MAI520 Computational Geometry 2003/2004
Peter Kumlin
Matematiska institutionen CTH& GU

Additional material for the course

"Computational Geometry" MAI520

The course "Computational Geometry" based on O'Rourke's book "Computational Geometry in C" was offered for the first time at Chalmers/GU during the late fall 2000. The following notes cover the additional material that was discussed in class. Some parts stem from a project, together with Björn Dahlberg, to collect the major ideas from the different fields of geometry – differential geometry, discrete geometry, computational geometry – and to present the material in a unified way, all aspects being present at the same time. The sudden death of Dahlberg made it impossible for such a project. Parts of these notes are some scattered debris from this undertaking. Needless to say, all shortcomings regarding these notes, are due to me.

Peter Kumlin

Handouts

- V.Chvatal: A Combinatorial Theorem in Plane Geometry, Journal of Combinatorial Theory (B) 18 (1975) 39 – 41
- N.Dyn/I.Goren/S.Rippa: Transforming triangulations in polygonal domains, Computer Aided Geometric Design **10** (1993) 531 – 536
- S.Rippa: Minimal Roughness Property of the Delaunay Triangulation, Computer Aided Geometric Design 7 (1990) 489 – 497
- D.G.Kirkpatrik/R.Seidel: The Ultimate Planar Convex Hull Algorithm?: SIAM J.Comput. 15 (1986) 287 – 299

Homework

- Exercises 1
- Exercises 2
- Exercises 3

Handbooks

- J.Goodman/J.O'Rourke (eds): Handbook of discrete and computational geometry, CRC Press 1997
- K.Rosen (ed): Handbook of discrete and combinatorial mathematics, CRC Press 2000
- J.-R.Sack/J.Urrutia (eds): Handbook of computational geometry, Elsevier 2000

Contents

1 Polygon Triangulation/Polygon Partitioning

At a conference in Stanford in 1973 Klee posed the problem of determining the minimum number of guards sufficient to cover the interior of an n–wall art gallery room. This problem was given in a response to a request from Chvatal for an interesting challenge. Chvatal soon supplied what has become known as Chvatal's Art Gallery Theorem, i.e. $\lfloor \frac{n}{3} \rfloor$ *guards are occasionally necessary and always sufficient to cover a polygon of n vertices.*

1.1 Art Gallery Theorems

The first proof of Chvátal's Art Gallery Theorem was of course given by Chvátal in 1975, see [14]. His proof starts with a triangulation of the polygon, as does Fisk's, but does not use graph coloring. Rather the theorem is proven directly by induction. Although Chvátal's proof is not as concise as Fisk's, it reveals aspects of the problem that are not brought to light by the coloring argument.

Define a **fan** as a triangulation with one vertex (the fan **center**) shared by all triangles. Chvátal took as his induction hypothesis this statement:

> *Induction Hypothesis*: Every triangulation of an n-gon can be partitioned into $g \leq \lfloor n/3 \rfloor$ fans.

For the basis, note that $n \geq 3$ since we start with an n-gon, and that there is just a single triangulation possible when $n = 3, 4$, and 5, each of which is a fan. Thus the induction hypothesis holds for $n < 6$.

Given a triangulation with $n \geq 6$, our approach will be to remove part of the triangulation, apply the induction hypothesis, and then put back the deleted piece. We know that there is always a diagonal (in fact, there are always at least two) that partitions off a single triangle. But note that this only reduces n by 1, and if we were unlucky enough to start with $n \equiv 1$ or 2 (mod) 3, then the induction hypothesis partitions into $g = \lfloor (n-1)/3 \rfloor = \lfloor n/3 \rfloor$ fans, and we will in general end up with $g+1$ fans when we put back the removed triangle. The moral is that, in order to make induction work with the formula $\lfloor n/3 \rfloor$, we have to reduce n by at least 3 so the induction hypothesis will yield less than g fans, allowing the grouping of the removed triangles into a fan.

So the question naturally arises: does there always exist a diagonal that partitions off 4 edges of the polygon, and therefore reduces n by 3? The answer is no! Chvátal's brilliant idea was to realize that there is always a diagonal that cuts off 4 or 5 or 6 edges:

Lemma 1 (Chvátal 1975) *For any triangulation of an n-gon with $n \geq 6$, there always exists a diagonal d that cuts off exactly 4, 5 or 6 edges.*

PROOF: Choose d to be a diagonal that separates off a minimum number of polygon edges that is at least 4. Let $k \geq 4$ be the minimum number, and label the vertices of the polygon $0, 1, \ldots, n-1$ such that d is $(0, k)$. d must support a triangle T whose apex is at some vertex t with $0 \leq t \leq k$. Since $(0, t)$ and (k, t) each cut off fewer than k edges, by the minimality of k we have $t \leq 3$ and $k - t \leq 3$. Adding these two inequalities yields $k \leq 6$.

Now the plan is to apply the induction hypothesis to the portion on the other side of the special diagonal d. Let G_1 be the triangulation partitioned off by d; it has $k+1$ boundary edges and hence is a $(k+1)$-gon. Let G_2 be the remainder of the original triangulation, sharing d; it has $n-k+1$ vertices. The induction hypothesis says that G_2 may be partitioned into $g' = \lfloor (n-k+1)/3 \rfloor$ fans. Since $k \geq 4, g' \leq \lfloor (n-3)/3 \rfloor = \lfloor n/3 \rfloor - 1$. Thus, in order to establish the theorem, we have to show that G_1 need only add one more fan to the partition. We will consider each possible value of k in turn.

Case 1 ($k = 4$). G_1 is a 5-gon. We already observed that every pentagon is a fan. Therefore, G has been partitioned into $\lfloor n/3 \rfloor - 1 + 1 = \lfloor n/3 \rfloor$ fans.

Case 2 ($k = 5$). G_1 is an 6-gon. Consider the triangle T of G_1 supported by d, with its apex at t. We cannot have $t = 1$ or $t = 4$, as then the diagonals $(0, t)$ or $(5, t)$ [respectively] would cut off just 4 edges, violating the assumed minimality of $k = 3$. The cases $t = 2$ and $t = 3$ are clearly symmetrical, so assume without loss of generality that $t = 2$. Now the quadrilateral $(2, 3, 4, 5)$ can be triangulated in two ways:

Case 2a. The diagonal $(2, 4)$ is present. Then G_1 is a fan, and we are finished.

Case 2b. The diagonal $(3, 5)$ is present. Form the graph G_0 as the union of G_2 and T. G_0 has $n-5+1+1 = n-3$ edges. Apply the induction hypothesis to it, partitioning it into $g' = \lfloor (n-3)/3 \rfloor = \lfloor n/3 \rfloor - 1$ fans. Now T must be part of a fan F in the partition of G_0, and the center of F must be at one of T's vertices:

Case 2b.1. F is centered at 0 or 2. Then merge $(0, 1, 2)$ into F, and make $(2, 3, 4, 5)$ its own fan. Now all of G is covered with $g' + 1 = \lfloor n/3 \rfloor$ fans.

Case 2b.2. F is centered at 5. Merge both $(2, 3, 5)$ and $(3, 4, 5)$ into F, and make $(0, 1, 2)$ a separate fan. The result is $g' + 1$ fans.

Case 3 ($k = 6$). G_1 is a 7-gon. The tip t of the triangle T supported by d cannot be at 1, 2, 4, 5, as then a diagonal would exist that cuts off

$4 \leq k < 6$ edges, contradicting the minimality of k. Thus $t = 3$. Each of the two quadrilaterals $(0, 1, 2, 3)$ and $(3, 4, 5, 6)$ has two possible triangulations, leading to four subcases.

Case 3a. The diagonals $(3, 1)$ and $(3, 5)$ are present. Then G_1 is a fan centered at 3, and we are finished.

Case 3b. The diagonals $(0, 2)$ and $(3, 5)$ are present. Join the quadrilateral $(0, 2, 3, 6)$ to G_2 to form a polygon G_0 with $n - 6 + 1 + 2 = n - 3$ vertices, which by the induction hypothesis can be partitioned into $g' = \lfloor n/3 \rfloor - 1$ fans. Let F be the fan of this partition to which the triangle $(0, 2, 3)$ belongs. The center of F must be at one of its vertices:

Case 3b.1. F is centered at 0 or 2. Merge $(0, 1, 2)$ into F and make $(3, 4, 5, 6)$ a separate fan.

Case 3b.2. F is centered at 3. Merge $(3, 4, 5, 6)$ into F, and make $(0, 1, 2)$ a separate fan.

In all cases, G is partitioned into $g' + 1 = \lfloor n/3 \rfloor$ fans.

Case 3c. The diagonals $(1, 3)$ and $(4, 6)$ are present. This is the mirror image of Case 3b.

Case 3d. The diagonals $(0, 2)$ and $(4, 6)$ are present. Merge T with G_2 to form a polygon G_0 of $n - 6 + 1 + 1 = n - 4$ vertices. Applying the induction hypothesis partitions G_0 into $g' = \lfloor (n - 4)/3 \rfloor \leq \lfloor n/3 \rfloor - 1$ fans. Let F be the fan of the partition containing T.

Case 3d.1. F is centered at 0. Merge the quadrilateral $(0, 1, 2, 3)$ into F, and make $(3, 4, 5, 6)$ a separate fan.

Case 3d.2. F is centered at 3. Since all of G_2 is behind the $d = (0, 6)$ diagonal, it is clear that we can just as well consider F to be centered at 0, falling into Case 3d.1.

Case 3d.3. F is centered at 6. This is the mirror image of Case 3d.1.

In all cases, G is partitioned into $g' + 1 = \lfloor n/3 \rfloor$ fans. This completes the proof. □

Placing guards at the fan centeres establishes the theorem:

Theorem 1 (Chvátal's Art Gallery Theorem 1975) $\lfloor n/3 \rfloor$ *guards are occasionally necessary and always sufficient to see the entire interior of a polygon of* n *edges.*

Note that both Chvátal's and Fisk's proofs incidentally establish by construction that the guards can be chosen to be vertex guards. However it is an easy task to find an example where one point guard is sufficient but more than one vertex guard is needed to cover the polygon.

The material in this section comes from O'Rourke [42].

1.2 Sorting/Order Statistics

We shall in this section briefly discuss the problem of sorting a finite set that has an order relation. By sorting a sequence we mean rearranging the elements so that they appear in non-increasing or non-decreasing order.

A significant portion of commercial data processing involves sorting large quantities of data. Therefore a large number of sorting algorithms have been developed, see Knuth [32].

There are two classes of sorting algorithms. The first class assumes no structure on the elements to be sorted. The basic operation is a comparison between a pair of elements. The second class of algorithms makes use of the structure of the set to be sorted. We will as an example of this discuss the sorting of a set of integers in a fixed range.

A **partial order** on a set S is a relation R such that for each a, b and c in S:

1. aRa is true (R is reflexive);

2. aRb and bRc imply aRc (R is transitive);

3. aRb and bRa imply $a = b$ (R is antisymmetric).

ruined the A **linear order** or **total order** on a set S is a partial order such that for every pair of elements a, b we have that aRb or bRa. The relation $\leq$ on the integers is a total order but set inclusion is a partial order that in general is not a total order.

The sorting problem can be formulated as follows. We are given a sequence of n elements $a_1, a_2, \ldots, a_n$ from a set having a total order, which we shall denote by $\leq$. We want to find a permutation π of these n elements such that $a_{\pi(1)}, a_{\pi(2)}, \ldots, a_{\pi(n)}$ appear in a non-decreasing order, i.e.

$$a_{\pi(1)} \leq a_{\pi(2)} \leq \ldots \leq a_{\pi(n)}.$$

Knuth has introduced a notational device that distinguishes between upper and lower bounds, which we adopt.

- $O(f(n))$ denotes the set of all functions $g(n)$ such that there exist positive constants C and n_0 with

$$|g(n)| \leq Cf(n)$$

for all $n \geq n_0$.

- $\Omega(f(n))$ denotes the set of all functions $g(n)$ such that there exist positive constants C and n_0 with

$$g(n) \geq Cf(n)$$

for all $n \geq n_0$.

- $\theta(f(n))$ denotes the set of all functions $g(n)$ such that there exist positive constants C_1, C_2 and n_0 with

$$C_1 f(n) \leq g(n) \leq C_2 f(n)$$

for all $n \geq n_0$.

- $o(f(n))$ denotes the set of all functions $g(n)$ such that for all positive constants C there is an n_0 with

$$|g(n)| \leq Cf(n)$$

for all $n \geq n_0$.

We first show that any algorithm which sorts by comparison only, must on some sequence of length n use at least $\Omega(n \log n)$ comparisons.

The problem of sorting by comparisons can be expressed in other ways. Given a set of distinct weights and a balance scale, we can ask for the least number of weightings necessary to completely rank the weights in order of magnitude, when the pans of the balance scale can each accommodate only one weight.

All n–elements sorting methods which satisfy the above constraints can be represented in terms of a binary tree structure. Each node contains two indices $[i, j]$ denoting a comparison of a_i versus a_j. The left outgoing edge corresponds to the outcome $a_i \leq a_j$ and the right edge represents the case $a_j \leq a_i$. We represent the termination of the algorithm by a box containing the permutation of $a_1, a_2, \ldots, a_n$ that leads to the ordering

$$a_{\pi(1)} \leq a_{\pi(2)} \leq \ldots \leq a_{\pi(n)}.$$

figure [1]

Theorem 2 *There is a constant $C > 0$ such that for any algorithm that sorts n elements by comparison only, there is a sequence of n elements for which the algorithm requires at least $C\, n \log n$ comparisons*

PROOF: In the above model for an algorithm we can think of the termination points of the algorithm as leaves. We claim that after k steps of the algorithm we have at most 2^k leaves. This is obviously true for $k = 1$ and if $k > 1$ we see that the tree is a union of two disjoint subtrees, each of which has height $k - 1$. An induction argument therefore establishes the claim. Let N be the minimum number of comparisons that are required for any algorithm. Since the number of permutations of n elements is $n!$ we must have

$$2^N \geq n! \geq (\frac{n}{2})^{\frac{n}{2}},$$

which shows the result. □

[1] In these lecture notes there are not yet any figures. The reader has do the drawings himself/herself. Hopefully the next version will contain proper figures.

We remark that using a better estimate of $n!$ like Stirling's formula, one can show that

$$N \geq \frac{n \log n + \frac{1}{2} \log n - n}{\log 2}.$$

Suppose S_1 and S_2 are two sequences that have been sorted in increasing order. We can easily sort the sequence one gets by concatenating S_1 and S_2 by successively selecting the smallest remaining element from S_1 and S_2. (In order to make this into a well-defined algorithm we break the ties in favour of S_1) This algorithm requires at most $m+n-1$ comparisons, where m, n are the number of elements in S_1, S_2. This remark is the basis for the following **divide-and-conquer**–method.

Theorem 3 *Let $T(n)$ be the smallest number of comparisons that are needed to sort a sequence of n elements in increasing order. Then*

$$T(n) = \theta(n \log n).$$

PROOF: If n is even it follows from the above remark that

$$T(n) \leq 2T(\frac{n}{2}) + n - 1.$$

If n is odd we have that

$$T(n) \leq T(\frac{n-1}{2}) + T(\frac{n+1}{2}) + n - 1.$$

An easy induction argument shows that $T(n) = O(n \log n)$. Using Theorem 2 concludes the proof. □

If a sequence to be sorted has some additional features it may be sorted in linear time. We will only discuss one example of this.

