Interner Bericht

Efficient Illumination by High Dynamic Range Images

T. Kollig, A. Keller

323/02

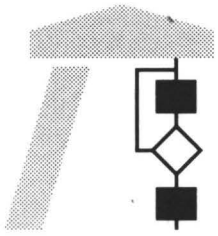

FACHBEREICH INFORMATIK

UNIVERSITÄT KAISERSLAUTERN

Postfach 3049 · D-67653 Kaiserslautern

Efficient Illumination by High Dynamic Range Images

T. Kollig, A. Keller

323/02

Universität Kaiserslautern
AG Numerische Algorithmen
Postfach 30 49
67653 Kaiserslautern
Germany

Dezember 2002

Herausgeber: AG Numerische Algorithmen
Leiter: Professor Dr. S. Heinrich

Efficient Illumination by High Dynamic Range Images

Thomas Kollig *
Kaiserslautern University

Alexander Keller †
Kaiserslautern University

Abstract

We present an algorithm for determining quadrature rules for computing the direct illumination of predominantly diffuse objects by high dynamic range images. The new method precisely reproduces fine shadow detail, is much more efficient as compared to Monte Carlo integration, and does not require any manual intervention.

1 Introduction

Rendering synthetic objects into real scenes requires their illumination by real world radiance, which can be captured as a high resolution spherical high dynamic range image [Debevec 1998; SpheronVR AG]. This technique has been introduced in [Debevec and Malik 1997; Debevec et al. 2002], is established in commercial products like LightWave, and has been used in many movies like the Oscar winning "The Matrix".

In this paper we address the efficient approximation of the radiance

$$L(x,\omega_o) \approx \int_{\Omega_x^+} L_{\text{hdr}}(\omega)V(x,\omega)f_r(\omega,x,\omega_o)\langle n(x),\omega\rangle d\omega \ ,$$

leaving a point x in direction ω_o. The spherical high dynamic range image L_{hdr} is of high resolution (up to 50 mega pixels). Ω_x^+ is the upper hemisphere around the surface normal $n(x)$ in point x. The visibility $V(x,\omega)$ is one if starting in point x no objects are seen in direction ω and zero otherwise.

For predominantly glossy bidirectional reflectance distribution functions f_r it is easy to evaluate the above integral by variance reduced Monte Carlo integration:

$$L(x,\omega_o) \approx \frac{1}{N}\sum_{i=0}^{N-1}\frac{L_{\text{hdr}}(\omega_i)V(x,\omega_i)f_r(\omega_i,x,\omega_o)\langle n(x),\omega_i\rangle}{p(\omega_i)} \ , \quad (1)$$

where the directions ω_i are p-distributed proportional to f_r. While this is an efficient procedure for highly specular surfaces, i.e. narrow cones of reflection, it utterly fails for predominantly diffuse surfaces due to high variance intrinsic to L_{hdr}. Choosing the directions ω_i proportional to L_{hdr} cannot be realized in an efficient way and approximations [Jensen 1995] usually do not perform any better than pure random sampling in this case. We also want to avoid variance reduction by manual stratification [Debevec et al. 2002]. On the other hand, projecting the integral operator into the spherical harmonics basis [Ramamoorthi and Hanrahan 2001; Ramamoorthi and Hanrahan 2002] allows for fast rendering. However, including detailed shadows in this model imposes considerable cost and restrictions [Sloan et al. 2002].

In this paper the above problem of illuminating predominantly diffuse surfaces by high resolution spherical high dynamic range images is addressed. Therefore a new method (sections 2 and 3) is presented that reliably captures all details of the high dynamic range image without manual intervention. The resulting rendering techniques (section 4) reproduce fine shadow details much more precise and are faster than previous approaches.

*e-mail:kollig@informatik.uni-kl.de
†e-mail:keller@informatik.uni-kl.de

2 Quadrature Rule Construction

The basic idea of our algorithm is to determine a quadrature rule $(\omega_i,B_i)_{i=0}^{N-1}$ only depending on the high dynamic range image L_{hdr}. Then the direct illumination can be approximated by

$$L(x,\omega_o) \approx \sum_{i=0}^{N-1} B_i\, V(x,\omega_i)f_r(\omega_i,x,\omega_o)\langle n(x),\omega_i\rangle \ , \quad (2)$$

where (ω_i,B_i) in fact corresponds to a directional light source from direction ω_i with radiosity B_i.

For a given partition $(\Omega_i)_{i=0}^{N-1}$ of the set of all unit directions Ω, the quadrature rule can be determined by

$$B_i := \int_{\Omega_i} L_{\text{hdr}}(\omega)d\omega \quad (3)$$

and choosing the directions ω_i as mass centroids of Ω_i, i.e.

$$\omega_i \in \Big\{\omega'\in\Omega \;\Big|\; \int_{\Omega_i}\left(\arccos(\langle\omega',\omega\rangle)\right)^2\|L_{\text{hdr}}(\omega)\|d\omega = \inf_{\omega''\in\Omega}\int_{\Omega_i}\left(\arccos(\langle\omega'',\omega\rangle)\right)^2\|L_{\text{hdr}}(\omega)\|d\omega\Big\} \ .$$

Thus the quadrature rule is based on the piecewise constant approximation

$$L_{\text{hdr}}(\omega) \approx \sum_{i=0}^{N-1}\frac{B_i}{|\Omega_i|}\chi_{\Omega_i}(\omega) \ ,$$

where χ_{Ω_i} is the characteristic function of $\Omega_i\subseteq\Omega$.

In the sequel it is shown how to construct a partition such that

- the partition $(\Omega_i)_{i=0}^{N-1}$ is the spherical Voronoi diagram of the set $(\omega_i)_{i=0}^{N-1}$ limiting the error if the visibility term V is neglected and
- $\max_{0\le i<N}\|B_i\|$ is as small as possible limiting the integration error of (2) made on each solid angle Ω_i.

Note that by the first constraint an implicit definition of Ω_i and ω_i is formed.

3 Determining the Quadrature Rule

The algorithm for determining the quadrature rule $(\omega_i,B_i)_{i=0}^{N-1}$ is based on Lloyd's relaxation method [Du et al. 1999] on the sphere:

1. Randomly select an initial set $(\omega_i)_{i=0}^{N-1}$ of directions.
2. Construct the Voronoi tesselation $(\Omega_i)_{i=0}^{N-1}$ associated to the directions $(\omega_i)_{i=0}^{N-1}$.

