Sampling and Reconstruction

Digital Image Synthesis

Yung-Yu Chuang

with slides by Pat Hanrahan, Torsten Moller and Brian Curless

Sampling theory

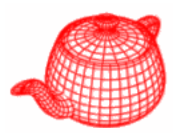

- Sampling theory: the theory of taking discrete sample values *(grid of color pixels)* from functions defined over continuous domains *(incident radiance defined over the film plane)* and then using those samples to reconstruct new functions that are similar to the original (reconstruction).

`Sampler`: selects sample points on the image plane

`Filter`: blends multiple samples together

Aliasing

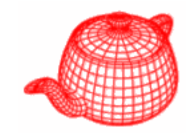

- Reconstruction generates an approximation to the original function. Error is called aliasing.

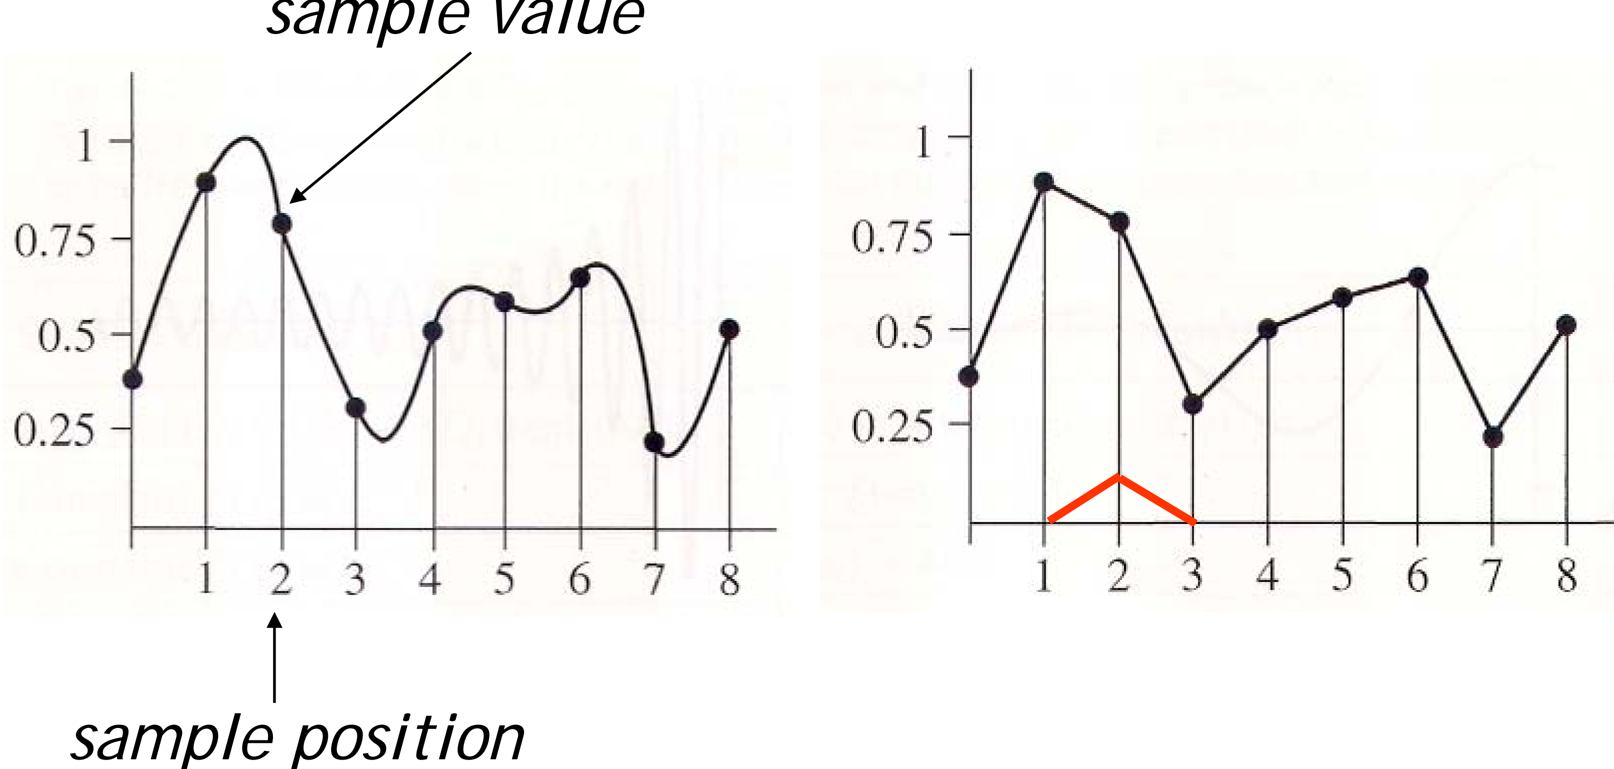

Sampling in computer graphics

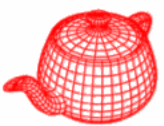

- Artifacts due to sampling - Aliasing
  - Jaggies
  - Moire
  - Flickering small objects
  - Sparkling highlights
  - Temporal strobing (such as Wagon-wheel effect)
- Preventing these artifacts - Antialiasing

Jaggies

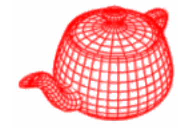

Retort sequence by Don Mitchell

Staircase pattern or jaggies

Moire pattern

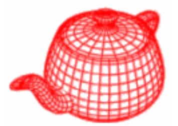

- Sampling the equation

  $\sin(x^2 + y^2)$

Fourier analysis

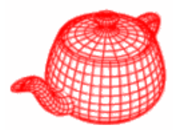

- Can be used to evaluate the quality between the reconstruction and the original.
- The concept was introduced to Graphics by Robert Cook in 1986. (extended by Don Mitchell)

Rob Cook

V.P. of Pixar
1981 M.S. Cornell
1987 SIGGRAPH Achievement award
1999 Fellow of ACM
2001 Academic Award with Ed Catmull and Loren Carpenter (for Renderman)

Fourier transforms

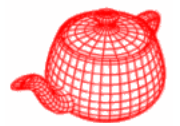

- Most functions can be decomposed into a weighted sum of shifted sinusoids.
- Each function has two representations
  - Spatial domain - normal representation
  - Frequency domain - spectral representation
- The *Fourier transform* converts between the spatial and frequency domain

Spatial Domain

$f(x)$

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

$$\Leftarrow f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} F(\omega) e^{i\omega x}\, d\omega \Leftarrow$$

Frequency Domain

$F(\omega)$

Fourier analysis

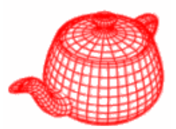

spatial domain *frequency domain*

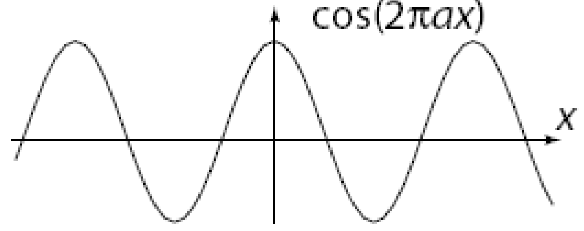

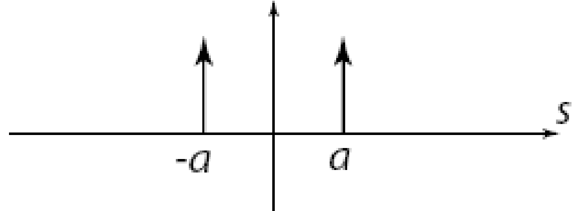

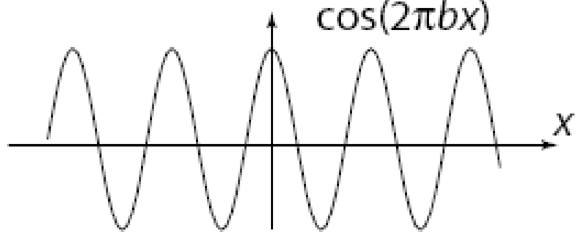

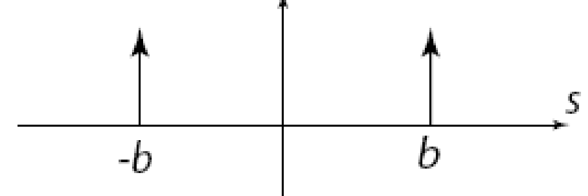

Fourier analysis

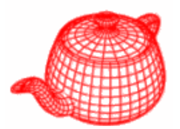

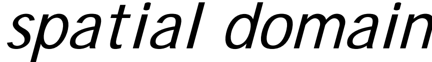

frequency domain

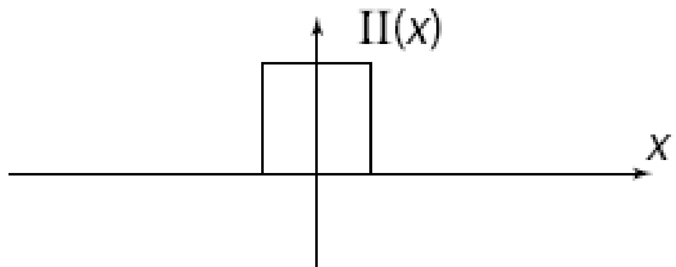

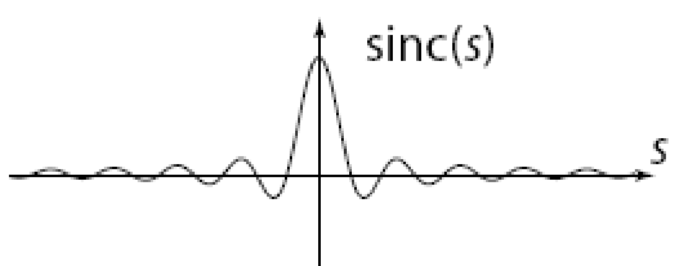

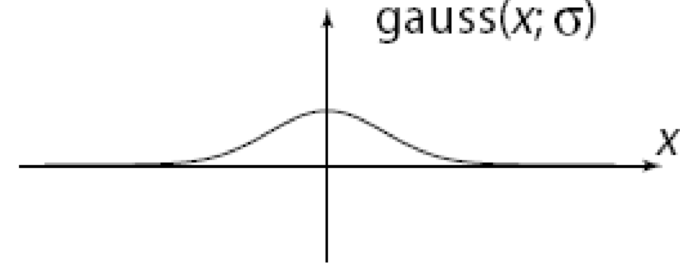

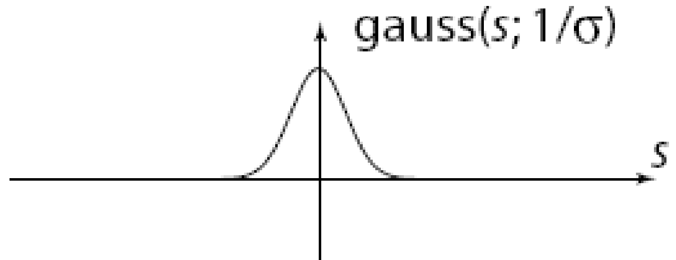

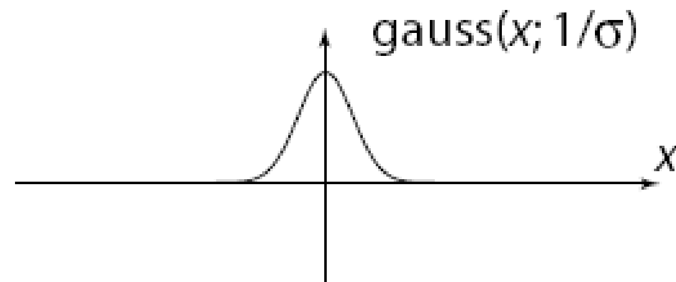

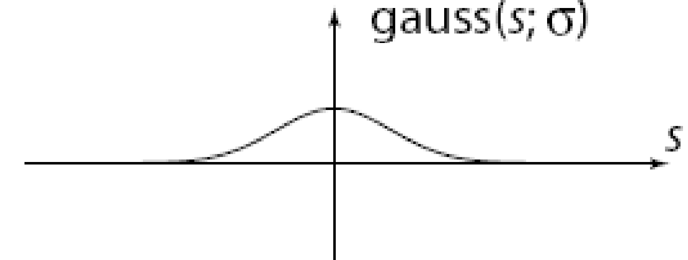

Fourier analysis

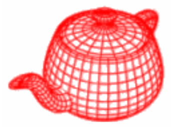

spatial domain

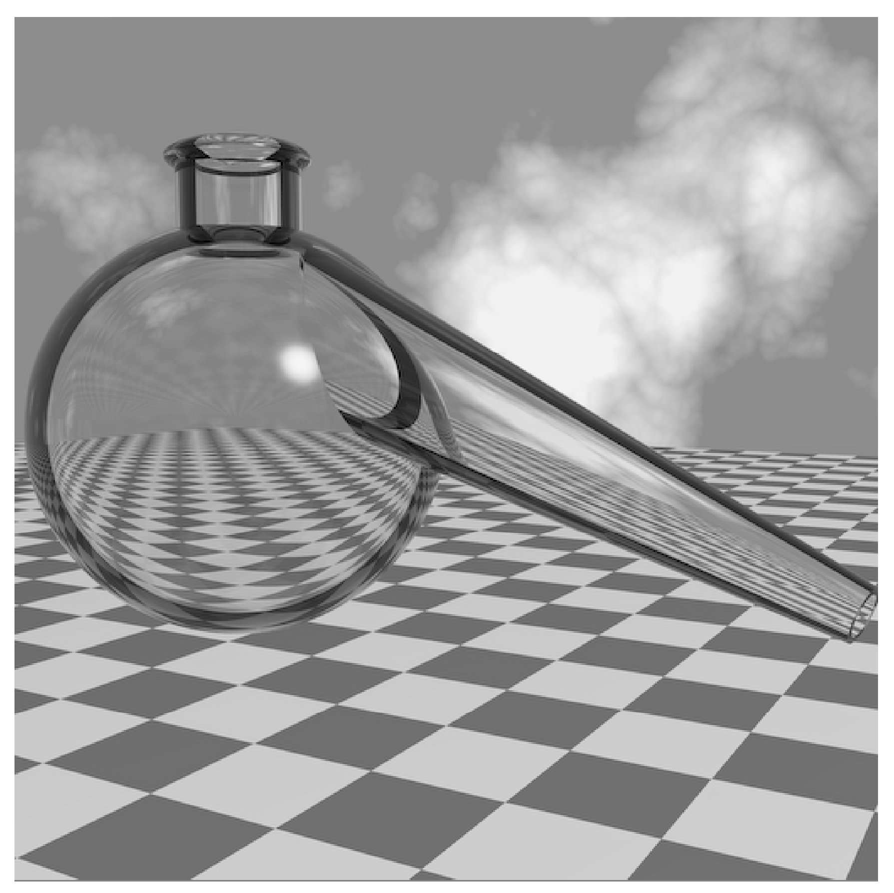

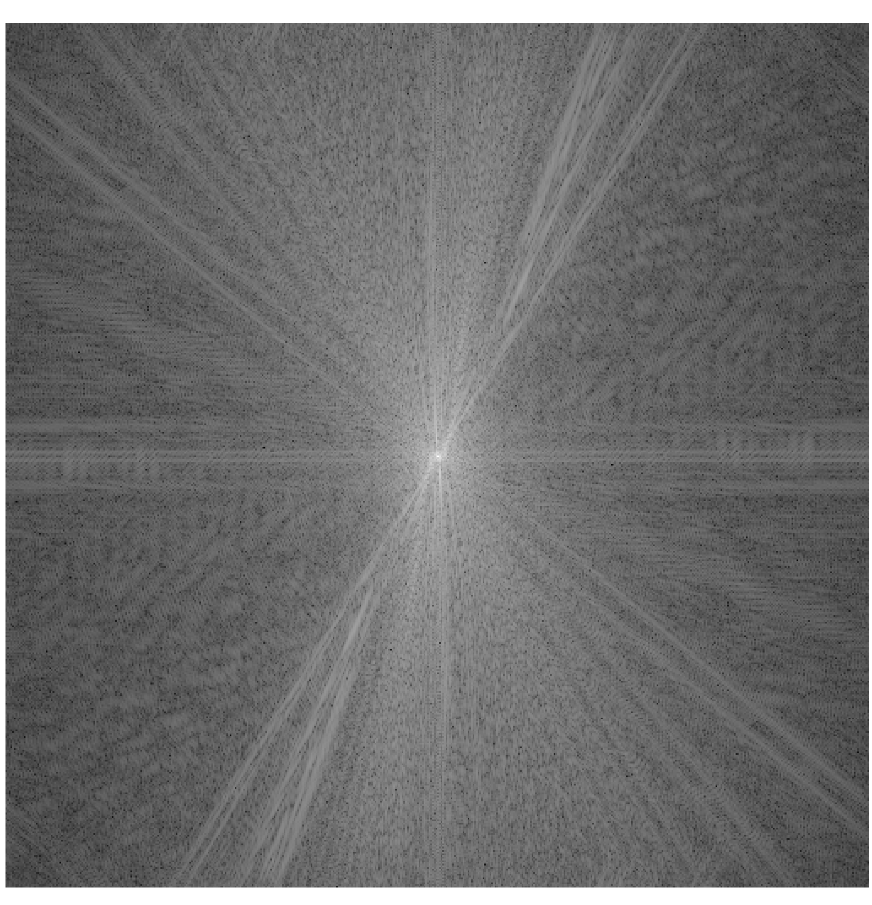

Convolution

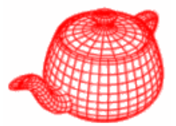

- *Definition*

$$h(x) = f \otimes g = \int f(x')g(x - x')\,dx'$$

- *Convolution Theorem*: Multiplication in the frequency domain is equivalent to convolution in the space domain.

$$f \otimes g \leftrightarrow F \times G$$

- *Symmetric Theorem*: Multiplication in the space domain is equivalent to convolution in the frequency domain.

$$f \times g \leftrightarrow F \otimes G$$

1D convolution theorem example

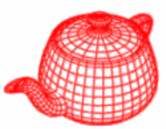

$f(x)$ x | $F(s)$ s

$*$ | $\times$

$h(x)$ x | $H(s)$ s

$\Downarrow$ | $\Downarrow$

$g(x)$ x | $G(s)$ s

2D convolution theorem example

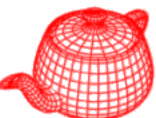

$f(x,y)$		$g(x,y)$		$h(x,y)$
	$*$		$\Rightarrow$	
	$\times$		$\Rightarrow$	
$F(s_x,s_y)$		$G(s_x,s_y)$		$H(s_x,s_y)$