Let S be a set of integers $a_1, a_2, \ldots, a_n$ in the range $\{0, 1, \ldots, m-1\}$, $m \geq 1, n \geq 1$. We shall now describe the **bucket sort**–method. For each $k \in \{0, 1, \ldots, m-1\}$ initialize a que – each such que is called a bucket. Scan the sequence from left to right and place the element a_i in the a_i–bucket. Concatenate the buckets by appending the content of $k+1$ to the end of que k. This yields the sorted sequence in $O(n+m)$ time. If $m = O(n)$ we can therefore sort the sequence in linear time $O(n)$.

We now turn to **order statistics**.

Let $S = \{a_j\}_{j=1}^n$ be a sequence of elements with a total ordering $\leq$. We will say that a_j is the kth **smallest element** if a_j is the smallest element in S with the property that there is a subset $J \subset \{1, 2, \ldots, n\}$ with $k-1$ elements such that

$$a_i \leq a_j \quad for \ \ all \ \ j \in J.$$

Notice that the index of the kth smallest element does not need to be unique.

For real numbers x let $\lceil x \rceil$ be the smallest integer ν such that $\nu \geq x$. ($\lfloor x \rfloor$ denotes the largest integer ν such that $\nu \leq x$) We recall that the choice $k = \lceil \frac{n}{2} \rceil$ corresponds to the **median** of S. By careful application of the divide-and-conquer strategy, often referred to as **prune and search**, we can find the kth smallest element of S in $O(n)$ time.

Theorem 4 *Let S be a sequence of n elements with a linear ordering $\leq$. Then there is an algorithm that finds the kth smallest element in S in linear time $O(n)$.*

PROOF: Let $\varphi(k, S)$ denote the kth smallest element in S. Let $\tau(k, n)$ be the shortest time for finding the kth smallest element and set

$$T(n) = \max\{\tau(k, n) : 1 \leq k \leq n\}.$$

Assume $n \geq 100$. Set $m = \lceil \frac{n}{5} \rceil$ and split S into sets with 5 elements each and possibly one with fewer elements than 5. Call these sets $E_1, \ldots, E_m$ with E_ν having 5 elements for $\nu = 1, 2, \ldots, m-1$. Let x_k be the median of E_k and set $X = \{x_k\}_{k=1}^m$. Let $q = \varphi(\lceil \frac{m}{2} \rceil, X)$ be the median of X and set $J_1 = \{l \in I_n : a_l \leq q \ \ and \ a_l \neq q\}$, $J_2 = \{l \in I_n : a_l = q\}$ and $J_3 = \{l \in I_n : q \leq a_l \ \ and \ \ a_l \neq q\}$. Here $I_n = \{1, 2, \ldots, n\}$. Let ν_i denote the number of elements of J_i, $i = 1, 2, 3$. We claim that ν_1 and ν_3 are each at most $\frac{3n}{4}$. Let $Q = \{j : 1 \leq j < m \ \ and \ \ q \leq x_j\}$. Notice that Q contains at least $\frac{n}{10}$ elements (almost). For each $j \in Q$ there are at least two additional elements in E_j that are $\geq q$. Hence the size of J_1 is at most $n - 3\frac{n}{10} \leq \frac{3n}{4}$. Similar arguments applies to J_3. We can now observe the following. Let $S_i = \{a_l : l \in J_i\}$, $i = 1, 2, 3$. If $\nu_1 \geq k$ then $\varphi(k, S) = \varphi(k, S_1)$. If $\nu_1 < k$ but $\nu_1 + \nu_2 \geq k$ then $\varphi(k, S) = q$. If $\nu_1 + \nu_2 < k$ then $\varphi(k, S) = \varphi(k - \nu_1 - \nu_2, S_3)$. Hence there is a constant $C > 0$ such that

$$T(n) \leq T(\lceil \frac{3n}{4} \rceil) + T(\lceil \frac{n}{5} \rceil) + Cn.$$

An easy induction argument yields that $T(n) = O(n)$. □

We remark that in most UNIX–systems one can find the quick-sort algorithm.

2 Triangulations

We will in this chapter be concerned with planar triangulations.

2.1 General Properties

In two dimensions, a **triangulation** $\mathcal{T}$ is a finite collection of closed triangles $\{T_i\}$ such that each T_i has a non-empty interior and the T_i:s have pairwise disjoint interiors. Moreover we also require that if two distinct triangles have more than one point in common the must share a common edge.

figure

We will say that $\mathcal{T}$ is a triangulation of the set E if

$$E = \bigcup_{T \in \mathcal{T}} T.$$

Important problems are the existence of triangulations with prescribed properties and efficient algorithms for their construction. A general reference for this section is [45].

Theorem 5 *Let Ω be a closed simply connected polygonal domain and let V_Ω denote its set of vertices. Assume $W \subset \mathrm{Int}(\Omega)$ is a finite set. Then there exists a triangulation $\mathcal{T}$ of Ω whose set of vertices equals $V_\Omega \bigcup W$.*

figure

A **chord** or a **diagonal** of a polygon is a line segment between two vertices that lies inside the polygon and does not intersect the polygon's boundary except at the end points.

Lemma 2 *Every polygon with more than three sides has a chord.*

Proof: (Theorem 5) We first see that there is a triangulation of Ω whose vertex set equals V_Ω. For once we have found a chord we can split the polygon in two and recursively triangulate each part. Next we add one point at a time. If P lies inside a triangle we split that triangle in three.

figure

If P falls on a common edge of two triangles we split each of them in two triangles.

figure

□

Every triangulation $\mathcal{T}$ gives rise to a graph on the vertex set of $\mathcal{T}$ by identifying edges of $\mathcal{T}$ with the edges of the graph. This graph is called the **skeleton** or **triangulation graph** of the triangulation.

Finally we remark that the number of edges in a triangulation of a polygonal domain Ω is uniquely determined by the number of the vertices at the boundary and the number of the inner vertices. More precisely, we have the following result which easily follows by applying Euler's formula.

Theorem 6 *Let Ω be a closed simply connected polygonal domain and let V_Ω denote its set of vertices. Assume $W \subset \mathrm{Int}(\Omega)$ is a finite set. Let $\mathcal{T}$ be any triangulation of Ω whose set of vertices equals $V_\Omega \bigcup W$. Then*[2]

[2]For a finite set U we let $|U|$ denote the number of elements in U.

1. *the number of triangles in $\mathcal{T}$ is equal to*

$$|V_\Omega| + 2|W| - 2.$$

2. *the number of edges in $\mathcal{T}$, diagonals and boundary edges, is equal to*

$$2|V_\Omega| + 3|W| - 3.$$

How fast can we triangulate a polygon with N vertices? Chazelle [9] has established a linear time algorithm for this.

2.2 Transformations of Triangulations

Let $\mathcal{T}$ be a triangulation and assume that $T_1, T_2 \in \mathcal{T}$ share a common edge and that $R = T_1 \bigcup T_2$ is a convex quadrilateral. By replacing the diagonal of R with the opposite diagonal we get a new triangulation with the same set of vertices.

figure

This transformation is called **swapping of diagonals or edges**.

Suppose now that $\mathcal{T}_1$ and $\mathcal{T}_2$ are two triangulations of a polygon with the same vertex sets. It is natural to ask if it is possible to transform $\mathcal{T}_1$ into $\mathcal{T}_2$ by only using diagonal swaps of convex quadrilaterals. The following result by Dyn/Goren/Rippa [18] answers this question.

Theorem 7 *Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two triangulations of a polygonal domain Ω. Assume $\mathcal{T}_1$ and $\mathcal{T}_2$ have the same set of vertices. Then it is possible to transform $\mathcal{T}_1$ into $\mathcal{T}_2$ by only swapping diagonals of convex quadrilaterals.*

We begin with the following easy lemma.

Lemma 3 *Let $\gamma = \{P_0, \dots, P_{N-1}\}$ be a simple closed polygon. Then γ has at least three interior angles strictly less that π.*

PROOF: Let $\alpha_0, \dots, \alpha_{N-1}$ denote the interior angles of γ. Then

$$\Sigma_{i=0}^{N-1} \alpha_i = (N-2)\pi.$$

Let M be the number of α_i:s that are strictly less than π. Then

$$(N-2)\pi \geq (N-M)\pi + \Sigma_{\alpha_i < \pi} \alpha_i > (N-M)\pi,$$

so $M > 2$. □

Lemma 4 *Let $\mathcal{T}$ be a triangulation of the closed set Ω. Denote by V the set of vertices of $\mathcal{T}$ and let $A, B \in V$ be two distinct vertices. Assume $(A,B) \subset \mathrm{Int}(\Omega)$ and $V \bigcap (A,B) = \emptyset$. Then there exists a triangulation $\mathcal{T}^*$ of Ω whose vertex set is V and such that $[A,B]$ is an edge of $\mathcal{T}^*$. The transformation from $\mathcal{T}$ to $\mathcal{T}^*$ can be achieved by swapping diagonals of convex quadrilaterals. Moreover only edges that intersect (A,B) are swapped during the transformation.*

PROOF: Assume that $[A,B]$ is not an edge[3] of $\mathcal{T}$. Set $\mathcal{F}_{\mathcal{T}} = \{T \in \mathcal{T} : \mathrm{Int}(T) \bigcap (A,B) \neq \emptyset\}$ and $W = \bigcup_{T \in \mathcal{F}_{\mathcal{T}}} T$. Then $[A,B] \subset W$. Since every point $P \in W$ can be connected to $[A,B]$ inside the triangle $T \in \mathcal{F}_{\mathcal{T}}$ that contains P we see that W is connected.

Let L be the line through A and B. Clearly $L \bigcap \partial W = \{A, B\}$ and $(L \setminus [A,B]) \bigcap W = \emptyset$. L separates the plane into two open half planes H_+ and H_-.

We now claim that W is simply connected. Let $\mathcal{F}_{\mathcal{T}} = \{F_i\}$ and notice that the closed intervals $J_k = T_k \bigcap [A,B]$ have pairwise disjoint interiors. We may assume that the ordering of the T_k:s has been chosen so that $\mathrm{dist}(A, J_k)$ increases with k. Set $W_n = \bigcup_{k=1}^{n} T_k$. The above argument shows that W_n is connected for every n. We also observe that $\partial T_{n+1} \setminus W_n$ is connected. Hence W_{n+1} is simply connected if W_n is simply connected. Therefore W is simply connected by induction.

Set $S = \partial W \bigcap V$. From the above follows that $W_+ = W \bigcap H_+$ is a polygon with vertices $(S \bigcap H_+) \bigcup \{A, B\}$.

The idea of the proof is now to construct a new triangulation $\mathcal{T}'$ such that $\mathcal{F}_{\mathcal{T}'}$ has fewer triangles than $\mathcal{F}_{\mathcal{T}}$.

Using Lemma 3 on the polygon W_+ we see that there is a $p \in S \bigcap H_+$ such that the interior angle to W at p is strictly less than π. Let $\mathcal{F}_{\mathcal{T}}(p)$ be

[3] We have adopted the convention that $[A,B]$ denotes the closed segment AB, i.e. with the endpoints A, B included, while (A,B) denotes the open segment $[A,B] \setminus \{A,B\}$.

the set of triangles $T \in \mathcal{F}_\mathcal{T}$ that have a vertex at p. Let $\mathcal{E}_\mathcal{T}(p)$ be the class of edges e in $\mathcal{T}$ such that $e \bigcap (A, B) \neq \emptyset$ and have end point at p. We have that $\mathcal{E}_\mathcal{T}(p) \neq \emptyset$ and if $e \in \mathcal{E}_\mathcal{T}(p)$ then e is the common edge of two triangles in $\mathcal{F}_\mathcal{T}(p)$.

Set $\mathcal{E}_\mathcal{T}(p) = \{e_1, \dots, e_n\}$ and let q_i be the end point of e_i that is different from p. Notice that it may happen that $q_i \in \mathrm{Int}(W)$. We now give ∂W the counterclockwise ordering. Let p_0, p_1 be the elements in S that are immediately before and after p. We now assume that $q_1, \dots, q_n$ are listed in counterclockwise order as end points of rays emanating from p. Since each q_i belongs to H_- and is contained inside the sector determined by p_0, p and p_1 it follows that $\{q_1, \dots, q_n, p_0, p_1\}$ is a simple polygon Q.

Using Lemma 3 on Q we see that there is a k, $1 \leq k \leq n$, such that the interior angle of Q at q_k is strictly less than π. Let T_1 and T_2 be the triangles in $\mathcal{F}_\mathcal{T}(p)$ with the common edge e_k and set $R = T_1 \bigcup T_2$. Then the interior angle of R at p is strictly less than π. If $1 < k < n$ then R is contained in the sector determined by q_{k-1}, q_k and q_{k+1}. Hence the interior angle of R at q_k is strictly less than π in this case. If k equals 1 then one of the triangles T_1, T_2 must equal the triangle pp_1q_1. Hence in this case R is contained in the sector determined by p_1, q_1 and q_2. If $k = n$ then R is contained in the sector determined by q_{n-1}, q_n and p_0. Therefore the interior angle of R at q_k is strictly less than π. Hence R is a convex quadrilateral. Swapping diagonals in R gives us a new triangulation $\mathcal{T}^*$ such that $\mathcal{E}_{\mathcal{T}^*}(p)$ has one edge less than $\mathcal{E}_\mathcal{T}(p)$. Repeating this construction leads to a triangulation $\mathcal{T}'$ with $\mathcal{E}_{\mathcal{T}'}(p) = \emptyset$. Hence $\mathcal{F}_{\mathcal{T}'}$ has fewer triangles than $\mathcal{F}_\mathcal{T}$.

Iterating the above procedure leads to a triangulation $\mathcal{T}_1$ with $\mathcal{F}_{\mathcal{T}_1} = \emptyset$, i.e. $[A, B]$ is an edge in $\mathcal{T}_1$. The lemma is proved. □

PROOF: (Theorem 7) Let $e_1, \dots, e_n$ be the edges in $\mathcal{T}_2$ that do not belong to the boundary of Ω. Let $\mathcal{T}^{(1)}$ be the triangulation one gets by using Lemma 4 to transform $\mathcal{T}_1$ into a triangulation that contains e_1. For $2 \leq k \leq n$ we let $\mathcal{T}^{(k)}$ be the triangulation one gets by using Lemma 4 to transform $\mathcal{T}^{(k-1)}$ into a triangulation that contains e_k. Since e_j and e_k do not intersect if $1 \leq j < k$ it follows from Lemma 4 that $e_1, \dots, e_{k-1}$ belong to $\mathcal{T}^{(k)}$, $2 \leq k \leq n$. Hence $\mathcal{T}^{(n)}$ equals $\mathcal{T}_2$ which yields the theorem. □

3 Convex Hulls

We complete O'Rourke's presentation of convex hull algorithms, see [41], with the algorithm of Kirkpatrik-Seidel which connects several concepts treated above.

3.1 Kirkpatrik-Seidel's Ultimate Convex Hull Algorithm

This algorithm was published in 1986 in [30].

The basic idea is to develop a divide and conquer algorithm that runs in $O(n \log h)$ time, where h is the size of the hull e.g. the number of extreme edges. For simplicity, we will only concentrate on finding the upper chain of the convex hull, which connects the leftmost point with the rightmost point.

