GRAPHICS & VISION SUMMER SCHOOL

Computer Graphics

Jenser **Henrik Wa**

Rasterization vs. Ray tracing

for each shape:
for each sample:
 get point where shape covers sample
 if point is closest point seen by sample:
 sample.colour = shade(point)

for each sample: for each shape: get point where shape covers sample if point is closest point seen by sample: sample.colour = shade(point)









Ray tracing

For each sample (x, y): Generate eye ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ Find the closest intersection Get shaded colour at intersection point Set sample colour to it



Ray-surface intersection

Given a ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$, find closest intersection i.e. minimum t

Return info needed for shading:

- Position **p**
- Normal **n**
- Object ID / material properties

(Roughly the same data you would need in a fragment shader)





 $\mathbf{t} = (\mathbf{n} \cdot (\mathbf{p}_0 - \mathbf{o}))/(\mathbf{n} \cdot \mathbf{d})$

Ray-triangle intersection

Intersect ray with plane, then check if it is inside triangle?













Ray-mesh intersection

Naïve approach: Test ray with all triangles, return the earliest hit.

Cost = O(#triangles)! Can we speed it up?

Construct a conservative **bounding volume**: all mesh triangles lie inside it

Super easy to reject rays that don't come close to intersecting the mesh.

In practice, we use bounding volume hierarchies to speed things up further.







Object-oriented raytracer design

We can ray trace any shape as long as it provides the following methods:

- bool hit(Ray o + td, real t_{min}, real t_{max}, HitRecord &rec) • Only consider intersections in the range $t_{\min} \leq t \leq t_{\max}$. Usually $[0, \infty]$ for eye rays • If hit, write the position, normal, material, etc. into the HitRecord
- Box bounding_box()
 - For early exit



Ray tracing

For each sample (x, y):

Generate eye ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$

Find the closest intersection (**p**, **n**, ...)

Get shaded colour at intersection point

Set sample colour to it



Shading the intersection point

Now we have position **p**, normal **n**, etc. We can compute other vectors v, l, h, etc. as usual

Apply your favourite reflection model, e.g. Blinn-Phong:

$$L = k_a I_a + \sum I_i (k_d \max(0, \mathbf{n} \cdot \boldsymbol{\ell}_i) + k_s \max(0, \mathbf{n} \cdot \mathbf{h}_i)^p)$$
$$f(\boldsymbol{\ell}_i, \mathbf{v})$$

 $f(\mathbf{l}, \mathbf{v}) = what fraction of incident light from direction <math>\mathbf{l}$ is reflected to direction \mathbf{v}





So far, what we have done is ray casting: shoot just one ray from the eye. This should give us exactly the same image as we'd get from rasterization!



Now we will also trace secondary rays to get additional effects.

Light should only be included if it is visible from **p**

- Shoot a "shadow ray" p + tl towards light source
- If no intersection within interval [0, t_{light}], add its contribution to L

Why [0, t_{light}] instead of [0, ∞]?

$L = k_a I_a + \sum_{i} I_i f(l_i, v)$







"Shadow acne"

- Why does this happen?
- **p** lies on surface
- $\mathbf{p} + t\mathbf{l}$ intersects surface at t = 0
- Floating-point arithmetic may give e.g. t =0.000001
- **p** thinks it's shadowed... by itself!

Solution: Pick a small positive ε ("bias") and only look for intersections in [ε , t_{light}].



Without bias



With bias

Ray tracing

For each sample (x, y): ray = makeRay(camera, x, y) hit = castRay(ray, scene) color = shade(hit, scene) image[x, y] = color



color = traceRay(ray, scene)





Reflection $\mathbf{r} = \mathbf{d} - 2\mathbf{d}_n$ $= \mathbf{d} - 2(\mathbf{n} \cdot \mathbf{d})\mathbf{n}$ Perfect mirror: L = traceRay(**p**, **r**, scene) Actually, k_r depends on $n \cdot d$ for many materials... Reflective surface: $L = [...Blinn-Phong...] + k_r' traceRay($ **p**,**r**, scene)

Again, don't forget ray bias!







 k_a , k_d coloured k_s , k_r colourless



Plastic

k_a, k_d zero k_s, k_r coloured

Metallic

Transmission

Suppose ray passes from one material to another with different indices of refraction.