The delta function

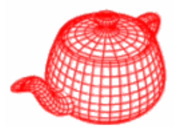

- Dirac delta function, zero width, infinite height and unit area

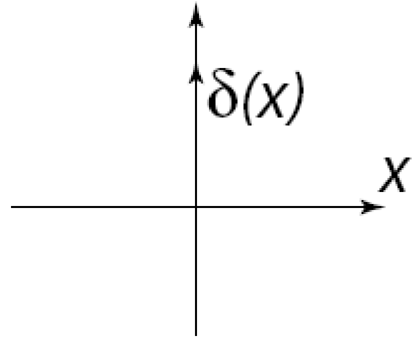

Sifting and shifting

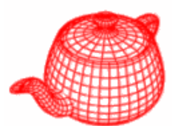

Sifting:

$$\int_{-\infty}^{+\infty} f(x)\delta(x-a)dx = \int_{a-\varepsilon}^{a-\varepsilon} f(x)\delta(x-a)dx = f(a)\int_{a-\varepsilon}^{a-\varepsilon} \delta(x-a)dx$$
$$= f(a)$$

$$f(x)\delta(x-a) = f(a)\delta(x-a)$$

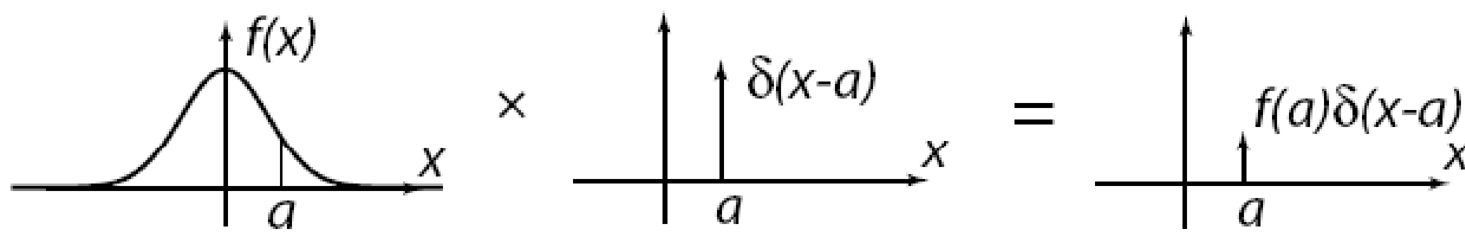

Shifting:

$$f(x) * \delta(x-a) = f(x-a)$$

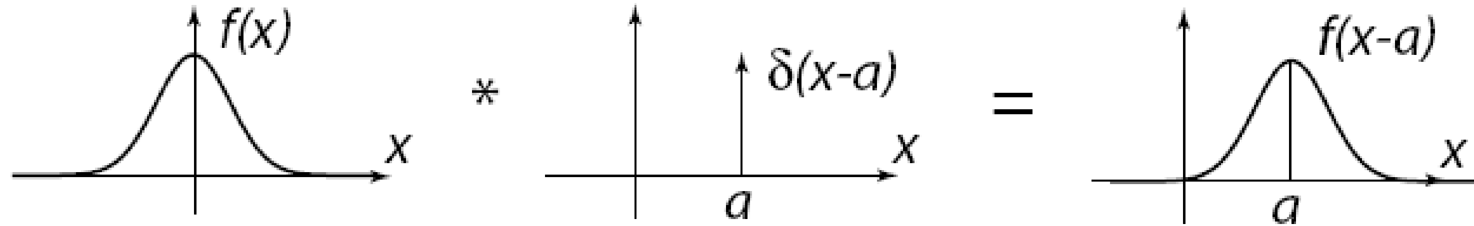

Shah/impulse train function

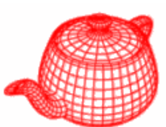

spatial domain

$$\text{III}(x) = \sum_{n=-\infty}^{\infty} \delta(x - nT)$$

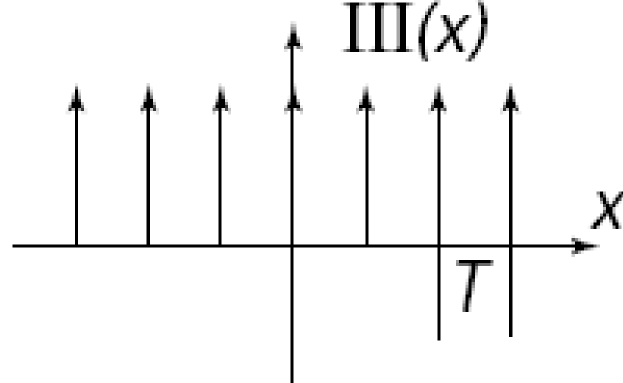

frequency domain

$$\text{III}(s) = \sum_{n=-\infty}^{\infty} \delta(s - ns_o), \quad s_o = 1/T$$

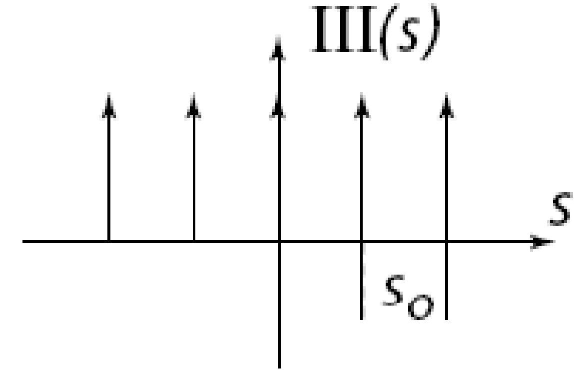

Sampling

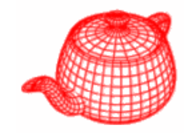

band limited

x

s

×

∗

III(x)

III(s)

x

s

T

s_o

⇓

⇓

x

s

Reconstruction

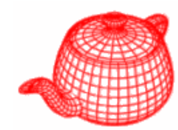

Reconstruction filter

The reconstructed function is obtained by interpolating among the samples in some manner

In math forms

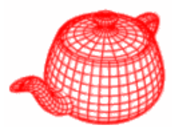

$$\widetilde{F} = (F(s) * \text{III}(s)) \times \Pi(s)$$

$$\widetilde{f} = (f(x) \times \text{III}(x)) * \text{sinc}(x)$$

$$\widetilde{f}(x) = \sum_{i=-\infty}^{\infty} \text{sinc}(x-i) f(i)$$

Reconstruction filters

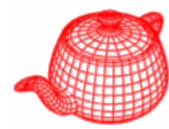

The sinc filter, while ideal, has two drawbacks:

- It has a large support (slow to compute)
- It introduces ringing in practice

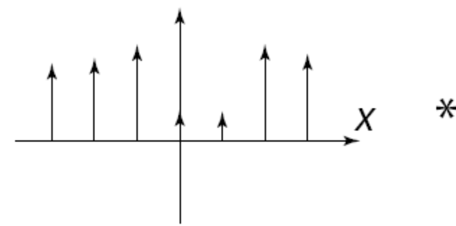

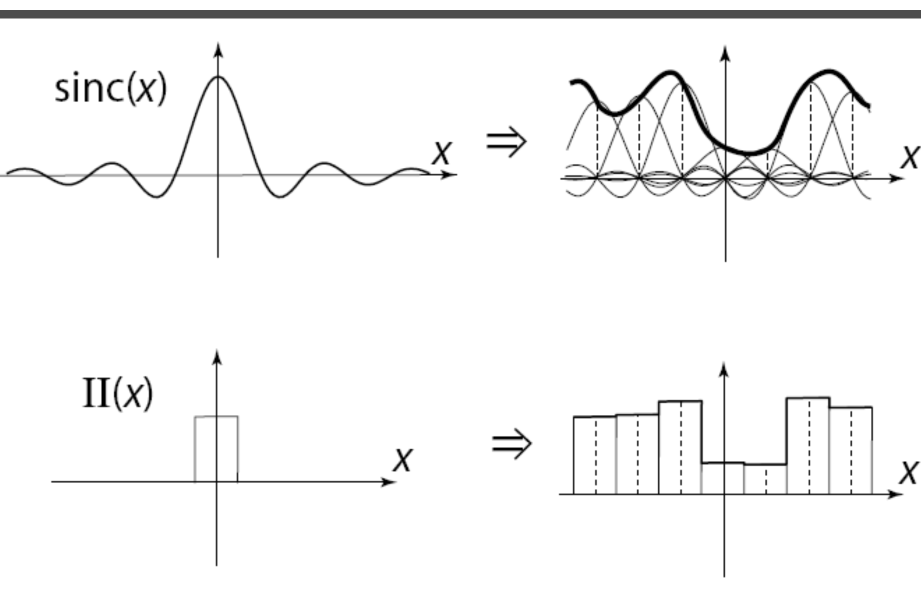

The box filter is bad because its Fourier transform is a sinc filter which includes high frequency contribution from the infinite series of other copies.

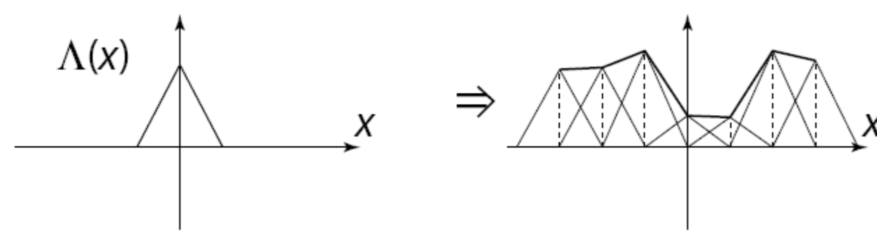

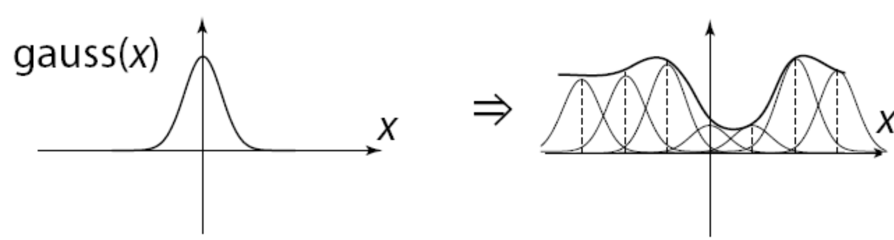

Aliasing

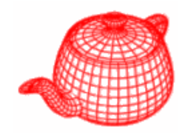

increase sample spacing in spatial domain

decrease sample spacing in frequency domain

$\times$

$*$

$\Downarrow$

$\Downarrow$

Aliasing

Aliasing

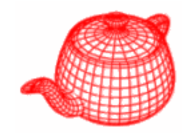

Aliasing

x

s

high-frequency details leak into lower-frequency regions

*

×

x

s

⇓

⇓

x

s

Sampling theorem

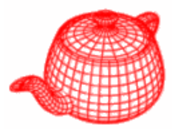

This result is known as the **Sampling Theorem** and is due to Claude Shannon who first discovered it in 1949:

> A signal can be reconstructed from its samples without loss of information, if the original signal has no frequencies above ½ the sampling frequency.

For a given **bandlimited** function, the minimum rate at which it must be sampled is the **Nyquist frequency**.

Sampling theorem

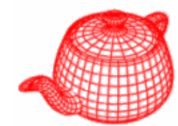

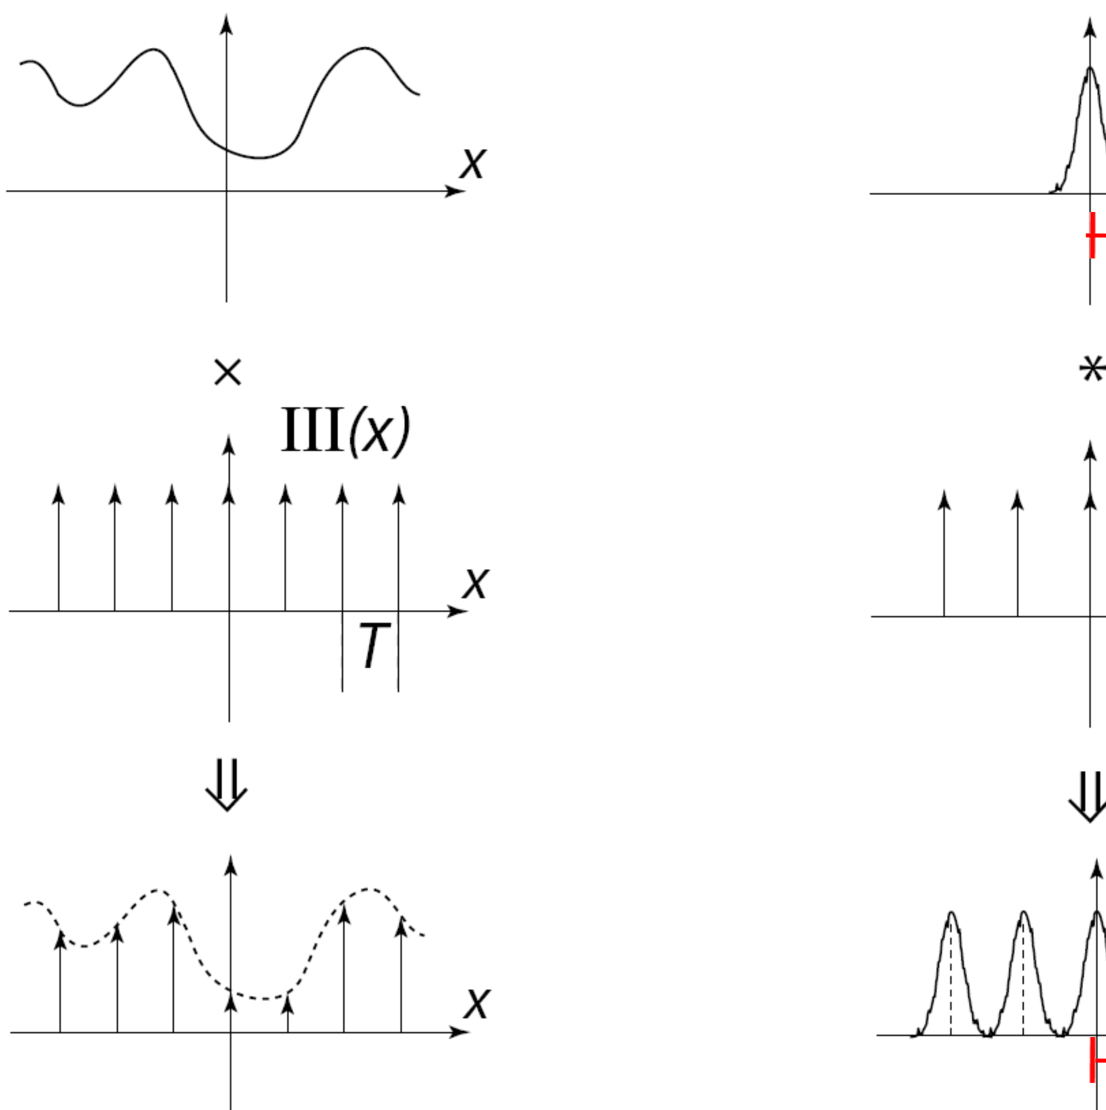

Aliasing due to under-sampling

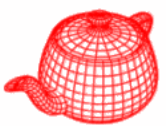

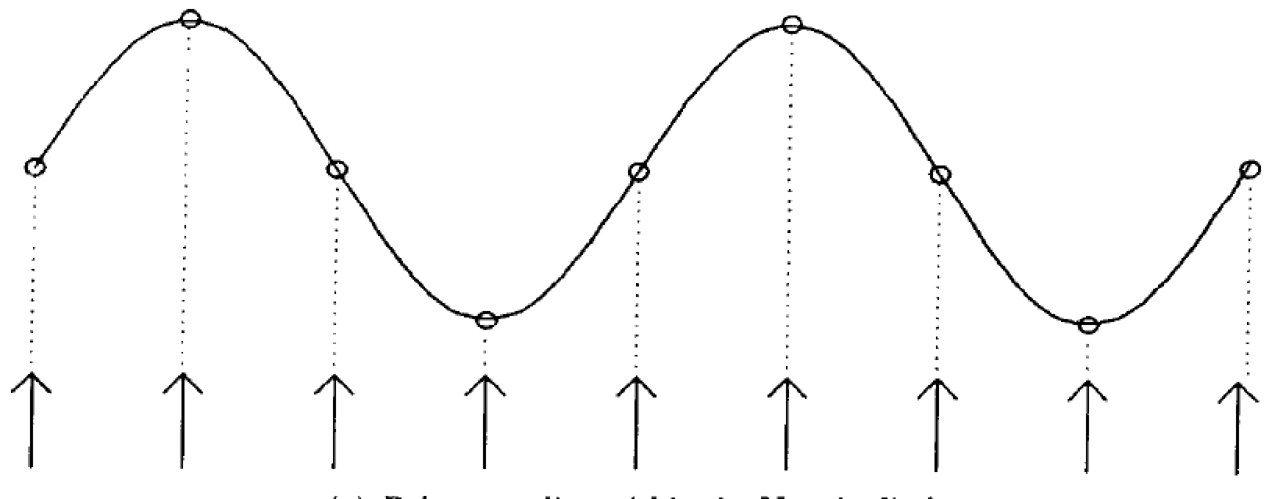

(a) Point sampling within the Nyquist limit

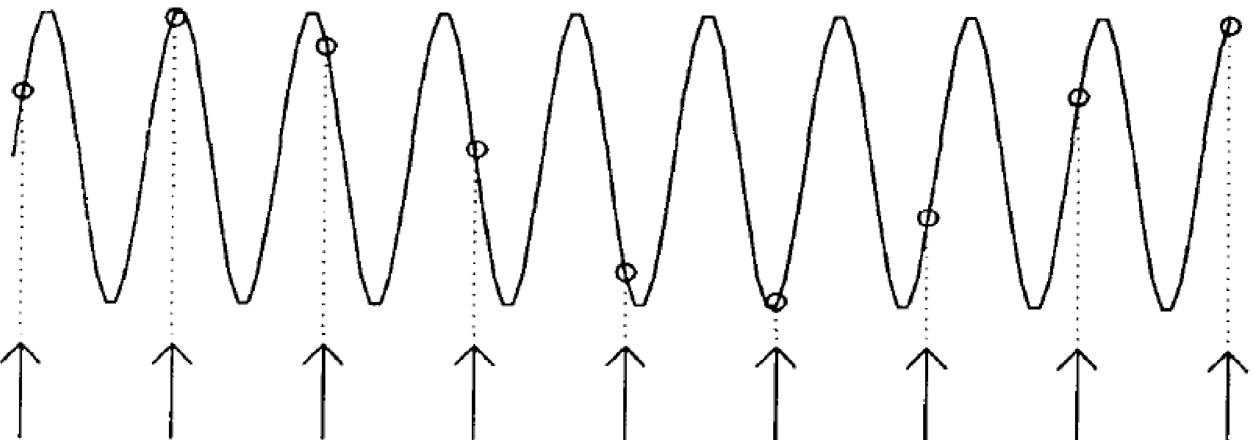

(b) Point sampling beyond the Nyquist limit

Sampling theorem

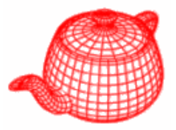

- For band limited functions, we can just increase the sampling rate
- However, few of interesting functions in computer graphics are band limited, in particular, functions with discontinuities.
- It is mostly because the discontinuity always falls between two samples and the samples provides no information about this discontinuity.

