct distance problem.

The organization of the paper is as follows. Section 2 develops the main components of the polynomial method. Section 3 proves the tight upper bound for the number of joints given by the polynomial method. Section 4 introduces the Erdős distinct distance problem and shows how to reduce it to an incidence counting problem. We introduce the incidence bound proposed by Guth and Katz and how it implies the distinct distance bound. Finally, Section 5 completes the proof of the bound given by Guth and Katz by showing how to bound the number of incidences in $\mathbb{R}^3$ using polynomial cell partitioning. We also show how polynomial cell partitioning gives an upper bound for incidences in $\mathbb{R}^2$ as a warm-up to the proof of the bound in $\mathbb{R}^3$.

2. Preliminaries

The problems in this paper involve sets of points in the real plane. For sets of points, a common way to introduce polynomials into the problem is finding a polynomial that vanishes on those points. The core of the polynomial method is bounding the attributes of the polynomial in a clever way that allows us to bound the original problem. Such attributes include the degree of the polynomial and the number of zeros.

In this section we develop these ideas which are in the background of the polynomial method. We show that given a finite set of points, we can find a small degree polynomial vanishing on those points. We are only interested in *non-zero polynomials*, where we say a polynomial is non-zero if it has at least one non-zero coefficient.

Theorem 2.1. *A non-zero univariate polynomial g(x) can have at most deg(g) real zeros.*

This bound is a well known fact, so no proof is given. It applies to univariate polynomials, which is sufficient for most of the paper. A variant of this bound for multivariate polynomials is introduced in Section 5.

If we can find a way to define a polynomial g such that all the points in a finite set S are zeros of g, we can bound the size of S by the degree of g. We show we can define such a polynomial g with a small degree given some finite set of points. We present the theorem for finite sets of points in $\mathbb{R}^n$, but we will only use it for points in $\mathbb{R}^2$.

Theorem 2.2. *Let $S \subset \mathbb{R}^n$ be a finite set. If $|S| < \binom{n+d}{d}$, then there exists a non zero polynomial $g \in \mathbb{R}[x_1, ..., x_n]$ of degree $\leq d$ such that $g(p) = 0$ for all $p \in S$.*

This theorem holds for any finite field $\mathbb{F}$, but for our applications in this paper we only see sets over the reals. The idea of the proof is that given a set of points $S \subset \mathbb{R}^n$ where $|S| < \binom{n+d}{d}$, we consider the family of polynomials with n variables $g(x_1, ..., x_n)$ with degree d. If we plug in each point $p \in S$, we can write $g(p) = 0$ and obtain a system of $|S|$ linear homogeneous equations in the coefficients of g. We find that g has more coefficients than $|P|$, meaning our system of equations has more variables than equations, which implies a non-zero solution for g.

Proof. Given a finite set $S \subset \mathbb{R}^n$, we consider the vector space of the polynomials over $\mathbb{R}^n$ that have n variables and fixed degree d such that $|S| < \binom{n+d}{d}$. These polynomials have form $g(x_1, ..., x_n)$. We can see g has $\binom{n+d}{d}$ monomials using a simple bars and stars argument. Restricting these polynomials to our set S gives a linear map $g \mapsto (g(s_1), g(s_2), ..., g(s_{|S|})) \in \mathbb{R}^{|S|}$ whose domain has dimension $\binom{n+d}{d}$, and whose range has dimension $|P| < \binom{n+d}{d}$. Then our linear map has a non trivial null space, implying there is a nonzero polynomial g that vanishes on S. □

Corollary 2.3. *There exists a constant C depending only on the dimension n such that for any finite set $S \subset \mathbb{R}^n$, there exists a degree $d \leq C|S|^{1/n}$ polynomial that vanishes on S.*

Proof. Plugging in $C|S|^{1/n}$ into the expression $\binom{n+d'}{d'}$, we obtain

$$\begin{aligned} \binom{n+d'}{d'} &\geq c' \cdot \frac{\lceil C|S|^{1/n} + n \rceil !}{\lceil C|S|^{1/n} \rceil ! n!} \\ &> c' \cdot \frac{(C|S|^{1/n})^n}{n!} \\ &= c' \cdot \frac{C^n |S|}{n!}. \end{aligned}$$

This quantity is at least $|S|$ for $C > (n!)^{1/n}/c'$, which depends only on n since c' is just a small constant to round $C|S|^{1/n}$. Then by Theorem 2.2, there exists a polynomial of degree $d \leq C|S|^{1/n}$ that vanishes on S. □

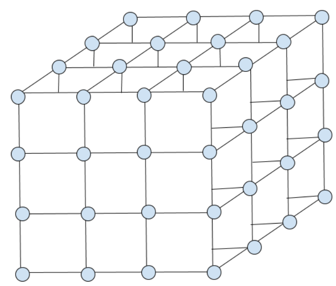

FIGURE 4. This $N \times N \times N$ grid of lines has N^3 joints.

The result from this theorem satisfies our need to find a vanishing polynomial g of small degree on any given set of points $S \subset \mathbb{R}^n$. Since we can then bound the number of zeros of g, we can bound the number of points $|S|$. This idea is the main component of the polynomial method, but in general the polynomial method always requires some additional algebraic claims that depend on the specific problem. This will become more apparent after seeing how the polynomial method works in the context of the joints problem in the next section.

3. The Joints Problem

The first problem we examine through the lens of the polynomial method is the joints problem. The joints problem concerns the intersections of lines in $\mathbb{R}^3$.

Definition 3.1. Let L be a set of lines in $\mathbb{R}^3$. A **joint** w.r.t the arrangement L is a point $p \in \mathbb{R}^3$ through which pass at least three, non coplanar, lines.

A basic question we can ask is what the maximal number of joints possible in an arrangement of lines L is. This question was first posed in the 90's in the context of computer graphics algorithms [1] [3]. The initial paper posing this question proved that the number of joints in a set of lines L is upper bounded by $c|L|^{7/4}$ for some constant c. There were a long line of papers incrementally improving this exponent using various complicated techniques until the question was ultimately solved completely by Guth and Katz using the polynomial method in 2015.

Theorem 3.2 (Guth-Katz)**.** *Let L be a set of lines in $\mathbb{R}^3$ and J be the set of joints defined by L. Then, $|J| \le C \cdot |L|^{3/2}$ for some constant C.*

We can see that this bound is optimal by considering an $N \times N \times N$ grid of lines, as shown in Figure 4. In other words, let L be the union of the following three sets, each containing N^2 lines:

$$L_{xy} = \{(i, j, t) | t \in \mathbb{R}\}, i, j \in [N]$$
$$L_{yz} = \{(t, i, j) | t \in \mathbb{R}\}, i, j \in [N]$$
$$L_{zx} = \{(i, t, j) | t \in \mathbb{R}\}, i, j \in [N]$$

Then $|L| = 3N^2$, and each point in $[N]^3$ is a joint. Then we have that the number of joints is $c|L|^{3/2}$, for $c = \frac{1}{3}$. Thus, we cannot give a smaller upper bound than the one given by Guth and Katz without making the bound invalid for some configurations.

The proof of the Guth Katz Theorem illustrates the power of the polynomial method. As mentioned in the previous section, the polynomial method generally requires additional algebraic claims dependent on the problem. For the proof of the Guth Katz theorem, we must restrict polynomials to the given set of lines L. To this end, we need to introduce some properties about the restrictions of polynomials to lines.

Definition 3.3. Let $g \in \mathbb{R}[x_1, \ldots, x_n]$ be a degree d polynomial and l be a line in $\mathbb{R}^n$. Parameterizing l as $l = a + tb$, where $a, b \in \mathbb{R}^n, t \in \mathbb{R}$, the **restriction** of g to l is $h(t) = g(a + tb)$.

This first property allows us to bound the degree of the restriction of a polynomial to a line.

Lemma 3.4. *Let $g \in \mathbb{R}[x_1, \ldots, x_n]$ be a polynomial, l be a line in $\mathbb{R}^n$, and h be the restriction of g to l. Then, $deg(h) \leq deg(g)$.*

Intuitively, $deg(h) \leq deg(g)$ since restricting the domain of a polynomial to a line could not yield a more expressive polynomial.

This next property allows us to bound the number of points a line can intersect a polynomial in.

Lemma 3.5. *Let $g \in \mathbb{R}[x_1, \ldots, x_n]$ be a polynomial, l be a line in $\mathbb{R}^n$, and h be the restriction of g to l. Then, l intersects g in at most $deg(g)$ points.*

Proof. From Lemma 3.4, $deg(h) \leq deg(g)$. Then, l contains $deg(h) \leq deg(g)$ roots of h. Since a root of h is also a root of g, l contains $\leq deg(g)$ roots of g. So, l intersects g in at most $deg(g)$ points. □

The final property allows us to represent the coefficients of the restriction in terms of gradients.

Lemma 3.6. *Let $g \in \mathbb{R}[x_1, \ldots, x_n]$ be a polynomial. Consider a line in $\mathbb{R}^n$, $l = a + tb$ where $a, b \in \mathbb{R}^n$. As in Definition 3.3, $h(t) = g(a + tb)$ is the restriction of g to l. Then, the coefficient of t in $h(t)$ is $\langle \nabla g(a), b \rangle$.*

Proof. To find the coefficient of t in h, we can take the derivative of h w.r.t. t and evaluate this derivative at $t = 0$. That is, we want to find

$$\frac{\partial}{\partial t} h(t) \Big|_{t=0} = \frac{\partial}{\partial t} g(a + tb) \Big|_{t=0}.$$

In order to evaluate this, we must find the partial derivatives of g. We can use the gradient of g to represent all the partial derivatives of g:

$$\nabla g = \begin{bmatrix} \frac{\partial g}{\partial x_1} & \frac{\partial g}{\partial x_2} & \cdots & \frac{\partial g}{\partial x_n} \end{bmatrix}.$$

Then, using the chain rule, we get

$$\begin{aligned} \frac{\partial}{\partial t} g(a + tb) &= \frac{\partial g}{\partial x_1} \frac{\partial x_1}{\partial t} + \ldots + \frac{\partial g}{\partial x_n} \frac{\partial x_n}{\partial t} \\ &= \frac{\partial g}{\partial x_1} b_1 + \ldots + \frac{\partial g}{\partial x_n} b_n \\ &= \langle \nabla g(a + tb), b \rangle. \end{aligned}$$

Evaluating at $t = 0$, we obtain our desired result

$$\langle \nabla g(a + tb), b\rangle = \langle \nabla g(a), b\rangle.$$

□

This lemma will allow us to determine what the gradient of g is from the restricted polynomial, which will be used to show a contradiction in the proof of the Guth and Katz bound.

Proof of Theorem 3.2. We use a proof by contradiction to show Guth and Katz's bound on the number of joints. We assume that the number of joints is larger than the bound $C \cdot |L|^{3/2}$ given in the theorem, i.e. $|J| > C \cdot |L|^{3/2}$ and consider the minimal degree nonzero polynomial g that vanishes on J. Then, we will show that g must also vanish on all lines in L. Finally, we will find a polynomial that contradicts either the minimality of g or the nonzero property of g, which implies that the Guth-Katz bound holds.

