Image Capture

Physical cameras create an image by integrating over an active area of a sensor

$$R = \int_A E(x,y)\,P(x,y)\,dx\,dy$$

Examples:

- **Vidicon: phosphors**
- **CCD array**
- **Retina: photoreceptors**

But, ray-tracers and z-buffer algorithms point sample ...

Image Generation

Physical displays create an image by using a finite sized source of light

Examples:

- **Printing: halftone dots**
- **Cathode ray tube: phosphors**
- **DAC's: sample and hold**

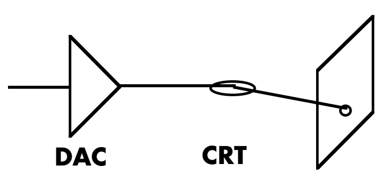

Fourier Analysis

The *Fourier transform* projects a function onto sines and cosines

$$F(\omega) = \int_{-\infty}^{\infty} f(x) e^{-i\omega x}\, dx$$

$$f(x) = \int_{-\infty}^{\infty} F(\omega) e^{i\omega x}\, d\omega$$

F(w) is the spectrum of the function

The spectrum quantifies the relative contributions of different frequencies

Filtering

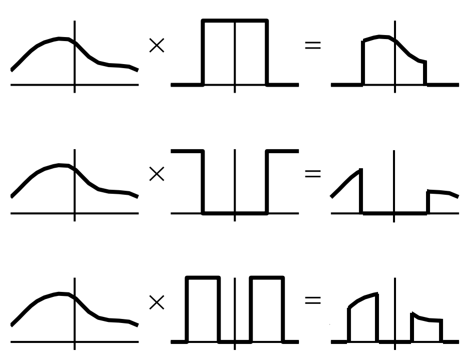

Convolution Theorem

***Theorem*: Multiplication in the frequency (space) domain is equivalent to convolution in the space (freq) domain.**

$$F \times G \quad \Leftrightarrow \quad f \otimes g$$

Where:

$$f \otimes g = \int f(y) g(x-y)\, dy$$

For example:

- **Multiply by filter and sum**
- **Moving average**

Sampling: Spatial Domain

Sampling: Frequency Domain

Reconstruction: Spatial Domain

Reconstruction: Frequency Domain

Sampling and Reconstruction

Sampling Theorem

This result if known as the *Sampling Theorem* and is due to Claude Shannon who first discovered it in 1949

> A signal can be reconstructed from its samples without loss of information, if the original signal has no frequencies above 1/2 the sampling frequency

For a given bandlimited function, the rate at which it must be sampled is called the *Nyquist Frequency*

Undersampling: Pre-Aliasing

Poor reconstruction: Post-Aliasing

Aliasing and Antialiasing

If the scene contains frequenciesgreater than the Nyquist Frequency, then we have a aliasing problem

Results of aliasing:

- Jaggies
- Moire
- Flickering small objects
- Sparkling highlights
- Temporal strobing

Preventing aliasing or antialiasing:

1. Prefilter the signal
2. Increase sampling rate

First computer graphics commandment: Thou shalt not make pictures with computational artifacts

Prefiltering

Recall that physical cameras integrates over an sensor area

Typically: area average or convolve with a box filter before sampling

Ideally: convolve with a sinc function to bandlimit the function at the Nyquist sampling rate.

Practically:

- **Constant colored polygonal fragments doable**
- **Complex environments not doable**

Supersampling

Increasing the sampling rate reduces, but does not eliminate, aliasing

Simplest form: Compute sample mean; area average

More generally: resample at pixel frequency

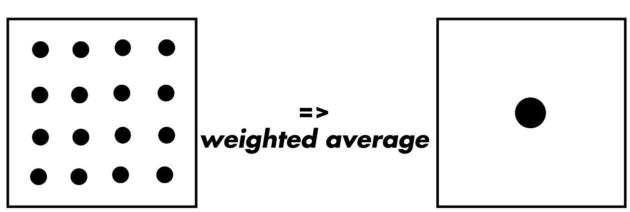

Non-uniform Sampling

In scanline, constant spaced samples

In ray tracing, may sample anywhere

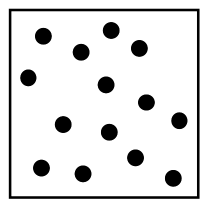

Spectral Consequences

Uniform: The spectrum of a uniform array of samples
is a uniform array of spikes

Non-uniform: The spectrum non-uniform array of samples, has a "DC" spike and "noise"

