Building Good Triangulations

1. Nets and thick triangulations
2. Triangulation of manifolds

Jean-Daniel Boissonnat
INRIA

Hamilton Mathematics Institute
17-19 June, 2018

Triangulations

A central subject since the early days of Computational Geometry

- Triangulations as data structures
- Triangulation of polygonal/polyhedral domains
- Optimal triangulations
- Delaunay triangulation

A central subject in Mesh Generation, Manifold Learning

- Quality of approximation
- Quality of elements
- Higher dimensions
- More general topological spaces

1. Delaunay-like complexes

2. Nets and Delaunay refinement

3. Thick triangulations

Voronoi diagrams

A set of points $\mathcal{P}$ in $(\mathbb{R}^d, \|.\|)$

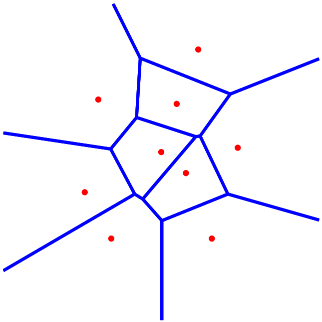

Voronoi cell $\quad V(p_i) = \{x : \|x - p_i\| \leq \|x - p_j\|, \ \forall j\}$

Voronoi diagram $(\mathcal{P})$ $\quad = \{$ set of cells $V(p_i), p_i \in \mathcal{P}$ $\}$

Delaunay Triangulations

Sur la sphère vide (On the empty sphere), Boris Delaunay (1934)

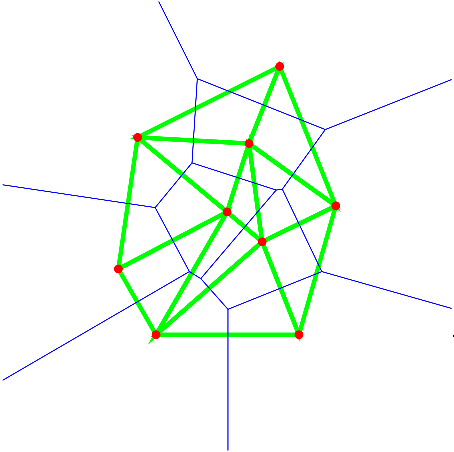

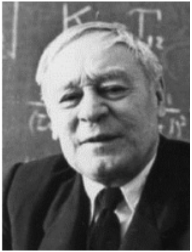

The Delaunay complex $\mathrm{Del}(\mathcal{P})$ is the nerve of $\mathrm{Vor}(\mathcal{P})$

Theorem

If $\mathcal{P}$ contains no subset of $d+2$ points on a same hypersphere, then $\mathrm{Del}(\mathcal{P})$ is a triangulation of $\mathcal{P}$

Delaunay Triangulations

Sur la sphère vide (On the empty sphere), Boris Delaunay (1934)

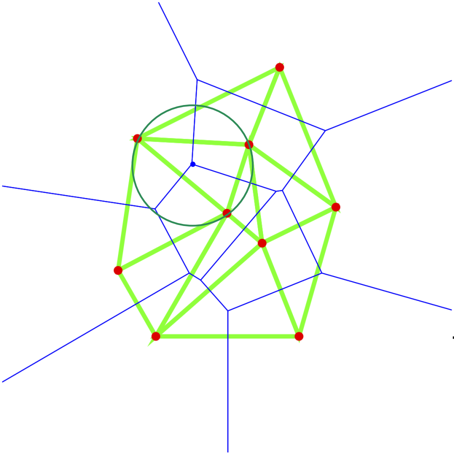

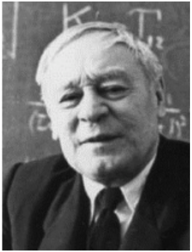

The Delaunay complex $\mathrm{Del}(\mathcal{P})$ is the nerve of $\mathrm{Vor}(\mathcal{P})$

Theorem

If $\mathcal{P}$ contains no subset of $d+2$ points on a same hypersphere, then $\mathrm{Del}(\mathcal{P})$ is a triangulation of $\mathcal{P}$

Delaunay Triangulations

Sur la sphère vide (On the empty sphere), Boris Delaunay (1934)

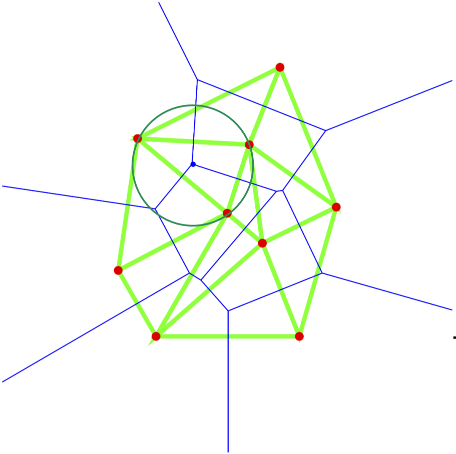

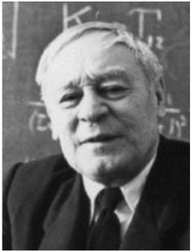

The Delaunay complex $\mathrm{Del}(\mathcal{P})$ is the nerve of $\mathrm{Vor}(\mathcal{P})$

Theorem

If $\mathcal{P}$ contains no subset of $d+2$ points on a same hypersphere, then $\mathrm{Del}(\mathcal{P})$ is a triangulation of $\mathcal{P}$

Correspondence between structures

$$h_{p_i} : x_{d+1} = 2p_i \cdot x - p_i^2 \qquad \hat{p}_i = (p_i, p_i^2) = h_{p_i}^*$$

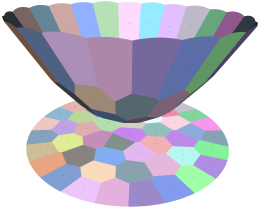

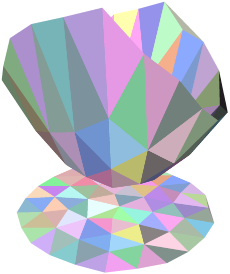

$$\mathcal{V}(\mathcal{P}) = h_{p_1}^+ \cap \ldots \cap h_{p_n}^+ \xrightarrow{\text{duality}} \mathcal{D}(\mathcal{P}) = \text{conv}^-(\{\hat{p}_1, \ldots, \hat{p}_n\})$$

$$\uparrow \qquad\qquad\qquad\qquad \downarrow$$

Voronoi diagram of $\mathcal{P}$ $\xrightarrow{\text{nerve}}$ Delaunay triang. of $\mathcal{P}$

The diagram commutes if $\mathcal{P}$ is in general position wrt spheres

The Voronoi diagram of n points of $\mathbb{R}^d$ has the same combinatorial complexity as the intersection of n half-spaces of $\mathbb{R}^{d+1}$

The Delaunay triangulation of n points of $\mathbb{R}^d$ has the same combinatorial complexity as the convex hull of n points of $\mathbb{R}^{d+1}$

The two complexities are the same by duality : $\Theta(n^{\lceil \frac{d}{2} \rceil})$ [Mc Mullen 1970]

Worst-case : points on the moment curve $\Gamma(t) = \{t, t^2, ..., t^d\} \subset \mathbb{R}^d$

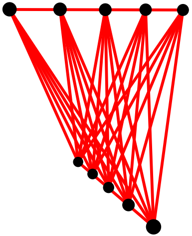

Quadratic in $\mathbb{R}^3$

Corollaries

Algorithmic complexity

Construction of $\mathrm{Del}(\mathcal{P})$, $\mathcal{P} = \{p_1, ..., p_n\} \subset \mathbb{R}^d$

1. Lift the points of $\mathcal{P}$ onto the paraboloid $x_{d+1} = x^2$ of $\mathbb{R}^{d+1}$:
   $p_i \to \hat{p}_i = (p_i, p_i^2)$
2. Compute $\mathrm{conv}(\{\hat{p}_i\})$
3. Project the lower hull $\mathrm{conv}^-(\{\hat{p}_i\})$ onto $\mathbb{R}^d$

Complexity : $\Theta(n \log n + n^{\lceil \frac{d}{2} \rceil})$ [Clarkson & Shor 1989] [Chazelle 1993]

Laguerre (power, weighted) diagrams

$\mathcal{B} = \{b_1, ..., b_n\}$

$D(x, b) = (x - p)^2 - r^2$

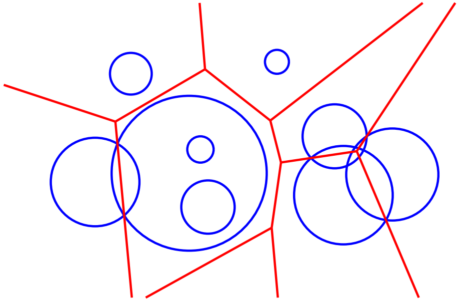

Voronoi cell : $V(b_i) = \{x : D(x, b_i) \leq D(x, b_j) \forall j\}$

Voronoi diagram of $\mathcal{B}$: $=$ { set of cells $V(b_i)$, $\ b_i \in \mathcal{B}$}

Delaunay triangulations of balls

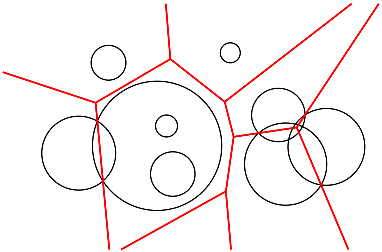

$\mathrm{Vor}(\mathcal{B})$

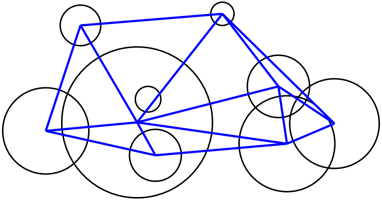

$\mathrm{Del}(\mathcal{B})$ is the nerve of $\mathrm{Vor}(\mathcal{B})$

Theorem

If the balls are in general position, then $\mathrm{Del}(\mathcal{B})$ is a triangulation of a subset $\mathcal{P}' \subseteq \mathcal{P}$ of the points

General position for balls :
no point of $\mathbb{R}^d$ has same power wrt $d+2$ balls

Delaunay triangulations of balls

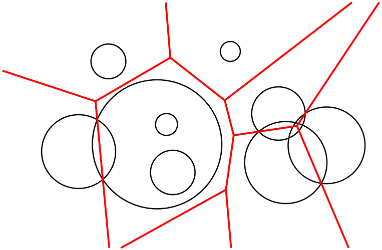

$\mathrm{Vor}(\mathcal{B})$

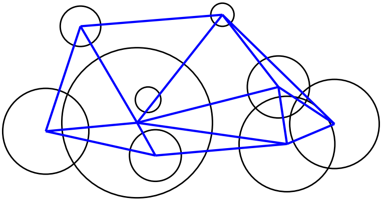

$\mathrm{Del}(\mathcal{B})$ is the nerve of $\mathrm{Vor}(\mathcal{B})$

Theorem

If the balls are in general position, then $\mathrm{Del}(\mathcal{B})$ is a triangulation of a subset $\mathcal{P}' \subseteq \mathcal{P}$ of the points

General position for balls :
no point of $\mathbb{R}^d$ has same power wrt $d+2$ balls

Correspondence between structures

$h_{b_i} : x_{d+1} = 2p_i \cdot x - p_i^2 + r_i^2$

$\hat{b}_i = (p_i, p_i^2 - r_i^2) = h_{b_i}^*$

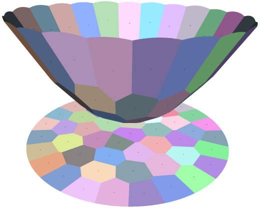

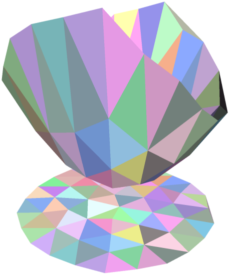

$$\begin{array}{ccc}
\mathcal{V}(\mathcal{B}) = h_{b_1}^+ \cap \ldots \cap h_{b_n}^+ & \xrightarrow{\text{duality}} & \mathcal{D}(\mathcal{B}) = \mathrm{conv}^-(\{\hat{b}_1, \ldots, \hat{b}_n\}) \\
\uparrow & & \downarrow \\
\text{Voronoi diagram of } \mathcal{B} & \xrightarrow{\text{nerve}} & \text{Delaunay triang. of } \mathcal{B}
\end{array}$$

The diagram commutes if $\mathcal{B}$ is in general position

Affine diagrams

Sites + distance functions s.t. the bisectors are hyperplanes

Theorem [Aurenhammer]

Any affine diagram of $\mathbb{R}^d$ is the Laguerre diagram of a set of balls of $\mathbb{R}^d$

Examples :

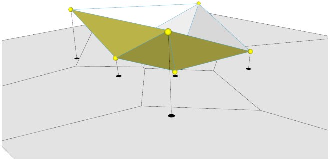

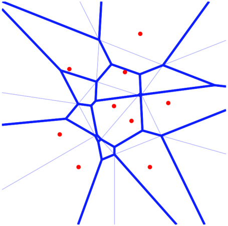

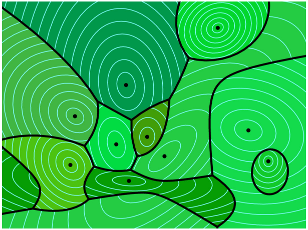

Delaunay triangulation restricted to a union of balls

Alpha-complex [Edelsbrunner et al.]

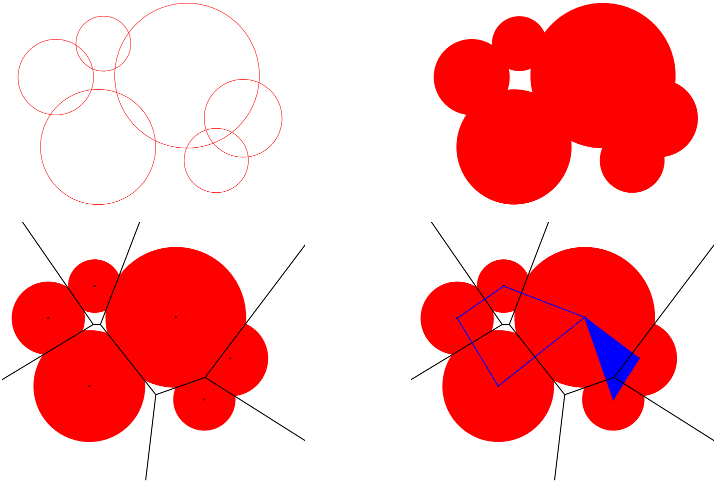

$$C(b) = b \cap V(b) \qquad \mathrm{Vor}_{|U}(\mathcal{B}) = \{f \in \mathrm{Vor}(\mathcal{B}),\quad f \cap U \neq \emptyset\}$$

$$U = \bigcup_{b \in \mathcal{B}} C(b) \qquad \mathrm{Del}_{|U}(\mathcal{B})$$

Discrete metric spaces : Witness Complex [de Silva]

L a finite set of points (landmarks) vertices of the complex

W a dense sample (witnesses) pseudo circumcenters

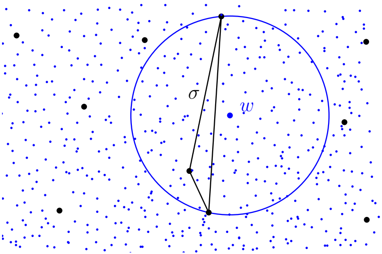

Let σ be a (abstract) simplex with vertices in L, and let $w \in W$. We say that w is a witness of σ if

$$\|w - p\| \leq \|w - q\| \qquad \forall p \in \sigma \text{ and } \forall q \in L \setminus \sigma$$

