CMSC 754: Lecture 13
Line Arrangements and Applications

Reading: Chapter 8 in the 4M's.

Line Arrangements: We have studied a number of the most fundamental structures in computational geometry: convex hulls, Voronoi diagrams and Delaunay triangulations. These are all defined over a finite set of points. As we saw earlier, points and lines in the plane are related to each other through the *dual transformation*. In this lecture, we will study a fundamental structure defined for a finite set of lines, called a *line arrangement*.

Consider a finite set $L = \{\ell_1, \ldots, \ell_n\}$ of lines in the plane. These lines naturally subdivide the plane into a cell complex, which is called the *arrangement* of L, and is denoted $\mathcal{A}(L)$ (see Fig. 1(a)). The intersection points form the vertices of the complex, the segments between consecutive intersection points form its edges, and the polygonal regions between the lines form the faces, also called *cells*. Although an arrangement contains unbounded edges and faces, as we did with Voronoi diagrams (from a purely topological perspective), it is possible to add a vertex at infinity and attach all these edges to this vertex to form a proper planar graph (see Fig. 1(b)). An arrangement can be represented using any standard data structure for cell complexes, such as a DCEL.

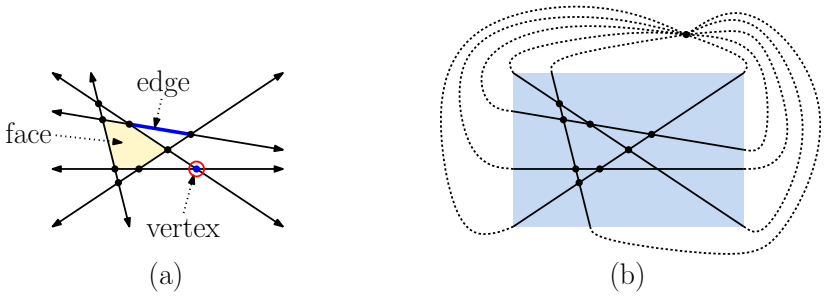


Fig. 1: Arrangement of lines; (a) the basic elements of an arrangement and (b) adding a vertex at infinity to form a proper planar graph.

As we shall see, arrangements have many applications in computational geometry. Using point-line duality, many of these applications involve sets of points. We will begin by discussing the basic geometric and combinatorial properties of arrangements and an algorithm for constructing them. Later we will discuss applications of arrangements to other problems in computational geometry. Although we will not discuss it, line arrangements in $\mathbb{R}^2$ can be generalized to hyperplane arrangements in $\mathbb{R}^d$. In such a case, the arrangement is a polyhedral cell complex.

Combinatorial Properties: The *combinatorial complexity* of an arrangement is the total number of vertices, edges, and faces in the arrangement. An arrangement is said to be *simple* if no three lines intersect at a common point. Through our usual general position assumption that no three lines intersect at a single point, it follows that we will be interested only in simple arrangements. We will also assume that no two lines are parallel. The following lemma shows that these quantities are all $\Theta(n^2)$ for simple planar line arrangements.

Lemma: Let $\mathcal{A}(L)$ be a simple arrangement of n lines L in the plan. Then:

(i) the number of vertices (not counting the vertex at infinity) in $\mathcal{A}(L)$ is $\binom{n}{2} = \frac{1}{2}(n^2-n)$.

(ii) the number of edges in $\mathcal{A}(L)$ is n^2

(iii) the number of faces in $\mathcal{A}(L)$ is $\binom{n}{2} + n + 1 = \frac{1}{2}(n^2+n+2)$.

Proof: The fact that the number of vertices is $\binom{n}{2}$ is clear from the fact that (since no two are parallel) each pair of lines intersects in a single point.

The number of edges follows from the fact that each line contains n edges. Each line is cut by the other $n-1$ lines (assuming no two are parallel), splitting it into n edges.

The number of faces follows from Euler's formula, $v-e+f=2$. To form a cell complex, recall that we added an additional vertex at infinity. Thus, we have $v = 1 + \binom{n}{2}$ and $e = n^2$. Therefore, the number of faces is

$$\begin{aligned} f &= 2 - v + e = 2 - \left(1 + \binom{n}{2}\right) + n^2 = 2 - \left(1 + \frac{n(n-1)}{2}\right) + n^2 \\ &= 1 + \frac{n^2}{2} + \frac{n}{2} = 1 + \frac{n(n-1)}{2} + n = \binom{n}{2} + n + 1, \end{aligned}$$

as desired.

Note that all these quantities are $\Theta(n^2)$. This generalizes to higher dimensions. The combinatorial complexity of an arrangement of n hyperplanes in $\mathbb{R}^d$ is $\Theta(n^d)$. Thus, these structures are practical only in relatively low dimensions when n is not too large.

Incremental Construction: Arrangements are used for solving many problems in computational geometry. But in order to use an arrangement, we first must be able to construct it.[1] We will present a simple incremental algorithm, which builds an arrangement by adding lines one at a time. Unlike the other incremental algorithms we have seen so far, this one is *not randomized*. Its worst-case asymptotic running time is $O(n^2)$, regardless of the insertion order. This is asymptotically optimal, since the arrangement has size $O(n^2)$. The algorithm also requires $O(n^2)$ space, since this is the storage needed for the final result.

