F Line Segment Intersection

F.1 Introduction

In this lecture, I'll talk about detecting line segment intersections. A line segment is the convex hull of two points, called the *endpoints* (or *vertices*) of the segment. We are given a set of n line segments, each specified by the x- and y-coordinates of its endpoints, for a total of $4n$ real numbers, and we want to know whether any two segments intersect.

To keep things simple, just as in the previous lecture, I'll assume the segments are in *general position*.

- No three endpoints lie on a common line.
- No two endpoints have the same x-coordinate. In particular, no segment is vertical, no segment is just a point, and no two segments share an endpoint.

This general position assumption lets us avoid several annoying degenerate cases. Of course, in any real implementation of the algorithm I'm about to describe, you'd have to handle these cases. Real-world data is *full* of degeneracies!

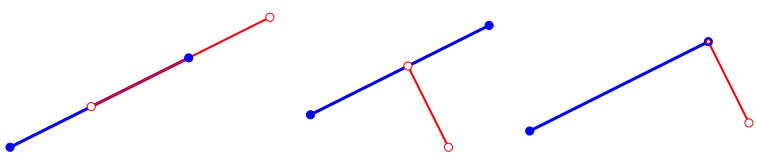

Degenerate cases of intersecting segments that we'll pretend never happen:
Overlapping colinear segments, endpoints inside segments, and shared endpoints.

F.2 Two segments

The first case we have to consider is $n = 2$. ($n \leq 1$ is obviously completely trivial!) How do we tell whether two line segments intersect? One possibility, suggested by a student in class, is to construct the convex hull of the segments. Two segments intersect if and only if the convex hull is a quadrilateral whose vertices alternate between the two segments. In the figure below, the first pair of segments has a triangular convex hull. The last pair's convex hull is a quadrilateral, but its vertices don't alternate.

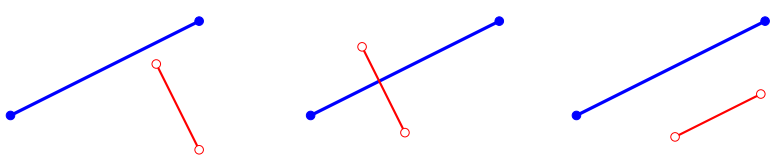

Some pairs of segments.

Fortunately, we don't need (or want!) to use a full-fledged convex hull algorithm just to test two segments; there's a much simpler test.

Two segments $\overline{ab}$ and $\overline{cd}$ intersect if and only if

- **the endpoints a and b are on opposite sides of the line $\overleftrightarrow{cd}$, and**
- **the endpoints c and d are on opposite sides of the line $\overleftrightarrow{ab}$.**

To test whether two points are on opposite sides of a line through two other points, we use the same counterclockwise test that we used for building convex hulls. Specifically, a and b are on opposite sides of line $\overleftrightarrow{cd}$ if and only if exactly one of the two triples a, c, d and b, c, d is in counterclockwise order. So we have the following simple algorithm.

```
INTERSECT(a, b, c, d):
  if CCW(a, c, d) = CCW(b, c, d)
      return FALSE
  else if CCW(a, b, c) = CCW(a, b, d)
      return FALSE
  else
      return TRUE
```

Or even simpler:

```
INTERSECT(a, b, c, d):
  return [CCW(a, c, d) ≠ CCW(b, c, d)] ∧ [CCW(a, b, c) ≠ CCW(a, b, d)]
```

F.3 A Sweep Line Algorithm

To detect whether there's an intersection in a set of more than just two segments, we use something called a *sweep line* algorithm. First let's give each segment a unique *label*. I'll use letters, but in a real implementation, you'd probably use pointers/references to records storing the endpoint coordinates.

Imagine sweeping a vertical line across the segments from left to right. At each position of the sweep line, look at the sequence of (labels of) segments that the line hits, sorted from top to bottom. The only times this sorted sequence can change is when the sweep line passes an endpoint or when the sweep line passes an intersection point. In the second case, the order changes because two adjacent labels swap places.[1] Our algorithm will simulate this sweep, looking for potential swaps between adjacent segments.

The sweep line algorithm begins by sorting the $2n$ segment endpoints from left to right by comparing their x-coordinates, in $O(n \log n)$ time. The algorithm then moves the sweep line from left to right, stopping at each endpoint.

We store the vertical label sequence in some sort of balanced binary tree that supports the following operations in $O(\log n)$ time. Note that the tree does not store any explicit search keys, only segment labels.

- **Insert** a segment label.
- **Delete** a segment label.
- Find the **neighbors** of a segment label in the sorted sequence.

$O(\log n)$ amortized time is good enough, so we could use a scapegoat tree or a splay tree. If we're willing to settle for an expected time bound, we could use a treap or a skip list instead.

