Lecture 20: Ray Tracing

Yuhao Zhu

http://yuhaozhu.com

CSC 292/572, Fall 2022 **Mobile Visual Computing**



Logistics

PA3 due 12/5, 11:30 AM.

Final project due 12/20, 11:30 AM.



The Roadmap

Theoretical Preliminaries

Human Visual Systems

Color in Nature, Arts, & Tech (a.k.a., the birth, life, and death of light)

Digital Camera Imaging

Modeling and Rendering

Applications

Modeling

Rasterization

Ray Tracing

Shading & Texture



Graphics







Visibility Problem

Two fundamental classes of visibility algorithms

- Object-centric (Rasterization)
- Image-centric (Ray tracing)







Basic Idea of Ray Tracing



© www.scratchapixel.com

foreach pixel in image
 ray = buildRay(camera, pixel)
 if (P = intersect(ray, mesh))
 pixel.color = shade(P)
 else
 pixel.color = backgroundColor



Basic Idea of Ray Tracing



```
Can sample multiple rays
                               per pixel for anti-aliasing
foreach pixel in image
  ray = buildRay(camera, pixel) 4
  if (P = intersect(ray, mesh))
    pixel.color = shade(P)
  else
    pixel.color = backgroundColor
```



Basic Idea of Ray Tracing



```
Can sample multiple rays
                               per pixel for anti-aliasing
foreach pixel in image
  ray = buildRay(camera, pixel) 4
  if (P = intersect(ray, mesh))
    pixel.color = shade(P)
  else
    pixel.color = backgroundColor
```



Generating Rays



What Defines a Ray?

A ray is defined by its original and the directional vector.

- We usually define ray as a segment, with a min and a max.
 - Its purpose will become clear later, but briefly it allows us to reduce computation in visibility test.





Generating a Ray in a Pinhole Camera

Very intuitive: connect a pixel and the pinhole to form a ray.

Remember in actual implementations the canvas is before the pinhole.







Many rays incident on one pixel. So for each pixel we need to sample the lens multiple times to trace **multiple rays**.

- Why trace many rays? Because the pixel color depends on all incident rays.
- This is not something achievable using only the perspective matrix.





11

Goal: given an arbitrary R_{out} how do we find R_{in} ?

- There is a unique ray, R_{out} , between a pixel **P** and a sample **L** on the lens
- There is a unique ray, R_{in} , going into the lens that generates R_{out}
- The closest point on R_{in} before the lens will hit ${\boldsymbol{\mathsf{P}}}$



ixel **P** and a sample **L** on the lens e lens that generates R_{out} will hit **P**



Use geometrical optics principles to determine R_{in} for a given R_{out} • Rays go through the lens center (chief ray) don't change their directions





Use geometrical optics principles to determine R_{in} for a given R_{out}

- Rays go through the lens center (chief ray) don't change their directions
- Rays parallel to the optical axis (parallel ray) pass through the lens focus





Steps:

- 1. use chief and parallel rays to find the intersection point **S** in the scene
- 2. find **L** on the lens from R_{out}
- 3. R_{in} is the ray between ${\bm S}$ and ${\bm L}$





Can Ray Tracing Capture DOF?

focus on the pixel.



What's described before doesn't rely on whether the closest hit is actually in-

• So it can inherent trace scene points that are out-of-focus, i.e., simulate depth of field.

Can Ray Tracing Capture DOF?

focus on the pixel.

What's described before doesn't rely on whether the closest hit is actually in-

• So it can inherent trace scene points that are out-of-focus, i.e., simulate depth of field.

Can Ray Tracing Capture DOF?

focus on the pixel.

What's described before doesn't rely on whether the closest hit is actually in-

• So it can inherent trace scene points that are out-of-focus, i.e., simulate depth of field.

Lens Sampling

many rays for a point to reduce noise (more in shading lecture).

• This is orthogonal to sampling multiple points per pixel.

