

Ray Tracing

**Adriana Schulz
CSE 457
Fall 2020**

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**.

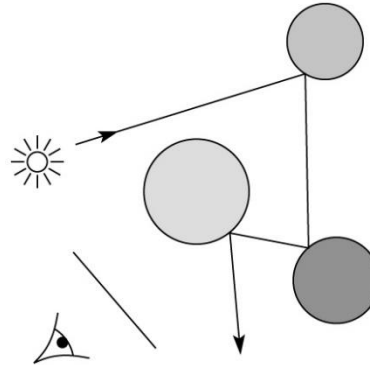
Here are the rules of geometric optics:

- ◆ 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



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**.

Whitted algorithm (cont'd)


Let's look at this in stages:




Primary rays



Shadow rays




Reflection rays




Refracted rays

Ray casting and local illumination

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



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

Ray-tracing pseudocode

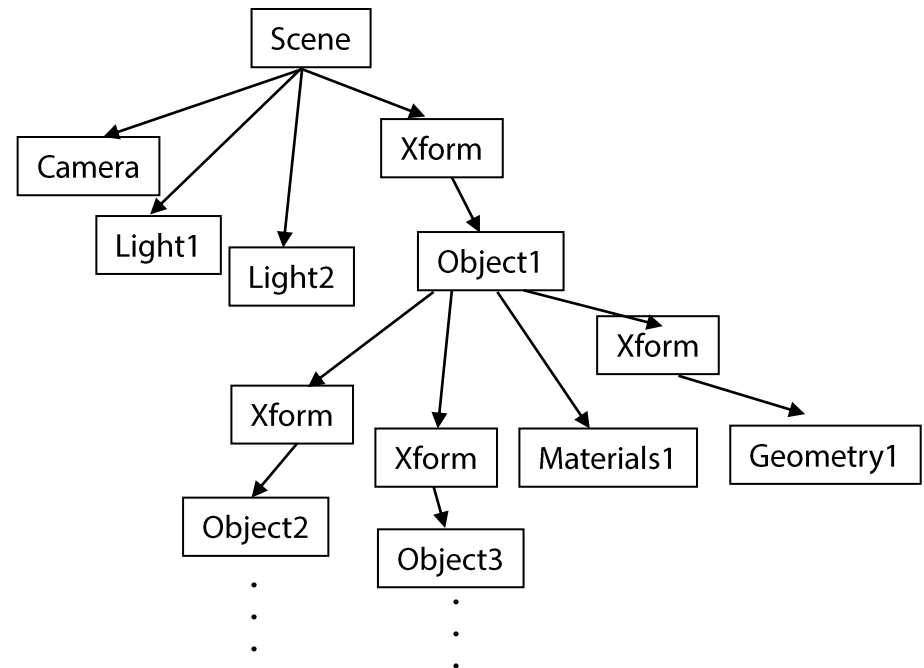

We build a ray traced image by casting rays through each of the pixels.

```

function tracelImage (scene):
  for each pixel (i, j) in image
     $A = \text{pixelToWorld}(i, j)$ 
     $P = C$ 
     $\mathbf{d} = (A - P) / \|A - P\|$ 
     $I(i, j) = \text{traceRay}(\text{scene}, P, \mathbf{d})$ 
  end for
end function
  
```

```

function traceRay (scene,  $P, \mathbf{d}$ ):
   $(t_{\cap}, \mathbf{N}, \text{mtrl}) \leftarrow \text{scene.intersect}(P, \mathbf{d})$ 
   $Q \leftarrow \text{ray}(P, \mathbf{d})$  evaluated at  $t_{\cap}$ 
   $I = \text{shade}( \dots )$ 
  return  $I$ 
end function
  
```



Shading pseudocode

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

function *shade* (mtrl, scene, Q , \mathbf{N} , \mathbf{d}):

$I \leftarrow \text{mtrl}.k_e$

for each light source Light **do**:

$\text{atten} = \text{Light} \rightarrow \text{distanceAttenuation} (\quad)$


$\mathbf{L} = \text{Light} \rightarrow \text{getDirection} (\quad)$

$I \leftarrow I + \text{ambient} + \text{atten} * (\text{diffuse} + \text{specular})$

end for


return I

end function



Ray casting with shadows

Now we'll add shadows by casting shadow rays:



Shading with shadows

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

function *shade* (mtrl, scene, Q , \mathbf{N} , \mathbf{d}):

$I \leftarrow \text{mtrl}.k_e$

for each light source Light **do**:

