Today : - Computing Delaunay triangulation
- Curve reconstruction

- Incremental algorithm for computing Del(P)

Let $P = \{p_1, \ldots, p_n\}$ be a point set in the plane

for simplicity: (general position) no four points on a circle

Let $P_K = \{p_1, \ldots, p_K\}$ for $3 \le K \le n$

Incremental Algorithm

Given $Del(P_K)$, add a new point p_{K+1} and compute $Del(P_{K+1})$

(1) Find the set of triangles of $Del(P_{K-1})$ whose circumcircles contain the new point p_K
Their union is a triangulated polygon

(2) Remove diagonals of this polygon and add edges from p_K to each of the vertices of the polygon

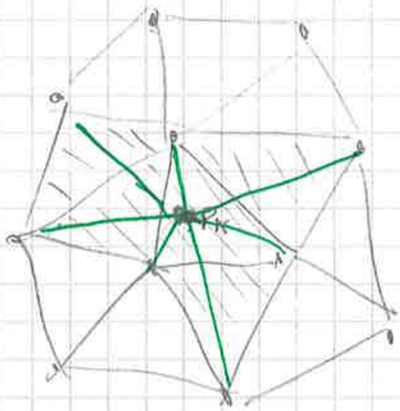

$Del(P_{K-1})$

The result is $Del(P_K)$ (Exercise)

Complexity : $O(n^2)$

# of triangles is linear on the number of vertices → finding triangle in which p_K is is linear

degree (p_K) at step K is at most $K-1$

Randomized Incremental Algorithm
Only sketched (see Comp. Geometry Book , Chapters 9.3 & 9.4)

(1) Add big triangle p_{-1}, p_{-2}, p_{-3}
(2) Add interior point from P randomly and compute Delaunay triangulation at each step.

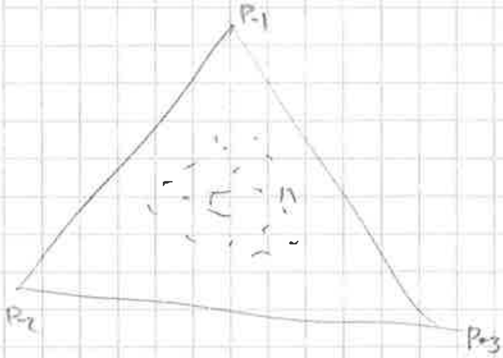

Expected # bad triangles = $9n+1$

point location : $O(\log n)$

Complexity : $O(n \log n)$.

Curve reconstruction

Motivation: Surface reconstruction from a sample of points on the surface of a 3D object. (i.e. scanning divices)

Method: Connect nearby points into some type of mesch, a surface of triangles.

Today, a simpler problem: Reconstruction of curves in $\mathbb{R}^2$

Problem: Given an (unknown) curve C in $\mathbb{R}^2$ and a finite sample P of points in C, compute the order of the points in C.

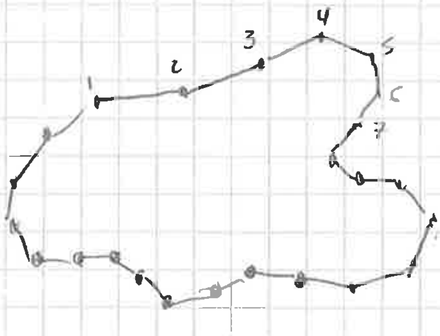

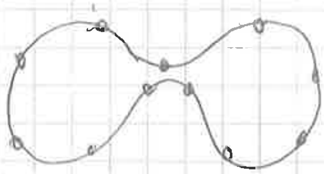

Not possible in general (several curves may have the same sample)

We need a "good sample": sufficiently dense.

We use the concept of medial axis.