* the artifacts from low sampling rate here is not aliasing; it's due to high variance in Monte Carlo integration.

Each sensor plane point receives infinitely many rays, so we need to sample

http://www.pbr-book.org/3ed-2018/Camera_Models/Projective_Camera_Models.html 17

Ray-Scene Intersection

Ray-Scene Intersection

Goal: calculate the [x, y, z] coordinates of the closest hit between the ray and the mesh.

Why closest hit?

Preserve visibility (like the z-buffer in rasterization)

Ray-Scene Intersection

Goal: calculate the [x, y, z] coordinates of the closest hit between the ray and the mesh.

Why closest hit?

Preserve visibility (like the z-buffer in rasterization)

Brute-force approach:

- Brute-force approach:
 - iterate all triangles

Brute-force approach:

- iterate all triangles
- test intersection for each triangle

Brute-force approach:

- iterate all triangles
- test intersection for each triangle
- return the closest hit, if any

Brute-force approach:

- iterate all triangles
- test intersection for each triangle
- return the closest hit, if any

Key task:

Brute-force approach:

- iterate all triangles
- test intersection for each triangle
- return the closest hit, if any

Key task:

• Ray-triangle intersection test and calculate the coordinates of the hit point, if any.

What Defines a Triangle?

A plane with three vertices.

• The vertices are guaranteed to be co-planar.

The plane that the triangles are in can be expressed as an implicit equation and can be calculated from the vertices.

$$\begin{cases} A(V1_x - V2_x) + B(V1_y - V2_y) + C(V1_z - V2_z) = 0 \\ A(V1_x - V3_x) + B(V1_y - V3_y) + C(V1_z - V3_z) = 0 \\ A \times V1_x + B \times V1_y + C \times V1_z = X \end{cases} \square$$

Ray-Triangle Intersection

 $P_x = O_x + D_x \times t$ $P_y = O_y + D_y \times t$ $P_{z} = O_{z} + D_{z} \times t$ $A \times P_x + B \times P_y + C \times P_z = 1$ $\frac{1 - (A \times O_x + B \times O_y + C \times O_z)}{A \times D_x + B \times D_y + C \times D_z}$

Three Caveats

1. The denominator is 0 if the normal is perpendicular to the direction of the ray (i.e., ray is parallel to the plane).

• Need a special test for whether the ray is parallel with the plane (before division).

Three Caveats

2. A ray doesn't intersect with a plane if the triangle plane is behind the origin of the ray

• i.e., t is negative.

Three Caveats

2. A ray doesn't intersect with a plane if the triangle plane is behind the origin of the ray

• i.e., t is negative.

3. Even if a real intersection point is found, the intersection point could be outside the triangle.

• Use barycentric coordinates to test whether a point is outside of a triangle.

Brute-Force Approach is Extremely Inefficient

Brute-force approach:

- iterate all triangles
- test intersection for each triangle
- return the closest hit, if any

Time complexity:

• O(# of rays x # of triangles)

Accelerating Ray-Scene Intersection

Prune the search space.

Only search part of the scene that does intersect the ray.



intersect(space, ray) {
if ray doesn't intersect space boundary:
 return



Prune the search space.

Only search part of the scene that does intersect the ray.



intersect(space, ray) {
if ray doesn't intersect space boundary:
 return



Prune the search space.

Only search part of the scene that does intersect the ray.



intersect(space, ray) {
if ray doesn't intersect space boundary:
 return



Prune the search space.