First, notice that lines that contain fewer than $\frac{|J|}{2|L|}$ joints contribute at most $|L| \cdot \frac{|J|}{2|L|} = \frac{|J|}{2}$ joints altogether. So, we can ignore these lines, since this would reduce $|J|$ by at most a factor of $\frac{1}{2}$.

Consider a nonzero polynomial $g(x, y, z)$ of minimal degree that vanishes on J. In other words, g is the smallest degree polynomial with at least one nonzero coefficient such that each joint in J is a root of g. Theorem 2.2 shows that a nonzero polynomial g with $deg(g) \leq 3|J|^{1/3}$ that vanishes on J exists (this can be verified by replacing S with J, n with 2, and d with $3|J|^{1/3}$ in the theorem).

Given this polynomial g, we now show that each line contains strictly more than $deg(g)$ joints. Since each line contains at least $\frac{|J|}{2|L|}$ joints and $deg(g) \leq 3|J|^{1/3}$, it suffices to show that $\frac{|J|}{2|L|} > 3|J|^{1/3}$.

Simplifying,

$$\begin{aligned} \frac{|J|}{2|L|} &> 3|J|^{1/3} \\ \iff |J|^{2/3} &> 6|L| \\ \iff |J| &> 6^{3/2}|L|^{3/2} \end{aligned}$$

That is, $\frac{|J|}{2|L|} > 3|J|^{1/3}$ is equivalent to $|J| > 6^{3/2}|L|^{3/2}$. By the initial assumption $|J| > C \cdot |L|^{3/2}$, the latter is true by choosing C large enough. So, each line contains strictly more than $deg(g)$ joints.

Given this, we will show that g vanishes on all lines in L. To do this, we will consider the restriction of g to a line and show that this resulting restricted polynomial vanishes on the line, which implies that g vanishes on the line.

Consider a line $l \in L$ and parameterize it as $l = a + tb$, where $a, b \in \mathbb{R}^3,\ t \in \mathbb{R}$. Restrict g to the line l using the univariate polynomial $f(t) = g(l) = g(a + tb)$.

We now show that f vanishes on l by contradiction. Assume f does not vanish on l. Then, by Lemma 3.5, l intersects g in at most $deg(g)$ points, so l contains at most $deg(g)$ roots of g. However, l must contain strictly more than $deg(g)$ roots of g. So, we have a contradiction, and thus f must vanish on l. The the same reasoning can be applied for each line $l \in L$ to show that g vanishes on all lines in L.

Finally, we will consider a single joint and three of the non-coplanar lines that pass through that joint. Then, we will use the fact that g vanishes on these three lines to show the existence of a polynomial that contradicts some property of g that we know to be true.

Consider a joint $p \in J$, and let l_1, l_2, l_3 be three non-coplanar lines that pass through p. Since l_1, l_2, l_3 are non-coplanar, their respective tangent vectors $u_1, u_2, u_3 \in \mathbb{R}^3$ are linearly independent. As before, we can parameterize l_i as $l_i = p + tu_i$ and restrict g to l_i using the polynomials $h_i(t) = g(l_i) = g(p + tu_i)$, where $t \in \mathbb{R}$, for $i = 1, 2, 3$. Since g vanishes on all three lines, $h_i(t)$ evaluates to 0 at all points on l. That is, $h_i(t)$ evaluates to 0 at any point in its domain. This implies that $h_i(t)$ is the zero polynomial, so all coefficients of h_i are 0. In particular, the coefficient of t is 0 in h_i. Applying Lemma 3.6, the coefficient of t in $h_i(t)$ is $\langle \nabla g(p), u_i \rangle = 0$. Since u_1, u_2, u_3 are linearly independent, they form a basis of $\mathbb{R}^3$. Then, $\nabla g(p) = 0$.

The same reasoning can be applied to all joints $p \in J$ to show that $\nabla g(p)$ vanishes on J. In addition, $deg(\nabla g) < deg(g)$. So, the three partial derivatives of g are all polynomials of smaller degree than g that vanish on J.

There are now two cases: $\nabla g(p)$ contains a nonzero polynomial (i.e. at least one of the partial derivatives of g is a nonzero polynomial), and $\nabla g(p)$ consists only of zero polynomials (i.e. all partial derivatives of g are the zero polynomial). We must find a polynomial that contradicts some known property of g in both cases.

Case 1: If $\nabla g(p)$ contains a nonzero polynomial, then g is not a minimal degree polynomial that vanishes on J. So, we have a contradiction for the minimality of g.

Case 2: If $\nabla g(p)$ contains only zero polynomials, then g has no nonzero partial derivatives. Then, g must be a constant polynomial. Since g is constant and evaluates to 0 at some point, then g must be the zero polynomial. Then we also have a contradiction for the nonzero property of g.

So, when $J > C \cdot |L|^{3/2}$, we can always find a contradiction for some property of the nonzero polynomial g of minimal degree that vanishes on J. Thus, we must have $|J| \leq C \cdot |L|^{3/2}$. □

This bound given by Guth and Katz is exciting because it solves a hard problem, which at first glance appears easy due to its simple problem statement. In the next section we examine another hard problem which is simple to state.

4. Erdős Distinct Distance Problem

The next problem we consider is the Erdős distinct distance problem, which seeks to lower bound the number of distinct distances defined by a set of points S in $\mathbb{R}^2$. Erdős posed several distance problems, of which the most famous is the distinct distance problem. These problems helped guide and shape the field of incidence geometry. We include this specific problem in our analysis of the polynomial method because its discussion involves several different types of uses of polynomials and illustrates a variety of techniques in the polynomial method. Furthermore, the Erdős distinct distance problem is still open, so it reveals the need for more refinements and ongoing research in the field of incidence geometry and the polynomial method.

Erdős initially posed the distinct distince problem in a 1946 paper, when he observed that a square grid of n points determines $cn/\sqrt{\log(n)}$ distinct distances, for some constant c. He conjectured that this bound was sharp, but proving this bound still remains an open problem.

In the most recent breakthrough, Guth and Katz showed that the number of distinct distances can be lower bounded by $cn/\log(n)$, almost obtaining the optimal bound conjectured by Erdős. This result was remarkable not only because it obtained the sharp exponent, but also because Guth and Katz used the polynomial method to prove the bound.

The lower bound shown by Guth and Katz is tight up to a factor of $\sqrt{\log(n)}$.

Theorem 4.1 (Guth-Katz)**.** *Let S be some set of points in $\mathbb{R}^2$ and $\boldsymbol{dist}(S) = \{||p-q|| \,|\, p,q \in S\}$. Then*

$$|\boldsymbol{dist}(S)| \geq \frac{c|S|}{\log(|S|)}.$$

The way Guth and Katz proved this result involved transforming the distance counting problem to an incidence counting problem between lines in $\mathbb{R}^3$, using a reduction via the framework established by Elekes and Sharir[6].

After reducing the problem, Guth and Katz showed the corresponding desired bound on the number of incidences between N lines in $\mathbb{R}^3$ using the polynomial method. In this section we show this reduction and introduce the desired incidence bound. The method of reducing the problem involves studying the problem in the terms of rigid motions of the plane.

First, note that the problem of counting distinct distances between pairs of points in S is equivalent to counting the number of repeated distances, since we can simply take the complement with respect to the total number of pairs of distances to obtain the number of unique distances. We are now interested in quadruples of points $a, b, c, d \in S$ where $||a-b|| = ||c-d||$.

Consider the problem in a new framework of *rigid motions*, which are transformations in $\mathbb{R}^2$ involving only translations and rotations. Note that for any four points $a, b, c, d \in S$, we have that $||a-b|| = ||c-d||$ if and only if the segment $\overline{ab}$ can be transformed via some rigid motion to be the same as the segment $\overline{cd}$, i.e. $\exists T : \mathbb{R}^2 \to \mathbb{R}^2$ such that $T(a) = c$ and $T(b) = d$, as shown in Figure 5.

Note that any rigid motion T can be expressed by three parameters: two for the translation and one for the rotation. Thus, if $\mathcal{R}$ is the set of possible rigid motions in $\mathbb{R}^2$, then it is a three-dimensional space in terms of the parameters used to define rigid motions. For each $a, b \in S$, define $L_{a,b} = \{T \in \mathcal{R} | T(a) = b\}$ as the set of all rigid motions that map a to b. Each $L_{a,b}$ will be some one-dimensional shape in $\mathcal{R}$ since specifying a starting and ending point for each rigid motion removes two degrees of freedom. This can be seen by observing that once a is translated to b, a can still be rotated freely around b. Also denote L as the set of all $L_{a,b}$ for $a, b \in S$.

Consider the intersection of some two $L_{a,c}$ and $L_{b,d}$ at a point $T \in \mathcal{R}$. This point represents a rigid motion, and since it is the intersection we have $T \in L_{a,c}$ and $T \in L_{b,d}$. This means that $T(a) = c$ and $T(b) = d$, which implies $||a-b|| = ||c-d||$ since we have a rigid transformation that now maps $\overline{ab}$ to $\overline{cd}$. In the reverse direction, any repeated distance implies the existence of an intersection between two lines $L_{a,c}$ and $L_{b,d}$ at a common rigid motion. Thus, an incidence between two lines in L represents a bijection with a repeated distance, as shown in Figure 6.

Our goal is to find some parameterization of the rigid motions such that the $L_{a,b}$ are actually lines in our 3D space $\mathcal{R}$. This allows us to count the number of distinct distances in our original set S by bounding the number of incidences between the

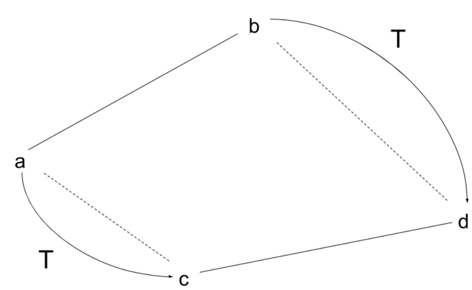

FIGURE 5. A rigid motion doesn't change the length of the segment

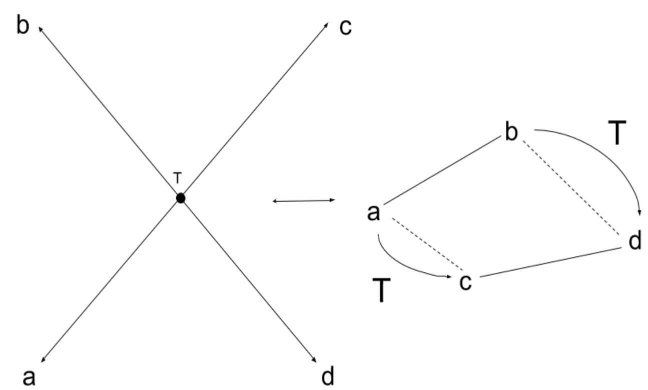

FIGURE 6. The same rigid motion T has $T(a) = c$ and $T(b) = d$

lines L. This is the fundamental basis of our reduction from a distance counting problem in $\mathbb{R}^2$ to an incidence problem in $\mathbb{R}^3$.

