CGI2014 manuscript No.
(will be inserted by the editor)

Real-time Rendering of Glossy Materials with Regular Sampling

Abstract Rendering view-dependent, glossy surfaces in order to increase the realism in real-time applications is a computationally complex task, that can only be performed by applying some approximations – especially when immediate changes in the scene in terms of material settings and object placement are a necessity. The use of environment maps is a common approach to this problem, but implicates performance problems due to costly pre-filtering steps or expensive sampling.

We therefore introduce a *regular sampling scheme* for environment maps that relies on an efficient MIP-map-based filtering step, and minimizes the number of necessary samples for creating a convincing real-time rendering of glossy material BRDFs.

Keywords Real-Time Rendering · BRDF Materials

1 Introduction

Extending the realism and quality of rendered images has always been one of the most driving forces in the field of computer graphics. Apart from numerous rendering techniques that have been developed over the last decades to produce more and more convincing images, various material models as well as the use of measured material data (e.g. BRDFs) have been suggested. While it is meanwhile possible (and even very common) to generate photo-realistic images and videos using offline rendering systems, the application of these material models in *real-time applications* is rare – mainly due to the high computational complexity involved in evaluating view-dependent, glossy materials, making it hard to achieve interactive frame rates in complex, modifiable scenes.

Still, the need for a realistic visualization of instantly modifiable materials in a scene is a requirement in numerous application scenarios, such as lighting design for architecture, or level design for games: it is necessary to handle more than just simple diffuse, perfectly reflecting, or transparent materials within complex scenes – without having to wait for an offline renderer to calculate the result within several minutes, or to "pre-bake" a specific, fixed material onto the objects in the scene.

Various approaches deal with the handling of glossy materials in real-time scenarios – but all of them are either limited in terms of scene complexity, interactions or modifiability, and/or require expensive precomputations (see Section 2 for an overview).

One of these approaches is derived from image-based lighting solutions, where a complete, high resolution environment map is computed for each object with a glossy material. With a combination of pre-filtering and sampling, an approximation of the material appearance can be computed at interactive rates. Our novel approach is related to this image-based method, but due to a MIP-mapping-based pre-filtering step with minimal costs (see Section 3), our method is applicable in fully dynamic and modifiable scenarios.

In order to obtain interactive or even real-time frame rates (i.e. at least 60 frames per seconds), the most crucial part is to sample and evaluate the view-dependent, complex BRDF material data as fast as possible, while keeping visually disturbing artifacts at a minimum. While it has been proposed by [Colbert and Kriv\'anek 2007] to sample the BRDF using a fixed importance-based pattern, we suggest:

- a novel MIP-map-based *regular* sampling scheme that achieves visual appearance comparable or superior to recent approaches with a lower number

of samples and minimal pre-computation overhead (i.e. the generation of the MIP-maps),
– an integration of our sampling scheme with the material model of Kurt et al. [Kurt et al 2010], and therefore allow
– the application of view-dependent, glossy BRDF materials in complex, modifiable scenarios.

2 Related Work

Real-Time shading. A number of methods that are used for real time shading approaches for image based lighting are based on spherical harmonics [Ramamoorthi and Hanrahan 2001, Kautz et al 2002, Ramamoorthi and Hanrahan 2002, Sloan et al 2002] or other basis functions [Ng et al 2004, Wang et al 2004, Liu et al 2004, Tsai and Shih 2006] and thus use significant amounts of pre-computation time. This makes them ill-suited for interactive application that require that both lighting and materials can be changed quickly. Even some of the methods that use environment maps such as the one by Kautz et al. [Kautz et al 2000] and McAlister et al.[McAllister et al 2002] are based on pre-filtering and thus preclude fast changes of the environment if they are used without any optimization. Ben-Artzi et al. [Ben-Artzi et al 2006] and Sun et al. [Sun et al 2007] showed that dynamic BRDF changes can be accomodated with an increased cost in both pre-computation time and storage.

Recently a number of methods based on micro environment maps [Ritschel et al 2009, Scherzer et al 2012] have been shown to produce interactive results for scenes with general glossy surfaces. Due to the limited precision of small environment maps, and due to the large storage requirements, these methods have problems dealing with small, very bright light sources such as spot lights. Moreover, they are limited in maximal scene complexity and do not allow interactive scene modifications.

The multi-pass rendering method from Diefenbach and Badler [Diefenbach and Badlert 1997] for reflections on large planar surfaces from is very popular. Fuhrmann et al. [Fuhrmann et al 2004] extended this method to simulate physically plausible glossy reflections by using an image space blur filter. A more general method that approximates different levels of glossiness from a perfect reflection image with a parameterized bilateral filter has been published by Robison and Shirley [Robison and Shirley 2009].

In order to accommodate completely dynamic lighting and changes in BRDFs, our goal is to reduce pre-computation time as much as possible. Similarly to the importance sampling method of Colbert and Krivanek [Colbert and Křivánek 2007, Křivánek and Colbert 2008] we only use MIP-mapped environment maps to sample the surface irradiance. We show that a carefully chosen regular sampling scheme yields superior results with a low sample budget to importance sampling for computing the contribution of the reflection lobe of glossy BRDFs.

