Incidences between Points and Lines
PRIMES Circle 2024

Ladia Khaing and Ada Kilian
Mentor: Paige Bright

Abstract

Suppose we are given a set of points, P, and a set of lines, $\mathcal{L}$, in Euclidean space. How large can

$$I(P, \mathcal{L}) = \#\{(p, \ell) \in P \times \mathcal{L} : p \in \ell\},$$

the number of incidences between P and $\mathcal{L}$ be? In this expository paper, we present three bounds for $I(P, \mathcal{L})$, discuss the sharpness (or lack thereof) of each bound, and ultimately derive the sharp bound over Euclidean space which is known as the Szemerédi–Trotter theorem. The proof of this theorem utilizes graph theory and probability, and is a powerful result in discrete geometry.

Contents

1 Introduction

Suppose we are given a set of n distinct points, P, and a set of m distinct lines, $\mathcal{L}$, in Euclidean space $\mathbb{R}^N$ ($N \geq 2$). One might then ask: how many point-line pairs can there be such that the point lies on the line? In other words, how can we possibly get a good upper bound on

$$I(P, \mathcal{L}) := \#\{(p, \ell) \in P \times \mathcal{L} : p \in \ell\}?$$

Such point-line pairs are known as *incidences*.

One can immediately see that for any set of points and lines, P and $\mathcal{L}$ respectively, we have $I(P, \mathcal{L}) \leq nm$,

$$\{(p, \ell) \in P \times \mathcal{L} : p \in \ell\} \subset P \times \mathcal{L},$$

and $|P \times \mathcal{L}| = nm$. We refer to this as the *trivial bound*. However, one may notice that for $n, m \geq 2$, this upper bound cannot occur - there is no sharp case - for any set of points and lines.

If there were nm incidences for a set of points and lines, this would imply that every point lies on every line. This is not possible, as any two unique points have at most one line through them, and any two unique lines have at most one point in common. The question thus becomes: how can we utilize geometric information about points and lines to improve upon the trivial bound?

In this paper, we will develop two bounds of this form. The first we refer to as the Cauchy–Schwarz bound, as it uses the Cauchy–Schwarz inequality to incorporate the geometric information we described above. This is explored in Section 2, and the following bound is obtained:

Theorem 1.1. *Let P be a set of n distinct points in $\mathbb{R}^N$ and let $\mathcal{L}$ be a set of m distinct lines in $\mathbb{R}^N$. Then,*

$$I(P, \mathcal{L}) \leq \min\{\sqrt{n}(m^2 + nm)^{1/2}, \sqrt{m}(n^2 + nm)^{1/2}\}.$$

As it turns out, this bound is not sharp in Euclidean space and can thus be improved upon. However, in different geometries (such as over finite fields) this bound is as good as one can hope for for arbitrary points and lines (see Section 2.1).

Thus, to improve the result over $\mathbb{R}^N$, however, we need use more information about Euclidean space. In particular, we will use the fact that there is an *ordering* to points on lines in Euclidean space, i.e., it makes sense to say that one point lies next to another on a line in $\mathbb{R}^N$. To use this information, we transform our points and lines into a graph and utilize graph theory. This bound is known as the Szemerédi–Trotter theorem, and is the best possible bound over Euclidean space:

Theorem 1.2 (Szemerédi–Trotter)**.** *Let P be a set of n distinct points in $\mathbb{R}^N$ and let $\mathcal{L}$ be a set of m distinct lines in $\mathbb{R}^N$. Then,*

$$I(P, L) \lesssim n^{2/3}m^{2/3} + n + m.$$

Remark 1.3. Note that here, $A \lesssim B$ simply denotes $A \leq CB$ for some positive constant C. For our purposes, we will not concern ourselves with this constant C, but if one kept careful track of details throughout this paper one can show the result with $C = 4$ (though this is not optimal).

We will develop the necessary graph theory in Section 3 (in which we utilize Bóna's book [Bón02]), and then prove the Szemerédi–Trotter bound as outlined in Iosevich's book [Ios07]. We lastly, and briefly, discuss the sharp examples for the Szemerédi–Trotter theorem in Section 3.3.

Acknowledgements. We would like to thank the MIT Math Department and PRIMES Circle for coordinating and managing this great learning opportunity! In particular, we would like to thank Mary Stelow and Marisa Gaetz for their comments and suggestions on our paper and presentation. We would like to thank Paige Bright for being a patient, dedicated mentor and for assisting us in our first-ever academic math conference and the formulation of rigorous mathematical proofs. Ladia would also like to thank her math teacher, Maureen Long, and her family for their extra help sessions and encouragement in pursuing higher maths with curiosity. Ada would like to thank her math teacher, Joel Patterson, for sharing this opportunity with her and fueling her curiosity in math beyond the curriculum in discussions throughout the school year.

2 The Cauchy–Schwarz Bound

To improve upon the trivial bound for incidences between points and lines, we will try to find a way to utilize geometric information about points and lines. In particular, that two lines can intersect at at most one point, or similarly that given two points, there is at most one line that passes through those two points. To utilize this geometric information, we will enumerate our set of points and lines.

We let $P = \{p_i\}_{i=1}^n$ and $\mathcal{L} = \{\ell_j\}_{j=1}^m$. Then, we define the following numbers (sometimes referred to as an incidence matrix):

$$a_{ij} = \begin{cases} 1 & \text{if } p_i \in \ell_j \\ 0 & \text{if } p_i \notin \ell_j. \end{cases}$$

This incidence matrix will prove helpful for encapsulating geometric information about points and lines algebraically.

Firstly, notice that by definition of the a_{ij},

$$I(P, \mathcal{L}) = \sum_{i=1}^{n} \sum_{j=1}^{m} a_{ij}.$$

In fact, this gives us another short proof of our trivial bound.