We will divide the points in P into two sets P_L and P_R of approximately equal size by using a vertical separating line. A few key observations are in order.

1. We can compute the upper supporting line, the **bridge**, *before* we compute the convex hulls of P_L and P_R, since these are functions of the point set, and not of the convex hulls of P_L and P_R. (Of course, the computation will be a little trickier since we cannot use the convex hull to find the bridge, as we did in the regular divide and conquer algorithm.)

2. We can delete all points that are immediately "below" the bridge. Specifically, if the bridge joins points p and q, we can delete all points with x coordinates between p and q. We recurse on the remaining points to find the upper chain of P'_L and P'_R (remaining sets of points). This will potentially reduce the number of points we recurse on.

There are a few things we need to do to implement the above ideas. The first is that we cannot afford to presort the points before starting the algorithm, so that for each divide step we have to invoke a linear time median finding algorithm to split S into equal pieces.

The second challenge is to find the bridge in $O(n)$ time. This will be done by a prune and search argument. In this method, we spend cn time, and we either find what we are looking for, or (for example) delete half the points from consideration, and recurse. The running time for such a scheme is

$$cn + c\frac{n}{2} + c\frac{n}{4} \ldots = O(n).$$

Observe that this actually works, as long as we eliminate a constant fraction of the points!

Before we get into details about how the bridge is actually computed, we will try to analyse the running time as a function of h, the size of the upper chain of the hull. If $h = 1$ in any recursive call, we compute a bridge and realize that the bridge connects the leftmost and rightmost points, and all other points can be eliminated immediately. Hence this takes $O(n)$ time.

If the upper chain of P'_R has size h_R and the upper chain of P'_L has size h_L, then we have that $h = 1 + h_R + h_L$. The running time of the algorithm can be expressed as $T(n, h)$.

$$T(n,h) \leq cn \quad \text{if} \quad h = 1$$
$$T(n,h) \leq cn + T\Big(\frac{n}{2}, h_L\Big) + T\Big(\frac{n}{2}, h_R\Big).$$

We need to argue that regardless of the split of h between h_L and h_R, the algorithm runs in $O(n \log h)$ time. This is easy to prove by induction. Assume $T(m, h') \leq cm \log h'$ for $m < n$ and $1 < h' < h$. This implies

$$T(n,h) \leq cn + c\frac{n}{2}\log h_L + c\frac{n}{2}\log h_R.$$

Now use the fact that $\frac{1}{2}(\log h_L + \log h_R) \leq \log \frac{h_L+h_R}{2}$, i.e. that the log-function is a concave function. This gives

$$T(n,h) \leq cn + c\frac{n}{2}\Big(2\log\Big(\frac{h_L+h_R}{2}\Big)\Big) \leq cn + cn(\log h - 1) \leq cn \log h.$$

It remains to indicate how to compute the **bridge** in $O(n)$ time.

Assume that the bridge has slope K^*. We will "guess" a value for the slope K, and try to figure out how K^* compares to K. Through each point p_i in P, let the line of slope K be ℓ_i. The line ℓ_j which has all the points of P "below" it is the line we are looking for – this line has the highest intercept with the dividing line (or the y axis, since they are all parallel). In $O(n)$ time we can compute which line has this property. (This line may not be unique.) If the line is defined by a point from P_L and a point from P_R this is the bridge we are looking for. If the line is defined *only* by points of P_L then $K^* < K$, and if the line is defined *only* by points of P_R then $K^* > K$.

Figure

In any case, this gives us some information about K^*. To apply the "prune and search" paradigm, we need to somehow discard a constant fraction of the points to continue looking for the bridge. We might "guess" the correct slope, and as soon as we do that the algorithm stops. If our guess is incorrect, we recurse on a constant fraction of the points. The running time is a decreasing geometric series and converges to $O(n)$.

We now outline the method for the pruning step. Consider a pair of points $a, b \in P$. Assume that a is to the left of b. Let the slope defined by this pair of points be K'. Suppose $K^* < K$ and $K \leq K'$, then $K^* < K'$. In this case we claim that a cannot be a bridge vertex. Suppose $K^* > K$ and $K \geq K'$, then $K^* > K'$. In this case we claim that b cannot be a bridge vertex. In either case, we are able to eliminate points as being candidate vertices.

How do we guarantee that we can find sufficiently many lines that have slope less than K, and sufficiently many lines with slope more than K? To achieve this, we arbitrarily pair up the points in set P, and compute their slopes. We then select the median slope of this set of $\frac{n}{2}$ lines, and use that as our guess for the slope of the bridge. We either guess the slope correctly and stop, or in either case we eliminate one point from $\frac{n}{4}$ pair's. We are left with at most $\frac{3n}{4}$ points and recurse on this set of points. This completes the prune and search algorithm.

Figure

4 Delaunay Triangulations/Voronoi Diagrams

Delaunay triangulations are named after the Russian mathematician Boris Delone. See [17] written in memory of George Voronoi.

4.1 Delaunay Triangulation: Existence and Uniqueness

Let $S \subset \mathbf{R}^2$ be a finite set of points that does not lie on a line. We say that $\mathcal{T}$ is a **Delaunay triangulation** if $\mathcal{T}$ is a triangulation of the convex hull of S with vertex set S and the circumcircles of the triangles in $\mathcal{T}$ newer contain a point from S in their interiors.

We shall now prove the existence of Delaunay triangulations. We will do this by taking the convex hull of the lift of S to the graph of $f(x,y) = x^2 + y^2$. The crucial property of the graph of f that we will need is the following observation.

Proposition 1 *Let $f(x,y) = x^2 + y^2$ and $g(x,y) = \langle (x,y), v \rangle + a$, $v \in \mathbf{R}^2$ and $a \in \mathbf{R}$. Then $f(x,y) < g(x,y)$ if and only if $|(x,y) - \frac{v}{2}|^2 < L$, where $L = a + \frac{|v|^2}{4}$.*

PROOF: Trivial. □

The significance of this result is that the set of points in $\mathbf{R}^3$ where the graph of f is below a plane always projects to the interior of a circle in $\mathbf{R}^2$.

Theorem 8 (Existence) *Let $S \subset \mathbf{R}^2$ be a finite set that does not lie on a line. Then S has a Delaunay triangulation.*

PROOF: Set $f(x,y) = x^2 + y^2$ and put $S_f = \{(x, y, f(x,y)) : (x,y) \in S\}$. Let now E be the convex hull of S_f in $\mathbf{R}^3$. The lower part of the boundary of E is the union of closed convex polygons F_i. For each F_i there is a $v_i \in \mathbf{R}^2$ and an $a_i \in \mathbf{R}$ such that if $g_i(x,y) = \langle (x,y), v_i \rangle + a_i$ then $f(x,y) \leq g_i(x,y)$ whenever $(x,y,z) \in F_i$ for some $z \in \mathbf{R}$. If $(x,y,z) \in F_i \bigcap S_f$ then $f(x,y) = g_i(x,y)$. Let G_i be the projection of F_i onto $\mathbf{R}^2$. From Proposition 1 follows that $G_i \bigcap S$ is contained in a circle C_i. Also C_i contains no point from S in its interior. Triangulating each G_i now yields the existence of a Delaunay triangulation. □

We shall next show how the edge-swapping transformation described in the previous section can be adapted to construct a Delaunay triangulation. The algorithm starts with any triangulation $\mathcal{T}$ of the convex hull of S. For a triangle T we let $C(T)$ and $\omega(T)$ denote the circumcircle of T and the closure of its interior. If e is an interior edge of $\mathcal{T}$ we let $R(e)$ denote the union of the triangles in $\mathcal{T}$ with the common edge e. If there is an interior edge e such that $R(e)$ is strictly convex and some vertex of $R(e)$ is contained in the interior of $C(T)$ for some triangle $T \in \mathcal{T}$ with edge e as an edge then we swap the edge. We call this the **flip algorithm**. We refer to [4] for the following result.

Theorem 9 *The flip algorithm terminates at a Delaunay triangulation.*

For the proof of the theorem we need the following results.

Lemma 5 *Let $T_1 = ABC$ and $T_2 = BCD$ be two triangles with disjoint interiors and a common edge BC. Assume that the quadrilateral $R = T_1 \bigcup T_2$ is not strictly convex. Then D lies outside $C(T_1)$ and A lies outside $C(T_2)$.*

figure

PROOF: From the assumption follows that R has at least one interior angle $\geq \pi$. Since the interior angles of R at A and D are strictly less that π we may assume that the interior angle at B is $\geq \pi$. We will argue by contradiction. Assume that D fall inside $C(T_1)$. Let D_1 be the intersection between $C(T_1)$ and the ray emanating from B that goes through D. Then the polygon ABD_1C has all its vertices on $C(T_1)$ and is therefore strictly convex. This contradicts the fact that $ABDC$ and ABD_1C have the same interior angle at B. By symmetry we are done. $\square$

Proposition 2 *Suppose T_1, T_2 are two triangles with disjoint interiors and a common edge e. Let A_1, A_2 be the vertices of T_1, T_2 that are not on e. Assume that A_1 falls outside $C(T_2)$ and that A_2 falls outside $C(T_1)$. Let L*

be the line through e. Assume that L separates the point P from T_1 and that P falls inside $C(T_1)$. Then P falls inside $C(T_2)$.

PROOF: We need only to treat the case when P does not belong to L. Let H be the half plane determined by L that contains P. By elementary geometry we have that $\omega(T_1) \bigcap H \subset \omega(T_2) \bigcap H$. □

PROOF: (Proof of Theorem 9) Let $\mathcal{E}(\mathcal{T})$ be the set of interior edges e such that if P is any vertex of $R(e)$ and if $T \in \mathcal{T}$ is any triangle with e as an edge then P belongs to the closure of the exterior of $C(T)$. It follows that $\mathcal{E}(\mathcal{T})$ increases by one element with every step of the flip algorithm. Hence the flip algorithm terminates.

We now let $\mathcal{T}$ be a triangulation that is left unchanged by the flip algorithm. Let $T_1, T_2 \in \mathcal{T}$ be two triangles with a common edge e. Let A_1, A_2 be the vertices of T_1, T_2 that do not lie on e. Since the flip algorithm has terminated it follows that A_1 is not inside $C(T_2)$ and A_2 is not inside $C(T_1)$.

Assume that $\mathcal{T}$ is not a Delaunay triangulation. Then there exists a $T \in \mathcal{T}$ and a vertex P for $\mathcal{T}$ such that P is in the interior of $C(T)$. Clearly $P \notin T$. Then there is an edge e of T such that if L denotes the line through e then L separates T from P. We notice that P can not lie on L. Let A be one of the end points of e. We may assume that (A, P) does not contain any vertex of $\mathcal{T}$ for otherwise we replace P by the vertex of $\mathcal{T}$ that is closest to A.

Let $T^* \in \mathcal{T}$ be the triangle different from T that also has e as an edge. Then T^* and P lie on the same side of L. From Proposition 2 follows that P is contained in the interior of $C(T^*)$.

We begin by assuming that $\text{Int}(T^*) \bigcap (A, P) \neq \emptyset$. In this case $[A, P]$ is not an edge for $\mathcal{T}$. Let $T_1, \ldots, T_n$ be the triangles in $\mathcal{T}$ that intersect (A, P). We let them be ordered such that $\text{dist}(T_k \bigcap (A, P), A)$ increases with k. Then $T_1 = T^*$. Let L_k be the line through the common edge of T_k and T_{k+1}. We observe that L_k separates P from T_k. By repeated application of Proposition 2 we see that P is in the interior of $C(T_k)$ for all k, $1 \leq k \leq n$. This contradicts the fact that $P \in T_n$.

Let W be the closed convex sector with apex A determined by e and $[A, P]$. Let $T_1, \ldots, T_n$ be the triangles in $\mathcal{T}$ with apex at A that are contained in W. We determine the ordering by requiring that $T_1 = T^*$ with T_k and T_{k+1} having a common edge e_k. Let L_k be the line determined by e_k. Then L_k separates P from T_k so P must be contained in the interior of $C(T_k)$

for all k, $1 \leq k \leq n$. Let e_n be the edge of T_n with an end in A that is different from e_{n-1}. Let $T_{n+1} \in \mathcal{T}$ be the triangle different from T_n that has e_n as an edge. Since P belongs to the interior of $C(T_n)$ it follows that $\text{Int}(T_{n+1}) \bigcap (A, P) \neq \emptyset$. This is impossible by our previous argument. This completes the proof of the theorem. □

Theorem 10 *Let $S \subset \mathbf{R}^2$ be a finite set that does not lie on a line. Suppose $a, b \in S$, $a \neq b$, and that there is a circle C such that $a, b \in C$ and there is no point from S in the interior of C. Then there is a Delaunay triangulation $\mathcal{T}$ of S such that $[a, b]$ is an edge in $\mathcal{T}$. Assume that $a, b, c \in S$ are the vertices of the triangle T. If $C(T)$ contains no points from S in its interior then T is a triangle in some Delaunay triangulation of S that can be constructed as in Theorem 8.*

Proof: We begin by establishing the second part of the theorem. Let P_a, P_b and P_c be the lifts of a, b and c to the graph of the function $f(x, y) = x^2 + y^2$. Then there is a unique plane H passing through P_a, P_b and P_c. From Proposition 1 follows that the lift S_f of S to the graph of f lies on one side of H. Hence T belongs to some Delaunay triangulation of S.

We can now prove the first part of the theorem by observing that the circle C can be modified so as to contain at least three points from S without having any point of S in its interior. □

Theorem 11 (Uniqueness) *Let $S \subset \mathbf{R}^2$ be a finite set that does not lie on a line. Assume that no four points in S belong to a circle.Then S has a unique Delaunay triangulation.*

Proof: Set $f(x, y) = x^2 + y^2$ and let S_f be the lift of S to the graph of f. The lower part of the boundary of the convex hull of S_f is the union of convex polygons F_i. From Proposition 1 follows that each F_i has exactly three vertices. Hence there is a unique Delaunay triangulation $\mathcal{T}^*$ constructed by Theorem 8. If $\mathcal{T}$ is any Delaunay triangulation of S then every triangle from $\mathcal{T}$ belongs to $\mathcal{T}^*$ by the previous theorem. □

We remark that by a similar argument one can show that any Delaunay triangulation is of the form given by Theorem 8. We leave the verification of this to the reader.

4.2 Delaunay Triangulation from a Stereographic Transformation

We shall in this section indicate how to construct Delaunay triangulations by stereographic transformations onto a sphere. A reference for this result is [7].

Let

$$M(\mathbf{x}) = \frac{\mathbf{x}}{|\mathbf{x}|^2}, \quad \mathbf{x} = (x, y, z) \in \mathbf{R}^3 \setminus \{0\}.$$

It is easily seen that M maps the plane $H = \{(x, y, -1) : (x, y) \in \mathbf{R}^2\}$ onto $\Omega \setminus \{0\}$, where Ω is the sphere

$$\Omega = \{\mathbf{x} \in \mathbf{R}^3 : |\mathbf{x} - (0, 0, -\frac{1}{2})| = \frac{1}{2}\}.$$

We also remark that M maps circles and lines in H onto circles in Ω and conversely.

We shall now construct a triangulation using the following procedure.