Let $L = \{\ell_1, \ldots, \ell_n\}$ denote the set of lines. We will add lines one by one and update the arrangement after each insertion. We show that we can insert the ith line in $O(i)$ time (regardless of insertion order). Summing over i, this yields a total running time proportional to $\sum_{i=1}^{n} i = O(n^2)$.

Suppose that the first $i-1$ lines have already been inserted. Consider inserting ℓ_i. We start by determining the leftmost (unbounded) face of the arrangement that contains this line. At $x = \infty$, the lines are sorted from top to bottom by increasing slope (irrespective of their y-intercepts). In $O(i)$ time, we can determine where the slope of ℓ_i falls relative to the slopes of the prior $i-1$ lines, which determines the leftmost face of the arrangement containing this line.[2]

[1] This is not quite accurate. For some applications, a plane sweep of the arrangement suffices. There is a plane-sweep algorithm specially tailored for line arrangements, called topological plane sweep.

[2] In fact, we could do this in $O(\log i)$ time by storing the slopes in an ordered dictionary, but this would not improve our overall running time.

The newly inserted line cuts through a sequence of $i-1$ edges and i faces of the existing arrangement. To process the insertion, we need to determine which edges ℓ_i cuts, then split each such edge and update the DCEL accordingly.

In order to determine which edges are cut by ℓ_i, we "walk" this line through the current arrangement, from one face to the next. When we enter a face, we need to determine which edge ℓ_i exits through. We answer this question with a simple strategy. We walk along the edges of the face, say counterclockwise, until we find the exit edge, that is, the next edge ℓ_i intersects. We then jump to the face on the other side of this edge and continue the trace with the neighboring face. This is illustrated in Fig. 2(a). The DCEL data structure supports such local traversals in time linear in the number of edges traversed. (You might wonder why we don't generalize the trapezoidal map algorithm. We could build a trapezoidal map of the arrangement and walk the new segment through a sequence of trapezoids. It turns out that this would be just as efficient.)

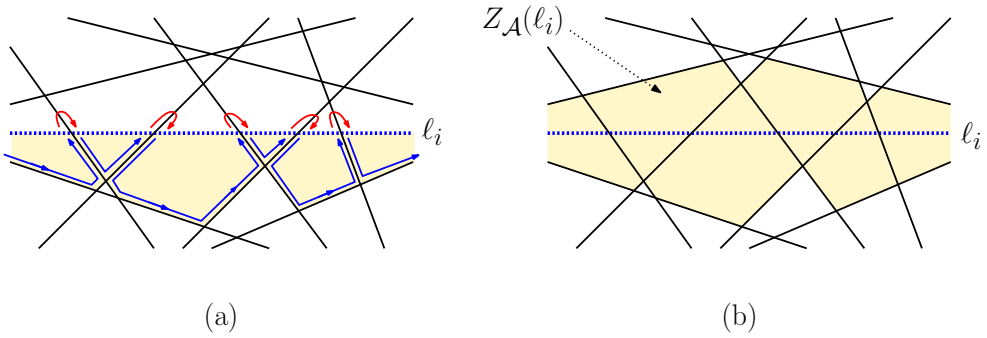


Fig. 2: Adding the line ℓ_i to the arrangement; (a) traversing the arrangement and (b) the zone of a line ℓ_i. (Note that only a portion of the zone is shown in the figure.)

Clearly, the time that it takes to perform the insertion is proportional to the total number of edges that have been traversed in this tracing process. A naive argument says that we encounter $i-1$ lines, and hence pass through i faces (assuming general position). Since each face is bounded by at most i lines, each facial traversal takes $O(i)$ time, giving a total of $O(i^2)$, which is much higher than the $O(i)$ time we promised earlier. Why is this wrong? It is based on a bound on the total complexity of the faces traversed. To improve this, we need to delve more deeply into an important combinatorial structure related to arrangements.

Zones and the Zone Theorem: Given an arrangement $\mathcal{A}$ of a set L of n lines, and given a line ℓ that is not in L, the *zone* of ℓ in $\mathcal{A}(L)$, denoted $Z_{\mathcal{A}}(\ell)$, is the set of faces of the arrangement that are intersected by ℓ (shaded in Fig. 2(b)). The running time of the incremental construction clearly grows linearly with the total number of edges of $Z_{\mathcal{A}}(\ell)$. (Actually, we only really care about the edges that lie below ℓ, but this won't affect the asymptotics.)

As argued above, the combinatorial complexity of a zone is clearly at most $O(n^2)$. However, we will show next that (remarkably!) the total complexity of any zone is at most $O(n)$.

Zone Theorem: Given an arrangement $\mathcal{A}(L)$ of n lines in the plane, and given any line ℓ in the plane not in L, the total number of edges in all the cells of the zone $Z_{\mathcal{A}}(\ell)$ is at most $6n$.

