Lecture 9: Monte Carlo Rendering

Chapters 4 and 5 in Advanced GI

Fall 2004

Kavita Bala

Computer Science

Cornell University

Homework

- HW 1out, due Oct 5
- Assignments done separately
  - Might revisit this policy for later assignments

Stochastic Ray Tracing

- Parameters?
  - # starting rays per pixel
  - # random rays for each surface point (branching factor)
- Path Tracing
  - Branching factor == 1

Algorithm so far ...

- Shoot # viewing rays through each pixel
- Shoot # indirect rays, sampled over hemisphere
  - Path tracing shoots only 1 indirect ray
- Terminate recursion using Russian Roulette

Performance/Error

- Want better quality with smaller number of samples
  - Fewer samples/better performance
  - Stratified sampling
  - Quasi Monte Carlo: well-distributed samples
- Faster convergence
  - Importance sampling: next-event estimation

Stratified Sampling

- Samples could be arbitrarily close
- Split integral in subparts

$$I = \int_{X_1} f(x)dx + \ldots + \int_{X_N} f(x)dx$$

- Estimator

$$\bar{I}_{strat} = \frac{1}{N}\sum_{i=1}^{N}\frac{f(\bar{x}_i)}{p(\bar{x}_i)}$$

- Variance: $\sigma_{strat} \le \sigma_{\sec}$

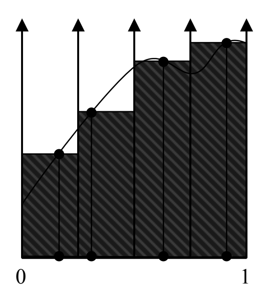

Higher Dimensions

- Stratified grid sampling:

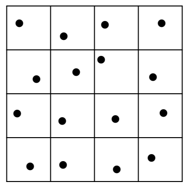

$\rightarrow N^d$ samples

- N-rooks sampling:

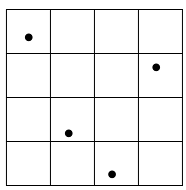

$\rightarrow N$ samples

Quasi Monte Carlo

- Eliminates randomness to find well-distributed samples
- Samples are determinisitic but “appear” random

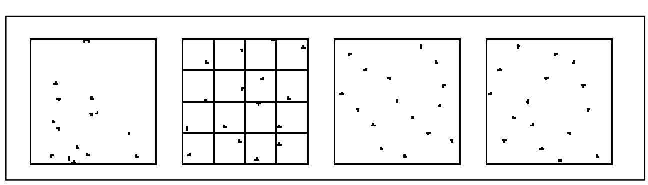

Performance/Error

- Want better quality with smaller number of samples
  – Fewer samples/better performance
  – Stratified sampling
  – Quasi Monte Carlo: well-distributed samples

- Faster convergence
  – Importance sampling: next-event estimation

Next Event Estimation

$$L(x \rightarrow \Theta) = L_e + L_{direct} + L_{indirect}$$

$$= L_e + \int_{\Omega_x}$$

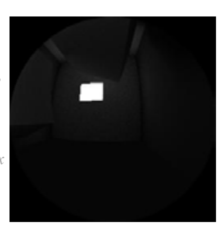

$$\cdot f_r \cdot \cos + \int_{\Omega_x}$$

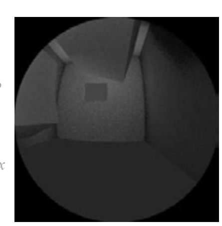

$$\cdot f_r \cdot \cos$$

- So … sample direct and indirect with separate MC integration

How to sample direct illumination

- Sampling a single light source
- Sampling for many lights

Direct Illumination

$$L(x \to \Theta) = \int_{A_{source}} f_r(x, -\Psi \leftrightarrow \Theta) \cdot L(y \to \Psi) \cdot G(x, y) \cdot dA_y$$

