



Monte Carlo Integration II

Digital Image Synthesis


Yung-Yu Chuang

with slides by Pat Hanrahan and Torsten Moller

variance = noise in the image



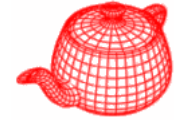
without variance reduction



with variance reduction

Same amount of computation for rendering this scene with glossy reflection

Variance reduction



- Efficiency measure for an estimator

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

- Although we call them variance reduction techniques, they are actually techniques to increase efficiency
 - Stratified sampling
 - Importance sampling

Russian roulette



- Assume that we want to estimate the following direct lighting integral

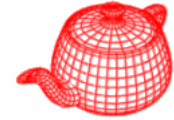
$$L_o(p, \omega_o) = \int_{\Omega} f_r(p, \omega_o, \omega_i) L_d(p, \omega_i) |\cos \theta_i| d\omega_i$$

- The Monte Carlo estimator is

$$\frac{1}{N} \sum_{i=1}^N \frac{f_r(p, \omega_o, \omega_i) L_d(p, \omega_i) |\cos \theta_i|}{p(\omega_i)}$$

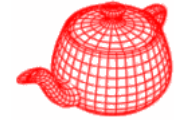
- Since tracing the shadow ray is very costly, if we somewhat know that the contribution is small anyway, we would like to skip tracing.
- For example, we could skip tracing rays if $|\cos \theta_i|$ or $f_r(p, \omega_o, \omega_i)$ is small enough.

Russian roulette



- However, we can't just ignore them since the estimate will be consistently under-estimated otherwise.
- Russian roulette makes it possible to skip tracing rays when the integrand's value is low while still computing the correct value on average.

Russian roulette



- Select some termination probability q ,

$$F' = \begin{cases} \frac{F - qc}{1 - q} & \xi > q \\ c & \text{otherwise} \end{cases}$$

$$E[F'] = (1 - q) \left(\frac{E[F] - qc}{1 - q} \right) + qc = E[F]$$

- Russian roulette will actually increase variance, but improve efficiency if q is chosen so that samples that are likely to make a small contribution are skipped. (if same number of samples are taken, RR could be worse. However, since RR could be faster, we could increase number of samples)

Careful sample placement




- Carefully place samples to less likely to miss important features of the integrand
- Stratified sampling: the domain $[0,1]^s$ is split into strata $S_1..S_k$ which are disjoint and completely cover the domain.

$$S_i \cap S_j = \phi \quad i \neq j \quad \bigcup_{i=1}^k S_i = [0,1]^s$$

$$|S_i| = v_i \quad \sum v_i = 1$$

$$p_i(x) = \begin{cases} 1/v_i & \text{if } x \in S_i \\ 0 & \text{otherwise} \end{cases}$$

Stratified sampling




$$V[\hat{I}_s] = \frac{1}{N} \sum_{i=1}^k v_i \sigma_i^2$$


$$V[\hat{I}_{ns}] = \frac{1}{N} \left[\sum_{i=1}^k v_i \sigma_i^2 + \sum_{i=1}^k v_i (\mu_i - I)^2 \right]$$

Thus, the variance can only be reduced by using stratified sampling.

Stratified sampling



without stratified sampling



with stratified sampling

Bias



- Another approach to reduce variance is to introduce bias into the computation.

$$\beta = E[F] - \int f(x)dx$$

- Example: estimate the mean of a set of random numbers X_i over $[0..1]$.

unbiased estimator $\frac{1}{N} \sum_{i=1}^N X_i$ variance (N^{-1})

biased estimator $\frac{1}{2} \max(X_1, X_2, \dots, X_N)$ variance (N^{-2})

Pixel reconstruction



$$I = \int w(x) f(x) dx$$

- Biased estimator $\hat{I}_b = \frac{\sum_{i=1}^N w(X_i) f(X_i)}{\sum_{i=1}^N w(X_i)}$
(but less variance)

- Unbiased estimator $\hat{I}_u = \frac{\sum_{i=1}^N w(X_i) f(X_i)}{N p_c}$

where p_c is the uniform PDF of choosing X_i

$$\begin{aligned} E[\hat{I}_u] &= \frac{1}{N p_c} \sum_{i=1}^N E[w(X_i) f(X_i)] \\ &= \frac{1}{N p_c} \sum_{i=1}^N \int w(x) f(x) p_c dx = \int w(x) f(x) dx \end{aligned}$$

Importance sampling



- The Monte Carlo estimator

$$F_N = \frac{1}{N} \sum_{i=1}^N \frac{f(X_i)}{p(X_i)}$$

converges more quickly if the distribution $p(x)$ is similar to $f(x)$. The basic idea is to concentrate on where the integrand value is high to compute an accurate estimate more efficiently.

