




Advanced Graphics





Implicit Surfaces, Voxels, Voronoi Diagrams and Particle Systems

Implicit surfaces

Implicit surface modeling⁽¹⁾ is a way to produce very ‘organic’ or ‘bulbous’ surfaces very quickly without subdivision or NURBS.

Uses of implicit surface modelling:

- Organic forms and nonlinear shapes
- Scientific modeling (electron orbitals, gravity shells in space, some medical imaging)
- Muscles and joints with skin
- Rapid prototyping
- CAD/CAM solid geometry




⁽¹⁾ AKA “metaball modeling”, “force functions”, “blobby modeling”...

How it works

The user controls a set of *control points*, like NURBS; each point in space generates a field of force, which drops off as a function of distance from the point (like gravity weakening with distance.)

This 3D field of forces defines an *implicit surface*: the set of all the points in space where some mathematical function (in this case, the value of the force field) has a particular key value.



Force functions

A few popular force field functions:

- “Blobby Molecules” – Jim Blinn


$$F(r) = a e^{-br^2}$$

- “Metaballs” – Jim Blinn


$$F(r) = \begin{cases} a(1 - 3r^2 / b^2) & 0 \leq r < b/3 \\ (3a/2)(1 - r/b)^2 & b/3 \leq r < b \\ 0 & b \leq r \end{cases}$$

- “Soft Objects” – Wyvill & Wyvill

$$F(r) = a(1 - 4r^6/9b^6 + 17r^4/9b^4 - 22r^2 / 9b^2)$$




Comparison of force functions



Discovering the surface

An *octree* is a recursive subdivision of space which “homes in” on the surface, from larger to finer detail.

- An octree encloses a cubical volume in space. You evaluate the force function $F(v)$ at each vertex v of the cube.
- As the octree subdivides and splits into smaller octrees, only the octrees which contain some of the surface are processed; empty octrees are discarded.



Polygonizing the surface


To display a set of octrees, convert the octrees into polygons.

- If some corners are “hot” (above the force limit) and others are “cold” (below the force limit) then the implicit surface crosses the cube edges in between.
- The set of midpoints of adjacent crossed edges forms one or more rings, which can be triangulated. The normal is known from the hot/cold direction on the edges.

To refine the polygonization, subdivide recursively; discard any child whose vertices are all hot or all cold.

Polygonizing the surface

Recursive subdivision (on a quadtree):




Polygonizing the surface


There are fifteen possible configurations (up to symmetry) of hot/cold vertices in the cube. →

- With rotations, that's 256 cases.

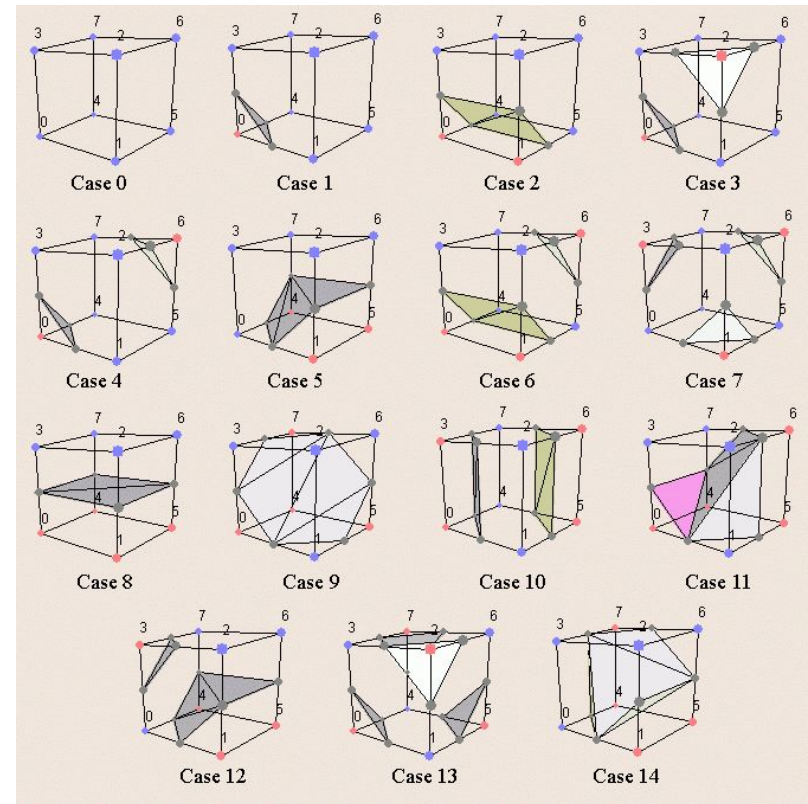
Beware: there are *ambiguous cases* in the polygonization which must be addressed separately. ↓



Break contour



Join contour



Images courtesy of [Diane Lingrand](#)

Polygonizing the surface

One way to overcome the ambiguities that arise from the cube is to decompose the cube into tetrahedra.