As with many combinatorial proofs, the key is to organize matter so that the counting can be done in an easy way. This is not trivial. We cannot count cell-by-cell, since some cells have high complexity and some low. We also cannot count line-by-line, because some lines contribute many edges to the zone and others just a few. The key in the proof is finding a (clever!) way to add up the edges so that each line appears to induce only a constant number of edges into the zone. (Note that our text counts zone edges a bit differently.)

Proof: The proof is based on a simple inductive argument. For the sake of illustration, let us rotate the plane so that ℓ is horizontal. By general position, we may assume that none of the lines of L are parallel to ℓ. We split the zone's edges into two groups: those that bound some face from the left and those that bound some face from the right. An edge of a face is said to be *left bounding* if the face lies in the right half-plane of the line defining this edge, and a face is *right bounding* if the face lies in the left half-plane of the line defining this edge (see Fig. 3). We will show that there are at most $3n$ left-bounding edges in the zone, and by a symmetrical argument there are at most $3n$ right-bounding edges. This will yield the desired total of $6n$ edges.

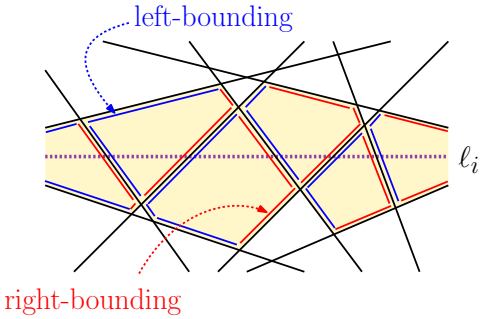


Fig. 3: Left- and right-bounding edges in the zone.

The proof is by induction on n. For the base case $n = 1$, ℓ has exactly one left-bounding edge in its zone, and $1 < 3 = 3n$. For the induction step, assume the induction hypothesis holds for any set of $n-1$ lines, and show it holds for an arrangement of n lines. Consider the rightmost line of the arrangement, which intersects ℓ. Call this ℓ_1 (see Fig. 4(b)). By the induction hypothesis, the zone of the remaining $n-1$ lines has at most $3(n-1)$ left-bounding edges.

Now, let us add ℓ_1 and see how many more left-bounding edges are generated. Because ℓ_1 is leftmost, it intersects the zone's rightmost face. All edges of this face are left-bounding edges. By convexity, ℓ_1 intersects the boundary of this face in two edges, denoted e_a and e_b, where e_a is above ℓ, and e_b is below (see Fig. 4(c)). Its insertion creates a new left-bounding edge along ℓ_1 between e_a and e_b, and splits each of e_a and e_b into two new left-bounding edges. Thus, there is a net increase of three edges, for a total of $3(n-1)+3 = 3n$ edges.

We assert that ℓ_1 cannot contribute any other left-bounding edges to the zone. This is because the lines containing e_a and e_b block this possibility. Thus, there are at most $3n$ left-bounding edges, as desired.

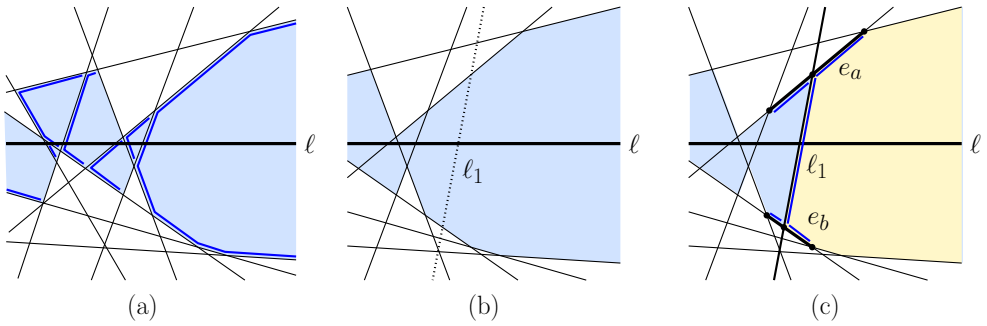


Fig. 4: Proof of the Zone Theorem.

Applications of Arrangements and Duality: When combined with point-line duality, line arrangements make it possible to solve a number of geometric computational problems. A number of examples are given below. Unless otherwise stated, we can solve all these problems in $O(n^2)$ time and $O(n^2)$ space by constructing a line arrangement. Alternately, they can be solved in $O(n^2 \log n)$ time and $O(n)$ space by applying plane sweep to the arrangement.

General-position test: Given a set of n points in the plane, determine whether any three are collinear.

Minimum area triangle: Given a set of n points in the plane, find the minimum-area triangle whose vertices are selected from these points.

Minimum k-corridor: Given a set of n points, and an integer k, determine the narrowest pair of parallel lines that enclose at least k points of the set. The distance between the lines can be defined either as the vertical distance between the lines or the perpendicular distance between the lines (see Fig. 5(a)).

Visibility graph: Given line segments in the plane, we say that two points are *visible* if the interior of the line segment joining them intersects none of the segments. Given a set of n non-intersecting line segments, compute the *visibility graph*, whose vertices are the endpoints of the segments, and whose edges are pairs of visible endpoints (see Fig. 5(b)).

