Surface mesh generation techniques with guaranteed properties

Trial lecture

Christopher Dyken

June 2, 2008

Problem statement

Given a surface $\mathcal{S} = \{x \in \mathbb{R}^3 : f(x) = 0\}$,
generate a mesh $\mathcal{M}$ that represents $\mathcal{S}$.

- Topological properties:
  - *Manifold surface*: Is the resulting surface $\mathcal{M}$ a valid surface?
  - *Topological type*: Are $\mathcal{S}$ and $\mathcal{M}$ the same type of shape?
- Geometric properties:
  - *Approximation error*. How good an approximation of $\mathcal{S}$ is $\mathcal{M}$?
  - *Triangle density*. We want a minimal amount of triangles, but "enough" triangles in difficult areas.
  - *Shape of triangles*. Avoid long thin triangles.
- Algorithmic properties:
  - *Termination*. For what kind of input is the algorithm guaranteed to terminate?

The notion that two objects are of "equal type"

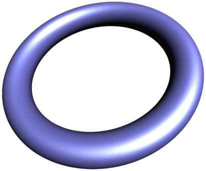

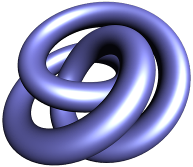

$\mathcal{M}$ and $\mathcal{S}$ are homeomorphic

A *homeomorphism* between $\mathcal{S}$ and $\mathcal{M}$ is a continuous bijection with a continuous inverse:

Points close on $\mathcal{S}$ correspond to points close on $\mathcal{M}$.

A homeomorphism doesn't imply a continuous *deformation* between $\mathcal{S}$ and $\mathcal{M}$, but an isotopy does:

$\mathcal{M}$ and $\mathcal{S}$ are isotopic

An *isotopy* is a continuous map $\gamma(\cdot, t)$,

$$\gamma : \mathcal{S} \times [0,1] \to \mathbb{R}^3, \quad \gamma(\mathcal{S}, 0) = \mathcal{S}, \quad \gamma(\mathcal{S}, 1) = \mathcal{M},$$

that is a homeomorphism for any fixed $t \in [0,1]$.

If $\gamma : \mathbb{R}^3 \times [0,1] \to \mathbb{R}^3$, γ is an *ambient* isotopy.

Regular subdivision of $\mathbb{R}^3$

Lorensen & Cline,
Marching cubes: A high res. 3D surface construction algorithm

Marching cubes assumes f is sampled on a regular 3D grid:

Grid cells have 8 neighbouring samples.

Label samples at corners:
$f < 0$: inside, $f \geq 0$: outside.

256 different cell label configurations,
can be reduced to 15 using symmetry.

Let configuration of labels determine
tessellation of surface inside cell.

Tessellation stored in a fixed table,
edge intersections found using linear interpolation.

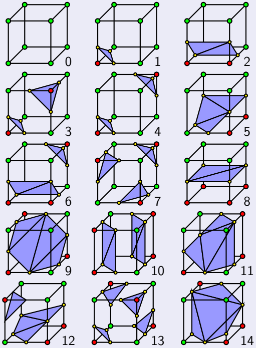

Labels alone cannot determine full topology in all cases:

E.g. case 6 is *ambigious*:
Two choices of connecting face.

Inconsistent choices produce holes in the resulting surface!

And should diagonally opposing corners be connected?

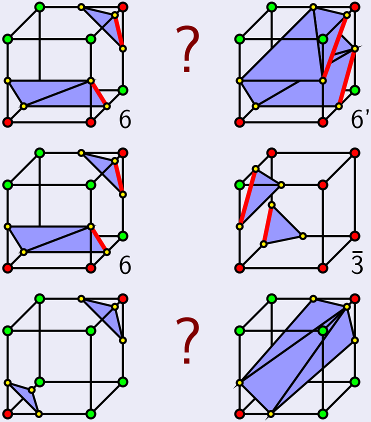

At least $\mathcal{M}$ should be a consistent 2-manifold.

Even better if $\mathcal{M}$ and $\mathcal{S}$ are homeomorphic, or even isotopic...

Montani, Scateni, & Scopigno,
A modified look-up table for implicit disambiguation of MC

- Sacrifice symmetry and extend table of cell configurations.

