EECS 487: Interactive Computer Graphics

Lecture 29:
- Distributed Ray Tracing

Ray Tracing

Introduction and context
- ray casting

Recursive ray tracing
- shadows
- reflection
- refraction

Ray tracing implementation

Distributed Ray Tracing
- anti-aliasing
- motion blur
- depth-of-field
- glossy surface
- translucency
- soft-shadows

Why Does Ray Tracing Look Obviously Computer Generated?

Crisp images . . . too "perfect":
- surfaces are perfectly shiny
- glass is perfectly clear
- everything in perfect focus
- every object is completely still
- even the shadows have perfect silhouettes
- but . . . up close, edges are jagged

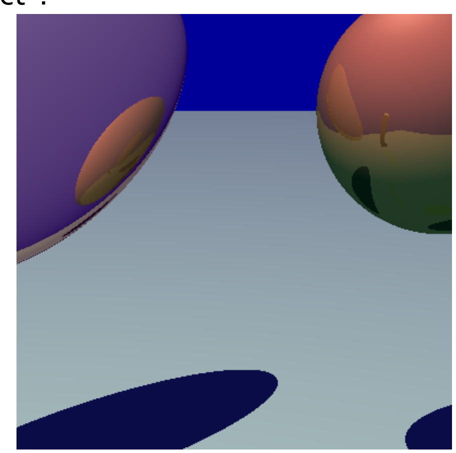

"Distributed" Ray Tracing

Replace each single ray with a distribution of multiple rays ⇒ a.k.a. *distribution* ray tracing
- average results together

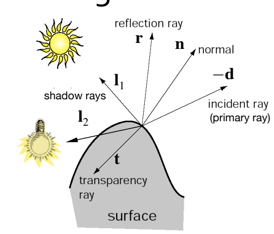

Multiple rays everywhere:
- multiple primary rays through a pixel:
  - supersampling: distribute rays spatially
  - motion blur: distribute rays temporally
  - depth of field: distribute rays through a lens
- multiple shadow rays to sample an area light
  - soft shadows
- multiple reflection rays
  - glossy surfaces (blurry reflection, rough specular)
- multiple refraction rays
  - translucency

(multiple secondary rays)

Supersampling

Instead of point sampling the color of a pixel with a ray, we cast multiple rays from eye (primary rays) through different parts of one pixel and average down the results

- for example, cast $n \times n$ sub-pixel rays, and average the results together:

$$\mathbf{c}_{pixel} = \frac{1}{n^2} \sum_{i=0}^{n-1} \sum_{j=0}^{n-1} \mathbf{c}_{subpixel(i,j)}$$

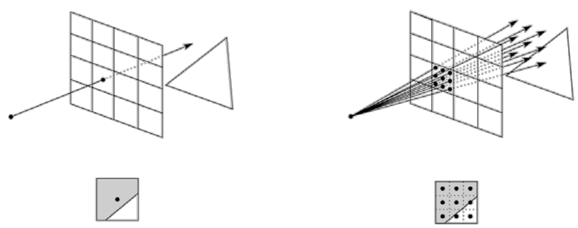

Curless08

Non-fixed Sampling Patterns

Super-sampling doesn't eliminate aliasing, it simply pushes it to higher frequencies

- supersampling captures more high frequencies, but frequencies above the supersampling rate are still aliased
- fundamentally, problem is that the signal is not bandlimited ⇒ aliasing happens

Other than regular, fixed sampling pattern, sampling can also be stochastic (a.k.a., random, probabilistic, or Monte Carlo)

Durand08

Stochastic Supersampling

Samples taken at non-uniformly spaced random offsets

Replaces low-frequency aliasing pattern by noise, which is less objectionable to humans

However, with random sampling, we could get unlucky, e.g., all samples in one corner

- over 80% of the samples are black while the pixel should be light grey

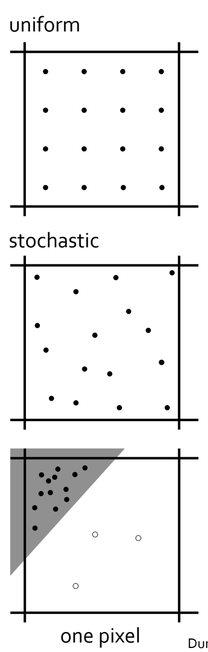

Durand08

Stratified Sampling

To prevent clustering of the random samples, divide domain (pixel) into non-overlapping regions (sub-pixels) called strata

Take one random sample per stratum

Jittered sampling is stratified sampling with per-stratum sample taken at an offset from the center of each stratum:

- one sample per stratum
- randomly perturb the sample location
- size of perturbation vector limited by the subpixel distance
- patented by Pixar!

