

Distributed Ray Tracing

CS319

Advanced Topics in Computer Graphics

John C. Hart

Why Ray Tracing Looks Fake

- Jagged edges
- Hard shadows
- Everything in focus
- Objects completely still
- Surfaces perfectly shiny
- Glass perfectly clear

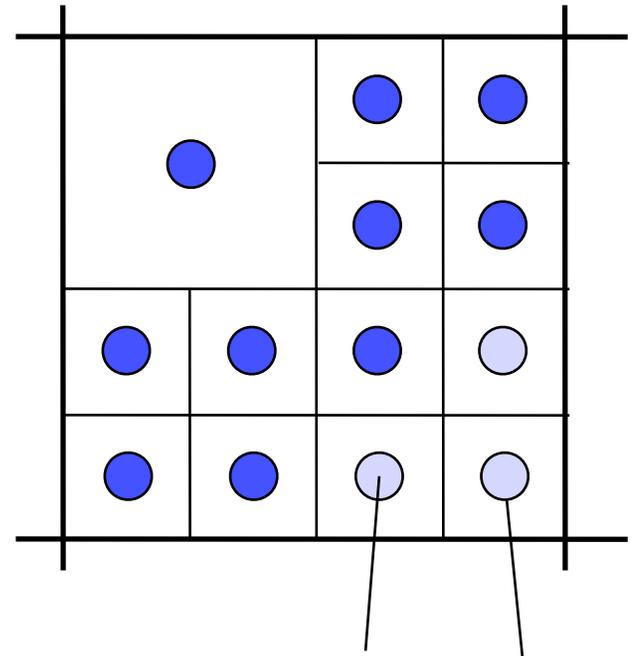


Distributed Ray Tracing

- Rob Cook, SIGGRAPH 84
- Replace single ray with distribution of rays
- Not just fat ray through pixel, but fat rays everywhere
- Multiple eye rays
 - supersampling
 - motion blur
 - depth of field
- Multiple shadow rays
 - soft shadows
- Multiple reflection rays
 - glossy surfaces
- Multiple refraction rays
 - diffuse glass

Supersampling

- Cast multiple rays from eye through different parts of same pixel
- Image function $f(x,y)$
 - Color that ray through (x,y) returns
- Sampling function $s(x,y)$
 - Controls where the samples occur within the pixel
 - One if sample at (x,y) , else zero
- Reconstruction filter $r(x,y)$
 - Computes the weighted average of resulting colors into a single color
 - Also weights according to area



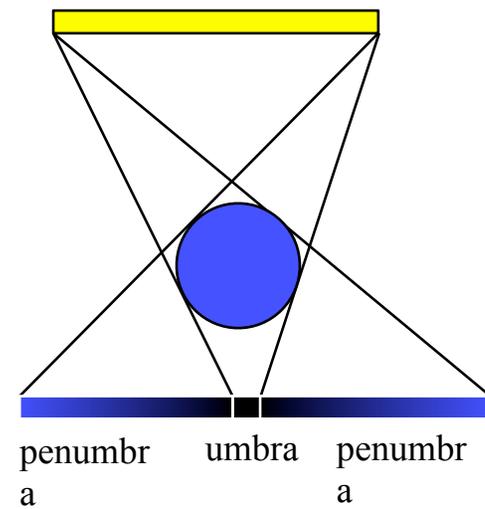
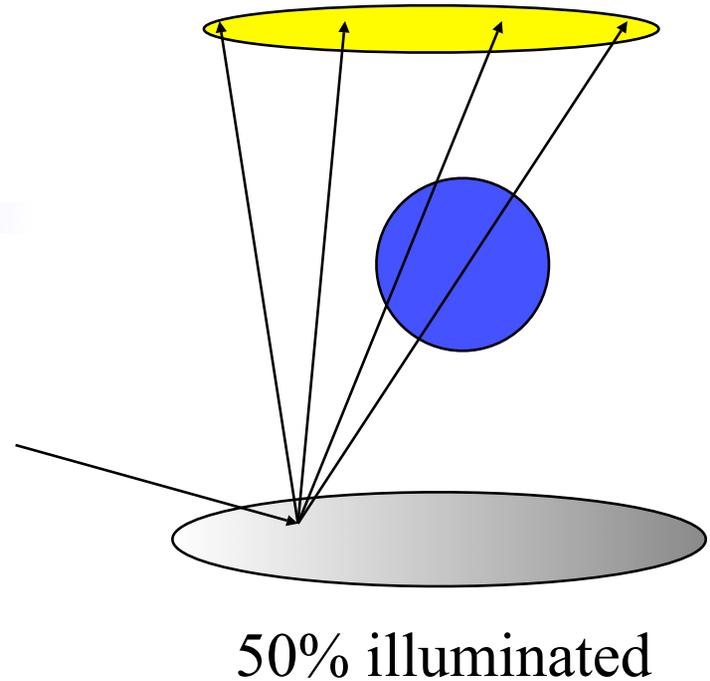
We must estimate the convolution of image with filter in space and time - by sampling

Soft Shadows

- Point light sources unrealistic
- Use area light source
- Cast shadow rays from surface to different locations on light
- Hits/rays = % illuminated
- Jitter locations to remove aliasing artifacts