The witness complex $\mathrm{Wit}(L, W)$ is the complex consisting of all simplexes σ such that for any simplex $\tau \subseteq \sigma$, τ has a witness in W

Discrete metric spaces : Witness Complex [de Silva]

L a finite set of points (landmarks) vertices of the complex

W a dense sample (witnesses) pseudo circumcenters

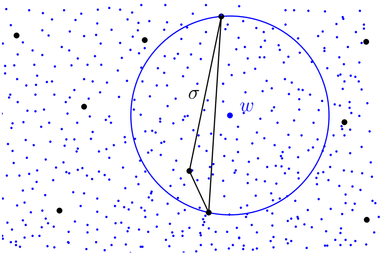

Let σ be a (abstract) simplex with vertices in L, and let $w \in W$. We say that w is a witness of σ if

$$\|w-p\| \leq \|w-q\| \quad \forall p \in \sigma \text{ and } \forall q \in L \setminus \sigma$$

The witness complex $\mathrm{Wit}(L, W)$ is the complex consisting of all simplexes σ such that for any simplex $\tau \subseteq \sigma$, τ has a witness in W

Discrete metric spaces : Witness Complex [de Silva]

L a finite set of points (landmarks) vertices of the complex

W a dense sample (witnesses) pseudo circumcenters

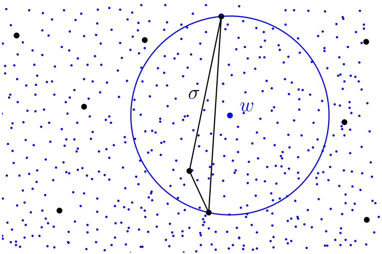

Let σ be a (abstract) simplex with vertices in L, and let $w \in W$. We say that w is a witness of σ if

$$\|w - p\| \leq \|w - q\| \quad \forall p \in \sigma \ \text{ and } \ \forall q \in L \setminus \sigma$$

The witness complex $\mathrm{Wit}(L, W)$ is the complex consisting of all simplexes σ such that for any simplex $\tau \subseteq \sigma$, τ has a witness in W

Easy consequences of the definition

- The witness complex can be defined for any metric space and, in particular, for discrete metric spaces
- If $W' \subseteq W$, then $\mathrm{Wit}(L, W') \subseteq \mathrm{Wit}(L, W)$
- $\mathrm{Del}(L) \subseteq \mathrm{Wit}(L, \mathbb{R}^d)$

Easy consequences of the definition

- The witness complex can be defined for any metric space and, in particular, for discrete metric spaces
- If $W' \subseteq W$, then $\mathrm{Wit}(L, W') \subseteq \mathrm{Wit}(L, W)$
- $\mathrm{Del}(L) \subseteq \mathrm{Wit}(L, \mathbb{R}^d)$

Easy consequences of the definition

- The witness complex can be defined for any metric space and, in particular, for discrete metric spaces
- If $W' \subseteq W$, then $\mathrm{Wit}(L, W') \subseteq \mathrm{Wit}(L, W)$
- $\mathrm{Del}(L) \subseteq \mathrm{Wit}(L, \mathbb{R}^d)$

Identity of $\mathrm{Del}(L)$ and $\mathrm{Wit}(L, \mathbb{R}^d)$ [de Silva 2008]

Theorem : $\mathrm{Wit}(L, W) \subseteq \mathrm{Wit}(L, \mathbb{R}^d) = \mathrm{Del}(L)$

Remarks

- Faces of all dimensions have to be witnessed

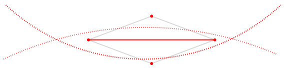

- $\mathrm{Wit}(L, W)$ is embedded in $\mathbb{R}^d$ if L is in general position wrt spheres

Theorem : $\mathrm{Wit}(L, W) \subseteq \mathrm{Wit}(L, \mathbb{R}^d) = \mathrm{Del}(L)$

Remarks

- Faces of all dimensions have to be witnessed

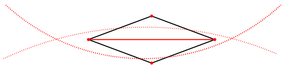

- $\mathrm{Wit}(L, W)$ is embedded in $\mathbb{R}^d$ if L is in general position wrt spheres

Identity of $\mathrm{Del}(L)$ and $\mathrm{Wit}(L, \mathbb{R}^d)$ [de Silva 2008]

Theorem : $\mathrm{Wit}(L, W) \subseteq \mathrm{Wit}(L, \mathbb{R}^d) = \mathrm{Del}(L)$

Remarks

- Faces of all dimensions have to be witnessed

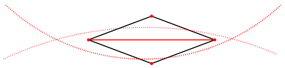

- $\mathrm{Wit}(L, W)$ is embedded in $\mathbb{R}^d$ if L is in general position wrt spheres

Proof of de Silva's theorem

Attali, Edelsbrunner, Mileyko 2007]

$\tau = [p_0, ..., p_k]$ is a k-simplex of $\mathrm{Wit}(L)$ witnessed by a ball B_τ (i.e. $B_\tau \cap L = \tau$)

We prove that $\tau \in \mathrm{Del}(L)$ by induction on k

Clearly true for $k = 0$

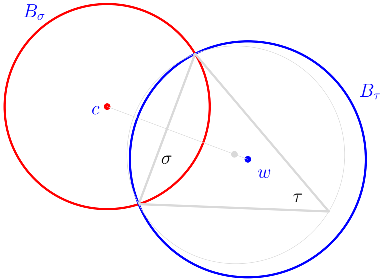

Hyp. : true for $k' \leq k - 1$

$B := B_\tau$

$\sigma := \partial B \cap \tau$, $l := |\sigma|$

// $\sigma \in \mathrm{Del}(L)$ by the hyp.

while $l + 1 = \dim \sigma < k$ do

$B \leftarrow$ the ball centered on $[cw]$ s.t.

- $\sigma \subset \partial B$,
- B witnesses τ
- $|\partial B \cap \tau| = l + 1$

$(B$ witnesses $\tau) \wedge (\tau \subset \partial B) \Rightarrow \tau \in \mathrm{Del}(L)$

Proof of de Silva's theorem

Attali, Edelsbrunner, Mileyko 2007]

$\tau = [p_0, ..., p_k]$ is a k-simplex of $\mathrm{Wit}(L)$ witnessed by a ball B_τ (i.e. $B_\tau \cap L = \tau$)

We prove that $\tau \in \mathrm{Del}(L)$ by induction on k

Clearly true for $k = 0$

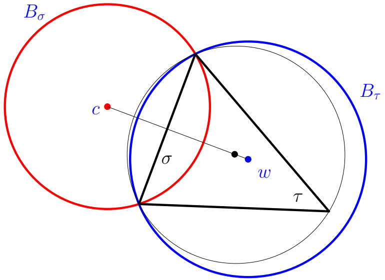

Hyp. : true for $k' \leq k - 1$

$B := B_\tau$

$\sigma := \partial B \cap \tau, \quad l := |\sigma|$

// $\sigma \in \mathrm{Del}(L)$ by the hyp.

while $l + 1 = \dim \sigma < k$ do

$B \leftarrow$ the ball centered on $[cw]$ s.t.