- So long as the random variables are sampled from a distribution that is similar in shape to the integrand, variance is reduced.

Informal argument



- Since we can choose $p(x)$ arbitrarily, let's choose it so that $p(x) \sim f(x)$. That is, $p(x) = cf(x)$. To make $p(x)$ a pdf, we have to choose c so that

$$c = \frac{1}{\int f(x)dx}$$

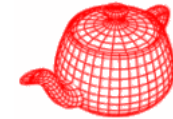
- Thus, for each sample X_i , we have

$$\frac{f(X_i)}{p(X_i)} = \frac{1}{c} = \int f(x)dx$$

Since c is a constant, the variance is zero!

- This is an ideal case. If we can evaluate c , we won't use Monte Carlo. However, if we know $p(x)$ has a similar shape to $f(x)$, variance decreases.

Importance sampling

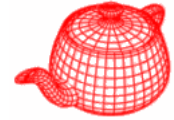


- Bad distribution could hurt variance.

$$I = \int_0^4 x dx = 8$$

method	Sampling function	variance	Samples needed for standard error of 0.008
importance	$(6-x)/16$	$56.8/N$	887,500
importance	$1/4$	$21.3/N$	332,812
importance	$(x+2)/16$	$6.4/N$	98,432
importance	$x/8$	0	1
stratified	$1/4$	$21.3/N^3$	70

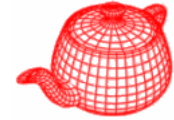
Importance sampling



- Fortunately, it is not too hard to find good sampling distributions for importance sampling for many integration problems in graphics.
- For example, in many cases, the integrand is the product of more than one function. It is often difficult to construct a pdf similar to the product, but sampling along one multiplicand is often helpful.

$$\int_{S^2} f(p, \omega_o, \omega_i) L_i(p, \omega_i) |\cos \theta_i| d\omega_i$$

Multiple importance sampling



$$L_o(p, \omega_o) = \int_{\Omega} f_r(p, \omega_o, \omega_i) L_d(p, \omega_i) |\cos \theta_i| d\omega_i$$

- If we sample based on either L or f , it often performs poorly.
- Consider a near-mirror BRDF illuminated by an area light where L 's distribution is used to draw samples. (It is better to sample by f .)
- Consider a diffuse BRDF and a small light source. If we sample according to f , it will lead to a larger variance than sampling by L .
- It does not work by averaging two together since variance is additive.

Multiple importance sampling



- To estimate $\int f(x)g(x)dx$, MIS draws n_f samples according to p_f and n_g samples to p_g , The Monte Carlo estimator given by MIS is

$$\frac{1}{n_f} \sum_{i=1}^{n_f} \frac{f(X_i)g(X_i)w_f(X_i)}{p_f(X_i)} + \frac{1}{n_g} \sum_{j=1}^{n_g} \frac{f(Y_j)g(Y_j)w_g(Y_j)}{p_g(Y_j)}$$

- Balance heuristic v.s. power heuristic

$$w_s(x) = \frac{n_s p_s(x)}{\sum_i n_i p_i(x)} \qquad w_s(x) = \frac{(n_s p_s(x))^\beta}{\sum_i (n_i p_i(x))^\beta}$$


Multiple importance sampling




- Assume a sample X is drawn from p_f where $p_f(X)$ is small, thus $f(X)$ is small if p_f matches f . If, unfortunately, $g(X)$ is large, then standard importance sampling gives a large value $\frac{f(X)g(X)}{p_f(X)}$
- However, with the balance heuristic, the contribution of X will be

$$\begin{aligned} \frac{f(X)g(X)w_f(X)}{p_f(X)} &= \frac{f(X)g(X)}{p_f(X)} \frac{n_f p_f(X)}{n_f p_f(X) + n_g p_g(X)} \\ &= \frac{f(X)g(X)n_f}{n_f p_f(X) + n_g p_g(X)} \end{aligned}$$


Importance sampling




Sample Light




Sample BRDF



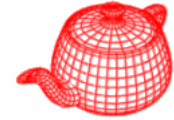
Multiple importance sampling



Result: better than either of the two strategies alone



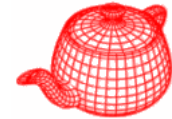
Monte Carlo for rendering equation



$$L_o(\mathbf{p}, \omega_o) = L_e(\mathbf{p}, \omega_o) + \int_{\Omega} f(\mathbf{p}, \omega_o, \omega_i) L_i(\mathbf{p}, \omega_i) |\cos \theta_i| d\omega_i$$

