Line Arrangements

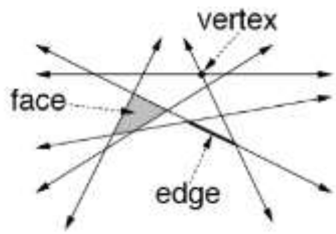

Let L be a set of n lines. Assumption: no 3 lines go through the same vertex. The arrangement defines a planar graph.
For simplicity we assume the infinite faces are bounded – done by connecting all edges of unbound faces to some point in infinity. The object we're dealing with are not necessarily segments (could be unbounded on some end).
For n line segments, the number of vertices will not be $O(n)$, but $O(n^2)$. That goes for the number of vertices, faces and edges.

Theorem

Given a set L of n lines in the plane:

- # vertices: $\frac{n(n-1)}{2}$
- # edges: n^2
- # faces: $\frac{n(n-1)}{2} + n + 1$

Vertices:

For n lines, every pair forms 1 intersection = vertex, yielding $\frac{n(n-1)}{2}$ vertices.

Edges:

By induction: when $n = 1$, we have 1 edge. Now, assume we have an arrangement of $n - 1$ lines, and by assumption there are $(n-1)^2$ edges. When we add a new line, it adds:

- $n - 1$: since each of the previous $n - 1$ edges is cut into 2 edges.
- n: since the new line is cut $n - 1$ times by the old ones – resulting with n edges.

The total: $(n-1)^2 + (n-1) + n = n^2$.

Faces:

By Euler's formula: $v - e + f = 2 \Rightarrow 1 + \frac{n(n-1)}{2} - n^2 + f = 2$.

We therefore would like to have an $O(n^2)$ algorithm (time **and** space) to construct this data structure, and will consider it efficient. The same can be done over curves (not straight lines).

Incremental Construction

Input: a set $L = \{l_1, \dots, l_n\}$ of n line segments, each represented by its equation.

Adding the line l_i should cost $O(i)$, and then the total running time for the data structure $\bigcup_{i=1}^{n}\{l_i\}$ is $\sum_{i=1}^{n} O(i) = O(n^2)$.

Assume $i-1$ lines that are already added. Those lines will intersect the bounding box in $2(i-1)$ points – 2 points for each line. Therefore when adding a new line l_i, we need to identify the edges of the bounding box l_i will intersect – that will take $O(n)$ time.

When found the entering edge, we immediately know the corresponding face and its chain, and from it (using one of the twins) we are led to the next face and so on. Along the traversal we update the data structure. We need to show that although at this step we have $O(i^2)$ edges, the traversal will take only $O(i)$.

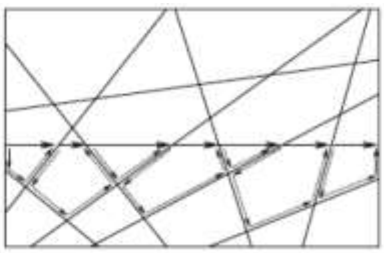

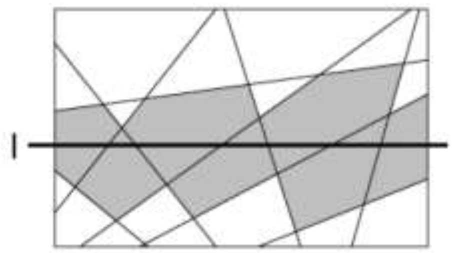

In worst case we go along all $O(i)$ faces, each consisting of $O(i)$ edges, the total complexity is $O(i^2)$. In reality the total checks (intersection queries) is $O(i)$, by the **Zone theorem**.

The Zone Theorem

Let $l \in L$ and let A be an arrangement (DCEL). We denote $Z_A(l)$ as the zone of l, which is the set of faces whose closure intersects l. The complexity of $Z_A(l_i)$ is $O(i)$:

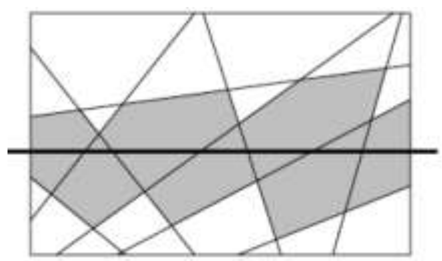

The theorem:

Assume $A(L)$ is an arrangement of n lines and l is added, then the number of edges of $Z_A(l)$ is $\leq 6n$.

Proof:

Based on an inductive argument: we assume l is horizontal (for simplicity). Next, if we have parallel lines to l we infinitesimally rotate them. Looking at any face in the arrangement, it is a convex object, constructed of left and right monotone chains. We will bound the number of chains on the left chain to $3n$, and the argument of the right chain is symmetric. Adding them up will give $6n$.

Look at the last line l' we hit before l leaves the box. l' is a left chain of that rightmost face. We want to prove the inductive step. When we remove l', $|L'| = |L - \{l'\}| = n - 1$, and by assumption $e' = 3(n-1)$ – the total number of edges in the left chain.

We now put l' back in the arrangement:

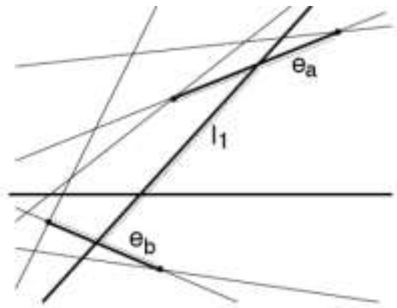

The face l' intersects is an open face. When we add l', it cuts at most 1 edge up an 1 edge in the bottom, denoted e_a, e_b. Both those edges are broken in 2 halves – one half is part of the new face, and 2 stay out. In addition to e_a, e_b, l' itself adds an edge to the left side of the chain. Note that the rest of l', e_a and e_b is invisible to l, therefore only at most 3 edges will be added by l', therefore there are now at most $3(n-1)+3=3n$ edges. The rest follows, and we get $Z_A(l) \le 6n$.

□

It follows that at the ith stage, the number of checks we will have to do for l_i is $O(i)$ – the number of edges in $Z_A(l_i)$ that we will (at most) query for intersection with l_i.

Applications of Arrangements and Duality

- Check that for n points in the plain, no 3 are collinear.
- Minimum area triangle: find the minimum area triangle among n points in the plain.
- Minimum k-corridor: identify the narrowest pair of parallel lines that enclose at least k points of a set of n points.
- Visibility graph: 2 points u, v are visible if the interior of the segment $\overline{uv}$ doesn't intersect anything.
- Maximum stabbing line: for n line segments in the plain, we want to find the line that intersects the maximum.
- Hidden surface removal: given an object and a visibility point, determine which parts of the object are seen.
- Ham sandwich cut: find a single line that bisects the sets of n red points and m blue points.

Duality:

In the $\mathbb{R}^2$ plain, each line is represented by the equation $y(x)=ax-b$, where a is the slope and $-b$ is the intercept.

For any given (a,b) we can construct the line. Assume we have the space where a is the x-axis, and b is the y-axis.

In this space, each line l is represented by a point l^* in that **dual** space. BUT, vertical lines cannot be represented (as the slope is ∞). We'll ignore those lines.

A point: characterized by infinite number of lines that intersect at the same point. Denote that point p, then $y_p = ax_p - b$, and as we change a, b we get different lines, but all go through p.

From that we get $b = x_p \cdot a - y_p$. In the dual space it is **a line**, with slope x_p and intercept y_p, and all the infinite lines that goes through p will become points on the line $b = x_p a - y_p$ in the dual space.

<u>Conclusion</u>: arrangements have dual arrangements in the dual space, and it simplifies some problems.

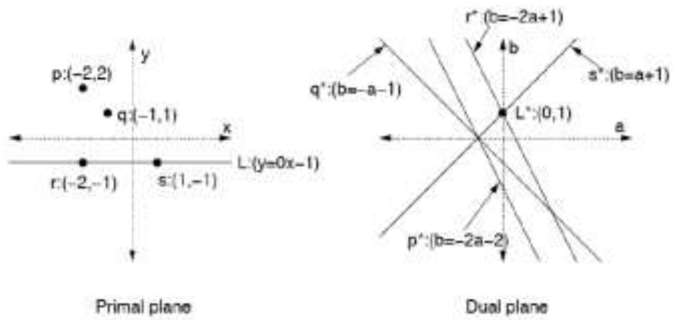

Properties:

- Self-Inverse: $(p^*)^* = p$
- Order reversing: p lies above/on/below l in the primal plane ⇔ line p^* passes below/on/above point l^* in the dual.
- Intersection preserving: l_1 and l_2 intersect in point p in the primal space ⇔ points l_1^*, l_2^* reside on the line p^* in the dual.
- Colinearity/coincidence: 3 points are collinear in the primal space ⇔ their dual lines intersect in the dual.

Checking collinearity of any 3 points:

- Transfer to the dual plane
- Solve the arrangement
- Check for vertices with degree ≥ 3 (⇔ there are 3 collinear points).