- A common decomposition is into five tetrahedra. →
- Caveat: need to flip every other cube. (Why?)
- Can also split into six.

Another way is to do the subdivision itself on tetrahedra—no cubes at all.





Image from the [Open Problem Garden](#)

Smoothing the surface

Improved edge vertices

- The naïve implementation builds polygons whose vertices are the midpoints of the edges which lie between hot and cold vertices.
- The vertices of the implicit surface can be more closely approximated by points linearly interpolated along the edges of the cube by the weights of the relative values of the force function.
 - $t = (0.5 - F(P1)) / (F(P2) - F(P1))$
 - $P = P1 + t (P2 - P1)$

Implicit surfaces -- demo



Marching cubes

An alternative to octrees if you only want to compute the final stage is the *marching cubes* algorithm (Lorensen & Cline, 1985):

- Fire a ray from any point known to be inside the surface.
- Using Newton's method or binary search, find where the ray crosses the surface.
 - Newton: derivative estimated from discrete local sampling
 - There may be many crossings
- Drop a cube around the intersection point: it will have some vertices hot, some cold.
- While there exists a cube which has at least one hot vertex and at least one cold vertex on a side and no neighbor on that side, create a neighboring cube on that side. Repeat.




Marching cubes is common in medical imaging such as MRI scans. It was first demonstrated (and patented!) by researchers at GE in 1984, modeling a human spine.

Voxels and volume rendering

A *voxel* (“volume pixel”) is a cube in space with a given color; like a 3D pixel.

- Voxels are often used for medical imaging, terrain, scanning and model reconstruction, and other very large datasets.
- Voxels usually contain color but could contain other data as well—flow rates (in medical imaging), density functions (analogous to implicit surface modeling), lighting data, surface normals, 3D texture coordinates, etc.
- Often the goal is to render the voxel data directly, not to polygonize it.



Voxels for deformable geometry

Voxels are uniquely well-suited to large-scale, dynamically deformable environments.

Geometry stored in a recursive data structure (“chunks”, arrays of cubes containing arrays of cubes) can be locally edited in real time.




*Fan art from the game Minecraft
(from Deviantart.com, Wallchan.com)*

Volume ray casting

If speed can be sacrificed for accuracy, render voxels with *volume ray casting*:

- Fire a ray through each pixel;
- Sample the voxel data along the ray, computing the weighted average (*trilinear* filter) of the contributions to the ray of each voxel it passes through;
- Compute surface gradient from of each voxel from local sampling; generate surface normals; compute lighting with the standard lighting equation;
- ‘Paint’ the ray from back to front, occluding more distant voxels with nearer voxels; this is the Painter’s Algorithm for hidden-surface removal.




Top: the steps of volume rendering
Bottom: a volume ray-cast skull.
Images from wikipedia.


Sampling in voxel rendering

Why trilinear filtering?

- If we just show the color of the voxel we hit, we'll see the exact edges of every cube.
- Instead, choose the weighted average between adjacent voxels.
 - Trilinear: averaging across X, Y, and Z




Your sample will fall somewhere between eight (in 3d) voxel centers. Weight the color of the sample by the inverse of its distance from the center of each voxel.



Voronoi diagrams

The *Voronoi diagram*⁽²⁾ of a set of points P_i divides space into ‘cells’, where each cell C_i contains the points in space closer to P_i than any other P_j .

The *Delaunay triangulation* is the dual of the Voronoi diagram: a graph in which an edge connects every P_i which share a common edge in the Voronoi diagram.



A Voronoi diagram (dotted lines) and its dual Delaunay triangulation (solid).

⁽²⁾ AKA “Voronoi tessellation”, “Dirichlet domain”, “Thiessen polygons”, “plesiohedra”, “fundamental areas”, “domain of action”...