BRDF models. Although striving for realism would suggest a physically based BRDF model [Cook and Torrance 1982, Torrance and Sparrow 1967, He et al 1991] as a basis for realistic rendering, the high evaluation cost is too expensive in real-time applications. For this reason there is a long history of empirical models that have been designed for fast evaluation beginning with Phong [Phong 1975], Ward [Ward 1992] and Schlick [Schlick 2003]. Recently Kurt et al. [Kurt et al 2010] introduced a model that combines the expressiveness of the previous two (anisotropy, Fresnel effect), is physically plausible under a larger number of parameter configurations and, due to an additionally introduced parameter, is particularly suited for fitting to material measurements. Our sampling method is based on this material model.

3 Environment Map Creation

As indicated in the introduction, an environment map is assigned to each scene object with glossy material surfaces (Figure 1).

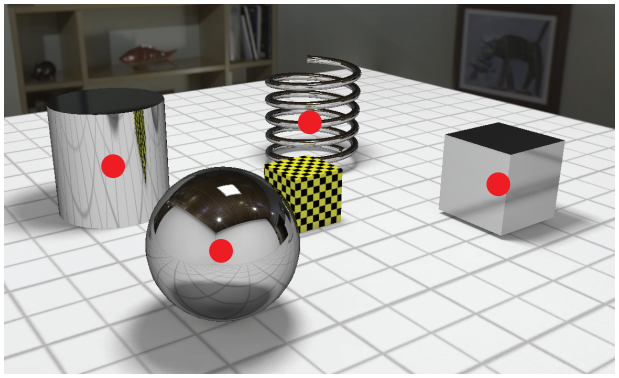

Fig. 1: An environment map is placed at the centroid of each object to approximate the corresponding irradiant light. Object repositioning or the change of material parameters in the scene leads to the need of an environment map update.

Such an environment map contains the complete irradiance at the centroid of the object P, and is used to approximate the incident light for all surface points P'

of the object (Figure 2). This is the standard approximation with assumption of the reflected environment at infinite distance [Blinn and Newell 1976]. GPU support enables fast run-time retrieval of incident light from arbitrary directions during rendering.

Fig. 2: Incident light approximation at arbitrary surface points P' of an object using the environment map generated at the object centroid P

MIP-maps are generated to provide pre-filtered versions of the perfect reflection. In order to generate convincing real-time renderings of surfaces with glossy materials, the pre-filtered reflections are used to sample the irradiance according to the BRDF of the material (see Section 4). Our sampling method covers typical reflection lobes with a minimum number of sample directions (see Section 5).

While the computation time of the environment maps depends on the scene complexity, the pre-filtering step of each environment map (i.e. the MIP-map generation) is only dependent on the chosen resolution.

The MIP-map pyramid provides average irradiance light of square regions with a sinlge sample lookup, however, the standard box-filter reconstruction may cause slightly shifted and distorted shapes of the specular highlights. By using a smoother image filter, the final rendering quality can be improved. In Table 1 we have listed the computation times of the pre-filtering step with different environment map resolutions – using standard MIP-mapping, and pyramids-of-gaussian [Adelson 1981] with 4×4 and 6×6 filter kernels.

Size	256	512	1024
Box 2×2	0.18ms	0.44ms	1.45ms
Gauss 4×4	0.52ms	1.23ms	3.86ms
Gauss 6×6	0.65ms	1.77ms	5.97ms

Table 1: Generation times of the MIP-map pyramid with different environment map resolutions and different image filters. The timings have been measured using the same hardware setup as used for the evaluation in Section 8.

Note that the pre-filtering does not depend on the BRDF, so that arbitrary materials can be visualized using the generated image pyramid. This makes it possible to render objects with multiple materials using a single environment map, reducing memory requirements as well as the number of updates after changes in the scene.

4 Material Model

The material model of Kurt et al. [Kurt et al 2010] represents the effects of glossy materials as a weighted sum of a selectable (small) number of glossy reflection lobes:

$$f_r(\omega_i, \omega_o) = \frac{k_d}{\pi} + \sum_{n=1}^{\#lobes} k_{s_n} \cdot f_{r_{lobe_n}}(\omega_i, \omega_o), \tag{1}$$

where k_d is the overall diffuse reflectivity and k_s the reflectivity of a specular lobe BRDF f_r. The geometry at the reflection point and the definition of the vectors used in this and the following equations are shown in Figure 3.

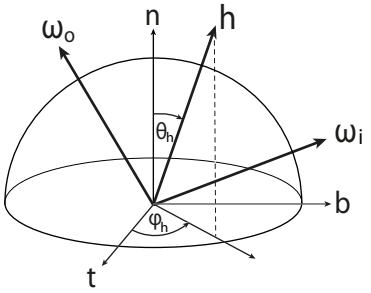

Fig. 3: Illustration of the used vectors and angles: ω_o is the outgoing vector with direction to the camera; the vector ω_i represents an incident direction; h is the half-way vector between ω_i and ω_o and θ_h, ϕ_h are its spherical coordinates with respect to the local tangent frame (t, b, n)

The partial BRDF of each of the reflection lobes is defined as follows:

$$f_{r_{lobe}}(\omega_i, \omega_o) = \frac{F(\omega_o \cdot h) D(h)}{4(\omega_o \cdot h)((\omega_i \cdot n)(\omega_o \cdot n))^{\alpha}} \tag{2}$$