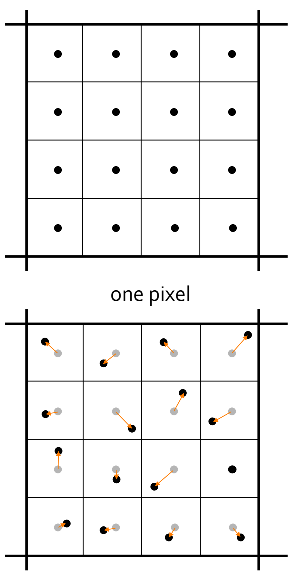

Durand

Temporal Aliasing: Motion Blur

Aliasing happens in time as well as in space

- the sampling rate is the frame rate: 30Hz (NTSC), 25Hz (PAL), 24Hz (film)
- if we point-sample time, objects have a jerky, strobbed look, e.g., sampling at ¼ rotation

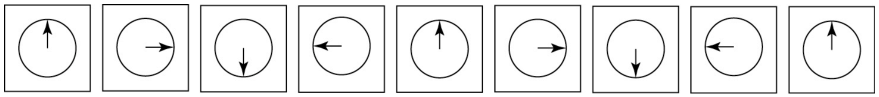

  - or move backward, e.g., sampled at ¾ rotation

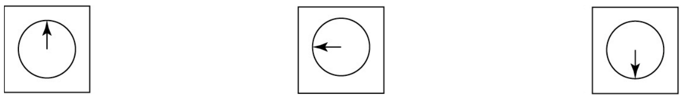

  - see http://www.michaelbach.de/ot/mot_wagonWheel/index.html
- fast moving objects move large distances between frames
  - be careful when doing collision detection!

Curless, Hanrahan

Temporal Aliasing: Motion Blur

Film automatically does temporal anti-aliasing

- photographic film integrates over the exposure time
- this shows up as motion blur in the photographs

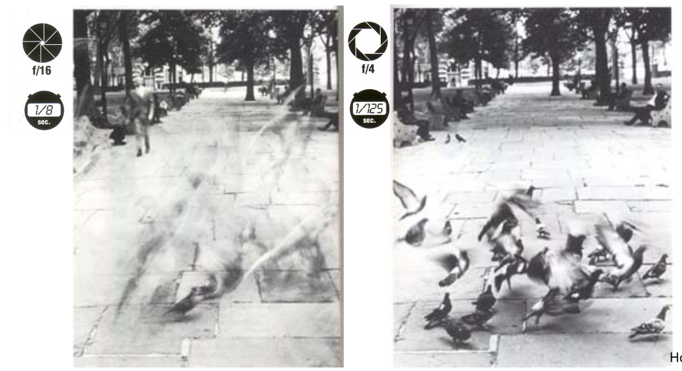

Hodginso6, London et al.

Temporal Aliasing: Motion Blur

To avoid temporal aliasing we need to average over time also

Sample objects temporally

- cast multiple rays from eye through the same point in each pixel
- each of these rays intersects the scene at a different time: $\mathbf{r}(\mathbf{e}, \mathbf{d}, t)$
- average out results

The result is still-frame motion blur and smooth animation

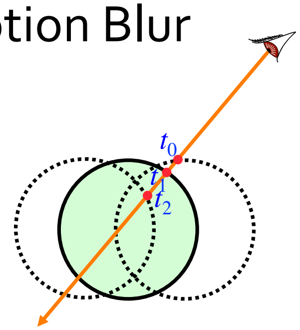

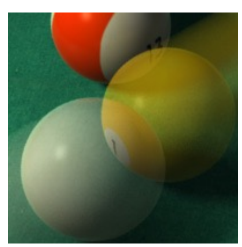

Hodginso6,Durando8,Merrello8

Depth of Field and Aperture

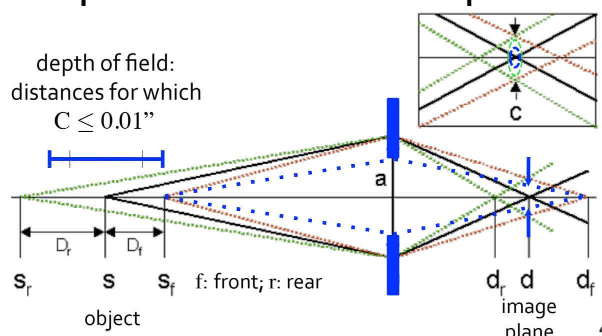

Thin lens equation $1/s + 1/d = 1/f$ where f is lens focal length

a = aperture
C = circle of confusion

Object is considered in focus if on an 8×10 print viewed at a distance of 10”, diameter of $C \leq 0.01$” (1930’s standard!)