Snell's law: $\eta_i \sin \theta_i = \eta_t \sin \theta_t$

$$\mathbf{t} = \left(\eta_r(\mathbf{n} \cdot \mathbf{i}) - \sqrt{1 - \eta_r^2 \left(1 - (\mathbf{n} \cdot \mathbf{i})^2\right)}\right) \mathbf{n} - \eta_r \mathbf{i}$$

where $\eta_r = \eta_i / \eta_t$.



Also depends on $n \cdot d$



Fresnel effect





image plane











image plane





Kanazawa and Ng



image plane





Kanazawa and Ng









Kanazawa and Ng







Recursion depth: 2





Recursion depth: 4

When to terminate the recursion?

- Fixed: stop if recursion depth > max
- Adaptive: stop if contribution of ray to final pixel colour < threshold
- ... or whichever comes first

On termination, return... diffuse colour? background colour? black?



What ray tracing can and can't do

Refraction

Reflection

Shadows





So far, we have learned how to make crude pictures of polygonal shapes.



How would we make photorealistic movies of complicated shapes? **RENDERING ANIMATION MODELING**



Tour of deeper aspects of computer graphics



Modelling





Rendering

Animation







How to define a unit circle in 2D?

Explicit:

{(cos θ , sin θ): $0 \le \theta < 2\pi$ }



Implicit:

{(x, y): $x^2 + y^2 - 1 = 0$ }



Explicit:

$\{(x(t), y(t)): t \in [a, b]\}$



When is it easy to **generate** an arbitrary point on the curve? When is it easy to **test** if a given point lies on the curve?

Implicit:

 $\{(x, y): f(x, y) = 0\}$



How to draw a curve given in one of these forms?

$\{(x(t), y(t)): t \in [a, b]\}$



Sample points at various values of t

Connect by polyline



Sample f at various points (x, y)

Draw boundary between + and – points


Representing geometry in 3D

Explicit:

- Polygon meshes
- Parametric curves and surfaces
- Subdivision surfaces
- Point clouds



Implicit:

- Algebraic surfaces, distance fields
- Constructive solid geometry
- "Blobby" surfaces
- Level sets













We will retain the same "interface" as polylines: user specifies a sequence of points. Now we want to define a smooth curve based on them.



Usually define parametrically: *x*(*t*), *y*(*t*) where *x*, *y* are piecewise polynomial functions a.k.a. **splines**





You all probably already know one way to fit a smooth function through multiple points: **polynomial interpolation**.



Hard to control: curve goes beyond the range of the control points

Very unstable for higher degrees!



Bézier curves

How can we guarantee the curve stays within the range of the control points? Construct the curve by recursive interpolation: **de Casteljau's algorithm** a.k.a.

Construct the curve by recursive inter "corner cutting"



$$\mathbf{b}_0^1 = \operatorname{lerp}(t, \mathbf{b}_0, \mathbf{b}_1)$$
$$\mathbf{b}_1^1 = \operatorname{lerp}(t, \mathbf{b}_1, \mathbf{b}_2)$$
$$\mathbf{b}_0^2 = \operatorname{lerp}(t, \mathbf{b}_0^1, \mathbf{b}_1^1)$$

$$b_0^1 = \text{lerp}(t, b_0, b_1)$$

$$b_1^1 = \text{lerp}(t, b_1, b_2)$$

$$b_2^1 = \text{lerp}(t, b_2, b_3)$$

$$\mathbf{b}_0^2 = \text{lerp}(t, \mathbf{b}_0^1, \mathbf{b}_1^1)$$

 $\mathbf{b}_1^2 = \text{lerp}(t, \mathbf{b}_1^1, \mathbf{b}_2^1)$

 $\mathbf{b}_0^3 = \text{lerp}(t, \mathbf{b}_0^2, \mathbf{b}_1^2)$











Piecewise Bézier curves (Bézier splines)

Chain together multiple Bézier curves of low degree (usually cubic)



Now we have local control: each control point only affects one or two segments Used basically everywhere (fonts, paths, Illustrator, PowerPoint, ...)

Another strategy to create smooth shapes from a coarse mesh of control points: subdivision

- Split each element by inserting new vertices
- Update positions of all vertices by local averaging
- Repeat...

