Arrangements

Definitions and Combinatorial Results

Let $\mathcal{L} = \{\ell_1, \cdots, \ell_n\}$ be s collection of n lines. An **arrangement** of $\mathcal{L}$, written $A(\mathcal{L})$, is a planar subdivision induced by $\mathcal{L}$. In $A(\mathcal{L})$,

- the vertices are the intersection points of $\mathcal{L}$;
- edges are maximal portions of lines between two adjacency vertices; and
- faces are the maximal (open) connected portions of $\mathbb{R}^2 \setminus \bigcup_{\ell \in \mathcal{L}} \ell$.

Assume that no three lines meet at a same point. Then there are $\binom{n}{2}$ vertices, n^2 edges, and $n^2/2 + n/2 + 1$ faces by Euler's formula $v - e + f = 1$.

A few more definitions:

- For a cell C, let its complexity be $|C|$ = the number of edges (equivalently, vertices) C has.
- For $p \in \mathbb{R}^2$, the **level** of p, written $\operatorname{lev}(p)$, is the number of lines in $\mathcal{L}$ that lie below p. Observe immediately that the level function is constant within a face or on an edge. Hence it is well-defined to consider the level of an edge/face too.
- We can further partition edges based on their levels: ***k*-level**, written $A_k(\mathcal{L})$, is the (closure of the) set of edges whose level is k. Consequently, each k-level is a polygonal chain that is monotone in the x-direction, bending only at vertices.
- Note that 0-level is the lower envelope: if ℓ_i defines the linear function f_i, then the 0-level is nothing but $L(x) = \min_{1 \leqslant i \leqslant n} f_i(x)$. Conversely, $(n-1)$-level is the upper envelope. In general, k^{th}-level encodes the k^{th} order Voronoi diagram.

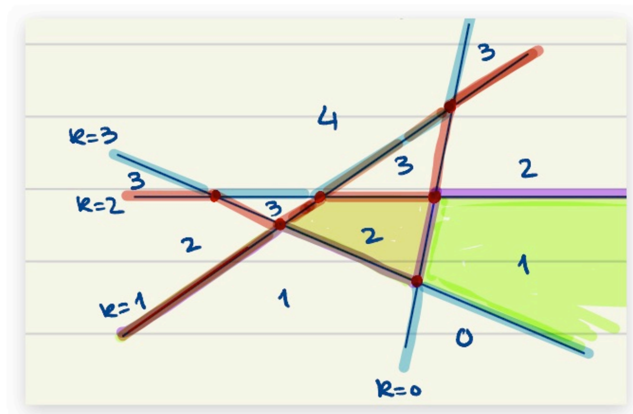

Figure 1: An arrangement of 4 lines, with 0-level (blue), 1-level (purple), 2-level (red), and 3-level (light blue). From lecture notes.

Theorem

$|A_k(\mathcal{L})| = \mathcal{O}(nk^{1/3})$ and also lower bounded by $\Omega(n \log k)$. It is unknown if a tighter upper bound exists.

One can also consider the union $A_{\leqslant k}(\mathcal{L}) = \bigcup_{j=0}^{k} A_j(\mathcal{L})$; it is known that $|A_{\leqslant k}(\mathcal{L})| = \Theta(nk)$. Therefore, *on average* a complexity of a level is $\mathcal{O}(n)$, but not much can be said about the worst case bound.

Given a line $\ell \notin \mathcal{L}$, we define the **zone** of ℓ given $\mathcal{L}$ to be the set of sells of $A(\mathcal{L})$ that intersect $\mathcal{L}$. Define $z(\ell, \mathcal{L}) = \sum_{C \in \text{zone}(\ell, \mathcal{L})} |C|$. It is clear that ℓ may intersect $\mathcal{O}(n)$ cells, but surprisingly, so is $z(\ell, \mathcal{L})$.

Theorem: Zone Theorem

For a line $\ell \notin \mathcal{L}$, $z(\ell, \mathcal{L}) = \mathcal{O}(n)$.
A corollary of the Zone Theorem is that $\sum_{C \in A(\mathcal{L})} |C|^2 = \mathcal{O}(n^2)$.

Computing the Arrangement

Naïvely, we can use a line sweep algorithm to compute the arrangement, which will take $\mathcal{O}(n^2 \log n)$ time because the size of the arrangement is $\mathcal{O}(n^2)$. Alternatively, randomized incremental algorithm which takes $\mathcal{O}(n^2)$ time.
We claim that a simple, deterministic incremental algorithm suffices, and the order of increment does not matter.

- We add lines one by one, say along $\ell_1, \cdots, \ell_n$. Define $\mathcal{L}_i = \{\ell_1, \cdots, \ell_i\}$.
- At each step, we add ℓ_{i+1}, then compute $A(\mathcal{L}_{i+1})$ from $A(\mathcal{L}_i)$ as follows:
  - Trace ℓ_i from left to right, tracing its intersection with $A(\mathcal{L}_i)$.
  - Find the leftmost face C intersected by ℓ_{i+1} (e.g. by checking the slope of ℓ_{i+1} with other lines in linear time).
  - Walk along the boundary of C to find the next intersection point. When this happens, split the encompassing cell into two, and split edge accordingly. Move to the newer intersection point, and iterate.

