Accelerated and anti-aliased ray tracing

Brian Curless
CSE 557
Fall 2015

Reading

Required:

- Shirley 12.3, 13.4.1

Further reading:

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

Speeding it up

Brute force ray tracing is really slow!

Consider rendering a single image with:

- m x m pixels
- n primitives
- average ray path length of d
- ℓ shadow ray per intersection
- 0, 1, or 2 rays cast recursively per intersection

Asymptotic # of intersection tests $\cong O\left(m^2 \cdot n \cdot 2^d \cdot \ell\right)$

For $m=1,000$, $n= 1,000,000$, $\ell =10$, $d= 10$…very expensive!!

In practice, some acceleration technique is almost always used.

We've already looked at reducing d with adaptive (early) ray termination.

Now we look at reducing the effect of n (number of primitives)…

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?

*bbox*

Hierarchical bounding volumes

We can generalize the idea of bounding volume acceleration with **hierarchical bounding volumes**.

Intersect with largest B.V...   ...then intersect with children...

...until you reach the leaf nodes - the primitives.

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

Uniform spatial subdivision

Another approach is **uniform spatial subdivision**.

Uniform subdivion in 2D

Uniform subdivion in 3D
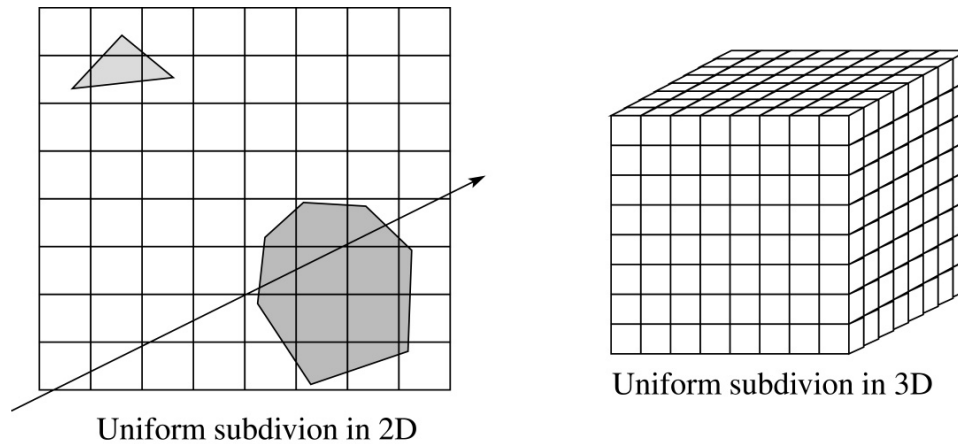
Idea:

- ◆ 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 a10^6 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

Still another approach is **non-uniform spatial subdivision**.

One of these is quadtrees / 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 ($k = 2$) k-d tree ($k = 3$)

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.

[Image credits: Wikipedia.]

8

Aliasing

Ray tracing is a form of sampling and can suffer from annoying visual artifacts...

Consider a continuous function $f(x)$. Now sample it at intervals Δ to give $f[i] = \text{quantize}[f(i\Delta)]$.

**Q**: How well does $f[i]$ approximate $f(x)$?

Consider sampling a sinusoid:

In this case, the sinusoid is reasonably well approximated by the samples.

Aliasing (con't)

Now consider sampling a higher frequency sinusoid

We get the exact same samples, so we seem to be approximating the first lower frequency sinusoid again.

We say that, after sampling, the higher frequency sinusoid has taken on a new "alias", i.e., changed its identity to be a lower frequency sinusoid.

Aliasing in rendering

One of the most common rendering artifacts is the "jaggies". Consider rendering a white polygon against a black background:

We would instead like to get a smoother transition:

Anti-aliasing

Q: How do we avoid aliasing artifacts?

1. Sampling: *more samples ... better display*

2. Pre-filtering: *analytic integration*

3. Combination: *Super-sampling + averaging down*

Example - polygon:

Polygon anti-aliasing

Without antialiasing

With antialiasing

Magnification

Antialiasing in a ray tracer

We would like to compute the average intensity in the neighborhood of each pixel.

When casting one ray per pixel, we are likely to have aliasing artifacts.

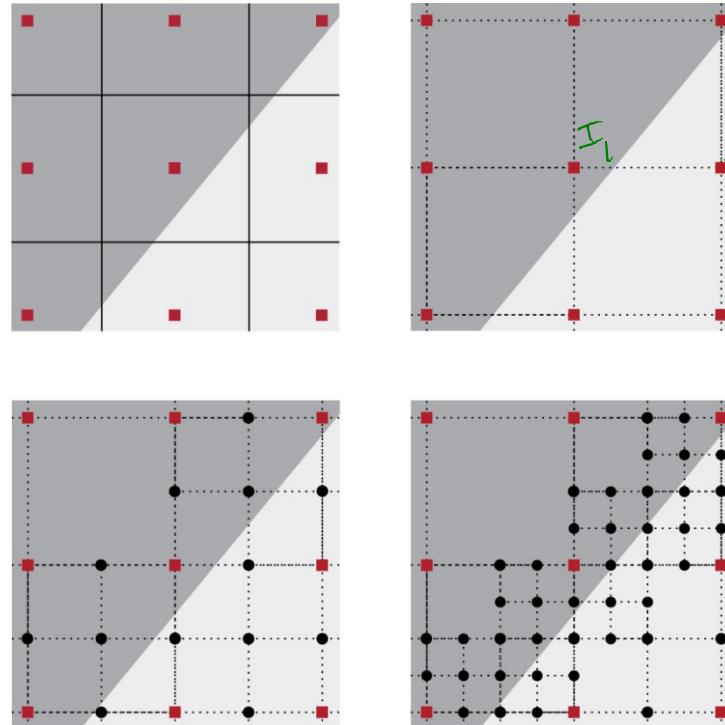To improve matters, we can cast more than one ray per pixel and average the result.

A.k.a., **super-sampling and averaging down**.

Antialiasing by adaptive sampling

Casting many rays per pixel can be costly. We need another acceleration trick!

If there are no rapid changes in intensity at the pixel, maybe only a few samples are needed.

Solution: **adaptive sampling**.

$$\| I_1 - I_2 \|_2 > thresh.$$

$$\| R_1 - R_2 \| > thresh_{red}$$

$$\| G_1 - G_2 \| > thresh_{green}$$

$$\vdots$$

**Q**: When do we decide to cast more rays in a particular area?

15

Temporal aliasing

Suppose we are rendering a "clock" with a fast turning hand:

What happens if we sample too infrequently? (This is sometimes called the "wagon wheel" effect.)

Another more common scenario is something moving quickly across the frame, e.g., a fast-moving particle:

Trajectory | Frames

How might we address these temporal aliasing effects?

55 AD

Motion blur

By averaging consecutive frames together within an exposure interval, we get **motion blur**: