Lower bounds for incidences

Talk by Dima Zakharov

Notes by Sanjana Das

September 5, 2024

This is joint work with Alex Cohen and Cosmin Pohoata.

§1 Approximate incidences

Dima's first goal is to explain and motivate the title. Usually in incidence geometry, you have a collection of points and lines — for example, the usual Szemerédi–Trotter theorem is a statement about the number of incidences between a collection of points and lines in the plane.

Definition 1.1. The number of incidences between $\mathcal{P}$ and $\mathcal{L}$ is defined as

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

Theorem 1.2 (Szemerédi–Trotter)
We have $I(\mathcal{P}, \mathcal{L}) \leq |\mathcal{P}|^{2/3} |\mathcal{L}|^{2/3} + |\mathcal{P}| + |\mathcal{L}|$.

We won't worry too much about this bound, but it's a classic result that started the study of incidence geometry (where you can study incidences between various types of objects — for example, replacing points or lines by something else).

In this talk, we're interested in *approximate* incidences — instead of measuring the number of pairs where the point p lies on the line ℓ, we allow some margin of error. So we'll still have a collection of points $\mathcal{P}$ and lines $\mathcal{L}$, but we'll count (p, ℓ) as an incidence whenever p is *close* to ℓ. We can imagine drawing the δ-neighborhood of each line ℓ; and if p lies in this δ-neighborhood, then we count this as an incidence.

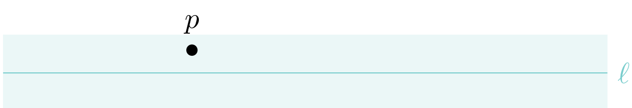

It's convenient to normalize the point set so that it lies in the unit square. Wo we have a collection of points $\mathcal{P} \subseteq [0, 1]^2$ and a collection of lines $\mathcal{L}$, and from $\mathcal{L}$ we form a collection of tubes $\mathcal{T}$ consisting of the δ-neighborhood of each of these lines. And we're interested in $I(\mathcal{P}, \mathcal{T})$.

This is an interesting problem in its own right, and it also has applications to the Kakeya problem and some problems in TCS (for example). There are some difficulties in transferring discrete results to continuous ones, but we've made progress over time. One classical result along these lines is the Ren–Wang Furstenberg estimate; we won't state this, but it's an approximate analog of the Szemerédi–Trotter theorem when $\mathcal{P}$ and

$\mathcal{T}$ are *well-separated*. For what well-separated means, we need some condition on what the points or lines look like — otherwise you could put all the points in a δ-ball and have all the lines intersect this δ-ball, and then $I(\mathcal{P}, \mathcal{T})$ would just be $|\mathcal{P}|\,|\mathcal{T}|$. So we need to forbid this picture somehow, and there are various ways to do this.

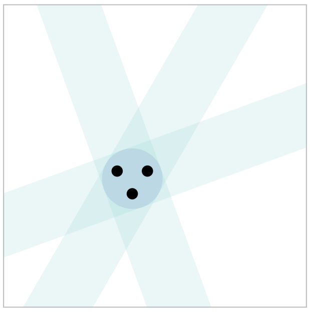

§2 Incidence lower bounds

Szemerédi–Trotter and the Furstenberg estimate are *upper* bounds on incidences, but the title promised *lower* bounds. First of all, we need to make sense of this question — naively, the only lower bound you can give is $I(\mathcal{P}, \mathcal{T}) \geq 0$. So our first goal is to fix this and try to formulate a sensible question (where we can hope there are interesting lower bounds).

§2.1 Incidences in finite fields

To motivate things, we'll move to finite fields — suppose that we have a set of points $\mathcal{P}$ and lines $\mathcal{L}$ in $\mathbb{F}_q^2$. Incidences in finite fields are also an interesting problem; some things are different (e.g., the Szemerédi–Trotter theorem is no longer true), and much less is known. One thing that *is* known is the following nice theorem.

Theorem 2.1 (Vinh 2011)
We have $\left|I(\mathcal{P}, \mathcal{L}) - \frac{1}{q}\,|\mathcal{P}|\,|\mathcal{L}|\right| \leq \sqrt{q\,|\mathcal{P}|\,|\mathcal{L}|}$.

Let's try to parse this bound. On the left-hand side, $I(\mathcal{P}, \mathcal{L})$ is the number of incidences, and $\frac{1}{q}\,|\mathcal{P}|\,|\mathcal{L}|$ is the *expected* number of incidences if we chose $\mathcal{P}$ and $\mathcal{L}$ randomly, or even just chose one of them randomly — a line has q points, so it occupies a $\frac{1}{q}$-fraction of the plane, and so we'd expect a $\frac{1}{q}$-fraction of points to lie on this line. Vinh's theorem is some quantification of this (saying that the actual number of incidences is close to this probabilistic heuristic).

Proof sketch. The proof is by the expander mixing lemma. You can draw a bipartite graph with all possible points on one side and lines on the other side, and you're counting edges between $\mathcal{P}$ and $\mathcal{L}$. You can write down the adjacency matrix for this graph and compute its eigenvalues, and other than the top one, they're all at most $\sqrt{q}$; then the expander mixing lemma gives this bound. □

What Theorem 2.1 says is that if the right-hand side is smaller than the expected number of incidences — e.g., if $|\mathcal{P}|, |\mathcal{L}| \gg q^{3/2}$ — then $I(\mathcal{P}, \mathcal{L})$ is well-approximated by the probabilistic heuristic, meaning that $I(\mathcal{P}, \mathcal{L}) \approx \frac{1}{q}\,|\mathcal{P}|\,|\mathcal{L}|$. There is an upper bound here, but there's also a lower bound. And this is a bit surprising — it means that if you have large sets of points and lines, then they determine lots of incidences.

§2.2 Finite fields vs. points and tubes

Let's try to formulate an analog of Theorem 2.1 for points and tubes. There's a general analogy between $[0,1]^2$ and $\mathbb{F}_q^2$, where lines in $\mathbb{F}_q^2$ look kind of similar to tubes in the unit squre. Specifically, we can think of lines in $\mathbb{F}_q^2$ as similar to tubes of width $\frac{1}{q}$, since q shifts of the line cover the entire plane.

Remark 2.2. In fact, lines in $\mathbb{F}_q^2$ really are tubes of width $\frac{1}{q}$ in the p-adic metric, so in some sense this correspondence is quite precise.

And the *number* of lines in $\mathbb{F}_q^2$ is q^2, while the number of different tubes in $[0,1]^2$ is roughly δ^{-2} (if we treat two tubes that overlap a lot as the same, then we can fix the direction of the tube in δ^{-1} ways and its shift in δ^{-1} ways, and this basically defines all distinct tubes). So the numbers match up; and the same occurs for points.

So incidences between points and tubes in $[0,1]^2$ are kind of analogous to incidences between points and lines in finite fields. And there are ways to make this analogy work — for example, there is an analog of the expander mixing lemma for tubes, which we use a lot.

If we look at this analogy and try to transfer Theorem 2.1, this motivates the following question.

Question 2.3. Suppose $|\mathcal{P}|$ and $|\mathcal{T}|$ are large. Is it true that $I(\mathcal{P},\mathcal{T}) \geq \delta\,|\mathcal{P}|\,|\mathcal{T}|$?

(Here each tube occupies a δ-fraction of the square, so if things were random, you'd expect the number of incidences to be a δ-fraction of all possible point-tube pairs.)

Unfortunately, this is false — somehow the two worlds are similar, but not quite. In the unit square, you can have a bunch of points in the top half and a bunch of tubes in the lower half; then you have almost the maximal number of points and tubes, but they don't interact at all, meaning that $I(\mathcal{P},\mathcal{T}) = 0$.

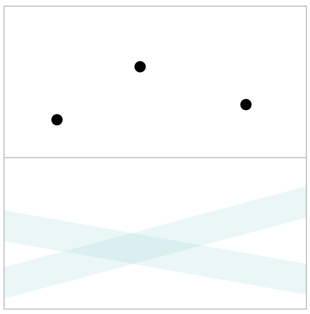

So this is sad; somehow a situation like this doesn't happen in finite fields, but it does for reals.