Later human study shows that C should be ≤ 0.003”

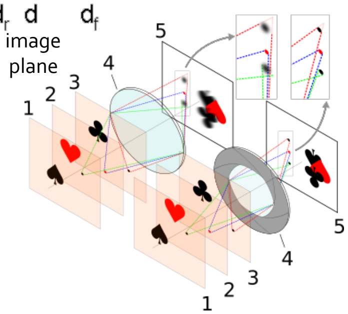

normankoren.com/Tutorials/MTF6.html

Depth of Field

Real cameras have lenses with focal lengths
- only one plane is truly in focus
- points away from the focus project as "circle of confusion"
- the further away from the focus the larger the circle

The range of distances that appear in focus is the depth of field
- smaller apertures (larger f-numbers) result in greater depth of field

Depth of field can be simulated by distributing primary rays through different parts of a lens assembly

Hart08

DoF: Implementation

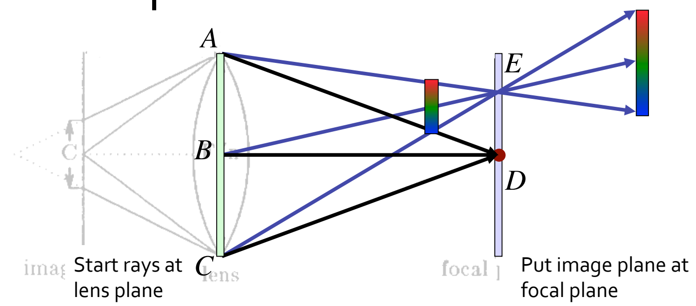

Standard ray tracing:
- all rays emanate from B
- pixel D uses ray BD
- pixel E, for object behind (or in front of) the focal plane, uses ray BE

Distributed ray tracing:
- rays emanate from lens plane
- pixel D uses rays AD, BD, CD
- pixel E averages rays AE, BE, CE
- to simulate more accurately, first refract primary rays through lens

Hart08

DoF: Implementation

Or simply select eye positions randomly from a square region

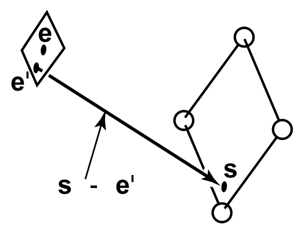

/Shirley02

Glossy Surfaces

Ray tracing simulates perfect specular reflection, true only for perfect mirrors and chrome surfaces

Most surfaces are imperfect specular reflectors:
- surface microfacets perturb direction of reflected rays
- reflect rays in a cone around perfect reflection direction
- Phong specular lighting tries to fake this with the m_{shi} exponent

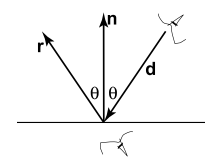

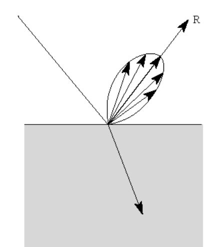

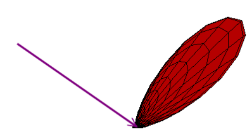

RTR,TP3

Glossy Reflections: Implementation

For each ray-object intersection
- instead of shooting one ray in the perfect specular reflection (mirror) direction,
- stochastically sample multiple rays within the cone about the specular angle
- strength of reflection drops off rapidly away from the specular angle,
- probability of sampling that direction should fall off similarly

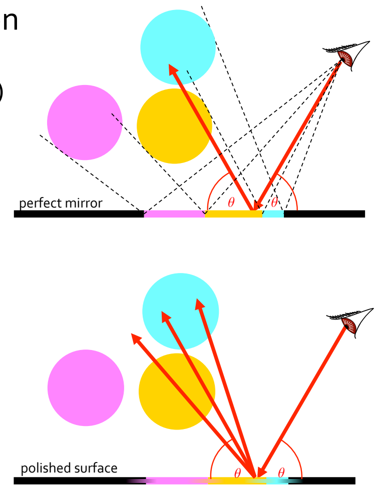

Hart08,Yu08,Durand08

Glossy Reflections

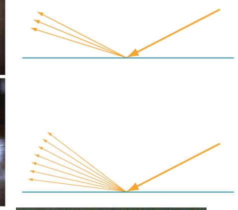

Instead of mirror images:
- highlights can be soft
- blurred reflections of objects

