Overview

Earlier lecture
- Statistical sampling and Monte Carlo integration

Last lecture
- Signal processing view of sampling
- Stochastic sampling

Today
- Sequential sampling
- Stratified sampling
- Importance sampling
- Multidimensional sampling patterns
- Discrepancy and Quasi-Monte Carlo

Latter
- Path tracing for interreflection

Rendering = Integrals

Cameras

$$R = \int_T \iint_\Omega \iint_A L(x,\omega,t)P(x)S(t)\cos\theta \, dA \, d\omega \, dt$$

Shutter

Motion Blur

Aperture

Depth of Field

Sensor

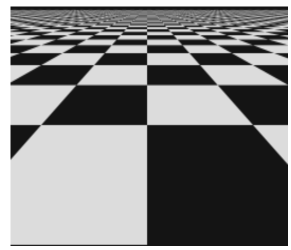

AntiAliasing

Cook, Porter, Carpenter, 1984 Mitchell, 1991

Four Views of Sampling

Sampling is fundamental to rendering

1. **Signal processing view of sampling**
2. **Monte Carlo integration**
3. **Discrepancy and Quasi-Monte Carlo**
4. **Numerical integration and quadrature rules**

Sequential Sampling

Variance

Definition

$$V[Y] \equiv E[(Y - E[Y])^2]$$
$$= E[Y^2 - 2YE[Y] + E[Y]^2]$$
$$= E[Y^2] - E[Y]^2$$

Variance decreases with sample size

$$V[\frac{1}{N}\sum_{i=1}^{N} Y_i] = \frac{1}{N^2}\sum_{i=1}^{N} V[Y_i] = \frac{1}{N} V[Y]$$

Decrease in Variance with N

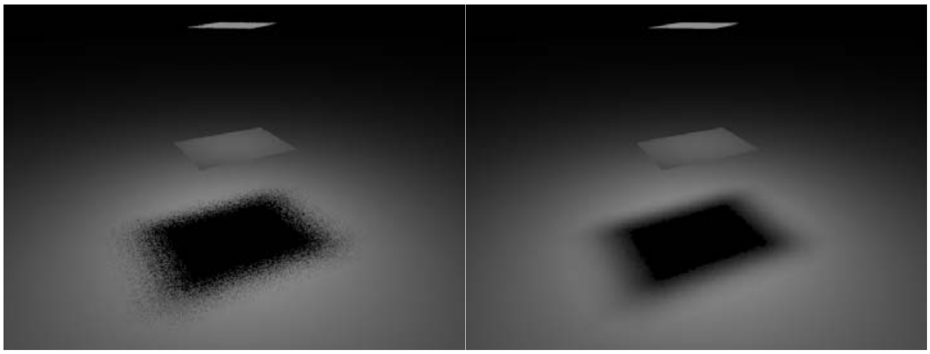

4 eye rays per pixel
1 shadow ray per eye ray

4 eye rays per pixel
16 shadow rays per eye ray

Convergence

Mean and standard deviation

$$\mu_N = \frac{1}{N}\sum_{i=1}^{N} Y_i$$

$$E[\mu_N] = E[\frac{1}{N}\sum_i Y_i] = \frac{1}{N}\sum_i E[Y_i] = E[Y]$$

$$V[\mu_N] = V[\frac{1}{N}\sum_{i=1}^{N} Y_i] = \frac{1}{N^2}\sum_{i=1}^{N} V[Y_i] = \frac{1}{N}V[Y]$$