Given this goal, we first proceed with analysis of our reduction to gain intuition on what we require of our parameterization. Consider the set $Q(S) = \{(a, b, c, d) \in S^4 \big| ||a - b|| = ||c - d||\}$, the set of points corresponding to repeated distances. By applying the Cauchy-Schwarz Inequality over all pairs of pairs of points determining

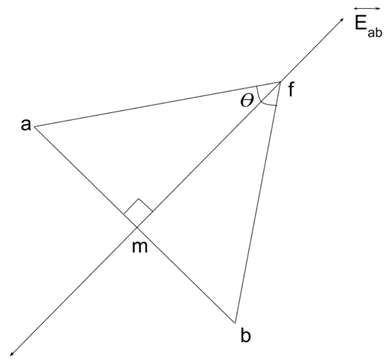

FIGURE 7. The rotation point f is located arbitrarily along E_{ab}

distances $||a-b||$ and $||c-d||$, we can obtain the inequality

$$|\mathbf{dist}(S)| \geq \frac{|S|^4}{|Q(S)|}$$

Thus, to show that the number of distinct distances satisfies

$$|\mathbf{dist}(S)| \geq \frac{c|S|}{\log(|S|)}$$

we can instead focus on bounding the number of repeated distances by showing

$$|Q(S)| \leq c'|S|^3 \log(|S|)$$

which is equivalent to bounding the number of incidences

$$|I(L)| = |\{(l,l') \in L^2 | l \cap l' \neq \phi\}|.$$

We now construct a parameterization that causes the $L_{a,b}$ to actually be lines in $\mathcal{R}$. The key intuition is to parameterize the angle of rotation θ so that it scales linearly with the possible transformations from some a to b. First, note that we can essentially ignore rigid motions that are pure translations. Any three points in S uniquely define a 4-tuple in $Q(S)$ since an a and b define a translation, with c then defining d under this translation. This means there are at most $|S|^3$ 4-tuples in $Q(S)$ arising from pure translations, so they already fit under our $|S|^3\log(|S|)$ desired asymptotic bound.

We now consider rigid motions that are comprised of nonzero rotations. Note that this is equivalent to rotations around a specific point. These can then be defined by some counter-clockwise rotation by some $\theta \in (0, 2\pi)$ around some fixed point $f = (f_x, f_y)$. Note that for any such rigid motion mapping point a to point b, the fixed point f must lie on the perpendicular bisector of $\overline{ab}$, as shown in Figure 7.

We define our rigid motions with the parameterization $\rho : \mathcal{R} \to \mathbb{R}^3$ where $\rho(T) = (f_x, f_y, \frac{1}{2}\cot(\frac{\theta}{2}))$. The core idea is to show that for any fixed motion T mapping point a to point b, all three coefficients will scale linearly along the perpendicular bisector of $\overline{ab}$, denoted as $E_{a,b}$.

Lemma 4.2. *For each $a, b \in S$, we have that $\rho(L_{a,b})$ is a line in $\mathbb{R}^3$.*

Proof of Lemma 4.2. Looking at $\triangle afm$, where m is the midpoint of $\overline{ab}$, we can calculate $\cot\left(\dfrac{\theta}{2}\right) = \dfrac{||f-m||}{\frac{||a-b||}{2}} \Rightarrow \frac{1}{2}\cot(\frac{\theta}{2}) = \frac{||f-m||}{||a-b||}$.

Now, let $d = (d_x, d_y)$ be some unit vector parallel to $E_{a,b}$. Since f must lie on $E_{a,b}$, we can express $f = m + ||f-m|| \cdot d = m + td$, where $t = ||f-m||$ denotes how far f is from m.

We then have $\rho(L_{a,b}) = \{(m_x, m_y, 0) + t(d_x, d_y, ||a-b||^{-1}) \big| t \in \mathbb{R}\}$, which is a line. □

Now that we have completed our reduction, we are primarily interested in bounding the number of incidences between N lines in $\mathbb{R}^3$. A successful bound for the number of incidences $I(L)$ in terms of the number of points in our set S is

$$|I(L)| \leq c|S|^3 \log(|S|),$$

as this would give the desired bound for the number of distinct distances. The number of lines is equal to the number of pairs of points in S, i.e. $|L| = \binom{|S|}{2} \leq |S|^2$. Then the following theorem shown by Guth and Katz gives us exactly the bound we desire.

Theorem 4.3 (Guth-Katz)**.** *Let L be a set of N lines in $\mathbb{R}^3$ such that no more than $\sqrt{N}$ lines intersect at a single point and no plane or doubly ruled surface contains more than $\sqrt{N}$ lines. Then the number of incidences of lines in L, $|I(L)|$, is at most $cN^{1.5} \log N$ for some constant c.*

Note that setting $N = |S|^2$ gives our desired bound. The theorem excludes certain scenarios of lines, and below we explain why it is possible to discard those cases.

First, if more than $\sqrt{N} = |S|$ lines intersect at a single point, then by the Pigeonhole Principle we can find some two pairs (a, b) and (a, b') where $b \neq b'$ such that $L_{a,b} \cap L_{a,b'} \neq \phi$. This means $\exists T$ such that $T(a) = b$ and $T(a) = b'$, which contradicts the fact that $b \neq b'$. Thus, no more than $\sqrt{N}$ lines intersect at a single point.

Second, if more than $\sqrt{N} = |S|$ lines are contained within a single plane, then we can again find two pairs (a, b) and (a, b') such that $T(a) = b = b'$ and $b \neq b'$ as before. This is because due to our lines' parameterization, no two lines are parallel, since each parameterization is uniquely defined by each line segment $\overline{ab}$. Thus, no more than $\sqrt{N}$ lines are contained in a single plane as well.

Finally, we consider doubly ruled surfaces. This case is more complex and beyond the scope of this paper, but it is handled in great detail by Guth and Katz[6].

We have shown that by applying Theorem 4.3, we can achieve our aforementioned required bound $|Q(S)| \leq c|S|^3 \log(|S|)$, which gives our desired bound that $|\mathbf{dist}(S)| \geq \frac{c|S|}{\log(|S|)}$ via a reduction to a 3D incidence problem. The next section shows how to bound incidences using the polynomial method and ultimately how to prove Theorem 4.3.

5. Bounding Incidences via Polynomial Cell Partitioning

Showing the bound on the number of incidences in $\mathbb{R}^3$ given in Theorem 4.3 involves partitioning $\mathbb{R}^3$, then bounding the number of incidences in each of the

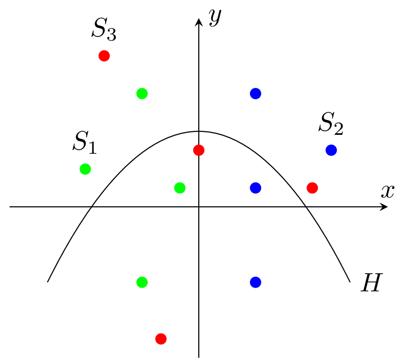

FIGURE 8. Hypersurface H bisects sets S_1 (green points), S_2 (blue points), and S_3 (red points) simultaneously.

"cells" created by the partition. In order for this to be possible, we require some way to partition $\mathbb{R}^3$ in a "balanced" way, meaning there are not too many lines in each cell and not too many cells. Guth and Katz showed that such a partition is possible, and that the partition can be created using a *hypersurface*.

Definition 5.1. A **hypersurface** is a set $H = \{x \in \mathbb{R}^n | h(x) = 0\}$, where $h(x)$ is a polynomial in n variables $x_1, ..., x_n$ of arbitrary degree d.

In other words, a hypersurface is an $n-1$ dimension subspace of $\mathbb{R}^n$. It is a generalization of a hyperplane, except a hyperplane is restricted to being defined by a degree 1 polynomial. We say the *degree* of a hypersurface H is the degree of the defining polynomial $h(x)$.

Guth and Katz showed that we can partition $\mathbb{R}^n$ using a low degree hypersurface that satisfies the other properties of a balanced partition mentioned above.

Theorem 5.2 (Polynomial Cell Partition)**.** *Let $S \subset \mathbb{R}^n$ be a finite set and let $t \geq 1$. Then there exists a hypersurface H that decomposes $\mathbb{R}^n$ into $\leq C_1 t$ cells, where hypersurface H has degree $d \leq C_2 t^{1/n}$ and each cell contains at most $|S|/t$ points from S for constants C_1, C_2.*

The idea for the proof of the Polynomial Cell Partition Theorem is to bisect the set S using a hypersurface H_1. We obtain two subsets of S, and we bisect both of them at the same time with one hypersurface H_2. Figure 8 illustrates an example of a single hypersurface, which is just a curve in $\mathbb{R}^2$, bisecting multiple sets simultaneously. We continue this process and iteratively bisect each of the i subsets of S using one hypersurface H_i. Each time, the number of cells doubles, and the number of elements in each cell is halved. After completing this process n times, we have hypersurfaces $H_1, H_2, ..., H_n$. We can take the union of all the hypersurfaces to obtain a single hypersurface H whose degree $d \leq \deg(H_1) + \deg(H_2) + ... + \deg(H_n)$ is small.

Then, the central idea to the proof of the Polynomial Cell Partition is that we can bisect multiple subsets of S at the same time using a single hypersurface. The following theorem allows us to do this and also bounds the degree of the bisecting hypersurface.

Theorem 5.3 (Discrete Polynomial Ham Sandwich)**.** *Let $S_1, ..., S_t \subset \mathbb{R}^n$ be t finite sets of points with $t < \binom{n+d}{d}$. Then there exists a degree d hypersurface H that bisects each of the sets $S_i, i \in [t]$.*

This theorem is an adaptation of The Polynomial Ham Sandwich Theorem, which claims that t *open* sets can be bisected with a single hypersurface. The Polynomial Ham Sandwich is used to prove the Discrete Polynomial Ham Sandwich, so we first begin with some discussion of the regular Polynomial Ham Sandwich Theorem.

5.1. **Polynomial Ham Sandwich.** The polynomial ham sandwich (abbreviated PHS) is an application of the polynomial method that allows us to bisect multiple open sets using a single hypersurface. The polynomial ham sandwich theorem was preceded by ham sandwich theorems, which involved multiple simultaneous bisections using hyperplanes rather than hypersurfaces. Historically, the "folklore" ham sandwich theorem makes the following claim:

Theorem 5.4 (Folklore Ham Sandwich)**.** *Every three bounded open sets in $\mathbb{R}^3$ can be simulataneously bisected using a single plane.*

We will first outline the practical implications of the Folklore Ham Sandwich Theorem before proceeding to generalizations and proofs. If we consider a real-world scenario where $\mathbb{R}^3$ is the observable universe and the three sets are viewed as two slices of bread and one slice of ham, then the Folklore Ham Sandwich Theorem states that if two friends wanted to split this sandwich equally (in terms of each ingredient), then one straight cut of a knife is always sufficient to split each in half. This is true regardless of orientation of the ingredients as well as how irregular the shapes get, since in this setting they can always be mathematically represented by bounded, open sets.