Proposition 2.1. *Let P be a set of n points in $\mathbb{R}^N$ and $\mathcal{L}$ be a set of m lines in $\mathbb{R}^N$. Then, $I(P,\mathcal{L}) \leq nm$.*

Proof. For every i and j, $a_{ij} \leq 1$ by definition. Hence,

$$I(P,\mathcal{L}) = \sum_{i=1}^{n}\sum_{j=1}^{m} a_{ij} \leq \sum_{i=1}^{n}\sum_{j=1}^{m} 1 = nm.$$

□

However, as we discussed before, it is impossible (for large n and m) for every a_{ij} to be 1, as this would mean that every point lies on every line. We can obtain a better bound by using the fact that two distinct lines in $\mathbb{R}^N$ intersect at at most one point. In other words, algebraically, we have that for $j \neq j'$

$$a_{ij} \cdot a_{ij'} = 1 \quad \text{for at most one value of } i. \tag{1}$$

In the rest of this section, we will develop machinery to utilize this geometric information to find a better bound than nm. In particular, we will use the Cauchy–Schwarz inequality to get a better bound.

Lemma 2.2 (Cauchy–Schwarz)**.** *Let a_i and b_i be real numbers for $1 \leq i \leq n$. Then,*

$$\sum_{i=1}^{n} a_i b_i \leq \left(\sum_{i=1}^{n} a_i^2\right)^{1/2} \left(\sum_{i=1}^{n} b_i^2\right)^{1/2}.$$

Proof. If a and b two real numbers, then $(a-b)^2 \geq 0$. One can rearrange this inequality to equivalently obtain

$$ab \leq \frac{a^2+b^2}{2}.$$

For ease of notation, let

$$A = \left(\sum_{i=1}^{n} a_i^2\right)^{\frac{1}{2}}, \quad B = \left(\sum_{i=1}^{n} b_i^2\right)^{\frac{1}{2}}.$$

Note that these are simply fixed constants given fixed sequences a_i and b_i.

Then, be factoring out constants, we have

$$\begin{aligned}\sum_{i=1}^{n} a_i b_i &= AB\sum_{i=1}^{n} \frac{a_i}{A}\cdot\frac{b_i}{B} \\ &\leq AB\cdot\frac{\sum a_i^2}{2A^2} + AB\cdot\frac{\sum b_i^2}{2B^2}.\end{aligned}$$

Here we used that $ab \leq \frac{a^2}{2} + \frac{b^2}{2}$ for every term $a_i b_i$. Recalling the definition of A and B, we see that

$$= \frac{AB}{2} + \frac{AB}{2} = AB.$$

This gives the desired inequality. □

We are now ready to use the algebraic information in (1). We will first start with $n = m$ to highlight the proof idea.

Theorem 2.3. *Let P be a set of n distinct points in $\mathbb{R}^N$ and let $\mathcal{L}$ be a set of n distinct lines in $\mathbb{R}^N$. Then,*

$$I(P,\mathcal{L}) \lesssim n^{3/2}.$$

Proof. Our setup is the same as the trivial bound, but now we utilize the Cauchy–Schwarz inequality to incorporate our geometric observation (namely that for $j \neq j'$, $a_{ij} \cdot a_{ij'} = 1$ for at most one value of i and must be zero for the rest). In particular, we have

$$\begin{aligned} I(P, \mathcal{L}) = \sum_{i=1}^{n} \sum_{j=1}^{n} a_{ij} &= \sum_{i=1}^{n} \left(\sum_{j=1}^{n} a_{ij} \right) \cdot 1 \\ &\leq \left(\sum_{i=1}^{n} 1 \right)^{1/2} \cdot \left(\sum_{i=1}^{n} \left(\sum_{j=1}^{n} a_{ij} \right)^2 \right)^{1/2}. \end{aligned}$$

Here, we used the Cauchy-Schwarz inequality with respect to the sum on i. Evaluating the first sum, we obtain that

$$= \sqrt{n} \cdot \left(\sum_{i=1}^{n} \left(\sum_{j=1}^{n} a_{ij} \right)^2 \right)^{1/2}.$$

Hence, at this point, we will simply show that

$$\sum_{i=1}^{n} \left(\sum_{j=1}^{n} a_{ij} \right)^2 \lesssim n^2.$$

To do so, we expand the square of the sum as a double sum, as seen below:

$$\sum_{i=1}^{n} \left(\sum_{j=1}^{n} a_{ij} \right)^2 = \sum_{i=1}^{n} \sum_{j=1}^{n} \sum_{j'=1}^{n} a_{ij} a_{ij'}.$$

Recall that given $j \neq j'$, $a_{ij} \cdot a_{ij'} = 1$ for at most one i (as two different lines can intersect at at most one point). Therefore, we break into cases: if $j \neq j'$ and otherwise:

$$= \sum_{i} \sum_{j \neq j'} a_{ij} a_{ij'} + \sum_{i=1}^{n} \sum_{j=1}^{n} a_{ij}^2$$

The first sum we bound by swapping the sum on i and the sum on $j \neq j'$. In particular, since $j \neq j'$, $\sum_i a_{ij} a_{ij'} \leq 1$. The second sum we bound the same as the trivial bound (using that $a_{ij}^2 = a_{ij}$ by definition), obtaining

$$\begin{aligned} &\leq \left(\sum_{j \neq j'} \sum_{i} a_{ij} a_{ij'} \right) + n^2 \\ &\leq \left(\sum_{j \neq j'} 1 \right) + n^2 \\ &\leq \#\{j, j' \in [1, n] : j \neq j'\} + n^2 \\ &\leq n^2 - n + n^2 \lesssim n^2. \end{aligned}$$

This completes the proof. $\square$

Using a nearly identical proof method, one can more generally obtain that (when there are n points and m lines, as opposed to $n = m$),

$$I(P, \mathcal{L}) \leq \sqrt{n}(m^2 + nm)^{1/2}.$$

Additionally, there is a symmetry between points and lines we can similarly utilize. In particular, given two different points, there is at most one line through those two points. Hence, algebraically, we have that for $i \neq i'$,

$$a_{ij} \cdot a_{i'j} = 1 \quad \text{for at most one value of } j.$$

Therefore, via a mild modification of our argument (squaring the sum over i as opposed to squaring over j), we obtain a similar bound:

$$I(P, \mathcal{L}) \leq \sqrt{m}(n^2 + nm)^{1/2}.$$

Which bound is better depends on the sizes of n and m. We record the above facts in the following theorem, though we won't go into more detail here.