$$\sigma_N^2 = \frac{1}{N-1}\sum_{i=1}^{N}(Y_i - \mu_N)^2$$

Variation in Estimators

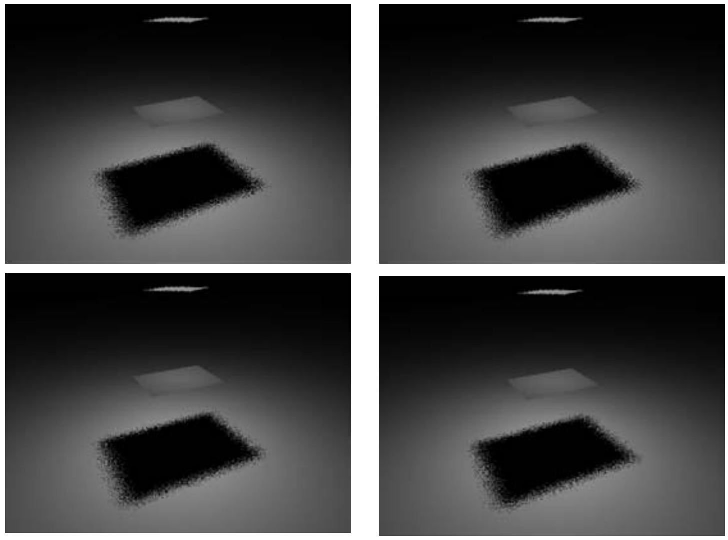

1 eye ray and 1 shadow ray per pixel

Sequential Sampling

Central Limit Theorem

$$\lim_{N\to\infty} \Pr\left\{\mu_N - E[Y] \le \frac{t\sigma}{N^{1/2}}\right\} = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{t} e^{-x^2/2}\,dx$$

Send rays until confidence in the estimate is high

- **Student t-distribution**
  Purgathofer (1987)
- **Chi-squared distribution**
  Lee, Redner, Uselton (1985)

Stratified Sampling

Stratified Sampling

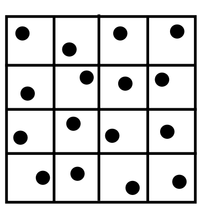

Stratified sampling like jittered sampling

Allocate samples per region

$$N = \sum_{i=1}^{m} N_i \quad F_N = \frac{1}{N} \sum_{i=1}^{m} N_i F_i$$

New variance

$$V[F_N] = \frac{1}{N^2} \sum_{i=1}^{m} N_i V[F_i]$$

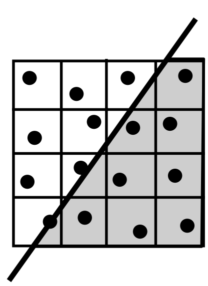

Thus, if the variance in regions is less than the overall variance, there will be a reduction in resulting variance

For example: An edge through a pixel

$$V[F_N] = \frac{1}{N^2} \sum_{i=1}^{\sqrt{N}} V[F_j] = \frac{V[F_E]}{N^{1.5}}$$

Shirley's Mapping

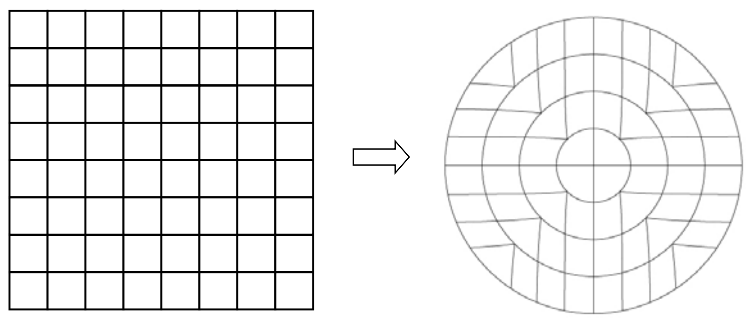

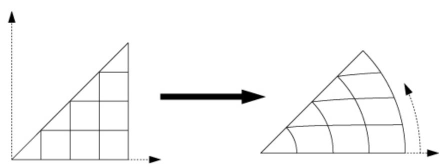

$$r = U_1$$

$$\theta = \frac{\pi}{4}\frac{U_2}{U_1}$$

High-dimensional Sampling

Stratified sampling (also numerical quadrature)