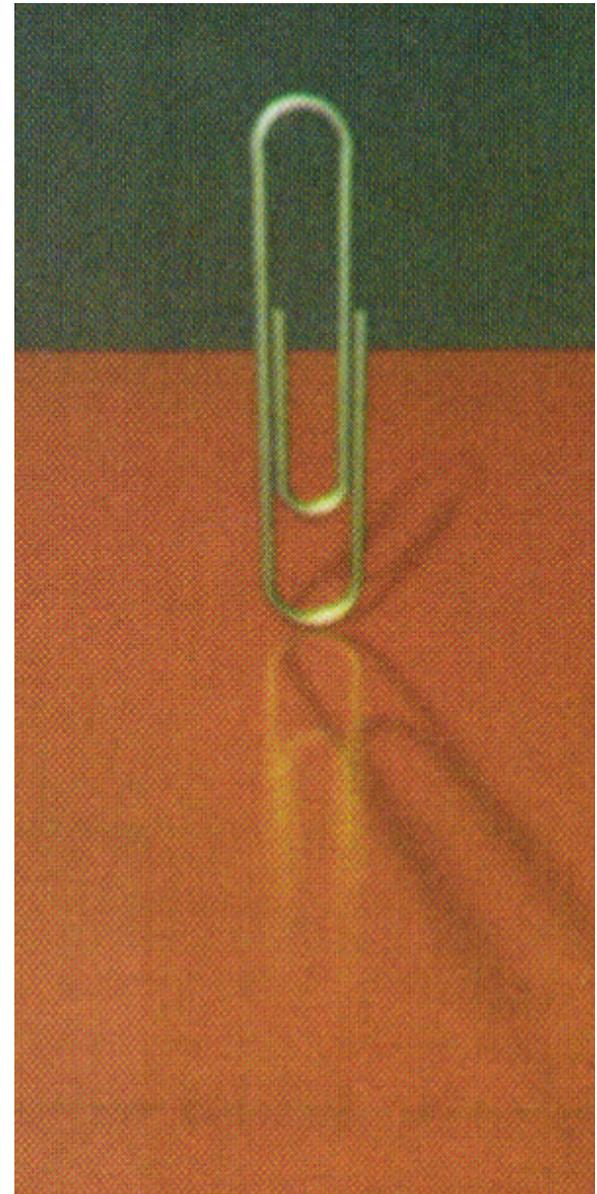
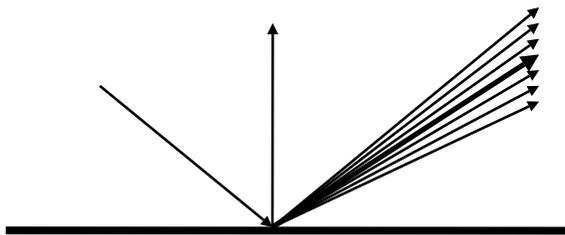
Integral is spatial average

Aliasing can occur and be very nasty, particularly with uniform grid



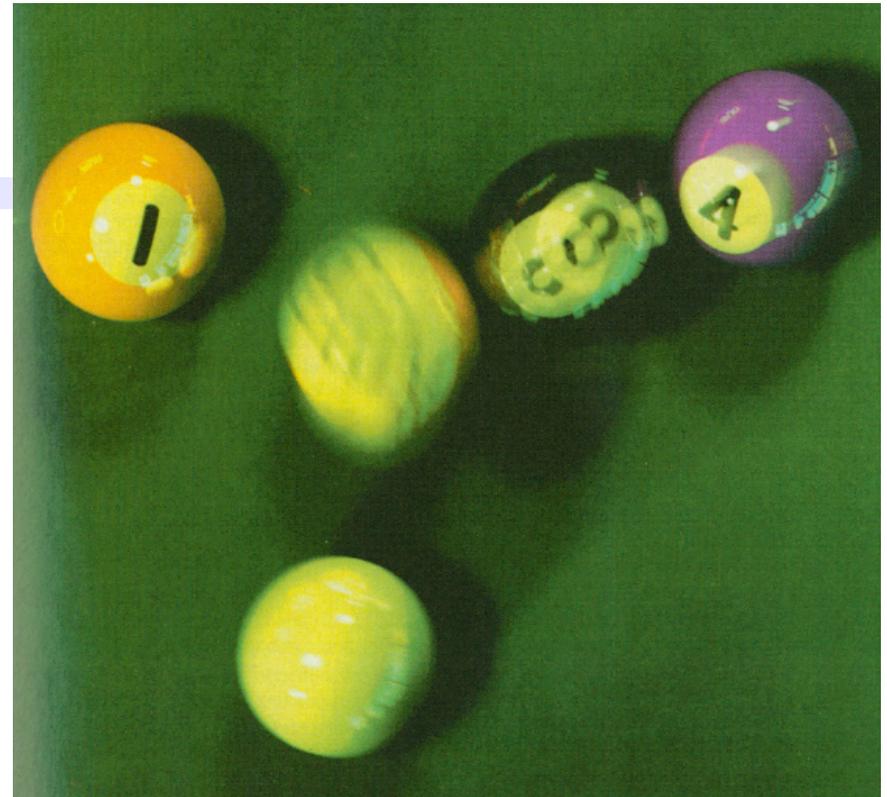
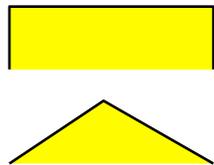
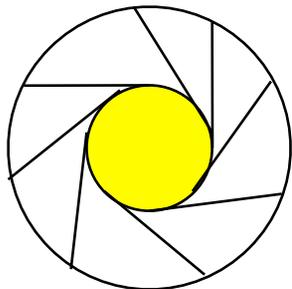
Glossy Surfaces