$\text{atten} = \text{Light} \rightarrow \text{distanceAttenuation}(Q) *$

$\text{Light} \rightarrow \text{shadowAttenuation}(\quad)$


$\mathbf{L} = \text{Light} \rightarrow \text{getDirection}(Q)$

$I \leftarrow I + \text{ambient} + \text{atten} * (\text{diffuse} + \text{specular})$

end for

return I

end function



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)

$L = \text{getDirection}(Q)$

$(t_{\cap}, \mathbf{N}, \text{mtrl}) \leftarrow \text{scene.intersect}(Q, \mathbf{L})$

Compute t_{light}

if $(t_{\cap} < t_{\text{light}})$ **then:**

$\text{atten} = (0, 0, 0)$

else

$\text{atten} = (1, 1, 1)$


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, $t_{\text{light}} = \infty$.



Shading in “Trace”

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

- ◆ Distance attenuation is clamped to be at most 1:

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

- ◆ Shadow attenuation A^{shadow} is included and is RGB-valued.


Here’s what the shading equation should look like:

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


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

Recursive ray tracing with reflection

Now we'll add reflection:



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
- ♦ $I_{\text{reflected}} = k_s I(Q, \mathbf{R})$

Remember that is a color value.

(Sometimes another variable, k_r , is used instead of k_s to allow for separate control of specular light reflection vs specular interreflection. We will not do this.)

Reflection




Law of reflection:

$$\theta_i = \theta_r$$

R is co-planar with d and N .

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 ( )  
    I ← I + mtrl.ks * traceRay(scene, Q, R)  
    return I  
end function
```



Terminating recursion


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

Possibilities:




Whitted ray tracing

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



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
- ◆ $I_{\text{reflected}} = k_s I(Q, \mathbf{R})$
- ◆ $I_{\text{transmitted}} = k_t I(Q, \mathbf{T})$

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

[Generally, for ideal specular surfaces, k_s 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.]


Refraction

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, **R** and **T** are co-planar with **d** and **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


Total Internal Reflection

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

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.



Marschner's notation

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

$$\mathbf{r} = \mathbf{R}$$

$$\mathbf{t} = \mathbf{T}$$

$$f = q_t$$


$$q = q_r = q_i$$

$$n = h_i$$

$$n_t = h_t$$

Ray-tracing pseudocode, revisited