The desired shape is what we converge to in the limit.



Subdivision curves







Subdivision surfaces

Connectivity of surfaces is more complicated. Many different subdivision schemes are possible:

- General polygon meshes: Catmull-Clark, Doo-Sabin, mid-edge [Peters] & Reif], ...
- Triangle meshes: Loop, modified butterfly [Zorin et al.], Sqrt(3) [Kobbelt], ...

Sqrt(3)











Catmull-Clark subdivision

- Split each *n*-sided face into *n* quads
- Update vertex positions by averaging:
- New face point = average of old face vertices
- New edge point = average of 2 old vertices and 2 new face points
- Updated vertex = $\frac{1}{n}(Q + 2R + (n-3)S)$ where Q = average of n new face points, R = average of n new edge points, S = old vertex







Examples













Rendering





Our goal: given light sources and scene geometry, find amount of light (i.e. radiance) incident on camera.

To do this, we need to know how surfaces transform **incident** radiance into **exitant** radiance





Real-Time Rendering

The BRDF

Light



`ωi



$f_r(\boldsymbol{\omega}_i \rightarrow \boldsymbol{\omega}_o) = L_o(\mathbf{x}, \boldsymbol{\omega}_o)/E(\mathbf{x})$

Computer Graphics Fundamentals of



Lambertian (diffuse) material

Simplest possible model: BRDF is a constant!

$$L_{o}(\boldsymbol{\omega}_{o}) = \int_{H^{2}} f_{r} L_{i}(\boldsymbol{\omega}_{i}) \cos(\boldsymbol{\theta}_{i}) d\boldsymbol{\omega}_{i}$$
$$= f_{r} E_{i}$$

To conserve energy, $f_r = \rho/\pi$ where albedo ρ is ≤ 1

Why? For constant radiance L, total flux density = $L \pi$





BRDF acquisition: Gonioreflectometer



actice

MERL BRDF database











Ray tracing





HENRIK WANN JENSEN 1999

Global illumination



Henrik Wann Jensen

Ray tracing revisited

For each sample:

Cast a ray into the scene

Find the closest intersection

Get exitant radiance at intersection point

Set sample colour to it

$$L_o(\mathbf{p}, \boldsymbol{\omega}_o) = L_e(\mathbf{p}, \boldsymbol{\omega}_o) + \int_{H^2}$$



 $f_r(\mathbf{p}, \boldsymbol{\omega}_i \rightarrow \boldsymbol{\omega}_o) L_i(\mathbf{p}, \boldsymbol{\omega}_i) \cos(\theta_i) d\boldsymbol{\omega}_i$

$$L_o(\mathbf{p}, \boldsymbol{\omega}_o) = L_e(\mathbf{p}, \boldsymbol{\omega}_o) + \int_{H^2}$$

- How to evaluate incident radiance from any direction (not just light sources)?
- How to compute the integral over a hemisphere?

$f_r(\mathbf{p}, \boldsymbol{\omega}_i \rightarrow \boldsymbol{\omega}_o) L_i(\mathbf{p}, \boldsymbol{\omega}_i) \cos(\theta_i) d\boldsymbol{\omega}_i$



What is $L_i(\mathbf{p}, \boldsymbol{\omega}_i)$? Simply exitant radiance from somewhere else!



Define tr(\mathbf{p} , $\boldsymbol{\omega}$) as the first surface point hit by the ray \mathbf{p} + $t\boldsymbol{\omega}$.

 $L_i(\mathbf{p}, \boldsymbol{\omega}_i) = L_o(tr(\mathbf{p}, \boldsymbol{\omega}_i), -\boldsymbol{\omega}_i)$

tr(**p**, **ω**)

$$L_o(\mathbf{p}, \boldsymbol{\omega}_o) = L_e(\mathbf{p}, \boldsymbol{\omega}_o) + \int_{H^2} f_r(\mathbf{p}, \boldsymbol{\omega}_o) d\boldsymbol{\omega}_o)$$

This is an **integral equation**! Unknown quantity L_o on both sides

Like ray tracing, we'll evaluate it recursively

$\boldsymbol{\omega}_i \rightarrow \boldsymbol{\omega}_o$) $L_o(tr(\mathbf{p}, \boldsymbol{\omega}_i), -\boldsymbol{\omega}_i) \cos(\theta_i) d\boldsymbol{\omega}_i$