- Surface microfacets perturb reflection ray directions
- Nearby objects reflect more clearly because distribution still narrow
- Farther objects reflect more blurry because distribution has spread
- Helmholtz reciprocity allows us to sample lobe of BRDF



Motion Blur

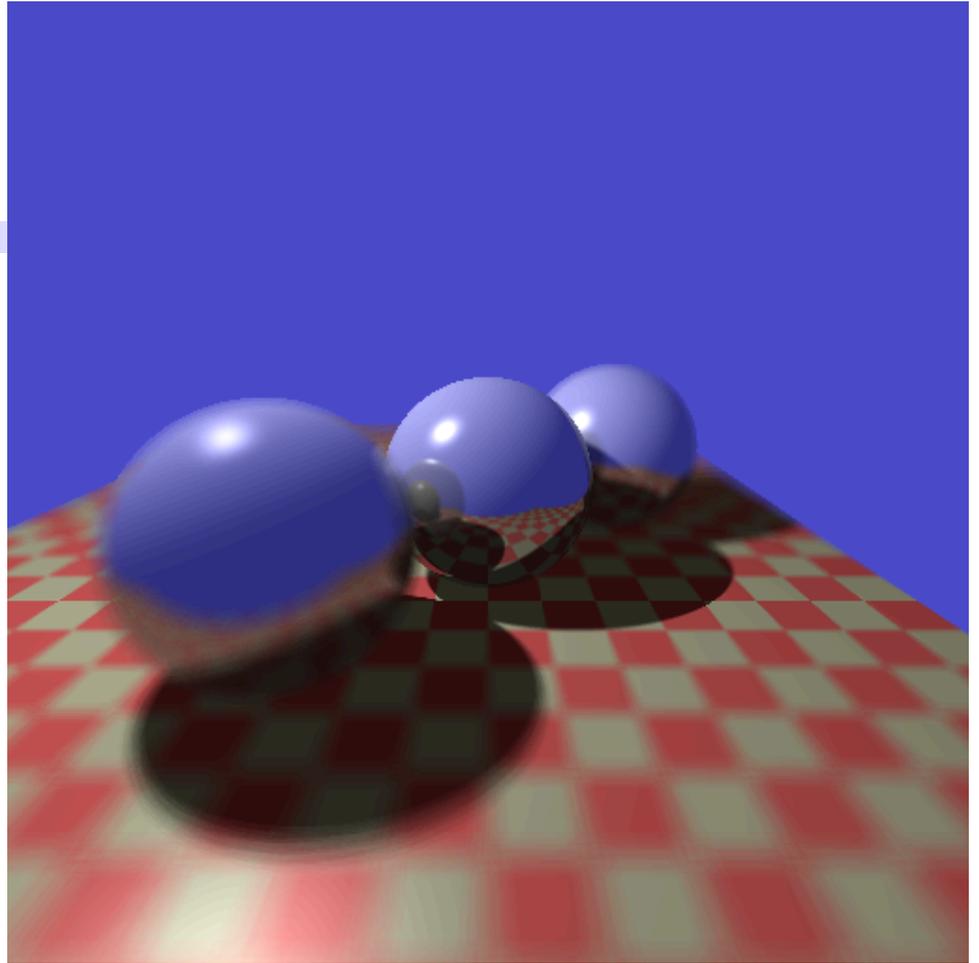
- Cast multiple rays from eye through same point in each pixel
- Each of these rays intersects the scene at a different time
- Reconstruction filter controls shutter speed, length
 - Box filter – fast shutter
 - Triangle filter – slow shutter