For a given error …

$$N \sim n^d$$

Random sampling

For a given variance …

$$V^{1/2} \sim \frac{1}{N^{1/2}}$$

Monte Carlo Integration better than numerical methods in high dimensional spaces

Importance Sampling

Importance Sampling

Biasing the sampling process

$$X_i \sim p(x)$$

$$Y_i = \tilde{f} = \frac{f(X_i)}{p(X_i)}$$

$$E[Y_i] = E\left[\frac{f(X_i)}{p(X_i)}\right]$$

$$= \int \left[\frac{f(X_i)}{p(X_i)}\right] p(x)\, dx$$

$$= \int f(x)\, dx$$

$$= I$$

Zero Variance with Perfect Biasing

$$\tilde{p}(x) = \frac{f(x)}{E[f]}$$

$$E[\tilde{f}^2] = \int \left[\frac{f(x)}{\tilde{p}(x)} \right]^2 \tilde{p}(x)\, dx$$

$$= \int \left[\frac{f(x)}{f(x)/E[f]} \right]^2 \frac{f(x)}{E[f]} dx$$

$$= E[f]^2$$

$$\Rightarrow \quad V[\tilde{f}] = E[\tilde{f}^2] - E[\tilde{f}]^2 = 0$$

But can't do this in reality

Sampling f is equivalent to knowing the integral of f

But, the closer to f, the lower the variance

Examples

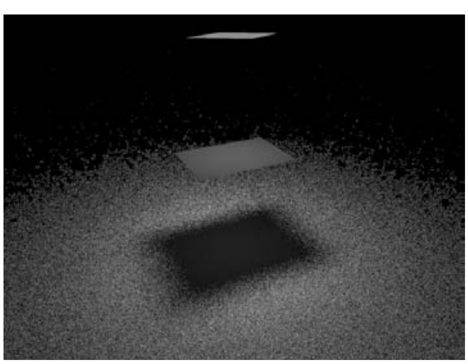

Projected solid angle

4 eye rays per pixel
100 shadow rays

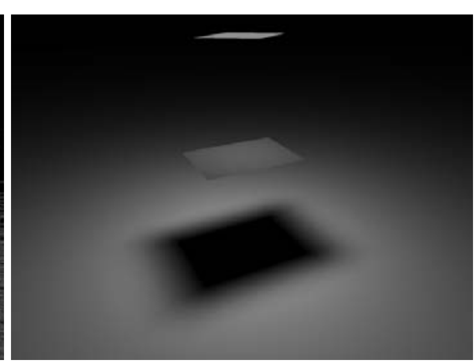

Area

4 eye rays per pixel
100 shadow rays

Variance Reduction

Efficiency measure

$$Efficiency \propto \frac{1}{Variance \bullet Cost}$$

Techniques

- Sequential sampling
- Importance sampling
  - Later, multiple importance sampling
- Stratified sampling
- …

Quasi-Monte Carlo View

Discrepancy

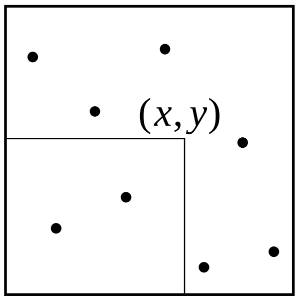

$A = xy$

$n(x, y)$ number of samples in A

$$\Delta(x,y) = \frac{n(x,y)}{N} - xy$$

$$D_N = \max_{x,y} \left| \Delta(x,y) \right|$$

Theorem on Total Variation

Theorem (Koksma 1942):

$$\left| \frac{1}{N} \sum_{i=1}^{N} f(X_i) - \int_0^1 f(x)\,dx \right| \le V(f) D_N$$

Proof: Integrate by parts

$$\int f(x) \quad [\frac{\delta(x - x_i)}{N} - 1]\,dx \qquad\qquad \frac{\partial \Delta(x)}{\partial x} = \frac{\delta(x - x_i)}{N} - 1$$

$$= \int f(x) \frac{\partial \Delta(x)}{\partial x}\,dx$$

$$= f\Delta\Big|_0^1 - \int \frac{\partial f(x)}{\partial x} \Delta(x) \cdots dx = -\int \frac{\partial f(x)}{\partial x} \Delta(x)\,dx$$

$$\le D_N \int \left| \frac{\partial f(x)}{\partial x} \right| dx = V(f) D_N$$

Halton and Hammersley Points

Radical inverse (digit reverse)

of integer *i* in integer base *b* $\phi_2(i)$

$$i = d_i \cdots d_2 d_1 d_0$$
$$\phi_b(i) \equiv 0.d_0 d_1 d_2 \cdots d_i$$

1	1	.1	1/2
2	10	.01	1/4
3	11	.11	3/4
4	100	.001	1/8
5	101	.101	5/8

Hammersley points

$(i/N, \phi_2(i), \phi_3(i), \phi_5(i), \cdots)$ $D_N = O(\frac{\log^{d-1} N}{N})$