Nearby objects reflect more clearly because distribution still narrow

Farther objects reflect more blurry because distribution has spread

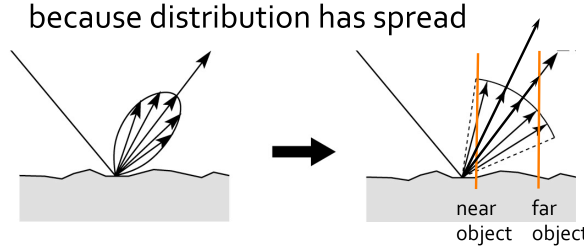

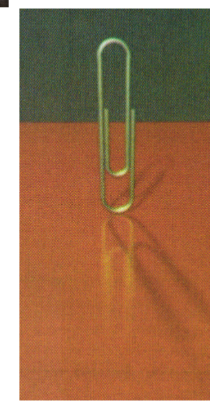

RTR, Curless09,Hart08

Translucency

Similar, but for refraction
- instead of distributing rays around the reflection ray, distribute them around the refracted ray

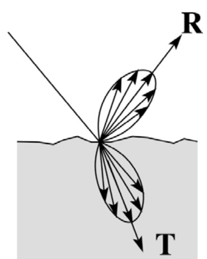

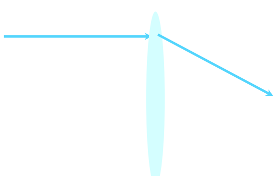

Transparent

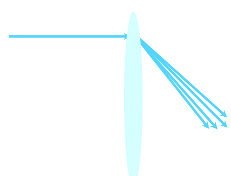

Translucent

Merrell08,Curless08

Shadows

Darkness caused when part or all of the illumination from a light source is blocked by an occluder (shadow caster)

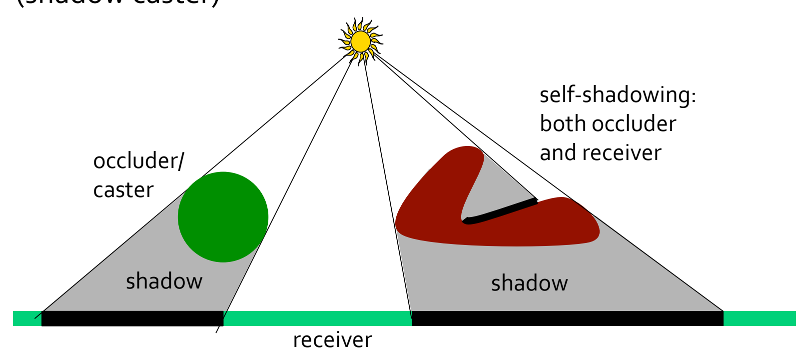

Akenine-

Hard and Soft Shadows

Point light sources give unrealistic hard shadows

Light sources that extend over an area (area light sources) cast soft-edged shadows

- some points see all the light: fully illuminated
- some points see none of the light source: the umbra
- some points see part of the light source: the penumbra

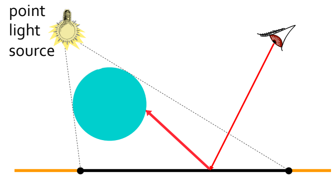

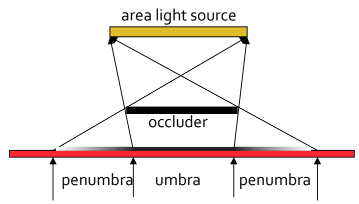

Durand,Chenney

Ray Tracing Area Light Source to Create Soft Shadows

Cast multiple shadow rays from surface, distributed across the surface of the light: each ray to a different point on the light source

- inside the umbra, no shadow rays get through to light
- inside the penumbra, some shadow rays get through and some parts of the are light is visible
- outside the penumbra, all shadow rays get through

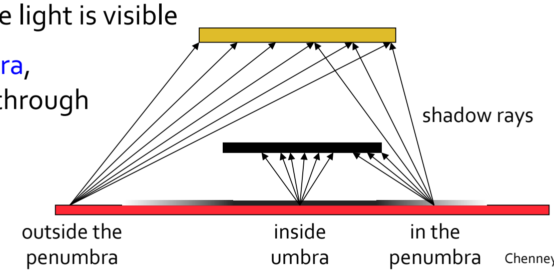

Chenney

Ray Tracing Area Light Source to Create Soft Shadows

At each point, sum the contributions of shadow rays from that point to find the strength of shadow: hits/rays = % illuminated