Nielson & Harmann,
The asymptotic decider: Resolving the ambiguity in MC

- Intersects asymptotic lines of iso-curves of the bilinear interpolant.

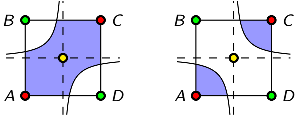

Nielson,
On marching cubes

- DeVella's necklace $DeV(T)$ is intersec. of asymptotic planes.
- If $DeV(T)$ exists and is inside cube $\rightarrow$ *tunnel*.
- 68 basic cases (uses only rotation),
  some need internal vertices, taken from $DeV(T)$.

Regular subdivision schemes:

- Modified table guarantees to a consistent 2-manifold.
- Asymptotic decider resolves ambiguities consistently and corners on faces are connected as the bilinear interpolant.
- On marching cubes consistently connect diagonally opposing corners as the trilinear interpolant.

Problems with regular subdivision:

- No guarantee that $\mathcal{M}$ and $\mathcal{S}$ are of the same type.
- Grid is too coarse: Miss small features.
- Grid is too fine: An excessive amount of triangles.

A strategy is then to *adaptively* subdivide the grid:

- Subdivide until $\mathcal{S}$ is "simple enough" inside a cell.
- Better control on approximation error (bounded by cell size).
- Better control on triangle distribution (governed by cell size).

Snyder's adaptive refinement algorithm

Subdivide boxes until $\mathcal{S}$ is *globally parameterizable* inside each box:

$\mathcal{S}$ is globally parameterizable over X if the planar projection of

$$\mathcal{S}|_X = \{(x_0, x_1, x_2) \in X : \quad f(x_0, x_1, x_2) = 0\}$$

along the axis x_i has no fold-overs, that is, $\frac{\partial f}{\partial x_i} \neq 0$ in X.

This condition is checked using *interval arithmetics*:

Instead of using a *single scalar* x, we calculate with an *interval* $[a, b]$.

Thus, $f([a, b])$ is the interval f takes on over the range $[a, b]$.

We *assume* that

If $[a, b] \to 0$ then $f([a, b]) \to 0$,

which is not always the case.

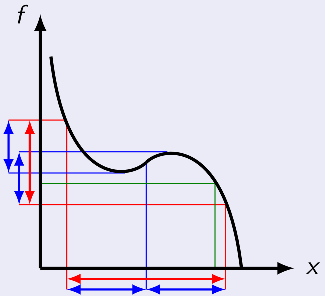

Initialize $\mathbb{A}$ with bounding box and while there exist a box X in $\mathbb{A}$:

- if $0 \notin f(X)$, then X is void of surface and is discarded.
- if $0 \notin \frac{\partial f}{\partial x_i}(\mathcal{S}|_X)$ then $\mathcal{S}|_X$ is globally parameterizable, put in $\mathbb{B}$.
- Otherwise, subdivide X and put subdivided boxes in $\mathbb{A}$.

Then, mesh each box x in $\mathbb{B}$, sorted from small to large:

- $\mathcal{S}|_X$ is globally parameterizable along x_i,
  can be projected onto the $\{x_0, x_1, x_2\} \backslash \{x_i\}$-plane w/o folding
- Intersect with side walls s.t. curves look as side was top-side.
- Triangulate projected regions.
- Propogate face-surface configuration to neighbouring boxes.

If no generated cube face are *tangent to $\mathcal{S}$*:

Snyder's algorithm *terminates*.

Adaptive refinement guided by small normal variation

Pantinga and Vegter,
Isotopic approximation of implicit curves and surfaces.

Small normal variation is a stronger condition,
but relaxes requirement on projection direction:

The small normal variation condition requires that

$$\langle \nabla f(a), \nabla f(b) \rangle \geq 0, \qquad \forall a, b \in \mathcal{S}|_X,$$

i.e., all normal vectors are inside a 90° cone.

Property implies global parameterizability along cone axis:

$\mathcal{S}|_X$ can be projected along cone axis onto a plane w/o foldover

Small normal variation and mean value theorem bound the curve:

The *curvature is limited*, and cells are *equally-sided cubes*, the surface cannot escape "too far" through a neighbouring cell.

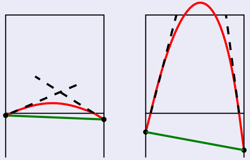

The surface can be *"isotopically pushed"* so that $\mathcal{M}$ intersects an edge only once.

