Lecture 16 (November 1, 2022)

Scribe: Joe Poremba

Overview of Today's Lecture:

- Crossing Number Inequality
- Ham Sandwhich Theorem
- Colinearity Testing

Part 1: Crossing Number Inequality

Recall, for a graph G its crossing number $Cr(G)$ is the minimum number of edge crossings in a drawing of G in the plane (curved edges are allowed in the embedding).

Recall this fact from last time, which follows from Euler.

Simple Bound: $Cr(G) \geq e - 3n$.

We will show the following lower bound, using a probabilistic proof.

Theorem: For $e > 4n$, we have that $Cr(G) \geq \frac{e^3}{64n^2}$.

> Proof: Take a drawing of G with $Cr(G)$ edges crossings. Choose a random (induced) subgraph H by independently and uniformly picking each vertex with probability p, for some p that we will specify later.
>
> Define the random variables $N_H = |V(H)|$ and $E_H = |E(H)|$. Obeserve that $\mathbb{E}[N_H] = pn$ (since each vertex is selected with probability p) and $\mathbb{E}[E_H] = p^2 e$ (since for an edge to appear, both of its ends must be selected, which happens with probability p^2).
>
> Now, define the random variable X_H to be the number of edge crossings in H (in the same drawing). For a given edge crossing of (this embedding of G) to appear in H, all four vertices must appear in H, which happens with probability p^4. Note that any edge crossing will have 4 distinct vertices involved, since otherwise it would be avoidable. (Draw a picture!). So $\mathbb{E}[X_H] = p^4 Cr(G)$.
>
> The magic happens now. From our simple bound, we have that $X_H \geq Cr(H) \geq E_H - 3N_H$. This holds for any outcome. Therefore, it also holds for their expectations. Hence, $p^4 Cr(G) \geq p^2 e - 3pn$. Now, selecting $p = 4n/e$ (which is in $(0, 1)$, since $e > 4n$), and isolating for $Cr(G)$ in the inequality, we obtain that $Cr(G) \geq \frac{e^3}{64n^2}$, as desired.

Two brief asides:

- Why did we select p this way? We could use any $p \in (0,1)$ and get a true result. A little optimization to get the largest lower bound yields this choice of p.
- There are slightly better crossing number bounds, but essentially all of them use similar ideas to this.

Part 2: Ham Sandwhich Theorem

After our brief foray into crossing numbers, we return to line arrangements and duality.

What is this weird title? The theorem is about slicing objects. Imagine you have a sibling, and you want to share a ham & cheese sandwich with them. Maybe the cheese and ham are not quite perfectly centred and flush with the bread. Regardless, you want to find a way to slice the sandwich such that you both end up with an equal amount of ham, cheese, and bread.

> Theorem (Ham Sandwich Theorem, 1938):
>
> In d dimensions, any d measurable objects can be cut in half by a single $(d-1)$-dimension hyperplane.

We will prove this theorem for $d=2$ using duality. More specifically, here is our setup.

> Theorem: Given any two point sets A and B in the plane, there exists a line that cuts both sets in half.

To be very formal with this, you have to be a little careful with what ``cut in half" means (e.g. what if your line goes a point?), but set that aside for now. We will also assume that no three points are colinear.

> Proof: Take the dual! What happens to our problem? Our points become lines, and (by the order-preserving property) we are now searching for a point that has half of the A^* lines above and half below, and half of the B^* lines above and half below.
>
> Consider just one of the line sets at a time. Just look at the line arrangement for A^*. What points have half of the lines above and half of the lines below? Say the *level* of a point is how many lines are above it. Go very far to the left (beyond all intersections). The A^* lines are ordered vertically. Call this the left-vertical order (we could similarly define the right-vertical order). Take any point on the middle line in the left-vertical order. Any such point has middle level. We can walk to the right along the line until we hit a vertex of the arrangment, and all the points we touch have middle-level. When we hit a vertex, the two lines crossing swap orders, so we swap lines and continue walking. These points also have middle-level. We continue walking in this way until we exit the A^* line arrangement. Call this the $$A^*$ middle-walk. It is not a straight line, but is continuous.