Voronoi diagrams

Given a set $S = \{p_1, p_2, \dots, p_n\}$, the formal definition of a Voronoi cell $C(S, p_i)$ is

$$C(S, p_i) = \{p \in R^d \mid |p - p_i| < |p - p_j|, i \neq j\}$$

The p_i are called the *generating points* of the diagram.

Where three or more boundary edges meet is a *Voronoi point*. Each Voronoi point is at the center of a circle (or sphere, or hypersphere...) which passes through the associated generating points and which is guaranteed to be empty of all other generating points.



Delaunay triangulations and *equi-angularity*

The *equiangularity* of any triangulation of a set of points S is a sorted list of the angles $(\alpha_1 \dots \alpha_{3t})$ of the triangles.

- A triangulation is said to be *equiangular* if it possesses lexicographically largest equiangularity amongst all possible triangulations of S .
- The Delaunay triangulation is equiangular.




Image from *Handbook of Computational Geometry*
(2000) Jörg-Rüdiger Sack and Jorge Urrutia, p. 227

Delaunay triangulations and *empty circles*

Voronoi triangulations have the *empty circle* property: in any Voronoi triangulation of S , no point of S will lie inside the circle circumscribing any three points sharing a triangle in the Voronoi diagram.




Image from *Handbook of Computational Geometry*
(2000) Jörg-Rüdiger Sack and Jorge Urrutia, p. 227


Delaunay triangulations and convex hulls

The border of the Delaunay triangulation of a set of points is always convex.

- This is true in 2D, 3D, 4D...

The Delaunay triangulation of a set of points in R^n is the planar projection of a convex hull in R^{n+1} .



- Ex: from 2D ($P_i = \{x, y\}_i$), loft the points upwards, onto a parabola in 3D ($P'_i = \{x, y, x^2 + y^2\}_i$). The resulting polyhedral mesh will still be convex in 3D.



Voronoi diagrams and the *medial axis*

The *medial axis* of a surface is the set of all points within the surface equidistant to the two or more nearest points on the surface.

- This can be used to extract a skeleton of the surface, for (for example) path-planning solutions, surface deformation, and animation.




[A Voronoi-Based Hybrid Motion Planner for Rigid Bodies](#)

M Foskey, M Garber, M Lin, DManocha

[Approximating the Medial Axis from the Voronoi Diagram with a Convergence Guarantee](#)

Tamal K. Dey, Wulue Zhao



[Shape Deformation using a Skeleton to Drive Simplex Transformations](#)


IEEE Transaction on Visualization and Computer Graphics, Vol. 14, No. 3, May/June 2008, Page 693-706

Han-Bing Yan, Shi-Min Hu, Ralph R Martin, and Yong-Liang Yang

Finding the Voronoi diagram


There are four general classes of algorithm for computing the Delaunay triangulation:

- Divide-and-conquer
- Sweep plane
 - Ex: Fortune's algorithm →
- Incremental insertion
- “Flipping”: repairing an existing triangulation until it becomes Delaunay




Fortune's Algorithm for the plane-sweep construction of the Voronoi diagram (Steve Fortune, 1986.)

This triangulation fails the circumcircle definition; we flip its inner edge and it becomes Delaunay. *(Image from Wikipedia.)*



Fortune's algorithm

1. The algorithm maintains a sweep line and a “beach line”, a set of parabolas advancing left-to-right from each point. The beach line is the union of these parabolas.
 - a. The intersection of each pair of parabolas is an edge of the voronoi diagram
 - b. All data to the left of the beach line is “known”; nothing to the right can change it
 - c. The beach line is stored in a binary tree
2. Maintain a queue of two classes of event: the addition of, or removal of, a parabola
3. There are $O(n)$ such events, so Fortune's algorithm is $O(n \log n)$






GPU-accelerated Voronoi Diagrams

Brute force:


- For each pixel to be rendered on the GPU, search all points for the nearest point

Elegant:

- Render each point as a discrete cone in isometric projection, let z-buffering sort it out



Voronoi cells in 3D



Silvan Oesterle, Michael Knauss

Particle systems

Particle systems are a monte-carlo style technique which uses thousands (or millions) or tiny graphical artefacts to create large-scale visual effects.

Particle systems are used for hair, fire, smoke, water, spores, clouds, explosions, energy glows, in-game special effects and much more.