Aliasing

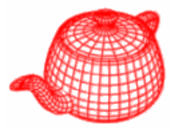

- Prealiasing: due to sampling under Nyquist rate
- Postaliasing: due to use of imperfect reconstruction filter

Antialiasing

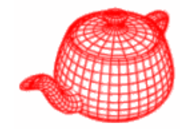

- Antialiasing = Preventing aliasing

1. Analytically prefilter the signal
   - Not solvable in general
2. Uniform supersampling and resample
3. Nonuniform or stochastic sampling

Antialiasing (Prefiltering)

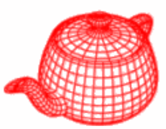

x   s

*   ×

x   s

⇓   ⇓

x   s

It is blurred, but better than aliasing

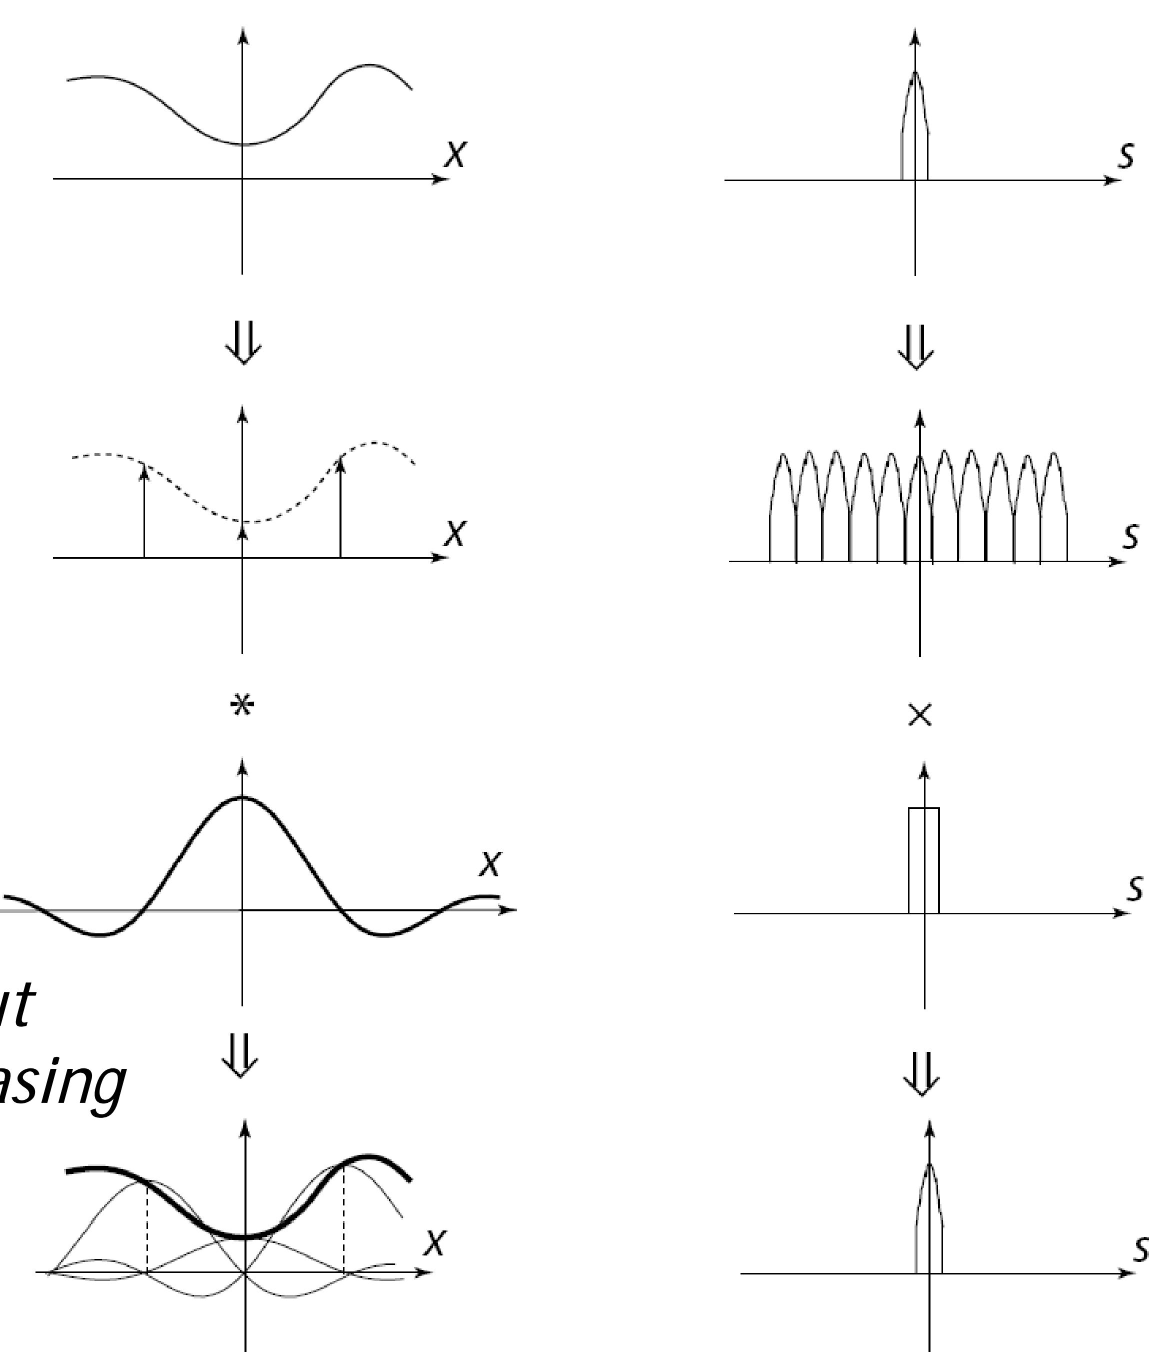

Uniform supersampling

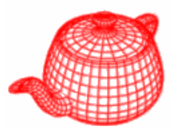

- Increasing the sampling rate moves each copy of the spectra further apart, potentially reducing the overlap and thus aliasing
- Resulting samples must be resampled (filtered) to image sampling rate

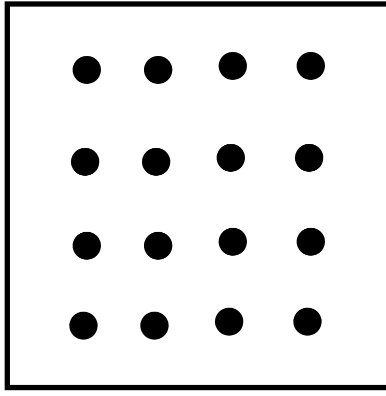

Samples

$$Pixel = \sum_{s} w_s \cdot Sample_s$$

Pixel

Point vs. Supersampled

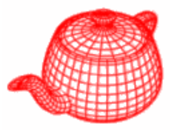

Point | **4x4 Supersampled**

Checkerboard sequence by Tom Duff

Analytic vs. Supersampled

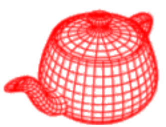

Exact Area

4x4 Supersampled

Non-uniform sampling

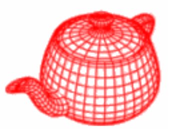

- Uniform sampling
  - The spectrum of uniformly spaced samples is also a set of uniformly spaced spikes
  - Multiplying the signal by the sampling pattern corresponds to placing a copy of the spectrum at each spike (in freq. space)
  - Aliases are coherent (structured), and very noticeable
- Non-uniform sampling
  - Samples at non-uniform locations have a different spectrum; a single spike plus noise
  - Sampling a signal in this way converts aliases into broadband noise
  - Noise is incoherent (structureless), and much less objectionable
- Aliases can't be removed, but can be made less noticeable.

Antialiasing (nonuniform sampling)

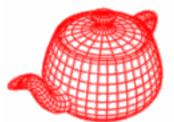

- The impulse train is modified as

$$\sum_{i=-\infty}^{\infty} \delta\left(x\text{-}\left(iT + \frac{1}{2} - \xi \right)\right)$$

- It turns regular aliasing into noise. But random noise is less distracting than coherent aliasing.

Jittered vs. Uniform Supersampling

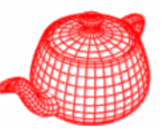

4x4 Jittered Sampling **4x4 Uniform**

Prefer noise over aliasing

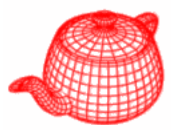

reference *aliasing* *noise*

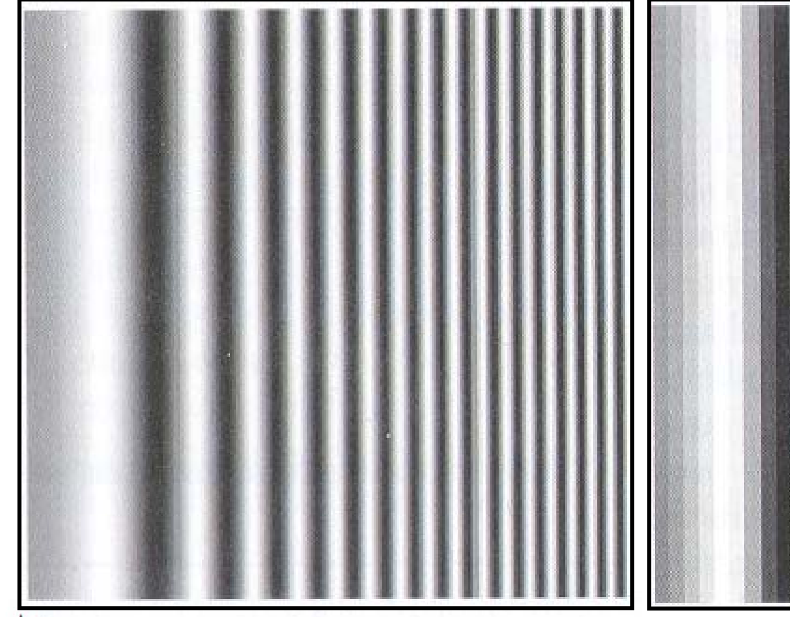

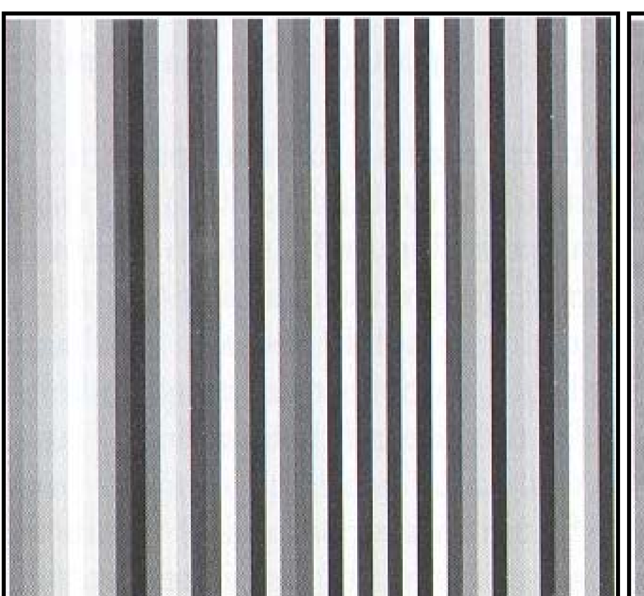

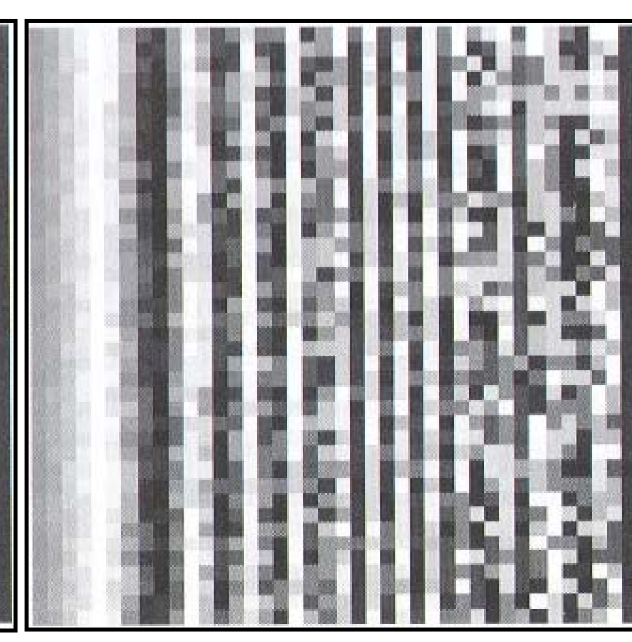

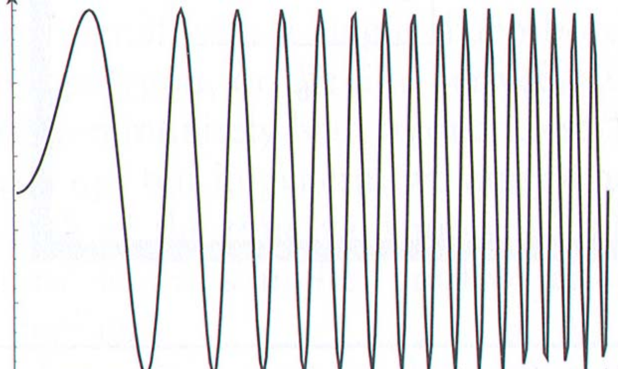

Jittered sampling

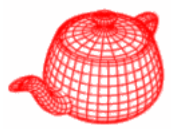

Add uniform random jitter to each sample

Poisson disk noise (Yellott)

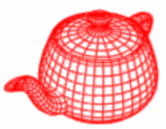

- Blue noise
- Spectrum should be noisy and lack any concentrated spikes of energy (to avoid coherent aliasing)
- Spectrum should have deficiency of low-frequency energy (to hide aliasing in less noticeable high frequency)

Distribution of extrafoveal cones

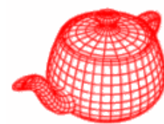

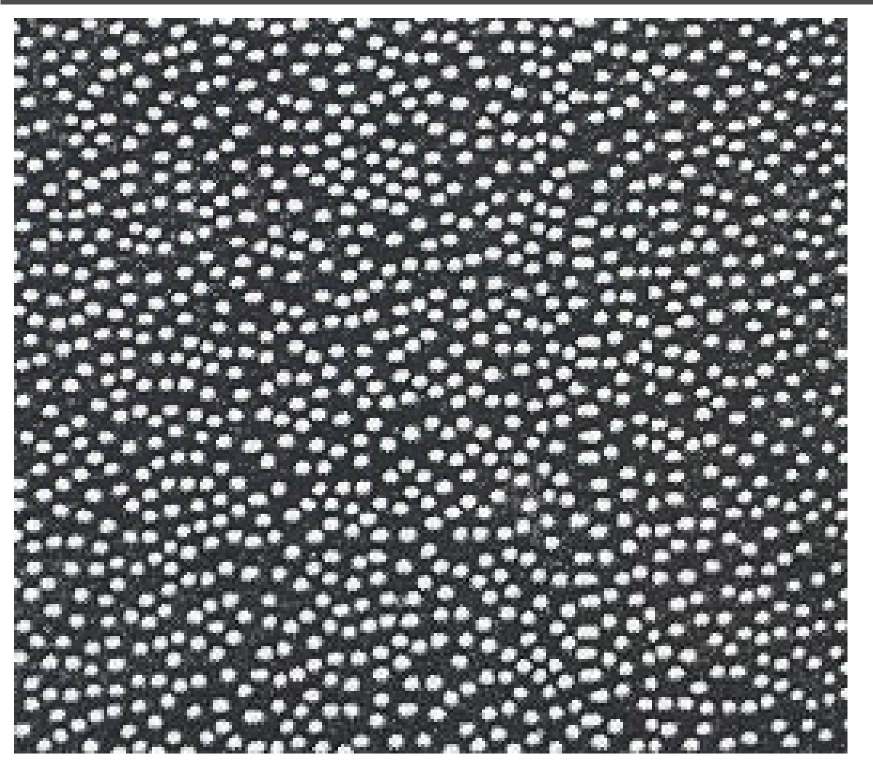

Monkey eye
cone distribution

Fourier transform

Yellott theory

- **Aliases replaced by noise**
- **Visual system less sensitive to high freq noise**