Quick probability recap

If X is a random variable with probability distribution p(x), its expected value or expectation is

Expectation is **linear**:

- $E[X_1 + X_2] = E[X_1] + E[X_2]$
- E[aX] = a E[X]

 $E[X] = \sum x_i p_i$ $E[X] = \int x p(x) dx$

(discrete)

(continuous)



The basic Monte Carlo method

If X is uniformly distributed in [a, b], then

E[f(X)] =

So, if I take N independent samples of X,

$$\frac{1}{N} \sum_{i=1}^{N} f(x_i) \approx E[f(X)] = \frac{1}{b-a} \int_{a}^{b} f(x) \, dx$$
$$\int_{a}^{b} f(x) \, dx \approx \frac{b-a}{N} \sum_{i=1}^{N} f(x_i)$$
Interpretation:
Integral = average value × domain size

$$\frac{1}{b-a} \int_{a}^{b} f(x) \, \mathrm{d}x$$



Monte Carlo rendering

With Monte Carlo, it's easy:

- Uniformly sample hemisphere of incident directions: $\mathbf{X}_i \sim U(H^2)$, probability density $p(\boldsymbol{\omega}) = 1/(2\pi)$
- Evaluate integrand $Y_i = f_r(\mathbf{p}, \mathbf{X}_i \rightarrow \boldsymbol{\omega}_o) L_i(\mathbf{p}, \mathbf{X}_i) \cos(\theta_i)$
- MC estimator is simply $F_N = 2\pi/N \sum_i Y_i$

We need to estimate the reflectance integral $\int f_r(\mathbf{p}, \boldsymbol{\omega}_i \rightarrow \boldsymbol{\omega}_o) L_i(\mathbf{p}, \boldsymbol{\omega}_i) \cos(\theta_i) d\boldsymbol{\omega}_i$





Blocker

Light

Incident lighting estimator uses different random directions in each pixel. Some of those directions point towards the light, others do not.

(Estimator is a random variable)

Keenan Crane

$$L_{o}(\mathbf{p}, \boldsymbol{\omega}_{o}) = L_{e}(\mathbf{p}, \boldsymbol{\omega}_{o}) + \int_{H^{2}}$$

incidentRadiance(**x**, **ω**):

 $f_r(\mathbf{p}, \boldsymbol{\omega}_i \rightarrow \boldsymbol{\omega}_o) L_i(\mathbf{p}, \boldsymbol{\omega}_i) \cos(\theta_i) d\boldsymbol{\omega}_i$

 $(\omega i, -\omega) * \cos_{\theta i} * 2\pi / N$
Problem: Exponential increase in number of samples per bounce



Solution: Just take one recursive sample per bounce!



But take many samples per pixel, and average them.

incidentRadiance($\mathbf{x}, \boldsymbol{\omega}$):

- $\mathbf{p} = intersectScene(\mathbf{x}, \boldsymbol{\omega})$
- $L = \mathbf{p}.emittedLight(-\omega)$
- $\omega i = \text{sampleDirection}(\mathbf{p}.\text{normal})$
- $L = incidentRadiance(\mathbf{p}, \boldsymbol{\omega}i) * \mathbf{p}.BRDF(\boldsymbol{\omega}i, -\boldsymbol{\omega}) * cos_{\theta}i * 2\pi$ return L
- This is called **path tracing**.

Each sample is tracing one possible path between the eye and a light source







Animation





Modeling



Character animation

https://www.youtube.com/watch?v=KDvfFzFIruQ



Physics-based animation



https://www.youtube.com/watch?v=chnS24QfgNY

nci et al. 2012

Animation is defined through a set of animation controls (degrees of freedom) whose values vary with time

For example:

- Character: joint angles, etc.
- Rigid body: translation and rotation
- Liquid: position/velocity of all particles(!)







Types of animation techniques



- Artist-specified (e.g. keyframing)
- Data-driven
 (e.g. motion capture)
- Procedural (e.g. simulation)

Less manual effort







Keyframe animation

In traditional (hand-drawn animation:

- Lead animator creates keyframes
- Assistant creates in-between frames ("tweening")







Thomas & Johnston, The Illusion of Life



In computer animation, keyframes = control points, tweening = splines!