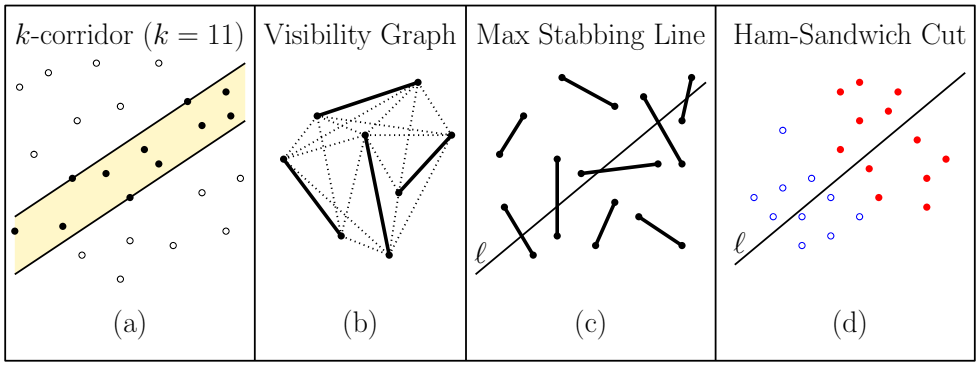


Fig. 5: Applications of arrangements.

Maximum stabbing line: Given a set of n line segments in the plane, compute the line ℓ that stabs (intersects) the maximum number of these line segments (see Fig. 5(c)).

Ham Sandwich Cut: Given n red points and m blue points, find a single line ℓ that simultaneously bisects these point sets. It is a famous fact from mathematics, called the *Ham-Sandwich Theorem*, that such a line always exists. If the two point sets are separable by a line (that is, the red convex hull and the blue convex hull do not intersect), then this can be solved in time $O(n+m)$ (see Fig. 5(d)).

In the remainder of the lecture, we'll see how problems like these can be solved using arrangements.

Sweeping Arrangements: Since an arrangement of n lines is of size $\Theta(n^2)$, we cannot expect to solve problems through the explicit use of arrangements in less than quadratic time. Most applications involve first constructing the arrangement, and then traversing it in some manner. In many instances, the most natural traversal is based on a plane sweep. (This is not the only way, however. Since a planar arrangement is a graph, methods such as depth-first and breadth-first search can be used.)

If an arrangement is to be built just so it can be swept, then maybe you don't need to construct the arrangement at all. You can just perform the plane sweep on the lines, exactly as we did for the line segment intersection algorithm. Assuming we sweep from left to right, the sweep line starts at $x = -\infty$ (which means sorting by slope). The sweep line status maintains the lines in, say, bottom-to-top order according to their intersection with the sweep line. The events are the vertices of the arrangement.

Note that the sweep-line status always contains exactly n entries. Whenever an intersection event occurs, we can update the sweep-line status by swapping two adjacent entries. Thus, instead of an ordered dictionary, it suffices to store the lines in a simple n-element array, sorted, say, from top to bottom. This means the **sweep-line updates can be performed in $O(1)$ time**, rather than $O(\log n)$(see Fig. 6(a)). We still need to maintain the priority queue, and each operation takes $O(\log n)$ time.

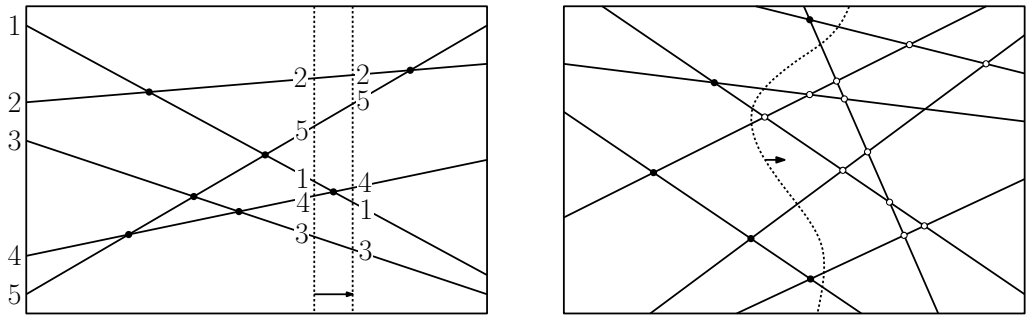


Fig. 6: Sweeping a line arrangement.

Sweeping an arrangement takes $O(n^2 \log n)$ time and $O(n)$ space. Because it is more space-efficient, this is often an attractive alternative to constructing the entire arrangement.

Topological Plane Sweep: (Optional) As mentioned above, the priority queue is the slowest part of plane sweeping an arrangement. Remarkably, there is a way to save this $O(\log n)$ factor, but we must abandon hope of sweeping events in purely left-to-right order. There is a somewhat more "relaxed" version of plane sweep, which works for line arrangements in the plane. The method is called *topological plane sweep*. It runs in $O(n^2)$ time (thus, eliminating an $O(\log n)$ factor from the running time) and uses $O(n)$ space.