Theorem 2.4. *Let P be a set of n distinct points in $\mathbb{R}^N$ and let $\mathcal{L}$ be a set of m distinct lines in $\mathbb{R}^N$. Then,*

$$I(P, \mathcal{L}) \leq \min\{\sqrt{n}(m^2 + nm)^{1/2}, \sqrt{m}(n^2 + nm)^{1/2}\}.$$

Notice that we have gotten an improvement over the trivial bound, as if $n = m$, the trivial bound gives that we have n^2 incidences whereas now we have $\sim n^{3/2}$ incidences. However, the Cauchy–Schwarz bound is not nearly as good as the advertised Szemerédi–Trotter bound. Still, as it turns out, the Cauchy–Schwarz bound (which only used basic assumptions about lines and points in Euclidean space) *is* sharp in different geometries. We prove this in the next section.

2.1 Sharpness over Finite Fields

In this section, we prove that the Cauchy–Schwarz bound cannot be improved upon in different, non-Euclidean geometries, in particular finite fields. To do so, we begin by defining modular arithmetic.

Definition 2.5. Given $a, q \in \mathbb{Z}$, by Euclid's algorithm, we know that we may (uniquely) write

$$a = kq + r$$

where $k, r \in \mathbb{Z}$, and $0 \leq r < q$. This is simply a form of visualizing the division of any number a by q with remainder r. Given $a = k \cdot q + r$, we denote

$$a \equiv r \pmod{q}.$$

Another way to view this is: $a \equiv r \pmod{q}$ if and only if $a - r$ is divisible by q. Here and throughout, we will let q denote a prime number.

As an example,

$$5 \pmod{5} = 0, \quad \text{as} \quad 5 = 5 + 0.$$

Similarly,

$$6 \pmod{5} = 1, \quad \text{as} \quad 6 = 5 + 1.$$

Definition 2.6. Let q be a prime number. Then, we define the finite field $\mathbb{F}_q$ as

$$\mathbb{F}_q = \{0, 1, 2, 3, \ldots, (q-1)\},$$

and define addition and multiplication between elements via modular arithmetic. I.e., given $a, b \in \mathbb{F}_q$, we have

$$\begin{aligned} a + b &:= (a + b) \pmod{q} \\ a \cdot b &:= (ab) \pmod{q}. \end{aligned}$$

With this definition in hand, we can now introduce notation/definitions for subtraction and "division."

Remark 2.7. We use the following notation over finite fields.

- Notice that, given any $a \in \mathbb{F}_q$, there exists a unique element $b \in \mathbb{F}_q$ such that $a + b = 0$, namely $b = q - a$. As such, we denote $q - a = -a$. Hence, we can define subtraction $a - b$ as $a + (-b)$, where the negative of b is defined uniquely above.

- Similarly, given $a \neq 0$, one can show via Fermat's little theorem that there exists a unique element b such that $a \cdot b = 1$. This is, as by Fermat's little theorem,

$$a^q = a \pmod q \iff a \cdot a^{q-2} = a^{q-1} = 1 \pmod q.$$

  As such, we denote $a^{-1} = a^{q-2}$. Note that this element is unique. Here, we crucially use that q is prime.

- Lastly, given $k \in \mathbb{N}$, we define $ka = \underbrace{a + a + \cdots + a}_{k \text{ times}} \pmod q$.

Before we define lines over finite fields, we need introduce the following proposition.

Proposition 2.8. *Let q be prime and let $0 \neq a \in \mathbb{F}_q$. Then,*

$$\#\{0, a, 2a, 3a, \ldots, (q-1)a\} = q.$$

Proof. Let us suppose that

$$\#\{0, a, 2a, 3a, \ldots, (q-1)a\} < q.$$

If this is the case, then there must exist $m \neq n$ such that $m \cdot a = n \cdot a$. This means that

$$(m - n) \cdot a = 0$$

Because $m \neq n$, we know that $m - n \neq 0$. Thus, there exists a multiplicative inverse $(m - n)^{-1}$. So, we have

$$(m - n) \cdot a = 0 \iff (m - n)^{-1} \cdot (m - n) \cdot a = a = 0.$$

This contradicts our hypothesis that $a \neq 0$ and concludes the proof. □

Now, we can define the plane over $\mathbb{F}_q$, $\mathbb{F}_q^2$, and define lines.

Definition 2.9. We define $\mathbb{F}_q^2$ as pairs of elements $(a, b) \in \mathbb{F}_q \times \mathbb{F}_q$ with addition and scalar multiplication defined element-wise. I.e.,

$$(a, b) + (c, d) = (a + c \pmod q, b + d \pmod q) \quad \text{and} \quad k(a, b) = (ka \pmod q, kb \pmod q).$$

We let 0 denote $(0, 0) \in \mathbb{F}_q^2$.