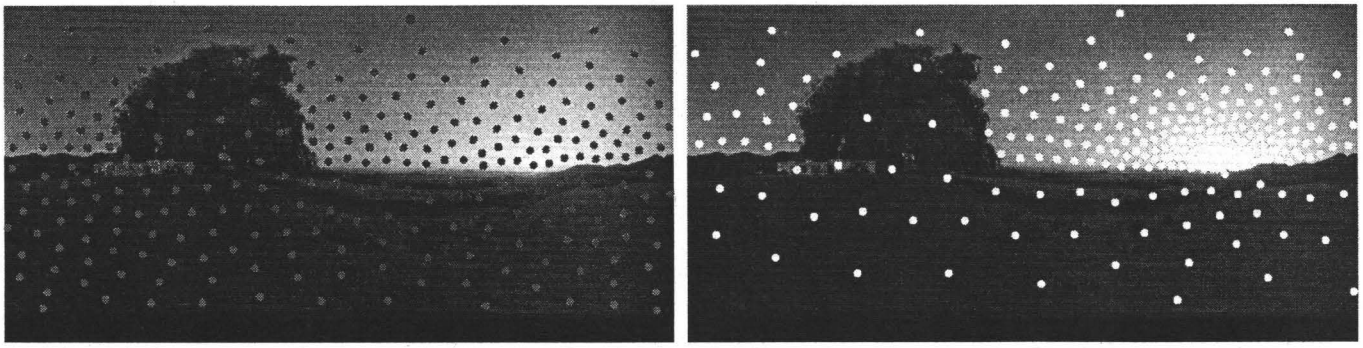

Figure 1: The $N = 256$ colored points in each image indicate the directions ω_i generated by Lloyd's relaxation algorithm on the left and our improved scheme on the right. For convenience the spherical images are displayed as 2:1 lattitude/longitude maps. Obviously the new approach captures the light distribution much more precisely resulting in a smaller integration error during rendering. Both images have been tone mapped for display.

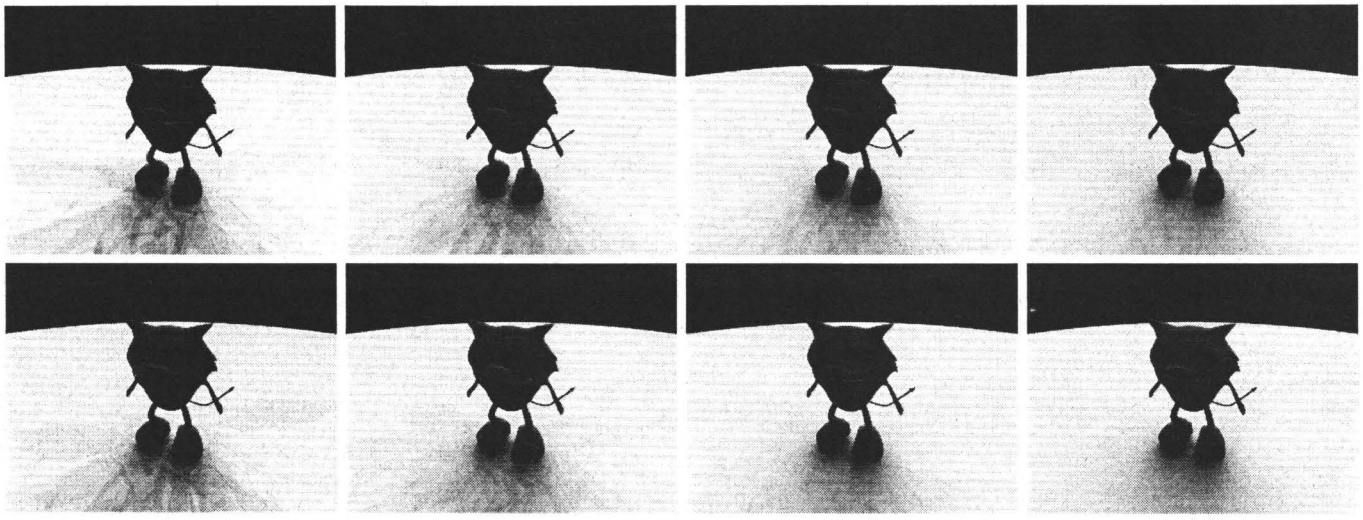

Figure 2: Images rendered using the quadrature rules generated by Lloyd's relaxation algorithm (top row) and our improved scheme (bottom row). The number of light sources is $N = 32, 64, 128, 256$ (from left to right). The shadow boundary artifacts caused by the directional light sources vanish much faster with our new scheme, which clearly indicates the faster decay of the integration error due to the more equalized and consequently smaller weights of the quadrature rule.

3. For each Voronoi region Ω_i replace ω_i by one of its mass centroidal directions.

4. If not terminated then go to step 2.

5. Compute the weights $(B_i)_{i=0}^{N-1}$ by (3).

The relaxation procedure is terminated if the maximum movement of all directions, i.e. the maximum angle of decentration, is below some threshold θ_T. For efficiency the Voronoi tesselation and the mass centroidal directions are approximated by uniformly sampling Ω. Thus the actual Ω_i implicitly are determined by the samples closest to ω_i [Hiller et al. 2001]. The average of these closest points is the corresponding mass centroid. Upon termination exactly these closest points are also used to approximate the quadrature weights B_i.

This straightforward algorithm is very sensitive to the initial choice of $(\omega_i)_{i=0}^{N-1}$ and becoming trapped by local minima is almost unavoidable. This is illustrated in the left image of figure 1, where far too many directions are located in unimportant regions.

