10. Line Arrangements

Lecture on Monday 3rd November, 2008 by Michael Hoffmann <hoffmann@inf.ethz.ch>

During the course of this lecture we encountered several situations where it was convenient to assume that a point set is "in general position". In the plane general position usually amounts to no three points being collinear and/or no four of them being cocircular. This raises an algorithmic question: How can we test for n given points whether or not three of them are collinear? Obviously, we can test all triples in $O(n^3)$ time. Can we do better? In order to answer this question, we will take a detour through the dual plane.

Recall the standard projective duality transform that maps a point $p = (p_x, p_y)$ to the line $p^* : y = p_x x - p_y$ and a non-vertical line $g : y = mx + b$ to the point $g^* = (m, -b)$. This map is ...

- Incidence preserving: $p \in g \iff g^* \in p^*$.
- Order preserving: p is above $g \iff g^*$ is above p^*.

Another way to think of duality is in terms of the parabola $\mathcal{P} : y = \frac{1}{2}x^2$. For a point p on $\mathcal{P}$, the dual line p^* is the tangent to $\mathcal{P}$ at p. For a point p not on $\mathcal{P}$, consider the vertical projection p' of p onto $\mathcal{P}$: the slopes of p^* and p'^* are the same, just p^* is shifted by the difference in y-coordinates.

The question of whether or not three points in the primal plane are collinear transforms to whether or not three lines in the dual plane meet in a point. This question in turn we will answer with the help of *line arrangements* as defined below.

10.1 Arrangements

The subdivision of the plane induced by a finite set L of lines is called the **arrangement** $\mathcal{A}(L)$. A line arrangement is **simple** if no two lines are parallel and no three lines meet in a point. Although lines are unbounded, we can regard a line arrangement a bounded object by (conceptually) putting a sufficiently large box around that contains all vertices. Such a box can be constructed in $O(n \log n)$ time for n lines. Moreover, we can view a line arrangement as a planar graph by adding an additional vertex at "infinity", that is incident to all rays which leave this bounding box. For algorithmic purposes, we will mostly think of an arrangement as being represented by a doubly connected edge list (DCEL), cf. Section 5.3.

Theorem 10.1 *A simple arrangement $\mathcal{A}(L)$ of n lines in $\mathbb{R}^2$ has $\binom{n}{2}$ vertices, n^2 edges, and $\binom{n}{2} + n + 1$ faces/cells.*

Proof. Since all lines intersect and all intersection points are pairwise distinct, there are $\binom{n}{2}$ vertices.

The number of edges we prove by induction on n. For $n = 1$ we have $1^2 = 1$ edge. By adding one line to an arrangement of $n-1$ lines we split $n-1$ existing edges into two and introduce n new edges along the newly inserted line. Thus, there are in total $(n-1)^2 + 2n - 1 = n^2 - 2n + 1 + 2n - 1 = n^2$ edges.

The number f of faces can now be obtained from Euler's formula $v - e + f = 2$, where v and e denote the number of vertices and edges, respectively. However, in order to apply Euler's formula we need to consider $\mathcal{A}(L)$ as a planar graph and take the symbolic "infinite" vertex into account. Therefore,

$$f = 2 - \left(\binom{n}{2} + 1\right) + n^2 = 1 + \frac{1}{2}(2n^2 - n(n-1)) = 1 + \frac{1}{2}(n^2 + n) = 1 + \binom{n}{2} + n\,.$$

□

The *complexity* of an arrangement is simply the total number of vertices, edges, and faces (in general, cells of any dimension).

10.2 Construction

As the complexity of a line arrangement is quadratic, there is no need to look for a subquadratic algorithm to construct it. We will simply construct it incrementally, inserting the lines one by one. Let $\ell_1, \ldots, \ell_n$ be the order of insertion.

At Step i of the construction, locate ℓ_i in the leftmost cell of $\mathcal{A}(\{\ell_1, \ldots, \ell_{i-1}\})$ it intersects. (The halfedges leaving the infinite vertex are ordered by slope.) This takes $O(i)$ time. Then traverse the boundary of the face F found until the halfedge h is found where ℓ_i leaves F. Insert a new vertex at this point, splitting F and h and continue in the same way with the face on the other side of h.