- Importance sampling: sample ω_i according to BxDF f and L (especially for light sources)
- If don't know anything about f and L , then use cosine-weighted sampling of hemisphere to find a sampled ω_i

Sampling reflection functions



```
Spectrum BxDF::Sample_f(const Vector &wo,
Vector *wi, float u1, float u2, float *pdf){
    *wi = CosineSampleHemisphere(u1, u2);
    if (wo.z < 0.) wi->z *= -1.f;
    *pdf = Pdf(wo, *wi);
    return f(wo, *wi);
}
```

For those who modified `Sample_f`, `Pdf` must be changed accordingly

```
float BxDF::Pdf(Vector &wo, Vector &wi) {
    return SameHemisphere(wo, wi) ?
        fabsf(wi.z) * INV_PI : 0.f;
} Pdf() is useful for multiple importance sampling.
```

Sampling microfacet model



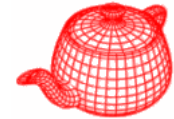
geometric attenuation G

microfacet distribution D Fresnel reflection F

$$f_r(\omega_i \omega_o) = \frac{D(\omega_h) G(\omega_i, \omega_o) F(\omega_i, \omega_h)}{4 \cos \theta_i \cos \theta_o}$$

Too complicated to sample according to f , sample D instead. It is often effective since D accounts for most variation for f .

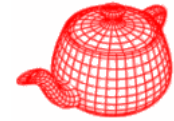
Sampling microfacet model



```
Spectrum Microfacet::Sample_f(const Vector &wo,
    Vector *wi, float u1, float u2, float *pdf) {
    distribution->Sample_f(wo, wi, u1, u2, pdf);
    if (!SameHemisphere(wo, *wi))
        return Spectrum(0.f);
    return f(wo, *wi);
}

float Microfacet::Pdf(const Vector &wo,
    const Vector &wi) const {
    if (!SameHemisphere(wo, wi)) return 0.f;
    return distribution->Pdf(wo, wi);
}
```


Sampling Blinn microfacet model




- Blinn distribution: $D(\cos \theta_h) = \frac{e+2}{2\pi} (\cos \theta_h)^e$
- Generate ω_h according to the above function

$$\cos \theta_h = e+1 \sqrt{\frac{\xi}{\zeta_1}}$$

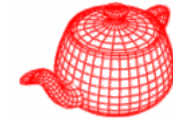
$$\phi_h = 2\pi \xi \zeta_2$$

- Convert ω_h to ω_i

$$\omega_i = -\omega_o + 2(\omega_o \cdot \omega_h)\omega_h$$



Sampling Blinn microfacet model



- Convert half-angle PDF to incoming-angle PDF, that is, change from a density in term of ω_h to one in terms of ω_i

$$\theta_i = 2\theta_h \text{ and } \phi_i = \phi_h$$

$$d\omega_i = \sin\theta_i d\theta_i d\phi_i$$


$$d\omega_h = \sin\theta_h d\theta_h d\phi_h$$

$$\frac{d\omega_h}{d\omega_i} = \frac{\sin\theta_h d\theta_h d\phi_h}{\sin\theta_i d\theta_i d\phi_i} = \frac{\sin\theta_h d\theta_h d\phi_h}{\sin 2\theta_h 2d\theta_h d\phi_h} = \frac{\sin\theta_h}{4 \cos\theta_h \sin\theta_h}$$

$$= \frac{1}{4 \cos\theta_h}$$



$$p(\theta) = \frac{p_h(\theta)}{4(\omega_o \cdot \omega_h)}$$



transformation method

$$p_y(y) = \left(\frac{dy}{dx} \right)^{-1} p_x(x)$$

Sampling anisotropic microfacet model

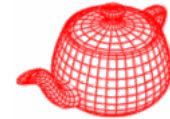
- Anisotropic model (after Ashikhmin and Shirley) for the first quadrant of the unit hemisphere

$$D(\omega_h) = \sqrt{(e_x + 1)(e_y + 1)} (\omega_h \cdot n)^{e_x \cos^2 \phi + e_y \sin^2 \phi}$$

$$\phi = \arctan \left(\sqrt{\frac{e_x + 1}{e_y + 1}} \tan \left(\frac{\pi \xi_1}{2} \right) \right)$$

$$\cos \theta_h = \xi_2 (e_x \cos^2 \phi + e_y \sin^2 \phi + 1)^{-1}$$

Estimate reflectance




$$\rho_{hd}(\omega_o) = \int_{\Omega} f_r(\omega_o, \omega_i) |\cos \theta_i| d\omega_i$$

```
Spectrum BxDF::rho(const Vector &w, int nSamples,
                  const float *samples) const
{
    Spectrum r = 0.;
    for (int i = 0; i < nSamples; ++i) {
        Vector wi;
        float pdf = 0.f;
        Spectrum f = Sample_f(w, &wi, samples[2*i],
                              samples[2*i+1], &pdf);
        if (pdf > 0.) r += f * AbsCosTheta(wi) / pdf;
    }
    return r / float(nSamples);
}
```

$$\frac{1}{N} \sum_{i=1}^N \frac{f_r(\omega_o, \omega_i) |\cos \theta_i|}{p(\omega_i)}$$