[1] Actually, if more than two segments intersect at the same point, there could be a larger reversal, but this won't have any effect on our algorithm.

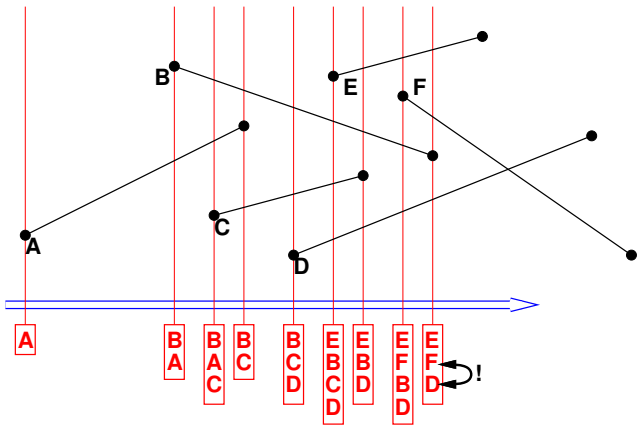

The sweep line algorithm in action. The boxes show the label sequence stored in the binary tree. The intersection between F and D is detected in the last step.

Whenever the sweep line hits a left endpoint, we insert the corresponding label into the tree in $O(\log n)$ time. In order to do this, we have to answer questions of the form 'Does the new label X go above or below the old label Y?' To answer this question, we test whether the new left endpoint of X is above segment Y, or equivalently, if the triple of endpoints left(Y), right(Y), left(X) is in counterclockwise order.

Once the new label is inserted, we test whether the new segment intersects either of its two neighbors in the label sequence. For example, in the figure above, when the sweep line hits the left endpoint of F, we test whether F intersects either B or E. These tests require $O(1)$ time.

Whenever the sweep line hits a right endpoint, we delete the corresponding label from the tree in $O(\log n)$ time, and then check whether its two neighbors intersect in $O(1)$ time. For example, in the figure, when the sweep line hits the right endpoint of C, we test whether B and D intersect.

If at any time we discover a pair of segments that intersects, we stop the algorithm and report the intersection. For example, in the figure, when the sweep line reaches the right endpoint of B, we discover that F and D intersect, and we halt. Note that we may not discover the intersection until long after the two segments are inserted, and the intersection we discover may not be the one that the sweep line would hit first. It's not hard to show by induction (hint, hint) that the algorithm is correct. Specifically, if the algorithm reaches the nth right endpoint without detecting an intersection, none of the segments intersect.

For each segment endpoint, we spend $O(\log n)$ time updating the binary tree, plus $O(1)$ time performing pairwise intersection tests—at most two at each left endpoint and at most one at each right endpoint. Thus, the entire sweep requires $O(n \log n)$ time in the worst case. Since we also spent $O(n \log n)$ time sorting the endpoints, the overall running time is $O(n \log n)$.

Here's a slightly more formal description of the algorithm. The input $S[1..n]$ is an array of line segments. The sorting phase in the first line produces two auxiliary arrays:

- label$[i]$ is the label of the ith leftmost endpoint. I'll use indices into the input array S as the labels, so the ith vertex is an endpoint of $S[\text{label}[i]]$.
- isleft$[i]$ is TRUE if the ith leftmost endpoint is a left endpoint and FALSE if it's a right endpoint.

The functions INSERT, DELETE, PREDECESSOR, and SUCCESSOR modify or search through the sorted label sequence. Finally, INTERSECT tests whether two line segments intersect.

```
AnyIntersections(S[1..n]):
  sort the endpoints of S from left to right
  create an empty label sequence
  for i ← 1 to 2n
      ℓ ← label[i]
      if isleft[i]
          Insert(ℓ)
          if Intersect(S[ℓ], S[Successor(ℓ)])
              return True
          if Intersect(S[ℓ], S[Predecessor(ℓ)])
              return True
      else
          if Intersect(S[Successor(ℓ)], S[Predecessor(ℓ)])
              return True
          Delete(label[i])
  return False
```

Note that the algorithm doesn't try to avoid redundant pairwise tests. In the figure below, the top and bottom segments would be checked $n-1$ times, once at the top left endpoint, and once at the right endpoint of every short segment. But since we've already spent $O(n \log n)$ time just sorting the inputs, $O(n)$ redundant segment intersection tests make no difference in the overall running time.

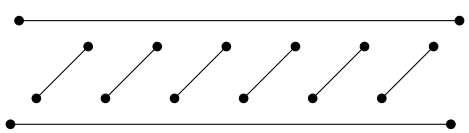

The same pair of segments might be tested $n-1$ times.