

Ray Tracing

Brian Curless
CSE 557
Autumn 2017

Reading

Required:

- ♦ Marschner and Shirley, Ch. 4, Section 13.1-13.2 (online handout)
- ♦ Triangle intersection (online handout)

Further reading:

- ♦ Shirley errata on syllabus page, needed if you work from his book instead of the handout, which has already been corrected.
- ♦ T. Whitted. An improved illumination model for shaded display. Communications of the ACM 23(6), 343-349, 1980.
- ♦ A. Glassner. An Introduction to Ray Tracing. Academic Press, 1989.
- ♦ K. Turkowski, "Properties of Surface Normal Transformations," Graphics Gems, 1990, pp. 539-547.

Geometric optics

Modern theories of light treat it as both a wave and a particle.

We will take a combined and somewhat simpler view of light – the view of **geometric optics**.


Here are the rules of geometric optics:

- ♦ Light is a flow of photons with wavelengths. We'll call these flows "light rays."
- ♦ Light rays travel in straight lines in free space.
- ♦ Light rays do not interfere with each other as they cross.
- ♦ Light rays obey the laws of reflection and refraction.
- ♦ Light rays travel from the light sources to the eye, but the physics is invariant under path reversal (reciprocity).


Eye vs. light ray tracing

Where does light begin?

At the light: light ray tracing (a.k.a., forward ray tracing or photon tracing)



At the eye: eye ray tracing (a.k.a., backward ray tracing)




We will generally follow rays from the eye into the scene.

Precursors to ray tracing


Local illumination

- ◆ Cast one eye ray, then shade according to light



Appel (1968)

- ◆ Cast one eye ray + one ray to light




5

Whitted ray-tracing algorithm

In 1980, Turner Whitted introduced ray tracing to the graphics community.

- ◆ Combines eye ray tracing + rays to light
- ◆ Recursively traces rays




Algorithm:

1. For each pixel, trace a **primary ray** in direction V to the first visible surface.
2. For each intersection, trace **secondary rays**:
 - ◆ **Shadow rays** in directions L_i to light sources
 - ◆ **Reflected ray** in direction R .
 - ◆ **Refracted ray or transmitted ray** in direction T .

6

Whitted algorithm (cont'd)


Let's look at this in stages:



7


Ray casting and local illumination

Now let's actually build the ray tracer in stages. We'll start with ray casting and local illumination:



8

Direct illumination



A ray is defined by an origin P and a unit direction \mathbf{d} and is parameterized by $t > 0$:

$$\mathbf{r}(t) = P + t \mathbf{d}$$

Let $I(P, \mathbf{d})$ be the intensity seen along a ray. Then:

$$I(P, \mathbf{d}) = I_{\text{direct}}$$

where

- I_{direct} is computed from the Blinn-Phong model

9

Shading in "Trace"

The Trace project uses a version of the Blinn-Phong shading equation we derived in class, with three modifications:

- Distance attenuation is clamped to be at most 1:

$$A_j^{\text{dist}} = \min \left\{ 1, \frac{1}{a_j + b_j r_j + c_j r_j^2} \right\}$$

- Shadow attenuation A^{shadow} is included.
- We will use k_d for the ambient reflection coefficient.

Here's what the shading equation should look like:

$$I = k_e + \sum_j k_d I_{L,j} + A_j^{\text{shadow}} A_j^{\text{dist}} I_{L,j} B_j \left[k_d (\mathbf{N} \cdot \mathbf{L}_j) + k_s (\mathbf{N} \cdot \mathbf{H}_j)^n \right]$$

This is the shading equation to use in the Trace project!

