



The Fundamental Problem




Given: model, material properties, eye/camera, lights

Generate 2-D image


3 Color Sensors (Cones)



How Cones See Spectrum



Graphics Pipeline



- transformation, rasterization
- originally fixed-function VLSI
- pipeline, parallelism
- GPUs evolve -> more powerful, programmable
- vertex shaders, fragment shaders

Shaders

- **Vertex Shaders:** programs that describe the traits (position, colors, and so on) of a vertex. The vertex is a point in 2D/3D space, such as the corner or intersection of a 2D/3D shape.
- **Fragment Shaders:** programs that deal with the per-fragment processing such as lighting. The fragment is a WebGL term that you can think of as a kind of pixel and contains color, depth value, texture coordinates, and more.

Coordinate Systems

- Model: where you define object
- World: place objects, define eye/camera, define light positions, perform lighting operations
- Eye: define view volume, perform lighting operations
- Canonical/Clip View Volume: clip
- Screen: device specific coordinates, most VSA

Want Matrix Representation

why?


$$(TM_3(TM_2(TM_1))) \begin{bmatrix} x \\ y \end{bmatrix} = (TM_3 TM_2 TM_1) \begin{bmatrix} x \\ y \end{bmatrix} = TM_{combined} \begin{bmatrix} x \\ y \end{bmatrix}$$

Homogeneous Coordinates

- want to represent all transformations with a matrix
- $P(x, y) \Leftrightarrow P(w \cdot x, w \cdot y, w)$, $w \neq 0$
- i.e. go up 1 more dimension
- we can always go back by dividing by w
- let's use $w = 1$
- eg. $P(3, 4) \Leftrightarrow P(3, 4, 1)$

Eye Coordinate System

- eye is at origin
- eye is looking along z axis (l.h.s.?)
- x-axis is horizontal
- y-axis is vertical




Viewing Transformation

given: $lookFrom$, $lookAt$, \overrightarrow{lookUp}

$$\hat{e}_z = \frac{lookAt - lookFrom}{\|lookAt - lookFrom\|}$$

$$\hat{e}_x = \frac{\hat{e}_z \times \overrightarrow{lookUp}}{\|\hat{e}_z \times \overrightarrow{lookUp}\|}$$

$$\hat{e}_y = \hat{e}_x \times \hat{e}_z$$



Direct Matrix Creation


$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \begin{bmatrix} ex_x & ex_y & ex_z & d_x \\ ey_x & ey_y & ey_z & d_y \\ ez_x & ez_y & ez_z & d_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$

$$d_x = -\hat{e}_x \cdot lookFrom$$


$$d_y = -\hat{e}_y \cdot lookFrom$$

$$d_z = -\hat{e}_z \cdot lookFrom$$

Projections



Perspective Projection




$$\frac{y_p}{D} = \frac{y_e}{z_e}$$

$$x_p = \frac{D \cdot x_e}{z_e}$$

$$y_p = \frac{D \cdot y_e}{z_e}$$

What if z_e is < 0 or $= 0$?


Clipping



- removing parts of object that are outside field of view
- related terms:
 - culling: quickly determining in/out
 - bounding box: axis-aligned box containing object
 - scissoring: combining clipping with scan conversion

Clipping Lines: Cohen-Sutherland

- compute labels for $p1$ & $p2$
- determine if total visible or trivial reject
- if $p1$ not outside, swap $p1$ & $p2$
- find edge $p1$ is out
- replace $p1$ with intersection of $p1$ - $p2$ and edge
- compute new label for $p1$



1001	1000	1100
0001	0000	0100
0011	0010	0110

if both labels 0
→ trivial accept


if $\text{label}(p1) \cap \text{label}(p2) \neq 0$
→ trivial reject

Hardware acceleration

Polygon Clipping: Sutherland-Hodgman

- convex clipping region
- clip against one edge at a time

P	C	Output
inside	inside	c
inside	outside	intersection point
outside	outside	-
outside	inside	intersection point; c




Bresenham's Alg

```

DrawLine(int x1,int y1,int x2,int y2){
    int x, y, dx, dy, error
    dx = x2-x1
    dy = y2-y1
    error = 2*dy-dx
    y = y1;
    for (x = x1; i <= x2; x++) {
        SetPixel(x, y)
        if (error > 0) {
            y++
            error = error - 2*dx
        } else
            error = error + 2*dy
    }
}

```




- multiply error by 2*dx (only care about sign)
- developed in 1960s
- pen plotters

Area Averaging


- contribution proportional to area within pixel square

$$I'(x_0, y_0) = \int_{x_0-0.5}^{x_0+0.5} \int_{y_0-0.5}^{y_0+0.5} I(x, y) dx dy$$



Convex Polygons

- find top vertex
- go down left and right sides
- compute intersections with scanline
- draw horizontal runs on each scan line



- *incremental*
- *amortize edge computations*

Newell & Sequin

- VLSI
- compute "winding number" as you go along scanline
- edge going up: +1
- edge going down: -1
- draw non-zero spans