For the fresnel term $F(\omega_o \cdot h)$, Schlick's approximation with the parameter r_0, the reflection coefficient at normal incidence, is used:

$$F(u) = r_0 + (1 - r_0) \cdot (1 - u)^5 \tag{3}$$

The microfacet distribution $D(h)$ of the model is an anisotropic normalized version of the Beckmann distribution, controlled by the roughness parameters m_x and m_y:

$$D(h) = \frac{q(h)}{\pi m_x m_y cos^4(\theta_h)} \tag{4}$$

$$q(h) = e^{-tan^2(\theta_h)(\frac{cos^2(\phi_h)}{m_x^2} + \frac{sin^2(\phi_h)}{m_y^2})} \tag{5}$$

Finally, the BRDF model incorporates the adjustable parameter $\alpha \in [0, 1]$ to control the shadowing/masking effect, allowing particular flexibility in representing very smooth and shiny materials. For $\alpha = 0$ the model is always physically plausible, but for larger values the energy conservation depends on the surface roughness $m_{x,y}$.

The material model also allows importance sampling of the BRDF with the transformation from uniform random variables $\xi_1, \xi_2 \in (0, 1]$ to sampling directions that follow the microfacet distribution (as first used in the Ward model):

$$\phi_h = arctan(\frac{m_y}{m_x} tan(2\pi\xi_2)) \tag{6}$$

$$\theta_h = arctan(\sqrt{\frac{-log(\xi_1)}{\frac{cos^2\,\phi_h}{m_x^2} + \frac{sin^2\,\phi_h}{m_y^2}}}) \tag{7}$$

For computing the reflected light at a surface point, the following integration has to be performed for each lobe:

$$L_{o_{spec}}(\omega_o) = \int_{\Omega_{HS}} L_i(\omega_i)\, f_{r_{lobe}}(\omega_o|\omega_i)\, |cos(\theta_i)|\, \mathrm{d}\omega_i \tag{8}$$

where Ω_{HS} is the hemisphere of directions around the surface normal, $|cos(\theta_i)|$ accounts for Lambert's cosine law, and $\mathrm{d}\omega_i$ is the differential solid angle around ω_i defined as $sin(\theta_i)\, d\theta_i\, d\phi_i$, as the integrated intervals densify towards the poles by a factor $\sin(\theta_i)$. $f_{r_{lobe}}$ is the specular part of the BRDF.

By using the importance-based sampling directions given in equations 6 and 7, the integral can be solved numerically with the Monte-Carlo integration by summing the correctly weighted contributions from each sampling direction [Szirmay-Kalos 2008]:

$$L_{o_{spec}}(\omega_o) = \frac{1}{n} \sum_{i=1}^{n} L_i(\omega_i)\, cos(\theta_i)\, \frac{f_{r_{lobe}}(\omega_o|\omega_i)}{p(\omega_i|\omega_o)} \tag{9}$$

The underlying sampling probability densitiy function (PDF) for each lobe of the material model is the following:

$$p(\omega_i|\omega_o) = \frac{q(h)}{4\pi\, m_x\, m_y\, cos^3(\theta_h)\, (\omega_o \cdot h)} \tag{10}$$

This material model can be directly used in the method of Colbert and Křivánek [Colbert and Krivánek 2007]. It turns out, that with a low number of samples this method produces unwanted artifacts (see Figures 8 and 10). Instead of an importance-based sample distribution, we propose a regular sample distribution, reducing these artifacts and better preserving the characteristics of the BRDF.

5 Regular Sample Pattern Generation

Our sample generation process places samples on equally spaced concentric circles around the half-way vector to cover the most significant sample directions according to the microfacet distribution given in Equation 5. The number of samples placed on each circle is scaled with the circumference. For anisotropic materials, which have an elliptical distribution, the same circular scheme is used and samples with low relevance are simply rejected.

The first step in the sample generation process is to find the angular offset θ_{max} where the PDF of the Beckmann distribution drops below a certain density ξ_0. Since the generation scheme is based on circles, the calculation is reduced to symmetric distribution with $\max(m_x, m_y)$. Deriving the importance transformation of θ (Equation 7) gives the following equation:

$$\theta_{max} = \arctan\left(\sqrt{-\log(\xi_0)}\ \max(m_x, m_y)\right) \tag{11}$$

In order to align n_c equally spaced circles from $\theta \in [0, \theta_{max}]$ there are two valid spacings s_1 and s_2. The spacing s_1 gives patterns where the first sample is placed at $\theta = 0$ and the offset spacing s_2 places the first samples on a circle around the pole of the half-vector space:

$$s_1 = \frac{\theta_{max}}{n_c + 1}, \quad s_2 = \frac{\theta_{max}}{n_c + 2} \tag{12}$$

Samples are then placed from $\phi \in [0, 2\pi)$ at equal offsets. The number of samples of the first circle n_0 is fixed at 6 for s_1 and 4 for s_2. The number of samples for the other circles n_i is increased by the ratio of the circumference to the first circle. In order to get perfect symmetric patterns for anisotropic materials, the number of samples per circle is rounded to the nearest multiple of two:

$$n_i = 2 \left\lfloor n_0 \frac{\sin(\theta_i)}{2\sin(\theta_0)} + \frac{1}{2} \right\rfloor \tag{13}$$

Figure 4 illustrates the generated sample patterns for both spacing with $n_c = 3$ and the threshold $\xi_0 = 0.1$. The pattern with s_1 has 19 and s_2 30 samples.