Estimate reflectance



```
Spectrum BxDF::rho(int nSamples, float *samples1,  
                  float *samples2) const
```

```
{
```

$$\rho_{hh} = \frac{1}{\pi} \int_{\Omega} \int_{\Omega} f_r(\omega_o, \omega_i) |\cos \theta_i \cos \theta_o| d\omega_i d\omega_o$$

```
Spectrum r = 0.;
```

```
for (int i = 0; i < nSamples; ++i) {
```

```
    Vector wo, wi;
```

```
    wo = UniformSampleHemisphere(samples1[2*i],  
                                  samples1[2*i+1]);
```

```
    float pdf_o = INV_TWOPI, pdf_i = 0.f;
```

```
    Spectrum f = Sample_f(wo, &wi, samples2[2*i],  
                          samples2[2*i+1], &pdf_i);
```

```
    if (pdf_i > 0.)
```

```
        r += f * AbsCosTheta(wi) * AbsCosTheta(wo)  
            / (pdf_o * pdf_i);
```

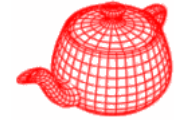
```
}
```

```
return r / (M_PI * nSamples);
```

```
}
```

$$\frac{1}{\pi} \frac{1}{N} \sum_{i=1}^N \frac{f_r(\omega_i', \omega_i'') |\cos \theta_i' \cos \theta_i''|}{p(\omega_i') p(\omega_i'')}$$

Sampling BSDF (mixture of BxDFs)



- We would like to sample from the density that is the sum of individual densities

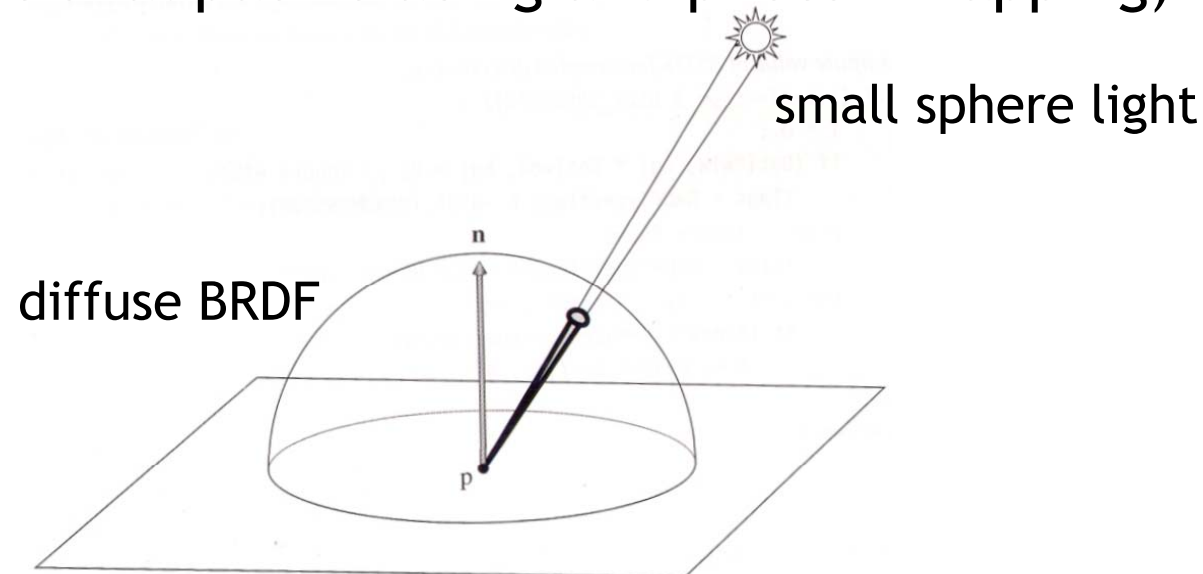
$$p(\omega) = \frac{1}{N} \sum_{i=1}^N p_i(\omega)$$

- Difficult. Instead, uniformly sample one component and use it for importance sampling. However, f and Pdf returns the sum.
- Three uniform random numbers are used, the first one determines which BxDF to be sampled (uniformly sampled) and then sample that BxDF using the other two random numbers