$$G(x, y) = \frac{\cos(n_x, \Theta)\cos(n_y, \Psi)Vis(x, y)}{r_{xy}^2}$$

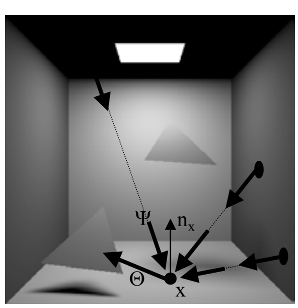

hemisphere integration

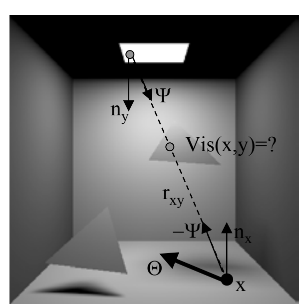

area integration

Estimator for direct lighting

- Pick a point on the light's surface with pdf

$$p(y)$$

- For N samples, direct light at point x is:

$$E(x) = \frac{1}{N}\sum_{i=1}^{N} \frac{f_r L_{source} \frac{\cos\theta_x \cos\theta_{\bar{y}_i}}{r^2_{x\bar{y}_i}} Vis(x,\bar{y}_i)}{p(\bar{y}_i)}$$

PDF for sampling light

- Uniform $p(y) = \frac{1}{Area_{source}}$

- Pick a point uniformly over light's area
  - Can stratify samples

- Estimator:

$$E(x) = \frac{Area_{source}}{N}\sum_{i=1}^{N} f_r L_{source} \frac{\cos\theta_x \cos\theta_{\bar{y}_i}}{r^2_{x\bar{y}_i}} Vis(x,\bar{y}_i)$$

Different pdfs

- Uniform

$$p(y) = \frac{1}{Area_{source}}$$

- Solid angle sampling
  - Removes cosine and distance from integrand
  - Better when significant foreshortening

$$E(x) = \frac{1}{N} \sum_{i=1}^{N} \frac{f_r L_{source} \frac{\cos\theta_x \cos\theta_{\bar{y}_i}}{r_{x\bar{y}_i}^2} Vis(x, \bar{y}_i)}{p(\bar{y}_i)}$$

Parameters

- How to distribute paths within light source?
  - Uniform
  - Solid angle
  - What about light distribution?
- How many paths (“shadow-rays”)?
  - Total?
  - Per light source? (~intensity, importance, …)

Formulation over all lights

- When M is large, each direct lighting sample is very expensive
- We would like to importance sample the lights
- Instead of M integrals

$$L(x \to \Theta) = \sum_{i=1}^{M} \int_{A_{source}} f_r(x, -\Psi \leftrightarrow \Theta) \cdot L_{source}(y \to -\Psi) \cdot G(x, y) \cdot dA_y$$

- Formulation over 1 integration domain

$$L(x \to \Theta) = \int_{A_{all\ lights}} f_r(x, -\Psi \leftrightarrow \Theta) \cdot L_{source}(y \to -\Psi) \cdot G(x, y) \cdot dA_y$$

Why?

- Do not need a minimum of M rays/sample
- Can use only one ray/sample
- Still need N samples, but 1 ray/sample
- Ray is distributed over the whole integration domain
  - Can importance sample the lights

Anti-aliasing

- Can piggy-back on the anti-aliasing of pixel

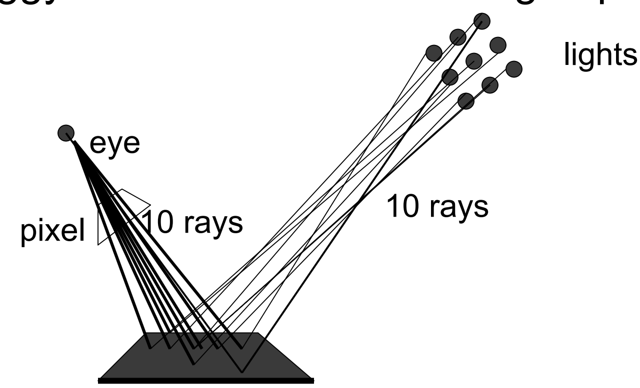

How to sample the lights?