Now, we consider a generalization of the Folklore Ham Sandwich Theorem which leads us to the claim of the main Ham Sandwich theorem. We define a hyperplane $H \in \mathbb{R}^n$ as defined by some degree one polynomial h in n variables, consisting of the set $\{x \in \mathbb{R}^n | h(x) = 0\}$. We also define the two sides of the hyperplane in different contexts:

Definition 5.5. A hyperplane H defines two sides of $\mathbb{R}^n$, denoted by H^+ and H^-, where

$$H^+ = \{x \in \mathbb{R}^n | h(x) > 0\}$$
$$H^- = \{x \in \mathbb{R}^n | h(x) < 0\}$$

Definition 5.6. For any bounded open set of points $U \in \mathbb{R}^n$, define the following intersections that allow us to split the set according to the two sides of a hyperplane:

$$U \cap H^+ = \{x \in U | h(x) > 0\}$$
$$U \cap H^- = \{x \in U | h(x) < 0\}$$

We can now define the generalized ham sandwich theorem.

Theorem 5.7 (Ham Sandwich)**.** *Let $U_1, U_2, \ldots, U_n \in \mathbb{R}^n$ be bounded open sets. Then there exists a hyperplane H such that for each $i \in \{1, 2, \ldots, n\}$, the two sets $U_i \cap H^+$ and $U_i \cap H^-$ have equal volume.*

To prove this, we will use the following theorem:

Theorem 5.8 (Borsuk-Ulam)**.** *Let $S^n \in \mathbb{R}^{n+1}$ be the n-dimensional unit sphere, and let $f : S^n \to \mathbb{R}^n$ be a continuous map such that $\forall x \in S^n, f(-x) = -f(x)$. Then $\exists x$ such that $f(x) = 0$.*

The proof of the Borsuk-Ulam theorem can be found using topology, as shown in Matoušek [5].

Proof of Theorem 5.7. Since scaling the coefficients in our polynomial h doesn't change the defined hyperplane H, WLOG we can scale down our polynomial $h(x) = h_0 + h_1x_1 + h_2x_2 + \ldots + h_nx_n$ such that $\sum_{i=0}^{n} h_i^2 = 1$. This makes the vector of coefficients $v_h = (h_0, h_1, \ldots, h_n)$ a unit vector, i.e. $v_h \in S^n$.

Now, define a function $f : S^n \to \mathbb{R}^n$ as:

$$f(v_h) = (|U_i \cap H^+| - |U_i \cap H^-|)_{i \in \{1,2,\ldots,n\}}$$

Then, we have $f(-v_h) = -f(v_h)$, since flipping the sign of the input vector flips the sides of the hyperplane. In addition, f is continuous, since small changes in v_h result in slight changes in the angle of the hyperplane so that the difference in volumes has no discontinuities. By applying the Borsuk-Ulam theorem, $\exists v_h$ such that $f(v_h) = 0$, which makes $|U_i \cap H^+| = |U_i \cap H^-|$ for all $i \in \{1, 2, \ldots, n\}$. Thus, we have shown there is some polynomial that bisects all the given sets simultaneously. □

We can also extend the ham sandwich theorem to apply to more than just n sets by utilizing hypersurfaces rather than just hyperplanes. Recall that a hypersurface is defined similarly to a hyperplane, except that the defining polynomial h is no longer restricted to just a degree of one.

The result involves an extension of our results from Theorem 2.2:

Theorem 5.9 (Polynomial Ham Sandwich)**.** *Let $U_1, U_2, \ldots, U_t \in \mathbb{R}^n$ be bounded open sets, where $t < \binom{n+d}{d}$. Then there exists a degree d hypersurface H such that for each $i \in \{1, 2, \ldots, t\}$, the two sets $U_i \cap H^+$ and $U_i \cap H^-$ have equal volume.*

Proof of Theorem 5.9. This proof is analogous to the proof of Theorem 5.7. Note that a degree d polynomial of n variables has at most $\binom{n+d}{d}$ coefficients, by a stars-and-bars combinatorial argument. We can then again identify each hypersurface with its unit vector of coefficients and define a function f that maps to the volume difference of each set's two sides with respect to the hypersurface. The constraint $t < \binom{n+d}{d}$ allows us to append extra zeros into the mapping result of f so that our domain and range of $f : S^a \to \mathbb{R}^b$ have $a = b$, i.e. $a = b = \binom{n+d}{d}$. This allows us to apply the Borsuk-Ulam theorem in a similar fashion, proving the generalized PHS. □

5.2. **Discrete PHS.** In order to adapt the Polynomial Ham Sandwich Theorem to obtain the The Discrete Polynomial Ham Sandwich Theorem (Theorem 5.3), we first modify our definition of bisection to apply to discrete sets. We say that a hypersurface $H = \{x \in \mathbb{R}^n | h(x) = 0\}$ bisects a set S if both sets $S^- = \{x \in S | h(x) < 0\}$ and $S^+ = \{x \in S | h(x) > 0\}$ have size at most $|S|/2$. In other words, at most half the points in S are in $\{h > 0\}$ and at most half the points are in $\{h < 0\}$. Our loosened definition of bisection allows for h to vanish on an arbitrary number of points in S.

The proof of the Discrete Polynomial Ham Sandwich is an exercise in real analysis. The idea is to take ϵ-balls around each element in each of sets S_i. The union of the ϵ-balls in S_i is an open set U_i. Then with t open sets we can apply the regular Polynomial Ham Sandwich and obtain a bisecting hypersurface H. By making the ϵ-balls smaller and smaller, we get progressively closer to the original discrete sets $S_1, ..., S_t$.

Proof of Discrete Polynomial Ham Sandwich. Let us denote $U_{i,\epsilon}$ as the union of the epsilon neighborhoods around each point in S_i. If we take open sets $U_{1,\epsilon}, ..., U_{t,\epsilon}$, we can apply the PHS to obtain a bisecting hypersurface H_ϵ. We can define a sequence $\epsilon_1, ..., \epsilon_i$ that converges to 0, e.g.

$$\epsilon_i = \frac{1}{2^i}.$$

For each ϵ_i we apply PHS to obtain a hypersurface H_{ϵ_i} that bisects the open sets that are the ϵ_i neighborhoods of each $S_1, .., S_t$. Then we have a sequence of hypersurfaces H_{ϵ_i} that are defined by polynomial h_{ϵ_i}. We can scale the polynomial without changing the sign, i.e. if we multiply each coefficient by some constant, the same points in $h(x) > 0$ will still be in the scaled $h'(x) > 0$. We can write the coefficients of the scaled polynomial as a vector, so we obtain a sequence of unit vectors $\bar{h}_{\epsilon_i}$ that still define the same hypersurfaces H_{ϵ_i}. These unit vectors lie on the unit sphere, which is compact. Then by compactness, there exists a subsequence of the unit vectors $\bar{h}_{\epsilon_{k\,i}}$ that converges uniformly to a limit $\bar{h}$. Let h be the polynomial with coefficients from vector $\bar{h}$. We claim that h bisects $S_1, ..., S_t$, i.e. the hyperplane H defined by h bisects $S_1, ..., S_t$. We show this by contradiction.

Suppose H does not bisect $S_1, ..., S_t$. Then for some S_i we can write WLOG

$$|S_i^+| > \frac{|S_i|}{2},$$

i.e. more than half the points in S_i are in $\{h > 0\}$. By continuity, we can choose some very small δ such that $h > \delta$ on the δ-ball around each point in S_i^+. Since $\bar{h}_{\epsilon_{k\,i}}$ converged uniformly to $\bar{h}$, we can find k_i large enough that $h_{\epsilon_{k_i}} > 0$ on the δ ball around each point in S_i^+. By making k_i large, we can make it so that $\epsilon_{k_i} < \delta$. Then $h_{\epsilon_{k\,i}} > 0$ on more than half of the ϵ_{k_i} balls in $U_{i,\epsilon_{k\,i}}$, the ϵ_{k_i} neighborhood of S_i, a contradiction. ☐

Now that we have shown the Discrete Polynomial Ham Sandwich, we can use it to obtain our Polynomial Cell Partition Theorem.

5.3. **Polynomial Cell Partitioning Proof.** The Polynomial Cell Partition Theorem (Theorem 5.2) claims that we can choose some $t \geq 1$ and partition $\mathbb{R}^n$ such that there are at most $C_1 t$ cells, each cell has at most $|S|/t$ points from a finite set $S \subset \mathbb{R}^n$, and that the partition is created by a single hypersurface H of degree at most $C_2 t^{1/n}$ for constants C_1, C_2. The proof involves iteratively applying the Discrete PHS Theorem. We choose the number of times we apply the Discrete PHS so that our partition satisfies the desired conditions.

To obtain the bound on the degree of the hypersurface H, we use the bound on the degree of the hypersurface given by the Discrete PHS. If we fix n, where we are working with space $\mathbb{R}^n$, then as the the number of sets t grows, we can always find a degree $d \leq C|S|^{1/n}$ hypersurface that bisects each of the t sets. This can be shown using an identical argument as in the proof of Corollary 2.3.

Proof of Polynomial Cell Partition Theorem. Applying the discrete PHS the first time, we obtain a hypersurface H_1 of degree $d_1 \leq c_1 \cdot 1^{1/n}$ that bisects the single set $|S|$ into 2 sets of size at most $|S|/2$ each (plus some points on the boundary). Applying PHS again on these two sets, we obtain a degree $d_2 \leq c_2 \cdot 2^{1/n}$ hypersurface H_2 that bisects both sets. This gives 4 cells with boundary in the hypersurface $H' = H_1 \cup H_2$ which has degree at most $d_1 + d_2$ since it is defined by the product of polynomials defining each hypersurface. Each cell has at most $|S|/4$ points from S. If we continue partitioning $\ell = \log_2 t$ times, we obtain ℓ hypersurfaces $H_1, H_2, ..., H_\ell$ with each H_j having degree $d_j \leq c_j \cdot 2^{j/n}$ and such that their union $H = \cup_{j \in [\ell]} H_j$ gives a partition into roughly t cells containing at most $|S|/t$ points each. The degree of H is bounded by the sum

$$\sum_{j=1}^{\ell} c_j 2^{j/n} \leq C t^{1/n}.$$

□

The Polynomial Cell Partition Theorem is crucial for showing the Guth Katz bound on the number of incidences in $\mathbb{R}^3$. Their proof involves partitioning $\mathbb{R}^3$ and bounding the number of incidences in each cell, then the number of incidences on the boundaries of the cells, the hypersurface. Polynomial cell partitioning can also be used to bound the number of incidences between points and lines in $\mathbb{R}^2$, providing an alternate proof to a well-established theorem that we introduce in the next section.

5.4. **Szemerédi Trotter Bound on Incidences in $\mathbf{R}^2$.** The following bound on the number of incidences between a set of points P and a set of lines L in $\mathbb{R}^2$ was given by Szemerédi and Trotter.