Then, we define affine lines in $\mathbb{F}_q^2$ as sets of the form

$$\ell = \{tu + v : t \in \mathbb{F}_q\},$$

where $u, v \in \mathbb{F}_q^2$ and $u \neq 0$. Lastly, we say that ℓ is a line through the origin if $0 \in \ell$

However, with this form in $\mathbb{F}_q^2$, many identical lines are formed for any $0 \neq u \in \mathbb{F}_q^2$ and $v \in \mathbb{F}_q^2$. It turns out that this overcounts the number of lines. So to prove that the Cauchy–Schwarz bound is sharp over finite fields, we need to enumerate the number of affine lines over $\mathbb{F}_q^2$.

Lemma 2.10. *There are $q + 1$ unique lines through the origin in $\mathbb{F}_q^2$.*

Proof. In this case, we are counting the lines that contain the element 0. This means, up to reordering of the set, lines through the origin are of the form

$$\ell = \{tu : t \in \mathbb{F}_q\}.$$

To figure out how many lines through the origin there are, we do a standard overcounting argument. Firstly, notice that $\mathbb{F}_q^2$ contains q^2-1 unique non-zero points. Hence, there are q^2-1 options for u to define a line ℓ. Fixing this line, we figure out how many options of u there are that result in the same set ℓ.

We claim that there are $q-1$ unique ways to represent ℓ via a $u \in \mathbb{F}_q^2 \setminus 0$. To see this, first notice that if $u' \notin \ell$, then $\{tu' : t \in \mathbb{F}_q\} \neq \ell$, as u' is in the latter but not the former. We claim that there are $(q-1)$ non-zero elements of $\mathbb{F}_q^2$ contained in ℓ, namely $u, 2u, 3u, \ldots, (q-1)u$. To see this, we simply need to reparameterize our lines. Notice that

$$\{t(2u) : t \in \mathbb{F}_q\} = \{(2t)u : t \in \mathbb{F}_q\} = \{tu : t \in \mathbb{F}_q\}$$

as multiplication is invertible over finite fields. The same follows for $3u, \ldots, (q-1)u$.

$$\begin{aligned}\ell &= \{tu : t \in \mathbb{F}_q\} \\ &= \{t(2u) : t \in \mathbb{F}_q\} \\ &= \vdots \\ &= \{t((q-1)u) : t \in \mathbb{F}_q\}.\end{aligned}$$

Hence, every line through the origin contains $q-1$ nonzero points, each of which provide another way to reparameterize ℓ. Since there are q^2-1 possible ways to write down a line with a nonzero element u, it follows that in total there are $\frac{q^2-1}{q-1} = q+1$ unique lines through the origin in $\mathbb{F}_q^2$.

We lastly just need to check that $u, 2u, \ldots, (q-1)u$ are all unique. However, this is followed by Proposition 2.8. Given $(u_1, u_2) = u \neq 0$, it must be that either u_1 or $u_2 \neq 0$. Suppose without loss of generality that $u_1 \neq 0$. Thus, by Proposition 2.8, $u_1, 2u_1, \ldots, (q-1)u_1$ are all distinct elements of $\mathbb{F}_q$, and thus $u, 2u, \ldots, (q-1)u$ are all distinct elements of $\mathbb{F}_q^2$. To see why we cannot contain any more points, notice that $qu = 0$ and $(q+1)u = u$ (and thus the sequence repeats itself). □

We use this fact to count the total number of affine lines (perhaps not through the origin) in $\mathbb{F}_q^2$.

Theorem 2.11. *The number of unique affine lines in $\mathbb{F}_q^2$ is $q(q+1)$.*

Proof. Technically, one must check that given two distinct lines in $\mathbb{F}_q^2$, they intersect at at most one point. We leave this proof to the reader but utilize this fact here.

We can break the set of affine lines into ones that are vertical (i.e. will not intersect the y-axis uniquely) and the ones that are nonvertical (i.e. the ones that will intersect the y-axis uniquely). Then, we perform casework.

Notice that there are q nonvertical lines through the origin (as there is only one vertical through the origin with $u = (0,1)$). I.e., there are q choices of u when $v = 0$ and ℓ is nonvertical. Then, we can count the set of all nonvertical affine lines (not through the origin) based on where they uniquely intersect the y-axis. Given there are q places for a nonvertical line to intersect the y-axis uniquely, it follows that there are q^2 unique and nonvertical lines in the set of all affine lines in $\mathbb{F}_q^2$. Lastly, there is only one vertical line through the origin, and there are q places where a vertical line can (uniquely) cross the x-axis. Thus, there are q unique vertical lines over $\mathbb{F}_q^2$.

Hence, in total, there are $q^2 + q = q(q+1)$ many unique lines in $\mathbb{F}_q^2$. □

This is enough to prove the sharpness of the Cauchy–Schwarz bound over finite fields (up to some constant).

Corollary 2.12. *Given a set of n distinct points, P, and m distinct lines, $\mathcal{L}$, in $\mathbb{F}_q^2$ (with q prime),*

$$I(P, \mathcal{L}) \leq \min\{\sqrt{n}(m^2 + nm)^{1/2}, \sqrt{m}(n^2 + nm)^{1/2}\}.$$

Furthermore, there exists a set of n points, P, in $\mathbb{F}_q^2$ and a set of n lines, $\mathcal{L}$, in $\mathbb{F}_q^2$ such that

$$I(P, \mathcal{L}) = n^{3/2}.$$

Proof. The first statement follows as the only information we need is that two distinct lines can intersect at at most one point, and similarly through two distinct points there is at most one line through them. The proof then follows precisely the same as over Euclidean space. We omit the proof here.

For the latter statement, let $P = \mathbb{F}_q^2$, the set of all $n = q^2$ points, and let $\mathcal{L}$ be any q^2 of the $q(q+1)$ unique lines over $\mathbb{F}_q^2$. Then, because every lines passes through q points in P (since P is the entire field), it follows that

$$I(P, \mathcal{L}) = \sum_{\ell \in \mathcal{L}} \#(P \cap \mathcal{L}) = \sum_{\ell \in \mathcal{L}} q = q^3.$$

By definition, $n^{3/2} = q^3$, proving the claim. $\square$

Remark 2.13. Note that here, $|P|$ and $|\mathcal{L}|$ are quite large. An interesting question, for instance, is if the incidence bound is possible to be achieved if $|P| = |\mathcal{L}| = q$. This is known as the finite field Furstenberg problem, though we won't explore this problem any further in this paper.