Example

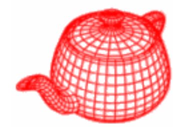

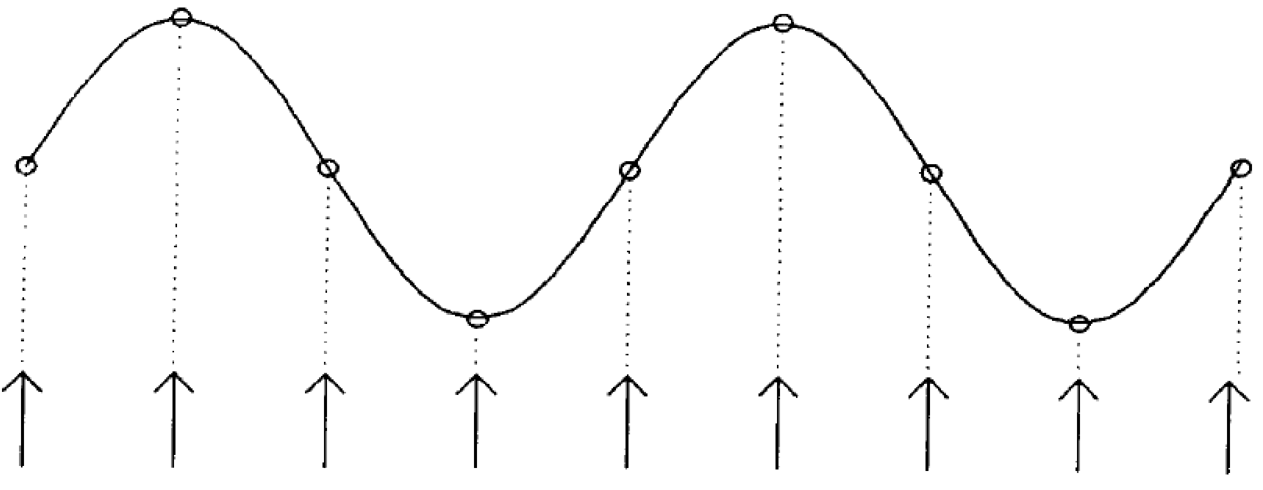

(a) Point sampling within the Nyquist limit

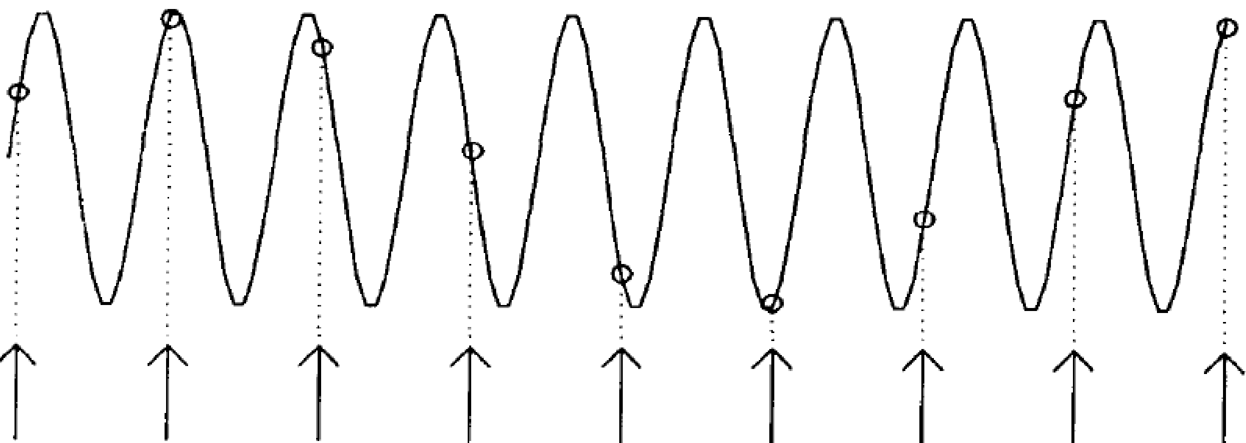

(b) Point sampling beyond the Nyquist limit

Aliasing

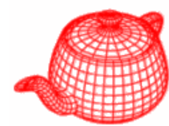

frequency domain

function (a)

(a) Original signal $F(x)$

(b) Sampling function $III(x)$

(c) Sampled signal $F(x) * III(x)$

(d) Ideal reconstruction filter $\Pi(x)$

(e) Final result

function (b)

(f) Original signal $F(x)$

(g) Sampling function $III(x)$

(h) Sampled signal $F(x) * III(x)$

(i) Ideal reconstruction filter $\Pi(x$

(j) Final result

alias=false frequency

Stochastic sampling

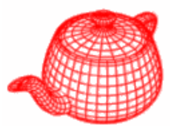

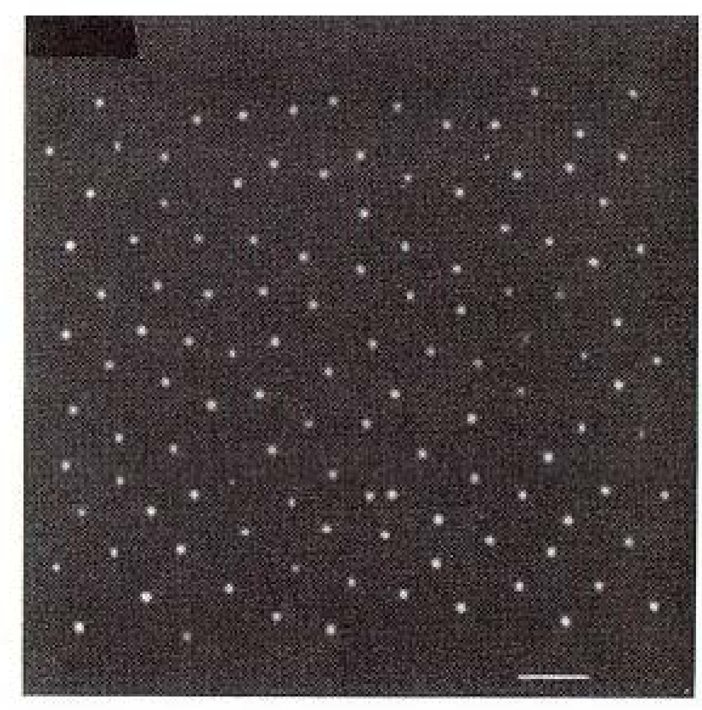

Fig. 3a. Monkey eye photoreceptor distribution.

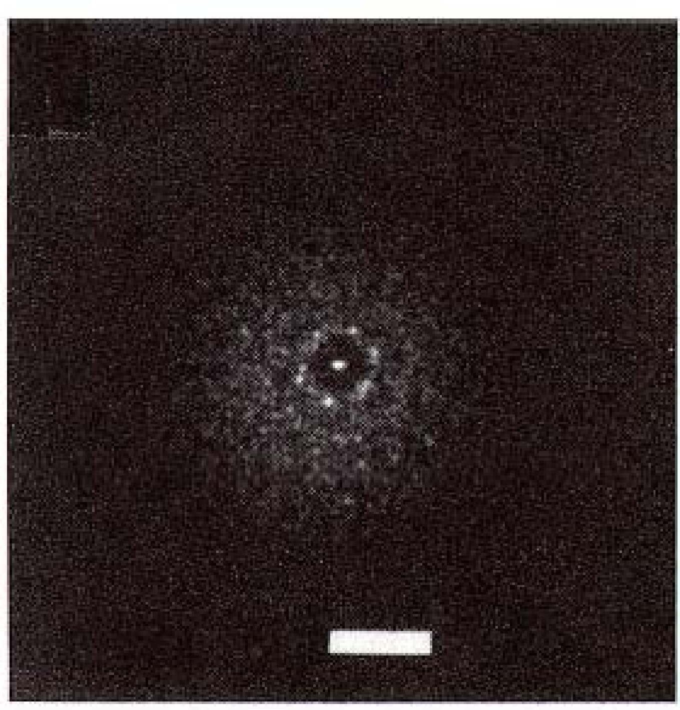

Fig. 3b. Optical transform of monkey eye.

Stochastic sampling

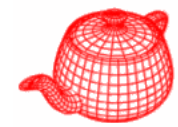

function (a)

(a) Original signal $F(x)$

(b) Sampling function $P(x)$

(c) Sampled signal $F(x) * P(x)$

(d) Ideal reconstruction filter $\Pi(x)$

(e) Final result

function (b)

(f) Original signal $F(x)$

(g) Sampling function $P(x)$

(h) Sampled signal $F(x) * P(x)$

(i) Ideal reconstruction filter $\Pi(x)$

(j) Final result

Replace structure alias by structureless (high-freq) noise

Antialiasing (adaptive sampling)

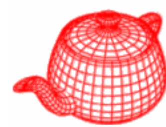

- Take more samples only when necessary. However, in practice, it is hard to know where we need supersampling. Some heuristics could be used.
- It only makes a less aliased image, but may not be more efficient than simple supersampling particular for complex scenes.

Application to ray tracing

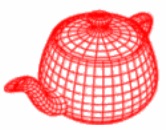

- Sources of aliasing: object boundary, small objects, textures and materials
- Good news: we can do sampling easily
- Bad news: we can't do prefiltering (because we do not have the whole function)
- Key insight: we can never remove all aliasing, so we develop techniques to mitigate its impact on the quality of the final image.

`pbrt` sampling interface

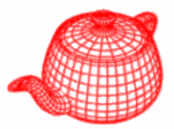

- Creating good sample patterns can substantially improve a ray tracer's efficiency, allowing it to create a high-quality image with fewer rays.
- Because evaluating radiance is costly, it pays to spend time on generating better sampling.
- **`core/sampling.*, samplers/*`**
- **`random.cpp, stratified.cpp, bestcandidate.cpp, lowdiscrepancy.cpp,`**

An ineffective sampler

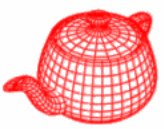

A more effective sampler

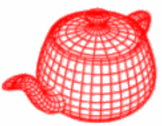

Main rendering loop for each task

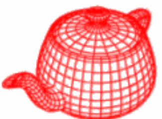

```
void SamplerRenderer::Run() {
  Sampler *sampler
    = mainSampler->GetSubSampler(taskNum, taskCount);
  ...
  // Allocate space for samples and intersections
  int maxSamples = sampler->MaximumSampleCount();
  Sample *samples=origSample->Duplicate(maxSamples);
  RayDifferential *rays=new RayDifferential[maxSamples];
  Spectrum *Ls = new Spectrum[maxSamples];
  Spectrum *Ts = new Spectrum[maxSamples];
  Intersection *isects = new Intersection[maxSamples]
  ...
```

Main rendering loop for each task

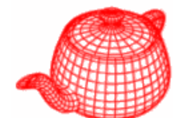

```
while ((sCnt=sampler->GetMoreSamples(samples,rng))>0){
  for (int i = 0; i < sCnt; ++i) {
    ...
    float rayWeight = camera->GenerateRayDifferential(
                          samples[i], &rays[i]);
    ...
    if (rayWeight > 0.f)
      Ls[i] = rayWeight * renderer->Li(scene, rays[i],
          &samples[i], rng, arena, &isects[i], &Ts[i]);
  ...
  } // end for
  if (sampler->ReportResults(samples, rays, Ls, …))
    for (int i = 0; i < sCnt; ++i) {
      ...
      camera->film->AddSample(samples[i], Ls[i]);
    }
} // end while
camera->film->UpdateDisplay(
  sampler->xPixelStart, sampler->yPixelStart,
  sampler->xPixelEnd+1, sampler->yPixelEnd+1);
```

Sampler

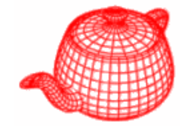

Generates a good pattern of multidimensional samples.

```
class Sampler {
  …
  virtual int GetMoreSamples(Sample *sample,
                 RNG &rng) = 0;
  virtual int MaximumSampleCount() = 0;
  virtual bool ReportResults(…);
  …
  const int xPixelStart, xPixelEnd;
  const int yPixelStart, yPixelEnd;
  const int samplesPerPixel;
  const float shutterOpen, shutterClose;
}
```

random number generator

for pre-allocating memory

report radiance for things like adaptive sampling

range of pixels

sample per pixel

Sampler

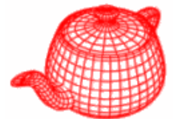

```
void Sampler::ComputeSubWindow(int num, int count,
 int *XStart, int *XEnd, int *YStart, int *YEnd) {
 int dx=xPixelEnd-xPixelStart,dy=yPixelEnd-yPixelStart;
 int nx = count, ny = 1;
 while ((nx & 0x1) == 0 && 2 * dx * ny < dy * nx) {
   nx >>= 1; ny <<= 1;
 }
 int xo = num % nx, yo = num / nx;
 float tx0=float(xo)/float(nx),tx1=float(xo+1)/float(nx);
 float ty0=float(yo)/float(ny),ty1=float(yo+1)/float(ny);
 *XStart = Floor2Int(Lerp(tx0,xPixelStart,xPixelEnd));
 *XEnd   = Floor2Int(Lerp(tx1,xPixelStart,xPixelEnd));
 *YStart = Floor2Int(Lerp(ty0,yPixelStart,yPixelEnd));
 *YEnd   = Floor2Int(Lerp(ty1,yPixelStart,yPixelEnd));
}
```

Sample

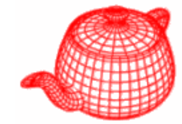

```
struct CameraSample {
  float imageX, imageY;
  float lensU, lensV;
  float time;
};
```

store required information for generating camera rays

store required information for one eye ray sample

```
struct Sample : public CameraSample {
  Sample(Sampler *sampler, SurfaceIntegrator *surf,
         VolumeIntegrator *vol, const Scene *scene);
  uint32_t Add1D(uint32_t num);
  uint32_t Add2D(uint32_t num);
  …
  // Sample Public Data
  vector<uint32_t> n1D, n2D;
  float **oneD, **twoD;
};
```

Note that it stores all samples required for one eye ray. That is, it may depend on depth.

Sample

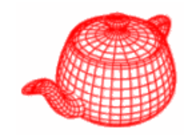

- Sample is allocated once in Render(). Sampler is called to fill in the information for each eye ray. The integrator can ask for multiple 1D and/or 2D samples, each with an arbitrary number of entries, e.g. depending on #lights. For example, WhittedIntegrator does not need samples. DirectLighting needs samples proportional to #lights.
- The structure of sample is initiated once and Sampler is responsible for filling in requested sample structure with well-behaved samples.

Data structure

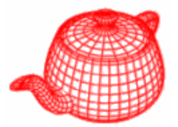

- Different types of lights require different numbers of samples, usually 2D samples.
- Sampling BRDF requires 2D samples.
- Selection of BRDF components requires 1D samples.

n1D 3 1 2 n2D 2 2 1 1 2 2 filled in by integrators

sample

oneD twoD allocate together to avoid cache miss

mem

1 2 3
lightNumOffset

1 3 5
lightSampleOffset

2 4 6
bsdfSampleOffset

PathIntegrator

Sample

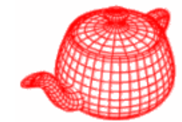

```
Sample::Sample(Sampler *sampler, SurfaceIntegrator
   *surf, VolumeIntegrator *vol, Scene *scene)
{
  if (surf) surf->RequestSamples(sampler,this,scene);
  if (vol) vol->RequestSamples(sampler, this, scene);
  AllocateSampleMemory();
}

void Sample::AllocateSampleMemory() {
  int nPtrs = n1D.size() + n2D.size();
  if (!nPtrs) {
    oneD = twoD = NULL; return;
  }
  oneD = AllocAligned<float *>(nPtrs);
  twoD = oneD + n1D.size();
```

Sample

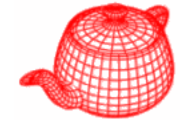

```
  int totSamples = 0;
  for (uint32_t i = 0; i < n1D.size(); ++i)
    totSamples += n1D[i];
  for (uint32_t i = 0; i < n2D.size(); ++i)
    totSamples += 2 * n2D[i];

  float *mem = AllocAligned<float>(totSamples);
  for (uint32_t i = 0; i < n1D.size(); ++i) {
    oneD[i] = mem; mem += n1D[i];
  }
  for (uint32_t i = 0; i < n2D.size(); ++i) {
    twoD[i] = mem; mem += 2 * n2D[i];
  }
}
```

PathIntegrator::RequestSamples

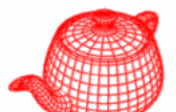

```
void PathIntegrator::RequestSamples(Sampler *sampler,
  Sample *sample, const Scene *scene)
{
 for (int i = 0; i < SAMPLE_DEPTH; ++i) {
   lightSampleOffsets[i]=LightSampleOffsets(1,sample);
  bsdfSampleOffsets[i]=BSDFSampleOffsets(1,sample);
   pathSampleOffsets[i]=BSDFSampleOffsets(1,sample);
 }
}
```

LightSampleOffsets

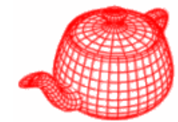

```
struct LightSampleOffsets {
  LightSampleOffsets(int count, Sample *sample);
  int nSamples, componentOffset, posOffset;
};

LightSampleOffsets::LightSampleOffsets(int count,
   Sample *sample) {
     nSamples = count;
     componentOffset = sample->Add1D(nSamples);
     posOffset = sample->Add2D(nSamples);
}
```

LightSampleOffsets $\longrightarrow$ LightSample

DirectLighting::RequestSamples

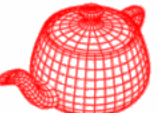

```
void DirectLightingIntegrator::RequestSamples(
  Sampler *sampler, Sample *sample, Scene *scene) {
 if (strategy == SAMPLE_ALL_UNIFORM) {
   uint32_t nLights = scene->lights.size();
   lightSampleOffsets=new LightSampleOffsets[nLights];
   bsdfSampleOffsets = new BSDFSampleOffsets[nLights];
   for (uint32_t i = 0; i < nLights; ++i) {
     const Light *light = scene->lights[i];
     int nSamples = light->nSamples;
     if (sampler) nSamples=sampler->RoundSize(nSamples);
     lightSampleOffsets[i]
       = LightSampleOffsets(nSamples, sample);
     bsdfSampleOffsets[i]
       = BSDFSampleOffsets(nSamples, sample);
   }
   lightNumOffset = -1;
 }
```

DirectLighting::RequestSamples

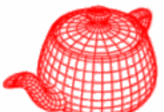

```
 else {
   lightSampleOffsets = new LightSampleOffsets[1];
   lightSampleOffsets[0]
     = LightSampleOffsets(1, sample);
   lightNumOffset = sample->Add1D(1);
   bsdfSampleOffsets = new BSDFSampleOffsets[1];
   bsdfSampleOffsets[0] = BSDFSampleOffsets(1, sample);
 }
}
```