Theorem 5.10 (Szemerédi Trotter). *Given a finite set of lines L and a finite set of points P in $\mathbb{R}^2$, we can bound the number of incidences $I(P, L)$*

$$|I(P, L)| \leq C(|P| \cdot |L|)^{2/3} + |L| + |P|,$$

for some constant C.

The original proof they gave was messy and complicated [1]. We will give a proof using polynomial cell partitioning that is much simpler, and the proof idea will serve as a warm-up to the proof of the Guth and Katz bound on the incidences in $\mathbb{R}^3$ as the ideas will be similar.

The idea of the proof is to use polynomial cell partitioning to partition $\mathbb{R}^2$ into cells. We bound the number of incidences in each cell and on the hypersurface that created the cells. In order to bound the number of incidences in cells and on the hypersurface, we need the following three lemmas. The first lemma extends the result of Lemma 3.5 to our framework of hypersurfaces. The second lemma bounds the number of lines that can be contained in the hypersurface. The third lemma is a simple bound on the number of incidences between L and P and will be applied in each cell.

Lemma 5.11. *Given any line $\ell \subset \mathbb{R}^2$ and hypersurface $H \subset R^2$ with degree d. Then either $\ell \subseteq H$ or $|\ell \cap H| \leq d$.*

Proof. Either the restriction of ℓ to the polynomial h defining hypersurface H is uniformly zero, implying $\ell \subseteq H$, or $|\ell \cap h| = |\ell \cap H| \leq d$ by Lemma 3.5. □

Lemma 5.12. *Given hypersurface $H \subset \mathbb{R}^2$ with degree d, there are at most d lines contained in H.*

This lemma is only true in two dimensions, but it can be generalized to higher dimensions by replacing the word "lines" with "hyperplanes." The proof of the lemma is a proof by contradiction.

Proof. Suppose there are $d+1$ distinct lines in H. Since $d+1$ is finite, we can choose some line ℓ not contained in H that is not parallel to any of the $d+1$ lines in H. Then ℓ intersects each of the $d+1$ lines, meaning ℓ intersects H at least $d+1$ times, contradicting Lemma 5.11. □

Lemma 5.13 (Cauchy Schwartz Bound)**.** *Given a finite set of lines L and points P in $\mathbb{R}^2$, we can bound the number of incidences $I(P,L)$*

$$|I(P,L)| \leq |P| \cdot |L|^{1/2} + |L|.$$

Also,

$$|I(P,L)| \leq |L| \cdot |P|^{1/2} + |P|.$$

Proof. We prove the first equation $|I(P,L)| \leq |P| \cdot |L|^{1/2} + |L|$. We express $|I(P,L)|$ using the indicator function $\mathbb{1}_{p \in \ell}$, which is equal to 1 if $p \in \ell$ for a point $p \in P$ and line $\ell \in L$, and 0 otherwise. We write

$$|I(P,L)| = \sum_{\ell \in L} \sum_{p \in P} \mathbb{1}_{p \in \ell}.$$

In order to bound this, we use the Cauchy Schwartz inequality, which says

$$\Big(\sum_{p=1}^{k} a_p \cdot b_p\Big)^2 \leq (\sum_{p=1}^{k} a_p^2) \cdot (\sum_{p=1}^{k} b_p^2).$$

Setting $a_p = 1, b_p = \sum_{p \in P} \mathbb{1}_{p \in \ell}$, if we square our expression for the number of incidences then by Cauchy Schwartz we have

$$\begin{aligned}
|I(P,L)|^2 &= \Big(\sum_{\ell \in L} \sum_{p \in P} \mathbb{1}_{p \in \ell}\Big)^2 \\
&\leq |L| \cdot \sum_{\ell \in L} \Big(\sum_{p \in P} \mathbb{1}_{p \in \ell}\Big)^2 \\
&= |L| \cdot \sum_{p_1, p_2 \in P} \sum_{\ell \in L} \mathbb{1}_{p_1 \in \ell} \cdot \mathbb{1}_{p_2 \in \ell} \\
&= |L| \cdot \Big(\sum_{p_1 = p_2 \in P} \sum_{\ell \in L} \mathbb{1}_{p_1 \in \ell} \cdot \mathbb{1}_{p_2 \in \ell} + \sum_{p_1 \neq p_2 \in P} \sum_{\ell \in L} \mathbb{1}_{p_1 \in \ell} \cdot \mathbb{1}_{p_2 \in \ell} \Big).
\end{aligned}$$

The first expression $\sum_{p_1 = p_2 \in P} \sum_{\ell \in L} \mathbb{1}_{p_1 \in \ell} \cdot \mathbb{1}_{p_2 \in \ell}$ is equivalent to the expression $|I(P,L)| = \Big(\sum_{\ell \in L} \sum_{p \in P} \mathbb{1}_{p \in \ell}\Big)^2$. For the second expression, where $p_1 \neq p_2$, we can view this as for every pair of points, we are checking if those two points are on a line. Otherwise, the value in the sum is 0. For every unique pair of points, they can only be on at most one line $\ell \in L$. Otherwise, if they were both incident to another line, the two lines would be the same line. Then the expression $\sum_{p_1 \neq p_2 \in P} \sum_{\ell \in L} \mathbb{1}_{p_1 \in \ell} \cdot \mathbb{1}_{p_2 \in \ell}$ is upper bounded by $\sum_{p_1 \neq p_2 \in P} 1$, which is upper bounded by $|P|^2$.

So our entire expression is upperbounded:

$$\begin{aligned} |I(P,L)|^2 &\le |L| \cdot (|I(P,L)| + |P|^2) \\ &\le |L|^2 + 2|L| \cdot |P|^2, \end{aligned}$$

where we use the following claim:

Lemma 5.14.

$$|I(P,L)| \le |P|^2 + |L|.$$

Before we prove Lemma 5.14, we note that our bound

$$|I(P,L)| \le |L|^2 + 2|L| \cdot |P|^2$$

implies the desired bound in Lemma 5.13. The second equation $|I(P,L)| \le |L| \cdot |P|^{1/2} + |P|$ uses an identical argument, where the roles of p and ℓ are swapped. □

Proof of Lemma 5.14. First we count lines that have at most one point in P on them. These lines contribute at most $|L|$ incidences (at most when all lines in L fall in this category). The rest of the lines have at least 2 points in P on each line. The total incidences on these lines is $\le |P|^2$, otherwise there would exist some point $p \in P$ that lies on $> |P|$ lines. Each of those lines would have another point in P on them by assumption, which implies there are more than $|P|$ points, a contradiction. □

Now that we have our three lemmas, we are ready to prove the Szemerédi Trotter Theorem (Theorem 5.10).

Proof of Theorem 5.10. First we assume that $|P|^{1/2} \ll |L| \ll |P|^2$. Otherwise the bound follows from Lemma 5.13. Next we apply the Polynomial Cell Partition Theorem with t to be chosen later. We obtain cells $C_1, C_2, ..., C_t$ with each cell having at most $\frac{|P|}{t}$ points. Note that the polynomial cell partition theorem says we will have $C_1 t$ cells, but we say we have t cells since the C_1 term will just be absorbed into the constant later. To see this, one can replace our usage of t with $C_1 t$ throughout this proof where it applies and find that it does not change our proof except some constants. The polynomial cell partition also gives us a hypersurface H of degree $d \le C_2 t^{1/2}$.

We bound the number of incidences in the cells then the number of incidences on the hypersurface. To do this, we introduce some notation. Let $P_0 = P \cap H$ be the points on the hypersurface, $L_0 = L \cap H$ be lines that intersect H, $P_i = P \cap C_i$ be points that are in cell C_i, and $L_i = L \cap C_i$ be lines that intersect cell C_i. Note that lines can be overcounted, i.e. the same line can appear in multiple L_i and also L_0. This overcounting does not matter since we only seek an upper bound. Then

$$I(P,L) \le I(P_0, L_0) + \sum_{i=1}^{t} I(P_i, L_i).$$

To bound the number of incidences in the cells, $\sum_{i=1}^{t} I(P_i, L_i)$, we apply the Cauchy Schwartz Bound (Lemma 5.13) to each cell. This gives

$$I(P_i, L_i) \le \frac{|P|}{t} \cdot |L_i|^{1/2} + |L_i|.$$

To bound the overall sum $\sum_{i=1}^{t} I(P_i, L_i)$, we need to bound sums $\sum_{i=1}^{t} |L_i|$ and $\sum_{i=1}^{t} |L_i|^{1/2}$. We observe that each $\ell \in L_i$ is not contained in H since if $\ell \subseteq H$, it would not intersect cell C_i. Then by Lemma 5.11, each line in L_i can intersect at most $d \leq C_2 t^{1/2}$ cells since the line must cross H to move from one cell to another. Then $\sum_{i=1}^{t} |L_i| \leq C_2 t^{1/2}|L|$. This, combined with the Cauchy Schwartz inequality, where we choose $a_i = |L_i|^{1/2}$ and $b_i = 1$ gives the bound $\sum_{i=1}^{t} |L_i|^{1/2} \leq C|L|^{1/2}t^{3/4}$.

Combining everything, we obtain

$$\begin{aligned}\sum_{i=1}^{t} I(P_i, L_i) &\leq C'\Big[\frac{|P|}{t} t^{3/4}|L|^{1/2} + t^{1/2}|L|\Big] \\ &= C'[t^{-1/4}|P||L|^{1/2} + t^{1/2}|L|.]\end{aligned}$$

Now that we have bounded $\sum_{i=1}^{t} I(P_i, L_i)$, we now want to bound $I(P_0, L_0)$. To do so, first split L_0 into two sets: let L_0' be the set of lines contained in H and let L_0'' be the set of lines not contained in H (but intersect H at some points). By Lemma 5.11, each line in L_0'' can intersect H in at most $d \leq C_2 t^{1/2}$ points. Then,

$$|I(P_0, L_0'')| \leq C_2 t^{1/2}|L_0''| \leq C_2 t^{1/2}|L|.$$

By Lemma 5.12, $|L_0'| \leq d \leq C_2 t^{1/2}$.

Using the Cauchy Schwartz Bound (Lemma 5.13), we get

$$|I(P_0, L_0')| \leq |L_0'||P_0|^{1/2} + |P_0| \leq t^{1/2}|P|^{1/2} + |P|$$

Since we initially assumed that $|P|^{1/2} << |L|$, we can write this bound as

$$|I(P_0, L_0')| \leq C''[t^{1/2}|L| + |P|].$$

Then,

$$\begin{aligned}|I(P_0, L_0)| &= |I(P_0, L_0')| + |I(P_0, L_0'')| \\ &\leq C_2 t^{1/2}|L| + C''[t^{1/2}|L| + |P|] \\ &\leq C'''[t^{1/2}|L| + |P|].\end{aligned}$$

Adding $|I(P_0, L_0)|$ and $\sum_{i=1}^{t} |I(P_i, L_i)|$, we get

$$\begin{aligned}|I(P, L)| &= |I(P_0, L_0)| + \sum_{i=1}^{t} |I(P_i, L_i)| \\ &\leq C'''[t^{1/2}|L| + |P|] + C'[t^{-1/4}|P||L|^{1/2} + t^{1/2}|L|] \\ &\leq C[t^{1/2}|L| + |P| + t^{-1/4}|P||L|^{1/2}].\end{aligned}$$

Finally, in order to obtain the desired bound on $|I(P, L)|$, we set

$$t = \frac{|P|^{4/3}}{|L|^{2/3}}.$$

Since we initially assumed $|L| << |P|^2$, we have $t \geq 1$. So, we are done. $\square$

5.5. **Bounding Incidences in $\mathbf{R}^3$.** The proof of the Szemerédi Trotter Theorem uses ideas that will be used in our proof of Theorem 4.3, which upper bounds the number of incidences in N lines in $\mathbb{R}^3$. The bound given by Guth and Katz is $cN^{1.5}\log N$ for some constant c. The proof is extremely simple using the following lemma, which lets us bound the number of incidences with at least k lines intersecting at that point.