- $\sigma \subset \partial B$,
- B witnesses τ
- $|\partial B \cap \tau| = l + 1$

(B witnesses τ) $\wedge$ ($\tau \subset \partial B$) $\Rightarrow \tau \in \mathrm{Del}(L)$

Case of sampled domains : $\mathrm{Wit}(L, W) \neq \mathrm{Del}(L)$

W a finite set of points $\subset \mathbb{T}^d = \mathbb{R}^d/\mathbb{Z}^d$ (the flat torus of dimension d)

$\mathrm{Wit}(L, W) \neq \mathrm{Del}(L)$, even if W is a dense sample of T^d

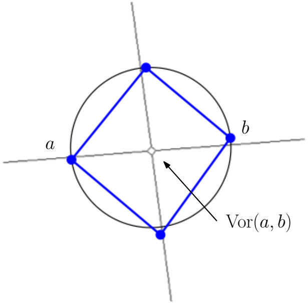

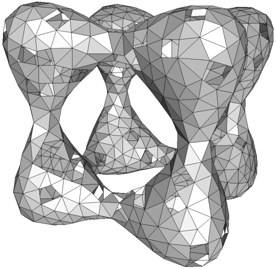

$$[ab] \in \mathrm{Wit}(L, W) \quad \Leftrightarrow \quad \exists p \in W, \quad \mathrm{Vor}_2(a, b) \cap W \neq \emptyset$$

Relaxed witness complex

Alpha-witness Let σ be a simplex with vertices in L. We say that a point $w \in W$ is an α-witness of σ if

$$\|w - p\| \leq \|w - q\| + \alpha \quad \forall p \in \sigma \quad \text{and} \quad \forall q \in L \setminus \sigma$$

Alpha-relaxed witness complex The α-relaxed witness complex $\mathrm{Wit}^{\alpha}(L, W)$ is the maximal simplicial complex with vertex set L whose simplices have an α-witness in W

$\mathrm{Wit}^{0}(L, W) = \mathrm{Wit}(L, W)$

Filtration : $\alpha \leq \beta \quad \Rightarrow \quad \mathrm{Wit}^{\alpha}(L, W) \subseteq \mathrm{Wit}^{\beta}(L, W)$

Lemma Assume that W is ε-dense in $\mathbb{T}^d$ and let $\alpha \geq 2\varepsilon$. Then

$$\mathrm{Wit}(L,W) \subseteq \mathrm{Del}(L) \subseteq \mathrm{Wit}^{\alpha}(L,W)$$

Proof

σ : a d-simplex of $\mathrm{Del}(L)$, c_σ its circumcenter

W ε-dense in $\mathbb{T}^d$ $\exists w \in W$ s.t. $\|c_\sigma - w\| \leq \varepsilon$

For any $p \in \sigma$ and $q \in L \setminus \sigma$, we then have

$$\begin{aligned} \forall p \in \sigma \;\text{ and }\; q \in L \setminus \sigma \qquad \|w-p\| &\leq \|c_\sigma - p\| + \|c_\sigma - w\| \\ &\leq \|c_\sigma - q\| + \|c_\sigma - w\| \\ &\leq \|w - q\| + 2\|c_\sigma - w\| \\ &\leq \|w-q\| + 2\varepsilon \end{aligned}$$

1 Delaunay-like complexes

2 Nets and Delaunay refinement

3 Thick triangulations

Definition of ε-nets (Delone sets)

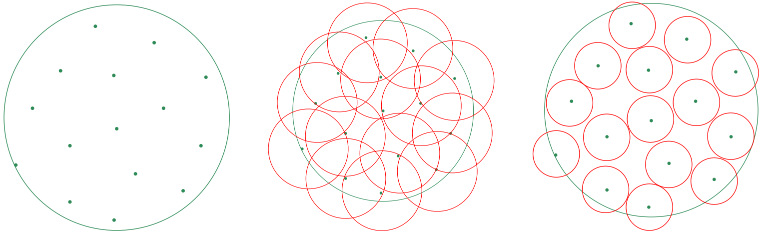

Definition

A finite set of points $\mathcal{P}$ is called an $(\varepsilon, \bar{\eta})$-net of a compact subset $\Omega \subset \mathbb{R}^d$ iff

Density : $\forall x \in \Omega, \exists p \in \mathcal{P} : \|x - p\| \leq \varepsilon$

Separation : $\forall p, q \in \mathcal{P} : \|p - q\| \geq \bar{\eta}\, \varepsilon$

$\mathcal{P}$ is simply called an ε-net if $\bar{\eta}$ is a cst that does not depend on ε

Existence of nets

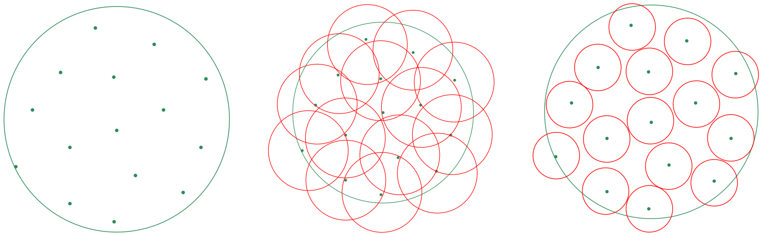

Lemma Any compact subset $\Omega \subset \mathbb{R}^d$ admits an $(\varepsilon, 1)$-net.

Proof

While there exists a point $p \in \Omega$, $d(p, \mathcal{P}) \geq \varepsilon$, insert p in $\mathcal{P}$

Size of $(\varepsilon, \bar{\eta})$-nets

Case of a ball of radius r of $\mathbb{R}^d$

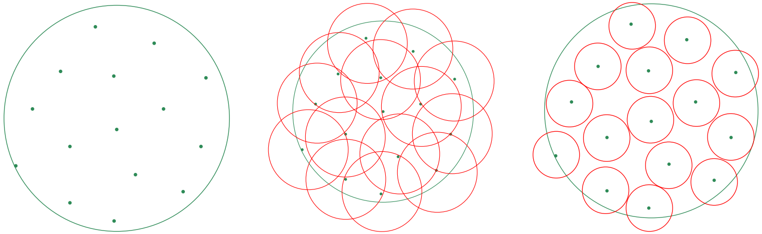

Lemma $$\left(\frac{r}{\varepsilon}\right)^d \leq n(\varepsilon, \bar{\eta}) \leq \left(\left(\frac{4}{\bar{\eta}}\right)^d \times \left(\frac{r}{\varepsilon}\right)^d\right)$$

Proof

Covering : $n \geq \frac{V_d \times r^d}{V_d \times \varepsilon^d} = \left(\frac{r}{\varepsilon}\right)^d$

Packing : $n \leq \frac{V_d \times (r+\frac{\eta}{2})^d}{V_d \times (\frac{\eta}{2})^d} \leq \left(\frac{4r}{\bar{\eta}\varepsilon}\right)^d$

Delaunay triangulations of ε-nets

Lemma

Let $\mathcal{P}$ be an ε-net of $\mathbb{T}^d = \mathbb{R}^d/\mathbb{Z}^d$ (the flat torus of dimension d)
The number of simplices of $\mathrm{Del}(\mathcal{P})$ is $O(|\mathcal{P}|)$

Proof

1. The radius of the CC-ball of any d-simplex is $\leq \varepsilon$

   all the neighbours of a vertex p are in $B(p, 2\varepsilon)$

   their number n_p is bounded by the previous lemma : $n_p \leq \left(\frac{8}{\bar{\eta}}\right)^d = 2^{O(d)}$

2. The number of simplices incident to p is at most the number of faces of the convex hull of n_p points of $\mathbb{R}^d$

$$n_p^{\lfloor \frac{d}{2} \rfloor} = 2^{O(d^2)}$$

Randomized Incremental Construction of the Delaunay triangulation of nets

[B., Devillers, Dutta, Glisse 2018]

$\boxed{\text{Z}}$ A bound on the size of $\mathrm{Del}(\mathcal{P})$ is not sufficient to bound the complexity of its construction

Theorem 1 The expected size of the Delaunay triangulation of a random subset S of a net $\mathcal{P}$ is linear in $|S|$

Theorem 2 If $\mathcal{P}$ is a net of $\mathbb{T}^d$, the standard RIC algorithm computes $\mathrm{Del}(\mathcal{P})$ in $O(|\mathcal{P}| \log |\mathcal{P}|)$ time

Randomized Incremental Construction of the Delaunay triangulation of nets

[B., Devillers, Dutta, Glisse 2018]

$\boxed{\text{Z}}$ A bound on the size of $\mathrm{Del}(\mathcal{P})$ is not sufficient to bound the complexity of its construction

Theorem 1 The expected size of the Delaunay triangulation of a random subset S of a net $\mathcal{P}$ is linear in $|S|$

Theorem 2 If $\mathcal{P}$ is a net of $\mathbb{T}^d$, the standard RIC algorithm computes $\mathrm{Del}(\mathcal{P})$ in $O(|\mathcal{P}| \log |\mathcal{P}|)$ time

Non-uniform nets

A finite set of points $P \subset \Omega$ is a ϕ-net of Ω if there exists two constants c and c' s.t.