This disadvantage easily can be overcome by incrementally determining the set of lighting directions. The key is to insert a new direction nearby the direction $\omega_{i'}$ with the maximum weight $\|B_{i'}\|$:

1. Set $n = 1$ and select a random direction ω_0.

2. Construct the Voronoi tesselation $(\Omega_i)_{i=0}^{n-1}$ associated to the directions $(\omega_i)_{i=0}^{n-1}$.

3. For each Voronoi region Ω_i replace ω_i by one of its mass centroidal directions.

4. If not terminated then go to step 2.

5. Compute the weights $(B_i)_{i=0}^{n-1}$ by (3).

6. If $n < N$ increase n by 1 and choose a new direction ω_{n-1} nearby the direction $\omega_{i'}$ with the maximum weight $\|B_{i'}\|$ and go to step 2.

3.1 Numerical Evidence

The images in figure 1 illustrate the placement of $N = 256$ directions. Clearly, the improved algorithm places more light sources in

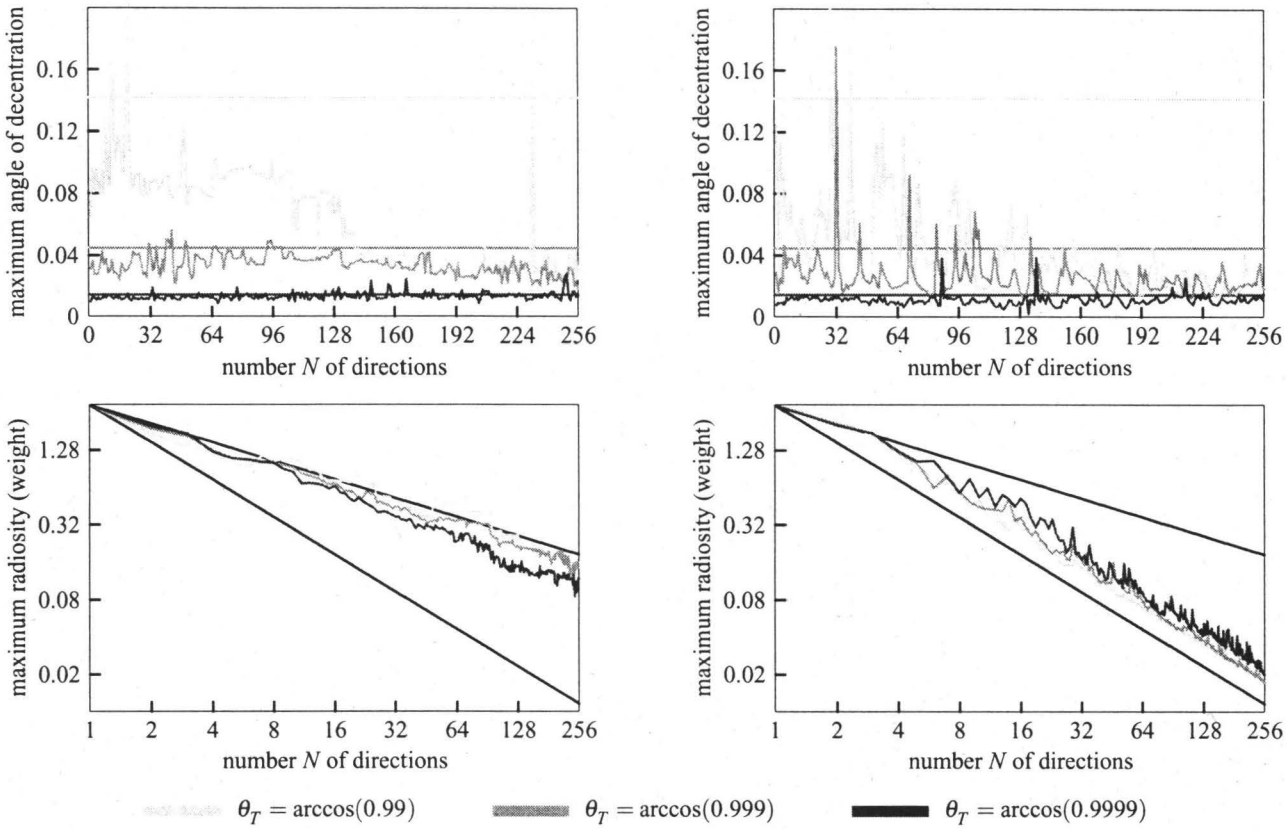

Figure 3: Comparison of the quadrature rules generated by Lloyd's relaxation algorithm on the left and our improved scheme on the right using different termination thresholds θ_T. The top row shows the maximum angle of decentration. Hardly noticeable the new scheme obtains a smaller decentration at a somewhat higher variance. However, as shown in the bottom row, the decay of the maximum radiosity is quadratically faster with the new approach. Instead of only $O\left(N^{-1/2}\right)$ almost the trivial lower bound of $O\left(N^{-1}\right)$ is achieved as indicated by the black lines.

the brighter regions of the high dynamic range image. The resulting quadrature rules have been used to render the images using (2) as shown in figure 2. While both quadrature rules are almost indistinguishable in unshadowed regions, the superiority of the improved quadrature rule becomes obvious in shadowed regions, which are rendered almost free of artifacts at already $N = 256$ directions. The increased preprocessing time reliably is compensated by the quality gain.

This result can be explained by comparing both quadrature rules with respect to the maximum radiosity $\max_{0\le i<N} \|B_i\|$ and the maximum angle of decentration after termination. This angle is given by $\max_{0\le i<N} \arccos\langle \omega_i, \gamma_i \rangle$, where γ_i is the centroidal direction of the Voronoi region corresponding to ω_i.

As shown in figure 3, forcing more iterations by lowering the termination threshold θ_T naturally decreases the maximum angle of decentration in the original approach. In the improved version the maximum angle of decentration hardly noticeable is smaller at a somewhat higher variance. At the same time the maximum radiosity practically remains unchanged independent of the approach.

However, the improved version almost achieves the trivial lower bound of $O\left(N^{-1}\right)$, which is quadratically better than the decay of the maximum radiosity in the original scheme. As already mentioned at the end of section 2, this efficiently decreases the integration error of (2).

4 Efficient Anti-Aliasing

Although our improved scheme results in much faster quadrature rules than previous approaches, using an identical quadrature rule for the whole image results in clearly visible shadow boundary artifacts if the number N of light sources is small (see figure 2). Therefore two methods are proposed that allow one to reduce the total number of shadow rays while efficiently preventing aliasing.