Lemma 5.15. *Let L be a set of N lines in $\mathbb{R}^3$ with the same conditions as in Theorem 4.3. Let $I_{\geq k}(L)$ denote the set of points that have at least k lines in L passing through them. Then for every $k \geq 2$,*

$$|I_{\geq k}(L)| \leq c\frac{N^{1.5}}{k^2},$$

for some constant c.

We will prove this lemma after we first show how the Guth Katz bound follows from Lemma 5.15. Let $I_{=k}(L)$ be the set of points with exactly k lines from L passing through them.

Proof of Theorem 4.3. By hypothesis, no more than $\sqrt{N}$ lines intersect at a single point. Also, each point in set $I_{=k}(L)$ contributes $\binom{k}{2} \leq k^2$ incidences. Thus, we can write the number of incidences $I(L)$ as

$$|I(L)| \leq \sum_{k=2}^{\sqrt{N}} |I_{=k}(L)| \cdot k^2.$$

We can rewrite the sum, using the observation that $|I_{=k}(L)| = |I_{\geq k}(L)| - |I_{\geq k+1}(L)|$. Next we apply Lemma 5.15 to upper bound it. We obtain

$$\begin{aligned}
\sum_{k=2}^{\sqrt{N}} |I_{=k}(L)| \cdot k^2 &= \sum_{k=2}^{\sqrt{N}} (|I_{\geq k}(L)| - |I_{\geq k+1}(L)|) \cdot k^2 \\
&\leq c' \sum_{k=2}^{\sqrt{N}} |I_{\geq k}(L)| \cdot k \\
&\leq c' \sum_{k=2}^{\sqrt{N}} c'' \frac{N^{1.5}}{k^2} \cdot k \\
&= c'c''N^{1.5} \cdot \sum_{k=2}^{\sqrt{N}} \frac{1}{k} \\
&\leq cN^{1.5}\log N,
\end{aligned}$$

as desired, where we absorb constants c', c'', and the implied third constant into one constant c in the last step. □

Then the proof of the Guth Katz bound on the number of incidences between lines in $\mathbb{R}^3$ follows quite simply from Lemma 5.15. The more difficult task is proving our lemma. To do this, we prove separately the two cases: when $k \geq 3$ and when $k = 2$.

To make our exponents easier to work with, we will scale the condition in our lemma by a factor of $\sqrt{N}$. So, we will prove the following scaled version of Lemma 5.15.

Lemma 5.16. *Let L be a set of N^2 lines in $\mathbb{R}^3$ such that no more than N lines intersect at a single point and no plane or doubly ruled surface contains more than N lines. Let $I_{\geq k}(L)$ denote the set of points that have at least k lines in L passing through them. Then for every $k \geq 2$,*

$$|I_{\geq k}(L)| \leq C\frac{N^3}{k^2},$$

for some constant C.

5.6. **The $k \geq 3$ case.** Note that the only surface that contains at least 3 distinct lines through each of its points is a plane (or a union of planes) [7]. So, the $k \geq 3$ case does not require any conditions on doubly ruled surfaces, and thus the only assumption we need is that no plane contains more than N lines in L. We also know $k < N$ since Theorem 4.3 requires that no more than N lines pass through a single point.

We will use proof by contradiction to show the Guth Katz bound. The outline is as follows.

Using the polynomial cell partition theorem, we partition the points in $I_{\geq k}(L)$ into cells whose boundary is a low-degree hypersurface. Then, we split our problem into two cases: a cellular case and an algebraic case. The cellular case corresponds to when the majority of the points in $I_{\geq k}(L)$ are inside the cells. In this case, we can apply the S-T Theorem in the 3-dimensional case in each cell and sum the resulting bounds. This proof will be similar to the cell partition proof of the S-T Theorem. The algebraic case corresponds to when the majority of the points in $I_{\geq k}(L)$ are on the boundary separating the cells. In this case, the proof will be similar to the joints problem proof. We will show that the hypersurface contains many lines in L, and that each of these lines contains many points. Then, we can use the assumption that $k \geq 3$ to analyze two special cases of lines. In both of these cases, we will use an argument similar to the joints problem proof to reach a contradiction.

We will make two regularity assumptions in order to make our proofs simpler. We can remove these assumptions through proof by induction as shown by Guth and Katz [6].

1. Each point in $I_{\geq k}(L)$ has at most $2k$ lines in L passing through it.
2. Each line in L is incident to at least $c \cdot \frac{k|I_{\geq k}(L)|}{|L|}$ lines for some constant c.

Let $S = I_{\geq k}(L)$. Assume for the sake of contradiction that $S \geq C \cdot \frac{N^3}{k^2}$ for some constant C to be chosen later.

We use the 3-dimensional case of the Polynomial Cell Partition Theorem (Theorem 5.2) and choose t so that the hypersurface H has degree $d = 3(\frac{N}{k})$. We can also assume that $k \ll N$, since otherwise the bound $|S| \leq C \cdot \frac{N^3}{k^2}$ follows trivially. Since $k \ll N$, we then have that $d \geq 1$. By the Polynomial Cell Partition Theorem, each cell contains at most $\frac{|S|}{t} \leq c \cdot \frac{|S|}{d^3}$ points for some constant c. Recall that in order for a line to intersect a cell, it must cross H. By Lemma 5.11, a line not contained in H can intersect H in at most $deg(H) = d$ points. So, each line in L can intersect at most d cells.

Let $S_H = S \cap H$ and $S_C = S \setminus S_H$. By the pigeonhole principle, either $|S_C| \geq \frac{|S|}{2}$ (the cellular case) or $|S_H| \geq \frac{|S|}{2}$ (the algebraic case).

We begin with the cellular case, when the majority of the points in S are contained inside the cells.

The cellular case

Assume $|S_C| \geq \frac{|S|}{2}$. We will find a contradiction to this assumption, thus showing that the cellular case actually cannot happen and that we will actually always be in the algebraic case. To do this, we will upper bound the number of points in S_C and show that this upper bound is smaller than the assumed lower bound of $\frac{|S|}{2}$.

We will first prove a corollary that follows nicely from the S-T Theorem. This corollary will allow us to bound the number of points in S_C in each cell.

Corollary 5.17. *Let P and L be sets of points and lines in $\mathbb{R}^2$. Let P_k be the set of points in P that have at least k lines passing through them, where $k \geq 2$. Then for some constant C,*

$$|P_k| \leq C\Big(\frac{|L|^2}{k^3} + \frac{|L|}{k}\Big).$$

Proof. Each point in P_k has at least k lines passing through it, so each point contributes at least k incidences. So, we have that $|I(P_k, L)| \geq k|P_k|$. Then, by Theorem 5.10 (the S-T Theorem),

$$k|P_k| \leq |I(P_k, L)| \leq C((|P_k| \cdot |L|)^{2/3} + |L| + |P_k|).$$

for some constant C. We thus know that at least one of the following statements must be true:

1. $k|P_k| \leq C(|P_k| \cdot |L|)^{2/3}$
2. $k|P_k| \leq C|L|$
3. $k|P_k| \leq C|P_k|$

We consider each case.

Case 1:

$$k|P_k| \leq C(|P_k| \cdot |L|)^{2/3}$$
$$|P_k|^{1/3} \leq C \cdot \frac{|L|^{2/3}}{k}$$
$$|P_k| \leq C \cdot \frac{|L|^2}{k^3}$$

So, the bound holds in this case.

Case 2:

$$k|P_k| \leq C|L|$$
$$|P_k| \leq C \cdot \frac{L}{k}$$

So, the bound holds in this case.

Case 3:

$$k|P_k| \leq C|P_k|$$
$$k \leq C$$

Every pair of lines can only intersect in at most one point. Then, $|L|$ lines can intersect in at most $\binom{|L|}{2}$ unique points. This means that at most $\binom{|L|}{2}$

points can have at least k lines passing through them (since $k \geq 2$). Then, $|P_k| \leq C|L|^2$, so the bound holds in this case as well.

Since the bound holds in all three cases, we are done. □

We can now use this corollary to bound the total number of points in all cells.

Let L_i be the set of lines in L that pass through the ith cell. Using the assumption that $|S_C| \geq \frac{|S|}{2}$ and applying Corollary 5.17 to each cell i to get

$$\frac{|S|}{2} \leq |S_C| \leq \sum_i \left(\frac{|L_i|^2}{k^3} + \frac{|L_i|}{k} \right). \tag{5.18}$$

Since each line can pass through at most d cells, we must have

$$\sum_i |L_i| \leq d|L|.$$

Recall that each cell contains at most $P_i \leq c \cdot \frac{|S|}{d^3}$ points for some constant c. Then, from our first regularity assumption (that each point has at most $2k$ lines passing through it), we have

$$\max_i |L_i| \leq 2k|P_i| \leq c \cdot \frac{2k|S|}{d^3}.$$

Then,

$$\begin{aligned}
\sum_i |L_i|^2 &\leq \sum_i \Big((\max_j |L_j|)|L_i| \Big) \\
&= \max_i |L_i| \sum_i |L_i| \\
&\leq c \cdot \frac{2k|S|}{d^3} \cdot d|L| \\
&= c \cdot \frac{2k|S|N^2}{d^2}
\end{aligned}$$

Using our bound in Equation 5.18, we have

$$\begin{aligned}
\frac{|S|}{2} &\leq \sum_i \left(\frac{|L_i|^2}{k^3} + \frac{|L_i|}{k} \right) \\
&\leq c \cdot \frac{2k|S|N^2}{k^3 d^2} + \frac{d|L|}{k} \\
&= c \cdot \frac{2|S|N^2}{k^2 d^2} + \frac{dN^2}{k}
\end{aligned}$$

Recall we set $d = 3 \cdot \frac{N}{k}$. Then,

$$\frac{|S|}{2} \leq c \cdot \frac{2|S|N^2}{k^2 d^2} + 3 \cdot \frac{N^3}{k^2} = c \cdot \frac{2|S|}{9} + 3 \cdot \frac{N^3}{k^2} \geq (2c \cdot \frac{C}{9} + 3) \cdot \frac{N^3}{k^2}.$$

Choosing C to be large enough and c to be small enough,

$$\frac{|S|}{2} \geq \frac{C}{2} \cdot \frac{N^3}{k^2} > (2c \cdot \frac{C}{9} + 3) \cdot \frac{N^3}{k^2} \geq \frac{|S|}{2}.$$

This is clearly a contradiction. This implies that we can never have $|S_C| \geq \frac{|S|}{2}$, and so we are always in the algebraic case.