Image Space

- for each pixel in image:
 - determine object closest to viewer
 - draw pixel appropriate color



Object Space

- for each object in scene:
 - determine parts of object that are unobstructed
 - draw those parts appropriate color




Techniques for Efficient Visible Surface Algorithms

- **coherence:** degree to which parts of environment or its projection exhibit local similarities
- examples of types of coherence: object, face, edge, scan-line, area, depth, frame



VSA: Z-Buffer (depth buffer)

- 2 buffers: color buffer, z-buffer
- compare during scan conversion:
 - if depth of new fragment is closer
 - update color buffer
 - update z-buffer



VSA: Painters Alg



- sort polygons back to front
- resolve any ambiguities (use extents, clip if necessary)
- scan-convert polygons back to front



BSP-Trees


Draw in priority order (back to front)

- create tree of subspaces (nodes store polygon, separating plane)
- recursively traverse tree using lookFrom point
- visit order:
 - far
 - plane
 - near




VSA: Scanline Alg

- scanline at a time
- reduce dimension: 3D->2D




VSA: Warnock's Alg

- recursively subdivide screen
- stop when "simple" or at pixel
- at pixel draw closest object




VSA: Weiler-Atherton

- fast sort of polygons by z
- select "closest" polygon
- use it to clip the rest
- if any poly inside clipping poly closer -> initial sort wrong
 - use it as clipping poly first
- otherwise discard those inside
- draw clipping poly



VSA: Ray Casting (Ray Tracing)



- shoot ray from eye through screen into world
- intersect objects with ray
- find closest intersection
- do shading/lighting calculation
- very floating-point intensive



Diffuse Reflection


(Lambertian Reflection)

- dull, matte surfaces
- reflect light equally in all directions
- light enters object, scatters internally
- eg: plastic, paint, paper, vegetation, snow, etc



Specular Reflection


- shiny surfaces, highlights



Diffuse + Specular

$$I = I_{light} \cdot (K_d \cdot \cos(\theta) + K_s \cdot \cos^n(\alpha))$$


[0...1]



$$\cos(\alpha) = \text{dot}(E, R)$$

Ambient Reflection


- modeling inter-reflection is hard
- parts in shadow are black (looks bad)
- approximate indirect lighting
- simplification:
use constant K_a



↖
[0...1]


Color

- K_a , K_d , K_s depend on λ , I_{light} depends on λ
- highlights: white for many objects (actually, color of light)




Gouraud Shading

- compute vertex normals
 - average of polygons around vertex
 - directly from model during tessellation
- perform lighting operation at vertex
- linearly interpolate resulting vertex color
(linear interpolation not correct)




Phong Shading

- compute vertex normals
- linearly interpolate vertex normals
(linear interpolation not correct)
- perform lighting operation per pixel



Shading Models for Polygons



Flat

Gouraud

Phong

Cook-Torrance

reflectance (Fresnel) is also a function of wavelength λ




Figure 5a. Reflectance (ρ) of a copper mirror as a function of wavelength (λ) and incidence angle (θ).



Red
Rubber

Lunar
Dust


Olive
Drab

Bronze

Tungsten

Copper

Texture Mapping



Bump Mapping

- texture maps can perturb surface normal
- illusion of bumps




Blinn, 1978



Problems with Texture

- small objects -> integrate over large number of texels





Solid Textures

- with complicated, curved 3D objects want a 3D texture to make sure textures match
- it needs lots of memory
- solution: procedural texture
- $f(x, y, z)$ evaluated by fragment shader
- as if you “carve” object out of solid texture

Perlin, 1985





Environment Maps



Shadows


- it's really a visible surface problem from the point of view of the light source
- one strategy: two-pass algorithms

hard part is to correctly share info from one pass to another



Shadow Map 2-Pass Z Buffer


- first pass: VSA from point of view of light src
- save z-buffer as shadow map (distance from light to object)



Williams, 1978

2-Pass Weiler Atherton



- Object-space
- first pass from light source
- result is lit polygons (not in shadow)
- use these polygons as surface detail polygons
- second pass from camera



Atherton, Weiler, Greenberg, 1981

Ray Casting

- during lighting calculation shoot ray towards light src
- if shadow ray hits object then in shadow




Ray Tracing

- a generalization of ray casting
- why?
 - visible surface
 - shadows
 - reflection
 - refraction




Whitted, 1980

Ray Tracing How





Ray Tracing How

$$I = \text{ambient} + \sum_{\text{lights}} (\text{diffuse} + \text{specular}) + K_r \cdot I_R + K_t \cdot I_T$$






Supersampling

- what can be done?
 - increase sample rate
 - increase sample rate and average over several samples (supersample, oversample) to get pixel
- expensive







Adaptive Supersampling

- heuristic: adaptive supersampling
- increase sample rate only in "troublesome" regions
- if difference in neighbours $>$ threshold
 - increase sample rate in neighbourhood



Non-uniform Sampling

- regular sampling pattern results in regular aliasing pattern
- non-uniform sampling results in noisy image
- noise less objectionable than regular aliasing pattern