We briefly note that anti-aliasing is further improved by the powerful technique of interleaved sampling [Keller and Heidrich 2001; Molnar 1991]. However, the application to both methods proposed in the sequel is straightforward and therefore omitted for the sake of clarity.

4.1 Randomly Perturbed Quadrature Rules

At low sampling rates the shadow boundary artifacts can be turned into less objectionable noise by stratified sampling using the partition $\left(\Omega_i\right)_{i=0}^{N-1}$. However, for maximum efficiency the strata Ω_i are approximated by inscribed spherical caps defined by cones of radius α_i centered in ω_i. Thus it is easy to generate a random direction inside each cone. In order to decorrelate the samples this random perturbation of the quadrature rule has to be performed each time (2) is evaluated. It is important to note that the variance of the original high dynamic range image L_{hdr} restricted to the strata Ω_i remains

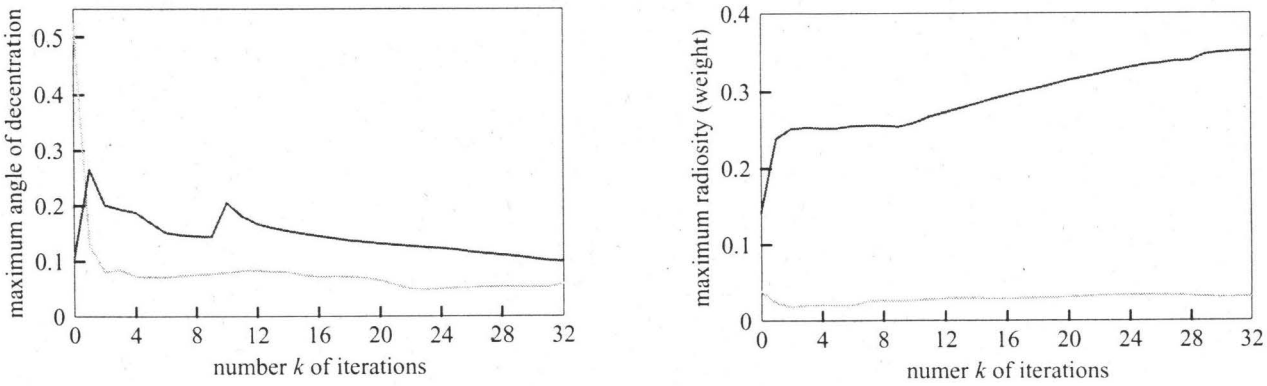

Figure 5: Quality of the interleaved quadrature rule after k iterations. The red lines indicate the maximum of the values of the $M = 4$ separate quadrature rules at $N = 32$ light sources, while the green lines indicate the values of the composite quadrature rule. The unavoidable loss of quality of the separate quadrature rules is kept minimal by our generation process.

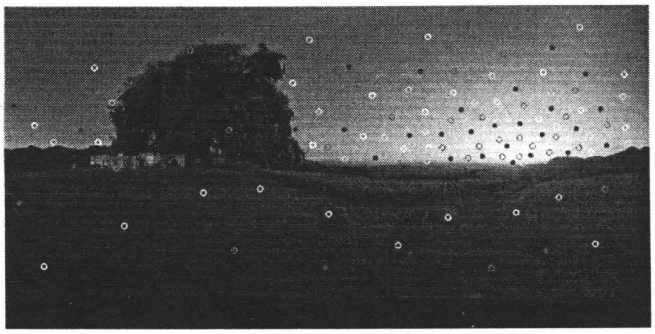

Figure 4: Interleaved quadrature rule consisting of $M = 4$ separate quadrature rules at $N = 32$ light sources. The colored points represent the directions $\omega_{j,i}$ of the light sources in the 2:1 lattitude/longitude map, while each color represents one quadrature rule. Our improved scheme precisely captures the light distribution by the total of $128 = 4 \cdot 32$ directions as well as every separate quadrature rule does as for example highlighted by the solid points. The image has been tone mapped for display.

high due to the fine image details. Therefore it has to be reduced, which is achieved by simply using the already filtered values B_i.

4.2 Interleaved Quadrature Rules

Due to the correlation coefficient of the integrand it pays off to send N shadow rays for a single eye ray, when computing the direct illumination. As shown in [Kollig and Keller 2002] maximum efficiency is achieved by correlated sampling.

We realize this kind of efficient multidimensional sampling for anti-aliasing by generating a separate quadrature rule $\left(\omega_{j,i}, B_{j,i}\right)_{i=0}^{N-1}$ for each eye ray $j = 0, \ldots, M-1$ and imposing the two constraints that both

- each separate quadrature rule $\left(\omega_{j,i}, B_{j,i}\right)_{i=0}^{N-1}$ for $j = 0, \ldots, M-1$ and
- the composite quadrature rule $\left(\omega_{j,i}, \widetilde{B}_{j,i}\right)_{j,i=0}^{M-1,N-1}$

fulfill the requirements of section 2. Note that both constraints use the same directions, however, different quadrature weights $B_{j,i}$ and $\widetilde{B}_{j,i}$ result from (3) as induced by the Voronoi tesselations of $\left(\omega_{j,i}\right)_{i=0}^{N-1}$ for $j = 0, \ldots, M-1$ and $\left(\omega_{j,i}\right)_{j,i=0}^{M-1,N-1}$, respectively.

Figure 4 shows an example of an interleaved quadrature rule composited out of $M = 4$ separate quadrature rules with $N = 32$ light sources each that has been generated using the following algorithm:

1. Generate an initial quadrature rule $\left(\omega_{0,i}, B_{0,i}\right)_{i=0}^{N-1}$ following section 3.
2. For $j = 1, 2, \ldots, M-1$ construct the quadrature rule $\left(\omega_{j,i}, B_{j,i}\right)_{i=0}^{N-1}$ by randomly perturbing the directions $\omega_{0,i}$.
3. For $j = 0, 1, \ldots, M-1$ construct the Voronoi tesselation $(\Omega_{j,i})_{i=0}^{N-1}$ of Ω associated to the directions $(\omega_{j,i})_{i=0}^{N-1}$.
4. For each Voronoi region $\Omega_{j,i}$ replace $\omega_{j,i}$ by one of its mass centroidal directions.
5. Construct the Voronoi tesselation $\left(\widetilde{\Omega}_{j,i}\right)_{j,i=0}^{M-1,N-1}$ of Ω associated to the directions $(\omega_{j,i})_{j,i=0}^{M-1,N-1}$.
6. For each Voronoi region $\widetilde{\Omega}_{j,i}$ replace $\omega_{j,i}$ by one of its mass centroidal directions.
7. k times iterate steps 3 through 6.
8. Compute the weights $\left(B_{j,i}\right)_{j,i=0}^{M-1,N-1}$ by (3).