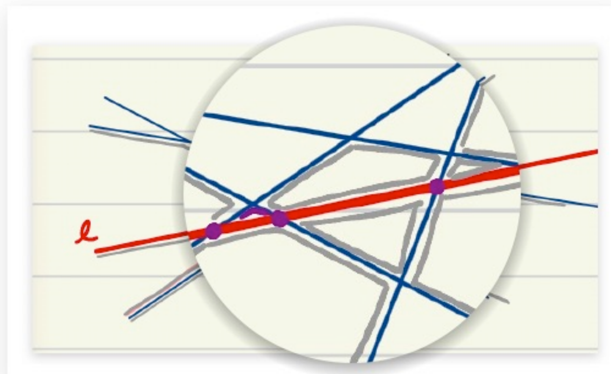

Figure 2: For each cell that ℓ intersects, we march counterclockwise from the first time it is intersected (from left to right) until we reach the other intersecting point. This tells us how to split the edges and the cell itself.

Assuming the Zone Theorem, time spent in inserting ℓ_{i+1} is $\mathcal{O}(z(\ell_{i+1}, \mathcal{L}_i)) = \mathcal{O}(i)$, so the total time of this algorithm is $\mathcal{O}(n^2)$ deterministically.

Applications

We consider two direct applications of geometric questions on $\mathbb{R}^2$ that can be related to arrangement via duality.
Degeneracy testing. Consider P a collection of n points in $\mathbb{R}^2$. Consider the question: Are any three of them colinear? A naïve solution would be to check every triple, yielding $\mathcal{O}(n^3)$ time.
Taking the dual, we transform the question into one about arrangement. We compute the arrangement $A(P^*)$ and check if it admits a vertex of degree > 4 (so three or more lines intersect at one point). Doing so only uses $\mathcal{O}(n^2)$ time, and it has been shown that this is the best achievable complexity.

Linear separability. Let $R, B \subset R^2$ be the sets of red/blue points, and we ask if R, B are linearly separable in $\mathbb{R}^2$.

Taking the dual, if R, B are separated by ℓ, then the lower/upper envelopes of R^*, B^* will contain ℓ^* in their intersection. Hence it suffices to check whether the two envelopes have a nonempty intersection.This relates back to checking appropriate k-levels of the arrangement.

Arrangements of 2-D Curves

Instead of linear functions, consider $\Gamma = \{\gamma_1, \cdots, \gamma_n\}$, a collection of x-monotone curves, and suppose each pair of curves intersects at $\leqslant s = \mathcal{O}(1)$ points. Then the arrangement $A(\Gamma)$ and the lower envelope $L_\Gamma(x) = \min_{1 \leqslant i \leqslant n} \gamma_i(x)$ can be defined similarly (note that the lower envelope need not be convex).
Question. How many breakpoints on $L_\Gamma(x)$ beyond the naïve upper bound is $s\binom{n}{2}$.
A combinatorial persepctive to this problem is by viewing each curve as a symbol. Then, we wish to write a sequence using the alphabet $\{1, \cdots, n\}$, subject to two rules: (i) no two adjacent characters are identical, and (ii) the number of alternations between any distinct $i, j \in \{1, \cdots, n\}$ is at most $s + 1$: the sequence cannot admit a subseuqnece of form $\{i, j, i, j, \cdots\}$ with a total length of $> s + 1$. Let $\lambda_s(n)$ denote the maximum length of admissible sequences. Then $\lambda_1(n) = n$ (as in the line arrangement case) and $\lambda_2(n) = 2n - 1$. Tight bounds for $s > 2$: $\lambda_3(n) = \Theta(n\alpha(n)), \lambda_4(n) = \Theta(n2^{\alpha(n)})$, and $\lambda_{2s+2}(n) = \Theta(n2^{\alpha^s(n)(1+o(1))})$, where $\alpha(n)$ is the inverse Ackermann function.

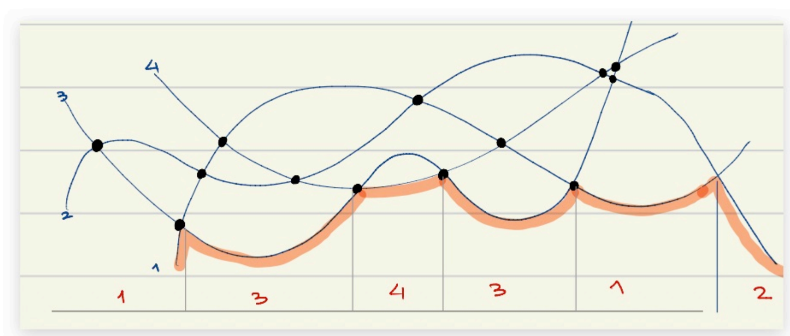

Figure 3: A lower envelope for the collection of curves. Observe that for example, curve 3 appears multiple times, something impossible if curves were replaced by lines, or if $s = 1$.