Sampling light sources



- Direct illumination from light sources makes an important contribution, so it is crucial to be able to generate
 - S_p : samples directions from a point p to the light
 - S_r : random rays from the light source (for bidirectional light transport algorithms such as bidirectional path tracing and photon mapping)



Interface




```
Spectrum Sample_L(Point &p, float pEpsilon,  
    LightSample &ls, float time, Vector *wi,  
    float *pdf, VisibilityTester *vis) = 0;
```

```
float Pdf(const Point &p, const Vector &wi)  
    const = 0;
```

```
float uPos[2], uComponent;
```


```
Spectrum Sample_L(Scene *scene, LightSample  
    &ls, float u1, float u2, float time,  
    Ray *ray, Normal *Ns, float *pdf) = 0;
```


Point lights




- S_p : delta distribution, treat similar to specular BxDF
- S_r : sampling of a uniform sphere

Point lights



```
Spectrum Sample_L(const Point &p, float u1, float u2,
    Vector *wi, float *pdf, VisibilityTester *vis)
{
    *pdf = 1.f; delta function
    return Sample_L(p, wi, visibility);
}
float Pdf(Point &, Vector &) for almost any direction, pdf is 0
{ return 0.; }
Spectrum Sample_L(Scene *scene, LightSample &ls,
    float u1, float u2, float time, Ray *ray,
    Normal *Ns, float *pdf) const {
    *ray = Ray(lightPos,
        UniformSampleSphere(ls.uPos[0], ls.uPos[1]),
        0.f, INFINITY, time);
    *Ns = (Normal)ray->d;
    *pdf = UniformSpherePdf();
    return Intensity;
}
```

Spotlights



- Sp: the same as a point light
- Sr: sampling of a cone (ignore the falloff)

$$p(\omega) = c \text{ over cone} \longrightarrow p(\theta, \phi) = c \sin \theta \text{ over } [0, \theta_{\max}] \times [0, 2\pi]$$

$$1 = \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\theta'} c \sin \theta d\theta d\phi = 2\pi c (1 - \cos \theta_{\max}) \longrightarrow p(\theta, \phi) = \frac{\sin \theta}{2\pi(1 - \cos \theta_{\max})}$$

$$p(\theta) = \int_{\phi=0}^{2\pi} \frac{\sin \theta}{2\pi(1 - \cos \theta_{\max})} d\phi = \frac{\sin \theta}{1 - \cos \theta_{\max}}$$

$$P(\theta) = \int_{\theta=0}^{\theta'} \frac{\sin \theta}{1 - \cos \theta_{\max}} d\theta = \frac{1 - \cos \theta'}{1 - \cos \theta_{\max}} = \xi_1 \longrightarrow \cos \theta = (1 - \xi_1) + \xi_1 \cos \theta_{\max}$$

$$p(\phi | \theta) = \frac{p(\theta, \phi)}{p(\theta)} = \frac{1}{2\pi} \longrightarrow P(\phi' | \theta) = \int_{\phi=0}^{\phi'} \frac{1}{2\pi} d\phi = \frac{\phi'}{2\pi} = \xi_2 \longrightarrow \phi = 2\pi\xi_2$$

Spotlights

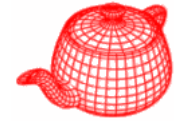


```
Spectrum Sample_L(Scene *scene, LightSample &ls,
    float u1, float u2, float time, Ray *ray,
    Normal *Ns, float *pdf) const
{
    Vector v = UniformSampleCone(ls.uPos[0],
        ls.uPos[1], cosTotalWidth);
    *ray = Ray(lightPos, LightToWorld(v), 0.f,
        INFINITY, time);
    *Ns = (Normal)ray->d;
    *pdf = UniformConePdf(cosTotalWidth);
    return Intensity * Falloff(ray->d);
}
```


Projection lights and goniophotometric lights

- Ignore spatial variance; sampling routines are essentially the same as spot lights and point lights


Directional lights



- S_p : no need to sample
- S_r : create a virtual disk of the same radius as scene's bounding sphere and then sample the disk uniformly.