An equally-sided bounding box is adaptively refined as an oct-tree:

The tree is balanced such that the size of adjacent cells maximally differ by a factor of two.

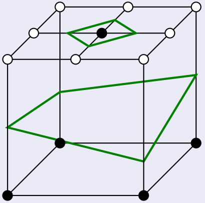

This reduces the number of configurations.

On faces, edge intersections are joined, and for all cases but one, the cell contains a *single loop*.

Start with bounding box, subdivide oct-tree until every leaf X is

- either void of $\mathcal{S}$
  (check if $f(X) \neq 0$ using interval arithmetics)
- or $\mathcal{S}|_X$ satisfies the small normal variation criterion
  (use interval arithmetics)
- and boxes are balanced
  (two adjacent boxes differ maximally in size by a factor of 2).

Then, build $\mathcal{M}$ by meshing the cells in the oct-tree:

- All edges with sign-changes get a vertex inserted.
- For each face in the oct-tree, connect vertices.
- For each leaf cell, connect loops of edges.

$\mathcal{S}$ continuous & non-degenerate & interval arithmetics converges:

Algorithm terminates and $\mathcal{M}$ and $\mathcal{S}$ are *isotopic*.

...but what about triangle shapes?

Delaunay refinement in the plane

Farthest point Delaunay refinement *improves triangle shapes*:

While a triangle T of low quality exists:

- Insert a vertex at the circumcircle of T.
- Circumcircle of T is no longer empty, T will be retriangulated.

Let r be the circumradius and l be the shortest edge of T.

Bounds on r/l implies bounds on the smallest angle θ of T:

$$\frac{r}{l} = \frac{1}{2\sin(\theta)}, \quad \text{and thus,} \quad \frac{r}{l} > B \quad \Longleftrightarrow \quad \theta < \theta_{\text{min}}.$$

When T is refined:

- Three new edges have length r, the rest are longer.
- New edges are at least Bl long.
- $B > 1$, i.e. $\theta_{\text{min}} < 30^\circ$: no new edges shorter than l.

How to define a Delaunay triangulation on a surface?

Chew,
Guaranteed-quality mesh generation for curved surface.

Generalize circle criterium using the surface Delaunay ball:

The *surface Delaunay ball* for a triangle T in a mesh $\mathcal{M}$ of set of points $P \subset \mathcal{S}$ is the sphere through the corners of T and with its center on $\mathcal{S}$.

Generalization is consistent for "reasonable surfaces":

If $\nabla \mathcal{S}$ over two adjacent triangles is inside a $\frac{\pi}{2}$-cone, apices of the two triangles are consistently outside their opposing circumcircle.

And calculating the criterium amounts to intersecting $\mathcal{S}$ with a line:

Centers of all spheres circumscribing T lies on a line.

Surface Delaunay refinement procedure:

- Initialize with a coarse mesh $\mathcal{M}$ of points $P \subset \mathcal{S}$.
- Build constrained Delaunay triangulation.
- While not finished:
  - Find triangles that either
    - has a minimal angle $< 30^\circ$, or
    - violates a user-specified size criterion.
  - Insert circumcentre of triangle with the largest circumcircle.
  - Update triangulation.

It is assumed that normals over triangles are inside a $\frac{\pi}{2}$-cone.

If algorithm halts:

- No interior triangle with minimum angle $< 30^\circ$.
- No triangle is larger than user-specified size criterion.

But *no* guarantees on topological relationships of $\mathcal{M}$ and $\mathcal{S}$:

Surface-based schemes have trouble with topology changes!

Using the 3D Delaunay triangulation

Boissonnat and Oudot,
Provably good sampling and meshing of surfaces.

The *restricted Delaunay triangulation (RDT)*
is a subset of the 3D Delaunay triangulation:

The RDT $\mathcal{M}$ for a set of points $P \subset \mathcal{S}$ is the set of faces from the $3D$ Delaunay triangulation whose dual Voronoi-edges intersects $\mathcal{S}$.

Surface Delaunay balls are empty in a RDT:

A triangle T in a RDT is characterized by that the sphere through the corners of T with center on $\mathcal{S}$ is empty.

We know how to get a surface from P...

Can we guarantee that $\mathcal{M}$ and $\mathcal{S}$ are isotopic?