§2.3 Setup and main result

The main problem we care about right now is a way to make our question make sense. The issue we saw above was that you can kind of separate the points and lines completely by putting them in different regions of the square; so we'll use a setup that forces them to kind of be together.

We'll have a collection of points $\mathcal{P} = \{p_1,\ldots,p_n\} \subseteq [0,1]^2$. And for each point p_j, we'll draw a $1\times\delta$ tube T_j passing through p_j; this gives a collection of tubes $\mathcal{T} = \{T_1,\ldots,T_n\}$.

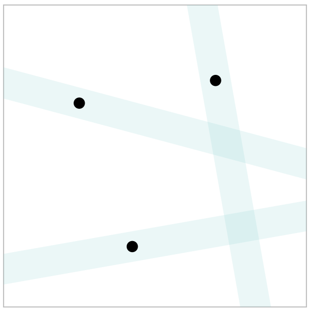

And now we can ask similar questions — what can we say about $I(\mathcal{P}, \mathcal{T})$? By design, there are at least n incidences (each point is in its own tube); so the first thing we might ask is, are there any *other* incidences?

Question 2.4. Is there some $p_i \in T_j$ where $i \neq j$?

We have two parameters n and δ, so the answer is going to depend on how these two are related.

Here's one possible picture: we can have a collection of points on a horizontal line, spaced out by roughly δ (this means $n \approx \delta^{-1}$); and we draw vertical tubes through each.

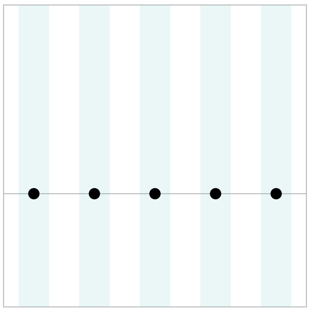

Then no tube hits any other point, so there are no (nontrivial) incidences. And with these parameters, you can actually make lots of examples with no incidences — you can place points however you like, and then find a tube through each point that avoids all the other points. So there's nothing special about this picture.

On the other hand, suppose that $n \gg \delta^{-2}$ (so we have a lot of points). Then by packing, we can find two points that are close together — specifically, with $d(p_i, p_j) < \delta$ (the best way to spread the points out is to make a grid). Then the tube through p_i necessarily has to contain p_j.

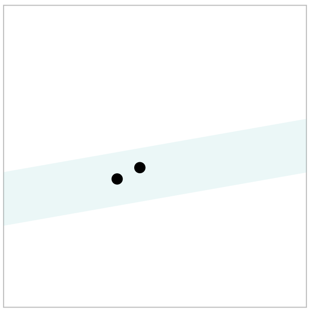

So for this question to be interesting, we need $\delta^{-1} \ll n \ll \delta^{-2}$. Here's the result the authors prove: we can't solve the problem for the *full* range, but we land somewhere in the middle.

Theorem 2.5 (Cohen–Pohoata–Zakharov 2024++)
If $n > \delta^{-3/2+o(1)}$, then there exist $i \neq j$ with $p_i \in T_j$.

By a standard subsampling argument, once you've guaranteed a *single* incidence, you can find a nontrivial incidence in any random subset, and this gives you lots of incidences.

Corollary 2.6
In this situation, we have $I(\mathcal{P}, \mathcal{T}) \gtrsim \delta^{3/2} |\mathcal{P}| |\mathcal{T}|$.

We won't say too much about the proof of Theorem 2.5, but we'll talk about some related things.

§3 Finite fields

First, what's the $\frac{3}{2}$ doing in Theorem 2.5? It turns out that it's the same $\frac{3}{2}$ from the finite fields setting (where we saw that if $|\mathcal{P}|, |\mathcal{L}| \gg q^{3/2}$, then $I(\mathcal{P}, \mathcal{L})$ is roughly what you'd get from the random heuristic).

So let's move back to finite fields for now, where we have a collection of points and lines. We'll use the same setup, where we have points $p_1, \ldots, p_n \in \mathbb{F}_q^2$ and lines ℓ_j through p_j. We can then ask the same question.

Question 3.1. For which n (as a function of q) does there always exist a nontrivial incidence (i.e., some $i \neq j$ with $p_i \in \ell_j$)?