Medial axis:

For a curve C, the medial axis $M(C)$ is the set of all points $x \in \mathbb{R}^2$ such that the closest point of C to x is not unique

C

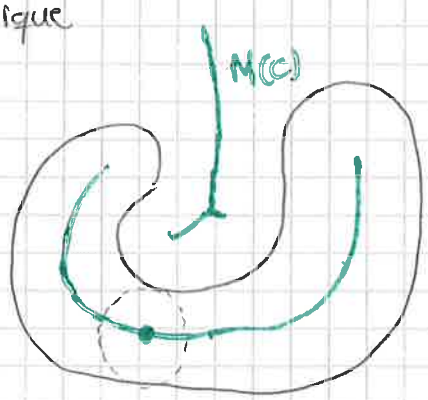

C

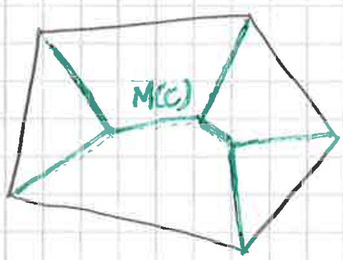

If C would be just a finite set of points:

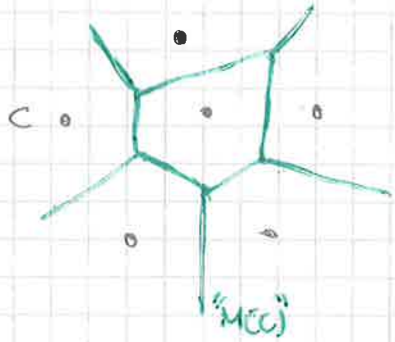

we get "M(C)" is the Voronoi diagram

So M(C) is like a generalization of Voronoi diagrams for lines/curves

Define the local feature size (lfs) of $p \in C$ is its distance to $M(C)$

Intuition: lfs is small for points in "complicated regions" of C

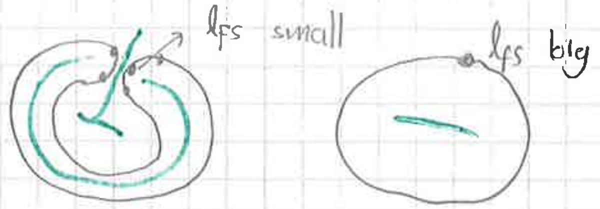

For $0 < \varepsilon \leq 1$, a point sample P of C is called <u>ε-sample</u> if for every $x \in C$, $\exists p \in P$ such that

$$\|x - p\| \leq \underbrace{\varepsilon \cdot lfs(x)}$$

↳ implies use more points in complicated regions.

<u>Theorem</u> (Amenta, Bern, Eppstein 1998) If P is a 0,252-sample ($\varepsilon \sim 1/4$) of C, then a reconstruction of C can be computed in $O(n \log n)$ time.

Algorithm is elegant an simple: It uses Delaunay triangulations and Voronoi diagrams.

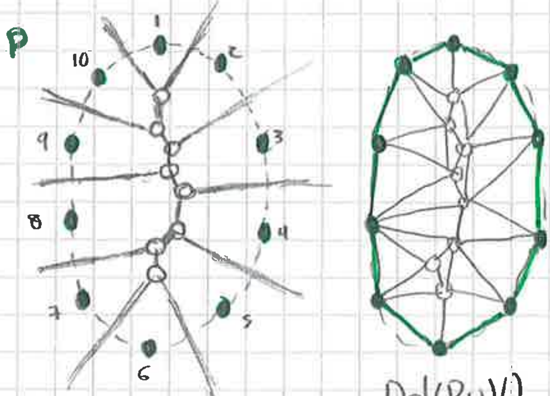

Vor(P) Del(P∪V)

<u>CRUST Algorithm</u> Let P be the set of sample points. (1) Compute Vor(P) and let V its set of vertices (2) Compute Del($P \cup V$) (3) The curve of P is composed of the edges of Del($P \cup V$) with both end points in P

Observe: Vertices of Vor(P) = V close to the medial axis!

<u>Complexity</u> (1) $O(n \log n)$
(2) $O(n \log n)$ because $|V|$ is linear

Total: $O(n \log n)$

Key observations for the proof

(1) Del (P) contains all edges between consecutive sample points.
("correct" edges)

sketch idea

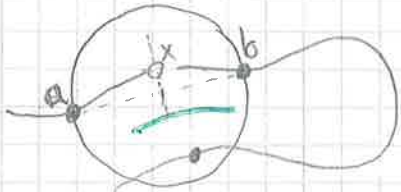

otherwise lfs(x) is too small but distance to a,b is large. (a,b too far apart)

diametral ball of consecutive points a,b∈P must be empty

(2) A circumscribing disk of an edge of Del(P) between non-consecutive points intersects the medial axis M(C)
("incorrect edges")

