Solutions to Homework 5

Solution 1: Given the point set P, recall that its *Euclidean minimum spanning tree*, denoted EMST(P), is the minimum spanning tree of the complete graph on P where the weight of an edge between two points is the distance between them. Let $\langle \ell_1, \ldots, \ell_{n-1} \rangle$ denote the lengths of the edges of the associated Euclidean minimum spanning tree of P sorted from smallest to largest. We will show below that the ith critical radius r_i equals $\ell_i/2$. Assuming this for now, it follows that we can compute the sequence $\langle r_1, \ldots, r_{n-1} \rangle$ in $O(n \log n)$ time as follows. Recall that EMST(P) is a subgraph of the Delaunay triangulation, DT(P) (proved in class). We first compute DT(P), and then compute the MST of the Delaunay triangulation by Kruskal's algorithm. We then sort the edge lengths and output the resulting sequence. Each of these steps can be performed in $O(n \log n)$ time.

To establish correctness, we will show that $r_i = \ell_i/2$. We first show that $r_i \leq \ell_i/2$. Let T_i denote the forest consisting of the i smallest edges of EMST(P). Clearly, T_i has $n - i$ connected components. Since every edge of T_i has length at most ℓ_i, balls of radius r_i centered at each edge endpoint cover every edge of T_i. Thus, the union of these balls covers each component of T_i, implying that the union has at least i connected components. This implies that $r_i \leq \ell_i/2$, as desired (see Fig. 1(a)).

Conversely, to show that $r_i \geq \ell_i/2$, suppose we run Kruskal's MST algorithm but restrict ourselves to edges strictly shorter than ℓ_i. Because Kruskal's algorithm adds edges in increasing order of length and each new edge reduces the number of connected components by exactly one, the algorithm will have $n-i$ connected components. By considering the subtrees of the partial spanning tree, it follows that it is possible to partition the point set into disjoint subsets $P_1, \ldots, P_{n-i} \subset P$, such that for any points p_i and p_j in different subsets, $\|p_i - p_j\| > \ell_i$ (see Fig. 1(b)). This implies that if we place disks of radius $\ell_i/2$ at all points in P, no ball from one group overlaps any ball from another group, so the balls are not connected.

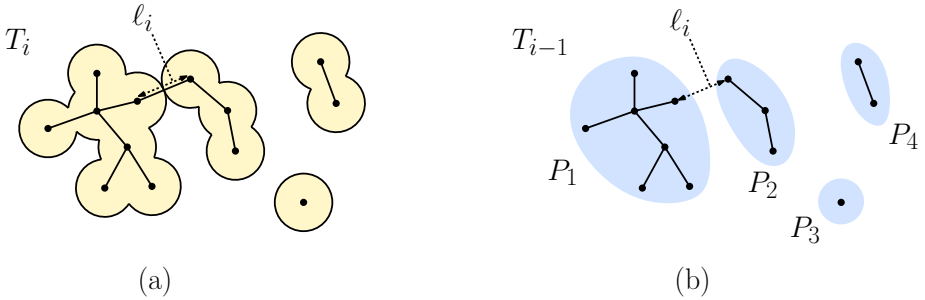


Figure 1: Correctness of the critical radius algorithm.

Solution 2:

(a) We prove that EMST(P) $\subseteq$ RNG(P) by contradiction. Let T = EMST(P), and let $w(T)$ denote the total weight of T. Let a and b be any two sites such that ab is an edge of T.

Suppose, to the contrary, that ab is not an edge of RNG(P). By definition of the RNG, there exists a site $c \in P$ such that $\max(|c-a|, |c-b|) < |a-b|$. Equivalently,

$$\max(|c-a|, |c-b|) - |a-b| \ < \ 0. \tag{1}$$

The removal of $\overline{ab}$ from the EMST splits the tree into two subtrees. Assume without loss of generality that c lies in the same subtree as a. Now, remove the edge $\overline{ab}$ from the EMST and replace it with the edge $\overline{bc}$ (see Fig. 2(a)). The result will be a spanning tree T' whose weight is

$$w(T') \ = \ w(T) + |c-b| - |a-b| \ \leq \ w(T) + \max(|c-a|, |c-b|) - |a-b|.$$

By Eq. (1), we have $w(T') < w(T)$, which contradicts the hypothesis that T is the EMST, completing the proof.