Yes! If the samples of P are dense enough.

The sample density of a *ψ-sample* is bound by $\psi : \mathcal{S} \rightarrow \mathbb{R}^+$:

For any point x on $\mathcal{S}$, there is a point $p \in P$ maximally $\psi(x)$ away.

The density of P is compared to the *local feature size*:

The $\mathrm{lfs}(x)$ is the Euclidean distance from x to the medial axis, and if density of P is bound by $\psi(x) = \epsilon\, \mathrm{lfs}(x)$, then P is an *ϵ-sample*.

Weak ϵ-samples only require condition on Delaunay balls centers.

The following theorem guarantees an isotopic mesh:

Theorem (Amenta & Bern, Boissonnat & Oudot)

If P is a weak ϵ-sample with $\epsilon < 0.1$ and $\mathcal{M} = \mathrm{RDT}(P)$
with at least a triangle on every component of $\mathcal{S}$,
then $\mathcal{M}$ is homeomorphic and ambient isotopic to $\mathcal{S}$.

Algorithm for building an ϵ-sample P:

- Initialize with at least one triangle on each component of $\mathcal{S}$.
- While not finished:
  - Intersect every Voronoi-edge with $\mathcal{S}$.
  - If intersection x exists,
    it is the center of a surface Delaunay ball with radius r.
  - If $r \geq \psi(x)$, insert x into P and update $\mathcal{M}$.

For 1-Lipschitz ψ where $0 < \psi(x) < \epsilon\,\mathrm{lfs}(x)$:

- The result is a weak ψ-sample.
- If $\epsilon < 0.1$ then $\mathcal{M}$ is ambient isotopic to $\mathcal{S}$
- Number of points bounded so the algorithm terminates:
  $$|P| < O(H(\psi, \mathcal{S})), \quad H(\psi, \mathcal{S}) := \int_{x \in \mathcal{S}} \frac{1}{\psi(x)^2} dx$$
- Combines with θ_{min}-predicate and terminates if $\theta_{\mathrm{min}} < \frac{\pi}{6}$.

need an explicit apriori lower bound on lfs(x)!

...can it be avoided?

Asserting the topological ball property

Cheng, Dey, Ramos, & Ray,
Sampling & meshing a surf. w/ guarant. topology and geometry.

An alternative requirement is the *topological ball property*:

A point set P on $\mathcal{S}$ has the *topological ball property* if any k-dim face of Vor(P) intersects $\mathcal{S}$ in a closed $(k-1)$-dim ball or is $\emptyset$.

And with this property, the topological space can be triangulated:

Theorem Edelsbrunner & Shah

If $P \subset \mathcal{S}$ satisfies the topological ball property,
and $\mathcal{M}$ is the restricted Delaunay triangulation of P,
then $\mathcal{M}$ is homeomorpic to $\mathcal{S}$.

Insert points on $\mathcal{S} = \{f(x) = 0\}$ where top. ball property fails:

1. For Voronoi edge e, $p = e \cap \mathcal{S}$ must be a single point:

   If e intersects $\mathcal{S}$ more than once, insert farthest intersection.

2. $\mathcal{M}$ is 2-manifold:

   If triangle fan of p contains *multiple cycles* or
   edge $[p, q]$ is not shared by *exactly two triangles*,
   insert farthest intersection of $\mathcal{S}$ with edges of Vor-cell of p/q.

3. For Voronoi face F, $s = \mathcal{S} \cap F$ must be a single segment:

   If s has a closed loop, s is at x tangent to a dir d in F.

   Insert intersection of $s \cap \ell(x, d)$ farthest from x.

   Tests 1 & 2 excludes that s is of multiple segments.

4. For Voronoi volume V, $m = \mathcal{S} \cap V$ must be a single disc:

   If the silhouette $\langle \nabla m, \nabla f(p) \rangle = 0$ is outside V, m is a disc.

   Break silhouette loops: insert points tangent to a $d' \perp \nabla \mathcal{S}(p)$.

   Insert points where silhouette intersects ∂V.

When algorithm terminates:

- P satisifes topological ball property:
  $\mathcal{M}$ *homeomorphic* to $\mathcal{S}$.
  (No guarantee for isotopy.)
- For smooth $\mathcal{S}$:
  New point q inserted is $\|p - q\| \leq 0.06 \operatorname{lfs}(p)$ for $p \in P$.
- Algorithm terminates.
- Not necessary to *know* lfs!
- Needs up to *second order derivatives* of f.