```
function traceRay (scene,  $P$ ,  $\mathbf{d}$ ):  
    ( $t_{\cap}$ ,  $\mathbf{N}$ , mtrl)  $\leftarrow$  scene.intersect ( $P$ ,  $\mathbf{d}$ )  
     $Q \leftarrow$  ray ( $P$ ,  $\mathbf{d}$ ) evaluated at  $t_{\cap}$   
     $I =$  shade (scene, mtrl,  $Q$ ,  $\mathbf{N}$ ,  $\mathbf{d}$ )  
     $\mathbf{R} =$  reflectDirection ( $\mathbf{N}$ ,  $\mathbf{d}$ )  
     $I \leftarrow I +$  mtrl. $k_s$  * traceRay (scene,  $Q$ ,  $\mathbf{R}$ )  
    if ray is entering object then  
         $\eta_i =$  index_of_air (=1.0003)  
         $\eta_t =$  mtrl.index  
    else  
         $\eta_i =$  mtrl.index  
         $\eta_t =$  index_of_air (=1.0003)  
    if (notTIR (  $\eta_i, \eta_t$  )) then  
         $\mathbf{T} =$  refractDirection ( $\mathbf{N}$ ,  $\mathbf{d}$ ,  $\eta_i, \eta_t$ )  
         $I \leftarrow I +$  mtrl. $k_t$  * traceRay (scene,  $Q$ ,  $\mathbf{T}$ )  
    end if  
    return  $I$   
end function
```



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


Terminating recursion, incl. refraction

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



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.


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.





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>

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_t to the ambient component, since that "ambient light" had to pass through the object to get there in the first place.




Geometric Aspects

Intersecting rays with spheres

Now we've done everything except figure out what that "scene.intersect(P , \mathbf{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 \mathbf{d} are:

$$x = P_x + td_x$$

$$y = P_y + td_y$$

$$z = P_z + td_z$$

- ◆ A sphere S of radius r centered at the origin defined by the equation:

Find: The t at which the ray intersects S .

Intersecting rays with spheres

Solution by substitution:

$$x^2 + y^2 + z^2 - r^2 = 0$$

$$(P_x + td_x)^2 + (P_y + td_y)^2 + (P_z + td_z)^2 - r^2 = 0$$

$$at^2 + bt + c = 0$$

where

$$a = d_x^2 + d_y^2 + d_z^2$$

$$b = 2(P_x d_x + P_y d_y + P_z d_z)$$


$$c = P_x^2 + P_y^2 + P_z^2 - r^2$$

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

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

Note: the Trace project only requires you to handle a sphere of radius $r = 0.5$. This sphere may be arbitrarily transformed when placed in the scene, but everything works out in the end. More later!

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 \mathbf{N} , how would we determine if another point X is on the plane?

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

This is the plane equation!

Ray-plane intersection (cont'd)




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

$$\mathbf{N} \times X = k$$

We can solve for the intersection parameter (and thus the point) by substituting X with the ray $P + t \mathbf{d}$:

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?


Given this normal, how would we compute k ?

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...

3D inside-outside test

One way to do this “inside-outside test,” is to see if Q lies on the left side of each edge as we move counterclockwise around the triangle.



How might we use cross and products to do this?


Barycentric coordinates

As we'll see in a moment, it is often useful to represent Q as an **affine combination** of A , B , and C :

$$Q = aA + bB + gC$$

where:

$$\underline{a + b + g = 1}$$




We call α , β , and γ , the **barycentric coordinates** of Q with respect to A , B , and C .

Computing barycentric coordinates

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

$$a = \frac{\text{Area}(QBC)}{\text{Area}(ABC)} \quad b = \frac{\text{Area}(AQC)}{\text{Area}(ABC)} \quad g = \frac{\text{Area}(ABQ)}{\text{Area}(ABC)}$$




How can cross products help here?

In the end, these calculations can be performed in the 2D projection as well!

Improvement: project down to 2D first

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



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

Why is this projection desirable?

Which axis should you “project away”? (Hint: consider the triangle normal.)

Interpolating vertex properties