Thus, the sampled spectrum will have a copy of f

And, noisy f's -> *Aliases are randomized*

Blue Noise

White Noise

Blue Noise

Properties:

- High frequencies may be removed during resampling
- Less objectionable perceptual artifact

History

- Yelliot (1982) observed in Monkey retinas
- Cook (1984) applied to computer graphics
- Poisson Disk ideal

Jittered Sampling

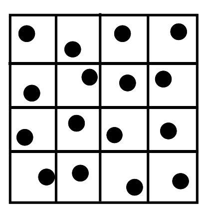

Divide pixel into a uniform grid

- Randomly perturb sample point in the grid
- "Decorrelate" from pixel to pixel

Variation:

- N-Rooks or N-Queens sampling

Statistical Sampling

Estimate pixel value by sample average:

$$\bar{f} = \frac{1}{N}\sum_{i=1}^{N} f(x_i, y_i)$$

The <u>Central Limit Theorem</u> states that the set of average values are normally distributed as N increases

The variance of this distribution is:

$$\sigma^2 \approx s^2 = \frac{1}{N-1}\sum_{i=1}^{N}(f(x_i, y_i) - \bar{f})^2$$

Variance Reduction

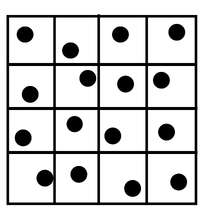

Stratified sampling like jittered sampling

$$\sigma_f^2 = \sum_{i=1}^{N}\frac{\sigma_i^2}{N^2} \le \frac{\sigma^2}{N}$$

Strata may individually have less variation

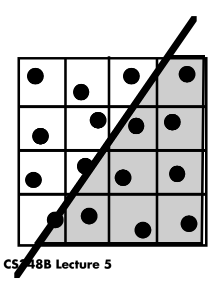

For example: An edge through a pixel

$$\sigma_f^2 = \sum_{i=1}^{\approx\sqrt{N}}\frac{\sigma_i^2}{N^2} \le \frac{\sigma^2}{N^{1.5}}$$

Sequential Sampling

Since averages are normally distributed,
Likelihood of lying outside the variance
may be estimated using student t distribution

Sample until probability is small

This is a form of *adaptive sampling*

Essentially the procedures used by pollsters

Adaptive Sampling

Whitted

Kajiya

Mitchell

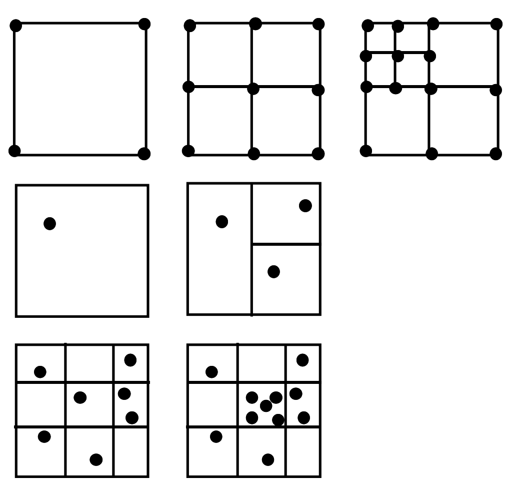

Four Views of Sampling

1. Signal processing
   - Aliasing and Antialiasing
   - Frequency viewpoint
2. Statistical view
   - Variance; central limit theorem
   - Monte Carlo
3. Discrepancy
   - Asymptotic efficiency
   - Quasi Monte Carlo
4. Numerical
   - Quadrature/Integration rules