So in our picture, we have a bunch of points, and we're trying to draw lines through these points that avoid all our other points.

First, by Vinh's estimate (Theorem 2.1), we know that if $n \gg q^{3/2}$, then the answer is yes — in this regime we know that the total number of incidences is roughly a $\frac{1}{q}$-fraction of the total number of pairs (p, ℓ); this is substantially more than the number of trivial incidences, which means that we can find some nontrivial incidence. This is the same kind of statement as Theorem 2.5 (recall that q is analogous to δ^{-1}).

And in the finite field setting, this is actually sharp — there is a nice example when $q = p^2$. We'll present this example because it's a nice object that it's good to know about. Intuitively, it looks kind of like taking a set of points on a circle and drawing the tangent line through each point (which avoids the other points on the circle).

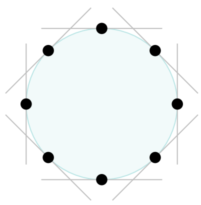

Example 3.2
For $q = p^2$, we can define a norm $\mathcal{N}(x) = x\overline{x} = x^{p+1}$ (where $\overline{x} = p^2$). We can then define a 'sphere' over the finite plane, which is called a *Hermitian unital*, as

$$\mathcal{P} = \{(x, y) \in \mathbb{F}_q^2 \mid \mathcal{N}(x) + \mathcal{N}(y) = 1\}.$$

Then for each point $p_i = (x, y)$, we define the tangent line $\ell_i = \{(x, y) + t(\overline{y}, \overline{x}) \mid t \in \mathbb{F}_q\}$.

A good way to think about the Hermitian unital is to imagine replacing $\mathbb{F}_p$ with $\mathbb{R}$ and $\mathbb{F}_q$ with $\mathbb{C}$. Then we're looking at pairs of complex numbers whose norms sum to 1; and if we write our complex numbers as pairs of reals, we get a 3-dimensional sphere.

We can check that in this example, we have $\ell_i \cap \mathcal{P} = \{p_i\}$ (so there are no nontrivial incidences). And we have $|\mathcal{P}| = q^{3/2}$ (there are q^2 options for x and y, and we're imposing one constraint $\mathcal{N}(x) + \mathcal{N}(y) = 1$; but since $\mathcal{N}(x)$ and $\mathcal{N}(y)$ live in $\mathbb{F}_p$, this constraint only restricts to a $\frac{1}{p}$-fraction of our points).

So this is a sharp example for the point-lines problem over finite fields, showing that $\frac{3}{2}$ is a special number for this problem. Interestingly, however, when q is prime we don't know of an example like this.

Question 3.3. If q is prime and $n > q^{3/2-\varepsilon}$, can we still find a nontrivial incidence $p_i \in \ell_j$ (with $i \neq j$)?

We don't know, and this is a really interesting open problem.

This ends the finite fields story for now.

§4 The first idea

We'll now say a bit about the proof of Theorem 2.5 — we can't say too much, but we'll talk about the first idea that starts the whole endeavor.

The proof is kind of by induction — we want to prove an exponent of $\frac{3}{2}$, and to get this, we start with an exponent of 2 (which is easy) and improve it bit by bit. And we combine this induction with the *symmetries* of the problem.

§4.1 Symmetries

What do we mean by symmetries? If you have a collection of points and tubes, you can look at various subsets of this configuration. For example, you can look at a smaller square and consider the points inside it; or you can look at a large tube and consider all the tubes inside it; or maybe you could look at something more complicated.

Here's what the full group of symmetries of the problem looks like. We start with the unit square $[0, 1]^2$, and we pick a rectangle $\mathcal{R}$ (not necessarily axis-aligned); we call its long side u and short side uw (for parameters $u, w > 0$). We can then try to restrict our problem to this rectangle.