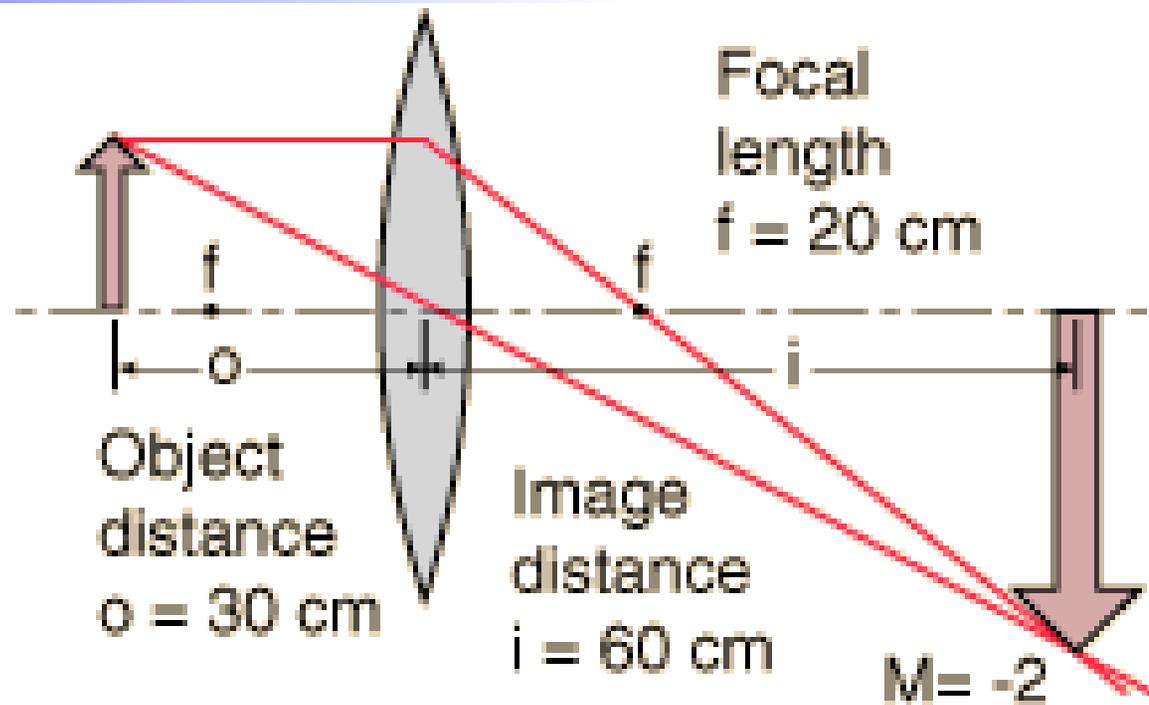
Filter is now temporal

Depth of Field

- Better simulation of camera model
 - f-stop
 - focus



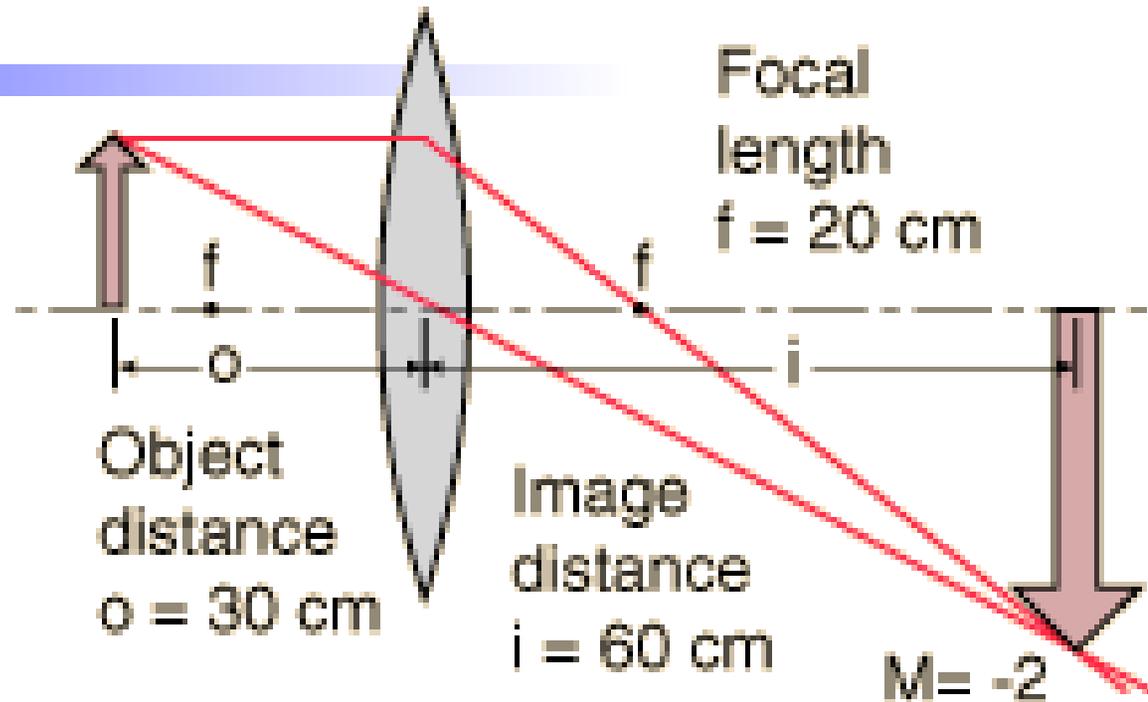
Thin lens



Key property of thin lens:

The family of rays leaving a point on the image plane (focal plane) to the lens arrives at a point on the focal plane (image plane)