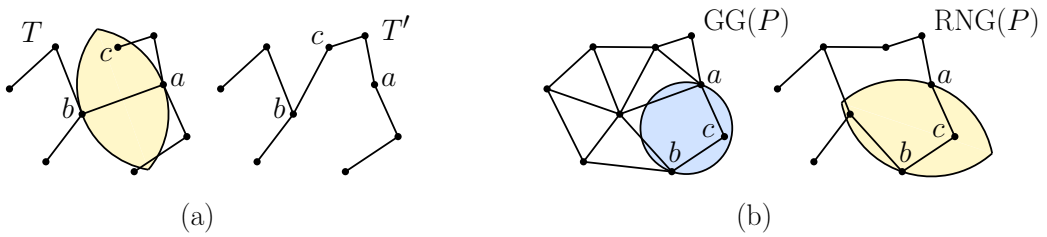


Figure 2: Neighborhood graph inclusions.

(b) To show that RNG(P) $\subseteq$ GG(P), it suffices to show that $ab \notin$ GG(P) implies that $ab \notin$ RNG(P). If $ab \notin$ GG(P), then there exists a site c that lies in the disk whose diameter is ab (see Fig. 2(b)). Since a and b are the farthest pair of points in the disk, we have $\max(|c-a|, |c-b|) < |a-b|$, which implies that $ab \notin$ GG(P), as desired.

(c) To prove that GG(P) $\subseteq$ DT(P), suppose that $ab \in$ GG(P). By definition of the Gabriel graph, the diametrical disk through a and b contains no other site. By the empty circle property, $ab \in$ DT(P), as desired.

Solution 3:

(a) Recall that we consider only left bounding edges. We claim that at most two tokens can be assigned to any intersection point of an arrangement line and ℓ, one from Rule 1 and one from Rule 2. Clearly, Rule 1 can assign at most one token, since no line can intersect ℓ twice.

To see that this is true for Rule 2, suppose to the contrary that two tokens were assigned to the same intersection point. Let ℓ'' be the intersecting line, and suppose that the tokens arrived from ℓ'_1 and ℓ'_2. Clearly, both of these lines must intersect ℓ'' above ℓ. Let p_1 and p_2 denote these intersection points (see Fig. 3(a)). We may assume, without loss of generality, that p_1 lies above p_2.

By hypothesis, ℓ'_1 is a left-bounding edge on the zone above ℓ. Thus, there exists a vertical line segment s joining a point on ℓ'_1 to ℓ, which is not intersected by any other line. In particular, it must lie entirely to the right of both ℓ'' and ℓ'_2. Consider the four sided figure, called Q,

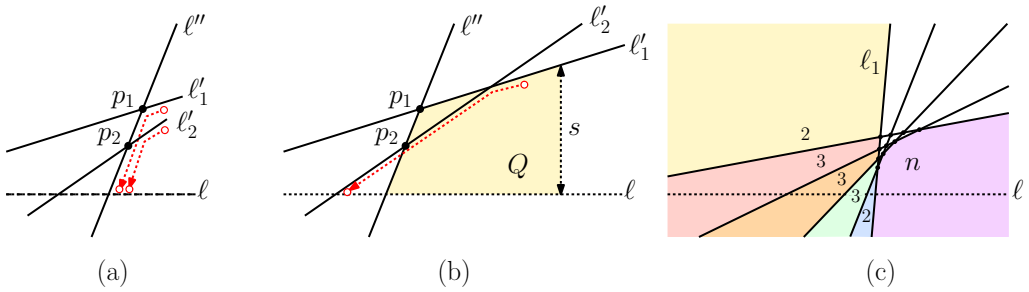


Figure 3: Bead-counting and the zone theorem.

bounded by s on the right, ℓ below, ℓ'' on the left, and ℓ_1' above. (This is shaded in Fig. 3(b).) Since p_2 lies below p_1, it follows that ℓ_2 must intersect Q. Since ℓ_2' is a left-bounding edge and it does not intersect s, it must intersect the upper edge of Q. However, if so, it would follow that the token associated with the edge of ℓ_1 on the zone would have been assigned to the intersection of ℓ_2' with ℓ (or at least, it would not be assigned to the intersection with ℓ'' as hypothesized). Thus, we have the desired contradiction.