The algebraic case

In the algebraic case, we now have $|S_H| \geq \frac{|S|}{2}$.

The cellular case proof showed that we cannot actually have $|S_C| \geq \frac{|S|}{2}$. We can replace the denominator in this bound with any constant c to get a contradiction to $|S_C| \geq \frac{|S|}{c}$, since this would only change the constants involved. This means that we cannot have $|S_C| \geq \frac{|S|}{c}$ for any constant c, so we must have $|S_C| \leq \frac{|S|}{c}$ for any constant c. To achieve this, we just need to set $d = D \cdot \frac{N}{k}$ for some constant $D > 3$ (as we set $D = 3$ previously).

Since $|S_C| \leq \frac{|S|}{c}$ for any constant c, we can have $|S_H| \geq (1-\varepsilon)|S|$ for any constant ε. If we take ε to be small enough, we can remove points that are not on H, since this will not change the bound on $|S|$. Then, we can reduce the algebraic case, where we know $|S_H| \geq \frac{|S|}{2}$, to the case where all points in S are in S_H, i.e. $S \subset H$, so all points in S are on H.

So, $|S|$ is a set of points with $|S| \geq C \cdot \frac{N^3}{k^2}$. All points in S lie on the hypersurface H with degree $d \leq D \cdot \frac{N}{k}$. In addition, $k \geq 3$ lines in L pass through each point in S, where $|L| = N^2$.

Choosing C large enough and using the second regularity assumption (that each line in L is incident to many lines, i.e. has many points in S on it), we get that the number of points in $S \subset H$ on each line in L is at least $c \cdot \frac{k|S|}{|L|}$.

$$\geq c \cdot C \cdot \frac{N^3}{k|L|} = c \cdot C \cdot \frac{N}{k} > 10 \cdot D \cdot \frac{N}{k} \geq 10d$$

for some constant c. Then by Lemma 5.11, each line in L must be completely contained in H. So, $L \subset H$.

By the assumption $k \geq 3$, we know each point in S has 3 lines passing through it. Split the points in S into two categories:

1. Critical points: points through which 3 non-coplanar lines pass (a joint).
2. Flat points: non-critical points through which 3 planar lines pass through.

Similarly, split the lines in L into two corresponding categories:

1. Critical lines: lines that contain at least $5d$ critical points.
2. Flat lines: lines that contain at least $5d$ flat points.

Since each line contains at least $10d$ points, each line must be either critical or flat by the pigeonhole principle. Then, also by the pigeonhole principle, there must be either $\geq \frac{N^2}{2}$ critical lines or $\geq \frac{N^2}{2}$ flat lines.

This divides the problem space nicely into two cases:

Case 1: There are $\geq \frac{N^2}{2}$ critical lines. In this case, we will use a proof very similar to the joints problem proof to find a contradiction to to the total number of lines in L.

Case 2: There are $\geq \frac{N^2}{2}$ flat lines. In this case, we will use a similar proof to find a contradiction to the condition that no plane contains more than N lines in L.

If we can find some contradiction in both cases, then we are done, since these contradictions mean that S cannot contain too many points, i.e. $|S| \leq C \cdot \frac{N^3}{k^2}$. We now consider each case.

Case 1: There are $\geq \frac{N^2}{2}$ critical lines.

Recall the joints problem proof. There, we showed that if a polynomial $g(x, y, z)$ vanishes on a joint p, then ∇g also vanishes on p. We want to make a similar statement about $H = \{h(x, y, z) = 0\}$ and ∇h.

Express the gradient of h as

$$\nabla h = \begin{bmatrix} \frac{\partial h}{\partial x} & \frac{\partial h}{\partial y} & \frac{\partial h}{\partial z} \end{bmatrix} = \begin{bmatrix} f_x & f_y & f_z \end{bmatrix}.$$

By Lemma 3.4, f_x, f_y, f_z all have degree $\leq deg(h) = d$.

If H contains a joint $p \in H$, then ∇h must also vanish on p. Note that each critical point is a joint w.r.t. $L \subset H$. Then, ∇h vanishes on all critical points. Since each critical line contains at least $5d$ critical points, ∇h vanishes on all critical lines as well.

Because H contains all lines in L, h vanishes on all lines in L. Then, since ∇h vanishes on all critical lines, H shares all critical lines with each of the hypersurfaces defined by its partial derivatives:

$$F_x = \{f_x(x,y,z) = 0\}, F_y = \{f_y(x,y,z) = 0\}, F_z = \{f_z(x,y,z) = 0\}.$$

Since at least half the lines in L are critical, H shares the majority of lines in L with each hypersurface F_x, F_y, F_z.

Since we want to find a contradiction to the number of lines in L, we can upper bound the number of lines that H shares with F_x, F_y, F_z and show that this upper bound is smaller than N^2. To do this, we will first show that two hypersurfaces that do not have a common factor share an upper bounded number of lines. Then, we will show that h does not share a common factor with at least one of its partial derivatives to conclude that the number of lines H shares with its partial derivatives is upper bounded.

Lemma 5.19. *Let $G_1 = \{g_1(x,y,z) = 0\}$ and $G_2 = \{g_2(x,y,z) = 0\}$ be two hypersurfaces such that g_1 and g_2 do not have a common factor. Then, $G_1 \cap G_2$ contains at most $deg(g_1)deg(g_2)$ lines.*

To prove this claim, we will need Bezout's Theorem.

Theorem 5.20 (Bezout)**.** *Let $f(x,y), g(x,y) \in \mathbb{R}[x,y]$ be two polynomials without a common factor. Then, f and g have at most $deg(f)deg(g)$ common roots.*

Since the proof of Bezout's Theorem involves resultants, it is beyond the scope of this paper, and can be found in Guth's lecture [8].

Proof of Lemma 5.19. Consider a plane $A \in \mathbb{R}^3$ and parameterize it as $A = au + bv + c$, where $u, v, c \in \mathbb{R}$ and $a, b \in \mathbb{R}$. Consider the restrictions $\hat{g}_1(u,v)$ and $\hat{g}_2(u,v)$ of g_1 and g_2 to A. Recall by Lemma 3.4 that $deg(\hat{g}_1) \leq deg(g_1)$ and $deg(\hat{g}_2) \leq deg(g_2)$. Since g_1 and g_2 do not share a factor, $\hat{g}_1$ and $\hat{g}_2$ must not either. Then, by Theorem 5.20 (Bezout's Theorem), $\hat{g}_1$ and $\hat{g}_2$ have at most $deg(\hat{g}_1)deg(\hat{g}_2) \leq deg(g_1)deg(g_2)$ common roots.

Note that g_1 and g_2 vanish on all points on their respective hypersurfaces G_1 and G_2. Then, g_1 and g_2 both vanish on points in $G_1 \cap G_2$, which means that $\hat{g}_1$ and $\hat{g}_2$ also both vanish on points in $G_1 \cap G_2$. In addition, we know $\hat{g}_1$ and $\hat{g}_2$ both vanish on A. Thus, each intersection p of A with a line in $G_1 \cap G_2$ is a root of both $\hat{g}_1$ and $\hat{g}_2$. Since A intersects each line in $G_1 \cap G_2$ at a distinct point, each line in $G_1 \cap G_2$ contributes a common root to $\hat{g}_2$ and $\hat{g}_2$. So, the number of lines contained in $G_1 \cap G_2$ is bounded by $deg(g_1)deg(g_2)$. □

In order to use Lemma 5.19, we must satisfy the condition it requires, namely that h does not share a factor with its partial derivatives. If h has repeated factors, then we can factor it as

$$h = \Pi_i r_i(x,y,z)^{\alpha_i},$$

where there is at least one i such that $\alpha_i \geq 2$. Note that we are interested in the zeros of h and not their multiplicity. Additionally, we want the degree of h to be small, so it is okay to reduce the degree of h. Since removing repeated factors of h (i.e. removing factors so that $\alpha_i = 1$ for all i in the factoring of h) does not change what the zeros of h are and reduces $deg(h)$, we can assume WLOG that h has no repeated factors. If h does not have repeated factors, then it must not share any common factors with at least one of f_x, f_y, f_z.

WLOG, assume that h does not share any factors with f_x. By Lemma 5.19, H and F_x share at most $deg(h)deg(f_x) \leq deg(h)^2 = d^2$ lines. Recall that in the algebraic case, $d \leq D \cdot \frac{N}{k}$. So, by choosing D small enough and because $k \geq 3$, H and F_x share at most $d^2 = D^2 \cdot \frac{N^2}{k^2} < \frac{N^2}{2}$ lines. Recall that H and F_x, F_y, F_z share all critical lines. This means that there are $< z\frac{N^2}{2}$ lines in L, which contradicts $|L| = N^2$.

Case 2: There are $\geq \frac{N^2}{2}$ flat lines.

We will prove this case by finding a contradiction to the assumption that no plane contains more than N lines. The proof will follow a similar argument to the proof of Case 1.

Recall from the first case that h and its gradient ∇h vanish on all critical points and that we showed h does not share a common factor with at least one of its partial derivatives. We want to say something similar about flat points and some polynomials.

To do so, we first define what it means for a polynomial to be plane-free.

Definition 5.21. A polynomial g is **plane-free** if no irreducible factor of g has degree 1.

Then, we can say the following about flat points and some group of 9 polynomials.

Lemma 5.22. *There exist 9 polynomials $\pi_1, \ldots, \pi_9 \in \mathbb{R}[x,y,z]$, each with degree at most $3d$, such that:*

1. *Each flat point is a root of all 9 polynomials π_i.*
2. *If $h(x,y,z)$ is plane-free, then h does not share a factor with at least one of the 9 polynomials π_i.*

This lemma is proved by Guth and Katz [6].

The first part of Lemma 5.22 implies that each polynomial π_i vanishes on every flat point. Then, each corresponding hypersurface $\Pi_i = \{\pi_i(x,y,z) = 0\}$ contains every flat point. Since each $deg(\pi_i) \leq 3d$ and each flat line contains at least $5d$ flat points, all flat lines are contained in each hypersurface Π_i. This is analogous to the statement from the first case that all critical lines are contained in the hypersurfaces F_x, F_y, F_z.

We can factor h into its plane-free component and non plane-free component as follows:

$$h(x,y,z) = h_p(x,y,z)h_n(x,y,z),$$

where h_p contains all the degree 1 irreducible components of h and h_n is plane-free. Then, the hypersurface $H_p = \{h_p(x,y,z) = 0\}$ is the union of all the planes contained in H.