The barycentric coordinates can also be used to interpolate vertex properties such as:

- ◆ material properties
- ◆ texture coordinates
- ◆ normals

For example:


$$k_d(Q) = ak_d(A) + bk_d(B) + gk_d(C)$$

Phong interpolated normals

Recall the idea of interpolating normal from the shading lecture, now updated to allow reflection and refraction.

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 interpolate \mathbf{N}_a , \mathbf{N}_b , \mathbf{N}_c to get \mathbf{N} ?

Interpolated normal in a ray tracer

As before, we will use the interpolated normal for shading, but a problem can arise when using this normal for other ray tracing purposes. Consider:



We see that the interpolated normal $\mathbf{N}_{\text{interp}}$ is of course different from the true geometric (triangle) normal \mathbf{N}_Δ . Here are a couple problems that can arise:



A similar problem can arise for refraction, in which the refracted direction \mathbf{T} is exiting the object, which is also wrong. There is no "right" answer for handling the discrepancy between normal and actual geometry.

Interpolated normal in a ray tracer, cont'd

We could play it “safe” and always use \mathbf{N}_Δ , but then we will not get nice curved reflections and refractions even when these rays are valid. For Trace, do the following...

Determining when entering/exiting object:

- ◆ Use the geometric/true normal (\mathbf{N}_Δ) when deciding whether you are entering/exiting an object.

Shading:

- ◆ Use $\mathbf{N}_{\text{interp}}$ for shading.

Reflection:

1. Start by using $\mathbf{N}_{\text{interp}}$ to compute reflection direction \mathbf{R} .
2. If \mathbf{R} is (incorrectly) entering the object, then re-compute \mathbf{R} using \mathbf{N}_Δ .

Refraction:

1. Start by using $\mathbf{N}_{\text{interp}}$ to check for Total Internal Reflection (TIR).
2. If TIR, then do not cast a refracted ray.
3. Else, use $\mathbf{N}_{\text{interp}}$ to compute refraction direction \mathbf{T} .
4. If \mathbf{T} is (incorrectly) exiting the object, then use \mathbf{N}_Δ to re-test for TIR and (if no TIR) to re-compute \mathbf{T} .

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?

Intersecting with xformed geometry

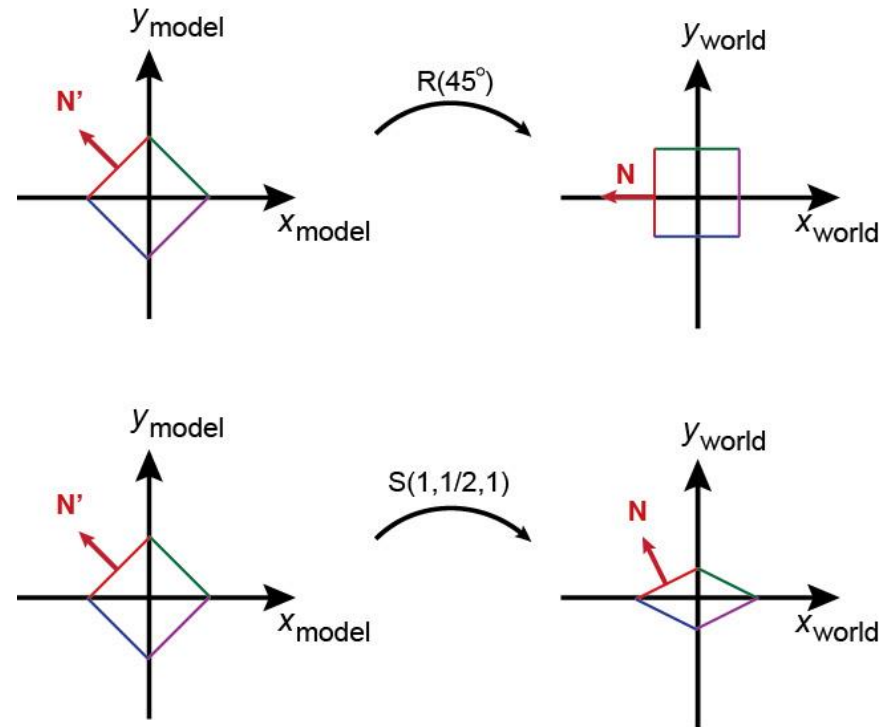
In general, objects will be placed using transformations. What if the object being intersected were transformed by a matrix M ?

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



Intersecting with xformed geometry

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



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.

Accelerated ray tracing

Reading

Required:


- ◆ Marschner and Shirley, Sections 12.3 (online handout)

Further reading:

- ◆ A. Glassner. An Introduction to Ray Tracing. Academic Press, 1989.

Faster ray-polyhedron intersection

Let's say you were intersecting a ray with a triangle mesh:




Straightforward method

- ◆ intersect the ray with each triangle
- ◆ return the intersection with the smallest t -value.

Q: How might you speed this up?

Bounding Volume Hierarchies (BVHs)

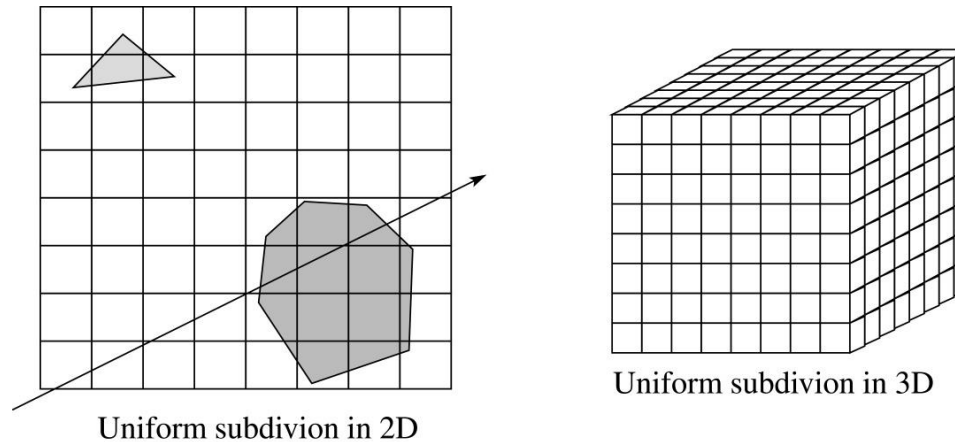
We can generalize the idea of bounding volume acceleration with **bounding volume hierarchies (BVHs)**.



Key: build balanced trees with *tight bounding volumes*.

Uniform spatial subdivision

Another approach is **uniform spatial subdivision**.




Idea:

- ◆ Partition space into cells (voxels)
- ◆ Associate each primitive with the cells it overlaps
- ◆ Trace ray through voxel array using fast incremental arithmetic to step from cell to cell


Q: Given a 10^6 triangle football stadium with a 10^6 triangle teapot on one of the seats, would a single uniform spatial subdivision be a good idea?

Non-uniform spatial subdivision: octrees

Another approach is **non-uniform spatial subdivision**. One version of this is octrees:




Quadtree in 2D




Octree in 3D

Non-uniform spatial subdivision: k -d trees

Another non-uniform subdivision is k -d
(k -dimensional) trees:



k -d tree ($d = 2$)



k -d tree ($d = 3$)

If the planes can be non-axis aligned, then you get BSP (binary space partitioning) trees.

Various combinations of these ray intersections techniques are also possible.

Summary

What to take home from this lecture:

- ◆ An intuition for how ray tracers can be accelerated.