No consideration for triangle quality, but can be extended:

Repeat until stable:

- Run a pass of e.g. Chew's algorithm (may destroy topology).
- Run a pass of steps 1–4 (may destroy triangle shape).

Extension terminates for smooth surfaces.

...what about surfaces with degenerate points?

Sweeping through space

Mourrain & Tecourt,
Isotopic meshing of a real algebraic surface.

The algorithm is based on the concept of the *polar variety*

The *polar variety* C is the set of points satisfying

$$f(x, y, z) = 0, \quad \frac{\partial f}{\partial z}(x, y, z) = 0.$$

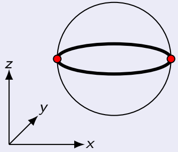

This is the silhouette along the z-axis and is usually a set of curves.

The polar variety is segmented using *slab points*:

Slab points are the x-coordinates where
C is singular or has tangent perpendicular to x-axis.

The polar variety slices the space into a set of vertical slabs, and the cross sections are meshed using a planar algorithm:

- Find critical points X of C:
  - C is tangent to y,
  - C intersects itself, or
  - C has another singularity.

  Between critical points C is x-monotonous.

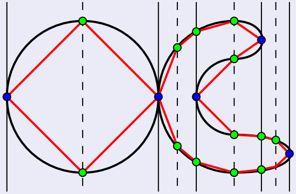

- Insert intermediate x-values into X between critical points.
- Find intersections of C and vertical lines through $x \in X$:
  - Intersections of critical points are multiple roots.
  - Intersections of intermediate points are simple roots.
- Connect intersections, based on multiplicity of roots.

The resulting curve is *isotropic* to C.

To build a mesh $\mathcal{M}$ from a surface $\mathcal{S}$:

1. Find all slab points X and insert intermediate x-coordinate.
2. For each $x_i \in X$,
   intersect $\mathcal{S}$ with the yz-plane through x_i.
   Build cross section using planar meshing.
3. Project cross sections onto xy-plane.
4. Connect critical points of polar variety.
   Regions are stacks of xy-monotone pieces.
5. Triangulate the regions using points from cross sections.
6. Multiplicate and raise the planar triangles to fill the 3D shape.

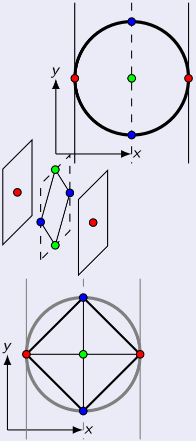

For *distinct slab-points* $\mathcal{M}$ is *ambient isotropic* to $\mathcal{S}$.

vertices is bound by $O(d^7)$ for an algebraic surface of degree d.

Summary

We have looked at some approaches,
each approach has different strength and weaknesses.

Approaches based on subdivision of space:

- Marching cubes approaches
  - produce consistent surface $\mathcal{M}$, but
  - cannot guarantee that $\mathcal{M}$ is of correct topological type.
- Snyder's adaptive algorithm
  - cannot handle singularities,
  - interval arithmetic must converge,
  - requires that generated faces are not tangent to $\mathcal{S}$.
- Pantinga & Vegter's small normal variation approach
  - creates an mesh isotopic to $\mathcal{S}$, but
  - cannot handle singularities,
  - interval arithmetic must converge.

Delaunay-based approaches:

- Chew's farthest point strategy
  - guarantees triangle size and shape, but
  - no guarantee that $\mathcal{M}$ is of correct topological type.
- Boissonnat & Oudots ϵ-sample strategy
  - creates an mesh isotopic to *non-singular* $\mathcal{S}$, but
  - and requires explicit apriori knowledge of lfs.
- Cheng, Dey, Ramos & Ray's topological ball approach
  - creates a mesh *homeomorphic* to *non-singular* $\mathcal{S}$,
  - without explicit knowledge of lfs, but
  - needs second order derivatives of f.

And finally an approach based on sweeping through space:

- Mourrain & Tecourt's space-sweeping approach
  - creates a mesh *isotopic* to $\mathcal{S}$ with degenerate points, but
  - requires distinct slab-points, and
  - no control of triangle size and shape.