Random sampler

Just for illustration; does not work well in practice

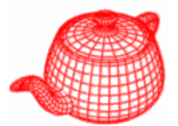

```
RandomSampler::RandomSampler(int xstart, int xend,
  int ystart, int yend, int ns,
  float sopen, float sclose) {
 xPos = xPixelStart;
 yPos = yPixelStart;
 nSamples = ns;
 imageSamples = AllocAligned<float>(5 * nSamples);
 lensSamples = imageSamples + 2 * nSamples;
 timeSamples = lensSamples + 2 * nSamples;
 // prepare samples for the first pixel
 RNG rng(xstart + ystart * (xend-xstart));
 for (int i = 0; i < 5 * nSamples; ++i)
   imageSamples[i] = rng.RandomFloat();
 for (int o = 0; o < 2 * nSamples; o += 2) {
   imageSamples[o]   += xPos;
   imageSamples[o+1] += yPos; }
 samplePos = 0;
}
```

private copy of the current pixel position

#samples consumed for current pixel

Random sampler

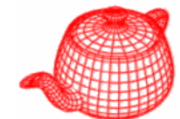

```
Sampler *RandomSampler::GetSubSampler(
                                      int num, int count)
{
  int x0, x1, y0, y1;
  ComputeSubWindow(num, count, &x0, &x1, &y0, &y1);
  if (x0 == x1 || y0 == y1) return NULL;
  return new RandomSampler(x0, x1, y0, y1, nSamples,
                                      shutterOpen, shutterClose);
}
```

Random sampler

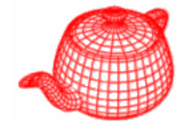

```
int RandomSampler::GetMoreSamples(
                    Sample *sample, RNG &rng) {
 if (samplePos==nSamples) {
   if(xPixelStart==xPixelEnd || yPixelStart==yPixelEnd)
            return 0;
   if (++xPos == xPixelEnd) {xPos=xPixelStart; ++yPos;}
   if (yPos == yPixelEnd) return 0;
```

generate all samples for one pixel at once

```
   for (int i = 0; i < 5 * nSamples; ++i)
     imageSamples[i] = rng.RandomFloat();
   for (int o = 0; o < 2 * nSamples; o += 2) {
     imageSamples[o]+=xPos; imageSamples[o+1]+=yPos;
   }
   samplePos = 0;
 }
```

Random sampler

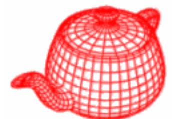

```
  sample->imageX = imageSamples[2*samplePos];
  sample->imageY = imageSamples[2*samplePos+1];
  sample->lensU = lensSamples[2*samplePos];
  sample->lensV = lensSamples[2*samplePos+1];
  sample->time = Lerp(timeSamples[samplePos],
                      shutterOpen, shutterClose);
  for (uint32_t i = 0; i < sample->n1D.size(); ++i)
    for (uint32_t j = 0; j < sample->n1D[i]; ++j)
      sample->oneD[i][j] = rng.RandomFloat();
  for (uint32_t i = 0; i < sample->n2D.size(); ++i)
    for (uint32_t j = 0; j < 2*sample->n2D[i]; ++j)
      sample->twoD[i][j] = rng.RandomFloat();
  ++samplePos;
  return 1;
}
```

Random sampling

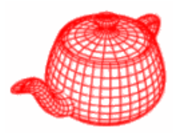

a pixel

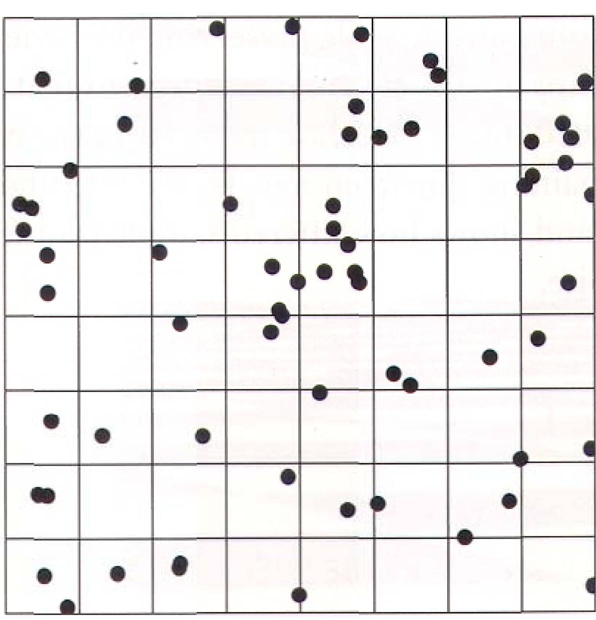

completely random

Stratified sampling

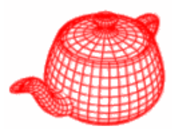

- Subdivide the sampling domain into non-overlapping regions (*strata*) and take a single sample from each one so that it is less likely to miss important features.

Stratified sampling

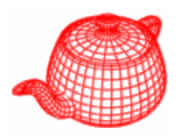

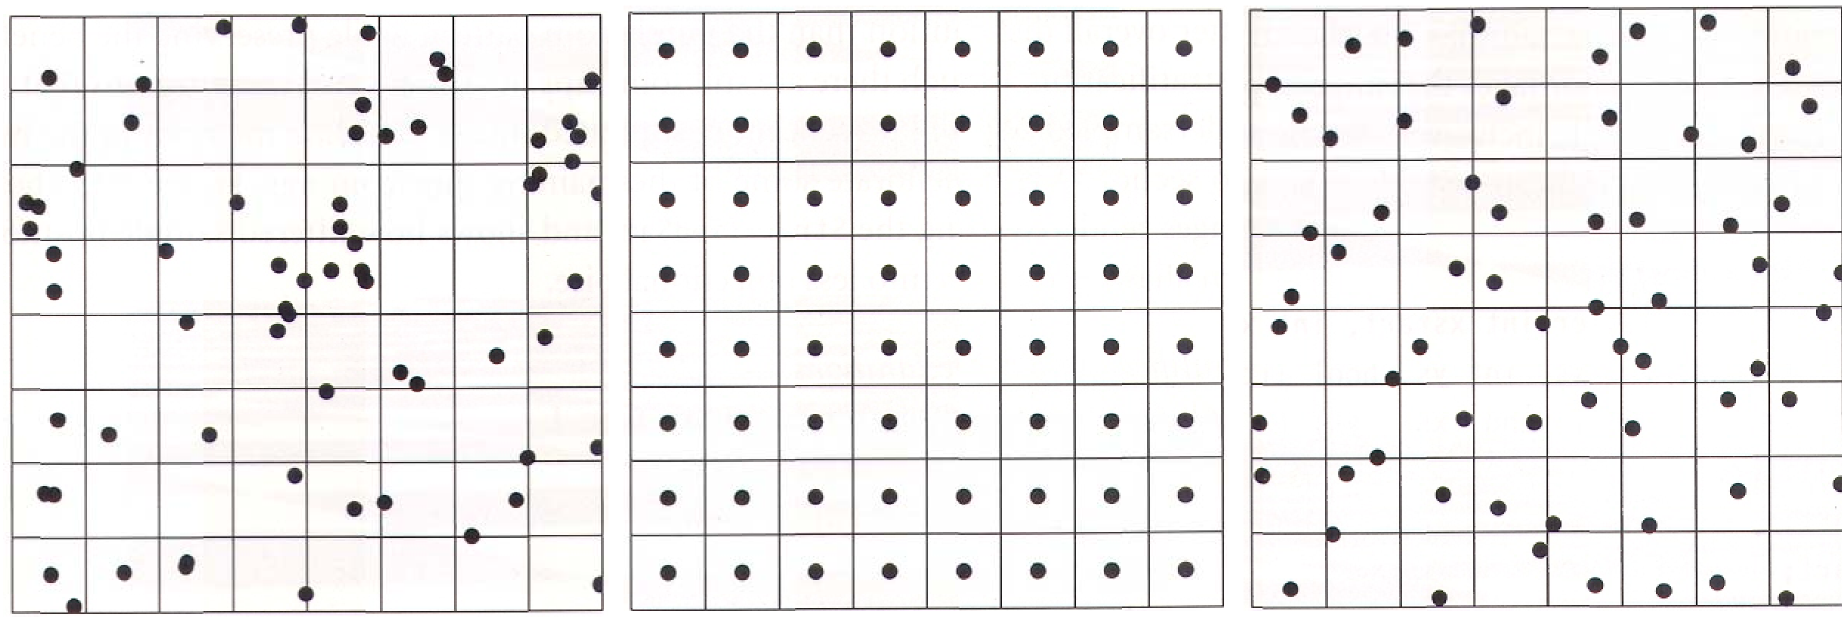

completely random

stratified uniform

stratified jittered

turns aliasing into noise

Comparison of sampling methods

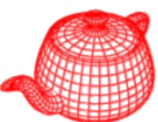

256 samples per pixel as reference

1 sample per pixel (no jitter)

Comparison of sampling methods

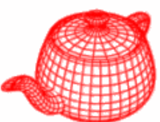

1 sample per pixel (jittered)

4 samples per pixel (jittered)

Stratified sampling

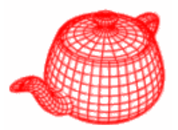

reference

random

stratified jittered

High dimension

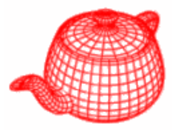

- D dimension means N^D cells.
- Solution: make strata separately and associate them randomly, also ensuring good distributions.

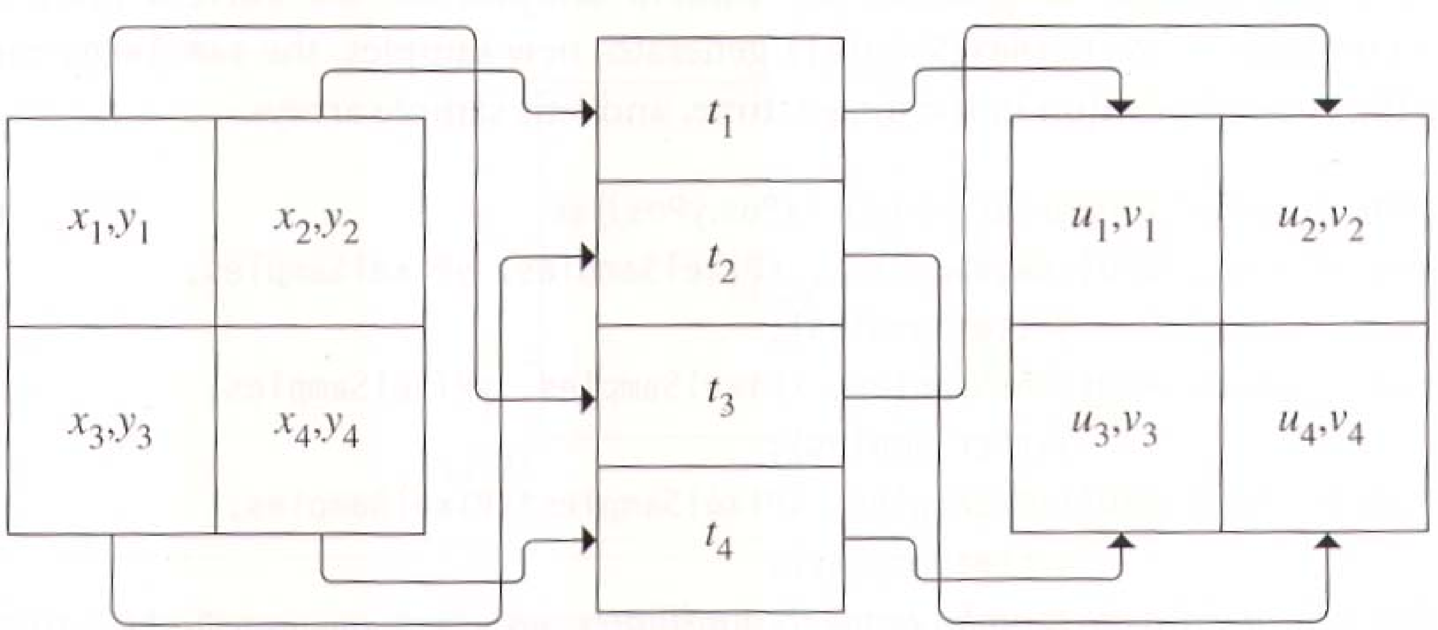

StratifiedSampler::GetMoreSamples

```
if (yPos == yPixelEnd) return 0;
int nSamples = xPixelSamples * yPixelSamples;

// Generate initial stratified samples
float *bufp = sampleBuf;
float *imageSamples = bufp; bufp += 2 * nSamples;
float *lensSamples = bufp;  bufp += 2 * nSamples;
float *timeSamples = bufp;
StratifiedSample2D(imageSamples, xPixelSamples,
                 yPixelSamples, rng, jitterSamples);
StratifiedSample2D(lensSamples, xPixelSamples,
                 yPixelSamples, rng, jitterSamples);
StratifiedSample1D(timeSamples, xPixelSamples *
                 yPixelSamples, rng, jitterSamples);

for (int o=0;o<2*xPixelSamples*yPixelSamples;o+=2){
  imageSamples[o]+=xPos; imageSamples[o+1]+=yPos;
}
```

StratifiedSampler::GetMoreSamples

```
Shuffle(lensSamples,xPixelSamples*yPixelSamples,2,rng);
Shuffle(timeSamples,xPixelSamples*yPixelSamples,1,rng);
for (int i = 0; i < nSamples; ++i) {
  samples[i].imageX = imageSamples[2*i];
  samples[i].imageY = imageSamples[2*i+1];
  samples[i].lensU = lensSamples[2*i];
  samples[i].lensV = lensSamples[2*i+1];
  samples[i].time = Lerp(timeSamples[i],
                         shutterOpen, shutterClose);
  for (uint32_t j = 0; j < samples[i].n1D.size(); ++j)
    LatinHypercube(samples[i].oneD[j],
                   samples[i].n1D[j], 1, rng);
  for (uint32_t j = 0; j < samples[i].n2D.size(); ++j)
    LatinHypercube(samples[i].twoD[j],
                   samples[i].n2D[j], 2, rng);
}
if (++xPos == xPixelEnd) {xPos = xPixelStart; ++yPos;}
return nSamples;
}
```

Stratified sampling

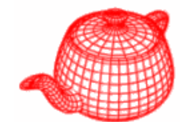

n stratified samples within [0..1]

```
void StratifiedSample1D(float *samp, int nSamples,
                        RNG &rng, bool jitter) {
   float invTot = 1.f / nSamples;
   for (int i = 0;  i < nSamples; ++i) {
     float delta = jitter ? rng.RandomFloat() : 0.5f;
    *samp++ = min((i+delta)*invTot, OneMinusEpsilon);
   }
}
```

nx*ny stratified samples within [0..1]X[0..1]

```
void StratifiedSample2D(float *samp, int nx, int ny,
                        RNG &rng, bool jitter) {
   float dx = 1.f / nx, dy = 1.f / ny;
   for (int y = 0; y < ny; ++y)
     for (int x = 0; x < nx; ++x) {
       float jx = jitter ? rng.RandomFloat() : 0.5f;
       float jy = jitter ? rng.RandomFloat() : 0.5f;
       *samp++ = min((x + jx) * dx, OneMinusEpsilon);
       *samp++ = min((y + jy) * dy, OneMinusEpsilon);
     }
}
```

Shuffle

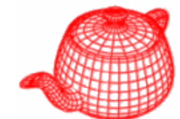

```
template <typename T>
void Shuffle(T *samp, int count, int dims, RNG &rng)
{
  for (int i = 0; i < count; ++i) {
    u_int other = i+(rng.RandomUInt()%(count-i));
    for (int j = 0; j < dims; ++j)
      swap(samp[dims*i + j], samp[dims*other + j]);
  }
}
```

d-dimensional vector swap

Latin hypercube sampling

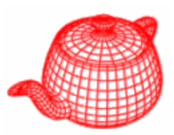

- Integrators could request an arbitrary n samples. nx1 or 1xn doesn't give a good sampling pattern.

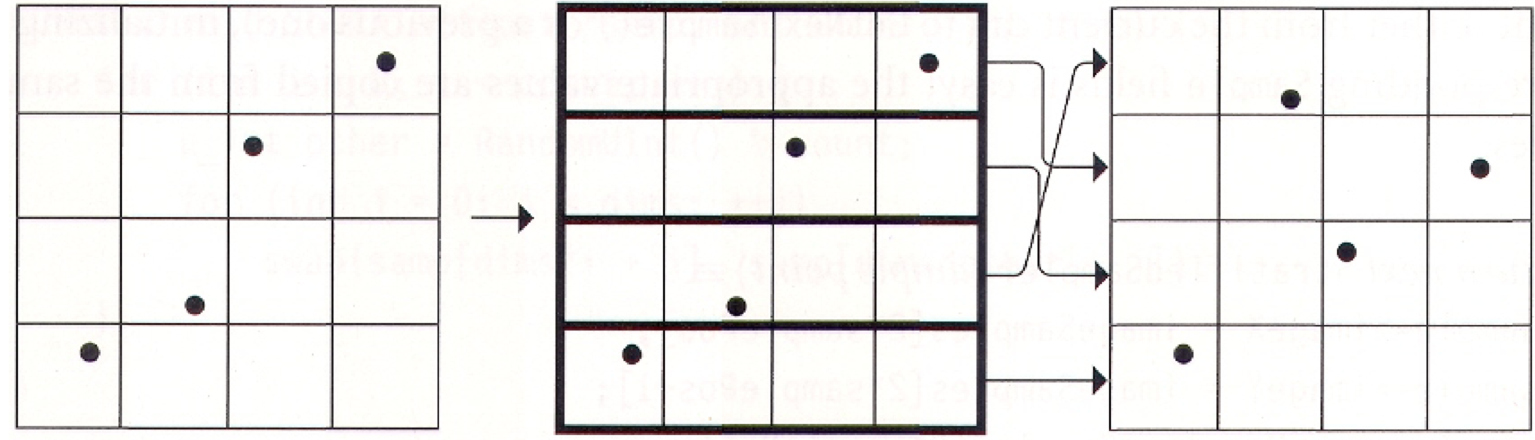

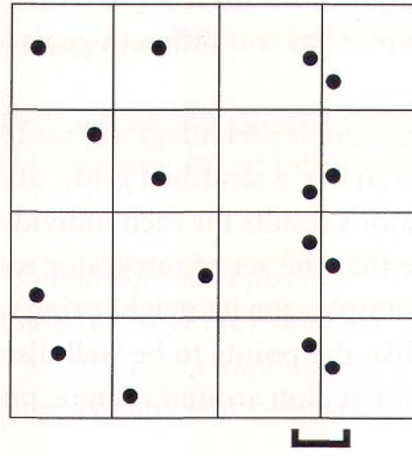

A worst case for stratified sampling

LHS can prevent this to happen

Latin Hypercube

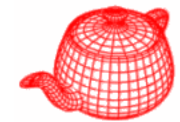

```
void LatinHypercube(float *samples,
                    int nSamples, int nDim, RNG &rng)
{
  // Generate LHS samples along diagonal
  float delta = 1.f / nSamples;
  for (int i = 0; i < nSamples; ++i)
    for (int j = 0; j < nDim; ++j)
      samples[nDim*i+j] = min((i+(rng.RandomFloat()))
                              *delta, OneMinusEpsilon);
  // Permute LHS samples in each dimension
  for (int i = 0; i < nDim; ++i) {
    for (int j = 0; j < nSamples; ++j) {
      u_int other=j+(rng.RandomUInt() % (nSamples-j));
      swap(samples[nDim * j + i],
           samples[nDim * other + i]);
      }
  }
}
```

note the difference with **shuffle**

Stratified sampling

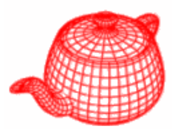

Stratified sampling

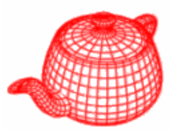

This is better because **`StratifiedSampler`** could generate a good LHS pattern for this case