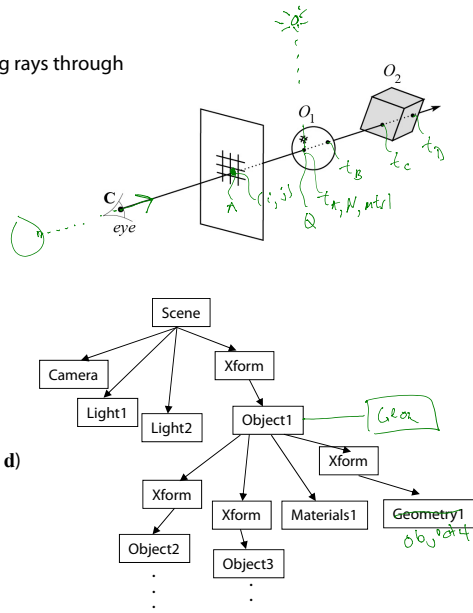
10

Ray-tracing pseudocode

We build a ray traced image by casting rays through each of the pixels.

```
function traceImage(scene):
  for each pixel (i, j) in image
    A = pixelToWorld(i, j)
    P = C
    d = (A - P) / ||A - P||
    I(i, j) = traceRay(scene, P, d)
  end for
end function
```

```
function traceRay(scene, P, d):
  (t, N, mtrl) ← scene.intersect(P, d)
  Q ← ray(P, d) evaluated at t
  I = shade(mtrl, N, scene, d, Q)
  return I
end function
```




11

Shading pseudocode

Next, we need to calculate the color returned by the *shade* function.


```
function shade(mtrl, scene, Q, N, d):
  I ← mtrl.k_e + mtrl.k_d * I_amb
  for each light source Light in scene do:
    atten = Light -> distanceAttenuation(Q)
    L = Light -> getDirection(Q)
    I ← I + atten * (diffuse term + specular term)
  end for
  return I
end function
```



12

Ray casting with shadows

Now we'll add shadows by casting shadow rays:




13

Shading with shadows

To include shadows, we need to modify the shade function:

```

function shade(mtrl, scene, Q, N, d):
  I ← mtrl.ke + mtrl.kd * IL
  for each light source Light in scene do:
    atten = Light -> distanceAttenuation(Q) *
    Light -> shadowAttenuation(Q, scene)
    L = Light -> getDirection(Q)
    I ← I + atten * (diffuse term + specular term)
  end for
  return I
end function
    
```



14

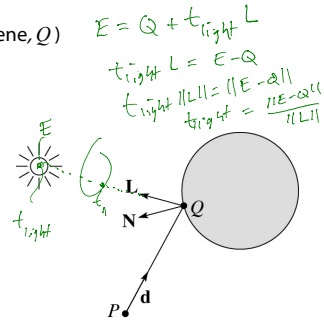
Shadow attenuation

Computing a shadow can be as simple as checking to see if a ray makes it to the light source.

For a point light source:

```

function PointLight::shadowAttenuation(scene, Q)
  L = getDirection(Q)
  (tr, N, mtrl) ← scene.intersect(Q, L)
  Compute tlight
  if (tr < tlight) then:
    atten = (0, 0, 0)
  else
    atten = (1, 1, 1)
  end if
  return atten
end function
    
```




Note: we will later handle color-filtered shadowing, so this function needs to return a *color* value.

For a directional light, $t_{\text{light}} = \infty$.

15


Recursive ray tracing with reflection

Now we'll add reflection:



16

Shading with reflection



Let $I(P, \mathbf{d})$ be the intensity seen along a ray. Then:

$$I(P, \mathbf{d}) = I_{\text{direct}} + I_{\text{reflected}}$$


where

- I_{direct} is computed from the Blinn-Phong model, plus shadow attenuation
- $I_{\text{reflected}} = k_r I(Q, \mathbf{R})$

Typically, we set $k_r = k_s$. (k_r is a color value.)

17

Reflection



Law of reflection:

$$\theta_i = \theta_r$$


\mathbf{R} is co-planar with \mathbf{d} and \mathbf{N} .

18

Ray-tracing pseudocode, revisited

```

function traceRay(scene, P, d):
    (t, n, mtrl) ← scene.intersect(P, d)
    Q ← ray(P, d) evaluated at t
    I = shade(scene, mtrl, Q, N, d)
    R = reflectDirection(-d, N)
    I ← I + mtrl.k_r * traceRay(scene, Q, R)
    return I
end function
    
```



19

Terminating recursion

Q: How do you bottom out of recursive ray tracing?

Possibilities:

do this → stop after N bounces


$\prod k_{r_i} \leq \text{thresh} \Rightarrow \text{stop}$

Adaptive (or early) ray termination

20


Whitted ray tracing

Finally, we'll add refraction, giving us the Whitted ray tracing model:



21

Shading with reflection and refraction



Let $I(P, \mathbf{d})$ be the intensity seen along a ray. Then:

$$I(P, \mathbf{d}) = I_{\text{direct}} + I_{\text{reflected}} + I_{\text{transmitted}}$$

where

- I_{direct} is computed from the Blinn-Phong model, plus shadow attenuation
- $I_{\text{reflected}} = k_r I(Q, \mathbf{R})$
- $I_{\text{transmitted}} = k_t I(Q, \mathbf{T})$

Typically, we set $k_r = k_s$ and $k_t = 1 - k_s$ (or $(0,0,0)$, if opaque, where k_i is a color value).