3 The Szemerédi–Trotter Bound

While the Cauchy–Schwarz bound is sharp over finite fields, it is not sharp over Euclidean space. The best bound over $\mathbb{R}^N$ is found to be Szemerédi–Trotter, which states:

Theorem 3.1. *Let P be a set of n distinct points in $\mathbb{R}^N$ and let $\mathcal{L}$ be a set of m distinct lines in $\mathbb{R}^N$. Then,*

$$I(P, L) \lesssim n^{2/3} m^{2/3} + n + m.$$

To improve over the Cauchy–Schwarz bound, we use the fact that the real numbers are ordered (i.e., for our purposes, it makes sense to say when two points are "next to" each other on a line in $\mathbb{R}^N$). We use this information when we turn our set of points and lines into a graph, as described in the following section.

3.1 Graph Theory

Given our set of points and lines, P and $\mathcal{L}$, we will transform these into a *graph*.

Definition 3.2 (Graph)**.** We define a graph $G = (V, E)$ to be a set of vertices V and a set of edges E. Every edge in E connects two vertices in V.

Remark 3.3. We denote $|V| = v$, $|E| = e$. If we need to specify the number of vertices, edges, or faces in a specific graph G, we will denote this as $v(G)$ and $e(G)$ respectively.

We describe the process in which we turn our points and lines into a graph, $G = (V, E)$, here. We let our set of points P become our vertices (i.e. $|V| = n$), and we will connect two vertices p_i and p_j $(i \neq j)$ by an edge if and only if p_i and p_j lie on a common line ℓ and they are adjacent to one another on the line. We call the resulting graph, G, our *incidence graph.* See Fig. 1

.

As it turns out, we can also find the number of edges in our graph G.

Remark 3.4. Here, we suppose that each line in our set of lines $\mathcal{L}$ contains at least one point of P. We can assume this is the case without loss of generality, as we can split $\mathcal{L}$ into lines with incidences, $\mathcal{L}_0$, and those without, $\mathcal{L}_1$. Notice that $|\mathcal{L}_0| \leq |\mathcal{L}|$, and

$$I(P, \mathcal{L}) = I(P, \mathcal{L}_0) + I(P, \mathcal{L}_1) = I(P, \mathcal{L}_0).$$

Hence, proving Theorem 3.1 for $\mathcal{L}_0$ implies the result for $\mathcal{L}$. In fact, this assumption is necessary for the following lemma.

Lemma 3.5. *Let G be the incidence graph described above, and suppose that each line in $\mathcal{L}$ contains at least one point in P. Then, we have that $e = I(P, \mathcal{L}) - m$.*

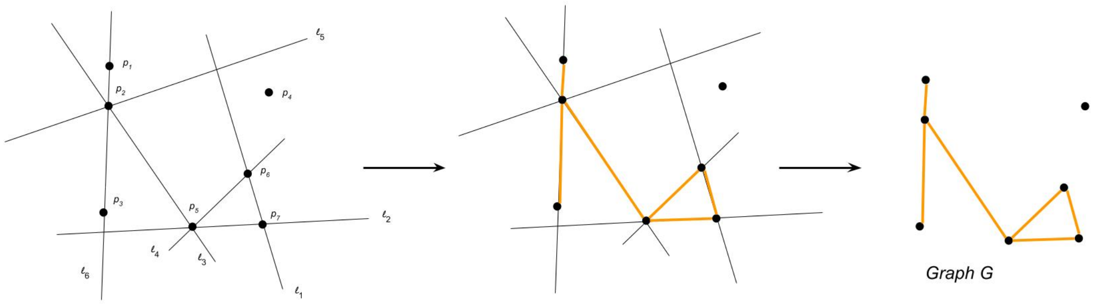

Figure 1: Turning points and lines P and $\mathcal{L}$ into an incidence graph G.

Proof. We number the lines in $\mathcal{L}$ as $\ell_1, \ell_2, \ldots, \ell_m$. We denote p_i to be the number of points from P on the line ℓ_i (by assumption $p_i \neq 0$ for each i). Notice that ℓ_i contributes p_i-many incidences to the set

$$\{(p, \ell) \in P \times \mathcal{L} : p \in \ell\}.$$

I.e.,

$$I(P, \ell_i) := \{(p, \ell_i) \in P \times \ell_i : p \in \ell_i\} = p_i.$$

Additionally, by definition,

$$I(P, \mathcal{L}) = \sum_{i=1}^{m} I(P, \ell_i).$$

Now, notice that ℓ_i contributes precisely $p_i - 1$ many edges to our incidence graph G. Note that here is where we use that each $p_i \neq 0$. Hence the total number of edges in G is

$$e = \sum_{i=1}^{m} (p_i - 1) = \left(\sum_{i=1}^{m} p_i \right) - m = \left(\sum_{i=1}^{m} I(P, \ell_i) \right) - m = I(P, \mathcal{L}) - m.$$

This completes the proof. □

So we have transformed our points and lines into a graph $G = (V, E)$ with $|V| = n$ and $|E| = I(P, \mathcal{L}) - m$. We now need to develop results from graph theory to relate the size of V and E to one another.

3.1.1 Euler's Theorem

To relate the number of vertices and edges to one another, we prove Euler's theorem for *planar* graphs.

Definition 3.6 (Planar, Connected)**.** A graph G is *planar* if it can be drawn on a plane surface so that no two of its edges intersect. Furthermore, G is a *connected* graph if, for each pair of vertices, there exists a path (i.e. a sequence of edges) which connects them. See Fig. 2.