- A discrete pdf $p_L(k_i)$ picks the light k_i
- A surface point is then picked with pdf $p(y_i|k_i)$
- Estimator with N samples:

$$E(x) = \frac{1}{N}\sum_{i=1}^{N}\frac{f_r L_{source} G(x, \overline{y}_i)}{p_L(k_i)\, p(y_i \mid k_i)}$$

Strategies for picking light

– Uniform $p_L(k) = \frac{1}{M}$

– Area $p_L(k) = \frac{A_k}{\sum A_k}$

– Power $p_L(k) = \frac{P_k}{\sum P_k}$

Don’t take visibility into account

Research on many lights

- Ward ‘91
- Shirley, Wang, Zimmerman ‘94
- Fernandez, Bala, Greenberg ‘02
- Wald and Slusallek ‘03

Alternative direct paths

- Shoot paths at random over hemisphere; check if they hit light source

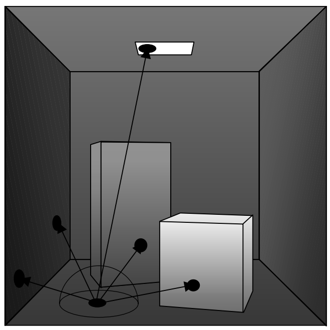

  - paths not used efficiently
  - noise in image
  - might work if light source occupies large portion on hemisphere

Alternative direct paths

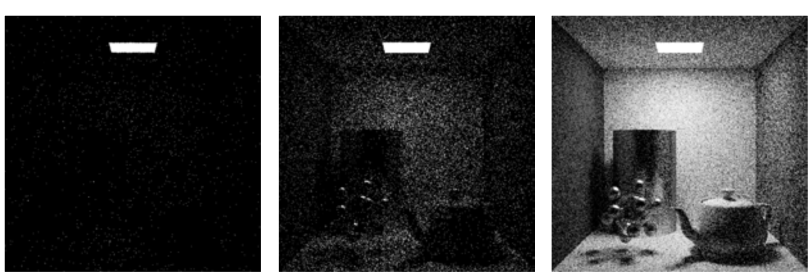

1 path / point | 16 paths / point | 256 paths / point

Direct paths

- Different path generators produce different estimators and different error characteristics
- Direct illumination general algorithm:

```
compute_radiance (point, direction)
        est_rad = 0;
        for (i=0; i<n; i++)
                p = generate_path;
                est_rad += energy_transfer(p) / probability(p);
        est_rad = est_rad / n;
        return(est_rad);
```


Stochastic Ray Tracing

- Sample area of light source for direct term
- Sample hemisphere with random rays for indirect term
- Optimizations:
  – Stratified sampling
  – Importance sampling
  – Combine multiple probability density functions into a single PDF

Indirect Illumination

- Paths of length > 1
- Many different path generators possible
- Efficiency depends on:
  - BRDFs along the path
  - Visibility function
  - ...

Indirect paths - surface sampling

- Simple generator (path length = 2):
  - select point on light source
  - select random point on surfaces

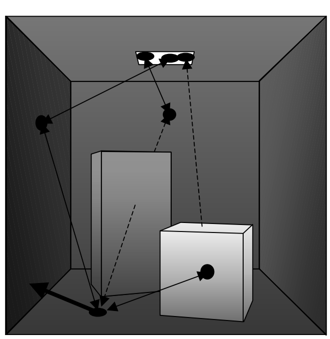

- per path:
  - 2 visibility checks

Indirect paths - surface sampling

- Indirect illumination (path length 2):

$$L(x \to \Theta) = \int_{A_{source}} \int_A L(y \to \Psi_1) f_r(z, -\Psi_1 \leftrightarrow \Psi_2) G(z,y) f_r(x, -\Psi_2 \leftrightarrow \Theta) G(z,x) dA_z dA_y$$

$$\langle L(x \to \Theta) \rangle = \frac{1}{N} \sum_{i=1}^{N} \frac{L(y_i \to \Psi_{1i}) f_r(z_i, -\Psi_{1i} \leftrightarrow \Psi_{2i}) G(z_i, y_i) f_r(x, -\Psi_{2i} \leftrightarrow \Theta) G(z_i, x)}{p_y(y_i) p_z(z_i)}$$