The basic idea:
“If lots of little things all do something the same way, our brains will see the thing they do and not the dots doing it.”





Fallout 4
(Bethesda)



Command and Conquer 3
(Electronic Arts)

History of particle systems

- 1962: Ships explode into pixel clouds in “Spacewar!”, the 2nd video game *ever*.
- 1978: Ships explode into broken lines in “Asteroid”.
- 1982: The Genesis Effect in “*Star Trek II: The Wrath of Khan*”.




Fanboy note: You can play the original Spacewar at <http://spacewar.oversigma.com/> -- the actual original game, running in a PDP-1 emulated in a Java applet.

Particle systems

How it works:

- Particles are generated an emitting source
 - Emitter position and orientation are specified discretely;
 - Emitter rate, direction, flow, etc are often specified as a bounded random range
 - This gives a Monte Carlo integration-style effect
- Time ticks; at each tick, particles move.
 - New particles are generated; expired particles are deleted
 - Forces (gravity, wind, etc) accelerate each particle
 - Acceleration changes velocity
 - Velocity changes position
- Particles are textured and rendered.



Transient vs persistent particles emitted to create a 'hair' effect (source: Wikipedia)

Particle systems—implementations

Closed-form function:

- Represent every particle as a parametric equation; store only the initial position p_0 , initial velocity v_0 , then apply fixed acceleration (such as gravity g .)
 - $p(t) = p_0 + v_0 t + \frac{1}{2} g t^2$
- No storage of state \rightarrow small memory footprint
- *Very* limited possibility of interaction
- Best for fire, projectiles, etc—non-responsive particles.

Discrete integration:

- Update every particle separately; this can be expressed as a loop over a list, or as a mutation of a texture (if using a GPU), or as a massive matrix multiplication operation (if using CUDA)



Particle systems—rendering

Can render particles as points, textured polys, or primitive geometry


- Minimize the data sent down the pipe!
- Polygons with alpha-blended images make pretty good fire, smoke, etc

Transitioning one particle type to another creates realistic interactive effects

- Ex: a ‘rain’ particle becomes an emitter for ‘splash’ particles on impact

Particles can be the force sources for a blobby model implicit surface

- Nice for simulating liquids



Hagit Schechter
<http://www.cs.ubc.ca/~hagitsch/Research/>



nvidia

“The Genesis Effect” – William Reeves
Star Trek II: The Wrath of Khan (1982)



References

Implicit modelling:

D. Ricci, *A Constructive Geometry for Computer Graphics*, Computer Journal, May 1973

J Bloomenthal, *Polygonization of Implicit Surfaces*, Computer Aided Geometric Design, Issue 5, 1988

B Wyvill, C McPheeters, G Wyvill, *Soft Objects*, Advanced Computer Graphics (Proc. CG Tokyo 1986)

B Wyvill, C McPheeters, G Wyvill, *Animating Soft Objects*, The Visual Computer, Issue 4 1986

<http://astronomy.swin.edu.au/~pbourke/modelling/implicitsurf/>

<http://www.cs.berkeley.edu/~job/Papers/turk-2002-MIS.pdf>

<http://www.unchainedgeometry.com/jbloom/papers/interactive.pdf>

<http://www-courses.cs.uiuc.edu/~cs319/polygonization.pdf>

Voxels:

J. Wilhelms and A. Van Gelder, *A Coherent Projection Approach for Direct Volume Rendering*, Computer Graphics, 35(4):275-284, July 1991.

Voronoi diagrams

M. de Berg, O. Cheong, M. van Kreveld, M. Overmars, “Computational Geometry: Algorithms and Applications”, Springer-Verlag,

<http://www.cs.uu.nl/geobook/>

<http://www.ics.uci.edu/~eppstein/junkyard/nn.html>

<http://www.iquilezles.org/www/articles/voronoinlines/voronoinlines.htm> // Voronois on GPU

Particle Systems:

William T. Reeves, “*Particle Systems - A Technique for Modeling a Class of Fuzzy Objects*”, Computer Graphics 17:3 pp. 359-376, 1983 (SIGGRAPH 83).

Lutz Latta, *Building a Million Particle System*, <http://www.2ld.de/gdc2004/MegaParticlesPaper.pdf>, 2004

<http://www.darwin3d.com/gamedev/articles/col0798.pdf>

http://mmacklin.com/pbf_sig_preprint.pdf