Directional lights



```
Spectrum Sample_L(Scene *scene, LightSample &ls,
    float u1, float u2, float time, Ray *ray,
    Normal *Ns, float *pdf) const {
    Point worldCenter;
    float worldRadius;
    scene->WorldBound().BoundingSphere(&worldCenter,
                                        &worldRadius);

    Vector v1, v2;
    CoordinateSystem(lightDir, &v1, &v2);
    float d1, d2;
    ConcentricSampleDisk(ls.uPos[0], ls.uPos[1], &d1, &d2);
    Point Pdisk = worldCenter + worldRadius
                    * (d1 * v1 + d2 * v2);
    *ray = Ray(Pdisk + worldRadius * lightDir,
              -lightDir, 0.f, INFINITY, time);
    *Ns = (Normal)ray->d;
    *pdf = 1.f / (M_PI * worldRadius * worldRadius);
    return L;
}
```

Area lights



- Defined by shapes
- Add shape sampling functions for `Shape`
- `Sp`: uses a density with respect to solid angle from the point `p`

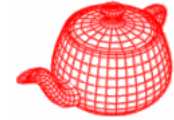
```
Point Shape::Sample(Point &P, float u1,  
float u2, Normal *Ns)
```

- `Sr`: generates points on the shape according to a density with respect to surface area

```
Point Shape::Sample(float u1, float u2,  
Normal *Ns)
```

- `virtual float Shape::Pdf(Point &Pshape)`
`{ return 1.f / Area(); }`

Area light sampling method

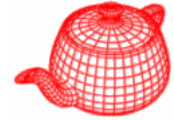


- Most of work is done by `Shape`.

```
Spectrum DiffuseAreaLight::Sample_L(Point &p, float
    pEpsilon, LightSample &ls, float time, Vector *wi,
    float *pdf, VisibilityTester *visibility) const {
    Normal ns;
    Point ps = shapeSet->Sample(p, ls, &ns);
    *wi = Normalize(ps - p);
    *pdf = shapeSet->Pdf(p, *wi);
    visibility->SetSegment(p, pEpsilon, ps, 1e-3f, time);
    Spectrum Ls = L(ps, ns, -*wi);
    return Ls;
}

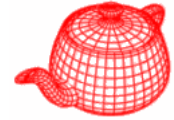
float DiffuseAreaLight:: Pdf(Point &p, Vector &wi) {
    return shapeSet->Pdf(p, wi);
}
```

Area light sampling method



```
Spectrum DiffuseAreaLight::Sample_L(Scene *scene,
    LightSample &ls, float u1, float u2, float time,
    Ray *ray, Normal *Ns, float *pdf) const
{
    Point org = shapeSet->Sample(ls, Ns);
    Vector dir = UniformSampleSphere(u1, u2);
    if (Dot(dir, *Ns) < 0.) dir *= -1.f;
    *ray = Ray(org, dir, 1e-3f, INFINITY, time);
    *pdf = shapeSet->Pdf(org) * INV_TWOPHI;
    Spectrum Ls = L(org, *Ns, dir);
    return Ls;
}
```


Sampling spheres



- Only consider full spheres

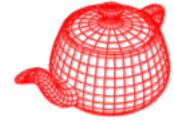
```
Point Sphere::Sample(float u1, float u2,
    Normal *ns) const
{
    Point p = Point(0,0,0) + radius
        * UniformSampleSphere(u1, u2);
    *ns = Normalize(
        (*ObjectToWorld)(Normal(p.x, p.y, p.z)));
    if (ReverseOrientation) *ns *= -1.f;
    return (*ObjectToWorld)(p);
}
```

Sampling spheres



$$\cos \theta = \sqrt{1 - \left(\frac{r}{|p - c|} \right)^2}$$

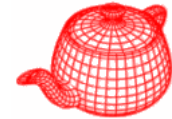
Sampling spheres



```
Point Sphere::Sample(Point &p, float u1,
    float u2, Normal *ns) const
{
    Point Pcenter
        = (*ObjectToWorld)(Point(0,0,0));
    Vector wc = Normalize(Pcenter - p);
    Vector wcX, wcY;
    CoordinateSystem(wc, &wcX, &wcY);

    if (DistanceSquared(p, Pcenter)
        - radius*radius < 1e-4f)
        return Sample(u1, u2, ns);
}
```


Sampling spheres



```
float sinThetaMax2 = radius*radius
                    / DistanceSquared(p, Pcenter);
float cosThetaMax = sqrtf(
    max(0.f, 1.f - sinThetaMax2));
DifferentialGeometry dgSphere;
float thit, rayEpsilon;
Point ps;
Ray r(p, UniformSampleCone(u1, u2,
    cosThetaMax, wcX, wcY, wc), 1e-3f);
if (!Intersect(r, &thit, &rayEpsilon, &dgSphere))
    thit = Dot(Pcenter - p, Normalize(r.d));
ps = r(thit);
*ns = Normal(Normalize(ps - Pcenter));
if (ReverseOrientation) *ns *= -1.f;
return ps;
}
```

It's unexpected.