Lastly, to understand Euler's theorem, we need to define the faces for a planar graph.

Definition 3.7 (Faces)**.** If $G = (V, E)$ is a planar graph, then G divides the plane into faces, which are enclosed regions surrounded by edges. The unbounded area outside the whole graph is always counted as one face. See Fig. 3.

Remark 3.8. We denote $|F| = f$. If we need to specify the number of faces in a specific graph G, we will denote this as $f(G)$.

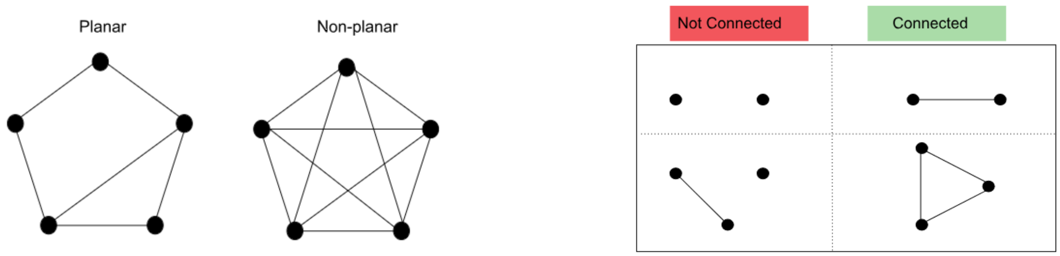

Figure 2: Depictions of planar and connected graphs.

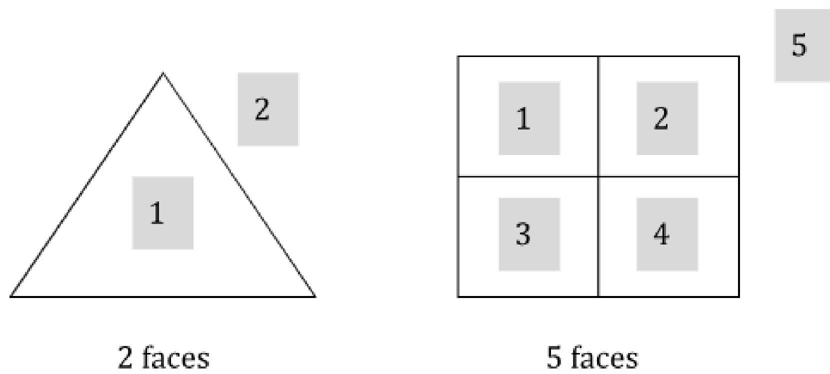

Figure 3: Depiction of faces of a planar graph.

We can now state Euler's theorem for planar connected graphs

Theorem 3.9 (Euler's theorem)**.** *Let G be a planar connected graph. Then,*

$$v(G) - e(G) + f(G) = 2.$$

We describe the proof outline here, but do not go into the complete details. For complete details, one can find the proof in [Bón02].

Proof Outline of Theorem 3.9. We will use induction on the number of edges in G.

Base Case: Let $e = 1$. Given G is connected, there are two possible graphs G that are planar and connected with $e = 1$ (A line segment or a loop). In either case, one can check directly that Euler's Theorem holds.

Inductive Hypothesis: Suppose the theorem holds for all connected, planar graphs with $e - 1$ edges.

Inductive Step: We want to prove that the theorem holds when a planar connected graph G has e edges. There are two cases.

Case 1: If we can remove any one edge from G and obtain a connected graph (which is still planar by assumption), then we call this new graph G' with $e - 1$ edges. By our induction hypothesis,

$$v(G') - e(G') + f(G') = 2. \tag{2}$$

Notice that we have removed no vertices and thus $v(G') = v(G)$. However, since we have removed an edge in a *planar* graph, the number of faces must have decreased by 1 (i.e. $f(G') = f(G) - 1$).

Plugging in these equalities into (2), we have

$$v(G) - (e(G) - 1) + (f(G) - 1) = 2,$$

which gives the theorem.

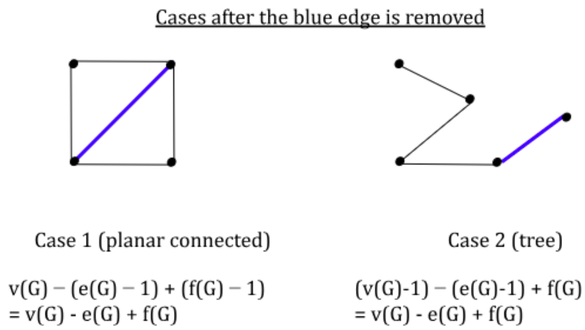

Case 2: If no such edge exists in G, it follows that G must be a tree. A graph is a tree if and only if between every pair of distinct vertices of a connected graph, there exist a unique path and no cycles. Then, one can remove a vertex (and corresponding edge) with degree 1– reducing the number of vertices and edges by 1, and keeping the number of faces the same. Then, applying the induction hypothesis, we see

$$v(G) - e(G) + f(G) = (v(G) - 1) - (e(G) - 1) + f(G) = 2$$

This completes the proof. $\square$

For our purposes, we will study graphs that are not necessarily connected (recall Fig. 1). To deal with this, we note the following generalization of Euler's theorem:

Definition 3.10. A graph G has k connected components if it can be written as a union of k connected graphs (called connected components), where there are no edges between any of the components. We denote the number of connected components of a graph G as $k(G)$. See [Bón02] for more details.

We can now state and prove a more general version of Euler's theorem for planar graphs.