Density : $\forall x \in \Omega, \exists p \in P: \quad \|x - p\| \leq c\,\phi(x)$

Separation : $\forall p, q \in P: \quad \|p - q\| \geq c' \max(\phi(p), \phi(q))$

Generation of ϕ-nets and of meshes

Delaunay refinement [Chew 1993, Ruppert 1995, Shewchuk 2002]

Domain : $\Omega = \mathbb{R}^d/\mathbb{Z}^d$ (periodic space)

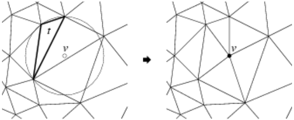

Sizing field $\phi : \Omega \to \mathbb{R}$ = function α-Lipschitz, $\alpha < 1$

$$|\phi(x) - \phi(y)| \leq \alpha \, \|x - y\|$$

$$\forall x \in \Omega, \; 0 < \phi_0 \leq \phi(x)$$

Bad simplex $\sigma : \; \|c_\sigma - p\| > \phi(c_\sigma)$ (c_σ = CC of σ)

Algorithm

INIT construct a (small) initial sample $\mathcal{P}_0 \subset \Omega$

WHILE $\exists$ a bad simplex σ

insert_in_Del(c_σ)

RETURN a ϕ-net and the associated Delaunay triangulation

Generation of ϕ-nets and of meshes

Delaunay refinement [Chew 1993, Ruppert 1995, Shewchuk 2002]

Domain : $\Omega = \mathbb{R}^d / \mathbb{Z}^d$ (periodic space)

Sizing field $\phi : \Omega \to \mathbb{R}$ = function α-Lipschitz, $\alpha < 1$

$$|\phi(x) - \phi(y)| \leq \alpha \, \|x - y\|$$

$$\forall x \in \Omega, \; 0 < \phi_0 \leq \phi(x)$$

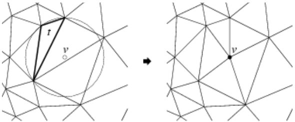

Bad simplex σ : $\|c_\sigma - p\| > \phi(c_\sigma)$ (c_σ = CC of σ)

Algorithm

INIT construct a (small) initial sample $\mathcal{P}_0 \subset \Omega$

WHILE $\exists$ a bad simplex σ

insert_in_Del(c_σ)

RETURN a ϕ-net and the associated Delaunay triangulation

The algorithm returns a ϕ-net of $\Omega = \mathbb{T}^d$

Separation : $\forall p, q \in \mathcal{P}$

$$\begin{aligned} \forall p \in \mathcal{P}, d(p, \mathcal{P} \setminus \{p\}) &= \|p - q\| \\ &\geq \min(\phi(p), \phi(q)) \\ &\geq \phi_0 \end{aligned}$$

$\Rightarrow$ the algorithm stops

Density : $\forall x \in \Omega,\ \exists\, \sigma, \quad d(x, P) \leq \|x - c_\sigma\| \leq \phi(c_\sigma) \overset{(*)}{\leq} \frac{\phi(x)}{1-\alpha}$

$$(*) \quad \phi(c_\sigma) \leq \phi(x) + \alpha\, \|x - c_\sigma\| \leq \phi(x) + \alpha\, \phi(C_\sigma) \quad \Rightarrow \quad \phi(c_\sigma) \leq \frac{\phi(x)}{1 - \alpha}$$

Size of the sample $= \Theta\left(\int_\Omega \frac{dx}{\phi^d(x)}\right)$

Upper bound

$B_p = B(p, r_p), \quad p \in \mathcal{P}, \quad r_p = \frac{\phi(p)}{2(1+\alpha)}$

$$\begin{aligned}
\int_\Omega \frac{dx}{\phi^d(x)} &\geq \sum_p \int_{B_p} \frac{dx}{\phi^d(x)} && \text{(the } B_p \text{ are disjoint)} \\
&\geq \frac{1}{(2+3\alpha)^d} \sum_p \frac{\mathrm{vol}(B_p)}{r_p^d} && (\phi(x) \leq \phi(p) + \alpha\,\|p - x\| \\
& && \quad \leq 2(1+\alpha)\,r_p + \alpha\,r_p) \\
&= \frac{V_d}{(2+3\alpha)^d}\,|\mathcal{P}| && (V_d = \text{vol unit ball of } \mathbb{R}^d) \\
&= C\,|\mathcal{P}|
\end{aligned}$$

Lower bound

- use the balls $B'_p(p, \frac{\phi(p)}{1-\alpha})$ that cover Ω

Size of the sample $= \Theta\left(\int_\Omega \frac{dx}{\phi^d(x)}\right)$

Upper bound

$B_p = B(p, r_p), \quad p \in \mathcal{P}, \quad r_p = \frac{\phi(p)}{2(1+\alpha)}$

$$\int_\Omega \frac{dx}{\phi^d(x)} \geq \sum_p \int_{B_p} \frac{dx}{\phi^d(x)} \qquad \text{(the } B_p \text{ are disjoint)}$$

$$\geq \frac{1}{(2+3\alpha)^d} \sum_p \frac{\mathrm{vol}(B_p)}{r_p^d} \qquad (\phi(x) \leq \phi(p) + \alpha\,\|p-x\| \leq 2(1+\alpha)\,r_p + \alpha\,r_p)$$

$$= \frac{V_d}{(2+3\alpha)^d}\,|\mathcal{P}| \qquad (V_d = \text{vol unit ball of } \mathbb{R}^d)$$

$$= C\,|\mathcal{P}|$$

Lower bound

- use the balls $B'_p(p, \frac{\phi(p)}{1-\alpha})$ that cover Ω

Bound on the angles (2d case)

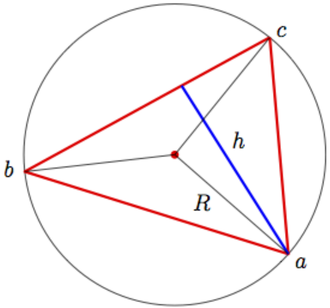

$\|a - c\| \geq \phi(a)$ (a inserted after c)

$R \leq \phi(c_{abc})$ (c_{abc} not inserted)

$\sin b = \frac{\|a-c\|}{2R} \geq \frac{\phi(a)}{2\phi(c_{abc})}$

$$\phi(c_{abc}) \leq \phi(a) + \alpha\,\|c_{abc} - a\| \leq \phi(a) + \alpha\,\phi(c_{abc}) \quad \Rightarrow \quad \phi(c_{abc}) \leq \frac{\phi(a)}{1-\alpha}$$