It achieves efficiency by relaxing the requirement that vertices be swept in strict left-to-right order. Rather, it uses a more "local" approach for deciding which vertex of the arrangement to sweep next.[3] This local approach ensures that vertices along each line are swept in the correct order, even though vertices on different lines are not necessarily swept in left-to-right order. Intuitively, we can think of the sweep line as a sort of *pseudoline* that intersects each line of the arrangement exactly once (see Fig. 6(b)). Although we will not justify it, the method applies to all the problems we discuss in today's lecture. *Note that topological plane sweep applies only to infinite lines, not line segments.*

Duality: Many of our applications will involve the dual transformation, which we introduced earlier in the semester. Recall that the dual of a point $p = (a, b)$ is the line $p^* : y = ax - b$, and the dual of a line $\ell : y = ax - b$ is the point $\ell^* = (a, b)$. Also recall the *order-reversing property* that the point p lies above line ℓ (at vertical distance h) if and only if the dual line p^* lies below dual point ℓ^* (also at vertical distance h).

Narrowest k-corridor: Given a set P of n points in the plane and an integer k, $1 \leq k \leq n$, we want to find the narrowest pair of parallel lines that enclose at least k points of P. (We call this a *slab* or *corridor*.) We define the *width* of the corridor to be the vertical distance between these. Our objective is to compute the minimum-width corridor that encloses k points (which may lie on the corridor's boundary). We can adapt the algorithm to minimize the perpendicular distance between the lines. We will make the usual general-position assumptions that no three points of P are collinear, and no two points have the same x-coordinate.

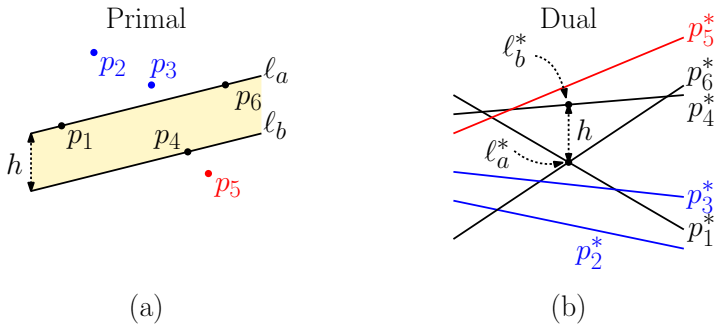
Consider any corridor defined by parallel lines ℓ_a (above) and ℓ_b (below) (see Fig. 7(a)). Since the lines are parallel, these points share the same a-coordinate, which implies that the line segment $\overline{\ell_a^* \ell_b^*}$ is vertical (see Fig. 7(b)). The vertical distance between the lines (i.e., the difference in their y-intercepts) equals the vertical distance between the dual points.

By the order-reversing property, points that lie above/below/within the corridor (shown in blue, red, and black in Fig. 7(a), respectively) are mapped to dual lines that pass below/above/through this segment (see Fig. 7(b)). Thus, we have the following equivalent dual formulation of this problem:

Shortest vertical k-stabber: Given an arrangement of n lines, determine the shortest vertical segment that stabs (intersects) k lines of the arrangement.

It is easy to show that the shortest vertical k-stabber may be assumed to have one of its endpoints on a vertex of the arrangement. (If the vertical line has endpoints in the interior of two edges, we can slide it left or right and decrease its length.)

[3] For details, see "Topologically sweeping an arrangement" by H. Edelsbrunner and L. J. Guibas, *J. Comput. Syst. Sci.*, 38 (1989), 165–194.

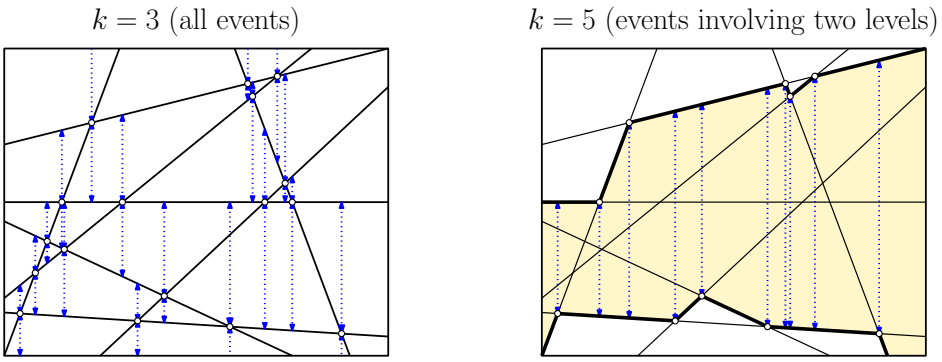


(a) (b)

Fig. 7: A 3-corridor in primal and dual forms. (Note that the corridor is not as narrow as possible.)

3-stabber: The 3-stabber is the simplest to describe. Perform a simple plane sweep of the arrangement (using a vertical sweep line). Whenever we encounter a vertex of the arrangement, we consider the distance from this vertex to the edge of the arrangement lying immediately above this vertex and the edge lying immediately below. (These are illustrated by the blue broken lines in Fig. 8(a).) We can solve this problem by plane sweep in $O(n^2 \log n)$ time and $O(n)$ space. (By using topological plane sweep, the extra $\log n$ factor in the running time can be removed.)

Note that we can use this to test whether points are in general position. It is easy to prove that a set of n points has three or more collinear points if and only if the dual arrangement's minimum 3-stabber is of length zero.