[Generally, k_r and k_t are determined by "Fresnel reflection," which depends on angle of incidence and changes the polarization of the light. This is discussed in Marschner's textbook and can be implemented for extra credit.]

22

Refraction


P
 a_d

Snell's law of refraction:

$$\eta_i \sin \theta_i = \eta_t \sin \theta_t$$

where η_i, η_t are **indices of refraction**.


In all cases, \mathbf{R} and \mathbf{T} are co-planar with \mathbf{d} and \mathbf{N} .



The index of refraction is material dependent.

It can also vary with wavelength, an effect called **dispersion** that explains the colorful light rainbows from prisms. (We will generally assume no dispersion.)

Medium	Index of refraction
Vacuum	1
Air	1.0003
Water	1.33
Fused quartz	1.46
Glass, crown	1.52
Glass, dense flint	1.66
Diamond	2.42



Index of refraction variation for fused quartz

23

Total Internal Reflection


The equation for the angle of refraction can be computed from Snell's law:

$$n_i \sin \theta_i = n_t \sin \theta_t \quad \theta_t = \sin^{-1} \left(\frac{n_i}{n_t} \sin \theta_i \right)$$

What "bad thing" can happen when $n_i > n_t$?

When θ_i is exactly 90° , we say that θ_i has achieved the "critical angle" θ_c .

For $\theta_i > \theta_c$, *no rays are transmitted*, and only reflection occurs, a phenomenon known as "total internal reflection" or TIR.



24

Marschner's notation

Marschner uses different symbols. Here is the translation between them:


$r = R$
 $t = T$
 $\phi = \theta_i$
 $\theta = \theta_r = \theta_t$
 $n = \eta_i$
 $n_t = \eta_t$

25

Ray-tracing pseudocode, revisited

```

function traceRay(scene, P, d):
    (t, N, mtrl) ← scene.intersect(P, d)
    Q ← ray(P, d) evaluated at t
    I = shade(scene, mtrl, Q, N, d)
    R = reflectDirection(N, d)
    I ← I + mtrl.kr * traceRay(scene, Q, R)
    if ray is entering object then
        ηi = index_of_air (=1.0003)
        ηt = mtrl.index
    else
        ηi = mtrl.index
        ηt = index_of_air (=1.0003)
    if (not TIR(φ, N, ηi, ηt)) then
        T = refractDirection(
            )
        I ← I + mtrl.kt * traceRay(scene, Q, T)
    end if
    return I
end function
    
```



Q: How do we decide if a ray is entering the object?

26

Terminating recursion, incl. refraction


Q: Now how do you bottom out of recursive ray tracing?

threshold on bounces
 Adaptive: $\{k_{s_i}, \alpha_i\}$

27

Shadow attenuation (cont'd)

Q: What if there are transparent objects along a path to the light source?




We'll take the view that the color is really only at the surface, like a glass object with a colored transparency coating on it. In this case, we multiply in the transparency constant, k_t , every time an object is entered or exited, possibly more than once for the same object.

28

Shadow attenuation (cont'd)


Another model would be to treat the glass as solidly colored in the interior. Marschner's textbook describes a the resulting volumetric attenuation based on Beer's Law, which you can implement for extra credit.





29

Photon mapping

Combine light ray tracing (photon tracing) and eye ray tracing:



...to get **photon mapping**.



Renderings by Henrik Wann Jensen:

<http://graphics.ucsd.edu/~henrik/images/caustics.html>


30

Normals and shading, reflection, and refraction when inside

When a ray is inside an object and intersects the object's surface on the way out, the normal will be pointing **away** from the ray (i.e., the normal always points to the outside by default).

You must **negate** the normal before doing any of the **shading, reflection, and refraction** that follows.

Finally, when shading a point inside of an object, apply k_i to the ambient component, since that "ambient light" had to pass through the object to get there in the first place.




31

Intersecting rays with spheres

Now we've done everything except figure out what that "scene.intersect(P, d)" function does.