$$\Rightarrow \quad \sin b \geq \frac{1-\alpha}{2}$$

Results

$$\phi(x) = \phi_0 + \alpha d(x, \partial\Omega)$$

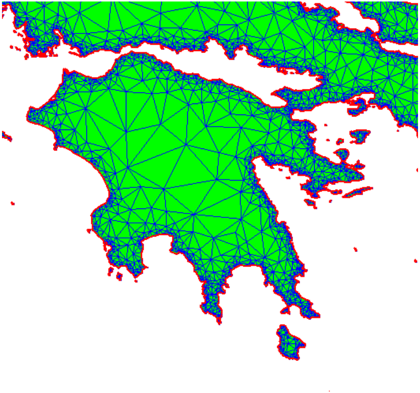

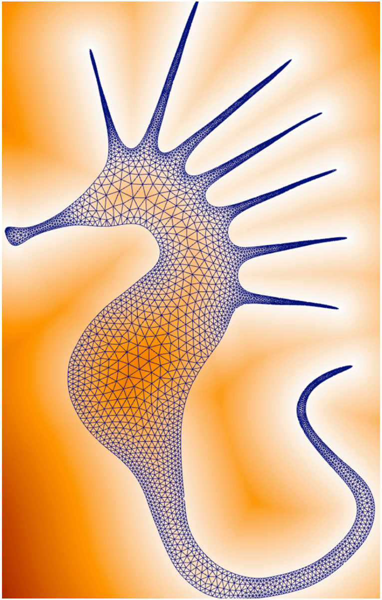

1 Delaunay-like complexes

2 Nets and Delaunay refinement

3 Thick triangulations

Delaunay refinement in higher dimensions

Returns nets but not necessarily thick simplices

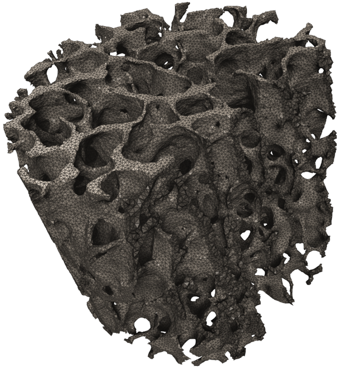

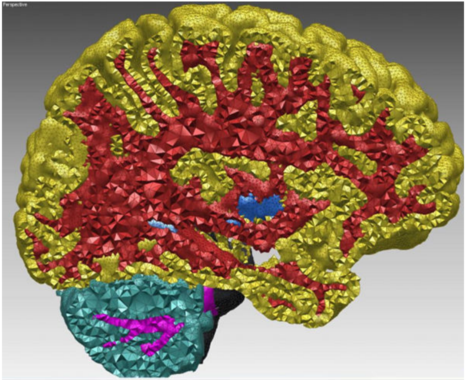

Delaunay refinement allows to control the circumradii of the simplices (density) and the separation

but not the thickness of the simplices except in 2d

A net whose Delaunay triangulation is not thick

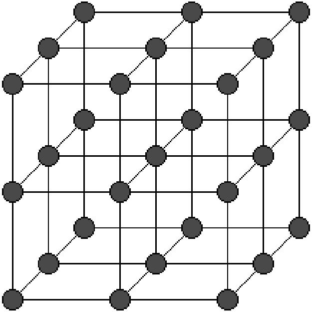

- Each squared face can be circumscribed by an empty ball
- This remains true if we slightly perturb the points creating arbitrarily flat simplices

The long quest for thick triangulations

Differential topology [Cairns], [Whitehead], [Whitney], [Munkres]

Differential geometry [Cheeger et al.]

Geometric theory of functions [Peltonen], [Saucan]

Quality of simplices

Altitudes

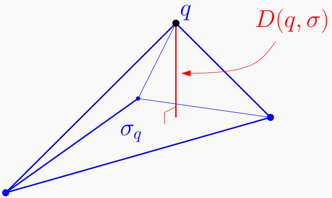

If σ_q is the face opposite to q in σ, the *altitude* of q in σ is

$$D(q, \sigma) = d(q, \mathrm{aff}(\sigma_q)),$$

Definition (Thickness)

The *thickness* of a j-simplex σ of diameter $\Delta(\sigma)$ is

$$\Theta(\sigma) = \begin{cases} 1 & \text{if } j = 0 \\ \min_{p \in \sigma} \frac{D(p, \sigma)}{j\Delta(\sigma)} & \text{otherwise} \end{cases}$$

Quality of simplices

Altitudes

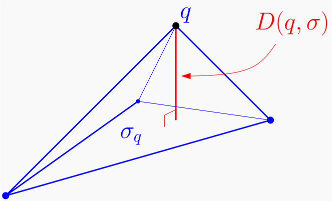

If σ_q is the face opposite to q in σ, the *altitude* of q in σ is

$$D(q,\sigma) = d(q, \mathrm{aff}(\sigma_q)),$$

Definition (Thickness)

The *thickness* of a j-simplex σ of diameter $\Delta(\sigma)$ is

$$\Theta(\sigma) = \begin{cases} 1 & \text{if } j = 0 \\ \min_{p\in\sigma} \frac{D(p,\sigma)}{j\Delta(\sigma)} & \text{otherwise} \end{cases}$$

Protection

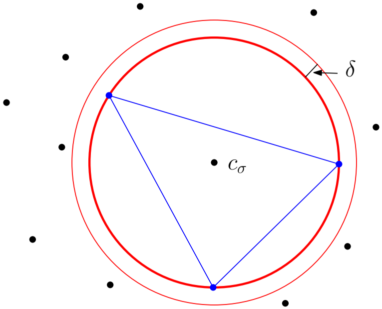

δ-protection

A Delaunay simplex $\sigma \subset L$ is δ-protected if

$$\|c_\sigma - q\| > \|c_\sigma - p\| + \delta \qquad \forall p \in \sigma \text{ and } \forall q \in L \setminus \sigma.$$

Protection implies separation and thickness

Soit $\mathcal{P}$ un $(\varepsilon, \bar{\eta})$-net, i.e.

- $\forall x \in \Omega, \quad d(x, P) \leq \varepsilon$
- $\forall p, q \in P, \quad \|p - q\| \geq \bar{\eta}\varepsilon$

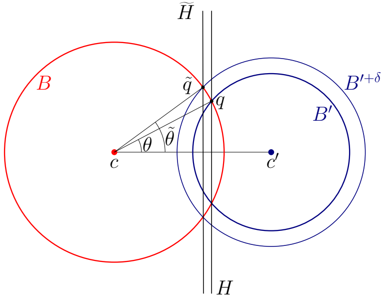

if all d-simplices de $\mathrm{Del}(P)$ are $\bar{\delta}\varepsilon$-protected, then

- Separation of P : $\bar{\eta} \geq \bar{\delta}$
- Thickness : $\forall \sigma \in \mathrm{Del}(P), \quad \Theta(\sigma) \geq \frac{\bar{\delta}^2}{8d}$

Bad configurations

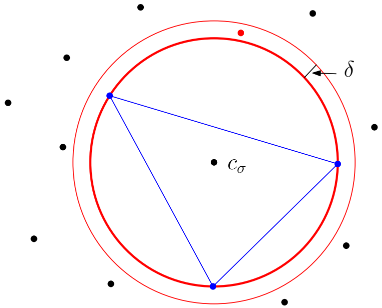

Bad configuration $\phi = (\sigma, p)$

- σ is a d-simplex with $R_\sigma \leq \varepsilon$
- $p \in Z_\delta(\sigma)$ where $Z_\delta(\sigma) = B(c_\sigma, R_\sigma + \delta) \setminus B(c_\sigma, R_\sigma)$

Protecting Delaunay simplices using random perturbations

Random variables : $\mathcal{P}'$ a set of random points $\{p'_1, ..., p'_n\}$ where $p'_i \in B(p_i, \rho), p_i \in \mathcal{P}$

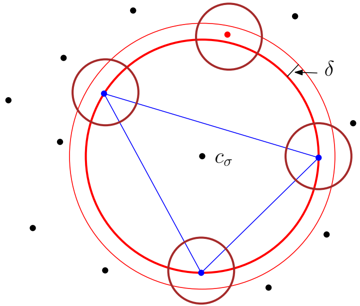

Algorithm

Input : $\mathcal{P} \in \mathbb{T}^d$, ρ, δ

while $\exists$ bad configuration $\phi' = (\sigma', p')$ **do**

resample the points of ϕ'

update $\mathrm{Del}(\mathcal{P}')$

Return P' and $\mathrm{Del}(\mathcal{P}')$

Let $A_1, ..., A_N$ be a set of bad events
each occurs with $\text{proba}(A_i) \leq p < 1$

Question : what is the probability that none of the events occur ?

The (easy) case of independent events

$$\text{proba}(\neg A_1 \wedge ... \wedge \neg A_N) \geq (1-p)^N > 0$$

What happens if the events are weakly dependent ?