Theorem 3.11. *Let G be a planar graph. Then,*

$$v(G) - e(G) + f(G) = k(G) + 1.$$

Proof of Theorem 3.11. We know that the result holds when $k(G) = 1$ by the proof for planar connected graphs. For a graph with k connected components, we simply apply Euler's theorem on each connected component. Let's call these components $G_1, \ldots, G_k$. Then, we obtain that

$$v(G_i) - e(G_i) + f(G_i) = 2$$

for each i. Then, roughly speaking, we sum all of these equations up obtaining that

$$\sum_{i=1}^{k} (v(G_i) - e(G_i) + f(G_i)) = 2k$$

In this process, notice that

$$v(G) = \sum_{i=1}^{k} v(G_i) \quad \text{and} \quad e(G) = \sum_{i=1}^{k} e(G_i)$$

However, notice that

$$f(G) = \left(\sum_{i=1}^{k} f(G_i) \right) - (k - 1)$$

as we end up overcounting the outside face $k - 1$ extra times.

Therefore,

$$v(G) - e(G) + f(G) = 2k - (k - 1) = k + 1.$$

$\square$

As a corollary of this result, note that we obtain the following relating $|V|$ and $|E|$ (which is our current goal for proving Szemerédi–Trotter).

Corollary 3.12. *If G is planar, then $e - 3v \leq 0$.*

Proof. Suppose G is planar. Then, Euler's theorem states that

$$v - e + f = k(G) + 1 \geq 0$$

where f is the number of faces, e is the number of edges, and v is the number of vertices of G. Notice that, as G is planar, each edge borders two faces and every face is created by at least three edges. Therefore, $f \leq \frac{2}{3}e$. Combining these two inequalities, we obtain the desired result. □

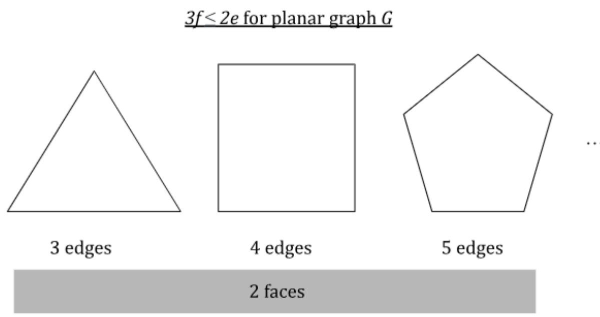

Hence, if our incidence graph, G is planar, then we have already proven the $I(P, \mathcal{L}) \leq n + m$ (using that $v = n$ and $e = I(P, \mathcal{L}) - m$). However, given the unregulated, random process by which cases of points, lines, and subsequent incidences are placed in Euclidean space, there is no guarantee that the graph G will be planar. Thus it is necessary to generalize this inequality in order to be able to relate edges and vertices in this way for any graph G.

3.1.2 The Crossing Number Lemma

Definition 3.13 (Crossing number)**.** Let G be a graph. Then, we define the crossing number of G, $\mathrm{cr}(G)$, to be the minimum number of crossing pairs of edges in any simple drawing of G.

In this sense, the crossing number of a graph measures how far away the graph is from being planar. In fact, one can see directly by the definitions that G is planar if and only if $\mathrm{cr}(G) = 0$. Using crossing numbers, we can obtain the following lemma.

Lemma 3.14. *Let G be a graph. Then, $e - 3v \leq \mathrm{cr}(G)$.*

Proof. We have seen this proven when $\mathrm{cr}(G) = 0$, i.e. when G is planar. If G is nonplanar, we transform G into a planar graph G' via the following method: embed G into the plane with the minimal number of crossings possible (namely, by definition, $\mathrm{cr}(G)$). Then, for each crossing, remove one of the two edges in the crossing and keep all the vertices. We denote the resulting graph $G' = (V', E')$ and notice that G' is planar. Furthermore, by construction, $|V'| = v$ and $|E'| = e - \mathrm{cr}(G)$. Thus, applying Corollary 3.12 to G', we see that

$$e' - 3v' \leq 0 \iff (e - \mathrm{cr}(G)) - 3v \leq 0.$$

Rearranging the inequality gives the result. □

As it turns out, the above lemma is not enough to obtain the Szemerédi–Trotter bound, but we can use it (and probability theory) to obtain the sharper bound known as the crossing number lemma.

Theorem 3.15 (Crossing Number Lemma)**.** *Let G be a graph with $e \geq 4v$. Then,*

$$\mathrm{cr}(G) \geq \frac{e^3}{64v^2}.$$

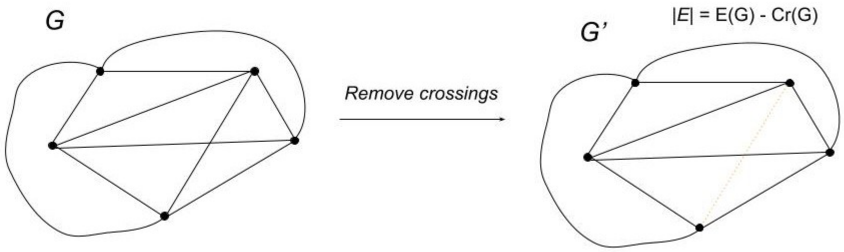

Proof. Let $0 < p \leq 1$. Let G_p be a random induced subgraph of G formed by including each vertex of G independently with probability p. We include an edge of G in G_p if and only if its endpoints are included in G_p.

By Lemma 3.14, we know that

$$E(G_p) - 3V(G_p) \leq \mathrm{cr}(G_p)$$

for every random induced subgraph G_p. Hence, we can take expectations of both sides (using linearity of expectation) to obtain

$$\mathbb{E}[E(G_p)] - 3\mathbb{E}[V(G_p)] \leq \mathbb{E}[\mathrm{cr}(G_p)]. \tag{3}$$

We now just need to relate the expectation for the number of edges, vertices, and crossings to the number of vertices, edges, and crossings in the initial graph G.