- 2 visibility values cause noise
  – which might be 0

Indirect paths - source shooting

- Shoot ray from light source, find hit location
- Connect hit point to receiver

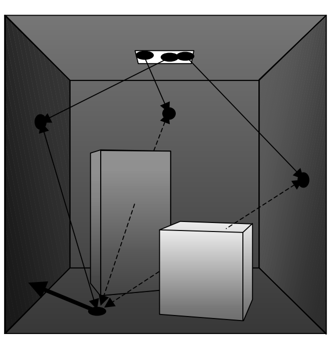

– per path:
  - 1 ray intersection
  - 1 visibility check

Indirect paths - receiver gathering

- Shoot ray from receiver point, find hit location
- Connect hit point to random point on light source

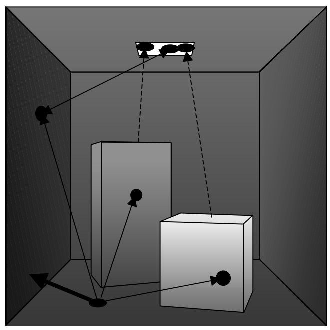

- per path:
  - 1 ray intersection
  - 1 visibility check

Indirect paths

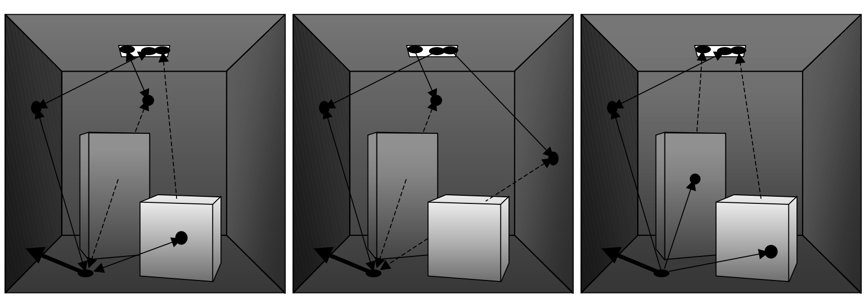

Surface sampling	Source shooting	Receiver gathering
- 2 visibility terms; can be 0	- 1 visibility term - 1 ray intersection	- 1 visibility term - 1 ray intersection

More variants ...

- Shoot ray from receiver point, find hit location
- Shoot ray from hit point, check if on light source

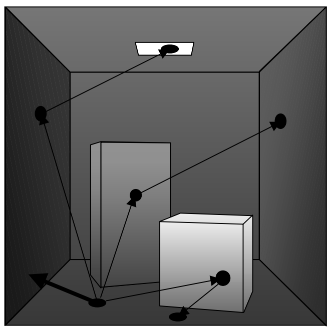

– per path:
  - 2 ray intersections
  - L_e might be zero

Indirect paths

- Same principles apply to paths of length > 2
  – generate multiple surface points
  – generate multiple bounces from light sources and connect to receiver
  – generate multiple bounces from receiver and connect to light sources
  – …

- Estimator and noise characteristics change with path generator

Indirect paths

```
compute_radiance (point, direction)
        est_rad = 0;
        for (i=0; i<n; i++)
                q = generate_indirect_path;
                est_rad += energy_transfer(q) / p(q);
        est_rad = est_rad / n;
        return(est_rad);
```


Stochastic Ray Tracing

- Sample area of light source for direct term
- Sample hemisphere with random rays for indirect term
- Optimizations:
  - Stratified sampling
  - Importance sampling
  - Combine multiple probability density functions into a single PDF

Sampling strategies

- Uniform sampling over the hemisphere

$$L(x \to \Theta) = \int_{\Omega_x} \underline{L(x \leftarrow \Psi) \cdot f_r(\Psi \leftrightarrow \Theta) \cdot \cos(\Psi, n_x)} \, d\omega_\Psi$$

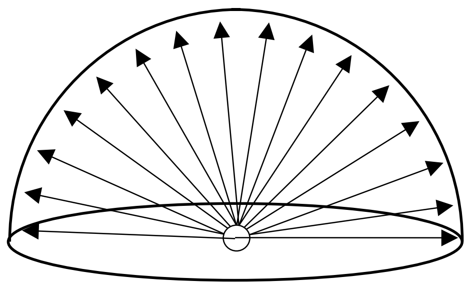

$$p(\Theta) = 1/(2\pi)$$

Sampling strategies

- Sampling according to the cosine factor

$$L(x \to \Theta) = \int_{\Omega_x} \underline{L(x \leftarrow \Psi) \cdot f_r(\Psi \leftrightarrow \Theta)} \, \cos(\Psi, n_x) \cdot d\omega_\Psi$$

