

Efficient Illumination by High Dynamic Range Images

Thomas Kollig¹ and Alexander Keller²

¹ Department of Computer Science, University of Kaiserslautern, Germany

² Department of Computer Science, University of Ulm, Germany

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^{2, 15}. This technique has been introduced by Debevec et al.^{4, 3}, is established in commercial products, and has been used in many movies.

In this context 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. 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 highly directional 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: On the one hand adaptive integration schemes like e.g. the VEGAS or MISER technique¹¹ are inefficient due to excessive additional book

keeping data structures. On the other hand importance sampling on the discrete set of directional light sources given by the pixels of the high dynamic range image still results in high variance as will be shown in section 5. In addition approximations⁷ usually do not perform any better than pure random sampling in this case. We also want to avoid variance reduction by manual stratification³. Projecting the integral operator into the spherical harmonics basis^{12, 13} allows for fast rendering. However, including detailed shadows in this model imposes considerable cost and restrictions¹⁴.

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. Although seemingly similar to LightGen¹, our technique yields an almost optimal integration scheme by removing some flaws of the straightforward approach. This is demonstrated by numerical evidence. Furthermore new superior anti-aliasing techniques are developed (section 4) that reproduce fine shadow details much more precisely and are faster than previous approaches.

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 .




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 latitude/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) as illustrated in figure 1. 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, clearly indicating the faster decay of the integration error due to the more equalized and consequently smaller weights of the quadrature rule.

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.

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

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 integration error of (2) if the visibility term V is neglected and
- $\max_{0 \leq i < N} \|B_i\|$ 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 and the second one bounds

$$\sum_{i=0}^{N-1} \left(B_i - \frac{1}{N} \sum_{i=0}^{N-1} B_i \right)^2 = \sum_{i=0}^{N-1} \left(B_i - \frac{1}{N} \int_{\Omega} L_{\text{hdr}}(\omega) d\omega \right)^2.$$

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 method⁵ on the sphere:

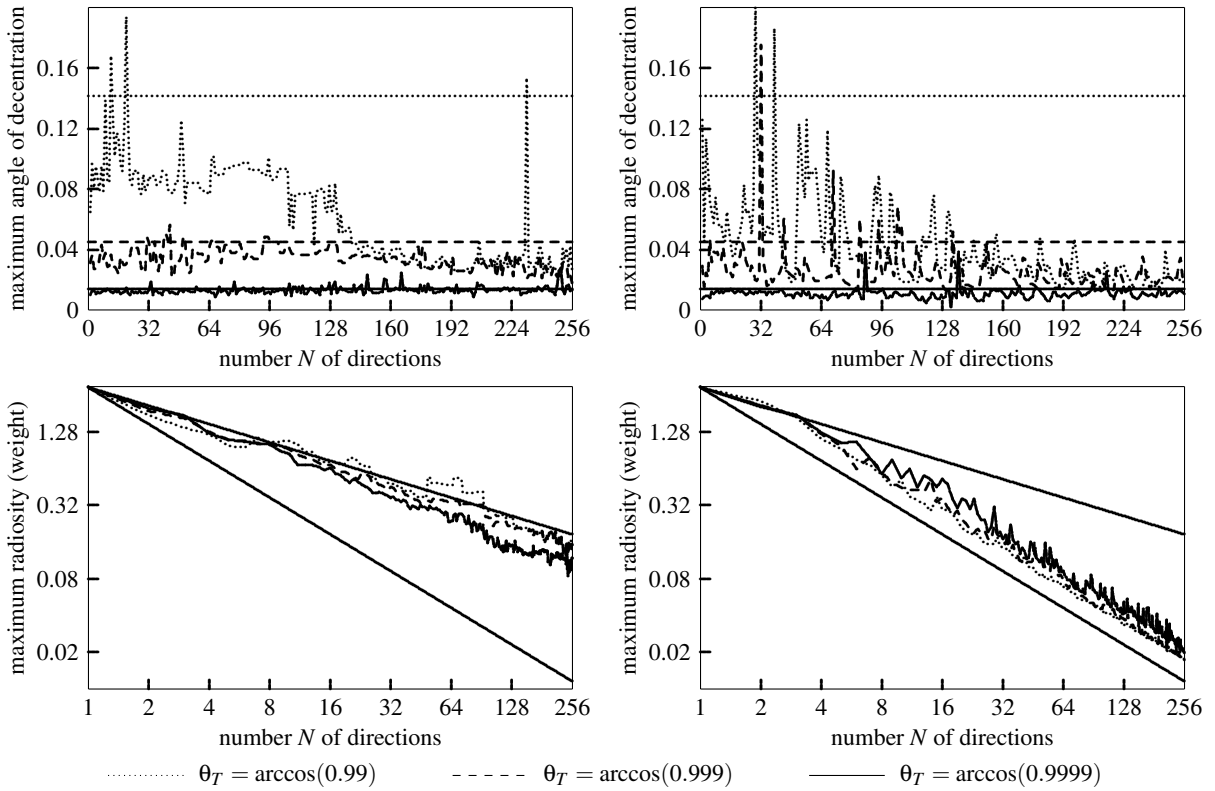


Figure 3: Comparison of the quadrature rules generated by Lloyd's relaxation algorithm on the left and our improved scheme on the right using different termination thresholds θ_T . The top row shows the maximum angle of decentration. Hardly noticeable the new scheme obtains a smaller decentration at 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 $\mathcal{O}(N^{-1/2})$ almost the trivial lower bound of $\mathcal{O}(N^{-1})$ is achieved as indicated.

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

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 tessellation and the mass centroidal directions are approximated by uniformly sampling Ω . Thus the actual Ω_i implicitly are determined⁶ by the samples closest to ω_i . 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 technique in fact is very similar to k -means clustering as used in LightGen¹.

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 tessellation $(\Omega_i)_{i=0}^{n-1}$ 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 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 $\max_{0 \leq i < N} \|B_i\|$ and the maximum angle of decentration after termination. This angle is given by $\max_{0 \leq i < N} \arccos\langle \omega_i, \gamma_i \rangle$, 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 is marginally 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 $\mathcal{O}(N^{-1})$, 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 (see section 3), 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.

Note that anti-aliasing is further improved by the powerful technique of interleaved sampling^{8, 10}. However, the application of interleaved sampling 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 $(\Omega_i)_{i=0}^{N-1}$. 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

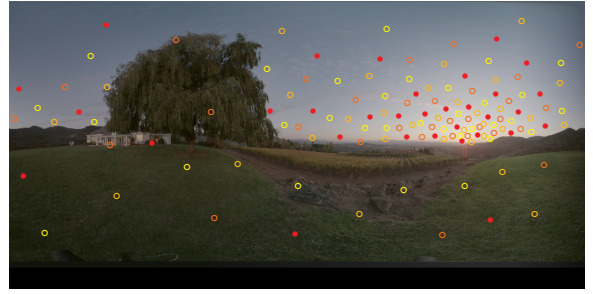


Figure 4: Interleaved quadrature rule consisting of $M = 4$ separate quadrature rules at $N = 32$ light sources. The colored points represent the directions $\omega_{j,i}$ of the light sources in the 2:1 latitude/longitude map, while each color represents one quadrature rule. Our improved scheme precisely 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.

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 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 with respect to the shadow rays⁹ it pays off to send N shadow rays for a single primary ray, when computing the direct illumination. Maximum efficiency is achieved by correlated sampling⁹.

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

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

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

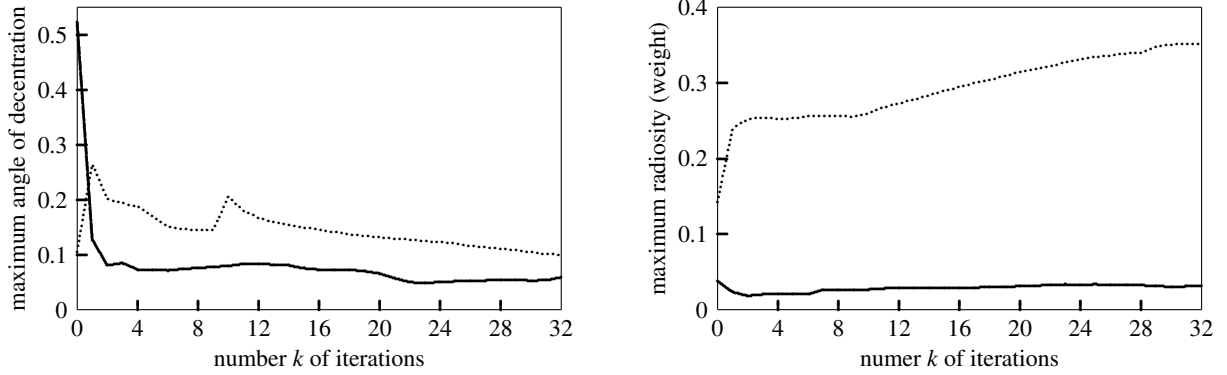


Figure 5: Quality of the interleaved quadrature rule after k iterations. The dotted lines indicate the maximum of the values of the $M = 4$ separate quadrature rules at $N = 32$ light sources, while the solid lines indicate the values of the composite quadrature rule. The unavoidable loss of quality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

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

5. Results

Figure 6 compares rendering methods for computing the direct illumination by a high dynamic range image L_{hdr} . For ease of comparison the sunset 2:1 latitude/longitude scan (see figures 1, 4, and 8) has been used.