Halton points (sequential)

$(\phi_2(i), \phi_3(i), \phi_5(i), \cdots)$ $D_N = O(\frac{\log^{d} N}{N})$

Zaremba-Hammersley Points

Folded radical inverse $\psi_r(i)$ $d_i' = (d_i + i) \bmod r$

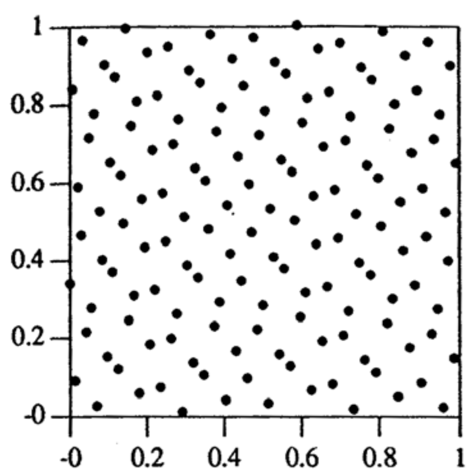

Zaremba-Hammersley Points

Edge Discrepancy

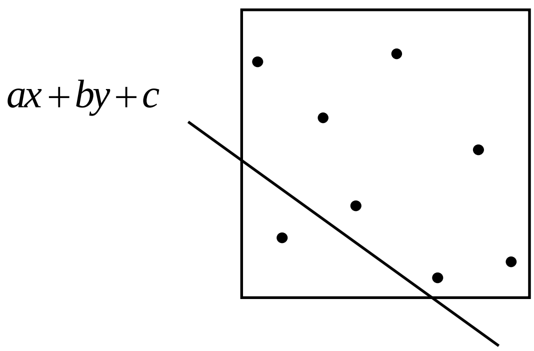

$$D_N = \Omega\left(\frac{1}{N^{3/4}}\right) \qquad D_N = O\left(\frac{\log^{1/2} N}{N^{3/4}}\right)$$

Edge Discrepancy

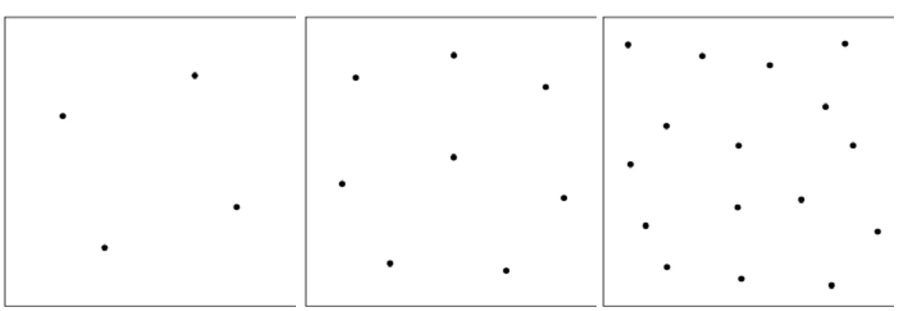

Mitchell low-discrepancy patterns

Edge Discrepancy

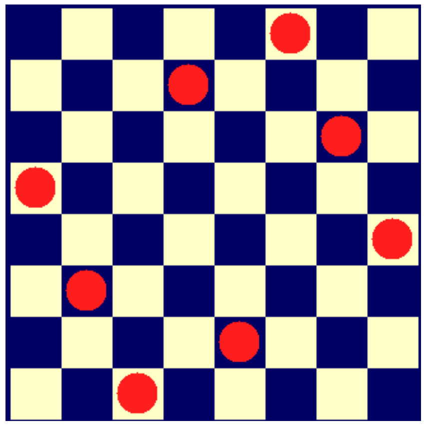

8-queens pattern

SGI IR Multisampling
- 8x8 subpixel grid
- 1,2,4,8 samples

Low-Discrepancy Patterns

Process	16 points	256 points	1600 points
Zaremba	0.0504	0.00478	0.00111
Jittered	0.0538	0.00595	0.00146
Poisson-Disk	0.0613	0.00767	0.00241
N-Rooks	0.0637	0.0123	0.00488
Random	0.0924	0.0224	0.00866