For the initialization step 2 the radius α of the cones of perturbation (see also section 4.1) is chosen small, e.g. $\alpha = 0.01$. The alternate application of Lloyd's relaxation step to the separate quadrature rules in steps 3 and 4 and to the composite quadrature rule in steps 5 and 6 achieves to satisfy both constraints as mentioned in the beginning of this section.

Figure 5 shows the development of the quality of both the separate quadrature rules and the composite quadrature rule. Upon initialization the maximum angle of decentration is small for the separate quadrature rules and due to a lack of correlation huge for the composite quadrature rule. Then the iteration process efficiently increases correlation forcing a higher but nevertheless decreasing maximum angle of decentration in the separate quadrature rules.

In consequence the maximum radiosities of the separate quadrature rules must be slightly increasing. At the same time, however, the maximum radiosity of the composite quadrature rule remains small. In practice only about $k = 4$ iteration steps are sufficient to obtain a high quality interleaved quadrature rule.

4.3 Numerical Evidence

Figure 6 compares our new rendering methods for computing the direct illumination by a high dynamic range image L_{hdr}. The high variance inherent to L_{hdr} even cannot be efficiently reduced by randomized quasi-Monte Carlo [Kollig and Keller 2002] when computing (1) and consequently strong noise remains visible. Obviously one single quadrature rule (see section 3) exposes clear shadow boundaries, which are transferred to less perceivable noise by randomly perturbed quadrature rules (see section 4.1). Finally the interleaved quadrature rule (see section 4.2) reduces the shadow artifacts to an imperceivable level. While randomly perturbed quadrature rules may be favored for rendering still images, interleaved quadrature rules are the better choice for animations, where coherent sampling over time avoids flicker.

Although each method has used 32 samples for computing the direct illumination of each primary ray in our new rendering methods about 15% of the shadow rays do not need to be shot, because they can be culled using the surface normal. In addition better memory coherence and less pseudo random number generator calls are the reasons for up to 50% reduced rendering times.

Note that pseudo random number generation can consume considerable amounts of time, which also is the reason for the 25% increase of rendering time when randomly perturbing the single quadrature rule. Considering the improved image quality at reduced rendering times the preprocessing of 23 seconds for a single or 75 seconds for an interleaved quadrature rule clearly pays off.

5 Conclusion

We presented a robust algorithm for the efficient computation of direct illumination of predominantly diffuse surfaces by spherical high dynamic range images. The scheme is designed for an industrial production environment, where high resolution (up to 50 mega pixels) spherical high dynamic range scans are used.

Without manual intervention our new scheme generates quadrature rules with almost equal weights resulting in a considerable noise reduction and improved anti-aliasing. As shown in the examples in figures 7 – 10 the light distribution is captured very precisely and such it is unnecessary to e.g. cut out light sources manually as in [Debevec et al. 2002] in order to reduce variance. The resulting quadrature rules consume negligible amounts of memory and can be stored along with the high dynamic range images saving repeated preprocessing. Similar to [Debevec 1998] the directional light sources can be projected back onto a reconstructed geometry thus allowing for even more precise shadowing.

An obvious improvement to our implementation is the acceleration of the computation of the mass centroids and the weights of the quadrature rules by graphics hardware. Although specular surfaces are simple to render, the efficient combination with our techniques is not straightforward for general e.g. non-additive reflection properties. This, as well as the generation of caustics, is subject to future research. We also will focus on combining quadrature rules from multiple high dynamic range images and on enabling our scheme for illumination by high dynamic range video.

Acknowledgements

The first author has been funded by the Stiftung Rheinland-Pfalz für Innovation. The research has been supported by Spheron VR AG.

References

DEBEVEC, P., AND MALIK, J. 1997. Recovering High Dynamic Range Radiance Maps from Photographs. In *SIGGRAPH 97 Conference Proceedings*, Addison-Wesley, T. Whitted, Ed., Annual Conference Series, ACM SIGGRAPH, 369–378.

DEBEVEC, P., FONG, N., AND LEMMON, D. 2002. Image-Based Lighting. In *SIGGRAPH 2002 Course Notes*, no. 5. ACM SIGGRAPH, July.

DEBEVEC, P. 1998. Rendering Synthetic Objects Into Real Scenes: Bridging Traditional and Image-Based Graphics With Global Illumination and High Dynamic Range Photography. In *SIGGRAPH 98 Conference Proceedings*, Addison-Wesley, M. Cohen, Ed., Annual Conference Series, ACM SIGGRAPH, 189–378.

DU, Q., FABER, V., AND GUNZBURGER, M. 1999. Centroidal Voronoi Tesselations: Applications and Algorithms. *SIAM Review 41*, 4 (Dec.), 637–676.

HILLER, S., DEUSSEN, O., AND KELLER, A. 2001. Tiled Blue Noise Samples. In *Proceedings of Vision, Modeling, and Visualization 2001*, IOS Press, 265–271.

JENSEN, H. 1995. Importance Driven Path Tracing Using the Photon Map. In *Rendering Techniques 95 (Proc. 6th Eurographics Workshop on Rendering)*, Springer, P. Hanrahan and W. Purgathofer, Eds., 326–335.

KELLER, A., AND HEIDRICH, W. 2001. Interleaved Sampling. In *Rendering Techniques 2001 (Proc. 12th Eurographics Workshop on Rendering)*, Springer, K. Myszkowski and S. Gortler, Eds., 269–276.

KOLLIG, T., AND KELLER, A. 2002. Efficient Multidimensional Sampling. *Computer Graphics Forum 21*, 3 (Sept.), 557–563.

MOLNAR, S. 1991. Efficient Supersampling Antialiasing for High-Performance Architectures. Tech. Rep. TR91-023, The University of Noth Carolina at Chapel Hill.