Advanced Optical Effects

- better camera models
- dull reflection
- frosted glass
- area light sources
- motion blur
- forward ray tracing


Distributed Ray Tracing

- all advanced optical effects require multiple rays
- multiple rays per pixel: 8 ... 64
- each ray can sample all effects independently



Radiosity

- ambient lighting is a hack
- want to model indirect lighting
- hard
- assume diffuse surfaces




Diffuse Environments

$$A_i F_{i \rightarrow j} = A_j F_{j \rightarrow i}$$

$$B_i = E_i + \rho_i \sum_j B_j F_{i \rightarrow j}$$

$$B_i - \rho_i \sum_j B_j F_{i \rightarrow j} = E_i$$




n simultaneous equations in n unknowns

Form Factor

- $F_{i \rightarrow j}$ = fraction of energy leaving patch i and arriving at patch j
- takes into account visibility



$$F_{i \rightarrow j} = \frac{1}{A_i} \iint_{A_i, A_j} \frac{\cos \theta_i \cos \theta_j}{\pi r^2} H_{ij} dA_j dA_i$$




- H_{ij} is either 0 or 1, depending if dA_i sees dA_j

Speeding Up Radiosity

- Gauss-Sidel: gathering
- progressive refinement: shooting
- adaptive subdivision of patches (hierarchical radiosity)




Bezier Curves






Splines

- piecewise polynomials that satisfy continuity conditions between the pieces
- 2 types:
 - interpolating (go through control points)
 - approximating (go near control points)
- work with 4 control points (**knots**) at a time




Bezier Patches



Subdivision

- smooth surface as the limit of a sequence of refinements of rougher mesh




Corner Cutting

- place 2 vertices, $1/4$ and $3/4$ between original vertices
- remove original vertices
- repeat




CSG

- Boolean set operations on simpler solids
- Union, Intersection, Difference, Negation






BSP Solids

- leaf nodes can store in/out




Voxels


- dice space into voxels
- voxel is either in or out
- uniform vs octree



SIMD

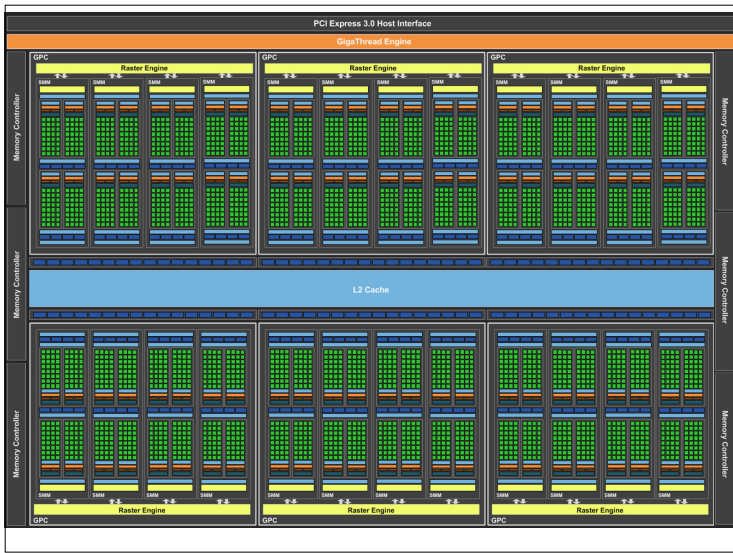


CPU




GPU

rip out parts that make CPU fast 128 ALUs in same space



Multi-threading hides latency

- work on many groups of vertices/fragments at same time
- when one group stalls, work on other group
- modern GPUs can have 128 threads!




Three key ideas

- use many “slimmed down cores” in parallel
- pack cores full of ALUs and share instruction streams
- avoid latency stalls by interleaving many groups of threads


Tile-based Rendering

- simpler primitives: points, lines, triangles
- frame buffer partitioned into tiles (eg 64x64) with just enough GPU rasterization hardware for 1 tile
- transform primitives and store them, noting which tiles they overlap (retained mode)
- work on one tile at a time




Antialiasing in CG

- assumption: intensity within single pixel is constant
- typical filter is the box filter (area average)
- better filters: triangle, gaussian, quadratic
- tradeoffs: computation, sharpness, aliasing, ringing



$$I'(x_0, y_0) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} I(x, y) F(x_0 - x, y_0 - y) dx dy$$



Antialiasing Approaches


(when you're forced to sample)

- supersampling
- adaptive supersampling
- non-uniform sampling

$$I'(x_0, y_0) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} I(x, y) F(x_0 - x, y_0 - y) dx dy$$


Adaptive Supersampling

- heuristic: adaptive supersampling
- increase sample rate only in “troublesome” regions
- if difference in neighbours > threshold
 - increase sample rate in neighbourhood



Non-uniform Sampling

- regular sampling pattern results in regular aliasing pattern
- non-uniform sampling results in noisy image
- noise less objectionable than regular aliasing pattern



Non-uniform Sampling Patterns

- Poisson
- Jitter
- Poisson Disk