Discrepancy of random edges, From Mitchell (1992)

Random sampling converges as $N^{-1/2}$
Zaremba converges faster and has lower discrepancy
Zaremba has a relatively poor blue noise spectra
Jittered and Poisson-Disk recommended for images

Multidimensional Sampling

Block Design

a	*d*	*c*	*b*
b	*a*	*d*	*c*
c	*b*	*a*	*d*
d	*c*	*b*	*a*

Latin Square

Block design

Alphabet of size *n*

Each symbol appears exactly once in each row and column

a			
		a	
	a		
			a

N-Rook Pattern

Incomplete block design

Replaced N^d samples with N samples

Permutations: $(\pi_1(i), \pi_2(i), \cdots \pi_d(i))$

Generalizations: N-queens, 2D projection

Space-time Patterns

6	10	2	13
3	14	12	8
15	0	7	11
5	9	4	1

Cook Pattern

15	8	5	2
4	3	14	9
10	13	0	7
1	6	11	12

Pan-diagonal Magic Square

Fully populate (x,y) samples

- **Recall blue noise good**
  - **Perceptually pleasing**
  - **Filtered during resampling**
- **Jitter to achieve blue noise**

Distribute t samples

- **Decorrelate space and time**
- **Nearby samples in space should differ greatly in time**

Mitchell (1991) designs

Spectrally Optimized Sampling

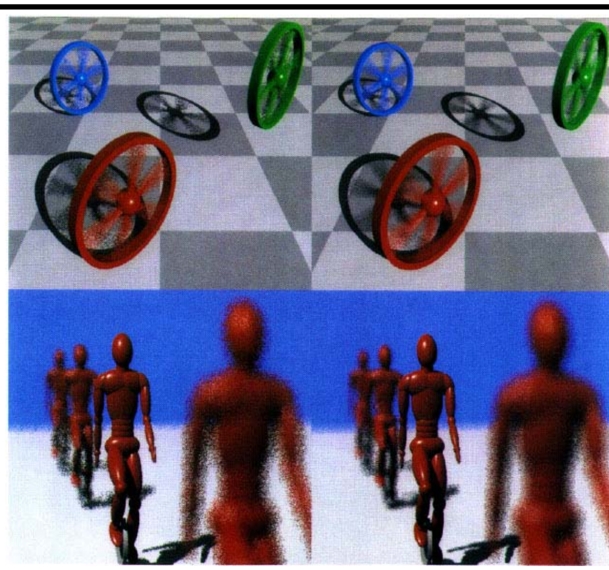

From Mitchell

Summary

4 Views of Sampling and Integration

Rendering requires high-dimensional integration

1. **Signal processing**
   - **Aliases and antialiasing**
   - **Antialiasing by supersampling**
   - **Stochastic sampling decorrelates aliases**
   - **Image plane is different than other dimensions**
     - **Blue noise good**
2. **Statistical sampling and Monte Carlo integration**
   - **Sequential, stratified and importance sampling**
   - **Variance reduction and efficiency**
   - **Asymptotic efficiency $N^{-1/2}$**
   - **Good for high dimensional sampling**

4 Views of Sampling

Rendering requires high-dimensional integration

1. Signal processing
2. Statistical sampling and Monte Carlo integration
3. Quasi Monte Carlo
   - Discrepancy and optimal sampling patterns
   - Lower bound on asymptotic efficiency in high dimensions
   - Artifacts in image space
4. Numerical
   - Quadrature rules for smooth functions
     - Bad for discontinuous functions; edges ...
   - Inefficient: N^d
     - Used in radiosity