RAMAMOORTHI, R., AND HANRAHAN, P. 2001. An Efficient Representation for Irradiance Environment Maps. In *Proceedings of ACM SIGGRAPH 2001*, E. Fiume, Ed., Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH, 497–500.

RAMAMOORTHI, R., AND HANRAHAN, P. 2002. Frequency Space Environment Map Rendering. *ACM Transaction on Graphics 21*, 3 (July), 517–526.

SLOAN, P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed Radiance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments. *ACM Transaction on Graphics 21*, 3 (July), 527–536.

SPHERONVR AG. http://www.spheron.com.

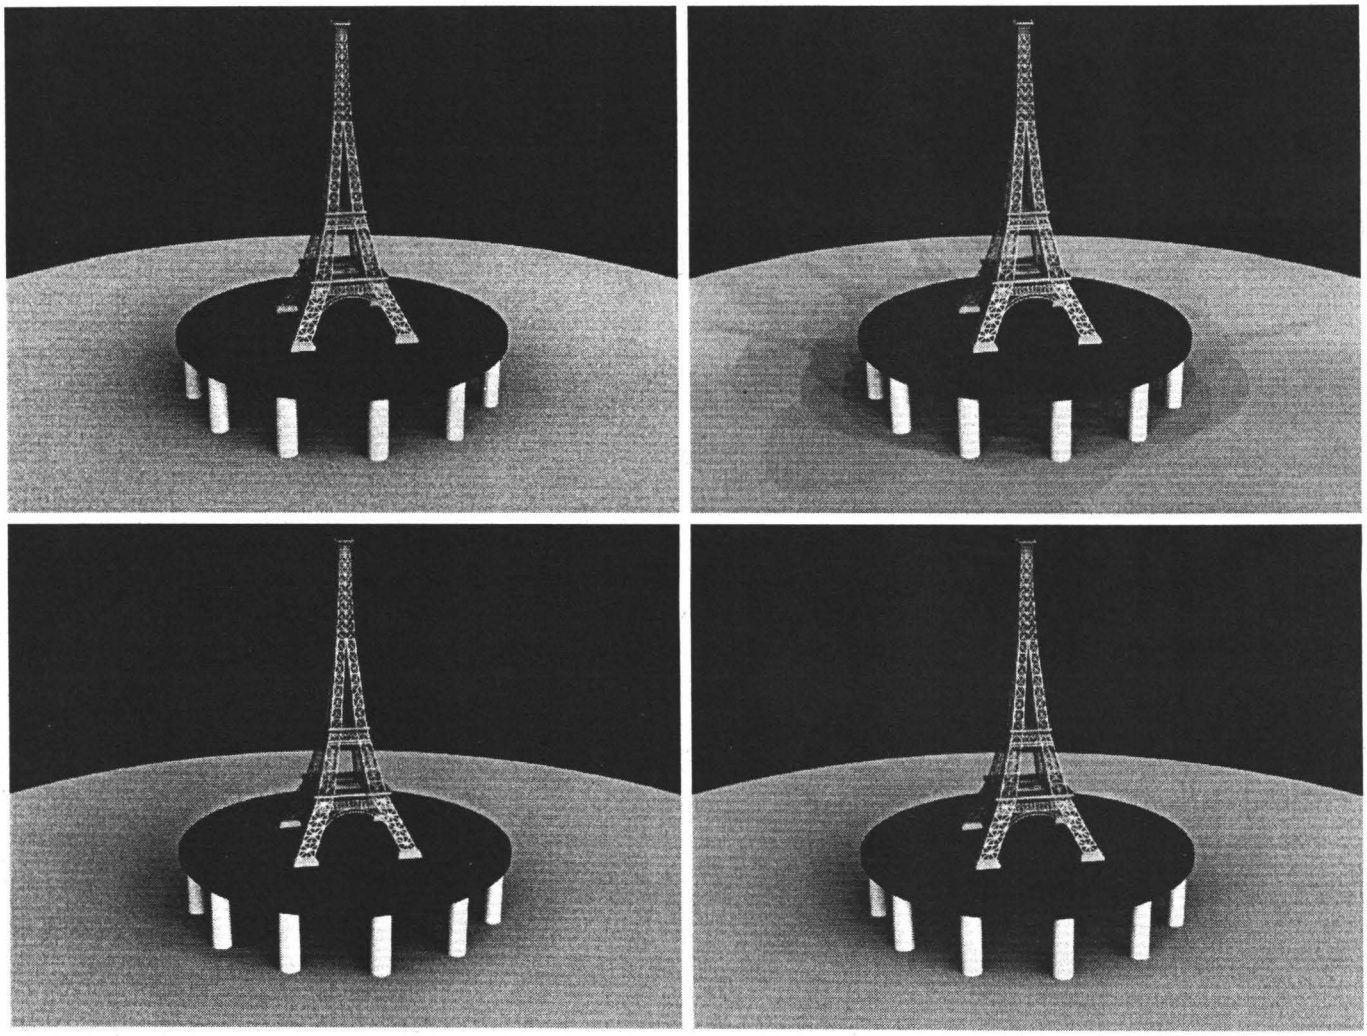

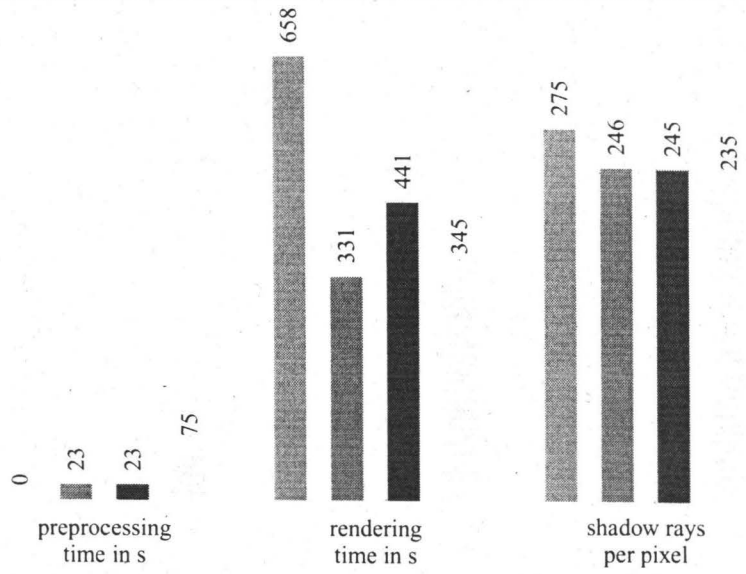

Figure 6: Comparison of direct illumination computed by randomized quasi-Monte Carlo (upper left, red bars), a single quadrature rule (upper right, cyan bars), a randomly perturbed quadrature rule (lower left, magenta bars), and an interleaved quadrature rule (lower right, yellow bars). For anti-aliasing 16 primary rays per pixel were traced. For each primary ray 32 shadow rays were shot into the hemisphere in the randomized quasi-Monte Carlo solution. The single and the randomly perturbed quadrature rule had $N = 32$ light sources and the interleaved quadrature rule consisted of $M = 16$ separate quadrature rules with $N = 32$ light sources each. On the left the preprocessing and rendering times on a 650 MHz Pentium III for the above images at resolution 640×480 and the number of shadow rays per pixel are displayed.

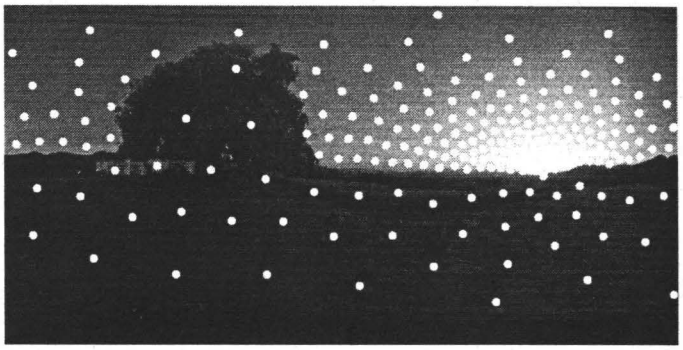

(a) single quadrature rule with $N = 256$

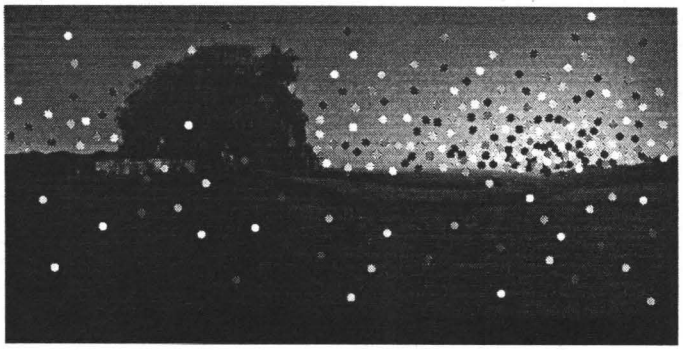

(b) interleaved quadrature rule with $M = 8$ and $N = 32$

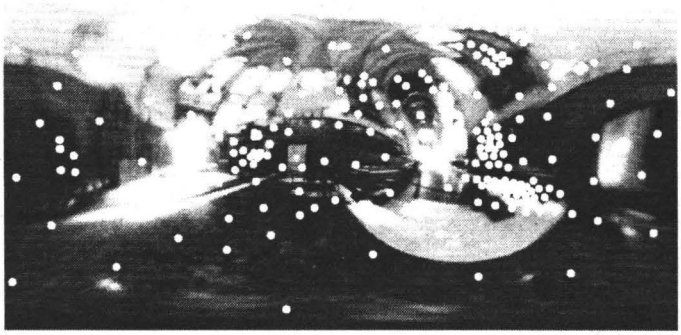

(c) single quadrature rule with $N = 256$

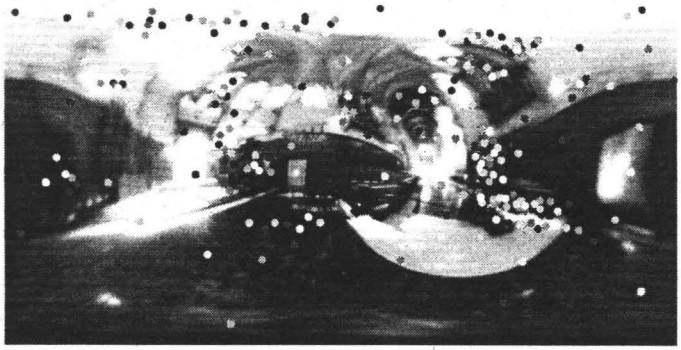

(d) interleaved quadrature rule with $M = 8$ and $N = 32$

Figure 7: The top row example demonstrates our method for a predominantly continuous light distribution. The Grace Cathedral example (courtesy P. Debevec) in the bottom row impressively shows that our new method reliably captures bright regions and discontinuities of the high dynamic range image. Note that no manual intervention like e.g. replacing the bright regions by area light sources (which would be even more tedious in the top row example) is required. All images have been tone mapped for display.