Let $A_1, ..., A_N$ be a set of bad events
each occurs with $\text{proba}(A_i) \leq p < 1$

Question : what is the probability that none of the events occur ?

The (easy) case of independent events

$$\text{proba}(\neg A_1 \wedge ... \wedge \neg A_N) \geq (1-p)^N > 0$$

What happens if the events are weakly dependent ?

Let $A_1, ..., A_N$ be a set of bad events
each occurs with $\text{proba}(A_i) \leq p < 1$

Question : what is the probability that none of the events occur ?

The (easy) case of independent events

$$\text{proba}(\neg A_1 \wedge ... \wedge \neg A_N) \geq (1-p)^N > 0$$

What happens if the events are weakly dependent ?

Lovász local lemma

[Lovász & Erdös 1975]

If, for $i = 1, ..., N$,

1. A_i is independent from all events except $\leq \Gamma$ of them
2. $\mathrm{proba}(A_i) \leq \frac{1}{e\,(\Gamma+1)}$ $\qquad e = 2.718...$

then

$$\mathrm{proba}(\neg A_1 \wedge ... \wedge \neg A_N) > 0$$

A constructive version of LLL

[Moser and Tardos 2010]

$\mathcal{P}$ a finite set of independent random variables

$\mathcal{A}$ a finite set pf events determined by the values of some of the variables of $\mathcal{P}$

Two events are independent iff they don't share any variable

Algorithm

for all $p \in \mathcal{P}$ **do**
  $v_p \leftarrow$ random evaluation of p ;

while some events of $\mathcal{A}$ occur for $(p = v_p, p \in \mathcal{P})$ **do**
  select (arbitrarily) such an event $A \in \mathcal{A}$;
  for all $p \in$ variables(A) **do**
    $v_p \leftarrow$ a new random evaluation of p ;

return $(v_p)_{p \in \mathcal{P}}$;

A constructive version of LLL

[Moser and Tardos 2010]

$\mathcal{P}$ a finite set of independent random variables

$\mathcal{A}$ a finite set pf events determined by the values of some of the variables of $\mathcal{P}$

Two events are independent iff they don't share any variable

Algorithm

for all $p \in \mathcal{P}$ **do**

  $v_p \leftarrow$ random evaluation of p ;

while some events of $\mathcal{A}$ occur for $(p = v_p, p \in \mathcal{P})$ **do**

  select (arbitrarily) such an event $A \in \mathcal{A}$;

  for all $p \in \text{variables}(A)$ **do**

    $v_p \leftarrow$ a new random evaluation of p ;

return $(v_p)_{p \in \mathcal{P}}$;

Moser and Tardos theorem

Hypothesis : for all events A_i, $i \in [1, N]$

1. A_i is independent from all events except $\leq \Gamma$ of them
2. $\text{proba}(A_i) \leq \frac{1}{e\,(\Gamma+1)}$ $\qquad e = 2.718...$

Theorem

- The algorithm assigns values to the variables $\mathcal{P}$ s.t. no event of $\mathcal{A}$ occurs
- The algorithm resamples an event $A \in \mathcal{A}$ at most $\boxed{\frac{1}{\Gamma}}$ times on expectation before finding such an assignment
- The expected total number of resamplings is at most $\boxed{\frac{N}{\Gamma}}$

Application to the protection algorithm

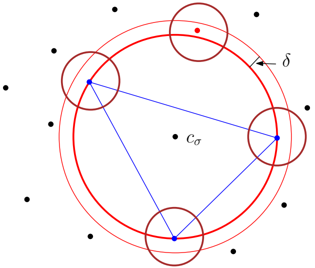

We need

- to bound the probability p that an event (bad configuration) occurs
- to bound the maximal number Γ of events that are not independent from a given event
- to satisfy : $p \leq \frac{1}{e\,(\Gamma+1)}$

Probability that (σ, p) is a bad configuration

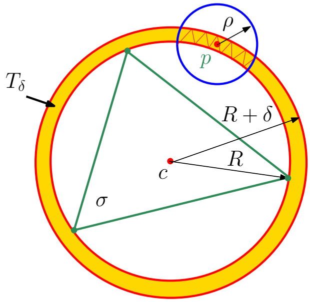

$B(c,R)$ the CC-ball of σ

Protection zone of σ
$T_\delta = B(c, R+\delta) \setminus B(c,R)$

Picking region of p : $B(p,\rho)$

$$\text{proba}((\sigma,p) \text{ is a bad configuration}) = \frac{\text{vol}(T_\delta \cap B(p,\rho))}{\text{vol}(B(p,\rho))}$$

$$\leq C\, \frac{\delta\, \rho^{d-1}}{\rho^d}$$

$$= C\, \frac{\delta}{\rho}$$

Bound on Γ

Lemma : An event is independent from all the events except at most Γ of them, where Γ is a constant that depends on $\bar{\eta}$, $\bar{\rho}$, $\bar{\delta}$ and d

Proof :

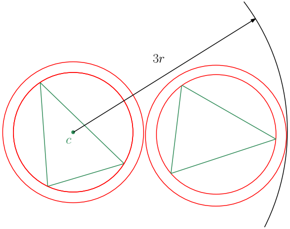

- **Locality** : if 2 configurations (σ, p) and (σ', p') have a vertex in common, all the vertices of σ' belong to the ball $B(c, 3r)$ where $r = \varepsilon + \rho + \delta$

- **Packing** : Since $\mathcal{P}$ is $(\eta - 2\rho)$-separated,

$$\Gamma = \left(\frac{3r + \frac{\eta - 2\rho}{2}}{\frac{\eta - 2\rho}{2}} \right)^{d(d+2)} = \left(1 + 6\, \frac{(1 + \bar{\rho} + \bar{\delta})\,(1 + \bar{\rho})}{\bar{\eta} - 2\bar{\rho}} \right)^{d(d+2)} = O\left(2^{d^2}\right)$$

Main result

Under the condition

$$\frac{eC}{\sqrt{d}}\left(\Gamma+1\right)\bar{\delta} \leq \bar{\rho} \leq \frac{\bar{\eta}}{4}$$

the algorithm stops

Guarantees on the output

- $d_H(\mathcal{P}, \mathcal{P}') \leq \rho$
- the d-simplices of $\mathrm{Del}(\mathcal{P}')$ are δ-protected
- and thus have a positive bound on thickness that does not depend on ε

Main result

Under the condition

$$\frac{eC}{\sqrt{d}}\,(\Gamma+1)\,\bar{\delta} \leq \bar{\rho} \leq \frac{\bar{\eta}}{4}$$

the algorithm stops

Guarantees on the output

- $d_H(\mathcal{P}, \mathcal{P}') \leq \rho$
- the d-simplices of $\mathrm{Del}(\mathcal{P}')$ are δ-protected
- and thus have a positive bound on thickness that does not depend on ε

Complexity of the algorithm

- The number of resampling executed by the algorithm is at most

$$\frac{C'n}{\Gamma} \leq C''n$$

  where C'' depends on $\bar{\eta}$, $\bar{\rho}$, $\bar{\delta}$ and (exponentially) d

- Each resampling consists in perturbing $O(1)$ points
- Updating the Delaunay triangulation after each resampling takes time $O(1)$
- The expected complexity is linear in the number of points

From theory to practice

- The bound on Γ is huge, which leads to weak guarantees on protection and thickness
- The behaviour of the algorithm is better in practice than predicted by the theory
- Usefull heuristics have been proposed

Energy

$$E_{ODT} = \|f - \hat{f}\|_{L^1}$$

(volume between the paraboloid and the approximation)

Minimizing E_{ODT} for given $\mathcal{P}$: $\mathrm{Del}(\mathcal{P})$

Optimal position of vertex x_i (the others being fixed)

$$\boxed{\boldsymbol{x}_i^* = \frac{\sum_{t \in N(\boldsymbol{x}_i)} |t|\, \boldsymbol{c}_t}{\sum_{t \in N(\boldsymbol{x}_i)} |t|}}$$