Let $S \subset H$ be a finite set. Let E_S denote the convex hull of the set $M(S) \subset \Omega \setminus \{0\}$. Then E_S is a convex polyhedron. We now let E_S^+ denote the class of faces $\{F_i\}$ of E_S such that 0 and E_S lie on the same side of the plane containing F_i.

We denote by V_i the vertex set of F_i, i.e. $V_i = F_i \bigcap M(S)$. Hence $M(V_i)$ is contained in a circle. Choose now a triangulation of the convex hull of $M(V_i)$. We let $\mathcal{T}_M$ be the collection of triangles one gets by letting F_i vary over all faces in E_S^+ and making one choice of the triangulation of the convex hull of $M(V_i)$. We call $\mathcal{T}_M$ a **triangulation** obtained by **stereographic projection**.

Theorem 12 *Let $S \subset \{(x, y, -1) : (x, y) \in \mathbf{R}^2\}$ be a finite set that does not lie on a line. Then every triangulation one gets by stereographic projection is a Delaunay triangulation.*

PROOF: Define

$$\tilde{M}(\mathbf{x}) = -\frac{\mathbf{x}}{z}, \quad \mathbf{x} = (x, y, z) \in \tilde{\Omega} \equiv \{(x, y, z) \in \mathbf{R}^3 : -1 \leq z < 0\}.$$

Notice that $M^{-1}(\mathbf{x}) = \tilde{M}(\mathbf{x})$ for all $\mathbf{x} \in \Omega \setminus \{0\}$. Since $\tilde{M}$ is one–to–one on the lower part E_S^+ of E_S and $\tilde{M}(F_i)$ equals the convex hull of $M(V_i)$ we see that $\mathcal{T}_M$ is a triangulation of the convex hull of S. From the mapping properties of M it follows that for all $T \in \mathcal{T}_M$ the intersection between S and the interior of the circumcircle of T is empty. □

4.3 Delaunay Triangulations minimize the Dirichlet Integral

We shall in this section study the Dirichlet integral of piecewise linear functions. Let $\mathcal{T}$ be a triangulation with vertex set S. Assume $f : S \to \mathbf{R}$ and let $f_{\mathcal{T}}$ denote the piecewise linear extension of f to $\bigcup_{T\in\mathcal{T}}$ that is linear on each $T \in \mathcal{T}$. We let

$$\mathcal{D}(f, \mathcal{T}) = \Sigma_{T\in\mathcal{T}} \int\int_T |\nabla f_{\mathcal{T}}|^2 \, dxdy$$

denote the **Dirichlet integral of** f.

We shall now establish the following remarkable extremal property of Delaunay triangulation due to Rippa [47].

Theorem 13 *Let $\mathcal{T}$ be a Delaunay triangulation of the set V. Let $f : V \to \mathbf{R}$ be any function. Assume that $\mathcal{T}^*$ is a triangulation of the convex hull of V with vertex set V. Then*

$$\mathcal{D}(f, \mathcal{T}) \leq \mathcal{D}(f, \mathcal{T}^*).$$

If $\mathcal{D}(f, \mathcal{T}) = \mathcal{D}(f, \mathcal{T}^*)$ the structure of f and the relationship between $\mathcal{T}$ and $\mathcal{T}^*$ can be analyzed. For an example of this see the proof of Theorem 17.

To give the proof of the theorem we need some lemmas. We give the following classical result for the convenience of the reader.

Lemma 6 (Ptolemy's First Theorem) *Suppose $ABCD$ is a quadrilateral inscribed in a circle. Then*

$$|[A, C]| \cdot |[B, D]| = |[A, B]| \cdot |[C, D]| + |[A, D]| \cdot |[B, C]|.$$

PROOF: We will use the following identity by Euler. Let $a, b, c, d \in \mathbf{C}$. Then

$$(d - a)(b - c) + (d - b)(c - a) + (d - c)(a - b) = 0.$$

We remark that if a, b, c, d lie on a line and are sequentially ordered then

$$|a - c| \cdot |b - d| = |a - d| \cdot |b - c| + |a - b| \cdot |c - d|.$$

We may assume that the circumcircle of the quadrilateral goes through the origin and that the origin is outside the quadrilateral. Setting

$$a = \frac{1}{A},\ b = \frac{1}{B},\ c = \frac{1}{C},\ d = \frac{1}{D}$$

yields the lemma. □

Lemma 7 *Suppose $ABCD$ is a strictly convex quadrilateral. Then*

$$|[A,C]| \cdot |[B,D]| \leq |[A,B]| \cdot |[C,D]| + |[A,D]| \cdot |[B,C]|$$

with equality if and only if the vertices fall on a circle.

PROOF: Assume that C lies outside the closed disk Δ, whose boundary contains A, B and D. Let C' be the intersection between $\partial\Delta$ and $[A,C]$.

figure

Select a coordinate system such that C' falls at the origin and C falls on the positive x-axis. For $x > 0$, let $P = (x, 0)$ and set

$$f(x) = |[A,B]| \cdot |[C,D]| + |[A,D]| \cdot |[B,C]| - |[A,C]| \cdot |[B,D]|.$$

Let θ and θ' be the angles APD and $AC'D$. Also let ϕ and ϕ' be the angles APB and $AC'B$. Then

$$0 < \theta < \theta' < \frac{\pi}{2}, \quad 0 < \phi < \phi' < \frac{\pi}{2}.$$

Now if $x > 0$ then

$$\frac{df}{dx} = |[A,B]| \cos\theta + |[A,D]| \cos\phi - |[B,D]|$$

$$> |[A,B]| \cos\theta' + |[A,D]| \cos\phi' - |[B,D]| = 0$$

since the angles ABD and ADB equal θ' and ϕ' respectively. Since $f(0) = 0$ the lemma follows. □

Lemma 8 *Let T be a triangle with vertices A, B, C and denote by A_m, C_m the midpoints of $[A, B]$ and $[B, C]$. Let ω_A and ω_C be the closed disks centered at A_m and C_m and passing through B. Denote by S the closed sector with apex B determined by T. Then*

$$\omega_A \bigcap \omega_C \bigcap S \subset T.$$

PROOF: We note that $\partial\omega_A \bigcap \partial\omega_C$ consists of exactly two points B and Q. Let Q_m be the midpoint of $[B, Q]$. Let L_m be the line through A_m and C_m.

figure

Then L_m intersects $[B, Q]$ perpendicularly at Q_m. If L denotes the line through Q that is parallel to L_m then A and C lie on L. Let H be the closed half plane determined by L that contains B. Then

$$\omega_A \bigcap \omega_C \bigcap S \subset H.$$

Hence

$$\omega_A \bigcap \omega_C \bigcap S \subset H \bigcap S = T.$$

□

If Q is the vertex of a polygon Γ we let $\omega(Q)$ denote the closed disk whose boundary goes through Q and its immediate neighbors. Also, $R(Q)$ will denote the radius of $\omega(Q)$.

Lemma 9 *Let $ABCD$ be a strictly convex quadrilateral. Assume that $C \notin \omega(A)$. Then*

$$\min(R(A), R(C)) \leq \min(R(B), R(D)). \tag{1}$$

Proof: For a vertex P we let z_P denote the center of $\omega(P)$. For two points P and Q we let $\Delta_{P,Q}$ be the closed disk that has $[P,Q]$ as a diameter.

From the assumptions follows that $A \notin \omega(C)$ but $\omega(B)$ and $\omega(D)$ contain all vertices.

Assume that (1) is false, say

$$R(B) < \min(R(A), R(C)).$$

Let L be the line through $[A,B]$ and let $H = H_{A,B}$ be the closed half plane determined by L that contains C. Let H_- be the complement of H. We now claim that

$$H_{A,B} \bigcap \omega(A) \bigcap \omega(B) \subset \Delta_{A,B}. \tag{2}$$

If z_A or z_B belong to H_- then (2) is obvious. If z_A and z_B both were in $H_{A,B}$ then clearly $H_{A,B} \bigcap \omega(B) \subset H_{A,B} \bigcap \omega(A)$. However $C \in H_{A,B} \bigcap \omega(B)$ but $C \notin H_{A,B} \bigcap \omega(A)$ so this is impossible. Hence (2) follows.

Let $H_{B,C}$ be the closed half plane that contains A and has the line through $[B,C]$ as its boundary. By the same argument we have

$$H_{B,C} \bigcap \omega(B) \bigcap \omega(C) \subset \Delta_{B,C}. \tag{3}$$

Hence $D \in \Delta_{A,B} \bigcap \Delta_{B,C}$. From the Lemma 8 follows that D belongs to the triangle ABC, which violates the assumption that $ABCD$ is strictly convex. By this contradiction (3) follows. □

Lemma 10 *Let $ABCD$ be a strictly convex quadrilateral. Assume that $C \notin \omega(A)$. Then*

$$\max(R(A), R(C)) \leq \max(R(B), R(D)). \tag{4}$$

Proof: Notice that $A \notin \omega(C)$. We shall argue by contradiction. Assume that the conclusion fails with, say,

$$R(A) > \max(R(B), R(D)).$$

For a vertex P let z_P denote the center of $\omega(P)$. Let L be the line through $[A,B]$ and denote by H_B the closed half plane determined by L that contains C. We claim that $z_A \notin H_B$. For if $z_A \in H_B$ then $H_B \bigcap \omega(B) \subset H_B \bigcap \omega(A)$ by the assumption that $R(A) > R(B)$. This is impossible since $C \notin \omega(A)$. Hence $z_A \in H_B^-$, where H_B^- is the complement of H_B.

By the same argument it follows that $z_A \in H_D^-$, where H_D^- is the plane determined by the line through $[A,D]$ that does not contain C. Hence $z_A \in H_B^- \bigcap H_D^-$, which is impossible. □

Lemma 11 *Let f be linear in the triangle ABC. Assume $f(A) = 1$, $f(B) = f(C) = 0$. Then*

$$\int\int_{ABC} |\nabla f|^2 \, dxdy = \frac{|[B,C]|^2}{4|ABC|},$$

where $|ABC|$ is the area of the triangle.

PROOF: Letting h be the height of the triangle from the vertex A one has that $|\nabla f| = \frac{1}{h}$. Hence

$$\int\int_{ABC} |\nabla f|^2 \, dxdy = \frac{|ABC|}{h^2} = \frac{|[B,C]|}{2h} = \frac{|[B,C]|^2}{4|ABC|}.$$

□

Lemma 12 *Let $ABCD$ be a strictly convex quadrilateral. Let $\mathcal{T}_1$ be the triangulation of $ABCD$ with triangles ABC and ADC. Let $\mathcal{T}_2$ be the triangulation of $ABCD$ with triangles ABD and BDC. Let $S = \{A, B, C, D\}$ and suppose $f : S \to \mathbf{R}$ satisfies $f(B) = f(C) = f(D) = 0$.*

1. *If all vertices A, B, C and D fall on a circle then*

$$\mathcal{D}(f, \mathcal{T}_1) = \mathcal{D}(f, \mathcal{T}_2).$$

2. *If $C \notin \omega(A)$ and $R(A) \leq \min(R(B), R(D))$ then*

$$\mathcal{D}(f, \mathcal{T}_1) < \mathcal{D}(f, \mathcal{T}_2).$$

PROOF: Let $\alpha = \frac{1}{2} f(A)$. Then

$$\mathcal{D}(f, \mathcal{T}_1) = \alpha^2 (\frac{|[B,C]|^2}{|ABC|} + \frac{|[C,D]|^2}{|ACD|})$$

and

$$\mathcal{D}(f, \mathcal{T}_2) = \alpha^2 \frac{|[B,D]|^2}{|ABD|}.$$

By Lemma 7

$$|[A,C]| \cdot |[B,D]| \leq |[B,C]| \cdot |[A,D]| + |[A,B]| \cdot |[C,D]|.$$

Dividing by $|[A,B]| \cdot |[A,C]| \cdot |[A,D]|$ gives

$$\frac{|[B,D]|}{|[A,B]| \cdot |[A,D]|} \leq \frac{|[B,C]|}{|[A,B]| \cdot |[A,C]|} + \frac{|[C,D]|}{|[A,C]| \cdot |[A,D]|}$$

or

$$\frac{|[B,D]|^2}{|[A,B]|\cdot|[A,D]|\cdot|[B,D]|}$$

$$\leq \frac{|[B,C]|^2}{|[A,B]|\cdot|[A,C]|\cdot|[B,C]|} + \frac{|[C,D]|^2}{|[A,C]|\cdot|[A,D]|\cdot|[C,D]|}.$$

For any triangle it is well-known that the product of its sides equals the area times the radius of the circumcircle times 4. Hence

$$\frac{|[B,D]|^2}{|ABD|} \leq \frac{R(A)}{R(B)}\frac{|[B,C]|^2}{|ABC|} + \frac{R(A)}{R(D)}\frac{|[C,D]|^2}{|ACD|}$$

with equality if and only if the vertices fall on a circle. Consequently *1.* is immediate and for *2.* we notice

$$\frac{|[B,D]|^2}{|ABD|} < \frac{|[B,C]|^2}{|ABC|} + \frac{|[C,D]|^2}{|ACD|}$$

from the assumption $R(A) \leq \min(R(B), R(D))$. □

Lemma 13 *Let $ABCD$ be a strictly convex quadrilateral. Assume that the triangles ABD and BCD form a Delaunay triangulation $\mathcal{T}$. Let $\mathcal{T}^*$ be the triangulation ABC and ADC. Suppose that $f : \{A, B, C, D\} \to \mathbf{R}$. Then*

$$\mathcal{D}(f, \mathcal{T}) \leq \mathcal{D}(f, \mathcal{T}^*).$$

If $\mathcal{D}(f, \mathcal{T}) = \mathcal{D}(f, \mathcal{T}^*)$ then $f_{\mathcal{T}^*}$ is linear or the vertices fall on a circle.

PROOF: We may assume that $R(A) \leq R(C)$. Let $h : \mathbf{R}^2 \to \mathbf{R}$ be the linear function that equals f on $\{B, C, D\}$. Set $S = \{A, B, C, D\}$ and define $g : S \to \mathbf{R}$ by $g(A) = f(A) - h(A)$ and $g(B) = g(C) = g(D) = 0$. Let p and q denote the piecewise linear extensions of g determined by $\mathcal{T}$ and $\mathcal{T}^*$ respectively. Then

$$\mathcal{D}(g, \mathcal{T}) = \int\!\!\int_{ABCD} |\nabla p|^2 \, dxdy \leq \mathcal{D}(g, \mathcal{T}^*) = \int\!\!\int_{ABCD} |\nabla q|^2 \, dxdy$$

by the previous lemma. However

$$\begin{aligned}\mathcal{D}(f, \mathcal{T}) - \mathcal{D}(g, \mathcal{T}^*) &= \int\!\!\int_{ABCD} (|\nabla(p+h)|^2 - |\nabla(q+h)|^2) \, dxdy \\ &\leq 2\int\!\!\int_{ABCD} \langle \nabla(p-q), \nabla h\rangle \, dxdy \\ &= \int_{\partial(ABCD)} (p-q)\frac{\partial h}{\partial n} \, d\sigma - \int\!\!\int_{ABCD} (p-q)\Delta h \, dxdy = 0\end{aligned}$$

since h is harmonic in $\mathbf{R}^2$ and $p = q$ on the boundary of $ABCD$. □

PROOF: (Proof of Theorem 13) The theorem is a direct consequence of the flip algorithm. □

For a triangle T let $R(T)$ denote the radius of its circumcircle. For a triangulation $\mathcal{T}$ we now set

$$R_{\max}(\mathcal{T}) = \max\{R(T) : T \in \mathcal{T}\}$$

and

$$R_{\min}(\mathcal{T}) = \min\{R(T) : T \in \mathcal{T}\}.$$

We now have the interesting result that these functionals are minimized by Delaunay triangulations. For the functional $R_{\min}(\mathcal{T})$ see [3].