(a) (b)

Fig. 8: The critical events in computing the shortest vertical 3-stabber (a) and 5-stabber (b).

k-stabber: Whenever we encounter a vertex in the plane sweep, we determine the distance to the lines of the arrangement that lie $k-2$ above and $k-2$ below (see the blue broken lines of Fig.8(b)). The "-2" accounts for two lines that pass through the vertex itself. Since the sweep-line status can be stored in a simple n-element array, we can access these entries in $O(1)$ time for any k.

Halfplane Discrepancy: Next we consider a problem derived from geometric sampling. Suppose that we are given a collection of n points P lying in a unit square $U = [0,1]^2$. We want to use

these points for random sampling purposes. In particular, the property that we would like these points to possess is that for any half-plane h, the fraction of points of P that lie within h should be roughly equal to the area of intersection of h with U. More precisely, define $\mu(h)$ to be the area of $h \cap U$, and $\mu_P(h) = |P \cap h|/|P|$. A sample is good if $\mu(h) \approx \mu_P(h)$, for any choice of h.

To make this more formal, we define the *discrepancy* (or more accurately, the *half-plane discrepancy*) of a finite point set P with respect to a half-plane h to be

$$\Delta_P(h) = |\mu(h) - \mu_P(h)|.$$

For example, in Fig. 9(a), the area of $h \cap U$ is $\mu(h) = 0.625$, and there are 7 out of 13 points in h, thus $\mu_P(h) = 7/13 = 0.538$. Thus, the discrepancy of h is $|0.625 - 0.538| = 0.087$. Define the *half-plane discrepancy* of P to be the maximum (or more properly the supremum, or least upper bound) of this quantity over all half-planes:

$$\Delta(P) = \sup_h \Delta_P(h).$$

Let's consider the problem of computing the discrepancy of a given point set P.

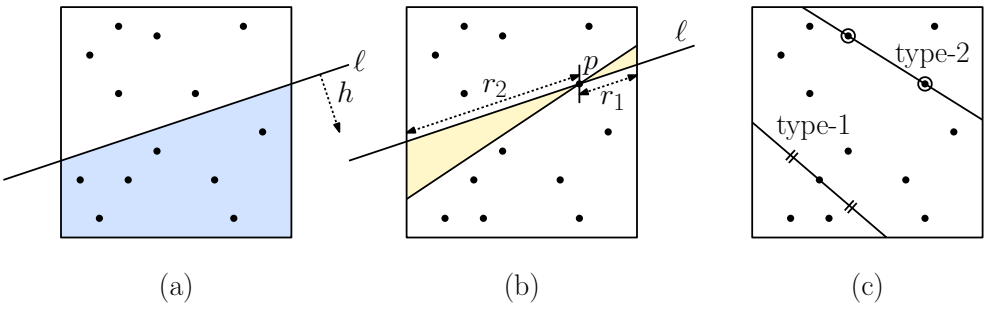


Fig. 9: Discrepancy of a point set.

Since there are an uncountably infinite number of half-planes that intersect the unit square, we should first derive some sort of *finiteness criterion* on the set of half-planes that might produce the greatest discrepancy.

Lemma: Let h denote the half-plane that generates the maximum discrepancy with respect to P, and let ℓ denote the line that bounds h. Then either:

(i) ℓ passes through one point of P, and this point is the midpoint of the line segment $\ell \cap U$, or

(ii) ℓ passes through two points of P.

Remark: If a line passes through one or more points of P, then should this point be included in $\mu_P(h)$? For the purposes of computing the maximum discrepancy, the answer is to either include or omit the point, whichever produces the larger discrepancy. The justification is that it is possible to perturb h infinitesimally so that it includes none or all of these points without altering $\mu(h)$.

Proof: We will show that any line can be moved until it satisfies either (i) or (ii) in such a manner that the discrepancy never decreases. First, if ℓ does not pass through any point of P, then (depending on which is larger $\mu(h)$ or $\mu_P(h)$) we can move the line up or down without changing $\mu_P(h)$ and increasing or decreasing $\mu(h)$ to increase their difference, until it does pass through a point of P. Next, if ℓ passes through a point $p \in P$, but is not the midpoint of the line segment $\ell \cap U$, then we claim that we can rotate this line about p and hence increase or decrease $\mu(h)$ without altering $\mu_P(h)$, to increase their difference.

To establish the claim, consider Fig. 9(b). Suppose that the line ℓ passes through point p and let $r_1 < r_2$ denote the two lengths along ℓ from p to the sides of the square. Observe that if we rotate ℓ through a small angle θ, then to a first-order approximation, the gain due to the area of the triangle on the right is $r_1^2\theta/2$, since this triangle can be approximated by an angular sector of a circle of radius r_1 and angle θ. The loss due to the area of the triangle on the left is $r_2^2\theta/2$. Thus, since $r_1 < r_2$, this rotation will decrease the area of the region lying below h infinitesimally. A rotation in the opposite direction increases the area infinitesimally. Since the number of points bounded by h does not change as a function of θ, the discrepancy cannot be achieved as long as such a rotation is possible.