1 camera sample and 16 shadow samples per pixel

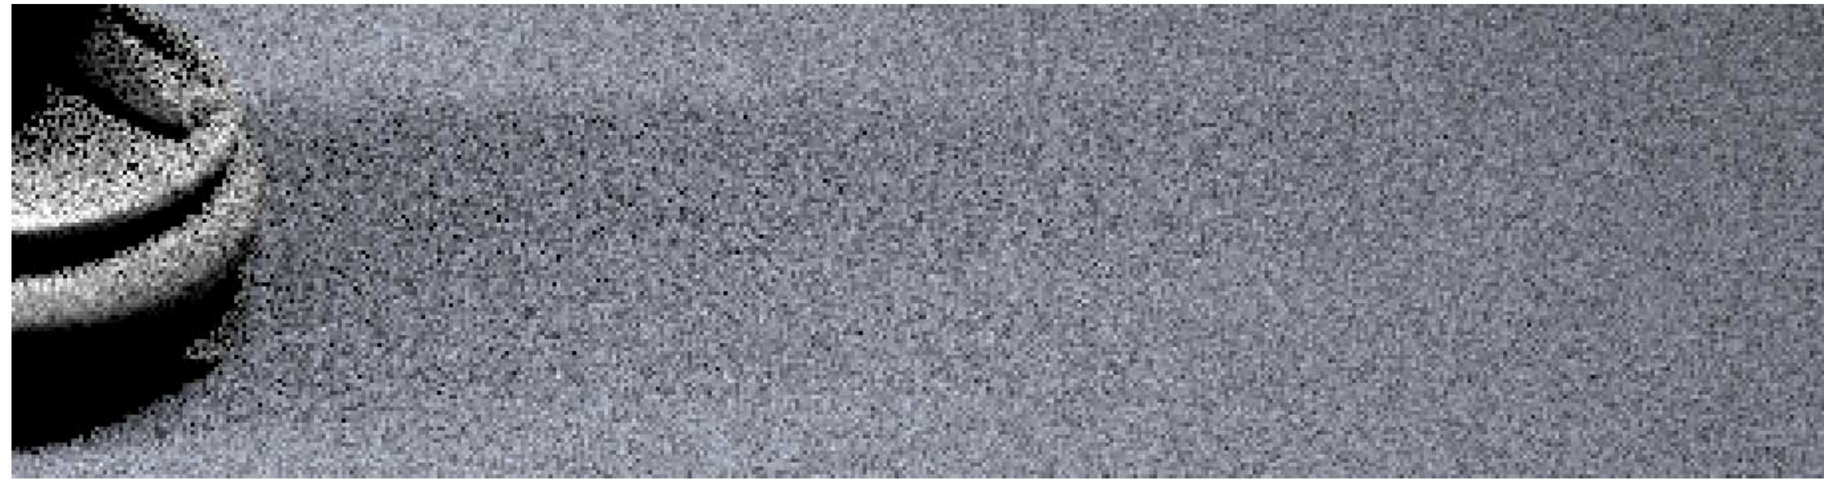

16 camera samples and each with 1 shadow sample per pixel

Low discrepancy sampling

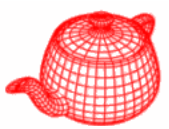

- A possible problem with stratified sampling

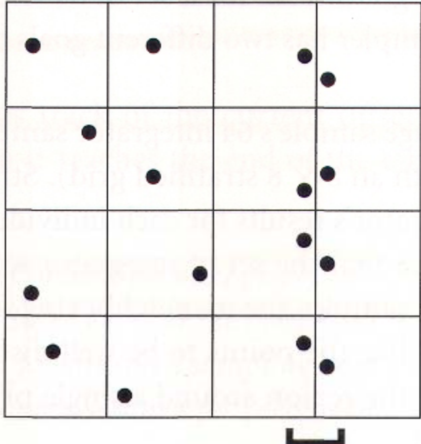

- Discrepancy can be used to evaluate the quality of patterns

Low discrepancy sampling

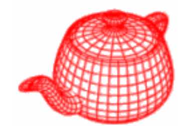

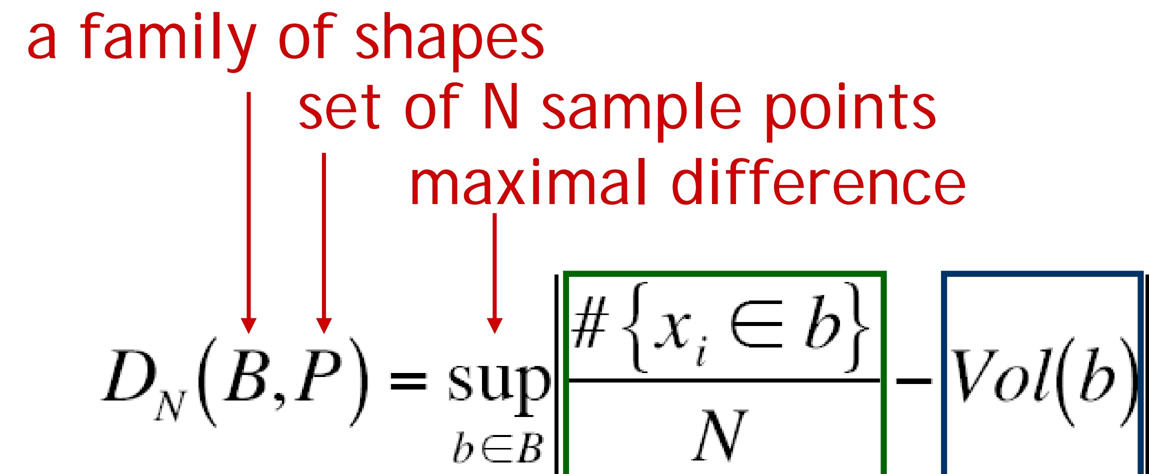

$$D_N(B,P) = \sup_{b \in B} \left| \frac{\#\{x_i \in b\}}{N} - Vol(b) \right|$$

volume estimated by sample number

real volume

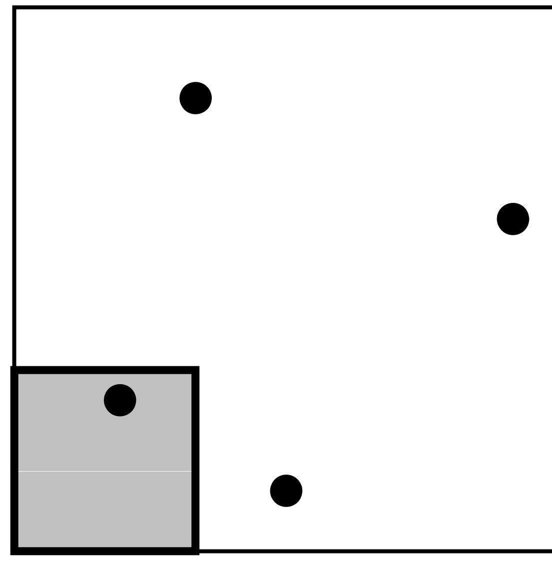

When B is the set of AABBs with a corner at the origin, this is called star discrepancy $D_N^*(P)$

1D discrepancy

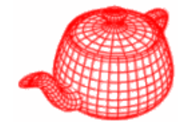

$$x_i = \frac{i}{N} \quad \Rightarrow \quad D_N^*(x_1,\ldots,x_n) = \frac{1}{N}$$

$$x_i = \frac{i-0.5}{N} \quad \Rightarrow \quad D_N^*(x_1,\ldots,x_n) = \frac{1}{2N}$$

$$x_i = general \quad \Rightarrow \quad D_N^*(x_1,\ldots,x_n) = \frac{1}{2N} + \max_{1\le i\le N}\left|x_i - \frac{2i-1}{2N}\right|$$

Uniform is optimal! However, we have learnt that irregular patterns are perceptually superior to uniform samples. Fortunately, for higher dimension, the low-discrepancy patterns are less uniform and works reasonably well as sample patterns in practice.
Next, we introduce methods specifically designed for generating low-discrepancy sampling patterns.

Radical inverse

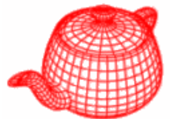

- A positive number n can be expressed in a base b as

$$n = a_k \ldots a_2 a_1 = a_1 b^0 + a_2 b^1 + a_3 b^2 + \ldots$$

- A radical inverse function in base b converts a nonnegative integer n to a floating-point number in *[0,1)*

$$\Phi_b(n) = 0.a_1 a_2 \ldots a_k = a_1 b^{-1} + a_2 b^{-2} + a_3 b^{-3} + \ldots$$

```
inline double RadicalInverse(int n, int base) {
  double val = 0;
  double invBase = 1. / base, invBi = invBase;
  while (n > 0) {
    int d_i = (n % base);
    val += d_i * invBi;
    n /= base;
    invBi *= invBase;
  }
  return val;
}
```

$a_k \ldots a_2 a_1$ → $0.a_1 a_2 \ldots a_k$

van der Corput sequence

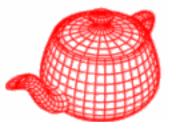

- The simplest sequence $x_i = \Phi_2(i)$
- Recursively split 1D line in half, sample centers
- Achieve minimal possible discrepancy

$$D_N^*(P) = O\left(\frac{\log N}{N}\right)$$

0 4 2 1 3

i	binary form of i	radical inverse	x_i
0	0	0.0	0
1	1	0.1	0.5
2	10	0.01	0.25
3	11	0.11	0.75
4	100	0.001	0.125
5	101	0.101	0.625
6	110	0.011	0.375

High-dimensional sequence

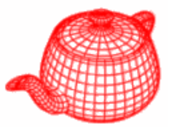

- Two well-known low-discrepancy sequences
  - Halton
  - Hammersley

Halton sequence

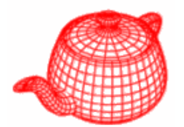

- Use relatively prime numbers as bases for each dimension

recursively split the dimension into p_d parts, sample centers

$$x_i = (\Phi_2(i), \Phi_3(i), \Phi_5(i), \ldots, \Phi_{p_d}(i))$$

- Achieve best possible discrepancy for N-D

$$D_N^*(P) = O\left(\frac{(\log N)^d}{N}\right)$$

- Can be used if N is not known in advance
- All prefixes of a sequence are well distributed so as additional samples are added to the sequence, low discrepancy will be maintained

Hammersley sequence

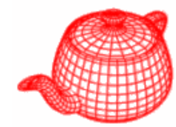

- Similar to Halton sequence.
- Slightly better discrepancy than Halton.
- Needs to know N in advance.

$$x_i = \left(\frac{i-1/2}{N}, \Phi_{b_1}(i), \Phi_{b_2}(i), \ldots, \Phi_{b_{d-1}}(i)\right)$$

Folded radical inverse

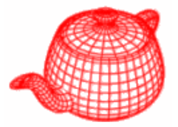

- Add the offset i to the ith digit d_i and take the modulus b.

$$\Phi_b(n) = \sum_{i=1}^{\infty} ((a_i + i - 1) \bmod b) \frac{1}{b^i}$$

- It can be used to improve Hammersley and Halton, called Hammersley-Zaremba and Halton-Zaremba.

Radial inverse

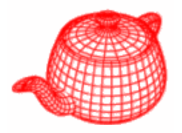

Halton

Hammersley

Better for that there are fewer clumps.

Folded radial inverse

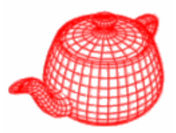

Halton

Hammersley

The improvement is more obvious

Low discrepancy sampling

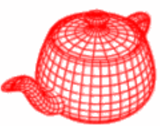

stratified jittered, 1 sample/pixel

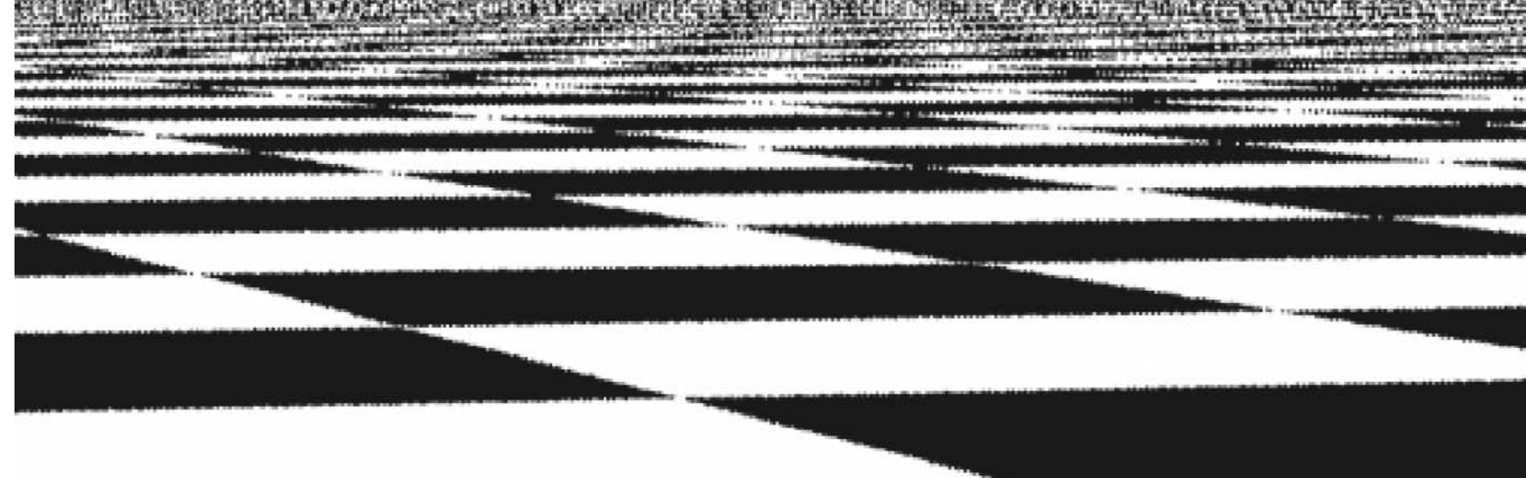

Hammersley sequence, 1 sample/pixel

Best candidate sampling

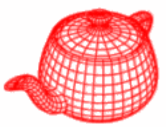

- Stratified sampling doesn't guarantee good sampling across pixels.
- *Poisson disk pattern* addresses this issue. The Poisson disk pattern is a group of points with no two of them closer to each other than some specified distance.
- It can be generated by *dart throwing*. It is time-consuming.
- *Best-candidate* algorithm by Dan Mitchell. It randomly generates many candidates but only inserts the one farthest to all previous samples.

Best candidate sampling

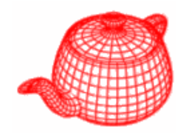

stratified jittered

best candidate

It avoids holes and clusters.