Infinite area lights

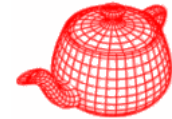


uniform sampling
4 image samples
X 8 light sample



importance sampling
4 image samples
X 8 light sample

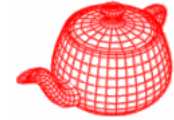
Infinite area lights



```
Spectrum Sample_L(Point &p, float pEpsilon,
    LightSample &ls, float time, Vector *wi,
    float *pdf, VisibilityTester *visibility)
{
    use importance sampling to find (u, v)
    float uv[2], mapPdf;
    distribution->SampleContinuous(ls.uPos[0],
    precomputed          ls.uPos[1], uv, &mapPdf);
    distribution2D
    if (mapPdf == 0.f) return 0.f;

    convert (u,v) to direction
    float theta=uv[1]*M_PI, phi=uv[0]*2.f* M_PI;
    float costheta = cosf(theta),
        sintheta = sinf(theta);
    float sinphi = sinf(phi),
        cosphi = cosf(phi);
```


Infinite area lights



```
*wi = LightToWorld(Vector(sintheta * cosphi,  
                           sintheta * sinphi, costheta));  
  
*pdf = mapPdf / (2.f * M_PI * M_PI * sintheta);  
if (sintheta == 0.f) *pdf = 0.f;  
  
visibility->SetRay(p, pEpsilon, *wi, time);  
Spectrum Ls = Spectrum(radianceMap->Lookup(uv[0],  
                                             uv[1]), SPECTRUM_ILLUMINANT);  
return Ls;  
}
```

$$g(u, v) = (\pi v, 2\pi u) \quad \left| J_g \right| = 2\pi^2$$

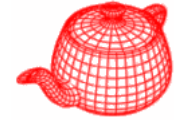
$$p(\theta, \phi) = \frac{p(u, v)}{2\pi^2} \quad p(\omega) = \frac{p(\theta, \phi)}{\sin \theta} = \frac{p(u, v)}{2\pi^2 \sin \theta}$$

Infinite area lights



```
Float Pdf(const Point &, const Vector &w)
{
    Vector wi = WorldToLight(w);
    float theta = SphericalTheta(wi),
           phi = SphericalPhi(wi);
    float sintheta = sinf(theta);
    if (sintheta == 0.f) return 0.f;
    float p =
        distribution->Pdf(phi*INV_TWOPI, theta*INV_PI)
        / (2.f * M_PI * M_PI * sintheta);
    return p;
}
```

Infinite area lights



```
Spectrum Sample_L(Scene *scene, LightSample &ls,
    float u1, float u2, float time, Ray *ray,
    Normal *Ns, float *pdf) const
{
    float uv[2], mapPdf;
    distribution->SampleContinuous(ls.uPos[0],
        ls.uPos[1], uv, &mapPdf);
    if (mapPdf == 0.f) return Spectrum(0.f);

    float theta = uv[1]*M_PI, phi = uv[0]*2.f*M_PI;
    float costheta = cosf(theta),
        sintheta = sinf(theta);
    float sinphi = sinf(phi),
        cosphi = cosf(phi);
    // sample a direction first
}
```

Infinite area lights

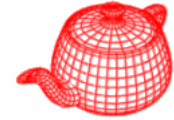


```
Vector d = -LightToWorld(Vector(sintheta * cosphi,  
                                sintheta * sinphi, costheta));  
*Ns = (Normal)d;
```

For the origin, do the similar as distant lights


```
Point worldCenter; float worldRadius;  
scene->WorldBound().BoundingSphere(&worldCenter,  
                                    &worldRadius);  
  
Vector v1, v2;  
CoordinateSystem(-d, &v1, &v2);  
float d1, d2;  
ConcentricSampleDisk(u1, u2, &d1, &d2);  
Point Pdisk = worldCenter + worldRadius *  
              (d1 * v1 + d2 * v2);  
*ray = Ray(Pdisk+worldRadius * -d, d, 0., ...);
```

Infinite area lights



```
float directionPdf = mapPdf
    / (2.f * M_PI * M_PI * sintheta);
float areaPdf = 1.f
    / (M_PI * worldRadius * worldRadius);
*pdf = directionPdf * areaPdf;
if (sintheta == 0.f) *pdf = 0.f;
Spectrum Ls = (radianceMap->Lookup(uv[0], uv[1]),
    SPECTRUM_ILLUMINANT);
return Ls;
}
```