We say that a line is *type-1* if it satisfies condition (i) and it is *type-2* if it satisfies condition (2) (see Fig. 9(c)). We will show that the discrepancy for each types of lines can be computed in $O(n^2)$ time.

Type-1: For each point $p \in P$, there are only a constant number of lines ℓ (at most two, I believe) through this point such that p is the midpoint of $\ell \cap U$. It follows that there are at most $O(n)$ type-1 lines. We can compute the discrepancy of each such line in $O(n)$ time, which leads to a total running time of $O(n^2)$.

Type-2: Consider a type-2 line ℓ that passes through two points $p_i, p_j \in P$. This line defines two half-planes, one above and one below. We'll explain how to compute the discrepancy of the lower half-plane, h^-, and the upper half-plane, h^+, is symmetric. First, observe that we can compute $\mu(h^-)$ in constant time, so all that remains is computing $\mu_P(h^-)$, that is, the number of points lying on or below ℓ.

If we apply our standard dual transformation, ℓ is mapped in the dual plane to a point ℓ^* at which the dual lines p_i^* and p_j^* intersect. This is just a vertex in the line arrangement $\mathcal{A}(P^*)$. By the order-reversing property of the dual transformation, the points lying on or below ℓ coincide with the dual lines that lie on or above this vertex.

We can compute this quantity in constant time for each vertex of the line arrangement. Recall that the sweep-line status is stored in a simple n-element array, sorted from top to bottom. The vertex in the arrangement corresponds to two consecutive entries in the sweep-line status, say at positions $k-1$ and k. The number of dual lines lying on or above the vertex is therefore just k (assuming that we index the array from 1 to n).

For example, consider the vertex being swept in Fig. 6(a). These intersecting lines are at indices $k-1 = 3$ and $k = 4$ in the sweep-line status, and hence there are 4 lines on or above this vertex, which implies that there are 4 points lying on or below the corresponding type-2 line $\overline{p_1p_4}$ in the primal configuration. Since we can compute

the discrepancy for each type-2 line in $O(1)$ time, the overall time to compute the discrepancies of all type-2 lines is $O(n^2)$.

Levels and k-Sets: Given a point set P, the vertices of its convex hull represent the "extreme" points of the set. We can generalize this concept of extreme objects to sets of points of larger size. Given an n-element point set P in the plane and an integer $1 \leq k \leq n-1$, a k-element subset $P' \subseteq P$ is called a *k-set* of P if there exists a half-plane h such that $P' = P \cap h$ (see Fig. 10(a)). Each vertex of the convex hull is a 1-set, and the vertices of each edge are a 2-set. Given P and k, we would like to compute all the k-sets of a point set.

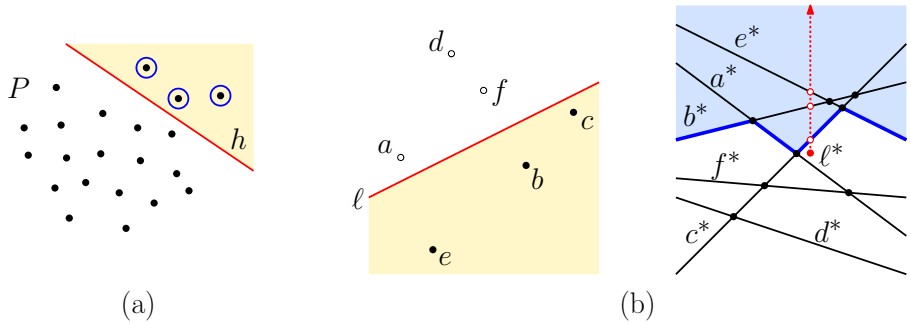


Fig. 10: (a) A 3-set and (b) the 3-set $\{b, c, e\}$ in the primal and dual.

How many k-sets does a given n-element set have? Clearly, the number cannot exceed $\binom{n}{k} = O(n^k)$, but given the half-plane constraint, the number is generally much smaller. We will see below that the number cannot exceed $O(n^2)$. However, obtaining a good bound is a major open problem. Currently, the best upper bound (due to Tamal Dey) is $O(nk^{1/3})$ and the best lower bound (due to Gabor Tóth) is $\Omega\big(n^{c\sqrt{\log k}}\big)$.

We will show that through the use of duality, we can enumerate all the k-sets in $O(n^2)$ time. (Outputting all the sets explicitly requires an additional factor of k, but our algorithm will output a concise representation, whose total size is proportional to the number of k-sets.) Let's discuss lower halfspaces, since upper halfspaces are symmetrical. By the order-reversing property of duality, any lower half-plane bounded by a line ℓ that contains k points of P corresponds to a dual point ℓ^* such that k lines of the dual arrangement $\mathcal{A}(P^*)$ lie on or above this point (see Fig. 10(b)).

Thus, in order to compute all the lower k-sets, it suffices to identify all the places in the dual arrangement where there are k-lines on or above. To make this formal, given any arrangement $\mathcal{A}$, for $1 \leq k \leq n$, we define *level* k of $\mathcal{A}$, denoted $\mathcal{L}_k(\mathcal{A})$ or simply $\mathcal{L}_k$, to be the set of edges of AA such that there are k lines on or above this edge (see Fig. 11).

