

Distribution Ray Tracing

Brian Curless
CSEP 557
Winter 2013

Reading



Required:

- ♦ Shirley, section 10.11

Further reading:

- ♦ Watt, sections 10.4-10.5
- ♦ A. Glassner. An Introduction to Ray Tracing. Academic Press, 1989. [In the lab.]
- ♦ Robert L. Cook, Thomas Porter, Loren Carpenter. "Distributed Ray Tracing." Computer Graphics (Proceedings of SIGGRAPH 84). 18 (3). pp. 137-145. 1984.
- ♦ James T. Kajiya. "The Rendering Equation." Computer Graphics (Proceedings of SIGGRAPH 86). 20 (4). pp. 143-150. 1986.

Pixel anti-aliasing



All of this assumes that inter-reflection behaves in a mirror-like fashion...

BRDF

The diffuse+specular parts of the Blinn-Phong illumination model are a mapping from light to viewing directions:

$$I = I_b \left[k_d (\mathbf{N} \cdot \mathbf{L}) + k_s \left(\mathbf{N} \cdot \frac{\mathbf{L} + \mathbf{V}}{\|\mathbf{L} + \mathbf{V}\|} \right)^{n_s} \right]$$


$$= I_b f_r(\mathbf{L}, \mathbf{V})$$

The mapping function f_r is often written in terms of incoming (light) directions ω_{in} and outgoing (viewing) directions ω_{out} :

$$f_r(\omega_{in}, \omega_{out}) \quad \text{or} \quad f_r(\omega_{in} \rightarrow \omega_{out})$$

This function is called the **Bi-directional Reflectance Distribution Function (BRDF)**.

Here's a plot with ω_{in} held constant:




BRDF's can be quite sophisticated...

Light reflection with BRDFs

BRDF's exhibit Helmholtz reciprocity:

$$f_r(\omega_{in} \rightarrow \omega_{out}) = f_r(\omega_{out} \rightarrow \omega_{in})$$

That means we can take two equivalent views of reflection. Suppose $\omega_{in} = \mathbf{L}$ and $\omega_{out} = \mathbf{V}$:



We can now think of the BRDF as weighting light coming in from all directions, which can be added up:

$$I(\mathbf{V}) = \int_H I(\mathbf{L}) f_r(\mathbf{L} \rightarrow \mathbf{V})(\mathbf{L} \cdot \mathbf{N}) d\mathbf{L}$$

Or, written more generally:


$$I(\omega_{out}) = \int_H I(\omega_{in}) f_r(\omega_{in} \rightarrow \omega_{out})(\omega_{in} \cdot \mathbf{N}) d\omega_{in}$$

5


Simulating gloss and translucency

The mirror-like form of reflection, when used to approximate glossy surfaces, introduces a kind of aliasing, because we are under-sampling reflection (and refraction).

For example:




Distributing rays over reflection gives:



6

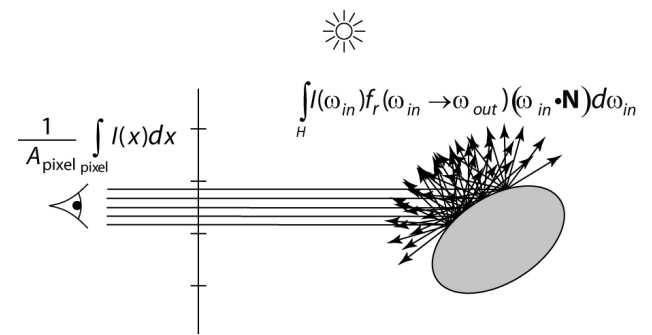
Reflection anti-aliasing



Reflection anti-aliasing

7


Pixel and reflection anti-aliasing



Pixel and reflection anti-aliasing

8

Full anti-aliasing



Full anti-aliasing...lots of nested integrals!


Computing these integrals is prohibitively expensive, especially after following the rays recursively.

We'll look at ways to approximate high-dimensional integrals...