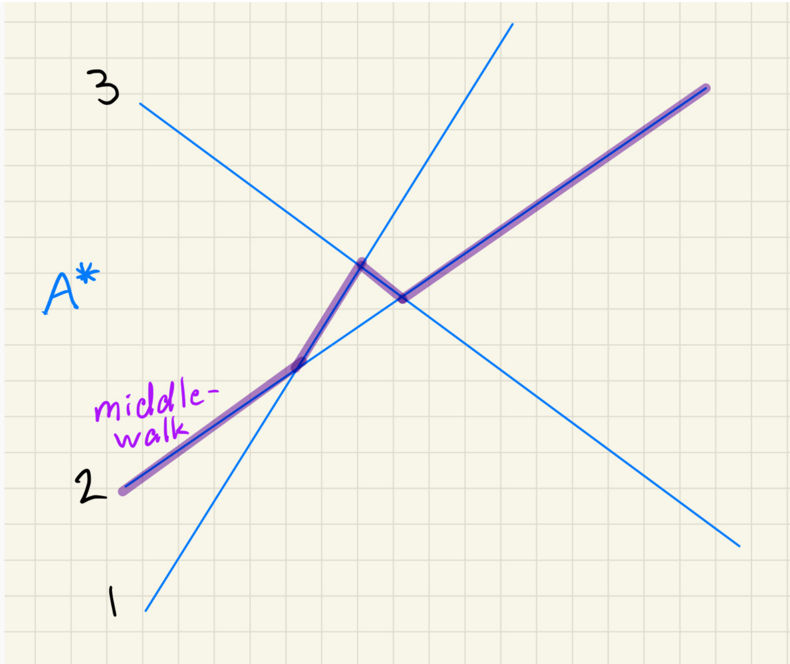

> Of course, we can do the same for B^*. What we want is a point that is in *both* the A^* and B^* middle-walks. Do these walks need to intersect? Yes! Observe two things:
>
> 1) The left-vertical order of A^* is the opposite of its right-vertical order. So, the middle-walk of A^* starts and finishes along the same line a^*. Similarly for B^* we get its start and end line b^*.
>
> 2) If you zoom out to look at the left- and right-vertical order of $A^* \cup B^*$, it has this same switching property. So a^* and b^* flip relative order on the left and right, like in this illustration.

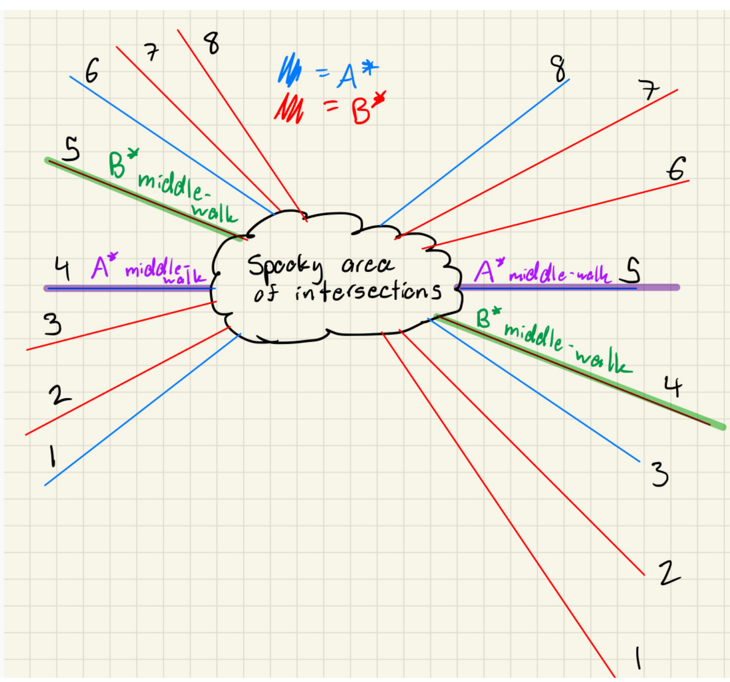

> The walks switch betwen lines inside the intersection cloud, but they are continuous! They have to cross somewhere in order to flip relative order! Hence, they intersect! □

Part 3: Colinearity Testing

We had assumed there are no 3 colinear points. In fact this is a degenerate case for many algorithms. But are we actually able to test for this efficiently?

Here's one simple algorithm:

- Dualize! We are now looking for the intersection of 3 lines.
- Just build the line arrangment and look for one! This takes $\Theta(n^2)$ time.

Hmm… but this seems a bit overkill for finding three points. Surely we can do better, right?

Right…?

☹️ This is the best known ☹️

We don't know that this is tight, but we have ``evidence'' that it is hard.

Recall the 3-SUM problem: Given n numbers, are there 3 that sum to zero?

It is* believed for 3-SUM that quadratic* is the best we can do. We say a problem P is 3-SUM-hard if an algorithm for P that runs in subquadratic* time implies a subquadratic* time solution to 3-SUM.

Theorem: 3-COLINEAR is 3-SUM-hard.

> Proof: Say we are a set S of n numbers and want to solve 3-SUM. Create the points (x, x^3) for $x \in S$. We claim (and leave as an exercise) that $a + b + c = 0$ if and only if the points $(a, a^3), (b, b^3), (c, c^3)$ are colinear. The result follows. □

Okay what's up with the *asterisks? Until around 2014, it was believed that 3-SUM required time $\Omega(n^2)$. However, that year it was proved by Grønlund & Pettie that there exists an algorithm that runs in time: $O(n^2/(\log n/\log\log n)^{\frac{2}{3}})$. This is smaller than quadratic, but barely. It is not even as good as $O(n^{1.99999})$ or $O(n^{2-\epsilon})$ for any constant $\epsilon > 0$. So we modify ``subquadratic'' above to mean $O(n^{2-\Omega(1)})$, and then 3-SUM-hard makes sense again.