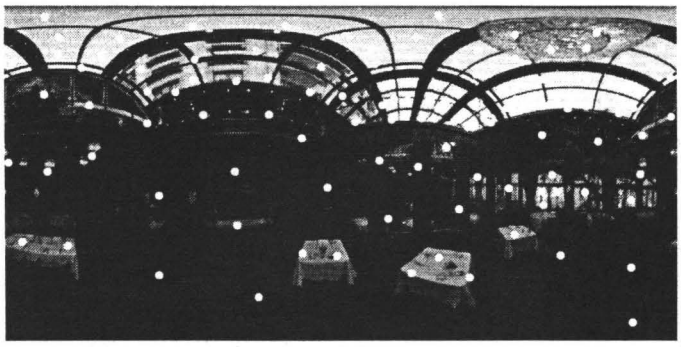

(a) single quadrature rule with $N = 128$

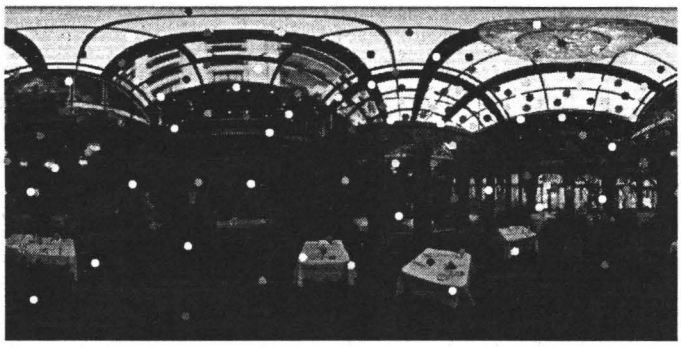

(b) interleaved quadrature rule with $M = 4$ and $N = 32$

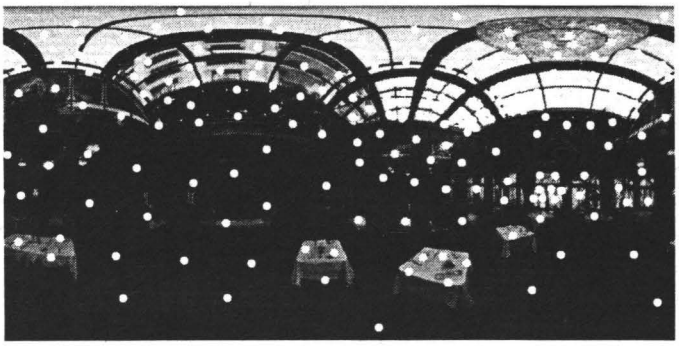

(c) single quadrature rule with $N = 256$

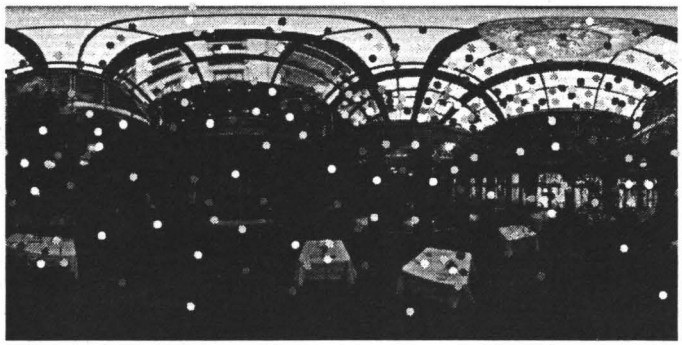

(d) interleaved quadrature rule with $M = 8$ and $N = 32$

Figure 8: More examples. All images have been tone mapped for display.

(a) single quadrature rule with $N = 128$

(b) interleaved quadrature rule with $M = 4$ and $N = 32$

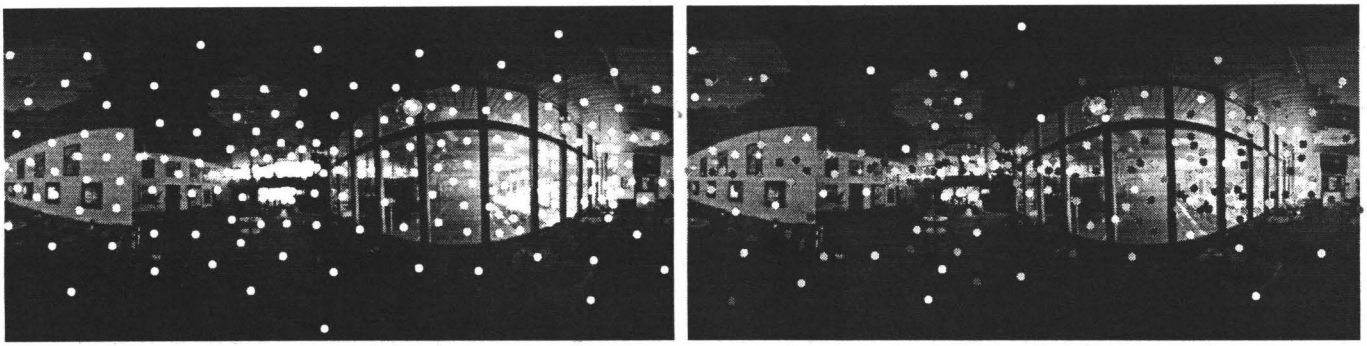

(c) single quadrature rule with $N = 256$

(d) interleaved quadrature rule with $M = 8$ and $N = 32$

Figure 9: More examples. Note that for $N = 256$ even the small spot lights, acting as point light sources, are captured. All images have been tone mapped for display.

(a) single quadrature rule with $N = 128$

(b) interleaved quadrature rule with $M = 4$ and $N = 32$

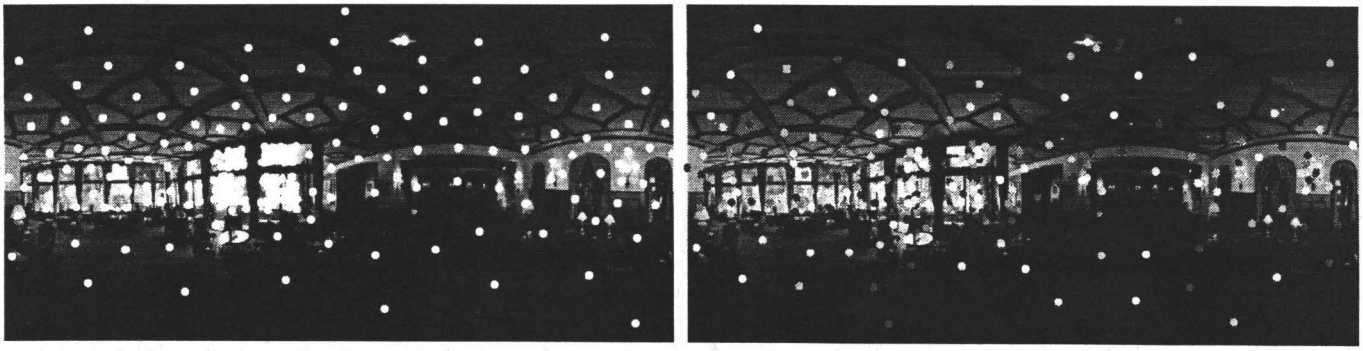

(c) single quadrature rule with $N = 256$

(d) interleaved quadrature rule with $M = 8$ and $N = 32$

Figure 10: More examples. All images have been tone mapped for display.