For anisotropic materials this process is extended by rejecting samples with $q(\phi_h, \theta_h)$ below the threshold

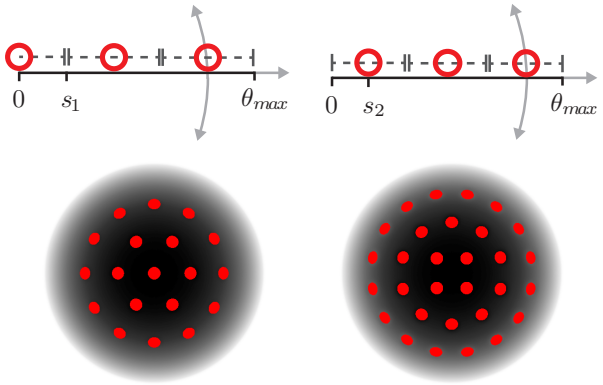

Fig. 4: Illustration of the two valid sample spacings s_1 (left) and s_2 (right) for equidistant patterns with $n_c = 3$. Spacing s_1 results in a pattern with 19 samples and s_2 30 samples.

ξ_0. Additionally, a second circle with half a sample offset is calculated and the circle with the highest sum of $q(\phi_h, \theta_h)$ of all samples is taken. This results in regular equally-spaced sample patterns for anisotropic materials. Figure 5 illustrates some examples.

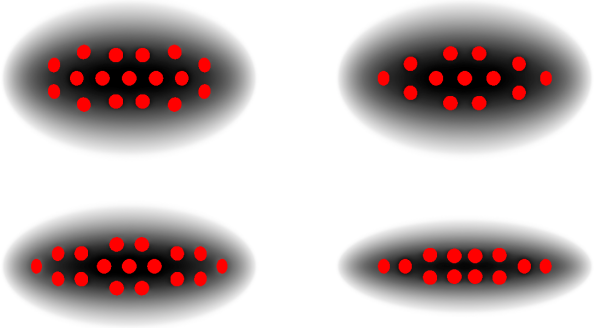

Fig. 5: Examples of equally-spaced sampling patterns for anisotropic materials.

Even though this sample generation process is generic, it does not allow the creation of an exact number of samples. Specifying the number of circles n_c may not seem very practical, and the resulting number of samples is highly dependent on the material anisotropy. Even though estimating the number of circles for a specified sample budget is not trivial, it is feasible to generate several patterns with a different number of circles, since the process is computationally inexpensive. The pattern with the closest number of samples to the specified budget can then be used.

6 Filtering of Evenly Spaced Samples

In order to avoid aliasing artifacts caused by undersampling, each sample is filtered using MIP-maps. Optimally, the kernel size is chosen such that the entire highlight gets covered by the union of all samples as shown in Figure 6 (left).

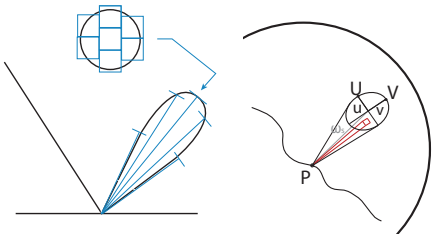

Fig. 6: Specular lobe covered by MIP-mapped samples (left). Ideally, the pyramidal cones defined by MIP-mapped samples should densely cover the lobe without overlapping. The cross-section of the lobe at unit distance from a shaded point p is assumed to be elliptical (right). Its vertices U, V are used to estimate the lobe area. Divided by the number of samples, we get the area that is covered by a single sample (red).

The cross-section of a specular lobe is assumed to be elliptical, as illustrated in Figure 6 (right). First, the area subtended by the entire lobe at unit distance is estimated. We take the extreme half-vector samples in x and y at the border ξ_0 and transform them into the incident directions ω_U and ω_V. Together with the lobe center ω_C, the vertices U and V of the elliptical lobe can be calculated:

$$\alpha_U = \arccos(\omega_C \cdot \omega_U) \quad \alpha_V = \arccos(\omega_C \cdot \omega_V) \tag{14}$$

Now the lengths of the ellipse semi-major and semi-minor axes u and v are obtained using:

$$u = \tan(\alpha_U) \quad v = \tan(\alpha_V) \tag{15}$$

With these lengths, the area of the elliptical cross section is computed using $A = uv\pi$. Now the area of a single sample at unit distance is estimated by dividing the unprojected area of the ellipse A by the number of samples. Since the reconstruction of the ellipse is dependent on the view-vector ω_o, it has to be performed in the pixel shader once before the specular lobe is sampled.

Finally, there are two effects to be considered individually for each sample. First, the fact that the transformation from half-vector space directions to incident directions no longer gives equally distributed samples in angular domain. This distortion causes samples at

grazing angles to appear sharper when the same estimated area is used for all samples. Therefore, the density transformation from half directions to sample directions (Equation 17) is used to adjust the approximated area covered by a sample with $\frac{1}{\omega_i \cdot h}$. Note that the factor $\frac{1}{4}$ is missing, since the area has already been calculated with the actual incident directions.