(b) Each of the n intersection points of ℓ with the lines of the arrangement can be assigned at most two tokens for left-bounding edges lying in the portion of the zone above ℓ. Adding to this the right-bounding edges and the portion of the zone below ℓ, we have a total of at most 8 tokens for each of the n intersection points, for a total of at most $8n$ tokens, implying that there are at most $8n$ edges on ℓ's zone.

(c) Note that if a line ℓ' intersects ℓ, it generates only one edge along the zone (one as a left-bounding edge and one as a right-bounding edge), but this edge is counted by four tokens, two lying above ℓ and two lying below ℓ. If we remove the lower token count from the argument, we have four tokens for the left and right bounding edges above ℓ and two tokens for the left and right bounding edges below ℓ, for a total of $6n$ tokens.

(d) Consider a circle centered on ℓ and create a collection of n lines that are tangent to this circle along its upper left side. (See Fig 3(c).) Observe that the rightmost face of the zone (containing c) has n edges. The leftmost face has two edges. The face immediately adjacent to the rightmost face has two edges. All the other $n-2$ faces are bounded by three edges: one on the left, one on the right, and the nearly vertical line ℓ_1. Thus, (summing from left to right) the total number of edges in the zone above ℓ is $2+3(n-2)+2+n = 4n-2$.

It is known that there is a tight bound of $5.5n$ on the total complexity of the zone. (See *this paper* by Rom Pinchasi. His count differs by $+4n$ because in his definition of the zone, he include ℓ and the edges cut by ℓ separately.)

Solution 4: In this problem, we will use two related facts from basic Euclidean geometry. First, given two adjacent triangles, $\triangle pqr$ and $\triangle pqs$, the sum of angles $\angle prq + \angle psq$ is greater/less/equal to π if and only if s lies inside/outside/on (respectively) the circumcircle of $\triangle pqr$ (see Fig. 4(a)). Second, a triangle $\triangle pqr$ is acute if and only if the center of its circumcircle lies within the triangle's interior (see Fig. 4(b)).

(a) We will prove the contrapositive, namely, if a triangulation is not Delaunay, then it is not acute. Let T be any triangulation that is not Delaunay. By Delaunay's Theorem, any triangulation that is locally Delaunay is also globally Delaunay, and therefore T is not locally Delaunay. This means that there exist two adjacent triangles, $\triangle pqr$ and $\triangle pqs$ such that s lies within the circumcircle of $\triangle pqr$. Therefore, by the first geometric fact, we have $\angle prq + \angle psq > \pi$, which implies that at least one of the two angles is not acute (see Fig. 4(a)). Therefore T is not acute, as desired.

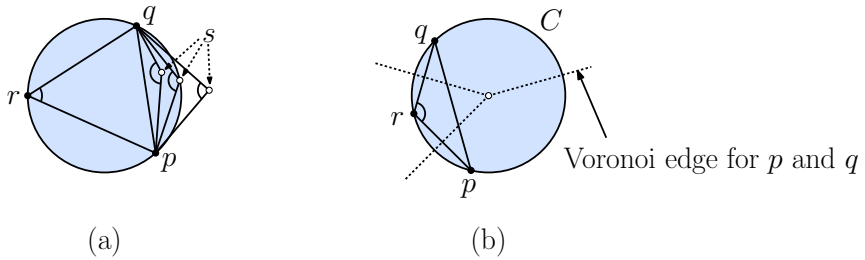


Figure 4: Acute triangulations and medial Voronoi diagrams.

(b) Again, we will show the contrapositive. Suppose that P's Delaunay triangulation is not acute. Consider a triangle $\triangle pqr$ that is not acute at some vertex r, and consider its circumcircle C (see Fig. 4(b)). From the second observation above, C's center (which is the Voronoi vertex of p, q, and r) lies outside the interior of $\triangle pqr$. Therefore, the edge $\overline{pq}$ does not intersect the Voronoi edge for p and q, implying that the triangulation is not medial.