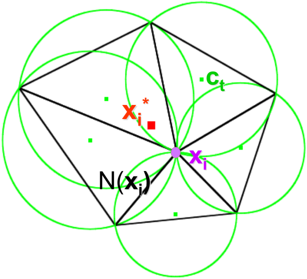

Optimal Delaunay triangulation

Lloyd's relaxation [Chen 2004], [Alliez et al. 2005]

Algorithm

1. Initialize $\mathcal{P}$ with n points in Ω
2. Compute $\mathrm{Del}(\mathcal{P})$
3. For each interior vertex x_i,
   $x_i \leftarrow x_i^*$ and update $\mathrm{Del}(\mathcal{P})$
4. If the error is small enough, stop ; otherwise go to Step 2

$$\boldsymbol{x}_i^* = \frac{\sum_{t \in N(\boldsymbol{x}_i)} |\boldsymbol{t}|\, \boldsymbol{c}_t}{\sum_{t \in N(\boldsymbol{x}_i)} |\boldsymbol{t}|}$$

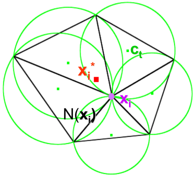

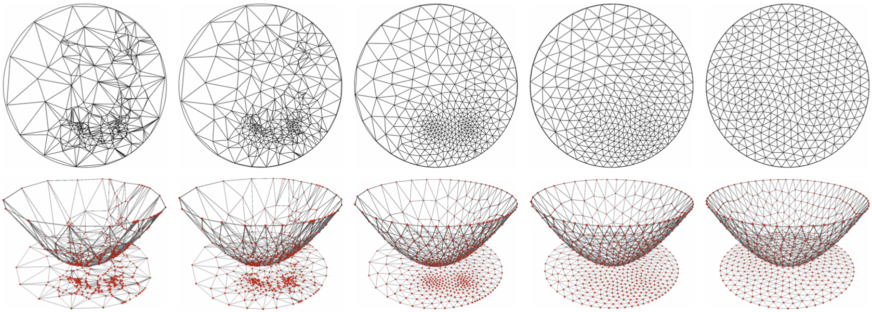

Combine perturbation of vertices and mesh optimization

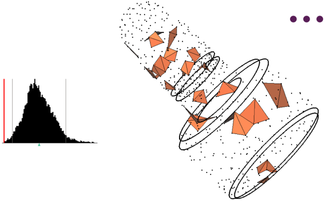

Interleaved
(angles > 4 deg)

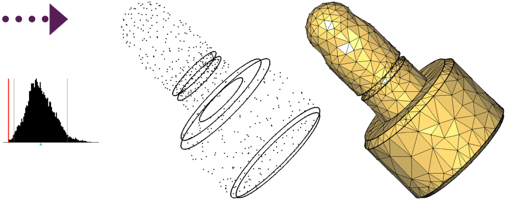

Slivers Removed
(angles > 10 deg)

[Tournois, Srinivasan, Alliez, 2009]

Results

[Tournois et al, 2009]

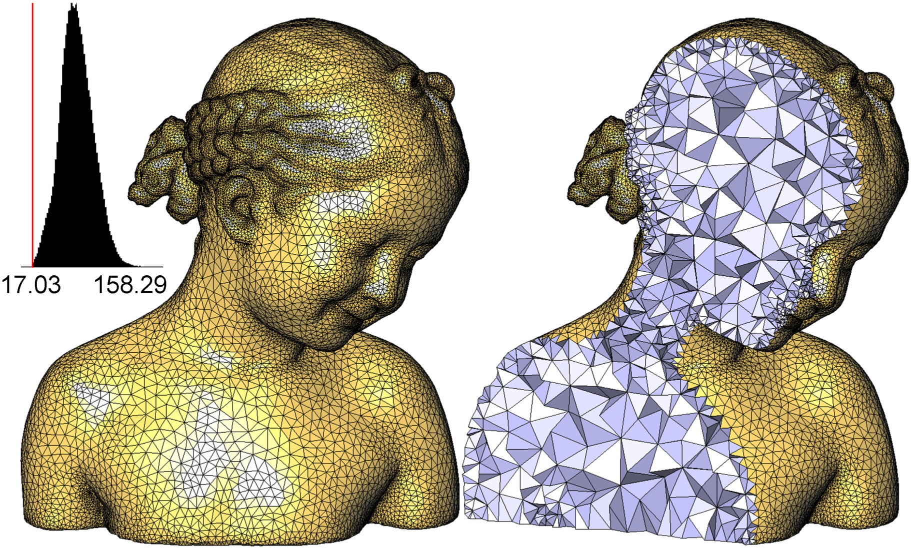

13% fewer points than DR

Results

[Tournois et al, 2009]

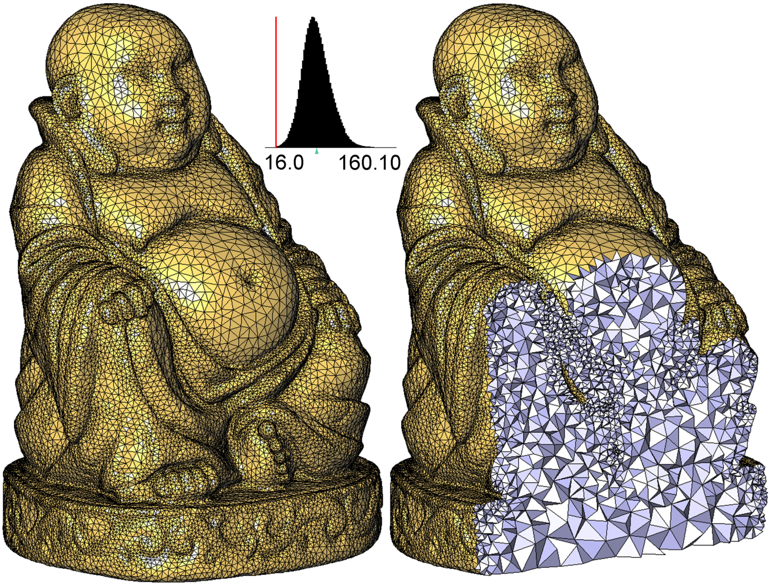

Results

[Tournois et al, 2009]

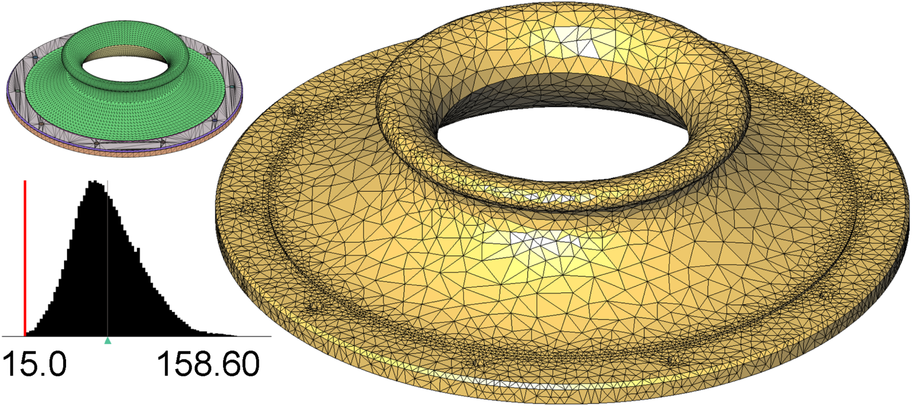

Conclusions

- The only use of LLL in Computational Geometry ?
- Instead of perturbing the points, one can weight the points (i.e. perturb the metric)
- A nice theoretical result but weak bounds
- Efficient heuristics
- Many applications
  - 3D mesh generation (sliver removal)
  - Construction of DT using only predicates of degree 2
  - More tomorrow

Open problems

- Can we obtain better bounds on protection ?
- Thick triangulation of a bounded domain (e.g. a simplex) ?