The expected value of $V(G_p)$ is pv, as we keep each vertex in G with probability p. Similarly, the expected value of $E(G_p)$ is p^2e, as each edge is included with probability p^2 (i.e. if both endpoints remain in the graph). Lastly, we claim that each crossing in G_p is obtained by two distinct edges (see [Gut] from Larry Guth's course on the polynomial method to verify this fact). For our purposes, this is enough to notice that $\mathbb{E}[\mathrm{cr}(G_p)] \leq p^4\mathrm{cr}(G)$. (Note the inequality here as when we remove vertices and edges, crossings may disappear altogether.) Thus, plugging this into (3), we see that

$$p^2e - 3pv \leq p^4\mathrm{cr}(G) \iff p^{-2}e - 3p^{-3}v \leq \mathrm{cr}(G).$$

We choose p to optimize the right-hand side, namely speaking letting $p = \frac{e}{4v}$. Here, we used that $e \geq 4v$ so that $p \in [0, 1]$. Plugging in this value of p, we get

$$\mathrm{cr}(G) \geq \frac{e^3}{64v^2}.$$

□

This inequality will in fact be enough to conclude the proof of the Szemerédi–Trotter bound.

3.2 Proof of Szemerédi–Trotter

In this section, we prove the Szemerédi–Trotter theorem (Theorem 3.1) which we state below for convenience.

Theorem 3.16. *Let P be a set of n distinct points in $\mathbb{R}^N$ and let $\mathcal{L}$ be a set of m distinct lines in $\mathbb{R}^N$. Then,*

$$I(P, \mathcal{L}) \lesssim n^{2/3}m^{2/3} + n + m.$$

At this point, we have transformed our points and lines into a graph $G = (V, E)$ with

$$v = n \quad \text{and} \quad e = I(P, \mathcal{L}) - m.$$

Now, we break into cases based on when we can apply the crossing number lemma (Theorem 3.15).

Case 1: $e < 4v$. If this holds, then by rearranging Lemma 3.5, we have

$$I(P, \mathcal{L}) = e + m < 4v + m = 4n + m \lesssim n + m.$$

This is part of our result for the Szemerédi-Trotter theorem and is a very good bound.

Case 2: $e \geq 4v$. Then, by the crossing number lemma, we have

$$\frac{e^3}{64v^2} \leq \operatorname{cr}(G).$$

To utilize this, we need notice a trivial upper bound for $\operatorname{cr}(G)$. Indeed, notice that every crossing in G must correspond to two edges intersecting in our graph, which means two lines overlap in $\mathcal{L}$.

This implies that

$$\operatorname{cr}(G) \lesssim m^2.$$

Comparing the upper and lower bounds for $\operatorname{cr}(G)$, we see that

$$\frac{e^3}{64v^2} \leq m^2 \implies I(P, \mathcal{L}) \lesssim n^{2/3}m^{2/3} + m.$$

Here, we used that $e = I(P, \mathcal{L}) - m$. Hence,

$$I(P, \mathcal{L}) \lesssim \begin{cases} n + m & e < 4v \\ n^{2/3}m^{2/3} + m & e \geq 4v \end{cases}.$$

Notice that in either case, the incidences are bounded by

$$\lesssim n^{2/3}m^{2/3} + n + m.$$

This concludes the proof of the Szemerédi–Trotter bound.

3.3 Sharpness over Euclidean Space

We lastly give a few remarks about the sharpness of the Szemerédi–Trotter bound for incidences between points and lines. In particular, we highlight three examples where each of the three terms dominates in Theorem 3.1 and the bound is sharp (perhaps up to some constant).

<u>The bounds of n and m</u>: Suppose we have n points, P, and 1 line, ℓ, such that each of the n points lie on ℓ. Then, $I(P, \ell) = n$ and n is the largest term on the right-hand side of Theorem 3.1. Similarly, suppose we have 1 point, p, and m lines, $\mathcal{L}$, such that every line passes through p. Then, $I(p, \mathcal{L}) = m$, and m is the largest term on the right-hand side of Theorem 3.1.

<u>The bound of $n^{2/3}m^{2/3}$</u>: Here, we present an example where the $n^{2/3}m^{2/3}$ term is the dominant term when $n \sim m$ (see [Wik24]). Consider the set of points and lines

$$\begin{aligned} P &= \{(a, b) \in \mathbb{Z}^2 : 1 \leq a \leq n, 1 \leq b \leq 2n^2\} \\ \mathcal{L} &= \{(x, mx + b) : m, b \in \mathbb{Z} \text{ and } 1 \leq m \leq n, 1 \leq b \leq n^2\}. \end{aligned}$$

In this example, it is clear that $|P| = 2n^3$ and $|\mathcal{L}| = n^3$. One can carefully check that each line contains n points (similar to the finite field example), and thus we have

$$I(P, \mathcal{L}) = n^3 \cdot n \sim (n^3)^{4/3}$$

This proves the sharpness of the Szemerédi–Trotter bound when $n \sim m$.

References

[Bón02] Miklós Bóna. *A Walk Through Combinatorics*. WORLD SCIENTIFIC, 2002. DOI: 10.1142/4918. eprint: https://www.worldscientific.com/doi/pdf/10.1142/4918. URL: https://www.worldscientific.com/doi/abs/10.1142/4918.

[Gut] Larry Guth. *The Polynomial Method Course Page*. URL: https://math.mit.edu/~lguth/PolynomialMethod.html.

[Ios07] Alex Iosevich. *A View from the Top: Analysis, Combinatorics, and Number Theory*. American Mathematical Society, 2007.

[Wik24] Wikipedia. *Szemerédi–Trotter theorem: Optimality — Wikipedia, The Free Encyclopedia*. http://en.wikipedia.org/w/index.php?title=Szemer%C3%A9di%E2%80%93Trotter%20theorem&oldid=1214703970. 2024.