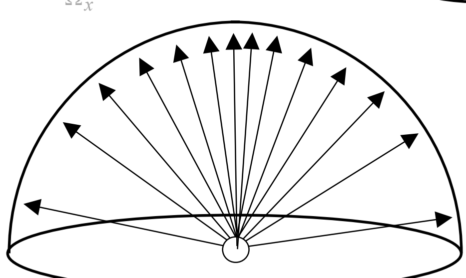

$$p(\Theta) = \cos\theta / \pi$$

Sampling strategies

- Sampling according to the BRDF

$$L(x \to \Theta) = \int_{\Omega_x} \underline{L(x \leftarrow \Psi)} \cdot f_r(\Psi \leftrightarrow \Theta) \cdot \underline{\cos(\Psi, n_x)} \cdot d\omega_\Psi$$

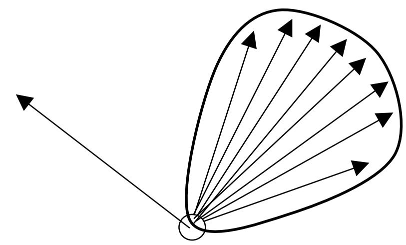

$$p(\Theta) \sim f_r(\Theta \leftrightarrow \Psi)$$

Example: sample according to BRDF

- Discrete pdf q_1, q_2, q_3 $\quad q_1 + q_2 + q_3 = 1$

$$L_{indirect} = L_{diffuse} + L_{specular}$$

$$\langle L_{indirect} \rangle = \begin{Bmatrix} \dfrac{L(x \leftarrow \Psi_i) k_d \cos(N, \Psi_i)}{q_1 p_1(\Psi_i)} \mid \xi < q_1 \\ \dfrac{L(x \leftarrow \Psi_i) k_s \cos^n(R, \Psi_i) \cos(N, \Psi_i)}{q_2 p_2(\Psi_i)} \mid q_1 \le \xi < q_2 \\ 0 \mid otherwise \end{Bmatrix}$$

Sampling strategies

- Sampling according to the BRDF times the cosine

$$L(x \to \Theta) = \int_{\Omega_x} \underline{L(x \leftarrow \Psi)} \cdot f_r(\Psi \leftrightarrow \Theta) \cdot \cos(\Psi, n_x) \cdot d\omega_\Psi$$

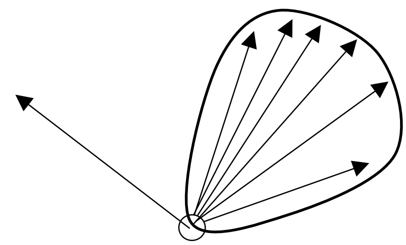

$$p(\Theta) \sim f_r(\Theta \leftrightarrow \Psi)\cos\theta$$

Comparison

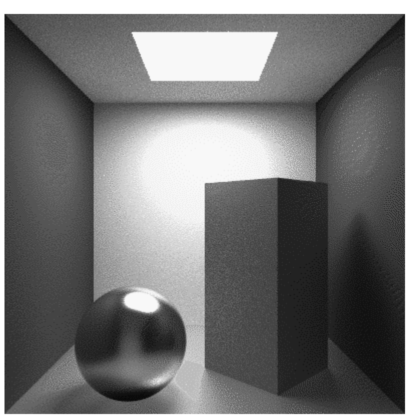

With importance sampling (brdf on sphere)

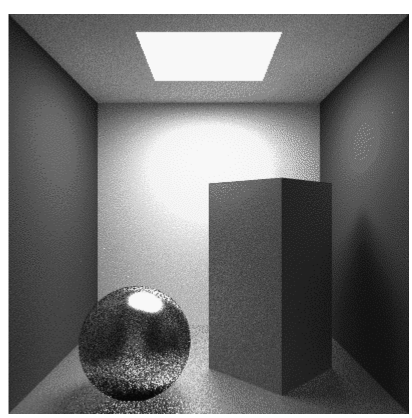

Without importance sampling