Theorem 14 *Let $\mathcal{T}$ be a Delaunay triangulation of the set S. Assume that $\mathcal{T}^*$ is a triangulation of the convex hull of S with vertex set S. Then*

$$R_{\min}(\mathcal{T}) \leq R_{\min}(\mathcal{T}^*) \tag{5}$$

and

$$R_{\max}(\mathcal{T}) \leq R_{\max}(\mathcal{T}^*). \tag{6}$$

PROOF: The theorem is a direct consequence of Lemma 9, Lemma 10 and the flip algorithm. □

For a triangulation $\mathcal{T}$ we let $R(\mathcal{T})$ denote the vector one gets by sorting the collection $\{R(T) : T \in \mathcal{T}\}$ in increasing order. Notice that $R_{\min}(\mathcal{T})$ and $R_{\max}(\mathcal{T})$ coincide with the first and last element in $R(\mathcal{T})$. We have the following sharpening of Theorem 14. If $\mathbf{x} = \{x_i\}_{i=1}^n$ and $\mathbf{y} = \{y_i\}_{i=1}^n$ are two sequences then we say that $\mathbf{x} \leq \mathbf{y}$ if $x_i \leq y_i$ for $1 \leq i \leq n$.

Theorem 15 *Let $\mathcal{T}$ be a Delaunay triangulation of the set S. Assume that $\mathcal{T}^*$ is a triangulation of the convex hull of S with vertex set S. Then*

$$R(\mathcal{T}) \leq R(\mathcal{T}^*). \tag{7}$$

PROOF: Suppose $\mathcal{T}^{**}$ is derived from $\mathcal{T}^*$ by an edge swap of a strictly convex quadrilateral. Since $\{R(T) : T \in \mathcal{T}^{**}\}$ is derived from $\{R(T) : T \in \mathcal{T}^*\}$ by decreasing two elements it follows that $\{R(T) : T \in \mathcal{T}^{**}\} \leq \{R(T) : T \in \mathcal{T}^*\}$. The theorem is a direct consequence of the flip algorithm. □

Suppose $\Phi : \mathbf{R}^n \to \mathbf{R}$ is **symmetric**, i.e. for every permutation π of $\{1, \ldots, n\}$ we have for all $\mathbf{x} \in \mathbf{R}^n$

$$\Phi(x_1, \ldots, x_n) = \Phi(x_{\pi(1)}, \ldots, x_{\pi(n)}).$$

If $\mathcal{T}$ is a triangulation and if Φ is symmetric we set

$$\Phi_R(\mathcal{T}) = \Phi(R(T_1), \ldots, R(T_n)),$$

where $T_1, \ldots, T_n$ are the triangles in $\mathcal{T}$. Clearly $\Phi_R(\mathcal{T})$ is well-defined. We now let $\mathcal{A}_{\text{sym}}$ denote the class of all symmetric functions Φ such that if $\mathbf{0} \leq \mathbf{x} \leq \mathbf{y}$ then $\Phi(\mathbf{x}) \leq \Phi(\mathbf{y})$.

Theorem 16 *Let $\mathcal{T}$ be a Delaunay triangulation of the set S. Assume that $\mathcal{T}^*$ is a triangulation of the convex hull of S with vertex set S. Then*

$$\Phi_R(\mathcal{T}) \leq \Phi_R(\mathcal{T}^*) \tag{8}$$

for all $\Phi \in \mathcal{A}_{\text{sym}}$.

PROOF: The result is an obvious consequence from Theorem 15. □

Notice that $\max\{x_i\}$, $\min\{x_i\}$, $\Sigma_i x_i$, $\Sigma_i x_i^2$, $\Pi_i x_i$ and $\sqrt[n]{\Pi_i x_i}$ all belong to the class $\mathcal{A}_{\text{sym}}$.

We shall now study the area of piecewise linear surfaces. Let $\mathcal{T}$ be a triangulation with vertex set S. Assume $f : S \to \mathbf{R}$ and let $f_{\mathcal{T}}$ denote the piecewise linear extension to $\bigcup_{T \in \mathcal{T}}$ that is linear on each $T \in \mathcal{T}$. Set

$$\mathcal{A}(f, \mathcal{T}) = \Sigma_{T \in \mathcal{T}} \int\int_T \sqrt{1 + |\nabla f_{\mathcal{T}}|^2}\, dxdy.$$

Notice that $\mathcal{A}(f, \mathcal{T})$ is the area of the graph of $f_{\mathcal{T}}$.

Theorem 17 *Let $\mathcal{T}$ be a Delaunay triangulation of the set S. Let $f : S \to \mathbf{R}$ be any function. Then there is a positive number $\epsilon_0 = \epsilon_0(f)$ such that if $|\epsilon| < \epsilon_0$ and if $\mathcal{T}^*$ is a triangulation of the convex hull of S with vertex set S then*

$$\mathcal{A}(\epsilon f, \mathcal{T}) \leq \mathcal{A}(\epsilon f, \mathcal{T}^*).$$

To prove the theorem we need the following lemma.

Lemma 14 *Let $\mathcal{T}, \mathcal{T}^*$ be as above. If $\mathcal{D}(f, \mathcal{T}) = \mathcal{D}(f, \mathcal{T}^*)$ then $\mathcal{A}(f, \mathcal{T}) = \mathcal{A}(f, \mathcal{T}^*)$.*

PROOF: From Lemma 13 it follows that if we carry out an edge swap in a strictly convex quadrilateral Q in the flip algorithm then $f_{\mathcal{T}^*}$ must be linear in Q. Hence the area is unchanged which establishes the lemma. □

PROOF: (Proof of Theorem 17) We will argue by contradiction. Assume the conclusion of the theorem fails. Since there are only finitely many triangulations with vertex set S there exists $\epsilon_j \to 0$ such that

$$\mathcal{A}(\epsilon_j f, \mathcal{T}) > \mathcal{A}(\epsilon_j f, \mathcal{T}^*).$$

Since

$$\mathcal{A}(\epsilon f, \mathcal{T}) - \mathcal{A}(\epsilon f, \mathcal{T}^*) = \frac{1}{2}\epsilon^2(\mathcal{D}(f, \mathcal{T}) - \mathcal{D}(f, \mathcal{T}^*)) + O(\epsilon^4),$$

it follows that $\mathcal{D}(f, \mathcal{T}) \geq \mathcal{D}(f, \mathcal{T}^*)$. However from Theorem 13 we have $\mathcal{D}(f, \mathcal{T}) = \mathcal{D}(f, \mathcal{T}^*)$. Hence $\mathcal{A}(\epsilon f, \mathcal{T}) = \mathcal{A}(\epsilon f, \mathcal{T}^*)$ by Lemma 14. □

An interesting problem here is to construct a polynomial time algorithm for finding piecewise linear interpolating surface of least area.

For a triangle T let $\theta_{\min}(T)$ denote the smallest angle of T. If $\mathcal{T}$ is a triangulation we set

$$\theta_{\min}(\mathcal{T}) = \min\{\theta_{\min}(T) : T \in \mathcal{T}\}.$$

Lemma 15 *Let $ABCD$ be a strictly convex quadrilateral. Assume that the triangles ABD and BCD form a Delaunay triangulation $\mathcal{T}$. Let $\mathcal{T}^*$ be the triangulation ABC and ADC. Then*

$$\theta_{\min}(\mathcal{T}) \geq \theta_{\min}(\mathcal{T}^*)$$

with equality if and only if the vertices all fall on a circle.

PROOF: Let $\{\alpha_i\}_{i=1}^4$ and $\{\beta_i\}_{i=1}^4$ be the angles in $\mathcal{T}$ and $\mathcal{T}^*$ that are pairwise opposite the same side.

figure

Let a, b, c, d denote the interior angles of the polygon at the vertices A, B, C, D. Notice that $a = \beta_2 + \beta_4$ so $a > \theta_{\min}(\mathcal{T}^*)$. Similarly $c > \theta_{\min}(\mathcal{T}^*)$. If the vertices fall on a circle then $\alpha_i = \beta_i$ for $1 \le i \le 4$. Hence

$$\theta_{\min}(\mathcal{T}) = \theta_{\min}(\mathcal{T}^*)$$

in this case.

Assume now that the vertices do not fall on a circle. Set $\theta^* = \theta_{\min}(\mathcal{T})$. If $\theta^* \in \{a, c\}$ then we are done by the above inequalities. If $\theta^* \in \{\alpha_1, \ldots, \alpha_4\}$, say $\theta^* = \alpha_1$, we have the following. Let C^* be the intersection between the circumcircle to ABD and the line segment $[A, C]$. Let $\{\beta_i^*\}_{i=1}^4$ be the angles in the polygon ABC^*D that correspond to the β_i:s. Then $\beta_1^* > \beta_1$, so $\theta^* > \beta_1 \ge \theta_{\min}(\mathcal{T}^*)$. □

Theorem 18 *Let $\mathcal{T}$ be a Delaunay triangulation of the set S. Assume that $\mathcal{T}^*$ is a triangulation of the convex hull of S with vertex set S. Then*

$$\theta_{\min}(\mathcal{T}) \ge \theta_{\min}(\mathcal{T}^*).$$

PROOF: The theorem is a direct consequence of Lemma 15 and the flip algorithm. □

A reference for the previous theorem is [34].

For a triangulation $\mathcal{T}$ we let $\Theta(\mathcal{T}) = \{\theta_i(\mathcal{T})\}$ denote the collection of all $\theta_{\min}(T)$, for triangles T in $\mathcal{T}$, sorted in increasing order. If $\mathcal{T}_1$ and $\mathcal{T}_2$ are two triangulations with the same number of triangles then we say that $\Theta(\mathcal{T}_2)$ is **lexicographically greater** than $\Theta(\mathcal{T}_1)$ if there is an integer m such that

$$\theta_j(\mathcal{T}_2) = \theta_j(\mathcal{T}_1) \ \text{ for } \ j < m$$

and

$$\theta_m(\mathcal{T}_2) > \theta_m(\mathcal{T}_1).$$

Corollary 1 *Let $\mathcal{T}, \mathcal{T}^*$ be two triangulations as in Theorem 18. If $\mathcal{T}^*$ is not a Delaunay triangulation then $\Theta(\mathcal{T})$ is lexicographically greater than $\Theta(\mathcal{T}^*)$.*

PROOF: It follows from Lemma 15 that the flip algorithm makes the angle sequence lexicographically greater. □

A reference for the previous result is [35].

4.4 Farthest Delaunay Triangulations

A farthest Delaunay triangulation is characterized by the property that the circumcircle of each triangle contains the vertex set of the triangulation. This is only possible if the vertices are the extreme points of their convex hull. More precisely we have the following definition.

Definition 1 *Assume that S is the set of vertices of a strictly convex polygon. Let $\mathcal{T}$ be a triangulation of* $\text{conv}(S)$ *with vertex set S. $\mathcal{T}$ is called a* **farthest Delaunay triangulation** *is $S \subset \Delta(T)$ for all $T \in \mathcal{T}$, where $\Delta(T)$ denotes the closed disk whose boundary contains the vertices of T.*

Let $f(x) = |x|^2$, $x \in \mathbf{R}^2$, and let $S_f = \{(p, f(p)) : p \in S\}$. Let $\{F_i\}$ be the faces of the upper part of the convex hull of S_f. Triangulating the projection of each of the faces F_i yields a farthest Delaunay triangulation.

Let R be a strictly convex quadrilateral and assume that the vertices of R do not lie on a circle. Then R has exactly two triangulations — one being the Delaunay triangulation and the other the farthest Delaunay triangulation.

We shall now describe the analogue of the flip algorithm for farthest Delaunay triangulations.

Let $\mathcal{T}$ be a triangulation and assume that e is the diagonal of a strictly convex quadrilateral R formed by two triangles in $\mathcal{T}$ with a common edge e. Assume that the vertices of R are not cocircular. The **reversed flip algorithm** consists of making a diagonal swap to obtain the farthest Delaunay triangulation of R. The flip algorithm consists of selecting the Delaunay triangulation of R.

Theorem 19 *Let S be the vertex set of a strictly convex polygon and let $\mathcal{T}$ be any triangulation of* $\text{conv}(S)$ *with vertex set S. Applying the reversed flip algorithm to $\mathcal{T}$ terminates in a farthest Delaunay triangulation of S.*

PROOF: Let $R(\mathcal{T})$ denote the radii of the circumcircles of the triangles in $\mathcal{T}$ sorted in increasing order. If $\mathcal{T}'$ is a triangulation obtained by applying the reversed flip algorithm then $R(\mathcal{T}) \le R(\mathcal{T}')$ and at least two components differ. Hence the algorithm must terminate.

We shall now prove that if the reversed flip algorithm has terminated with a triangulation $\mathcal{T}^*$ then $\mathcal{T}^*$ is a farthest Delaunay triangulation. Let N be the number of elements in S. The result is obvious if $N \le 4$. For T a triangle let $\Delta(T)$ be the closed disk whose boundary contains the vertices of T. Suppose that there exists a $T_0 \in \mathcal{T}^*$ such that for some $p \in S$ we have $p \notin \Delta(T_0)$. Let e be an edge of T_0 that separates p from T_0. Let H be the closed half plane determined by e that contains p. Set $S^* = S \bigcap H$ and $\mathcal{T}^{**} = \{T \in \mathcal{T}^* : T \subset H\}$. Since S^* has fewer elements than S it follows by the induction assumption that $\mathcal{T}^{**}$ is a farthest Delaunay triangulation of S^*. Let $T_1 \in \mathcal{T}^{**}$ be the triangle that has e as an edge. Then $R = T_0 \bigcup T_1$ is strictly convex. Since $p \in \Delta(T_1)$ but $p \notin \Delta(T_0)$ we have that the vertices of R are not on a circle. Also $T_0 \subset \Delta(T_1)$ and $\mathrm{Int}(T_0) \bigcap H = \emptyset$. This is impossible. □

Using our previous notation we have the following properties of farthest Delaunay triangulations.

Theorem 20 *Let S be the vertex set of a strictly convex polygon. Let $a, b \in S$, $a \neq b$, and assume there is a closed disk B such that $S \subset B$ and $a, b \in \partial B$. Then there is a farthest Delaunay triangulation $\mathcal{T}$ of S such that $[a, b]$ is an edge of $\mathcal{T}$. Assume there is a closed disk Δ such that $S \subset \Delta$ and $a, b, c \in \partial\Delta$. Then the triangle with vertices a, b, c is a triangle in some farthest Delaunay triangulation of S.*

Theorem 21 *Let S be the vertex set of a strictly convex polygon. Let $\mathcal{T}$ be a farthest Delaunay triangulation of S. Let $\mathcal{T}^*$ be any triangulation of* $\mathrm{conv}(S)$ *with vertex set S. Then the following holds:*

1. $R_{\max}(\mathcal{T}^*) \le R_{\max}(\mathcal{T})$.
2. $R_{\min}(\mathcal{T}^*) \le R_{\min}(\mathcal{T})$.
3. $R(\mathcal{T}^*) \le R(\mathcal{T})$.
4. *If $\Phi \in \mathcal{A}_{\mathrm{sym}}$ then*
$$\Phi_R(\mathcal{T}^*) \le \Phi_R(\mathcal{T}).$$
5. *If $f : S \to \mathbf{R}$ is any function then*
$$\mathcal{D}(f, \mathcal{T}^*) \le \mathcal{D}(f, \mathcal{T}).$$

6. $\theta_{\min}(\mathcal{T}) \leq \theta_{\min}(\mathcal{T}^*)$.