Second, when sampling in environment maps, the projection of an area from the unit sphere to the surrounding cube is distorted and the number of pixels increases in the corners. The correction factor of this distortion is given by the distance where the lookup direction ω_i intersects the cube, which is given by the length of the lookup direction normalized by its maximum. Since the environment map cube surrounds the unit sphere, the cube-face area is 4. This gives the following final equation for the estimated number of covered texels per sample, including the environment map and the half-vector distortion:

$$pixels_{sample} = \frac{AN}{4n_{samples}} \left\| \frac{\omega_i}{\|\omega_i\|_\infty} \right\| \frac{1}{\omega_i \cdot h}, \tag{16}$$

where N is the number of the cube-face pixels.

The corresponding MIP-level can be calculated by $\log_4(pixels_{sample})$, as the number of averaged pixels increases by a factor 4 in each level of the MIP-map pyramid.

7 BRDF Integral with Regular Samples

In order to use our regular sampling pattern to approximate the BRDF integral with the Monte Carlo sampling scheme (Equation 9), we need to derive a suitable probability density function. Assuming equally distributed samples, their PDF would be constant and dependent on the portion of the sampled area on the hemisphere. However, since our sample vectors represent half-way directions h, only the probability density in half-space p_h is constant and the density transformation to the out vector space p_o has to be considered. The derivation of this transformation has been shown by Walter in Section 3.2 [Walter 2005] and results in the following equation:

$$p_o(o) = \frac{p_h(h)}{4(h \cdot i)} \tag{17}$$

This gives a dependency on the relation between the half direction h and the incident direction i. The PDF $p_h(h)$ that describes the probability density for equally spaced half-way vectors is given by the average weight of the Beckmann distribution (Euqation 4) with m_x and m_y of the material roughness to preserve the overall energy transport. Therefore, the PDF constant is precomputed with the actual sample vectors as follows:

$$p_h(h) = \frac{1}{N} \sum_{i=1}^{N} D(h_i) \tag{18}$$

The specular lobe integral can now be evaluated in a shader with the following pseudo code:

```
float3 EvaluateSpecularLobe(float3 lightOutVec)
{
  // Equation 14, 15 & Area
  float a = estimateSampleArea(lightOutVec);

  float3 brdfSum = 0;
  for (int i = 0; i < numSamplesLobe; i++)
  {
    float3 h = tangentToWorldSpace(SampleVectors[i]);
    float3 lightInVec = reflect(-lightOutVec, h);

    // Equation 16
    float mipLevel = calcMipLevel(a, lightInVec, lightOutVec);

    // Environment map lookup
    float3 L = cubeMapTexture.SampleLevel(lightInVec, mipLevel);

    // Equation 2
    float brdf = evaluateBrdfKarl(lightInVec, lightOutVec);

    // Equation 17 & 18
    float pdf = PDFConstant / (4 * dot(h, lightInVec));

    // Equation 9 (inner sum)
    brdfSum += L * brdf / pdf * dot(lightInVec, n);
  }

  return SpecularCoefficient * brdfSum / numSamplesLobe;
}
```

8 Evaluation

In this section we will demonstrate the performance of our material rendering method. We provide an extensive evaluation of the quality compared to the importance-based approach of Colbert and Křivánek [Colbert and Krivánek 2007] as well as a reference solution in scenarios with different materials. The benchmark system consists of an Intel i7-920 processor and a Fermi-based NVIDIA Quadro 6000 graphics card.

The basic requirement for a real-time application is to deliver at least 60 frames per second. Sampling an environment-map as excessively done in our rendering method is primarily limited by the memory bandwidth of the GPU. Figure 7 (bottom) shows the correlation of the number of samples per pixel and the frame rate in high-definition (1920×1080) resolution, measured in the evaluation scene of Figure 8. The different graphs represent scenarios where a quarter (green), half (red) and all pixels (blue) of the screen have been shaded using our technique. This evaluation shows that a budget of about 20 samples is available for complex materials when large portions of the screen are covered.

The following evaluations will use a reference rendering with >10000 samples accumulated over several

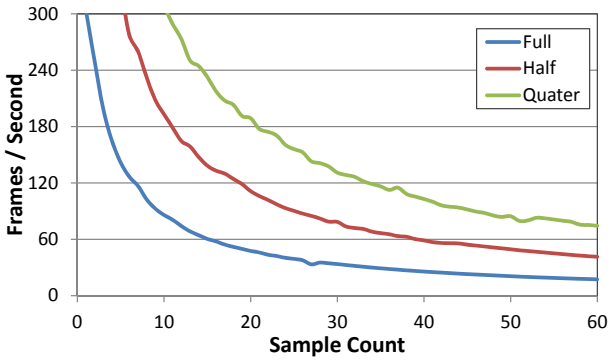

Fig. 7: Frame-rate dependent on the number of samples in a full-HD rendering where: 100% (blue), 50% (red) and 25% (green) pixels on the screen are sampled as complex material.

seconds by a Monte Carlo method. The correlation is analyzed by the root mean square (RMS) error between the images. Since we are only interested in the accuracy of sampling the specular lobe, the Fresnel and shadowing/masking effect of the selected materials have been disabled ($r_0 = 1.0, \alpha = 0.0$). In order to simulate a demanding scenario, the environment map of the evaluation setup contains tiny bright spots with $2 - 10°$ diameter and a ratio of relative local intensities up to $1 : 3000$. The environment map has a size of 1024, and a 6×6 Gaussian filter has been used for the construction of the MIP-map pyramid. It is used in our regular sampling and the compared importance-based sampling method.