9

Glossy reflection revisited

Let's return to the glossy reflection model, and modify it – for purposes of illustration – as follows:




We can visualize the span of rays we want to integrate over, within a pixel:




10

Whitted ray tracing

Returning to the reflection example, Whitted ray tracing replaces the glossy reflection with mirror reflection:




Thus, we render with anti-aliasing as follows:




11

Monte Carlo path tracing

Let's return to our original (simplified) glossy reflection model:




An alternative way to follow rays is by making random decisions along the way – a.k.a., Monte Carlo path tracing. If we distribute rays uniformly over pixels and reflection directions, we get:




12

Importance sampling

The problem is that lots of samples are “wasted.”
Using again our glossy reflection model:




Let’s now randomly choose rays, but according to a probability that favors more important reflection directions, i.e., use **importance sampling**:




13

Stratified sampling

We still have a problem that rays may be clumped together. We can improve on this by splitting reflection into zones:




Now let’s restrict our randomness to within these zones, i.e. use **stratified sampling**:



14

Stratified sampling of a 2D pixel

Here we see pure uniform vs. stratified sampling over a 2D pixel (here 16 rays/pixel):



The stratified pattern on the right is also sometimes called a **jittered** sampling pattern.

One interesting side effect of these stochastic sampling patterns is that they actually injects noise into the solution (slightly grainier images). This noise tends to be less objectionable than aliasing artifacts.

15

Distribution ray tracing

These ideas can be combined to give a particular method called **distribution ray tracing** [Cook84]:

- ♦ uses non-uniform (jittered) samples.
- ♦ replaces aliasing artifacts with noise.
- ♦ provides additional effects by distributing rays to sample:
 - Reflections and refractions
 - Light source area
 - Camera lens area
 - Time

[This approach was originally called “distributed ray tracing,” but we will call it distribution ray tracing (as in probability distributions) so as not to confuse it with a parallel computing approach.]

16

DRT pseudocode

TraceImage() looks basically the same, except now each pixel records the average color of jittered sub-pixel rays.

```

function traceImage(scene):
  for each pixel (i, j) in image do
    I(i, j) ← 0
    for each sub-pixel id in (i, j) do
      s ← pixelToWorld(jitter(i, j, id))
      p ← COP
      d ← (s - p).normalize()
      I(i, j) ← I(i, j) + traceRay(scene, p, d, id)
    end for
    I(i, j) ← I(i, j)/numSubPixels
  end for
end function
  
```

A typical choice is numSubPixels = 5*5.


17

DRT pseudocode (cont'd)

Now consider *traceRay()*, modified to handle (only) opaque glossy surfaces:


```

function traceRay(scene, p, d, id):
  (q, N, material) ← intersect(scene, p, d)
  I ← shade(...)
  R ← jitteredReflectDirection(N, -d, material, id)
  I ← I + material.kr * traceRay(scene, q, R, id)
  return I
end function
  
```




18

Pre-sampling glossy reflections (Quasi-Monte Carlo)




19

Soft shadows




Distributing rays over light source area gives:




20

The pinhole camera, revisited

Recall the pinhole camera:




We can equivalently turn this around by following rays from the viewer:




21

The pinhole camera, revisited

Given this flipped version:



how can we simulate a pinhole camera more accurately?




22

Lenses

Pinhole cameras in the real world require small apertures to keep the image in focus.

Lenses focus a bundle of rays to one point => can have larger aperture.



For a "thin" lens, we can approximately calculate where an object point will be in focus using the the Gaussian lens formula:

$$\frac{1}{d_i} + \frac{1}{d_o} = \frac{1}{f}$$



where f is the **focal length** of the lens.

23

Depth of field

Lenses do have some limitations. The most noticeable is the fact that points that are not in the object plane will appear out of focus.


The **depth of field** is a measure of how far from the object plane points can be before appearing "too blurry."




24

Simulating depth of field

Consider how rays flow between the image plane and the in-focus plane:



We can model this as simply placing our image plane at the in-focus location, in *front* of the finite aperture, and then distributing rays over the aperture (instead of the ideal center of projection):



25


Simulating depth of field, cont'd



26

Chaining the ray id's


In general, you can trace rays through a scene and keep track of their id's to handle *all* of these effects:



27

DRT to simulate _____

Distributing rays over time gives:



28

Summary

What to take home from this lecture:

1. The limitations of Whitted ray tracing.
2. How distribution ray tracing works and what effects it can simulate.