CS 274
Computational Geometry (Spring 2017)
Homework 1

Homework rules: All homeworks in this class should be treated like take-home exams. Each homework is comprised mainly of individual questions, but a few harder questions (marked with a **G**) are group questions.

For individual questions, you are on your honor to follow these rules:

- Do not seek or obtain assistance from anyone but me.
- Do not help your classmates with the homeworks.
- Do not read any homework solutions or "sample solutions," including those produced by students (except yourself), instructors, or TAs, related to this or any other course on Computational Geometry or a closely related topic. Some of the homework questions I'll use this semester have been used before by me or other teachers at other institutions, and you're on your honor not to consult the solutions.

For group questions, you may work with one or two other students in the class. The three rules above still apply, except that you work together with your group members. Group questions must be handed in *separately* from questions you solve individually, with the names of all the group members on a single submission. (Only if you solve a group question alone may you include it with your individual solutions.)

If you wish, you may work with a different group for each group question. The maximum permitted group size is three students. If you and a classmate discuss a group question, you are required to submit a solution together.

Outside of your group, or outside of group questions, giving or receiving hints or answers is grounds for an immediate F in the course.

However, you may consult any books, papers, web pages, or other inanimate information sources you please, so long as they're not other people's homework solutions or sample solutions. If you use information from such a source, please cite it. I don't believe that the solutions to many of the problems I'll give are available in easy-to-find references, but if you do find one, it's fair game as long as you cite it. After all, a good literature search should be rewarded. You may not, however, let your classmates know about your discovery or sources until the deadline has passed.

You are welcome to ask me to clarify anything in the readings, the lectures, or this homework.

I prefer typeset solutions, on paper. I will accept handwritten solutions with the proviso that I reserve the right to give a zero to a solution if there is any word or variable I can't read clearly. (Sorry, I just don't have time to infer from context whether that vague squiggle is an *i* or a *j*.) Any such decision is final and irrevocable. By contrast, typeset solutions reserve the right to argue with my grading after the fact.

On the other hand, handwritten *figures* will be gratefully accepted. Anything to encourage you to draw figures! A simple figure can often mean the difference between my spending five minutes or two hours grading an answer.

You may give me your solutions directly, or put them in my CS department mailbox or under my office door (529 Soda) before the deadline. Emailed solutions will be accepted grudgingly, but please use email only if you're out of town or it's otherwise really inconvenient to come to Soda Hall. Snail-mailed solutions (sent to my office; see my web page for the address) will be accepted if they are postmarked the day before the deadline. Solutions postmarked on the day of the deadline will be accepted if the postmark is clearly from well outside the Bay Area. Late homeworks are not accepted.

Homework 1 is **due at the start of class (2:40 pm) on Wednesday, February 15, 2017**.

[1G] Convexity (4 points). Let P be a closed point set (not necessarily finite) in d-dimensional Euclidean space (E^d). Let conv P denote the convex hull of P.

(i) Given a point $q \in E^d \setminus$ closure conv P, describe how to mathematically construct a hyperplane h that separates q from P—to be precise, q is on one side of h, all of P is on the other side, and neither q nor P intersects h. If P is finite, your construction should suggest an algorithm for computing h (assuming you already have a subroutine for computing convex hulls of finite point sets). If P is not finite, your construction should be valid mathematically, though it might not be computable.

(ii) Prove that the hyperplane you construct in part (i) separates q from P.

(iii) Prove that for every point $q \in$ conv P, no hyperplane exists that separates q from P.

Hints: Look for a very simple answer to part (i). I gave three definitions of *convex hull* in class. The definition that makes part (ii) easiest isn't necessarily the definition that makes part (iii) easiest.

[2] Right turn strictly prohibited (4 points). Let P be a set of n points in the plane in general position—that is, no three points in P are collinear.

(i) In pseudocode, write an algorithm that reorders P into a sequence S such that the path that visits the points of S in order (following straight edges from point to point) makes left turns only. In other words, any three consecutive points in S occur in counterclockwise order. Your pseudocode should explicitly give all the orientation tests; e.g. "**if** ORIENT2D$(p, q, r) > 0$".

(ii) Explain in plain English why your algorithm works.

[3] Maximality (6 points). Let P be a set of n points in the plane. A point $p \in P$ is *maximal* if no point in P is both above and to the right of p.

(i) Give a four-point example P that has both a maximal point not on the convex hull, and a convex hull point that is not maximal.

(ii) For any P, give an $O(n)$-time algorithm to compute the maximal points of P. Be careful to handle degeneracies correctly (i.e., points with the same x-coordinates or the same y-coordinates).

(iii) Suppose we are restricted to a decision-tree model of computation. Give an $O(n \log h)$-time algorithm to compute the h maximal points of P. (Adapt Chan's Shattering algorithm.)

[4] Face topology reconstruction (2 points). Suppose we are restricted to a decision-tree model of computation. Practically, this means we cannot use radix sort, bucketing, or other techniques that make k-way decisions for large k; we are only permitted to make binary decisions (i.e., **if** statements). It's well known, for instance, that no comparison-based sort can have a better worst-case running time than $\Theta(n \log n)$, even though radix sort (which is not comparison-based) runs in $O(n)$ time.
(See http://www.cs.berkeley.edu/~jrs/61b/lec/34 if you're not familiar with this kind of lower bound result.)

Suppose we are trying to compute the overlay of two subdivisions, and we have completed the sweepline intersection step. We now have a DCEL with the faces missing, and we want to determine the face topologies next (i.e., match outer boundaries of faces with inner boundaries of their holes). However, we have foolishly not retained any information from the sweepline stage. Explain why the worst-case running time of every algorithm that determines the face topologies is in $\Omega(h \log h)$ (in the decision-tree model), where h is the complexity of the DCEL.

[5G] Few turn vertices (6 points). The sweepline polygon triangulation algorithm (Chapter 3 of the Dutch Book) begins by sorting the vertices of the polygon along one coordinate axis. Suppose the input polygon has n vertices including r turn vertices, where r might be much smaller than n.

  (i) Describe an algorithm that sorts the vertices in $O(n \log r)$ time, even in the decision-tree model of computation (which rules out radix sort).
 (ii) Argue that if you use this sorting algorithm, the sweepline polygon triangulation algorithm runs in $O(n \log r)$ time (with no other changes).
(iii) Suppose you use radix sort, which runs in $O(n)$ time. Explain how to remove all tree operations from the treatment of regular vertices, so the sweepline triangulation algorithm runs in $O(n + r \log r)$ time. You may not switch to an entirely new algorithm, but you can make changes to the sweepline algorithm outside of handling regular vertices, such as: changing the way edges are stored in the balanced tree data structure, or changing the tree data structure itself; using additional simple data structures; and augmenting the vertex handling routines to store extra information.

[6] Red-black closest pair (3 points). Let P and Q be two disjoint point sets in the plane. (Think of them as a red point set and a black point set.) Let $p \in P$ and $q \in Q$ be two points from these sets that minimize the Euclidean distance $|pq|$. Prove that $\overline{pq}$ is an edge of DT$(P \cup Q)$. (This observation leads easily to an $O(n \log n)$ algorithm for finding p and q, the red-black closest pair.) Hint: there is a simple answer to this question, and almost everybody gives a much more complicated answer.