7. $\Theta(\mathcal{T}^*)$ *is lexicographically greater than* $\Theta(\mathcal{T})$.

4.5 General Properties of Voronoi Diagrams

The Voronoi diagram is a basic structure for the understanding of point sets in a metric space. It is a natural tool for analyzing the geometric distribution and internal relations of data points. Here is the general definition. Let S be a finite set of points in a metric space X with metric δ. Set

$$V_S(p) = \{x \in X : \delta(x,p) \leq \min\{\delta(x,q) : q \in S \setminus \{p\}\}\},$$

i.e. $V_S(p)$ is the set of points x that are not closer to any other point in S than to p. We call $V_S(p)$ the **Voronoi cell** associated to p.

Proposition 3 *The Voronoi cells have the following properties:*

1. $p \in V_S(p)$ *for all* $p \in S$

2. $\bigcup_{p\in S} V_S(p) = X$

3. *If* $x \in V_S(p) \bigcap V_S(q)$ *then* $\delta(x,p) = \delta(x,q)$.

We call the collection of all Voronoi cells the **Voronoi diagram** of S. As an indication of its usefulness, the Voronoi diagrams has independently been "discovered" in many fields, for a discussion of this see [2] and [40], in particular pp 6 – 10 in [40].

There are several natural extensions. For instance, we can associate to every subset $Y \subset S$, $Y \neq \emptyset$, the **generalized Voronoi cell** $V_S(Y)$ by letting

$$V_S(Y) = \{x \in X : \text{for all } p \in Y \text{ and all } q \in S \setminus Y \;\; \delta(x,p) \leq \delta(x,q)\}.$$

That is $V_S(Y)$ is the locus of all points x such that each point in Y is at least as close to x as to any point in $S \setminus Y$. Of course, it may happen that $V_S(Y)$ is empty. The Voronoi diagram is of **order** k if it is the collection of all (non-empty) generalized Voronoi cells of subsets of S with k elements.

Recently, there has been an interest in a still more general construction. Let $E_1, E_2, \ldots, E_N$ be N closed subsets of X and let

$$d_k(x) = \inf\{\delta(x,y) : y \in E_k\}.$$

We consider the cells defined by

$$\{x \in X : d_k(x) \leq \min\{d_j(x) : j \neq k\}\}.$$

For more about this see [40].

4.6 Classical Voronoi Diagrams in the Euclidean Plane

We shall in this section let $S \subset \mathbf{R}^2$ be a finite set of points that is not included in any line. In particular, S has at least three elements. In this case the Voronoi cell

$$V_S(p) = \{x \in \mathbf{R}^2 : |x - p| \leq \min\{|x - q| : q \in S \setminus \{p\}\}\}$$

is a possibly unbounded polygonal domain. Therefore $V_S(p)$ is often called the **Voronoi polygon** associated to p.

Let p and q be two distinct points in $\mathbf{R}^2$ and let $b(p,q)$ be the line that perpendicularly bisects the line segment $[p,q]$. We call $b(p,q)$ the **bisector** to p and q. The bisector divides the plane into two half planes and we denote by $H(p,q)$ that closed half plane that contains p.

Proposition 4 *For all $p \in S$ we have that*

$$V_S(p) = \bigcap_{q \in S \setminus \{p\}} H(p,q),$$

and

$$\bigcup_{p \in S} V_S(p) = \mathbf{R}^2.$$

If $p, q \in S$, $p \neq q$, then

$$\mathrm{Int}(V_S(p)) \bigcap \mathrm{Int}(V_S(q)) = \emptyset.$$

$V_S(p)$ is a closed convex polygonal domain with non-empty interior. Moreover, $V_S(p)$ is unbounded if and only if $p \in S \bigcap \partial\mathrm{conv}(S)$.

PROOF: If $p \in S \bigcap \mathrm{Int}(\mathrm{conv}(S))$ then there is a triangle T with vertices in $\mathrm{ext}(S)$ such that $p \in \mathrm{Int}(T)$. It is easily seen that

$$V^* = \bigcap_{q \in \mathrm{ext}(T)} H(p,q)$$

is bounded. Hence $V_S(p)$ is bounded in this case, which easily yields the proposition. □

We call the edges of the Voronoi polygons for **Voronoi edges**. We will say that a Voronoi edge e is **generated** by $p, q \in S$ if $e = V_S(p) \bigcap V_S(q)$. The vertices of the Voronoi polygons are called **Voronoi vertices**. The following proposition is an immediate consequence of Proposition 4.

Proposition 5 *All Voronoi edges are half lines or bounded line segments. Let the Voronoi edge e be generated by p and q. Then e is bounded if and only if*

$$[p, q] \bigcap \mathrm{Int}(\mathrm{conv}(S)) \neq \emptyset.$$

Moreover, e is unbounded if and only if

$$[p, q] \subset \partial(\mathrm{conv}(S)).$$

Proposition 6 *Let $p \in S$ and let $q \in S \setminus \{p\}$ be the point in $S \setminus \{p\}$ that is closest to p. Then p and q generate a Voronoi edge.*

PROOF: Let $\eta = \dfrac{p+q}{2}$. It is enough to show that $[p, \eta]$ does not intersect any Voronoi edge generated by p and some $w \in S \setminus \{p, q\}$. If $\eta_m \in [p, \eta]$ belongs to the Voronoi edge generated by p and some $w \in S \setminus \{p, q\}$ then

$$\frac{|p - w|}{2} < |p - \eta_m| \leq |p - \eta| = \frac{|p - q|}{2}.$$

Since q was closest to p we must have $|p - w| = |p - q| = |p - \eta_m|$. Hence $\eta = \eta_m$, which is impossible because $w \neq q$. □

Let Q_S be the collection of all Voronoi vertices for S. Since the Voronoi polygons are convex, at least three Voronoi edges meet at any Voronoi vertex.

Proposition 7 *For $q \in Q_S$ let Δ_q be the maximal closed disk Δ centered at q such that* $\mathrm{Int}(\Delta) \bigcap S = \emptyset$. *Then $\partial \Delta_q \bigcap S$ has at least three elements.*

PROOF: For $q \in Q_S$ let

$$V^*(q) = \{p \in S : V_S(p) \text{ has a vertex at } q\}.$$

If $p_1, p_2 \in V^*(q)$ share an edge then $|q - p_1| = |q - p_2|$. Let r be the common value of $\{|p - q| : p \in V^*(q)\}$. Pick $p_0 \in \partial \Delta_q \bigcap S$. From the definition of $V_S(p_0)$ follows that $q \in V_S(p_0)$, i.e. $p_0 \in V^*(q)$. Hence $\partial \Delta_q \bigcap S = V^*(q)$ and since $V^*(q)$ has at least three elements we are done. □

For $q \in Q_S$ let $V_S^*(q)$ be the points $p \in S$ for which $V_S(p)$ has a vertex at q. We denote by $W_S(q)$ the closed convex hull of $V_S^*(q)$.

Theorem 22 *We have that*

$$\bigcup_{q \in Q_S} W_S(q) = \mathrm{conv}(S). \tag{9}$$

If $q_1, q_2 \in Q_S$ with $q_1 \neq q_2$ then

$$\mathrm{Int}(W_S(q_1)) \bigcap \mathrm{Int}(W_S(q_2)) = \emptyset. \tag{10}$$

Moreover, $V_S^*(q)$ is contained in a circle centered at q.

PROOF: Let Δ_q be the maximal disk centered at q for which $\mathrm{Int}(\Delta_q) \bigcap S = \emptyset$. If $q \in Q_S$ then $V_S^*(q) \subset \partial \Delta_q$ by Proposition 7. Hence $W_S(q) \subset \Delta_q \bigcap \mathrm{conv}(S)$.

We shall now establish (10). Let $q_1, q_2 \in Q_S$ with $q_1 \neq q_2$. If $\mathrm{Int}(\Delta_{q_1}) \bigcap \mathrm{Int}(\Delta_{q_2}) = \emptyset$ then (10) follows trivially. Assume now that

$$\mathrm{Int}(\Delta_{q_1}) \bigcap \mathrm{Int}(\Delta_{q_2}) \neq \emptyset.$$

Neither Δ_{q_1} nor Δ_{q_2} is contained in the other. Hence $\partial \Delta_{q_1} \bigcap \partial \Delta_{q_2}$ consists of two points z_1 and z_2. Let l be the line through z_1 and z_2. Then l separates q_1 and q_2 and let H_i be the closed half plane determined by l such that $q_i \in H_i$. Clearly $\mathrm{Int}(H_1) \bigcap \mathrm{Int}(H_2) = \emptyset$. Clearly $V_S^*(q_1) \bigcap \mathrm{Int}(\Delta_{q_2}) = \emptyset$. Since $V_S^*(q_1) \subset \Delta_{q_1}$ if follows that

$$V_S^*(q_1) \subset \Delta_{q_1} \setminus (\mathrm{Int}\Delta_{q_2}) \subset H_1.$$

Similarly

$$V_S^*(q_2) \subset H_2.$$

Hence $W_S(q_1) \subset H_1$ and $W_S(q_2) \subset H_2$, which yields (10).

Let $X = \bigcup_{q \in Q_S} W_S(q)$. Clearly $X \subset \mathrm{conv}(S)$ but assume $X \neq \mathrm{conv}(S)$. Hence there is a non-empty disk B such that $B \subset \mathrm{conv}(S)$ but $B \bigcap X = \emptyset$. Let l be any line such that $l \bigcap B \neq \emptyset$, $l \bigcap S = \emptyset$ but $l \bigcap \mathrm{Int}(W_S(v)) \neq \emptyset$ for some $v \in Q_S$. Let x be any point in $B \bigcap l$ and let $w \in l \bigcap X$ be a point that is closest to x. Then $w \in W_S(q_1)$ for some $q_1 \in Q_S$. In addition, $l \bigcap \mathrm{Int}(W_S(q_1)) \neq \emptyset$ and w belongs to a unique edge $[p_1, p_2]$ of $W_S(q_1)$. Of course $p_1, p_2 \in S$.

figure

Let e be the Voronoi edge that crosses $[p_1, p_2]$. Then q_1 is an end point of e. Furthermore, e is generated by p_1 and p_2. Let l_1 be the line through p_1 and p_2. Let $p_3 \in V_S^*(q_1)$ be different from both p_1 and p_2. Then l_1 separates x and p_3 so $[p_1, p_2] \bigcap \text{Int}(\text{conv}(S)) \neq \emptyset$. From Proposition 7 we have that e must be bounded. Let q_2 be the end point of e that is different from q_1. Now $p_1, p_2 \in V_S^*(q_2)$. Hence $w \in \text{Int}(W_S(q_1) \bigcap W_S(q_2))$, so $w \in \text{Int}X$. This contradicts the definition of w, which yields the theorem. □

Triangulating each $W_S(q)$ gives now a Delaunay triangulation.

Corollary 2 *Let $\mathcal{T}$ be a triangulation of* conv(S) *with vertex set S such that each $T \in \mathcal{T}$ is contained in some $W_S(q)$, $q \in Q_S$. Then $\mathcal{T}$ is a Delaunay triangulation for S.*

PROOF: Immediate consequence of Theorem 22. □

For $p \in \mathbf{R}^2$ let h_p be the half space

$$\{(x, y) : x \in \mathbf{R}^2, y \in \mathbf{R} \ \text{ and } \ y \geq 2 < x, p > -|p|^2\}.$$

Then the boundary of h_p is the tangent plane to the paraboloid $y = |x|^2$ at $x = p$. Let $\pi : \mathbf{R}^3 \to \mathbf{R}^2$ be the projection that sends (x, y) into x. An easy computation shows that $x \in \pi(\partial h_p \bigcap h_q)$ if and only if $|x - p| \leq |x - q|$. Hence we have the following method for constructing the Voronoi diagram in $\mathbf{R}^2$.

Proposition 8 *Let h_S be the polyhedron*

$$h_S = \bigcap_{p \in S} h_p.$$

Then the Voronoi polygons for S are the projections of the faces of h_S.

4.7 Pattern Recognition and Subgraphs of Delaunay Triangulations

For pattern recognition one has developed several graphs associated to a point set. We will in this section describe some of these and their relationships to Delaunay triangulations.

Throughout this section we let S denote a finite point set in $\mathbf{R}^2$ that is not contained in a line. In particular, there are at least three points in S.

Setting $f(x) = |x|^2$, $x \in \mathbf{R}^2$, we let $S_f = \{(p, f(p)) : p \in S\}$ be the lift of S. Denote by L_S the lower part of the convex hull of S_f, see section 4.1. We let $DG(S)$ be the **Delaunay graph** of S, i.e. the graph with vertex set S and pq is an edge in $DG(S)$ if and only if $(p, f(p))$ and $(q, f(q))$ are the end points of an edge of L_S. Notice that $DG(S)$ is a subgraph of the skeleton of any Delaunay triangulation.

Let S be the vertex set of a convex polygon and let U_S denote the upper part of the convex hull of S_f. We denote by $FDG(S)$ the **farthest Delaunay graph**, i.e. the graph with vertex set S and $p, q \in S$ form an edge if and only if $(p, f(p))$ and $(q, f(q))$ form an edge in U_S. For a general set S we let $FDG(S)$ denote $FDG(\text{ext}(S))$.

The Gabriel graph, see [23], [37] and [58], is used for defining adjacency in point patterns. If $p, q \in S$, $p \neq q$, then pq is an edge in the **Gabriel graph** $GG(S)$ if and only if

$$\Delta_{pq} \bigcap (S \setminus \{p, q\}) = \emptyset,$$

where Δ_{pq} is the closed disk that has $[p, q]$ as a diameter.

Theorem 23 *The Gabriel graph $GG(S)$ of S is a subgraph of the Delaunay graph $DG(S)$ of S.*

Proof: From Theorem 10 follows that every edge in $GG(S)$ is the edge of some triangle in some Delaunay triangulation of S. Assume that pq is an edge in $GG(S)$ but pq is not an edge in $DG(S)$. Let T be a triangle in a Delaunay triangulation of S such that $[p, q]$ is an edge of T. Let Δ be the closest disk whose boundary contains the vertices of T. Then $\text{Int}(\Delta) \bigcap S = \emptyset$ and $\partial\Delta \bigcap (S \setminus \{p, q\}) \neq \emptyset$. hence $[p, q]$ is not a diameter of S. Let h_+, h_- be the two open half planes determined by the line through p and q. Since $[p, q]$ does not belong to $DG(S)$ we must have $\Delta \bigcap S \bigcap h_+ \neq \emptyset$ and $\Delta \bigcap S \bigcap h_- \neq \emptyset$. Since $\Delta \neq \Delta_{pq}$ we must have that $\Delta \bigcap h_+$ or $\Delta \bigcap h_-$ is contained in $\text{Int}(\Delta_{pq})$. This contradiction establishes the theorem. $\square$

If $p \neq q$ we let L_{pq} denote the **lune**

$$L_{pq} = \{w \in \mathbf{R}^2 : |w - p| \leq r\} \bigcap \{w \in \mathbf{R}^2 : |w - q| \leq r\},$$

where $r = |p - q|$.