Autodesk Maya's Graph Editor



Motion capture











<u>https://www.youtube.com/watch?v=4NU9ikjqjC0</u>

CAESAR

Physics-based animation (a.k.a. simulation)

- Solve the equations of motion to automatically get physically realistic motion.
- e.g. Rigid bodies
- Degrees of freedom: position, rotation

$$\frac{\mathrm{d}^2 \mathbf{x}}{\mathrm{d}t^2} = \mathbf{f}_{\mathrm{ext}}/m$$
$$\frac{\mathrm{d}^2 \mathbf{R}}{\mathrm{d}t^2} = \cdots$$

• Challenges: collisions, frictional contact, stacking





Deformable bodies, cloth, etc.

Every vertex can move independently! But deformation causes internal elastic forces

- Physically accurate: finite element method
- Simplified approximation: mass-spring systems (just a bunch of particles connected by springs)





Fluids (smoke, water, fire, etc.)

Described by systems of partial differential equations

Velocity field v(x, t): every point has its own velocity!

$$\frac{\partial \mathbf{v}}{\partial t} = \left[\dots \text{ somethin} \right]$$





- ng involving $\mathbf{v}, \frac{\partial \mathbf{v}}{\partial \mathbf{x}}, \text{ etc. } \dots$





Physics in character animation

- Flesh
- Hair

• • •

• Clothing





https://vimeo.com/245424174



Mass-spring systems









Selle et al. 2008



Recall springs in 1 dimension from physics classes. Hooke's law: force is proportional to displacement $F = -k x = -k (\ell - \ell_0)$

Potential energy:

$$U = \frac{1}{2} k (\ell - \ell_0)^2$$

In fact $F = -dU/d\ell$





Let's first define the potential:

 $U = \frac{1}{2} k (||\mathbf{x}_i - \mathbf{x}_i|| - \ell_0)^2$

Then $\mathbf{f}_{ij} = -\partial U / \partial \mathbf{x}_i \Rightarrow$

 $f_{ii} = -k (||x_i - x_i|)$

Similarly $\mathbf{f}_{ii} = -\partial U / \partial \mathbf{x}_i$ (but it's also just $-\mathbf{f}_{ii}$)

Exercise: Derive this expression from $-\partial U/\partial \mathbf{x}_i$. Optional: Look up multivariable calculus identities, chain rule, etc. so you don't have to differentiate componentwise.

In 3D, suppose a spring connects particles *i* and *j*. What should be the force f_{ij} on *i* due to *j*?

$$k \left(\|\mathbf{x}_{i} - \mathbf{x}_{j}\| - \ell_{0} \right) \frac{\mathbf{x}_{i} - \mathbf{x}_{j}}{\|\mathbf{x}_{i} - \mathbf{x}_{j}\|}$$
$$= -k \left(\|\mathbf{x}_{ij}\| - \ell_{0} \right) \mathbf{\hat{x}}_{ij}$$



Time stepping

Equations of motion:

 $\frac{\mathrm{d}\mathbf{x}_i}{\mathrm{d}t} = \mathbf{v}_i$

Take small time steps from time t_0 to t_1 , then t_1 to t_2 , and so on... $\mathbf{v}_i(t_{n+1}) \approx \mathbf{v}_i(t_n)$

 $\mathbf{x}_{i}(t_{n+1}) \approx \mathbf{x}_{i}(t_{n}) + \Delta t \mathbf{v}_{i}(t_{n+1})$

Tradeoff: Smaller $\Delta t \rightarrow$ more accurate, but more computation to reach any desired t.

$$\frac{\mathrm{d}\mathbf{v}_i}{\mathrm{d}t} = m_i^{-1} \sum_{ij} \mathbf{f}_{ij}$$

$$h_{n}$$
) + $\Delta t m_{i}^{-1} \sum \mathbf{f}_{ij}(t_{n})$



Witkin & Bara ff 2001

• Structural springs



- Structural springs
- Shear springs



- Structural springs
- Shear springs
- Bending springs



- Structural springs
- Shear springs
- Bending springs







Skeel Lee





Skeel Lee



Skeel Lee



Jensei Henrik Wann