Sorting all angular sequences:

For each point p in a set of n points we want to perform an angular sweep.

Naïve:

For each point sort the segment angles $\phi_1 \leq \cdots \leq \phi_{n-1}$. Per vertex the cost is $\Theta(n \lg n)$, and for all vertices: $\Theta(n^2 \lg n)$.

Improvement:

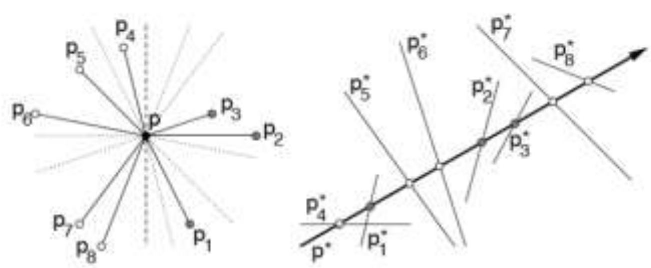

We transfer the problem from the primal space to the dual space in linear time.

The order is maintained, however lines may be discovered $180°$ away – but then we can look at the x value of the point whose line we check – if it is $<$ than p's, it is on the left side, otherwise – on the right. We can then process and extract the correct order of intersection. We then transfer the intersection points back to the lines (angles) in the primal plane. The total is $O(n^2)$.

Note: the sorting here is linear, using the assumption that we can do that when the numbers are bounded – and they are (angles are at most 2π).

Maximum Discrepancy:

Assume S is a set of n points in a unit square $U = [0,1]^2$.

Cross a line l that cuts U in two, and denote the lower half plane as $h(l)$. The Density is defined: $\mu_S(h) = \frac{|h \cap S|}{|S|}$

We can measure the area of $h \cap U$, and define that quadrilateral's area as $\mu(h) = \frac{|Q|}{|U|}$

If $\Delta(h) := |\mu(h) - \mu_S(h)| = 0$, we say this density is a good representation of the area. This is a measurement of how good the distribution S mimics the area.

The maximum discrepancy will tell us how good S is as a distribution representation of the area. If that max is close to 0, S is good. The ideal sampling is that of uniform distribution.

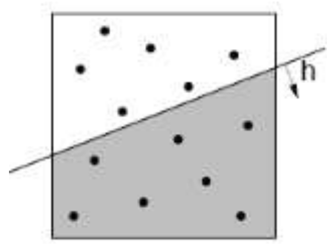

There are infinitely many lines l that give infinitely $h(l)$s. But, the structure gives an algorithm.

Main lemma:

Let h denote the half-plane that gives the maximum discrepancy. The corresponding line l either goes through two points $p_1, p_2 \in S$, or it goes through a single point – a point that is equidistant from two edges of the box U.

So l is either:

i. Passes through 2 points of S

ii. Passes through one point of S that is a midpoint of the line segment $l \cap U$.

Why l can't just go through any points in S?

Assume that is the case; we will show that we can increase/decrease the discrepancy by moving that line up or down.

Case ii:

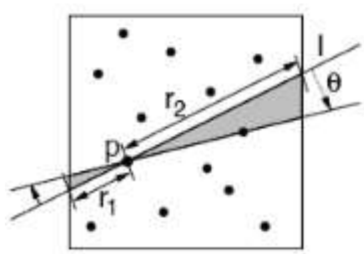

Assume $\mu_S(h) > \mu(h)$, then we can decrease $\mu(h)$ by moving the line down without hitting any points, thus increasing the discrepancy, and then we get a contradiction.

If we moved the line and hit one points $p \in S$, and assume the segments r_1, r_2 hold $r_1 < r_2$. We can approximate the gray areas. Denote the angle θ, then the areas are: $\frac{1}{2}r_1^2\theta, \frac{1}{2}r_2^2\theta$ – and we can choose either of them, but if $r_1 = r_2$ then it won't matter. Slide 40 explains it well.

<u>Case i</u>:

In the primal space there are n points in the plane, we want to find the line l that passes through any 2 points. We need to do those for any 2 points, and count how many points are above those lines, and how many below – from that we can easily calculate Δ.

We map it to the dual space: we get n lines. Each pair of lines here corresponds to a pair of points in the primal space, and each line here corresponds to a point (intersection point) in the primal.

We now only need to check how many lines in the dual are <u>above</u> this point, and this corresponds to how many points in the primal are <u>below</u> the line.

To do that we compute the arrangement in the dual space.