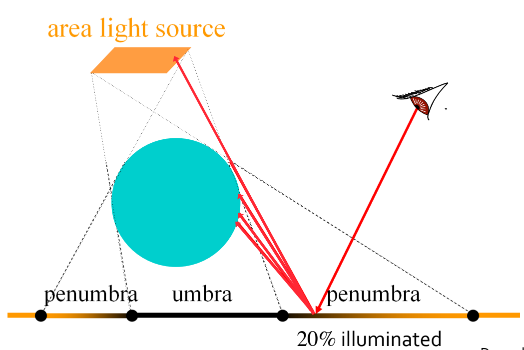

Durand

Sampling Area Light

Anti-aliasing:

- break a pixel into a grid of sub-pixels and distribute rays over the sub-pixels

Soft-shadows:

- break an area light into a grid of $N = n \times n$ point lights, each with $1/N$-th the intensity of the base light
- sample the light, not the pixel
- each primary ray generates multiple shadow rays per light

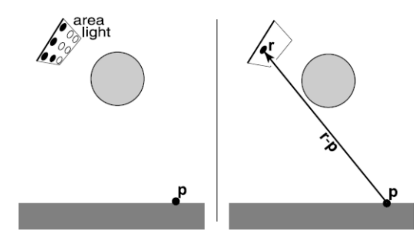

Shirley02

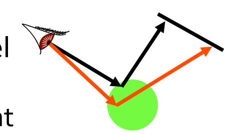

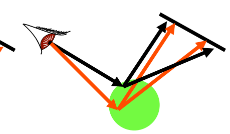

Classic Recursive Ray Tracing

```
render() {
  for each pixel {
    generate primary/viewing ray
    pixel_color = trace(primary_ray)
  }
}
```

```
trace(ray) {
  (point, normal) = ray.intersect(scene);
  return shade(point, normal)
}

shade(pt, normal) {
  color = 0
  for each light source {
    if(!intersect(shadow_ray, scene))
      color += direct_illumination
  }
  if(specular)
    color += F * trace(reflected_ray)
  // also add transmitted color…
  return color
}
```

Distributed Area-Light Sampling

```
render() {
  for each pixel {
    generate primary/viewing ray
    pixel_color = trace(primary_ray)
  }
}
```

```
trace(ray) {
  (point, normal) = ray.intersect(scene);
  return shade(point, normal)
}

shade(pt, normal) {
  color = 0
  for each light source {
    generate N random shadow rays
    foreach (shadow ray) {
      if(!intersect(shadow_ray, scene))
            color += direct_illumination*1/N
    }
  }
  if(specular)
    color += F * trace(reflected_ray)
  // also add transmitted color…
  return color
}
```

Sample Distribution?

Uniform distribution gives rise to sharp transitions/patterns inside penumbra

4 eye rays per pixel in both cases:

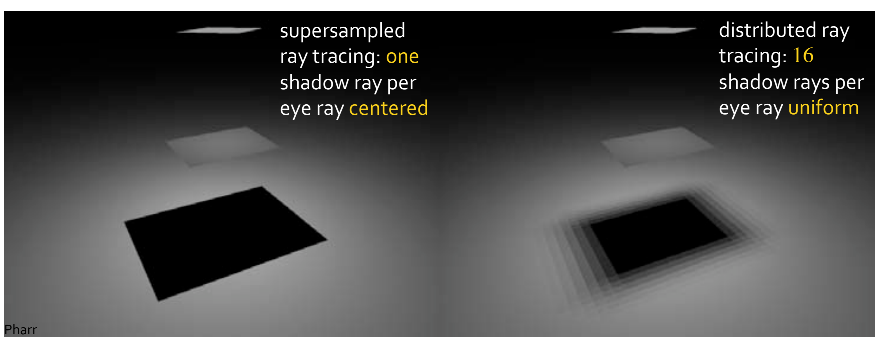

Stochastic Sampling of Area Light

Area light represented as a rectangle in 3D, each ray-object intersection samples the area-light at random: $\mathbf{r} = \mathbf{c} + \xi_1 \mathbf{a} + \xi_2 \mathbf{b}$, where ξ_1 and ξ_2 are random variables

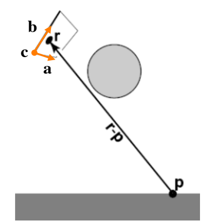

Stratified sampling of the area light with samples spaced uniformly plus a small perturbation $\{\ \mathbf{r}(i,j), 0 \leq i,j \leq n{-}1\ \}$

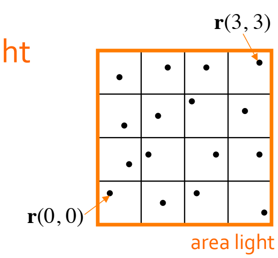

Shirley

Sampling Rate

As with stochastic super-sampling for anti-aliasing, light sampling rate must be high, otherwise high-frequency noise becomes visible