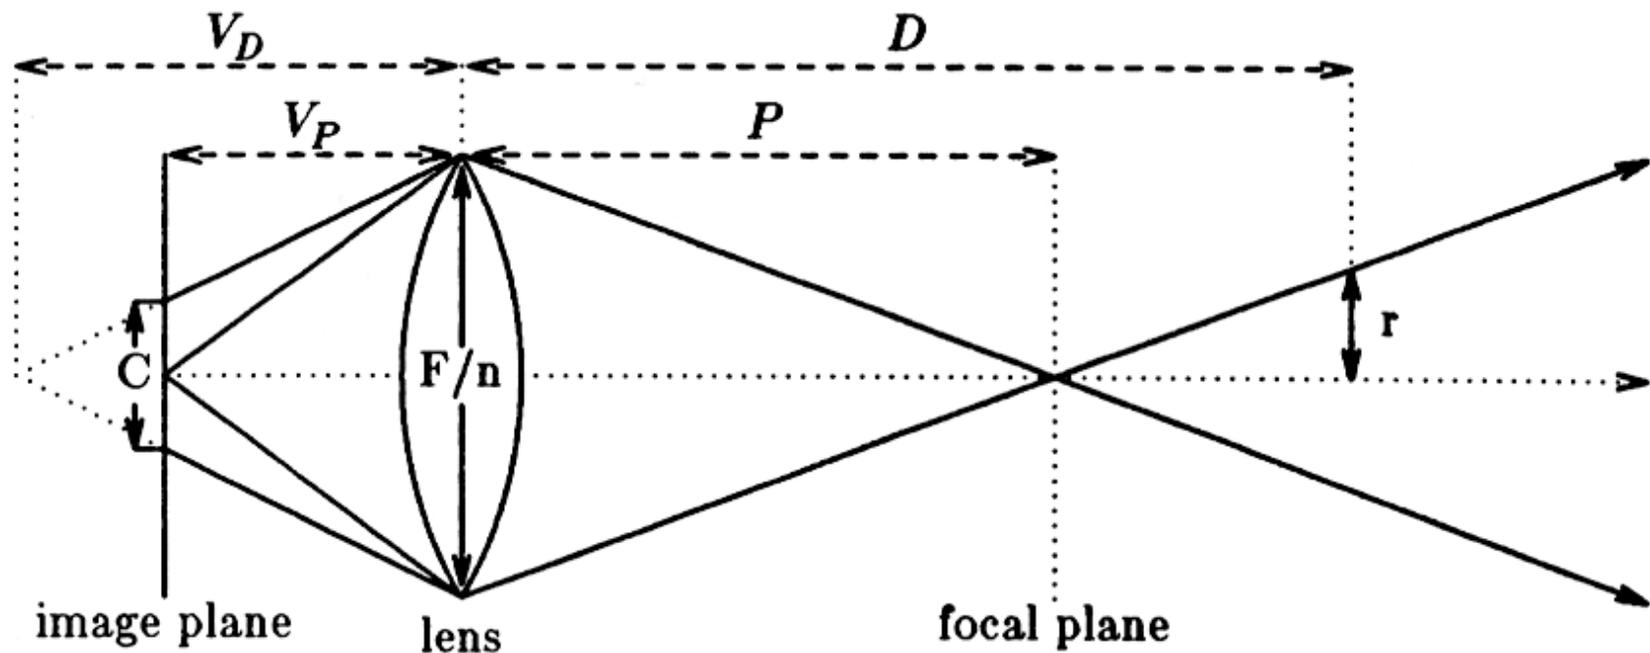
Thin Lens



Thin lens equation:

$$\frac{1}{o} + \frac{1}{i} = \frac{1}{f}$$
$$\frac{1}{\text{object distance}} + \frac{1}{\text{image distance}} = \frac{1}{\text{focal length}}$$

Depth of field



Points that lie off the focal plane image to “circles” (polygons in real lens systems)
“acceptable” circle=circle of confusion