Mostly, it calls each object to find out the t -value at which the ray intersects the object. Let's start with intersecting spheres...



Given:

- The coordinates of a point along a ray passing through P in the direction d are:

$$\begin{aligned} x &= P_x + td_x \\ y &= P_y + td_y \\ z &= P_z + td_z \end{aligned}$$

$P+td$ implicit surfaces

- A unit sphere S centered at the origin defined by the equation:

$$x^2 + y^2 + z^2 = 1$$

$$f(x, y, z) = x^2 + y^2 + z^2 - 1$$

$$\nabla f = \begin{bmatrix} 2x \\ 2y \\ 2z \end{bmatrix}$$

Find: The t at which the ray intersects S .

32

Intersecting rays with spheres

Solution by substitution:

$$\begin{aligned} x^2 + y^2 + z^2 - 1 &= 0 \\ (P_x + td_x)^2 + (P_y + td_y)^2 + (P_z + td_z)^2 - 1 &= 0 \\ at^2 + bt + c &= 0 \end{aligned}$$

where

$$\begin{aligned} a &= d_x^2 + d_y^2 + d_z^2 \\ b &= 2(P_x d_x + P_y d_y + P_z d_z) \\ c &= P_x^2 + P_y^2 + P_z^2 - 1 \end{aligned}$$


Q: What are the solutions of the quadratic equation in t and what do they mean?

$$t = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

discriminant $\Delta = b^2 - 4ac$


Q: What is the normal to the sphere at a point (x, y, z) on the sphere?

$$N = \begin{bmatrix} x \\ y \\ z \end{bmatrix}$$



33

Ray-plane intersection



Next, we will consider intersecting a ray with a plane.

To do this, we first need to define the plane equation.

Given a point S on a plane with normal N , how would we determine if another point X is on the plane?


(Hint: start by forming the vector $X - S$.)

$$\begin{aligned} N \cdot (X - S) &= 0 \\ N \cdot X - N \cdot S &= 0 & N \cdot X &= k \\ N \cdot X &= \frac{N \cdot S}{1} & & \end{aligned}$$

This is the plane equation!

34

Ray-plane intersection (cont'd)



Now consider a ray intersecting a plane, where the plane equation is:

$$N \cdot X = k$$

We can solve for the intersection parameter (and thus the point) by substituting X with the ray $P + t d$.


$$\begin{aligned} N \cdot (P + t d) &= k \\ N \cdot P + t N \cdot d &= k \\ t &= \frac{k - N \cdot P}{N \cdot d} \end{aligned}$$

$$N \cdot d = 0 \Rightarrow d \parallel \text{plane}$$

↓
no \cap

35

Ray-triangle intersection



To intersect with a triangle, we first solve for the equation of its supporting plane.

How might we compute the (un-normalized) normal?

$$N \approx (B - A) \times (C - A) \quad \hat{N} = \frac{N}{|N|}$$

Given this normal, how would we compute k ?

$$N \cdot X = k \quad N \cdot A = k = N \cdot B = N \cdot C$$

Using these coefficients, we can intersect the ray with the triangle to solve for Q .


Now, we need to decide if Q is inside or outside of the triangle...

36

Computing barycentric coordinates

Recall that given a point Q that is inside of triangle ABC , we can solve for Q 's barycentric coordinates in a simple way:

$$\alpha = \frac{\text{SArea}(QBC)}{\text{SArea}(ABC)} \quad \beta = \frac{\text{SArea}(AQC)}{\text{SArea}(ABC)} \quad \gamma = \frac{\text{SArea}(ABQ)}{\text{SArea}(ABC)}$$



$$\alpha A + \beta B + \gamma C$$

$$u \cdot v = \|u\| \|v\| \sin \theta \hat{N}$$


Before, we computed signed areas with 2D cross products. How can we compute signed area with 3D cross products?

dot all X-products w/ \hat{N}

How can we use the barycentric coordinates to decide if Q is inside ABC ? *α or β or $\gamma < 0 \Rightarrow$ no \hat{N}*

Improvement: project down to 2D first

Without loss of generality, we can make this determination after projecting down a dimension:



$$Q = \alpha A + \beta B + \gamma C$$

$$Q' = \alpha A' + \beta B' + \gamma C'$$

If Q' is inside of $A'B'C'$, then Q is inside of ABC .


Why is this projection desirable? *faster - fewer calculations in 2D*

Which axis should you "project away"? (Hint: *one w/ max dot product w/ \hat{N}*)

Per-face normals for triangle meshes

We will be shading and calculating reflections and refractions based on surface normals.

For a triangle mesh, we can make the natural assumption that each triangle has a constant normal (the normal of its supporting plane):



Recall the Blinn-Phong shading equation for a single light source (no ambient or emissive):

$$I = I_L B \left[k_d (\mathbf{N} \cdot \mathbf{L}) + k_s (\mathbf{N} \cdot \mathbf{H})_+^n \right]$$

Typically, \mathbf{L} and \mathbf{V} vary only a small amount over each triangle, if at all. *$L, V \approx \text{const}$*


Q: If material properties (k_d, k_s, n_s) are constant over the mesh, how will shading vary within a triangle? *Not*

Phong interpolated normals

For better results, we can "hallucinate" smoothly varying within a triangle, a process called **Phong interpolation**. (Not the same as "Phong shading.")

Here's how it works:

1. Compute normals at the vertices.
2. Interpolate normals and normalize.
3. Shade, reflect, and refract using the interpolated normals.



Q: How do we compute per vertex normals? *average of neighboring face normals*

Q: How do we interpolate N_a, N_b, N_c to get N ? *$N = \alpha N_a + \beta N_b + \gamma N_c$
 $\hat{N} = \frac{N}{\|N\|}$*

Interpolating vertex properties

The barycentric coordinates can be used to interpolate a variety of vertex properties such as:

- material properties
- texture coordinates
- normals

For example:

$$k_d(Q) = \alpha k_d(A) + \beta k_d(B) + \gamma k_d(C)$$


Interpolating normals, known as **Phong interpolation**, gives triangle meshes a smooth shading appearance. (Note: don't forget to normalize interpolated normals.)

41

Epsilons

Due to finite precision arithmetic, we do not always get the exact intersection at a surface.

Q: What kinds of problems might this cause?



Q: How might we resolve this?


$t > \text{RAY_EPSILON} \Rightarrow \cap$
else no \cap

42

Intersecting with xformed geometry

In general, objects will be placed using transformations. What if the object being intersected were transformed by a matrix M ?

Apply M^{-1} to the ray first and intersect in object (local) coordinates! **Note:** do not normalize d !



$$M^{-1}(Q = P + t_n d)$$

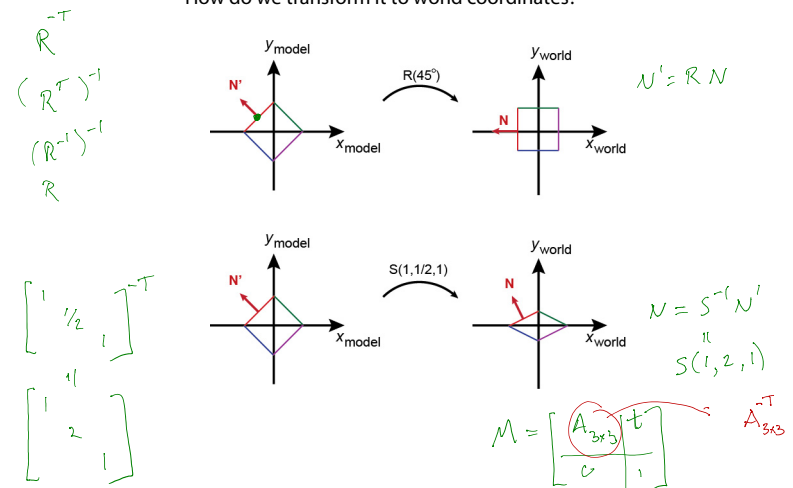
$$M^{-1}Q = M^{-1}P + M^{-1}(t_n d)$$

$$Q' = P' + t_n d'$$

43

Intersecting with xformed geometry

The intersected normal is in object (local) coordinates. How do we transform it to world coordinates?



44

Summary

What to take home from this lecture:

- ♦ The meanings of all the boldfaced terms.
- ♦ Enough to implement basic recursive ray tracing.
- ♦ How reflection and transmission directions are computed.
- ♦ How ray-object intersection tests are performed on spheres, planes, and triangles
- ♦ How barycentric coordinates within triangles are computed
- ♦ How ray epsilons are used.