Evaluating the material impressions (Figure 8) of the isotropic and anisotropic material of our novel approach (a) compared to the reference solution (c), the specular highlights have a well-defined shape and preserve the characteristics of the BRDF. In comparison, the PDF-based MIP-map lookup of the importance-based method (b) results in brightening of actually dark areas when a low number of samples is used. Furthermore, the shape of the specular highlights either appears irregularly (with no MIP-bias) or too blurred and with a non-determinable shape. These effects are mostly visible with anisotropic materials (also see Figure 10).

A false color visualization of the difference image shows the error distribution around the specular highlights (Figure 9). When the MIP-bias is tweaked to yield the lowest possible error, our regular sampling methods balances the error around each sample and the sampling pattern becomes visible. With increased number of samples the pattern dissolves and the result evenly converges to the reference image. In comparison, the importance-based method distributes the error over a

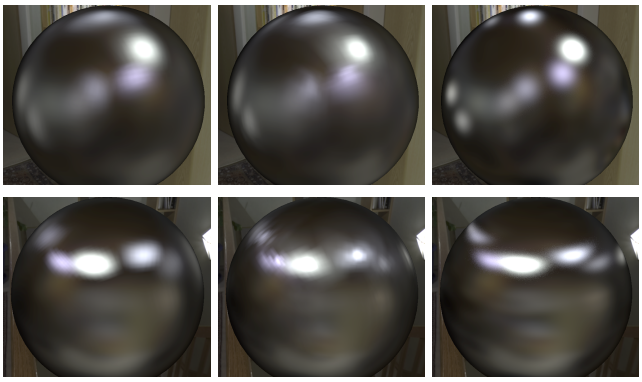

(a) Regular (b) Importance (c) Reference

Fig. 8: Material with $m_x = m_y = 0.1$ approximated with 14 samples (top) and an anisotropic material with $m_x = 0.2$, $m_y = 0.075$ approximated with 17 samples (bottom). Regular samples with ξ=0.2 produce a good impression of the material with well defined specular highlights for both materials (a). While the importance-based method has similar quality with the isotropic material (b, top), the specular highlights of the anisotropic materials are irregularly distorted (b, bottom).

larger region and therefore does not achieve a similarly low average error, but the uneven distribution results in a smoother image.

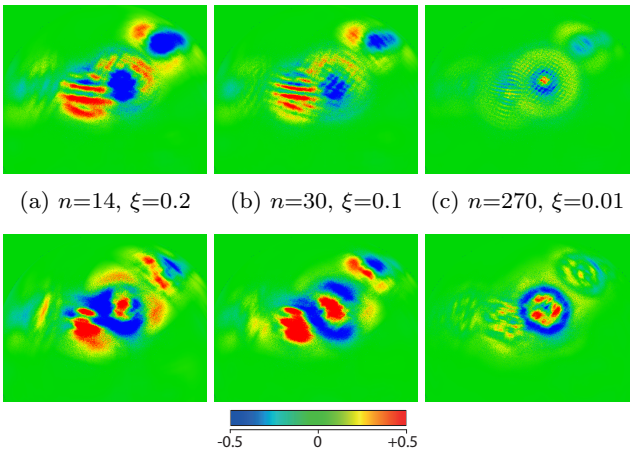

(a) n=14, ξ=0.2 (b) n=30, ξ=0.1 (c) n=270, ξ=0.01

Fig. 9: Error visualization of material rendering with $m_x = m_y = 0.15$ with 14 (left), 30 (middle) and 270 (right) samples with regular sampling (top) and importance sampling (bottom). Regular sampling balances the error at each sample and converges evenly to the reference solution, while importance sampling has a higher average error that is distributed more smoothly.

The result images with 270 samples in Figure 10 show that our sampling method provides a detailed

BRDF appearance for isotropic and anisotropic materials with low difference to the reference image. It turns out that the regular sample alignment together with a defined MIP-level lookup is very robust in terms of high frequencies and intensity spikes in the irradiance image and is therefore well suited for high-quality renderings. The importance-based method of Colbert and Křivánek shows good convergence for isotropic materials as well, but does not maintain the same quality with anisotropic BRDFs. Adjusting the MIP-bias causes a significant distortion towards an isotropic BRDF and introduces too much blur on the far ends of the specular highlights.