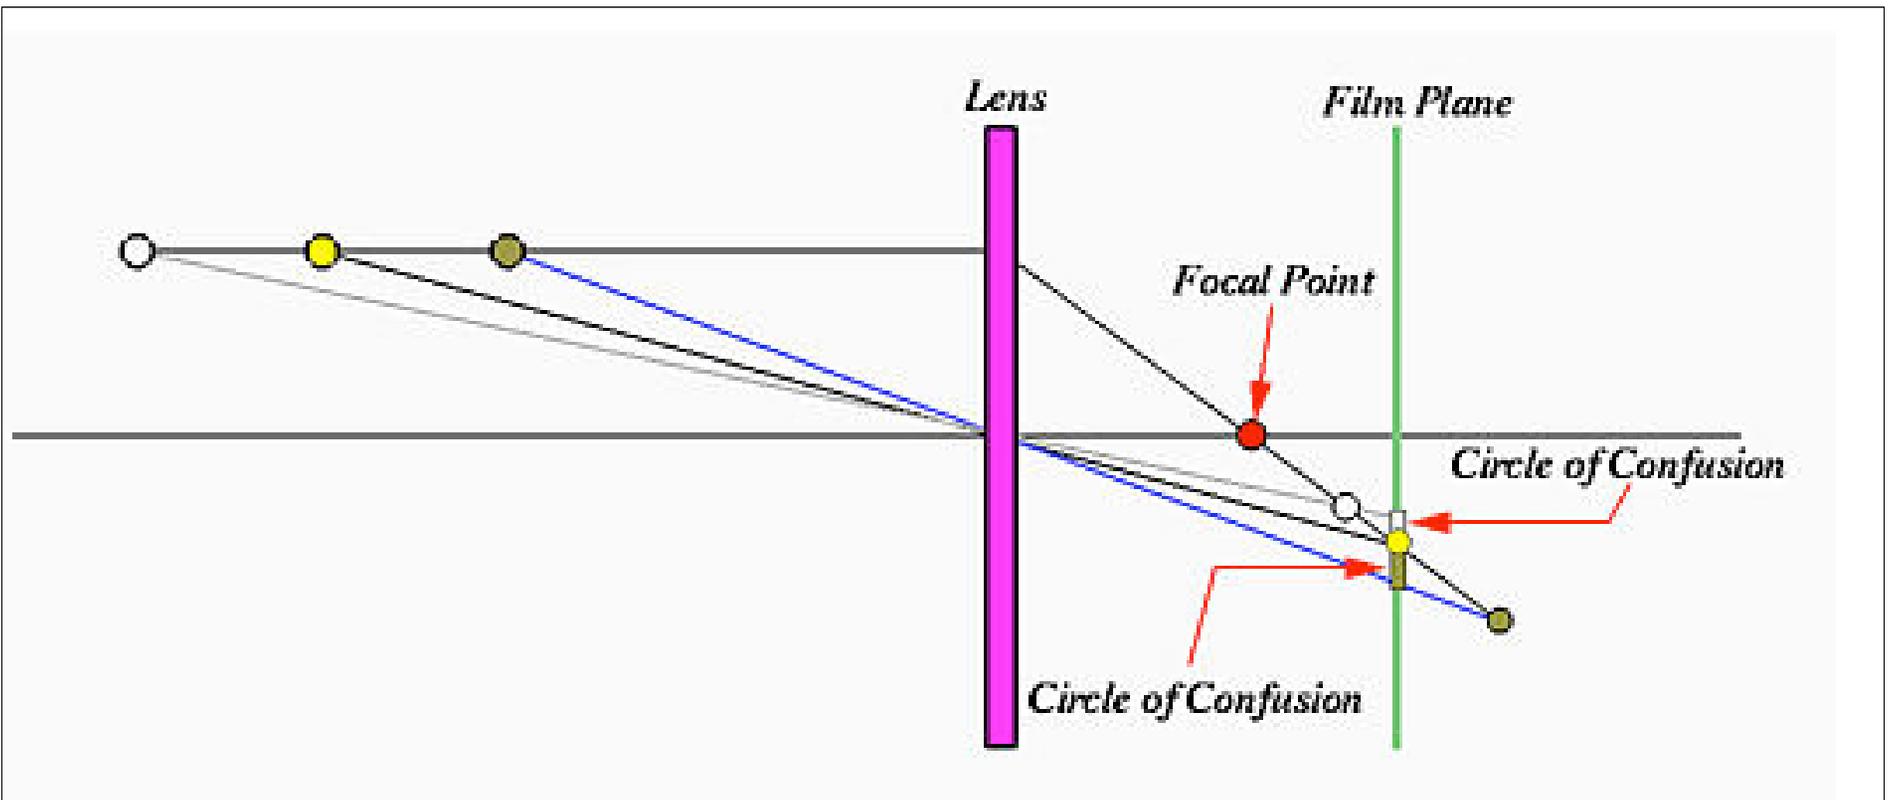
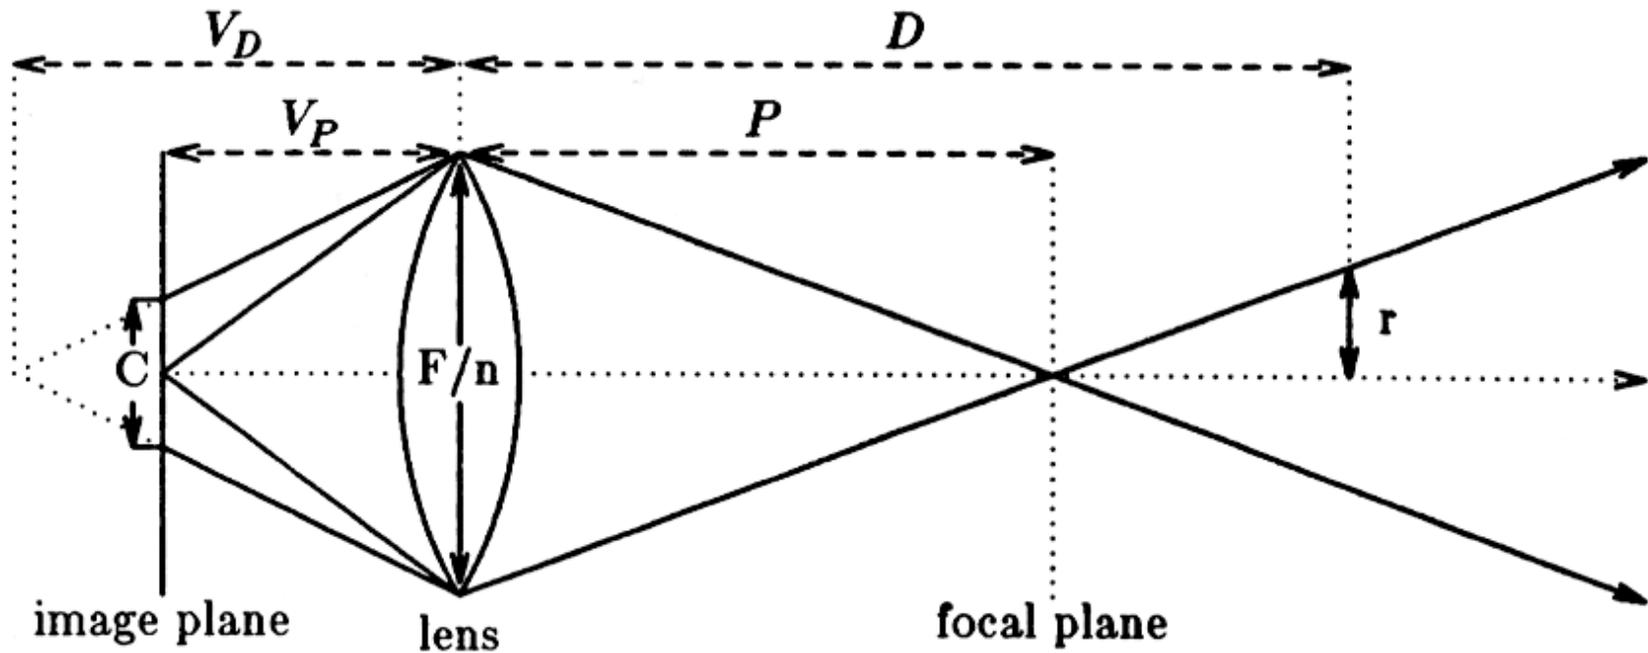