Best candidate sampling

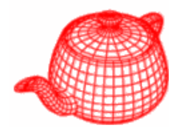

- Because of it is costly to generate best candidate pattern, pbrt computes a “tilable pattern” offline (by treating the square as a rolled torus).
- **`tools/samplepat.cpp`**

  → **`sampler/bestcandidate.out`**

Best candidate sampling

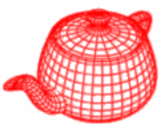

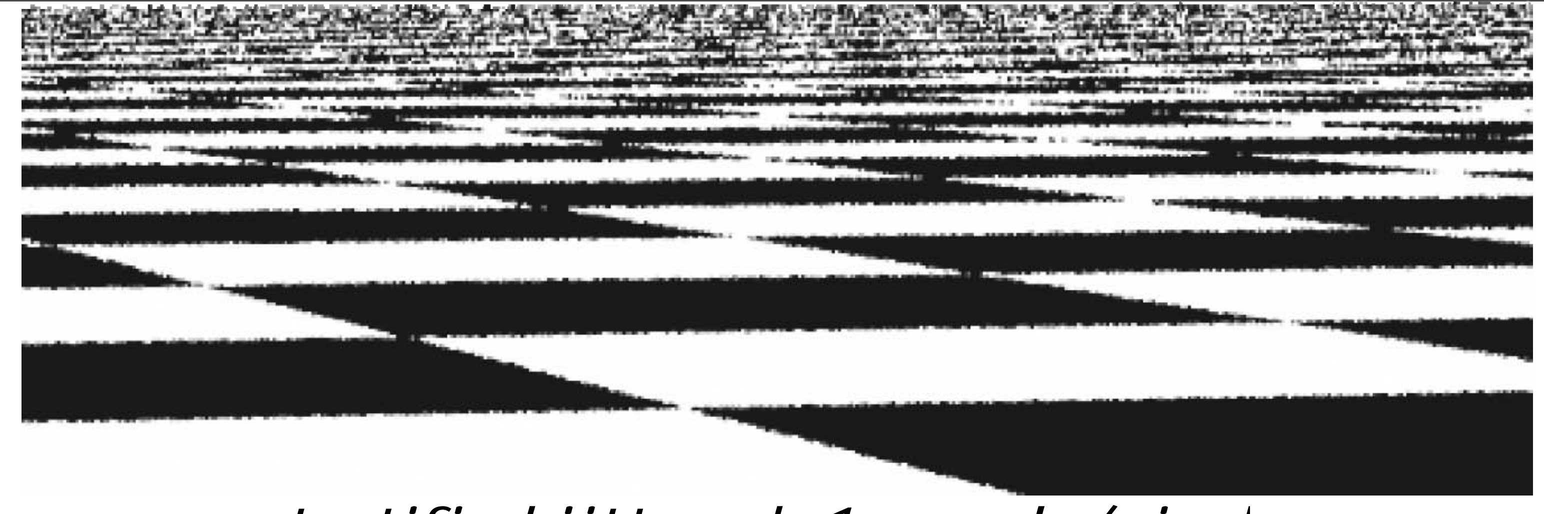

stratified jittered, 1 sample/pixel

best candidate, 1 sample/pixel

Best candidate sampling

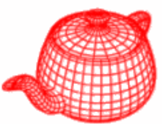

stratified jittered, 4 sample/pixel

best candidate, 4 sample/pixel

Comparisons

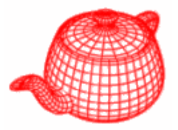

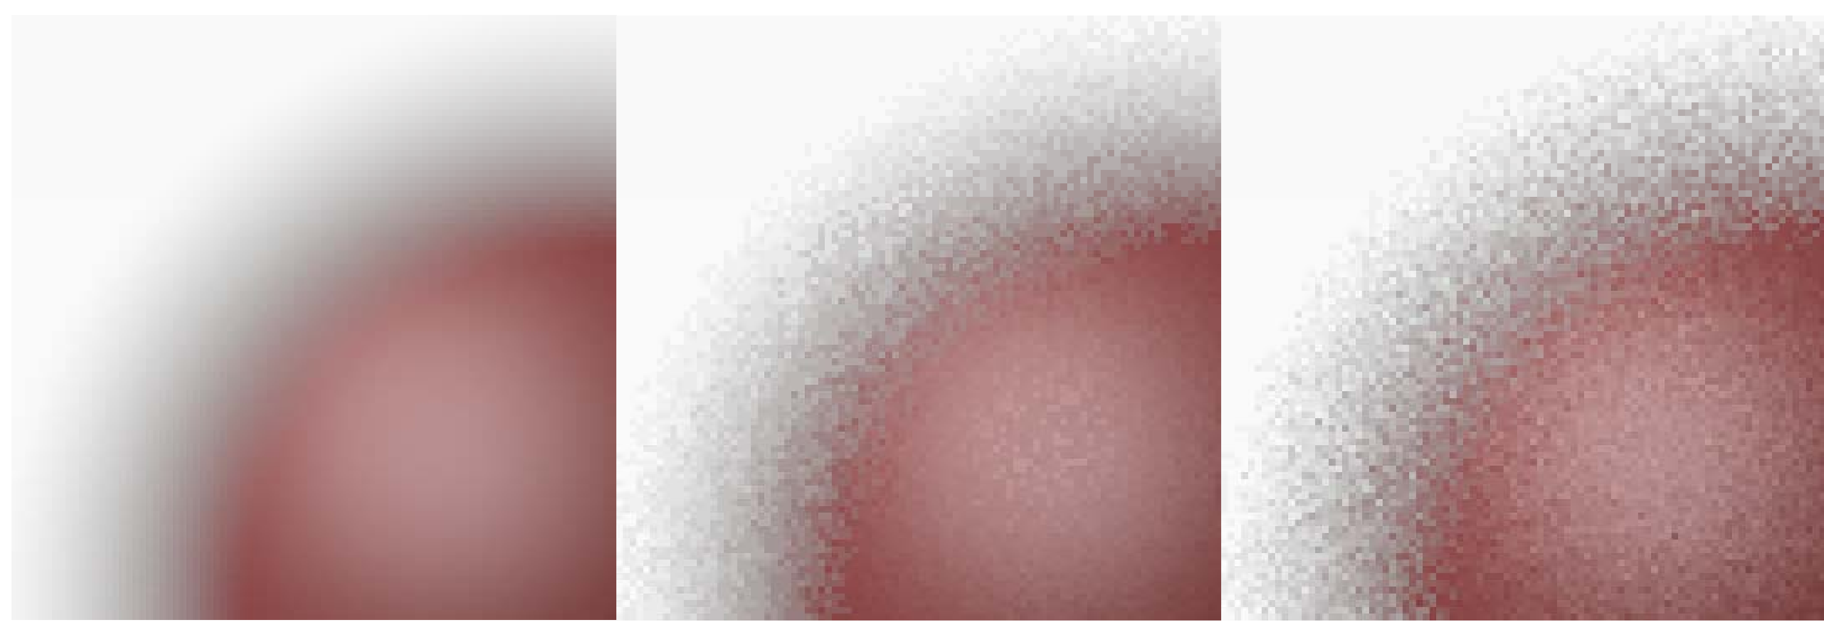

reference *low-discrepancy* *best candidate*

Adaptive sampling

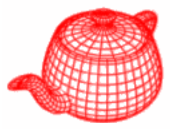

- More efficiently generate high-quality images by adding extra samples in parts of the image that are more complex than others.
- pbrt supports two kinds of simple refinement criteria: (1) to check to see if different shapes are intersected by different samples, indicating a likely geometric discontinuity and (2) to check for excessive contrast between the colors of different samples.

Main rendering loop for each task

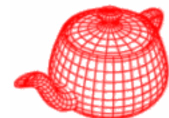

```
while ((sCnt=sampler->GetMoreSamples(samples,rng))>0){
  for (int i = 0; i < sCnt; ++i) {
    ...
    float rayWeight = camera->GenerateRayDifferential(
                           samples[i], &rays[i]);
    ...
    if (rayWeight > 0.f)
      Ls[i] = rayWeight * renderer->Li(scene, rays[i],
          &samples[i], rng, arena, &isects[i], &Ts[i]);
  ...
  } // end for
  if (sampler->ReportResults(samples, rays, Ls, …))
    for (int i = 0; i < sCnt; ++i) {
      ...
      camera->film->AddSample(samples[i], Ls[i]);
    }
} // end while
camera->film->UpdateDisplay(
  sampler->xPixelStart, sampler->yPixelStart,
  sampler->xPixelEnd+1, sampler->yPixelEnd+1);
```

AdaptiveSampler

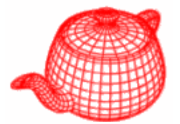

```
class AdaptiveSampler : public Sampler {
  …
  int xPos, yPos; current position
  int minSamples, maxSamples; least and max number
  float *sampleBuf;                of samples

  enum AdaptiveTest { ADAPTIVE_COMPARE_SHAPE_ID,
                      ADAPTIVE_CONTRAST_THRESHOLD };
  AdaptiveTest method; which criterion to use
  bool supersamplePixel;
};      whether the current pixel
        needs extra samples
```

AdaptiveSampler

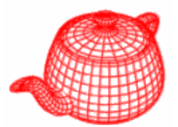

1. **supersamplePixel** is set to **false** initially
2. The initial set of **minSamples** is generated by **GetMoreSamples**.
3. Rendering loop evaluates these samples and report them back to **ReportResults**.
4. If more samples are needed, **ReportResults** sets to **true** and leave **(xPos, yPos)** unchanged.
5. The next call to **GetMoreSamples** generates a new set of **maxSamples** samples
6. When more samples are not needed or **maxSamples** samples have been used, **supersamplePixel** is set **false** and **(xPos, yPos)** is advanced.

AdaptiveSampler

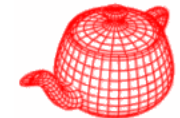

```
int AdaptiveSampler::GetMoreSamples(…) {
  …
  if (supersamplePixel) {
    LDPixelSample(xPos, yPos, shutterOpen,
      shutterClose, maxSamples, samples,
      sampleBuf, rng);
    return maxSamples;
  } else {
    if (yPos == yPixelEnd) return 0;
    LDPixelSample(xPos, yPos, shutterOpen,
      shutterClose, minSamples, samples,
      sampleBuf, rng);
    return minSamples;
  }
}
```

AdaptiveSampler

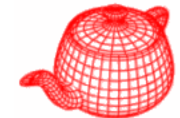

```
bool AdaptiveSampler::ReportResults(…) {
  if (supersamplePixel) {
    supersamplePixel = false;
    if (++xPos == xPixelEnd) {
      xPos = xPixelStart; ++yPos; }
    return true;
  } else if (needsSupersampling(…)) {
    supersamplePixel = true;
    return false;
  } else {
    if (++xPos == xPixelEnd) {
      xPos = xPixelStart; ++yPos;
    }
    return true;
  }
}
```

needsSupersampling

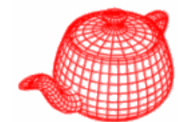

```
bool AdaptiveSampler::needsSupersampling(
  Sample *samples, const RayDifferential *rays,
  const Spectrum *Ls, const Intersection *isects,
  int count)
{
 switch (method) {
  case ADAPTIVE_COMPARE_SHAPE_ID:
   for (int i = 0; i < count-1; ++i)
    if (isects[i].shapeId != isects[i+1].shapeId ||
     isects[i].primitiveId != isects[i+1].primitiveId)
    return true;
   return false;
```

Efficient but fails to capture cases like (1) coplanar triangles with different ids but without edges (2) a parametric patch Could fold over and need more samples (3) shadows, textures …

needsSupersampling

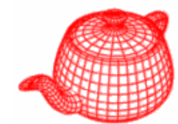

```
  case ADAPTIVE_CONTRAST_THRESHOLD:
   float Lavg = 0.f;
   for (int i = 0; i < count; ++i)
     Lavg += Ls[i].y();
   Lavg /= count;
   const float maxContrast = 0.5f;
   for (int i = 0; i < count; ++i)
     if (fabsf(Ls[i].y() - Lavg) / Lavg > maxContrast)
       return true;
   return false;
 }
 return false;
}
```

Not always successful. An example is ImageTexture which has been filtered For antialiasing. Even if the samples have high contrast, it probably does Not need more samples.

Adaptive sampling

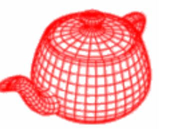

Adaptive sampling (geometry)

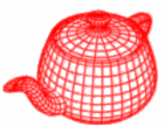

Adaptive sampling (contrast)

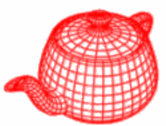

Reconstruction filters

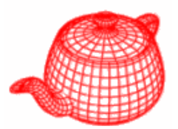

- Given the *chosen* image samples, we can do the following to compute pixel values.
  1. reconstruct a continuous function L′ from samples
  2. prefilter L′ to remove frequency higher than Nyquist limit
  3. sample L′ *at pixel locations*
- Because we will only sample L′ at pixel locations, we do not need to explicitly reconstruct L′s. Instead, we combine the first two steps.

Reconstruction filters

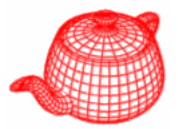

- Ideal reconstruction filters do not exist because of discontinuity in rendering. We choose nonuniform sampling, trading off noise for aliasing. There is no theory about ideal reconstruction for nonuniform sampling yet.
- Instead, we consider an interpolation problem

filter — sampled radiance

$$I(x,y)=\frac{\sum_i f(x-x_i,y-y_i)L(x_i,y_i)}{\sum_i f(x-x_i,y-y_i)}$$

final value

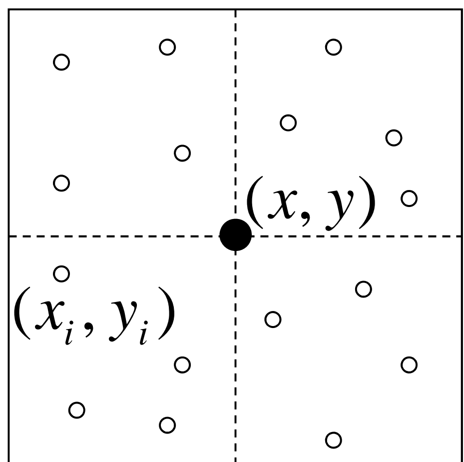

Filter

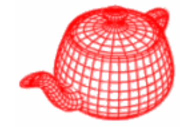

- provides an interface to $f(x,y)$
- **`Film`** stores a pointer to a filter and use it to filter the output before writing it to disk.

width, half of support

```
Filter::Filter(float xw, float yw)
float Evaluate(float x, float y);
```

$f(x, y)$

x, y is guaranteed to be within the range; range checking is not necessary

- **`filters/* (box, gaussian, mitchell, sinc, triangle)`**

Box filter

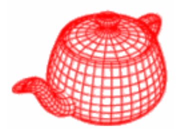

- Most commonly used in graphics. It's just about the worst filter possible, incurring postaliasing by high-frequency leakage.

```
float BoxFilter::Evaluate(float x, float y)
{
  return 1.;
}
```

no need to normalize since the weighted sum is divided by the total weight later.

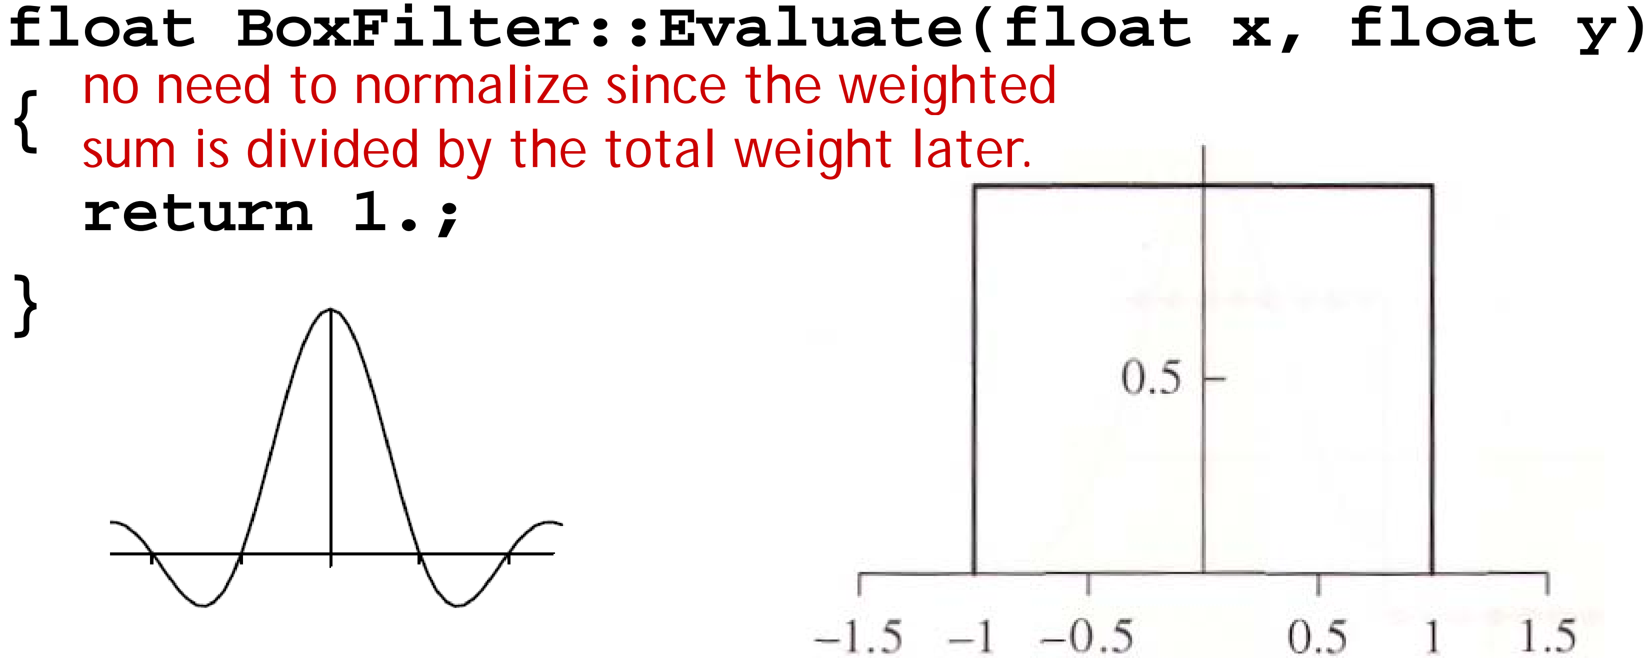

Sinc

Triangle filter

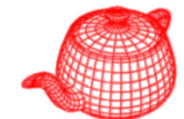

```
float TriangleFilter::Evaluate(float x, float y)
{
  return max(0.f, xWidth-fabsf(x)) *
         max(0.f, yWidth-fabsf(y));
}
```

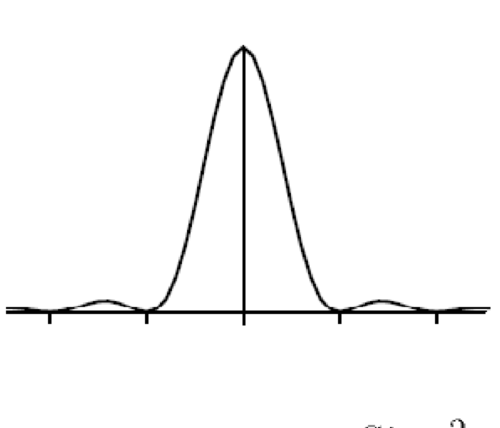

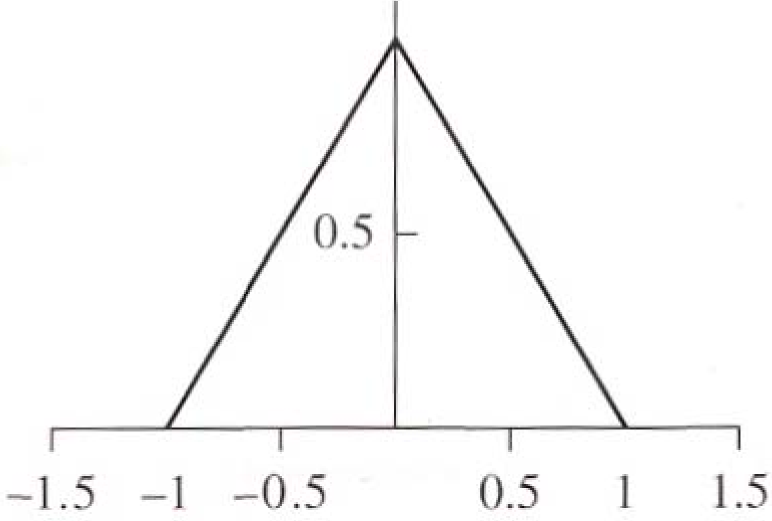

Gaussian filter

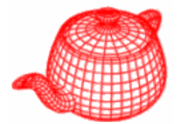

- Gives reasonably good results in practice

```
float GaussianFilter::Evaluate(float x, float y)
{
  return Gaussian(x, expX)*Gaussian(y, expY);
}
```

Gaussian essentially has a infinite support; to compensate this, the value at the end is calculated and subtracted.