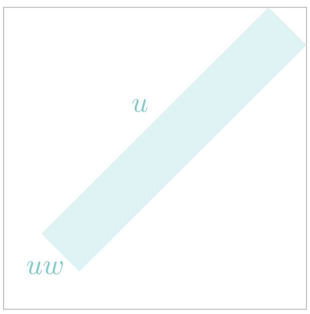

First, we have points and tubes which are paired together, so it's convenient to think of them in terms of the set of point-line pairs; so we define

$$X = \{(p_j, \ell_j) \mid j \in [n]\}.$$

(We can think of X as a subset of $\mathbb{R}^3$, since you're choosing a point and then a line through it.)

Now we have a rectangle $\mathcal{R}$, and we want to restrict to pairs $(p, \ell) \in X$ satisfying certain conditions. For one thing, we want $p \in \mathcal{R}$. But that's not all — we also want ℓ to be aligned with $\mathcal{R}$. More formally, what this means is that the slope of ℓ should be close to the slope of $\mathcal{R}$ — specifically,

$$|\text{slope}(\ell) - \text{slope}(\mathcal{R})| \leq w.$$

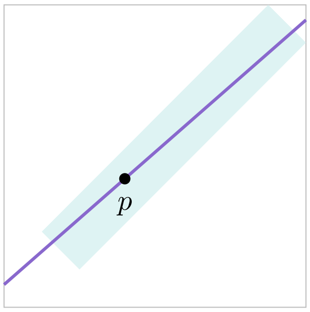

(This makes sense because the aspect ratio of $\mathcal{R}$ is w, so we should allow lines to wiggle in a w-window.)

Definition 4.1. We define $X \cap \mathcal{R} = \{(p, \ell) \mid p \in \mathcal{R}, \ell \text{ aligned with } \mathcal{R}\}$.

Now the key point is that we've restricted to this small rectangle $\mathcal{R}$, and we can blow $\mathcal{R}$ back up to the unit square to get the same problem, but with different parameters.

We're starting with a $uw \times u$ rectangle; we first stretch its short side to make it a $u \times u$ rectangle, and then dilate it to the unit square. We'll call this map $\psi_{\mathcal{R}} : \mathcal{R} \to [0, 1]^2$ — so $\psi_{\mathcal{R}}$ is the affine map that stretches our rectangle $\mathcal{R}$ back to the unit square.

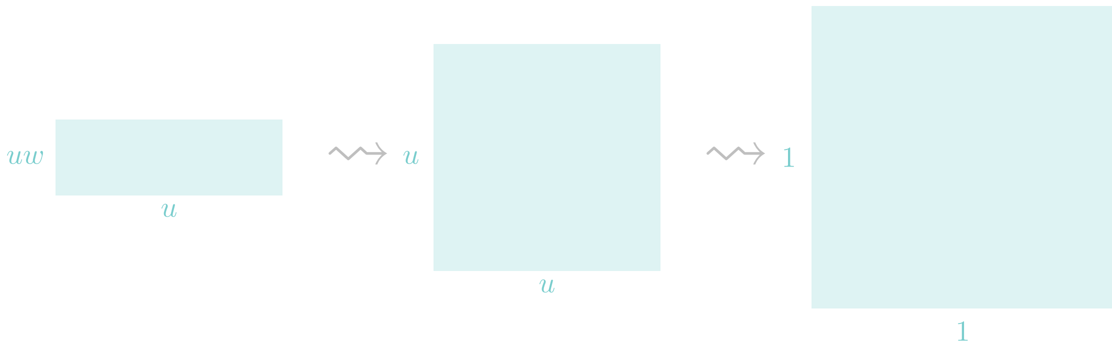

Then we can consider $X_{\mathcal{R}} = \psi_{\mathcal{R}}(X \cap \mathcal{R})$ — so we're looking at all the point-line pairs that land inside our rectangle where the line is also aligned with the rectangle, and blow them back up to the unit square. This stretches the tubes by a factor of roughly $1/uw$. So if we assume the original set had no nontrivial incidences at scale δ, then the new set $X_{\mathcal{R}}$ has no nontrivial incidences at scale δ/uw.

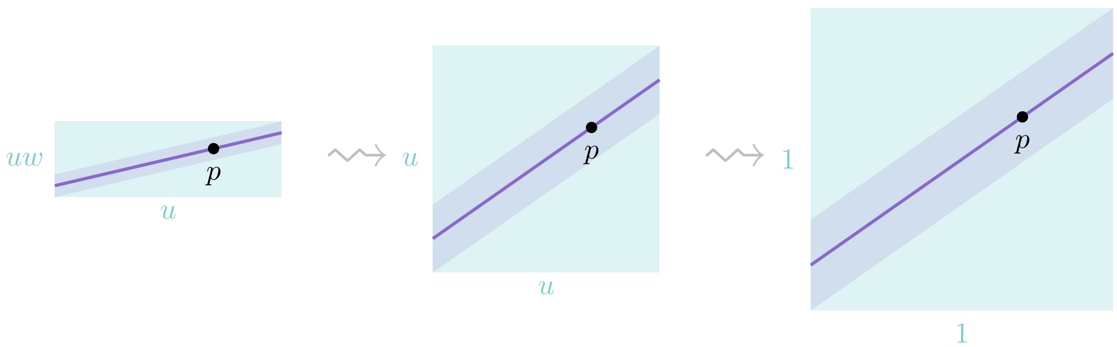

§4.2 Induction

So we've gotten another incidence problem, but with a different parameter δ, which is closer to 1. And now we can use induction to show that $X_{\mathcal{R}}$ has to be small. Specifically, our goal is to show that $|X| \leq \delta^{-\gamma}$, where we're trying to use induction to make γ approach $\frac{3}{2}$. Applying such an inductive statement to our subset $X_{\mathcal{R}}$ gives

$$|X_{\mathcal{R}}| \leq (\delta/uw)^{-\gamma} \approx u^{\gamma} w^{\gamma} |X| . \tag{4.1}$$

So just using this inductive assumption, you can get some sort of decay statement — that if we choose any rectangle, we get this power decay for its intersection with our set X.

It might be unclear how to use this, but this is the basic starting point. Dima wants to emphasize that it's important that we're using a *two-dimensional* family of symmetries. When you think about this problem, it's not surprising that you'd restrict to *subsquares*; but we actually need to use rectangles of all possible aspect ratios. This means there's a two-dimensional family of possible ways to restrict, which really helps with the analysis (without this, the authors don't know how to make things work). Then you end up getting a 2-dimensional Lipschitz function and analyzing it, though we won't go into that.

We'll now state the main theorem of the paper, which is a theorem about collections of point-line pairs (i.e., points together with lines through them) which satisfy a decay assumption like (4.1); the rest of the work goes into proving this.

Theorem 4.2

Let $\gamma > \frac{3}{2}$, and let X be a set of point-line pairs such that

$$|X_{\mathcal{R}}| \leq u^{\gamma} w^{\gamma} |X|$$

for all $uw \times u$ rectangles $\mathcal{R}$. Then $I(\mathcal{P}, \mathcal{T}) \approx \delta |\mathcal{P}| |\mathcal{T}|$, where $\mathcal{P}$ is the set of points in X, and $\mathcal{T}$ is the set of $1 \times \delta$ tubes around lines in X.

(Here we interpret $\approx$ quite liberally; it's really $\delta^{1+o(1)}$.)

This is a way to prove an analog of Vinh's estimate from before (Theorem 2.1) — the $\frac{3}{2}$ that goes into this decay property is kind of similar to the $q^{3/2}$ from Vinh's theorem.

The proof uses Fourier analysis. Very roughly, the δ-error tolerance is kind of like blurring the points and lines by δ. Then if you take the Fourier transform, the points and lines live in the ball of radius δ, which you can dyadically decompose into shells. Then the L^2 norms of these shells correspond to Vinh's error term (the right-hand side of Theorem 2.1) — Vinh's theorem only has one scale while here there are dyadically many, but it's still similar. You need to prove an initial estimate that you have the correct number of incidences on some large scale, and then pass from that large scale to some small scale. And it's a combination of this with the rectangle stuff that makes things work.

Remark 4.3. The exponent $\frac{3}{2}$ is sharp for at least the *upper* bound on $I(\mathcal{P}, \mathcal{T})$, because of a Szemerédi–Trotter-like example. For the lower bound, we're not sure whether it's sharp — it might be possible that even if γ is very close to 1, we still get at least the correct number of incidences; but proving this seems out of reach.