figure from <http://www.cs.mtu.edu/~shene/DigiCam/User-Guide/950/depth-of-field.html>

A point that images larger than the circle of confusion is out of focus. By choosing the coc, we choose the depth of field. Coc radius is typically chosen to be 0.01 in on an 8x10 in print.

Depth of field is affected by aperture

Camera Model



F – focal length

n – aperture number

C – circle of confusion

$$V_P = FP/(P-F)$$

$$V_D = FD/(D-F)$$

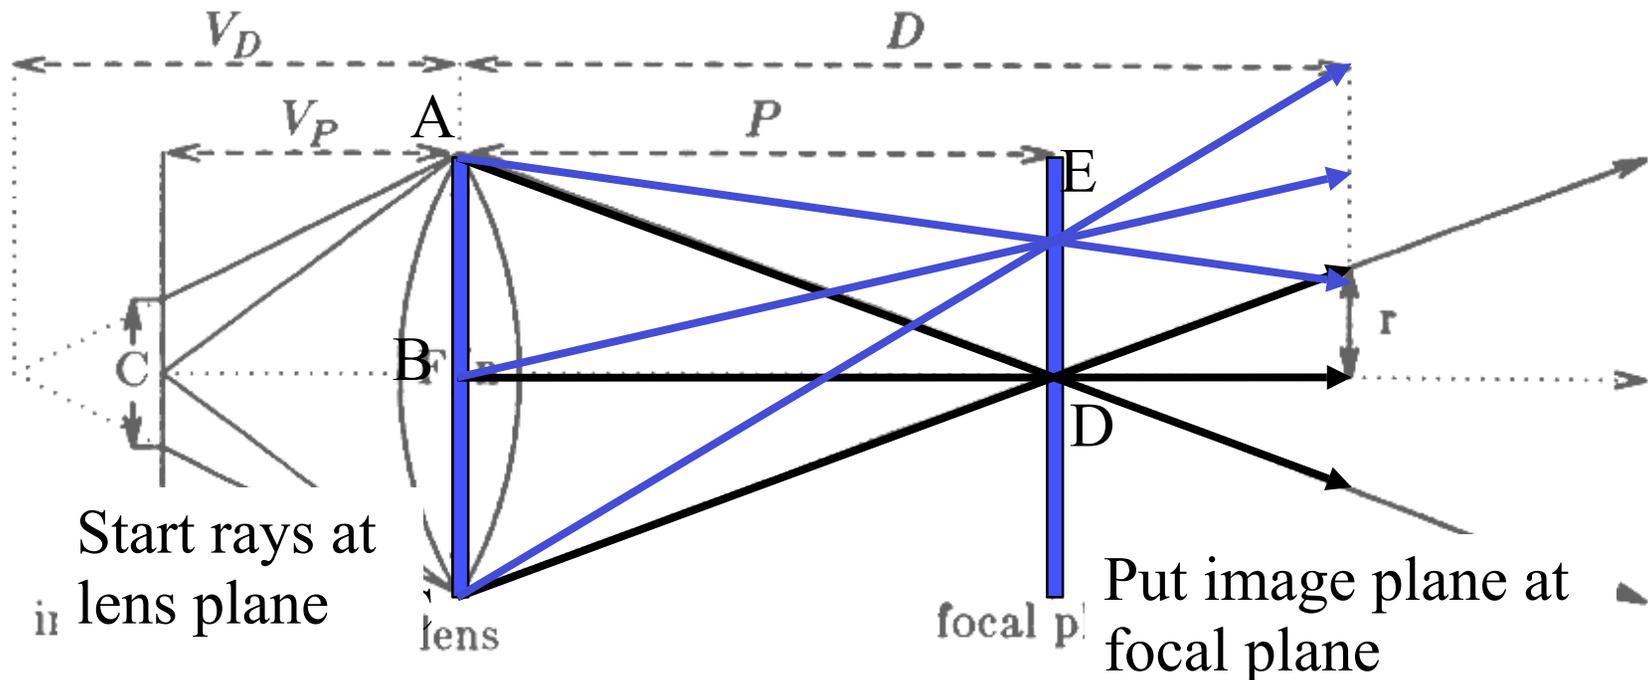
$$C = (|V_D - V_P|/V_D) (F/n)$$

$$r = \frac{1}{2} (F/n) (D-P)/P$$

$$R = (-V_P/D) r$$

$$R = \frac{1}{2} C$$

Implementation



Standard ray tracing:

Pixel D uses ray BD

Pixel E uses ray BE

All rays emanate from B

Distributed ray tracing:

Pixel D uses rays AD, BD, CD

Pixel E uses rays AE, BE, CE

Rays emanate from lens plane

Advanced Camera Models

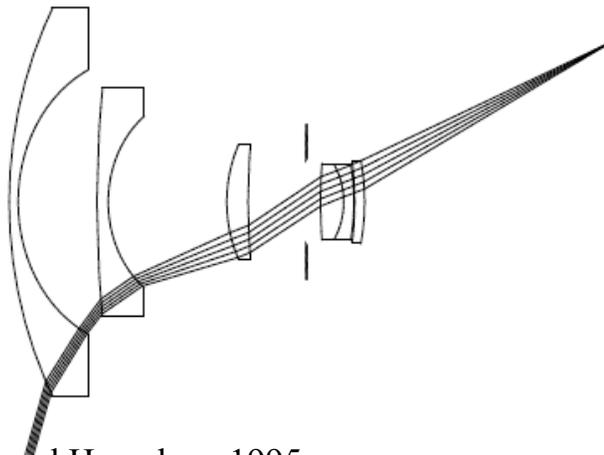
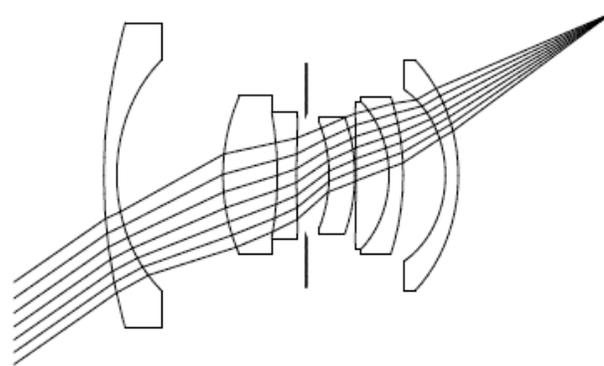


figure from Kolb, Mitchell and Hanrahan, 1995

Advanced Camera Models

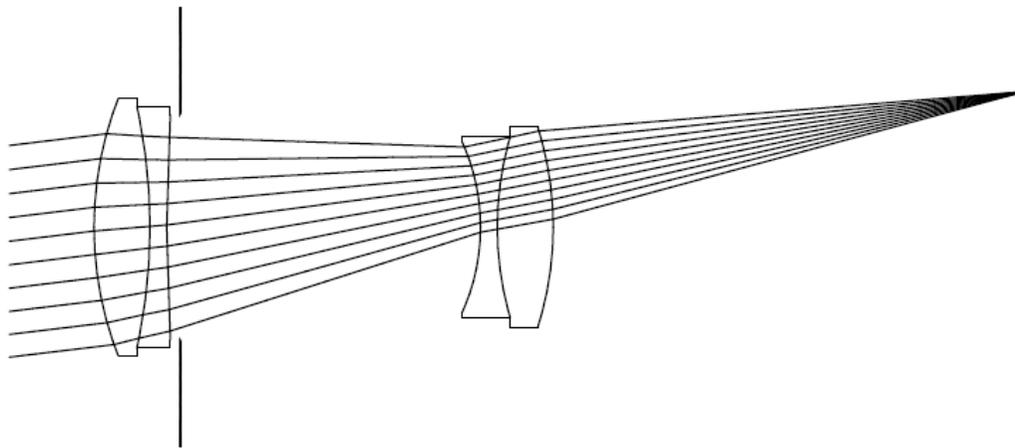
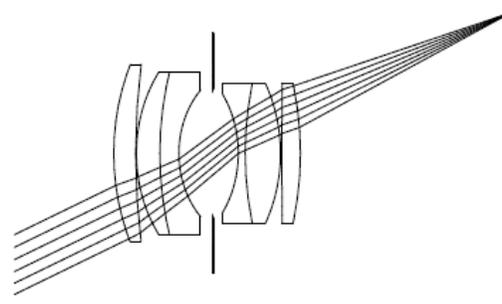


figure from Kolb, Mitchell and Hanrahan, 1995

Distributed Ray Sampling

- Mitchell SIGGRAPH '91
- Supersampled antialiasing
 - Jittering removed jaggies, Moires
 - Ordered jitter, Poisson balanced distribution
- Decorrelation inhibits aliasing
- Need to decorrelate samples
 - across all parameter axes
 - between all parameter axes