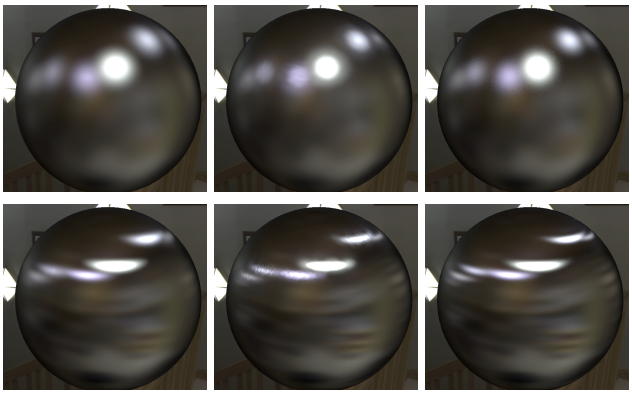

(a) Regular (b) Importance (c) Reference

Fig. 10: Result comparison with 270 samples: The regular sample pattern with ξ=0.01 results in a clear shape of the isotropic and anisotropic specular highlights (a) and closely matches the reference solution (c). The importance-based sample pattern yields a similar result for the isotropic material (top, b), but shows distorted and irregular specular highlights in the result of the anisotropic material (bottom, b).

9 Application

In order to prove the applicability of our novel approach, we have integrated the proposed BRDF filtering and sampling scheme in a light design and modeling application, where real-time interactions as well as an immediate physically plausible visualization of the light distribution in the scene are both important requirements. While the diffuse light propagation of the global illumination solution is not view-dependent and can hence be stored in light maps, glossy materials need to be evaluated and rendered each frame. We do this for the last "bounce" of the illumination, i.e. glossy materials are not reflected multiple times.

The application of our novel technique poses no problem even in large, complex architectural scenes: Interactive or even real-time frame rates are achieved in scenarios with up to over 100 objects with glossy materials (i.e. over 100 environment maps), making modeling and immediate parameter changes possible. We refer the reader to the supplemental video, where we demonstrate the real-time rendering as well as the interactions in the scenes, as well as to Figure 11 for example screenshots.

10 Conclusion

We have proposed a regular sampling method for evaluating glossy material BRDFs using environment maps. It reduces the number of necessary samples compared to previous methods, and thus improves the rendering performance for comparable or even superior visualization quality in real-time applications. Especially the visualization of anisotropic materials is significantly closer to the reference solution than previous approaches with the same number of samples. The used material model of Kurt et al. allows a high variety of material types and provides close fitting to measured BRDF data.

Since the environment maps can be generated and pre-filtered within a few milliseconds, immediate material changes and scene modifications pose no problem for our approach. Moreover, a single environment map can be used for objects with multiple materials, making both updates and memory storage efficient. Both has been demonstrated by applying our novel method in a light planning software framework.

References

Adelson EH (1981) Image data compression with the laplacian pyramid. In: In Proceedings of the conference on pattern recognition and image processing, IEEE Computer Society Press, pp 218–223

Ben-Artzi A, Overbeck R, Ramamoorthi R (2006) Real-time brdf editing in complex lighting. In: ACM Transactions on Graphics (TOG), ACM, vol 25, pp 945–954

Blinn JF, Newell ME (1976) Texture and reflection in computer generated images. Commun ACM 19(10):542–547

Colbert M, Krivánek J (2007) Gpu-based importance sampling. GPU Gems 3:459–476

Cook RL, Torrance KE (1982) A reflectance model for computer graphics. ACM Transactions on Graphics (TOG) 1(1):7–24

Diefenbach PJ, Badlert NI (1997) Multi-pass pipeline rendering: Realism for dynamic environments. In: Symposium on Interactive 3D Graphics, pp 59–70

Fuhrmann AL, Tobler RF, Maierhofer S (2004) Real-time glossy reflections on planar surfaces. In: Proc. of the 3rd internat. Conf. on Computer graphics, virtual reality, visualisation and interaction in Africa, ACM, pp 87–91