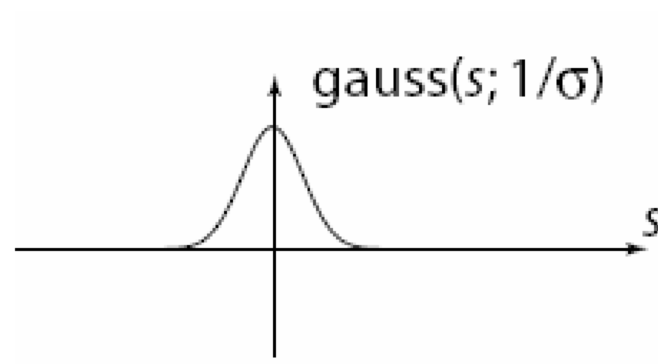

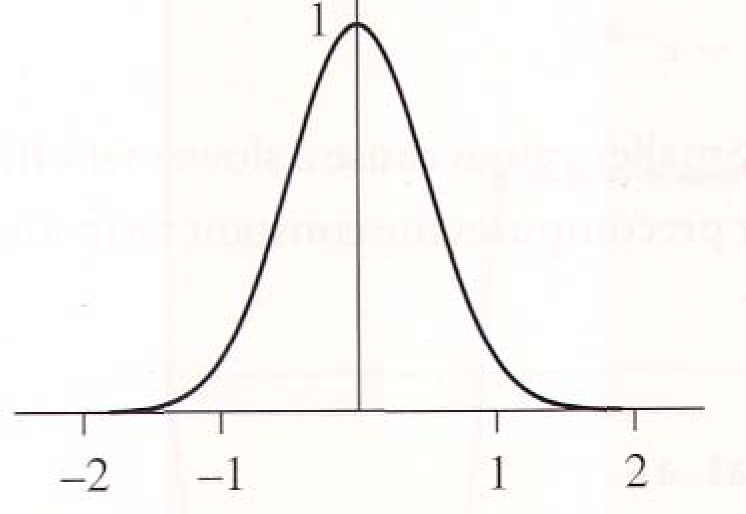

Mitchell filter

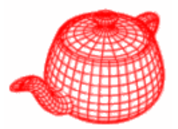

- parametric filters, tradeoff between ringing and blurring
- Negative lobes improve sharpness; ringing starts to enter the image if they become large.

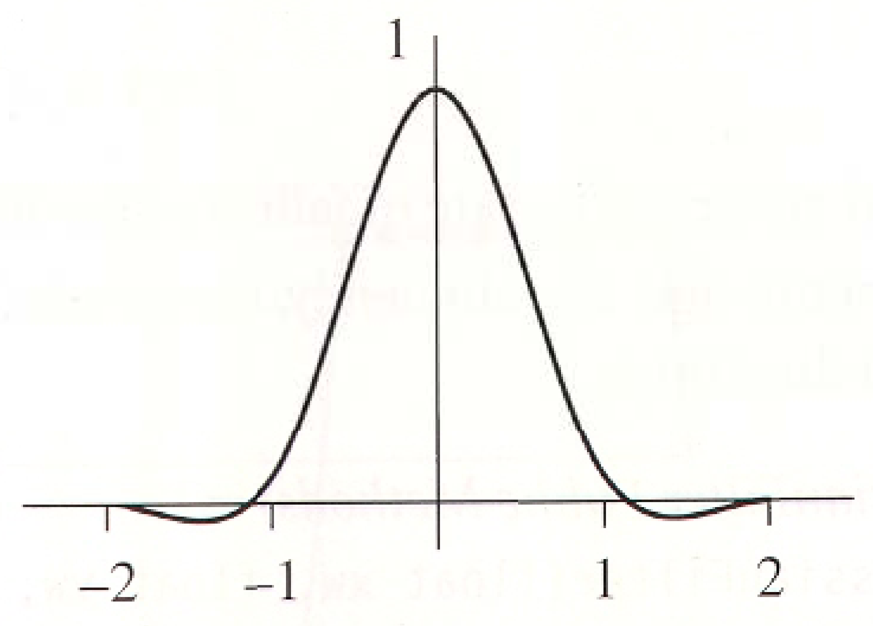

Mitchell filter

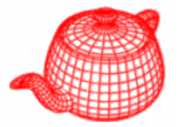

$$h(x) = \frac{1}{6}\begin{cases} (12-9B-6C)x^3 + (-18+12B+6C)x^2 + (6-2B) & |x|<1 \\ (-B-6C)x^3 + (6B+30C)x^2 + (-12B-48C)x + (8B+24C) & 1<|x|<2 \\ 0 & otherwise \end{cases}$$

- Separable filter
- Two parameters, B and C, B+2C=1 suggested

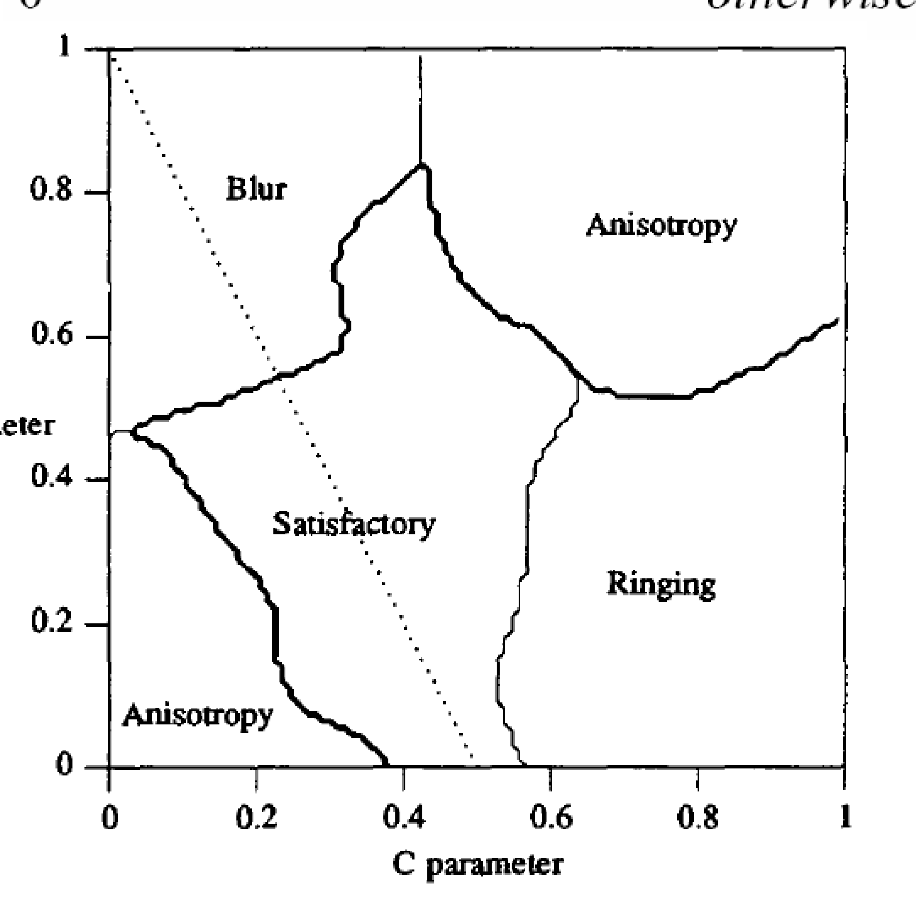

FFT of a cubic filter. Mitchell filter is a combination of cubic filters with C^0 and C^1 Continuity.

Sinc^4

Windowed sinc filter

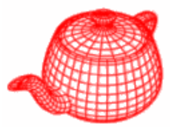

$$Lanczos \quad w(x) = \frac{\sin \pi x / \tau}{\pi x / \tau}$$

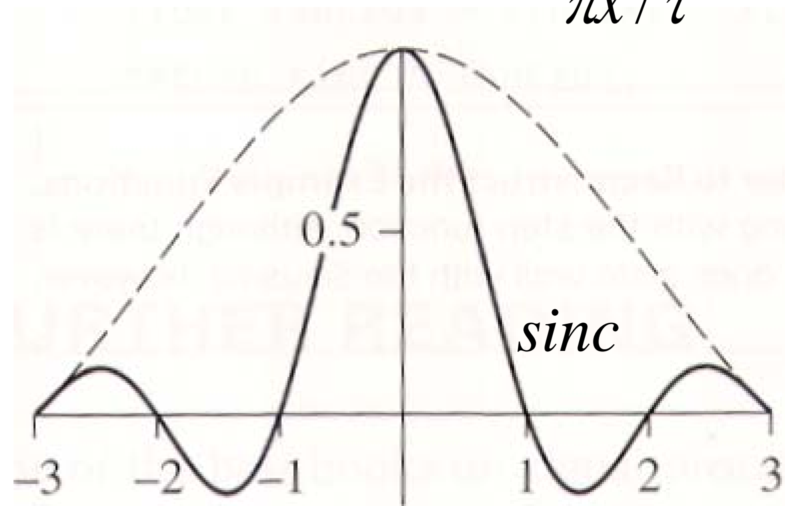

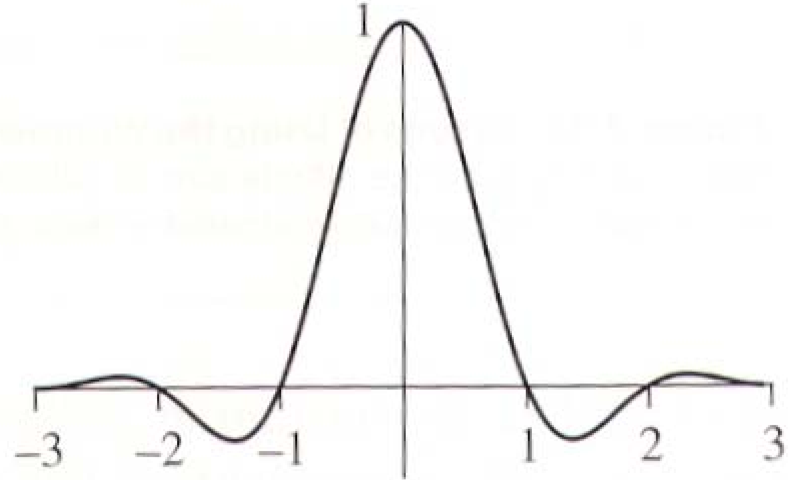

Comparisons

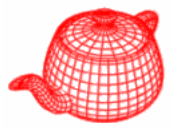

box

Mitchell

Comparisons

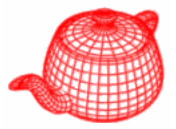

windowed sinc

Mitchell

Comparisons

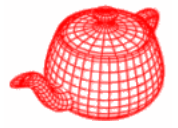

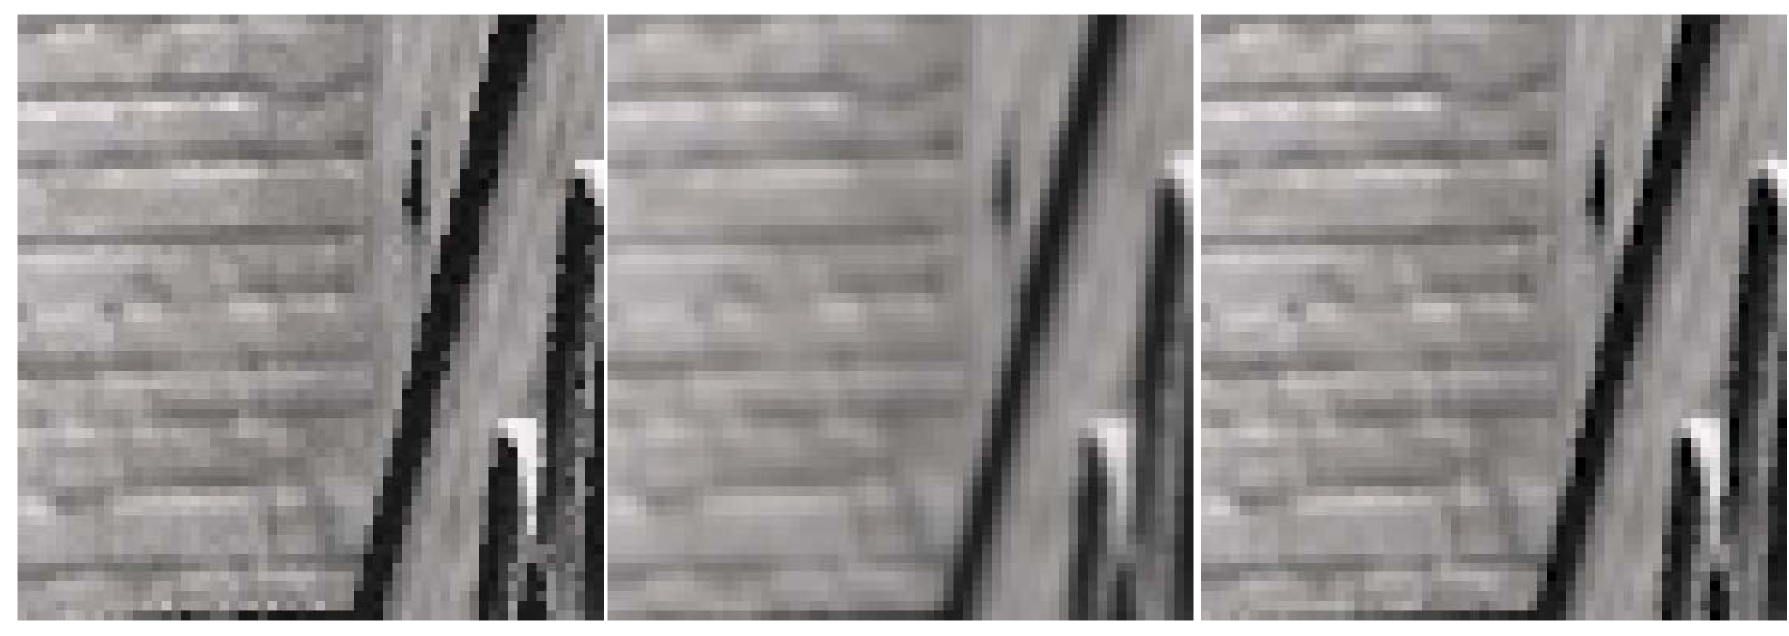

box *Gaussian* *Mitchell*

Film

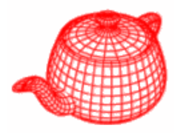

- **`Film`** class simulates the sensing device in the simulated camera. It determines samples' contributions to the nearby pixels and writes the final floating-point image to a file on disk.
- Tone mapping operations can be used to display the floating-point image on a display.
- **`core/film.*`**

Film

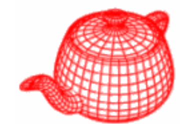

```
class Film {
public:
  Film(int xres, int yres)
      : xResolution(xres), yResolution(yres) { }
```

add samples for later reconstruction by weighted average

```
  virtual void AddSample(const CameraSample &sample,
                         const Spectrum &L) = 0;
```

simply sum samples' contributions, not average them. Pixels with more samples will be brighter. It is used by light transport methods such as MetropolisRender

```
  virtual void Splat(const CameraSample &sample,
                     const Spectrum &L) = 0;
```

Film

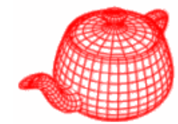

the sample extent could be a bit larger than the pixel extent

```
virtual void GetSampleExtent(int *xstart,
  int *xend, int *ystart, int *yend) const = 0;
virtual void GetPixelExtent(int *xstart,
  int *xend, int *ystart, int *yend) const = 0;
```

be notified when a region has been recently updated. Do nothing as default.

```
virtual void UpdateDisplay(int x0, int y0, int x1,
  int y1, float splatScale = 1.f);
```

generate the final image for saving to disk or displaying. It accepts a scale factor.

```
virtual void WriteImage(float splatScale = 1.f)=0;

const int xResolution, yResolution;
};
```

ImageFilm

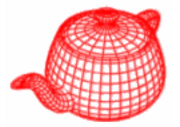

- **film/image.cpp** implements the only film plug-in in pbrt. It filters samples and writes the resulting image to disk.

```
ImageFilm::ImageFilm(int xres, int yres,Filter *filt,
  float crop[4],string &filename, bool openWindow)
{
  ...
  pixels = new BlockedArray<Pixel>(xPixelCount,
                                   yPixelCount);
  <precompute filter table>
}
```

in NDC space. useful for debugging, or rendering on different computers and assembling later (crop)

on some system, it can be configured to open a window and show the image as it's being rendered (openWindow)

AddSample

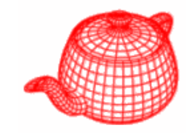

extent

sample

precomputed
Filter table

grid of pixels

$$I(x,y)=\frac{\sum_i f(x-x_i,y-y_i)L(x_i,y_i)}{\sum_i f(x-x_i,y-y_i)}$$

find the nearest neighbor
in the filter table

Recent progresses on Poisson sampling

- On-the-fly computing
  - Scalloped regions [SIGGRAPH 2006]
- Tile-based
  - Recursive Wang tile [SIGGRAPH 2006]
- Parallel
  - Li-Yi Wei [SIGGRAPH 2008]
- Show three videos for them

Fast Poisson-Disk Sampling

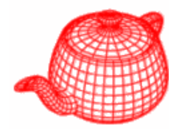

Fast Poisson-Disk Sampling

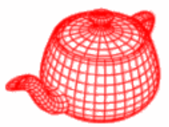

Recursive Wang Tiles for Blue Noise

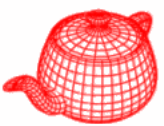

32,965 points, 22.75ms
1,449,011 points per second

34,897 points, 15.7ms
2,222,739 points per second

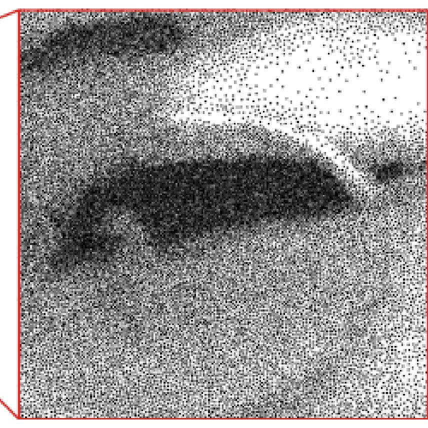

22,748 points, 11.67ms
1,949,272 points per second