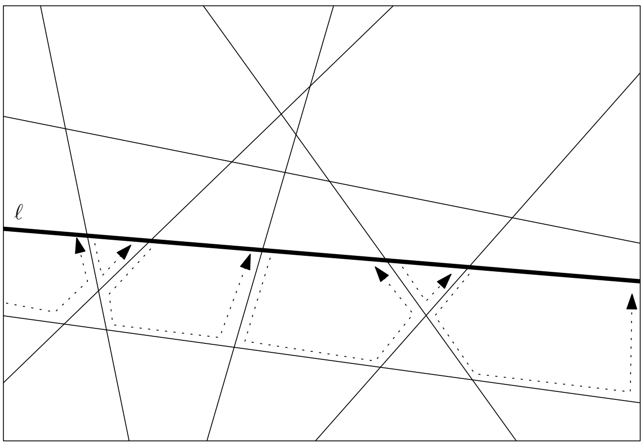

Figure 10.1: *Incremental construction: Insertion of a line ℓ.*

What is the time needed for this traversal? The complexity of $\mathcal{A}(\{\ell_1, \ldots, \ell_{i-1}\})$ is $\Theta(i^2)$, but we will see that the region traversed by a single line has linear complexity only.

10.3 Zone Theorem

For a line ℓ and an arrangement $\mathcal{A}(L)$, the **zone** $Z_{\mathcal{A}(L)}(\ell)$ of ℓ in $\mathcal{A}(L)$ is the set of faces from $\mathcal{A}(L)$ whose closure intersects ℓ.

Theorem 10.2 *Given an arrangement $\mathcal{A}(L)$ of n lines in $\mathbb{R}^2$ and a line ℓ (not necessarily from L), the total number of edges in all cells of the zone $Z_{\mathcal{A}(L)}(\ell)$ is at most $6n$.*

Proof. Without loss of generality suppose that ℓ is horizontal and that none of the lines from L is horizontal. Split the edges of $Z_{\mathcal{A}(L)}(\ell)$ into two groups, those with positive slope are called *left-bounding* and those with negative slope are called right-bounding. (Another way to think of this separation is to split each cell at its topmost and at its bottommost vertex). We will show that there are at most $3n$ left-bounding edges by induction on n.

For $n = 1$, there is exactly one left-bounding edge in $Z_{\mathcal{A}(L)}(\ell)$ and $1 \leq 3n = 3$. Assume the statement is true for $n - 1$.

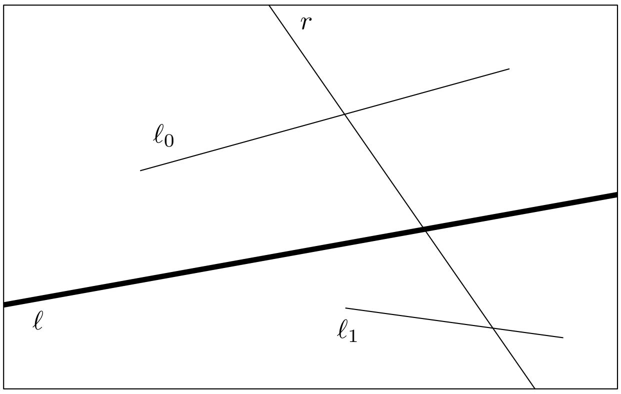

Figure 10.2: *At most three new left-bounding edges are created by adding r to $\mathcal{A}(L \setminus \{r\})$.*

Consider the rightmost line r from L intersecting ℓ and the arrangement $\mathcal{A}(L \setminus \{r\})$. By the induction hypothesis there are at most $3n-3$ left-bounding edges in $Z_{\mathcal{A}(L\setminus\{r\})}(\ell)$. Adding r back adds at most three new left-bounding edges: At most two existing left-bounding edges (call them ℓ_0 and ℓ_1) of the rightmost cell of the zone are intersected by r and thereby split in two, and r itself contributes one more left-bounding edge to that cell. The line r cannot contribute a left-bounding edge to any cell other than the rightmost: to the left of r, the edges induced by r form right-bounding edges only and to the right of r all other cells touched by r (if any) are shielded away from ℓ by one of ℓ_0 or ℓ_1. Therefore, the total number of edges in $Z_{\mathcal{A}(L)}(\ell)$ is bounded from above by $3 + 3n - 3 = 3n$. □

Corollary 10.3 *The arrangement of n lines in $\mathbb{R}^2$ can be constructed in $O(n^2)$ time and this is optimal.*

Corresponding bounds in $\mathbb{R}^d$: Complexity of arrangements in $\Theta(n^d)$, zone of a hyperplane is $O(n^{d-1})$.

10.4 The Power of Duality

The real beauty and power of line arrangements becomes apparent in context of projective point $\leftrightarrow$ line duality. The following problems all can be solved in $O(n^2)$ time and space by constructing the dual arrangement.