By Lemma 5.22, h_n does not share a common factor with some π_j. Then by Lemma 5.19, $H_n = \{h_n(x, y, z) = 0\}$ and Π_j share at most

$$deg(h_n)deg(\pi_j) \leq d(3d) = 3d^2$$

lines. Recall that $L \subset H$, so H contains all flat lines. Since we know Π_j also contains all flat lines, H and Π_j must share at least $\geq \frac{N^2}{2}$ lines. Note that $H = H_p \cup H_n$. Then, H and H_p share at least $\geq \frac{N^2}{2} - 3d^2$ lines. Recall $d \leq D \cdot \frac{N}{k}$. Choosing D small enough and because $k \geq 3$, we have that H and H_p share at least

$$\geq \frac{N^2}{2} - 3D \cdot \frac{N^2}{k^2} \geq cN^2$$

lines for some constant c.

Note that $deg(h_p) \leq deg(h) = d$, so h_p contains at most d degree 1 components. Then, H_p contains at most d planes. Then, by the pigeonhole principle, there must be at least $\geq \frac{cN^2}{d} > N$ lines in some plane. This contradicts our assumption that no plane contains more than N lines, so we are done.

5.7. **The $k = 2$ case.** We will prove the $k = 2$ case using contradiction. Let $S = I_{\geq 2}(L)$ be the set of points of intersection of at least 2 lines in L. We want to show that $|S| \leq C_1 \frac{N^3}{k^2}$, but for the $k = 2$ case, k is small enough that it can be absorbed into the constant. Then our contradiction assumption will be $|S| > C \cdot N^3$ for some large constant C to be chosen later.

The idea of the proof is to find a polynomial f of low degree that vanishes on each line in L. Then each line in L is contained in the hypersurface H defined by our vanishing polynomial. We factorize the polynomial as the product of four polynomials f_1, f_2, f_3, f_4, chosen carefully so that each polynomial has a specific property. They define corresponding hypersurfaces H_1, H_2, H_3, H_4 such that $H = H_1 \cup H_2 \cup H_3 \cup H_4$. We then consider that each line in L is in one of the H_1, H_2, H_3, or H_4. Then incidences are either between lines in the same hypersurface or between lines in different hypersurfaces. We consider both cases, showing the bound Lemma 5.16 holds in both, thus arriving at a contradiction.

Then in order to complete the proof, we need to show that it is possible to find a low degree polynomial that vanishes on set L. The following lemma shows we can find a degree $d \leq |N|^{1/2}$ polynomial that vanishes on N lines.

Lemma 5.23. *Let $\ell_1, \ell_2, ..., \ell_t$ be t lines in $\mathbb{R}^3$. Then there exists a non-zero polynomial of degree $d \leq 10t^{1/2}$ that vanishes on all lines ℓ_i. In other words, the restriction of the polynomial to all lines is identically zero.*

The proof of Lemma 5.23 is almost the same as the proof of Theorem 2.2.

Proof of Lemma 5.23. A polynomial $f(x, y, z)$ in $\mathbb{R}^3$ of degree $10t^{1/2}$ has $\binom{10t^{1/2}+3}{3} > 10t^{1.5}$ monomials and, therefore, coefficients. Each constraint of the form $f|_{\ell_i} \equiv 0$, which says that f vanishes on ℓ_i, gives at most $\deg(f) + 1 \leq 10t^{1/2} + 1$ homogeneous linear equations in the coefficients of f. Each of those linear equations comes from the vanishing of one of the coefficients of the univariate restriction to the line ℓ_i. Then there are more coefficients than equations, so there is a non-trivial solution. □

Unfortunately Lemma 5.23 is not strong enough for our purposes, as we will need a much smaller degree polynomial that vanishes on L. We can use this lemma, however, to prove the existence of a vanishing polynomial of smaller degree. This depends on the fact that L is not an arbitrary set. We assumed for our proof that $|S| > CN^3$, so L is a set of lines with many intersections. In fact, we can conclude that some fraction of lines must have at least $CN/10$ points of intersection on them. Absorbing the $1/10$ into our constant C and throwing away some fraction of lines, we can assume WLOG that each line in L has at least CN distinct points of intersection on it. With this additional hypothesis, we can prove there exists a much lower degree vanishing polynomial.

Lemma 5.24. *Suppose c is a large enough constant. Let L be a set of N^2 lines in $\mathbb{R}^3$ such that each line in L intersects at least cN other lines in distinct points. Then, there exists a non-zero polynomial of degree $d \leq \frac{N}{\sqrt{c}}$ that vanishes on all lines in L.*

Proof. We take a random subset L' of L by choosing each line to be included in L' independently with probability $1/C$. Also, with high probability, each line in our original set L will still intersect at least $N/2$ lines in L'. By Lemma 5.23, we can find a polynomial $f(x, y, z)$ of degree $10\sqrt{L'} \leq CN/\sqrt{C}$ that vanishes on L'. Polynomial f must also vanish on L since the restriction of f to each line in L has at least $N/2 > \deg(f)$ zeros when C is large. □

Now that we have our low degree polynomial that vanishes on N lines, we are ready to begin our proof of the $k = 2$ case of Lemma 5.16.

Proof of $k = 2$ case. Suppose, for contradiction, that $|S| > C \cdot N^3$ for some large constant C to be chosen later.

By Lemma 5.24, we can find a polynomial $f(x, y, z)$ of degree $d \leq N/\sqrt{c}$ such that f vanishes on all lines in L. We can write $f = \prod_i f_i(x, y, z)$, where f_i are irreducible polynomials. We can assume that f is square free, meaning no f_i is repeated (otherwise if there were repeated f_i's, we could remove the duplicates and the remaining product polynomial would still vanish on all lines in L). Then if we let F be the hypersurface defined by $f(x, y, z)$, F is the union of the hypersurfaces F_i defined by the different f_i's. Also if we denote d_i as the degree of f_i, we can write $d = \sum_i d_i$.

We split the f_i's into 4 groups: let f_{pl} be the product of all f_i's that are degree one, which correspond to F_i's which are hyperplanes; let f_{dr} be the product of of the doubly-ruled components; let f_{sr} be the product of the singly-ruled components; let f_{nr} be the product of the remaining f_i's that are non-ruled components. We also let $F_{pl}, F_{dr}, F_{sr}, F_{nr}$ be the hypersurfaces defined by the polynomials $f_{pl}, f_{dr}, f_{sr}, f_{nr}$ respectively.

Since f was defined to vanish on L, each line $\ell \in L$ is contained in F. Then each line $\ell \in L$ must be contained in one of the irreducible factors of F. Then we can split our set of incidences S into two cases: incidences between lines in different factors of F and incidences between lines in the same factor.

We first consider incidences between lines in diferent factors. We consider WLOG that a line ℓ in factor F_{pl} is, by construction of our factors, not contained in the other factor hypersurfaces F_{dr}, F_{sr}, F_{nr}. Then since ℓ is not contained in F_i for $i \in \{dr, sr, nr\}$, by Lemma 5.11, ℓ can intersect hypersurface F_i in at most d points. Then ℓ can only intersect lines in F_i in at most d points. The total incidences between lines in different factors is then bounded by $|L|d \leq N^2 \cdot N/\sqrt{c} \leq CN^3$, for some constant C. Then now we only need to consider incidences between lines in the same factor.

We consider each of the factors then. For the F_{pl} factor, which is the union of planes, we use the hypothesis that there are at most N lines in each plane. Then we have at least N^2 intersections in each plane. Since there are at most $d \leq N$ planes in F_{pl} due to the bound on the overall degree of F, the total number of incidences in F_{pl} is at most N^3. The same argument works for the incidences in F_{dr} since by hypothesis there are at most N lines in every doubly-ruled surface. The last two factors to examine are F_{sr} and F_{nr}.

To examine incidences in F_{sr} and F_{nr}, we will introduce two lemmas, one about singly ruled surfaces and the other about non-ruled surfaces, without proof. The proofs for the two lemmas can be found in [6]. The first lemma bounds intersections in a singly ruled surface.

Lemma 5.25. *Let $S \subset \mathbb{R}^3$ be a singly ruled surface. Then, every line in S, with the exception of at most 2 lines, can intersect at most $\deg(S)$ other lines in S.*

If there were 3 lines in S that had more than deg(3) intersections with other lines in S, then S would be doubly ruled. Using this lemma, we know that each singly ruled hypersurface $F_i \subset F_{sr}$, meaning each F_i included in the union of hypersurfaces that comprise F_{sr}, can have at most two "exceptional" lines that can have up to $|L| = N^2$ intersections each. Then each F_i can have at most $2N^2$ incidences contributed by their two "exceptional" lines. There are at most $d \leq N$ components F_i, so "exceptional" lines contribute at most N^3 incidences. Each non-"exceptional" line in F_i contributes at most $\deg(F_i) \leq d \leq N$ intersections. Then overall, lines in F_{sr} contribute at most CN^3 incidences.

Lastly, we consider incidences in F_{nr}, comprised of the union of the non-ruled surfaces in F. To do this, we need the following lemma that bounds the number of lines in a non-ruled surface.

Lemma 5.26. *A non-ruled surface $S \subset \mathbb{R}^3$ can contain at most $\deg(S)^2$ lines.*

This means that F_{nr} can contain at most $d^2 \leq N^2/C$ lines. We assume that there are more than AN^3 incidences contributed by lines in F_{nr}, where A is a large constant to be chosen later. We can choose constants A and C to be as large as we want. We can argue using induction to bound the number of incidences in F_{nr}. Specifically, we assume that Lemma 5.16 holds for $(N-1)^2$ lines, then use the assumption on the lines of F_{nr}. We omit the details here, but it requires a careful choice of constants A and C, and can be found in [6]. Then since there are $\leq CN^3$ incidences contributed by lines in each factor $F_{pl}, F_{dr}, F_{sr}, F_{nr}$, which together contain all lines in L, we are done with the $k = 2$ case.

□

Acknowledgments

Thanks to our instructor Dr. Nir Gadish for introducing us to the initial survey paper and guiding us throughout this semester! Also thank you to our communications advisor Dr. Malcah Effron, who also guided us through this paper writing process. Their help has been immense and we are extremely grateful.

References

[1] Dvir, Zeev. "Incidence Theorems and Their Applications." *Princeton University*, 2012. https://arxiv.org/pdf/1208.5073.pdf

[2] Tao, Terrence. "The Guth-Katz bound on the Erdős distance problem." https://terrytao.wordpress.com/2010/11/20/the-guth-katz-bound-on-the-erdos-distance-problem/

[3] Guth, Larry. http://math.mit.edu/ lguth/PolyMethod/lect4.pdf

[4] Guth, Larry. "Unexpected Applications of Polynomials in Combinatorics." http://math.mit.edu/ lguth/Exposition/erdossurvey.pdf

[5] Matoušek, Jiří. Using the Borsuk-Ulam Theorem: Lectures on Topological Methods in Combinatorics and Geometry. Springer Publishing Company, Incorporated, 2007.

[6] Guth, Larry. Katz, Nets. "On the Erdős Distinct Distance Problem in the Plane." 2011. https://arxiv.org/pdf/1011.4105.pdf

[7] Fuchs, Dmitry. Tabachnikov, Serge. "Mathematical Omnibus: Thirty Lectures on Classic Mathematics." 2007.

[8] Guth, Larry. http://math.mit.edu/ lguth/PolyMethod/lect13.pdf