We can now define the **relative neighbourhood graph** $RNG(S)$ of the set S. For reference see [57] and [56]. If $p, q \in S$ then pq is an edge of $RNG(S)$ if and only if

$$\text{Int}(L_{pq}) \bigcap S = \emptyset.$$

Since $\Delta_{pq} \subset L_{pq}$ and $\Delta_{pq} \setminus \{p, q\} \subset \text{Int}(L_{pq})$ we have the following result.

Theorem 24 *The relative neighbourhood graph $RNG(S)$ is a subgraph of the Gabriel graph $GG(S)$.*

A **Euclidean minimum spanning tree** of S, the set of such graphs is denoted by $EMST(S)$, is a tree having the vertex set S and for which the sum of the length of all edges attains the minimum over all trees having the vertex set S. See [25] and [48].

Theorem 25 *An edge of an Euclidean minimum spanning tree for S is an edge of $RNG(S)$.*

PROOF: We will argue by contradiction. Assume pq is an edge of a Euclidean spanning tree G that is not an edge of $RNG(S)$. Removing the edge pq from G splits G into two subtrees G_1 and G_2, say p is in G_1 and q is in G_2. Our assumption means that there is a $w \in \text{Int}(L_{pq}) \bigcap S$. We may without loss of generality assume that $w \in G_1$. We now get a connected graph G_3 by taking the union of G_1 and G_2 with the edge wq. Since the total length of the edges in G_3 is smaller than the total length of the edges in G we get a contradiction. □

The above subgraphs of the Delaunay triangulation are all undirected. We shall now introduce a directed graph of S, the **nearest neighbour graph** $NNG(S)$, by letting $\vec{pq}$, $p \neq q$, be an edge if and only if

$$|p - q| = \min\{|p - w| : w \in S \setminus \{p\}\}.$$

For reference see [33], [50] and [43].

Theorem 26 *Let $p, q \in S$, $p \neq q$. If $\vec{pq}$ is an edge in $NNG(S)$ then pq is an edge in $RNG(S)$. If there is no other edge in $NNG(S)$ having initial vertex p then pq is an edge of every $EMST(S)$.*

PROOF: The first statement is obvious. We shall now prove the second part. Assume to the contrary that pq is not an edge of a Euclidean spanning tree G. Let G_1 be the graph one gets by adding pq as an edge to G. Then G_1 has a cycle Γ that contains pq. Hence there is a $w \in S \setminus \{p, q\}$ such that pw is an edge of G that is on Γ. Let G_2 be the graph one gets by removing the edge pw from G_1. Then G_2 is a spanning tree for S. Since the edge pq is shorter than pw we have that the total length of the edges in G_2 is smaller than that of G. We get a contradiction. □

We next quote a remarkable result from [29].

Theorem 27 *Let $\mathcal{T}$ be a Delaunay triangulation of S. For $p, q \in S$ let $D(p, q)$ denote the length of the shortest chain of edges of $\mathcal{T}$ connecting p and q. Then*

$$D(p, q) \leq C|p - q|,$$

where $C = \dfrac{2\pi}{3\cos(\frac{\pi}{6})} \approx 2.42.$

For a proof see [29].

We shall now describe a generalization of the convex hull of S. This generalization leads to a one-parametric family of subdivisions which intuitively connects the "crude" and "fine" shape of S.

The generalization of the convex hull is based on the notion of α–**disks** with α any real number. This is

1. closed disks of radii $\dfrac{1}{\alpha}$ if $\alpha > 0$;

2. closed half planes if $\alpha = 0$;

3. the complements of closed disks of radii $-\dfrac{1}{\alpha}$ if $\alpha < 0$.

We now define the α–**hull** of S as the intersection of all α–disks that contain S. We denote the α–hull of S by $\text{Hull}_\alpha(S)$. We will use the convention that $\text{Hull}_\alpha(S)$ is the entire plane if S is not contained in any α–disk. We have the following result.

Proposition 9 *If* $\alpha_1 < \alpha_2$ *then* $\text{Hull}_{\alpha_1}(S) \subset \text{Hull}_{\alpha_2}(S)$.

PROOF: Assume A_2 is an α_2–disk that contains S. Since A_2 is the intersection of all α_1–disks that contain A_2 we have that

$$\text{Hull}_{\alpha_1}(S) \subset A_2.$$

Taking the intersection of all α_2–disks that contain S yields the proposition.

□

We shall now introduce the notion of α-shape. A general reference for α-shape is [20] and [19] chapter 13.2.6.

We say that a point $p \in S$ is α-**extreme** if for all sufficiently small $\epsilon > 0$ there is an $(\alpha+\epsilon)$-disk $B_{\alpha+\epsilon}$ such that $S \subset B_{\alpha+\epsilon}$ and $S \bigcap \partial B_{\alpha+\epsilon} = \{p\}$. Assume now that $p, q \in S$, $p \neq q$. We will say that an α-disk B_α **leans on** the ordered pair (p,q) if $p, q \in \partial B_\alpha$ and an infinitesimal movement of B_α normal to and towards the left of the vector $q-p$ puts p and q outside B_α.

figure

Furthermore the ordered pair (p,q) is said to be α-**exposed** if for all sufficiently small $\epsilon > 0$ there is an $(\alpha+\epsilon)$-disk $B_{\alpha+\epsilon}$ such that $S \subset B_{\alpha+\epsilon}$ and $B_{\alpha+\epsilon}$ leans on (p,q).

Let l_{pq} denote the line through p and q that is directed from p to q. We note that (p,q) is 0-extreme if and only if $[p,q]$ is an edge of $\text{conv}(S)$ and no point of S lies to the right of l_{pq}.

figure

Proposition 10 *Suppose $p, q \in S$, $p \neq q$. Then the following holds:*

1. *There is a real number A_p such that p is α-extreme if and only if $\alpha < A_p$.*

2. *Suppose (p, q) is α-exposed for some α. Then there are real numbers A_{pq} and B_{pq} such that (p, q) is α-exposed if and only if $A_{pq} \leq \alpha < B_{pq}$.*

3. *If (p, q) is α-exposed then both p and q are α-extreme.*

4. *If $\alpha \geq 0$ and p is α-extreme then there is a $w \in S \setminus \{p\}$ such that (p, w) is α-exposed.*

PROOF: We now give the arguments for the statements in the order they are stated.

1. Let I_p be the set of all α such that p is α-extreme. Clearly I_p is bounded from above and $I_p \neq \emptyset$. By elementary geometry we see that if $\beta \in I_p$, then there is an $\epsilon > 0$ such that $(-\infty, \beta + \epsilon) \subset I_p$.

2. Let I_{pq} be the set of all α such that (p, q) is α-exposed. By assumption $I_{pq} \neq \emptyset$. Clearly I_{pq} is a bounded set. Let I^*_{pq} be the set of $\rho \in \mathbf{R}$ such that there is a ρ–disk B_ρ with $S \subset B_\rho$ and B_ρ leans on (p, q). Assume $\alpha_1 \leq \alpha_2$ and $\alpha_1, \alpha_2 \in I_{pq}$. By elementary geometry we have that if $\epsilon > 0$ is small enough then $[\alpha_1, \alpha_2 + \epsilon)$ is contained in I_{pq}. Hence I_{pq} is an interval. Let $A_{pq} = \inf I_{pq}$ and $B_{pq} = \sup I_{pq}$. Clearly $(A_{pq}, B_{pq}) \subset I^*_{pq}$. Hence $A_{pq} \in I_{pq}$ so $I_{pq} = [A_{pq}, B_{pq})$.

3. This is an immediate consequence of the definition of α–exposed.

4. Since the result is obvious for $\alpha = 0$ we will assume $\alpha > 0$. Let B be the smallest closed disk such that $S \subset B$ and $p \in \partial B$. Pick $p^* \in \partial B$ such that $[p, p^*]$ is a diameter of B. Let λ be the ray that originates in p and passes through p^*. Let Q be the circle with radius $\frac{1}{\alpha}$ and center p. If ρ denotes the radius of B then $\rho < \frac{1}{\alpha}$. Let p_0 be the intersection between Q and λ. Let Δ_q be the disk centered at q with radius $\frac{1}{\alpha}$. Notice that $B \subset \Delta_{p^*}$. In particular, $S \subset \Delta_{p^*}$ and $(S \setminus \{p\}) \bigcap \partial\Delta_{p^*} = \emptyset$.

figure

We have that $p \in \partial\Delta_q$ for all $q \in Q$. Starting at $q = p^*$ we now move q in the counter clockwise direction along Q until $(S \setminus \{p\}) \bigcap \partial\Delta_q \neq \emptyset$. Again let q denote this point on Q and set $E = (S \setminus \{p\}) \bigcap \Delta_q$. Let $w \in E$ be the furthest from p. Since $S \subset \Delta_q$ it is easily seen that (p, w) is α–exposed.

□

Let S_α be the points $p \in S$ such that for some $q \in S$ we have that (p, q) or (q, p) is α–exposed. We define the α–graph $EG_\alpha(S)$ as the graph on S_α such that $p, q \in S_\alpha$ form an edge if and only if (p, q) or (q, p) is α–exposed.

Theorem 28 *Let $\alpha \in \mathbf{R}$ and assume e is an edge in $EG_\alpha(S)$. If $\alpha \geq 0$ then e is an edge in $FDG(S)$. If $\alpha < 0$ then e is an edge of $DG(S)$.*

PROOF: This is an immediate consequence of Theorem 20 and Theorem 10.
□

5 Arrangements

Arrangements form one of the important structures in computational geometry. We slightly touch on the notion of duality and use results on Schinzel-Davenport sequences to prove the zone theorem.

5.1 Duality

We study two applications of geometric duality. The first one yields a fast algorithm for computing the minimum area triangle from among n points in the plane, see [12]. The second one gives a preprocessing algorithm for answering "ray-shooting" queries quickly, see [11].

Min-Area Triangles

To find the minimum area triangle, we first focus our attention on the problem of finding a min area triangle that has two fixed points P_i and P_j chosen as the "base" of the triangle. Our task reduces to finding for each pair of points, P_i and P_j, the point P_k whose normal distance to the line through P_i and P_j is the least. Say the line ℓ through P_iP_j has slope m and y intercept c. For all lines having slope m, through other points, the line through P_k has y intercept nearest to c. Call its intercept c'. (This follows from the fact that the normal distance between the lines is proportional to their y intercepts.)

In the dual, this line will be the first point vertically below (above) the point of intersection of the duals of P_i and P_j. This intersection point will have coordinates, (m, c'), and the line closest to this, and vertically below (above), is the dual of P_k, passing through the point (m, c').

To compute point P_k, for every pair of points P_iP_j, we construct the arrangement of the dual graph. While constructing the arrangement, we keep track of which line is immediately below which point. This can can easily be done while constructing the arrangement itself. (Recall that in the proof of the zone theorem, we do a "walk" around each face that the new line is incident to. While doing this weak we can perform the required updates.) This gives us all the information we need. To compute the min area triangle, we minimize over all P_i and P_j, which can be done in $O(n^2)$ time, where n is the number of points.

Figure

Ray-Shooting

Given a simple polygon P of n vertices, we wish to preprocess this polygon to answer ray-shooting queries of the following form: given a query point $q \in P$ and a direction $\vec{d}$, which is the *first* edge of the polygon that is hit when we shoot a ray from q? We will show that these queries can be answered in $O(\log^2 n)$ time. (Chazelle and Guibas provide a more complex $O(\log n)$ query time algorithm as well.)

The main idea is to decompose the polygon into a "hierarchy" of polygons. We use the "polygon-cutting theorem" of Chazelle's to do this. The polygon cutting theorem states that every simple polygon with n vertices ($n > 3$) contains a diagonal such that adding the diagonal, partitions the polygon into two polygons P_1 and P_2, with each P_i having at most $\frac{2}{3}n$ vertices.

We now construct a tree using this decomposition. The vertices of this tree correspond to polygons. The root corresponds to the polygon P. We apply the polygon cutting theorem to P, and the children of this vertex correspond to polygons P_1 and P_2 respectively. We now apply the polygon cutting theorem recursively to each node, until the corresponding polygon is a triangle. Observe that the height of this tree is $O(\log n)$, since we are reducing the size of the polygons by a constant factor at each step. For any polygon P, let $P(e)$ denote the diagonal that is used to decompose the polygon. We assume that each node in the tree has the diagonal information available.

The polygon cutting theorem can be proved using the fact that the dual of the triangulation is a tree, and we can prove that each such tree (vertices with degree at most 3) always contains an edge that partitions the tree into smaller trees with at most $\frac{2}{3}n$ vertices.

We solve the query $hit(P, q, \vec{d})$ as follows. We assume that $q \in P$. If $q \in P_1$ then the only way the ray can exit the polygon from δP_2 is if it crosses edge $P(e)$ without crossing any polygon boundary of P_1.

Figure

Step 1. (Assume that $q \in P_1$.) If the ray from point q with direction $\vec{d}$ does not intersect edge $P(e)$, then recursively solve $hit(P_1, q, \vec{d})$.

Step 2. If the ray appears to pierce $p(e)$ at point p, then we have to check if the ray is able to reach $P(e)$ from q. If so, we recursively solve $hit(P_2, p, \vec{d})$ else we recursively solve $hit(P_1, q, \vec{d})$.

We still need to check if the ray from q reaches $P(e)$. We do this by shooting a ray from p towards q. If the ray reaches q, then the ray from q must reach $P(e)$. Compute $o = reach(P_1, p, -\vec{d})$ – if the segment $\overline{po}$ contains q then the ray from q in direction d reaches p, otherwise it crosses the boundary at an earlier point. This reduces the problem to another ray shooting query. Why is this query any easier to solve?

We will show that such queries are simpler, since they are being shot from a particular edge on the boundary of P_1, namely the edge $P(e)$. Chazelle and Guibas make use of a duality transformation that establishes a correspondence between points (on a two sided plane) and rays.

Any ray r_i shot from $P(e)$ becomes a point p_i in dual space. We decompose the dual space into "regions" with the property that if two rays r_1 and r_2 from $P(e)$ hit the same edge e' (perhaps at different places) then the points p_1 and p_2 are in the same "region". Moreover, they prove that this region is *convex*. The main idea is that any query ray r' shot from edge $P(e)$ can be transformed into a point p', and by doing a *point-location* query in $O(\log n)$ time, we can determine the edge that ray r' hits (based on the region it falls into). Once we know the edge, it is easy to compute the intersection point. This takes $O(\log n)$ time.

The entire algorithm takes $O(\log^2 n)$ time, since we start at the root of the polygon tree, and proceed down the tree doing a point location query at each level. The tree has $O(\log n)$ levels, and each point location takes $O(\log n)$ time. Further data structure tricks are used to reduce the query time to $O(\log n)$.