Observe that $\mathcal{L}_1$ and $\mathcal{L}_n$ are the upper and lower envelopes of the set of lines, respectively. Generally, $\mathcal{L}_k$ is an x-monotone polygonal curve. While the upper and lower envelopes can have no more than n edges, the other levels can have more, since the same line may appear multiple times on a given level. Also observe that two levels meet at each vertex in a "knocked-kneed" manner, bouncing off each other. By our previous observations, the lower k-sets are naturally associated with $\mathcal{L}_k$, and the upper k-sets are associated with the kth level from the bottom, that is $\mathcal{L}_{n-k+1}$.

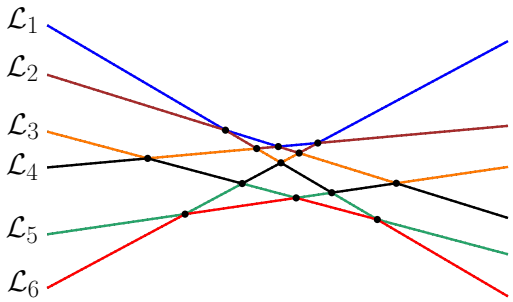


Fig. 11: Levels in an arrangement.

It is an easy matter to modify plane sweep to classify all the edges of the arrangement with their associated level numbers. Observe that for the sake of computing k-sets, whenever we sweep a vertex of the arrangement in the dual setting, this corresponds in the primal to a single point entering the k-set and a single point leaving the k-set. Thus, we can succinctly encode all the k-sets by computing the initial k-set at $-\infty$ (corresponding to the k-set to the right of a vertical line), and for each subsequent k-set represent it by the pair of points being swapped. Therefore, we can report all the k-sets in $O(n^2)$ time and $O(n)$ space.

Sorting all Angular Sequences: Earlier, we introduced the problem of computing visibility graphs. We will not explicitly discuss the solution of that problem here, but we will discuss a fundamental subroutine in this algorithm. Consider a set of n points in the plane. For each point p in this set, we want to sort the remaining $n-1$ points in cyclic order. Clearly, we can compute the cyclically sorted order about any given point in $O(n \log n)$ time, and this leads to an overall running time of $O(n^2 \log n)$. We will show that we can do this in $O(n^2)$ time. (This is rather surprising. Lower bounds on sorting imply that we cannot sort n sets of $n-1$ numbers faster than $\Omega(n^2 \log n)$ time. But here we are exploiting the special structure of the cyclically ordered point sets.)

Here is how we do it. Suppose that p is the point around which we want to sort, and let $\langle p_1, \ldots, p_n \rangle$ be the points in final angular order about p (see Fig. 12(a)). Consider the arrangement defined by the dual lines p_i^*. How does this order manifest itself in the arrangement?

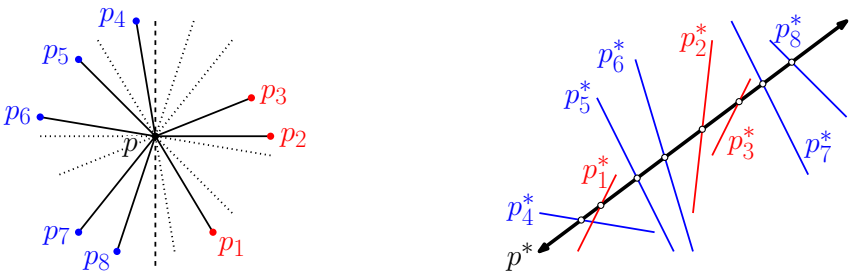


Fig. 12: Arrangements and angular sequences.

Consider the dual line p^*, and its intersection points with each of the dual lines p_i^*. These

form a sequence of vertices in the arrangement along p^*. Consider this sequence ordered from left to right. It would be nice if this order were the desired circular order, but this is not quite correct. It follows from the definition of our dual transformation that the a-coordinate of each of these vertices in the dual arrangement is the slope of some line of the form $\overline{pp_i}$ in the primal plane. Thus, the sequence in which the vertices appear on the line is a *slope ordering* of the points about p_i, which is not quite the same as the *angular ordering*.

However, given this slope ordering, we can simply test which primal points lie to the left of p (shown in blue in Fig. 12(a)), and separate them from the points that lie to the right of p (shown in red in Fig. 12(a)). We partition the vertices into two sorted sequences, and then concatenate these two sequences, with the points on the right side first, and the points on the left side later. For example, in Fig. 12, we partition the slope-sorted sequence $\langle 4, 1, 5, 6, 2, 3, 7, 8\rangle$ into the left sequence $\langle 4, 5, 6, 7, 8\rangle$ and the right sequence $\langle 1, 2, 3\rangle$, and then we concatenate them to obtain the final angle-sorted sequence $\langle 1, 2, 3, 4, 5, 6, 7, 8\rangle$.

Thus, once the arrangement has been constructed, we can reconstruct each of the angular orderings in $O(n)$ time, for a total of $O(n^2)$ time. (Topological plane sweep can also be used, but since the output size is $\Omega(n^2)$, there no real benefit to be achieved by using topological plane sweep.)