Only search part of the scene that does intersect the ray. Key: how to partition the space?



intersect(space, ray) {
if ray doesn't intersect space boundary:
 return



Object vs. Space Partitioning

Space partitioning: One object could be in different partitions



Object partitioning: different partitions could overlap in space











Find the bounding box of the scene





Find the bounding box of the scene Generate a uniform grid





- Find the bounding box of the scene
- Generate a uniform grid
- Find intersecting cells





- Find the bounding box of the scene
- Generate a uniform grid
- Find intersecting cells





- Find the bounding box of the scene
- Generate a uniform grid
- Find intersecting cells
- For each intersecting cell:





- Find the bounding box of the scene
- Generate a uniform grid
- Find intersecting cells
- For each intersecting cell:
 - Iterate over all the containing triangles





- Find the bounding box of the scene
- Generate a uniform grid
- Find intersecting cells
- For each intersecting cell:
 - Iterate over all the containing triangles
 - Get the closet intersection within the cell





- Find the bounding box of the scene
- Generate a uniform grid
- Find intersecting cells
- For each intersecting cell:
 - Iterate over all the containing triangles
 - Get the closet intersection within the cell
 - Update the global closet intersection



Grid Resolution



Too few cells:

• No speedup

Too many cells:

Many empty cells to check and to store

A useful heuristics:

- The number of cells should be proportional to the number of triangles
- #cell in each dimension = $n^{1/3}$



When Uniform Grid Works



Small objects roughly uniformly distributed in space



When Uniform Grid Fails



Objects sparsely distributed in space ("teapot in a stadium")

































Quadtree (2D)





Octree (3D)









С



http://groups.csail.mit.edu/graphics/classes/6.838/S98/meetings/m13/kd.html 35





Building K-D Tree



- Recursively using axis-aligned planes to split the space
- Stop when certain terminating conditions are met
 - # of objects in a cell < threshold
 - Max tree depth met
- Organize the splits using a tree
- Find the closest hit by traversing the tree



































С



http://groups.csail.mit.edu/graphics/classes/6.838/S98/meetings/m13/kd.html















С



http://groups.csail.mit.edu/graphics/classes/6.838/S98/meetings/m13/kd.html





























Binary Space Partitioning Tree



- K-D tree is a special case of binary space partitioning (BSP) tree, which recursively split the space with planes (3D) or lines (2D) • Arbitrary split planes here
- Useful when objects are large and non-axisaligned, in which case K-D tree will split objects into different partitions
 - Good reference: Ray Tracing with the BSP Tree [Ize, Wald, Parker, 2008]





Binary Space Partitioning Tree



- K-D tree is a special case of binary space partitioning (BSP) tree, which recursively split the space with planes (3D) or lines (2D) • Arbitrary split planes here
- Useful when objects are large and non-axisaligned, in which case K-D tree will split objects into different partitions
 - Good reference: Ray Tracing with the BSP Tree [Ize, Wald, Parker, 2008]




BVH Tree





BVH Tree





BVH Tree







BVH Tree







BVH Tree







BVH Tree







BVH Tree







BVH Tree







BVH Tree







BVH Tree







BVH Tree







BVH Tree







Object vs. Space Partitioning

Space partitioning: One object could be in different partitions



Object partitioning: different partitions could overlap in space







• A, B, C, D, E are the bounding volumes, which are Axis-Aligned Bounding Boxes (AABBs) here. Other (irregular) bounding volumes are possible.







ClosestHit = NA



Intersection Test Using BVH

Current

Ray-AABB Intersection Test

Stack





ClosestHit = NA



Intersection Test Using BVH

- Current
 - Stack







Intersection Test Using BVH

ClosestHit = NA



- Current
 - Stack







ClosestHit = NA



Intersection Test Using BVH

E

- Current
 - Stack







Intersection Test Using BVH

ClosestHit = NA



- Current
 - Stack

E







ClosestHit = 2



Intersection Test Using BVH

Current

Ray-AABB Intersection Test IFI

Stack





ClosestHit = 2

Distance to E > Distance to 2; Stop!



Intersection Test Using BVH

Current

Ray-AABB Intersection Test I E I

Stack





Ray-AABB Intersection

Ray: $\mathbf{O} + t\mathbf{D}$, $t_{min} \le t \le t_{max}$ 0 t_{min} $\mathsf{t}_{\mathsf{hit}}$ t_{max} D









Should this be counted as a hit?