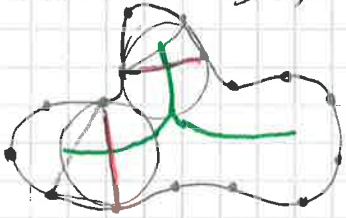

The center of the disk or a point close by lies on M(C)

Consequence : If we know M(C), we could detect and remove incorrect edges
But : C is unknown and so is M(C)

(3) The Voronoi vertices of Vor(P) are close to M(C)

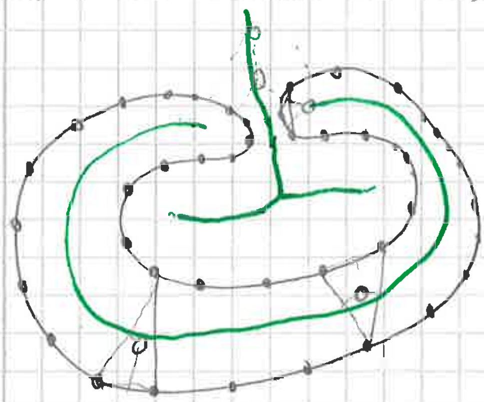

Follows from (2) :
Given p_i, p_j, p_k forming a triangle of Del(P)

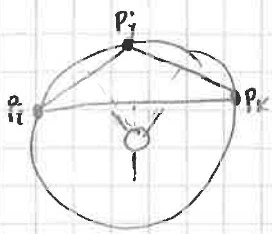

at least one edge is incorrect
so the center of the circumscribing disk of p_i, p_j, p_k (which is a vertex of Vor(P))
is close to the medial axis.

Let V be the set of all Voronoi vertices of $Vor(P)$
Consider the Delaunay triangulation of $P \cup V$

$$Del(P \cup V)$$

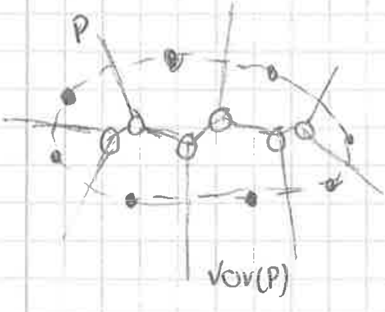

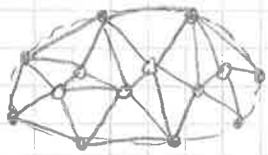

idea:
Voronoi Vertices "sample" the medial axis

(4) Every "correct edge" is still contained in $Del(P \cup V)$

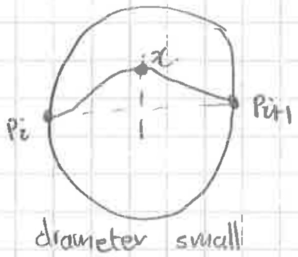

$\|x - p\| \leq 0{,}25 \varepsilon \, lfs(x) \leq \frac{1}{4} lfs(x)$

$\rightsquigarrow$ Disk does not intersect the medial axis

(5) No incorrect edge $p_i p_j$ of $Del(P)$ appears in $Del(P \cup V)$

Follows also from (2) :

Let $p_i p_j$ be an incorrect edge of $Del(P)$.
Any circle with $p_i p_j$ on the boundary with no point of P inside contains a point of the medial axis, close to its center.
Therefore, a point of V. (Voronoi vertices "sample" the medial axis)
So $p_i p_j$ is not an edge of $Del(P \cup V)$

As a consequence of (4) & (5) :

The correct edges of $Del(P)$ are precisely the edges of $Del(P \cup V)$ with both end points in P.

Thus, the CRUST algorithm reconstructs the curve C.

Varres variant algorithms: NN-CRUST, Power-CRUST.

Relax the sample condition : increasing ε
(less sample points
farther apart)

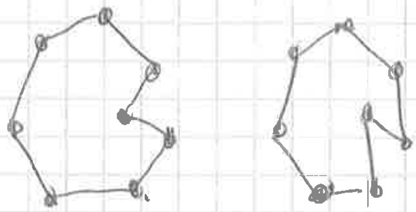

two possible curves for the same ε-sample
ε to big

Thm - $\varepsilon = 0{,}252$ suffices
$\varepsilon = 1$ impossible

Bjerkevik 2022 : 0,66 - sample suffices
0,72 - sample impossible.

Question: What is the "true" (biggest) ε that suffices ?