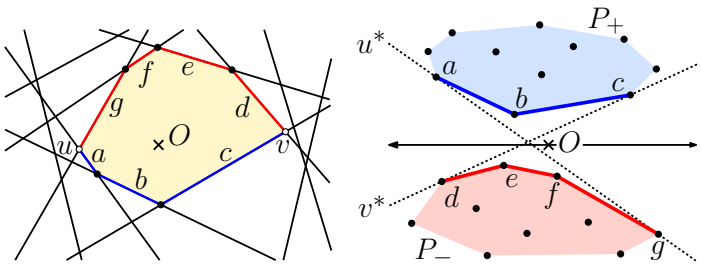
Solution to the Challenge Problem: Throughout, we will make the general-position assumptions that no point P lies on either of the coordinate axes, and all the x- and y-coordinates are distinct.

(a) Let f_0 denote the face of $\mathcal{A}(P^*)$ containing the origin. Let us assume first that f_0 is bounded. Let u and v be the leftmost and rightmost vertices of f_0 (see Fig. 5(a)). These vertices partition the edges bounding f_0 into two groups, those bounding its upper boundary and those bounding its lower boundary. Define P_+ and P_- to be the subsets of P lying above and below the origin, respectively. If either set is empty, then f_0 is either the upper or lower face of the arrangement. It follows from the remarks in Lecture 6 (on the relationship between convex hulls and envelopes) that the edges of the face are in 1–1 correspondence with the vertices of the upper or lower convex hull of P.

Otherwise, suppose that both sets are nonempty. Consider the two lines that define the crossing supporting lines of $\text{conv}(P_+^*)$ and $\text{conv}(P_-^*)$ (see Fig. 5(b)). We claim that:

- These lines are the duals u^* and v^* of the leftmost and rightmost vertices of f.
- The edges bounding f_0's lower boundary from u to v consist of the points lying on the portion of the lower hull of P_- from its intersection with u^* to its intersection with v^*.
- The edges bounding f_0's upper boundary from v to u consist of the points lying on the portion of the upper hull of P_- from its intersection with v^* to its intersection with u^*.

Notice that the portion of the boundary of the convex hulls of P_+ and P_- does not include the entire upper and lower hulls, just the portion between the supporting lines.

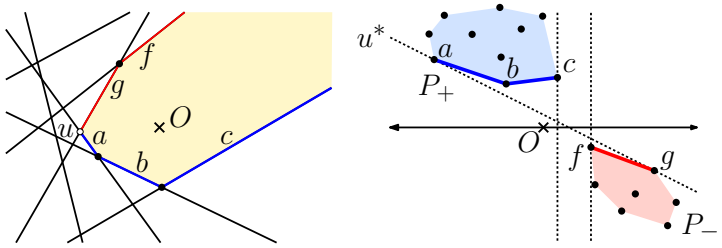


(a) (b)

Figure 5: Edges bounding a face of an arrangement (bounded face).

To see this, observe that the points of P_+ have positive y-coordinates and those of P_- have negative y-coordinates. Thus, the dual lines of P_+^* pass below the origin and those of P_-^* pass above the origin. Therefore, the only dual lines that can contribute to the lower boundary of f_0 are from P_+^*, and those that can contribute to the upper boundary of f_0 are from P_-^*. It follows that the lower boundary of f_0 consists only of edges from the upper envelope of P_+^* and the upper boundary of f_0 consists only of edges from the lower envelope of P_-^*. These envelopes intersect each other in the primal at u and v, and this implies that in the dual, the relevant portions of the hulls are bounded between the supporting points for u^* and v^*.

If the face is unbounded, the two envelopes fail to intersect on one side or the other. (In Fig. 6(a), we show the case where v does not exist.) This means that there exist supporting lines of arbitrarily high slope, that is, vertical lines, separating the two hulls in the dual (see Fig. 6(b)). The answer still consists of a portion of the lower hull of P_+ and the upper hull of P_-, but the extreme points are bounded by these two vertical supporting lines.



(a) (b)

Figure 6: Edges bounding a face of an arrangement (unbounded face).

(b) The general case for a point q lying in an arbitrary face can be transformed to case (a) by applying a translation that maps q to the origin. The principal difference is that the role of the x-axis in defining P_+ and P_- is replaced by the dual line q^*. P_+ consists of the points lying above q^* and P_- are the point below.