Sampling lights




- Structured Importance Sampling of Environment Maps, SIGGRAPH 2003

irradiance

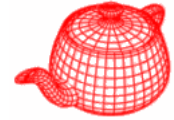
$$E(x) = \int_{\Omega_{2\pi}} L_i(\vec{\omega}) S(x, \vec{\omega}) (\vec{\omega} \cdot \vec{n}) d\vec{\omega}$$

binary visibility

Infinite area light; easy to evaluate



Importance metric



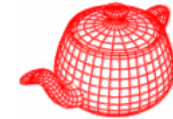
illumination of a region

$$\Gamma(L, \Delta\omega) = L^a \Delta\omega^b$$

solid angle of a region

- Illumination-based importance sampling (a=1, b=0)
- Area-based stratified sampling (a=0, b=1)

Variance in visibility



- After testing over 10 visibility maps, they found that variance in visibility is proportional to the square root of solid angle (if it is small)

$$V[S, \Delta\omega] \approx \frac{\theta}{3T}$$

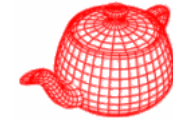
visibility map $\Delta\omega = \pi\theta^2$ parameter typically between 0.02 and 0.6

- Thus, they empirically define the importance as

$$\Gamma[L, \Delta\omega] = L \cdot (\min(\Delta\omega, \Delta\omega_0))^{1/4}$$

set as 0.01

Hierarchical thresholding



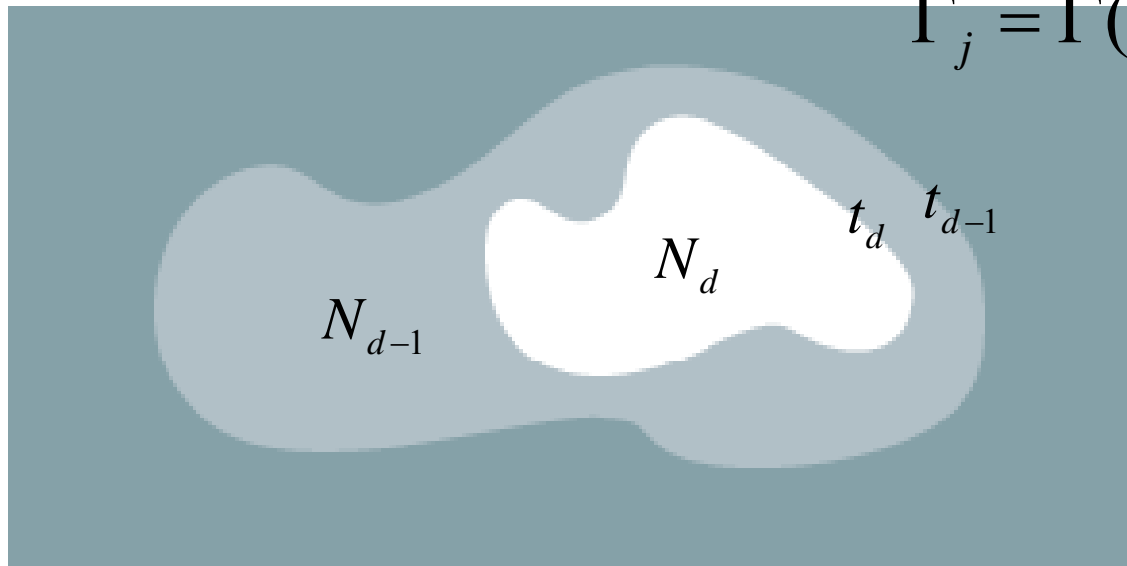
$$t_i = i\sigma \quad i = 0, \dots, d-1$$

standard deviation of
the illumination map

$$d=6$$


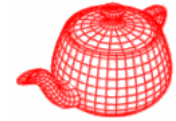
$$\Gamma_{4\pi} = \Gamma(\sum L, \Delta\omega_0) = L\Delta\omega_0^{1/4}$$

$$\Gamma_j = \Gamma(\sum_{i \in C_j}^i L_i, \sum_{i \in C_j} \Delta\omega_i)$$




$$N_j = N \frac{\Gamma_j}{\Gamma_{4\pi}}$$

Hierarchical stratification



Results



Importance w/ 300 samples






Importance w/ 3000 samples




Structured importance w/ 300 samples

Sampling BRDF




http://www.cs.virginia.edu/~jdl/papers/brdfsamp/lawrence_sig04.ppt


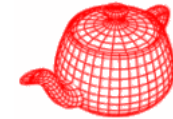
Sampling product





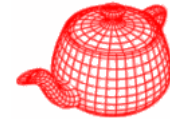
- Wavelet Importance Sampling: Efficiently Evaluating Products of Complex Functions, SIGGRAPH 2005.






Sampling product





Wavelet decomposition




Sample warping



Results



Results



BRDF
importance
sampling

Structured
importance
sampling

Our algorithm

3 samples

10 samples

30 samples

60 samples

100 samples

Results