Should this be counted as a hit?







Should this be counted as a hit? Yes; any ray segment that originates from within an AABB must be treated as intersecting.





Various Trade-offs Worth Considering

Time to build the tree vs. time to search.

- Incrementally update a tree (e.g., scene) slowing changing in an animation)?
- Can we built the tree offline?

Shape of the bounding volume.

• Tight bounding volumes provide more precise intersect test, but are costly to build and to search.

Tree structures take memory.





https://www.scratchapixel.com/lessons/advanced-rendering/introduction-acceleration-structure/bounding-volume-hierarchy-BVH-part1 58









Recursive Ray Tracing



• To implement realistic shading.





- To implement realistic shading.
- The color of an exiting ray depends on the colors of all incident rays.





- To implement realistic shading.
- The color of an exiting ray depends on the colors of all incident rays.
 - color here really means radiance.





- To implement realistic shading.
- The color of an exiting ray depends on the colors of all incident rays.
 - color here really means radiance.
 - also depends on the surface material (diffuse vs. specular vs. ...); later.





- To implement realistic shading.
- The color of an exiting ray depends on the colors of all incident rays.
 - color here really means radiance.
 - also depends on the surface material (diffuse vs. specular vs. ...); later.
- How do we know the color of an incident ray? Cast more rays!





• To implement realistic shading.





- To implement realistic shading.
- The color of an exiting ray depends on the colors of all incident rays.

Secondary Ray

Secondary Ray Secondary Ray

Valette, et al. [TVCG'08]


Why Recursive Ray Tracing?

- To implement realistic shading.
- The color of an exiting ray depends on the colors of all incident rays.
 - color here really means radiance.

Secondary Ray

Secondary Ray Secondary Ray

Valette, et al. [TVCG'08]



Why Recursive Ray Tracing?

- To implement realistic shading.
- The color of an exiting ray depends on the colors of all incident rays.
 - color here really means radiance.
 - also depends on the surface material (diffuse vs. specular vs. ...); later.



Valette, et al. [TVCG'08]



Why Recursive Ray Tracing?

- To implement realistic shading.
- The color of an exiting ray depends on the colors of all incident rays.
 - color here really means radiance.
 - also depends on the surface material (diffuse vs. specular vs. ...); later.
- How do we know the color of an incident ray? Cast more rays!

Secondary Ray Secondary Ray Secondary Ray

Valette, et al. [TVCG'08]



Simple Whitted-Style Recursive Ray Tracing

A simplification of <u>Whitted-style ray tracing</u>, assuming purely transparent surface.

castRay(ray, mesh) { if (P = nearestIntersect(ray, mesh)) reflectRay = buildReflectRay(P) refractRay = buildRefractRay(P) **reflectColor** = castRay(reflectRay, mesh)) **refractColor** = castRay(refractRay, mesh)) float kr fresnel(dir, N, hitObject->ior, kr) P.color = reflectionColor * kr + **refractionColor** * (1 - kr) else P.color = backgroundColor



https://www.scratchapixel.com/lessons/3d-basic-rendering/ray-tracing-overview/light-transport-ray-tracing-whitted





Simple Whitted-Style Recursive Ray Tracing







https://blogs.nvidia.com/blog/2018/08/01/ray-tracing-global-illumination-turner-whitted/ 64





Things to Remember

Ray tracing makes it easy (conceptually) to implement realistic shading.

Compared to rasterization, ray tracing is much more time consuming, dominated by ray-scene intersection test, which is exacerbated by the need for recursive ray tracing.

We can accelerate the testing using acceleration structures that prune the search space. BVH is the most common acceleration structure.

Modern GPUs, while traditionally optimized for rasterization, now have hardware support for ray tracing (e.g., BVH traversal, ray-AABB/triangle intersection test).