General position test. Given n points in $\mathbb{R}^2$, are any three of them collinear? (Dual: do three lines meet in a point?)

Minimum area triangle. Given n points in $\mathbb{R}^2$, what is the minimum area triangle spanned by any three of them? For any vertex of the dual arrangement (primal: line through two points p and q) find the closest point vertically above/below through which an input line passes (primal: closest line above/below parallel to the line through p and q that passes through an input point).

10.5 Ham Sandwich Theorem

Suppose two thieves have stolen a necklace that contains rubies and diamonds. Now it is the time to distribute the prey. Both, of course, should get the same number of rubies and the same number of diamonds. On the other hand, it would be a pity to completely disintegrate the beautiful necklace. Hence they want to use as few cuts as possible to achieve a fair gem distribution.

To phrase the problem in a geometric (and somewhat more general) setting: Given two finite sets R and D of points, construct a line that bisects both sets, that is, in either halfplane defined by the line there are about half of the points from R and about half of the points from D. To solve this problem, we will make use of the concept of levels in arrangements.

Definition 10.4 *For an arrangement $\mathcal{A}(L)$ induced by a set L of n lines in the plane, we say that a point p is on the k-***level** *in $\mathcal{A}(L)$ if and only p lies on some line from L and there are at most $k-1$ lines below and at most $n-k$ lines above p. The 0-level is also referred to as the* **lower envelope**.

Theorem 10.5 *Let $R, D \subset \mathbb{R}^2$ be finite sets of points. Then there exists a line that bisects both R and D. That is, in either open halfplane defined by ℓ there are no more than $|R|/2$ points from R and no more than $|D|/2$ points from D. Moreover, a bisecting line can be found in $O(n \log n)$ time, where $n = |R| + |D|$.*

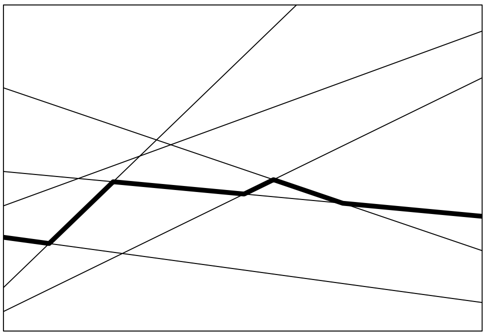

Figure 10.3: *The 2-level of an arrangement.*

Proof. Without loss of generality suppose that both $|R|$ and $|D|$ are odd. (If, say, $|R|$ is even, simply remove an arbitrary point from R. Any bisector for the resulting set is also a bisector for R.) We may also suppose that no two points from $R \cup D$ have the same x-coordinate. (Otherwise, rotate the plane infinitesimally.)

Let R^* and D^* denote the set of lines dual to the points from R and D, respectively. Consider the arrangement $\mathcal{A}(R^*)$. The median level of $\mathcal{A}(R^*)$ defines the bisecting lines for R. As $|R| = |R^*|$ is odd, both the leftmost and the rightmost segment of this level are defined by the same line ℓ_r from R^*, the one with median slope. Similarly there is a corresponding line ℓ_d in $\mathcal{A}(D^*)$.

Since no two points from $R \cup D$ have the same x-coordinate, no two lines from $R^* \cup D^*$ have the same slope, and thus ℓ_r and ℓ_d intersect. Consequently, the median level of $\mathcal{A}(R^*)$ and the median level of $\mathcal{A}(D^*)$ intersect an odd number of times, in particular, they intersect. Any point that lies on both median levels correspond to a primal line that bisects both point sets simultaneously.

A point of intersection can be found by sorting the intersections of ℓ_r with all lines from R^* by x-coordinate and processing them from left to right. Initially, there are some number k of lines from D^* above ℓ_r (k can be computed in linear time). At each intersection, k either increases by one or decreases by one, and at some point, k must be $(|D| - 1)/2$. □

How can the thieves use Theorem 10.5? If they are smart, they drape the necklace along some convex curve, say, a circle. Then using Theorem 10.5 they construct a line ℓ which simultaneously bisects the set of diamonds and the set of rubies. As any line intersects the circle at most twice, the necklace is cut at most twice.

You can also think of the two point sets as a discrete distribution of a ham sandwich that is to be cut fairly, that is, in such a way that both parts have the same amount of ham and the same amount of bread. That is where the name “ham sandwich cut” comes from. The theorem also holds in $\mathbb{R}^d$, saying that any d finite point sets (or finite Borel measures, if you want) can simultaneously be bisected by a hyperplane. This implies that the thieves can fairly distribute a necklace consisting of d types of gems using at most d cuts.