5.2 Davenport-Schinzel Sequences

A Davenport-Schinzel sequence is a sequence $U = (u_1, \ldots, u_m)$ constructed over an alphabet of size n. We define a **Davenport-Schinzel sequence** of order s as follows (referred to as $DS(n, s)$). There are two conditions that these sequences are required to satisfy.

1. $u_i \neq u_{i+1}$ for each $i < m$.
2. There do not exist $s+2$ indices $1 \leq i_1 < i_2 < \ldots < i_{s+2} \leq m$ such that

$$u_{i_1} = u_{i_3} = u_{i_5} = \ldots = a$$

$$u_{i_2} = u_{i_4} = u_{i_6} = \ldots = b$$

   and $a \neq b$.

This means that the presence of long alternations of any pair of distinct symbols in Davenport-Schinzel sequences are forbidden. As an example we see that 122341 is not a DS sequence since 2 occurs consecutively, and that 14234132 is a DS sequence of order 3, but not of order 2, since there is an alternating subsequence $1 \ldots 2 \ldots 1 \ldots 2$.

Let $\lambda_s(n)$ be the length of a longest $DS(n, s)$ sequence. We are interested in deriving good upper bounds on $\lambda_s(n)$, as function of n. These sequences are interesting since they can be used to derive upper bounds on the complexity of lower envelopes of line segments, lower envelopes of more general functions, as well as obtaining an alternate proof of the zone theorem. They are very often used to obtain simpler proofs of the complexity of various combinatorial structures as well. In short, they are very useful! For the origin of this notion see [15] and [16].

We first give some results for $\lambda_s(n)$ for $s = 1, 2, 3$.

Lemma 16 $\lambda_1(n) = n$

PROOF: We note that $(1, 2, \ldots, n)$ is a $DS(1, n)$-sequence so $\lambda_1(n) \geq n$. To prove the reverse inequality let U be a $DS(1, n)$-sequence. U cannot take the form $(\ldots a \ldots b \ldots a \ldots)$ for any $a \neq b$ and therefore the elements in U must be distinct. Hence $|U| \leq n$. □

Lemma 17 $\lambda_2(n) = 2n - 1$

PROOF: We note that $(1, 2, \ldots, n-1, n, n-1, \ldots, 2, 1)$ is a $DS(2, n)$-sequence so $\lambda_2(n) \geq 2n - 1$. The reverse inequality is proved by induction. The case $n = 1$ is obvious. Assume the statement is true for $n - 1$ and let U be a $DS(2, n)$-sequence (over the alphabet consisting of $1, 2, \ldots n$). Without loss of generality we can assume that the leftmost occurrence of i is to the left of j iff[4] $i < j$. This implies that there can be only one occurrence of n in U otherwise there must be a subsequence of the form $(i, \ldots, n, \ldots, i, \ldots, n)$ which is forbidden. Remove the single occurrence of n, and if the symbols adjacent to n are the same remove also one of these from U. The resulting sequence is a $DS(2, n - 1)$-sequence and it is one or two elements shorter than U. The induction hypothesis yields $|U| \leq 2(n - 1) - 1 + 2 = 2n - 1$. □

Lemma 18 $\lambda_3(n) \leq 2n(1 + \log n)$

PROOF: We prove this by induction. Let a be the least frequently occurring character in the sequence. An upper bound on the number of times it occurs

[4] iff="if and only if"

is $\frac{\lambda_3(n)}{n}$. Deleting all copies of a yields a sequence that can potentially have consecutive characters. However, at most 2 such consecutive characters are generated (one each, from the first and last occurrence of a). Otherwise, the sequence has the form $\ldots a \ldots bab \ldots a \ldots$ and is not a DS sequence or order 3. The remaining sequence is a $DS(n-1,3)$ sequence of length $\lambda_3(n-1)$. Hence, an upper bound on $\lambda_3(n)$ is $\lambda_3(n-1) + \frac{\lambda_3(n)}{n} + 2$. Simplifying, we get $(1-1/n)\lambda_3(n) \leq \lambda_3(n-1) + 2$.

This is the same as $\frac{\lambda_3(n)}{n} \leq \frac{\lambda_3(n-1)}{n-1} + \frac{2}{n-1}$. Expanding the recurrence gives the upper bound. □

Zone Theorem

We now study some applications of these sequences. The first application is to obtain a new proof of the zone theorem.

Let ℓ be the new line that is being inserted. We would like to prove that the total complexity of the faces that ℓ intersects is $O(n)$. We do this as follows. Assume ℓ is horizontal. We will show that the complexity of the portion of the faces "above" ℓ is $4n$. Assume no two lines are parallel. With this assumption the highest vertex on any face is unique (since no other line is horizontal). As we traverse the faces in clockwise order, we output the line number each edge belongs to. Lines are labelled with distinct numbers. We create two sequences, a left edge sequence and a right edge sequence. For the edges before the highest point in the clockwise traversal of a face, the line numbers are appended to the left sequence (LS), and the edges after the highest point in the traversal, the line numbers are appended to the right sequence (RS).

We claim that each list is $DS(n,2)$ sequence. The first condition is easy to establish. The second condition is established by a "picture-proof". Consider the right sequence. If it contains $a \ldots b \ldots a$ then it cannot contain b again. Make a picture!

Figure

This proves that the total complexity of the faces above ℓ is at most $4n$. This gives us a bound of about $8n$. Since the lines that cross ℓ were counted twice, we can reduce this to $6n$. However, the tight upper bound is $\frac{11}{2}n$.

Closest Pair/Nearest Neighbors of Moving Points

Let P be set of n points in the plane that are moving along predefined trajectories. Let $x_i(t)$ and $y_i(t)$ denote the position of point p_i at time t. Assume that these functions are polynomials of a fixed degree s. We fix our attention on point p_i, and let $g_{ij}(t)$ be the square of the distance from p_i to p_j at time t. Clearly, g_{ij} is a polynomial of degree $2s$. Let $G_i(t) = \min_{j \neq i} g_{ij}(t)$. The function $G_i(t)$ is a minima of $n-1$ functions. Each pair of functions g_{ij} and g_{ik} pairwise intersect in at most $2s$ points, since that is the degree of the polynomial defining them. We now claim that the number of intersection points on the lower envelope is at most $\lambda_{2s}(n-1)$. This is because if two functions appear $2s+2$ times alternately on the lower envelope is at most $\lambda_{2s}(n-1)$. This is because if two functions appear $2s+2$ times alternately on the lower envelope, then they must have $2s+1$ intersection points. This is an upper bound on the number of times the nearest neighbor of p_i can change.

References

[1] Aho,A.V./Hopcroft,J.E./Ullman,J.D., *The design and analysis of computer algorithms*, Addison-Wesley 1974

[2] Aurenhammer,F., Voronoi diagrams – A survey of a fundamental geometric data structure, *ACM Computing Surveys* **23** (1991) 345 – 405

[3] D'Azevedo,E.F., On optimal interpolation triangle incidences, *SIAM J. Sci. Stat. Comput.* **10** (1989) 1063 – 1075

[4] Bern,M./Eppstein,D., Mesh generation and optimal triangulation, in *Computing in Euclidean geometry*, 23 – 90, Lecture Notes Ser. Comput., 1, World Sci. Publishing, River Edge, NJ, 1992.

[5] Blaschke,H., Über Paare umfangsgleichen Eilinien, *Jahresber. Deutsch. Math.-Verein.* **48** (1938) 69 – 74

[6] Bol,G., Über ein Satz von Cauchy, *Jahresber. Deutsch. Math.-Verein.* **48** (1938) 74 – 76

[7] Brown,K.Q., Voronoi diagrams from convex hulls, *Information Processing Letters* **9**(5) 223 – 228

[8] Cauchy,A.L., Sur les polygones et les polyedres, Werke (2) 1 (Paris 1905) 26 – 38

[9] Chazelle,B., Triangulating a simple polygon in linear time, *Disc. and Comp. Geometry* **6** (1991) 485 – 524

[10] Chazelle,B./Edelsbrunner,H., Optimal solutions for a class of point retrieval systems, *J. Symbolic Comput.* **1** (1985) 47 – 56

[11] Chazelle,B./Guibas,L., Visibility and intersection problems in plane geometry, *Discrete Comput. Geom.* **4** (1989) 551 – 581

[12] Chazelle,B./Guibas,L./Lee,D., The power of geometric duality, *BIT* **25** (1985) 76 – 90

[13] Chazelle,B./Palios,L., Triangulating a non-convex polytope, *Disc. and Comp. Geometry.* **5** (1990) 505 – 526

[14] Chvatal,V., A combinatorial theorem in plane geometry, *J. Combinatorial Theory (B)*, **18** (1975) 39 – 41

[15] Davenport,H./Schinzel,A., A combinatorial problem connected with differential equations, *Amer. J. Math.*, **87** (1965) 684 – 689

[16] Davenport,H., A combinatorial problem connected with differential equations II, *Acta Arithmetica*, **17** (1971) 363 – 372

[17] Delaunay,B. Sur la sphère vide, *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk* **7** (1934) 793 – 800

[18] Dyn,N./Goren,I./Rippa,S., Transforming triangulations in polygonal domains, *Comput. Aided Geom. Design* **10** (1993) 531 – 536

[19] Edelsbrunner,H., *Algorithms in combinatorial geometry*, Springer-Verlag 1987

[20] Edelsbrunner,H/Kirkpatrik,D.G./Seidel,R., On the shape of a set of points in the plane, *IEEE Trans. Inform. Theory* **IT-29** (1983) 551 – 559

[21] Edelsbrunner,H./Tan,T.S., An upper bound for conforming Delaunay triangulations, *Disc. and Comp. Geometry* **10** (1993) 197 – 213

[22] Eppstein,D., Approximating the minimum weight triangulation, *Disc. and Comp. Geometry* **11** (1994) 163 – 191

[23] Gabriel,K.R./Sokol,R.R., A new statistical approach to geographic variation analysis, *Systematic Zoology*, **18** (1969) 259 – 278

[24] Goodman,J./O'Rourke,J.(eds), *Handbook of discrete and computational geometry*, CRC Press 1997

[25] Gower,J.C./Ross,G.J.S., Minimum spanning trees and single linkage cluster analysis, *Applied Statistics* **18** (1969) 54 – 64

[26] Guggenheimer,H., On plane Minkowski geometry, *Geom. Dedicata* **12** (1982) 371 – 381

[27] Hwang,F.K./Richards,D./Winter,P., *The Steiner tree problem*, New York, Elsevier

[28] Ivanov,A./Tuzhilin,A., *Minimal networks*, CRC Press

[29] Keil,J.M./Gutwin,C.A., The Delaunay closely approximates the complete Euclidean graph. In: *Proceedings of the First Workshop on Algorithms and Data Structures*, Lecture Notes in Computer Science volume 382 (1989) 47 – 56

[30] Kirkpatrik,D.G./Seidel,R., The ultimate planar convex hull algorithm?, *SIAM J. Comput.*, **15** (1986) 287 – 299

[31] Klincsek,G.T., Minimal triangulations of polygonal domains, *Ann. Disc. Math.* **9** (1980) 121 – 123

[32] Knuth,D.E., *The art of computer programming, Vol. 3: Sorting and searching*, Addison-Wesley 1973

[33] Kolars,J.F./Nystuen,J.D., *Human Geography*, McGraw-Hill, New York 1974

[34] Lawson,C.L., Software for C^1 surface interpolation. In: Rice,J.(ed), *Mathematical Software III*, Academic Press, New York, 161 – 194

[35] Lee,D.T./Lin,A., Generalized Delaunay triangulation for planar graphs, *Disc. and Comp. Geometry*, **1** (1986) 201 – 217

[36] Lin,M.C./Manocha,D.(eds), *Applied computational geometry*, Lecture Notes in Computer Science 1148, Springer-Verlag 1996

[37] Matula,D.W./Sokol,R.R., Properties of Gabriel graphs relevant to geographic variation research and clustering of points in the plane, *Geographic Analysis* **12** (1980) 205 – 222

[38] Mezey,P.G., The α, β, γ-hull and the T-hull of a point set, tools for the analysis of shapes and relative orientation of objects in 3-space, *J. Math. Chem.* **8** (1991) 91 – 102

[39] Nishizeki,T./Chiba,N., *Planar graphs: theory and algorithms*, Nort-Holland 1988

[40] Okabe,A./Boots,B./Sugihara,K., *Spatial tessellations: Concepts and applications of Voronoi diagrams*, Wiley 1992

[41] O'Rourke,J., *Computational geometry in C*, Cambridge University Press 1998

[42] O'Rourke,J., *Art gallery theorems and algorithms*, Oxford University Press 1987

[43] Pielou,E., *Mathematical Ecology*, Wiley-Interscience, New York 1977

[44] Plaisted,D.A./Hong,J., A heuristic triangulation algorithm, *J. Algorithms* **8** (1987) 405 – 437

[45] Preparata,F.P./Shamos,M.I., *Computational Geometry: an introduction*, Springer-Verlag, New York 1985

[46] Rajan,V.T., Optimality of the Delaunay triangulation in $\mathbf{R}^d$, *Proc. 7th ACM Symp. Comp. Geometry* (1991) 357 – 363

[47] Rippa,S., Minimal roughness property of the Delaunay triangulation, *Comput. Aided Geom. Design* **7** (1990) 489 – 497

[48] Rohlb,F.J., Hierarchical clustering using the minimum spanning tree, *The Computer Journal* **16** (1973) 93 – 95

[49] Rosen,K.(ed), *Handbook of discrete and combinatorial mathematics*, CRC Press 2000

[50] Saaty,T.L., *Optimization in Integers and Related Problems*, McGraw-Hill, New York 1970

[51] Sack,J.-R./Urrutia,J.(eds), *Handbook of computational geometry*, Elsevier 2000

[52] Schatteman,A., A four-vertex theorem for polygons and its generalization to polytopes, *Geom. Dedicata* **34** (1990) 229 – 242

[53] Schmitt,D./Spehner,J.-C., On equiangularity of Delaunay diagrams in every dimension, *Proc. 5th Caanadian Conf. Comp. Geometry* (1993) 346 – 351

[54] Schoenberg,I.J./Zaremba,S.K., On Cauchy's lemma concerning convex polygons, *Canad. J. Math.* **19** (1967) 1062 – 1071

[55] Selig,J.M., *Geometrical Methods in Robotics*, Springer-Verlag 1996

[56] Supowit,K.J., The relative neighborhood graph with an application to minimum spannning trees, *Journal of ACM* **30** (1983) 428 – 447

[57] Toussaint,G.T., Pattern recognition and geometrical problems. In: *Proceedings of the Fifth International Conference on Pattern Recognition*, Miami Beach, Volume 2, 1324 – 1347

[58] Urquhart,R., Graph theoretical clustering based on limited neighbourhood sets, *Pattern Recognition* **15** (1982) 173 – 187

[59] Voronoi,G., Nouvelles applications des paramètres continues à la théorie des formes quadratiques. Premier Mémoire: Sur quelques propriétés des formes quadratiques positives parfaites, *J. Reine Angew. Math.*, **133** (1907) 97 – 178

[60] Voronoi,G., Nouvelles applications des paramètres continues à la théorie des formes quadratiques. Deuxième Mémoire: Recherches sur les parallélloèdres primitifs, *J. Reine Angew. Math.*, **134** (1908) 198 – 287

[61] Ziegler,G., *Lectures on polytopes*, Springer-Verlag, Berlin 1995