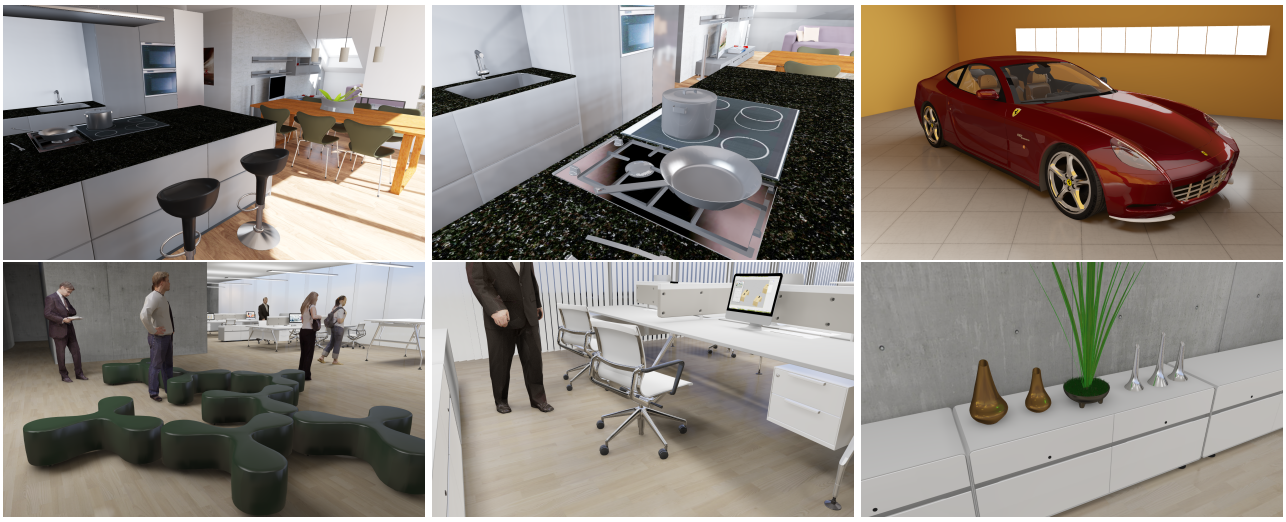

Fig. 11: Results of our environment map-based BRDF material rendering method in our real-time architectural visualization and light planning application with three scenes (the diffuse part of the global illumination solution is pre-computed and applied via light maps): The kitchen/living room interior (top row, left & middle) shows 8 glossy material types and uses 80 environment maps for objects with glossy materials. The office room (bottom row) has 7 different glossy materials with over 100 environment maps. The car showroom (top right) has a single high resolution environment map and demonstrates a wide range of glossy materials (car paint, glass, chrome, plastic, rubber, leather, ...). For a demonstration of the real-time visualization and modifications of view-dependent materials, please see the accompanying video.

He XD, Torrance KE, Sillion FX, Greenberg DP (1991) A comprehensive physical model for light reflection. In: ACM SIGGRAPH Computer Graphics, ACM, vol 25, pp 175–186

Kautz J, Vázquez PP, Heidrich W, Seidel HP (2000) A unified approach to prefiltered environment maps. In: EG Workshop on Rendering, Springer. Berlin, vol 6

Kautz J, Sloan PP, Snyder J (2002) Fast, arbitrary brdf shading for low-frequency lighting using spherical harmonics. In: Proc. of the 13th Eurographics workshop on Rendering, EG, pp 291–296

Křivánek J, Colbert M (2008) Real-time shading with filtered importance sampling. In: Computer Graphics Forum, ACM, vol 27, pp 1147–1154

Kurt M, Szirmay-Kalos L, Křivánek J (2010) An anisotropic brdf model for fitting and monte carlo rendering. SIGGRAPH Comput Graph 44(1):3:1–3:15

Liu X, Sloan PP, Shum HY, Snyder J (2004) All-frequency precomputed radiance transfer for glossy objects. In: Eurographics Symposium on Rendering, EG, vol 1

McAllister DK, Lastra A, Heidrich W (2002) Efficient rendering of spatial bi-directional reflectance distribution functions. In: Proc. of the ACM SIGGRAPH/EUROGRAPHICS Conf. on Graphics hardware, EG, pp 79–88

Ng R, Ramamoorthi R, Hanrahan P (2004) Triple product wavelet integrals for all-frequency relighting. In: ACM Trans. on Graphics (TOG), ACM, vol 23, pp 477–487

Phong BT (1975) Illumination for computer generated pictures. Communications of the ACM 18(6):311–317

Ramamoorthi R, Hanrahan P (2001) An efficient representation for irradiance environment maps. In: Proc. of the 28th annual Conf. on Computer graphics and interactive techniques, ACM, pp 497–500

Ramamoorthi R, Hanrahan P (2002) Frequency space environment map rendering. In: ACM Transactions on Graphics (TOG), ACM, vol 21, pp 517–526

Ritschel T, Engelhardt T, Grosch T, Seidel HP, Kautz J, Dachsbacher C (2009) Micro-rendering for scalable, parallel final gathering. In: ACM Transactions on Graphics (TOG), ACM, vol 28, p 132

Robison A, Shirley P (2009) Image space gathering. In: HPG '09: Proceedings of the Conference on High Performance Graphics 2009, ACM, New York, NY, USA, pp 91–98

Scherzer D, Nguyen CH, Ritschel T, Seidel HP (2012) Pre-convolved radiance caching. In: Computer Graphics Forum, ACM, vol 31, pp 1391–1397

Schlick C (2003) An inexpensive brdf model for physically-based rendering. In: Computer graphics forum, ACM, vol 13, pp 233–246

Sloan PP, Kautz J, Snyder J (2002) Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In: ACM Transactions on Graphics (TOG), ACM, vol 21, pp 527–536

Sun X, Zhou K, Chen Y, Lin S, Shi J, Guo B (2007) Interactive relighting with dynamic brdfs. ACM Transactions on Graphics (TOG) 26(3):27

Szirmay-Kalos L (2008) Monte Carlo Methods in Global Illumination - Photo-realistic Rendering with Randomization. VDM Verlag, Saarbrücken, Germany, Germany

Torrance KE, Sparrow EM (1967) Theory for off-specular reflection from roughened surfaces. JOSA 57(9):1105–1112

Tsai YT, Shih ZC (2006) All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. In: ACM Transactions on Graphics (TOG), ACM, vol 25, pp 967–976

Walter B (2005) Notes on the ward brdf. Tech. rep., Technical Report PCG-05-06, Cornell Program of Computer Graphics

Wang R, Tran J, Luebke D (2004) All-frequency relighting of non-diffuse objects using separable brdf approximation. In: EG Symposium on Rendering, pp 345–354

Ward GJ (1992) Measuring and modeling anisotropic reflection. In: ACM SIGGRAPH Computer Graphics, ACM, vol 26, pp 265–272