The images on the left of figure 6 show the results of sampling the hemisphere (mid-left) and importance sampling the high dynamic range image (lower left) identifying each pixel as one directional light source. The high variance of both estimators even cannot be sufficiently reduced by randomized quasi-Monte Carlo⁹ and consequently strong noise remains visible. However, as expected importance sampling the discrete set of light sources (lower left) is less noisy and avoids disturbing spike noise (see the enlargements in figure 6).

The images on the right of figure 6 are rendered using our new techniques. Obviously one single quadrature rule (see section 3) exposes clear shadow boundary artifacts (upper right), which are transferred to less perceivable noise (mid-right) by randomly perturbed quadrature rules (see section 4.1). Finally the interleaved quadrature rule (see section 4.2) reduces the shadow artifacts to an imperceptible level (lower right). 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 up to 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 perturb-

ing 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. On the other hand obtaining the same quality without interleaved quadrature rules would require a single quadrature rule with much more light sources that in consequence would take much longer time to generate.

6. 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 industrial production, where high resolution spherical high dynamic range scans are used (see figure 7).

Without manual intervention our new scheme generates quadrature rules with minimized weights resulting in a considerable noise reduction and improved anti-aliasing. As shown in figure 8 the light distribution is captured very precisely. Thus it becomes redundant to e.g. cut out light sources manually³ 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. The directional light sources can be projected back onto a reconstructed geometry² 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 reflections properties, which cannot be represented as a weighted sum of basic bidirectional distribution functions. 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 AG¹⁵.

References

1. Cohen, J., and Debevec, P. "Light-Gen" HDRShop Plugin. 2001. <http://www.ict.usc.edu/~jcohen/lightgen/lightgen.html>.
2. P. Debevec. Rendering Synthetic Objects Into Real Scenes: Bridging Traditional and Image-Based Graphics With Global Illumination and High Dynamic Range Photography. In M. Cohen, editor, *SIGGRAPH 98 Conference Proceedings*, Annual Conference Series, pages 189–378. ACM SIGGRAPH, Addison-Wesley, July 1998.
3. P. Debevec, N. Fong, and D. Lemmon. Image-Based Lighting. In *SIGGRAPH 2002 Course Notes*, number 5. ACM SIGGRAPH, July 2002.
4. P. Debevec and J. Malik. Recovering High Dynamic Range Radiance Maps from Photographs. In T. Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 369–378. ACM SIGGRAPH, Addison-Wesley, August 1997.
5. Q. Du, V. Faber, and M. Gunzburger. Centroidal Voronoi Tesselations: Applications and Algorithms. *SIAM Review*, 41(4):637–676, December 1999.
6. S. Hiller, O. Deussen, and A. Keller. Tiled Blue Noise Samples. In *Proceedings of Vision, Modeling, and Visualization 2001*, pages 265–271. IOS Press, 2001.
7. H. Jensen. Importance Driven Path Tracing Using the Photon Map. In P. Hanrahan and W. Purgathofer, editors, *Rendering Techniques 95 (Proc. 6th Eurographics Workshop on Rendering)*, pages 326–335. Springer, 1995.
8. A. Keller and W. Heidrich. Interleaved Sampling. In K. Myszkowski and S. Gortler, editors, *Rendering Techniques 2001 (Proc. 12th Eurographics Workshop on Rendering)*, pages 269–276. Springer, 2001.
9. T. Kollig and A. Keller. Efficient Multidimensional Sampling. *Computer Graphics Forum*, 21(3):557–563, September 2002.
10. S. Molnar. Efficient Supersampling Antialiasing for High-Performance Architectures. Technical Report TR91-023, The University of North Carolina at Chapel Hill, 1991.
11. W. Press, W. Vetterling, S. Teukolsky, and B. Flannery. *Numerical Recipes in C++: the Art of Scientific Computing*. Cambridge University Press, 2nd edition, 2002.
12. R. Ramamoorthi and P. Hanrahan. An Efficient Representation for Irradiance Environment Maps. In E. Fiume, editor, *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, pages 497–500. ACM SIGGRAPH, August 2001.
13. R. Ramamoorthi and P. Hanrahan. Frequency Space Environment Map Rendering. *ACM Transaction on Graphics*, 21(3):517–526, July 2002.
14. P. Sloan, J. Kautz, and J. Snyder. Precomputed Radiance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments. *ACM Transaction on Graphics*, 21(3):527–536, July 2002.
15. SpheronVR AG. <http://www.spheron.com>.



Figure 7: Reality check: A virtual car has been placed into a real set recorded by one high resolution high dynamic range image. The scanned image is used both as background shot and for our fully automatic illumination technique.



(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 8: 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.

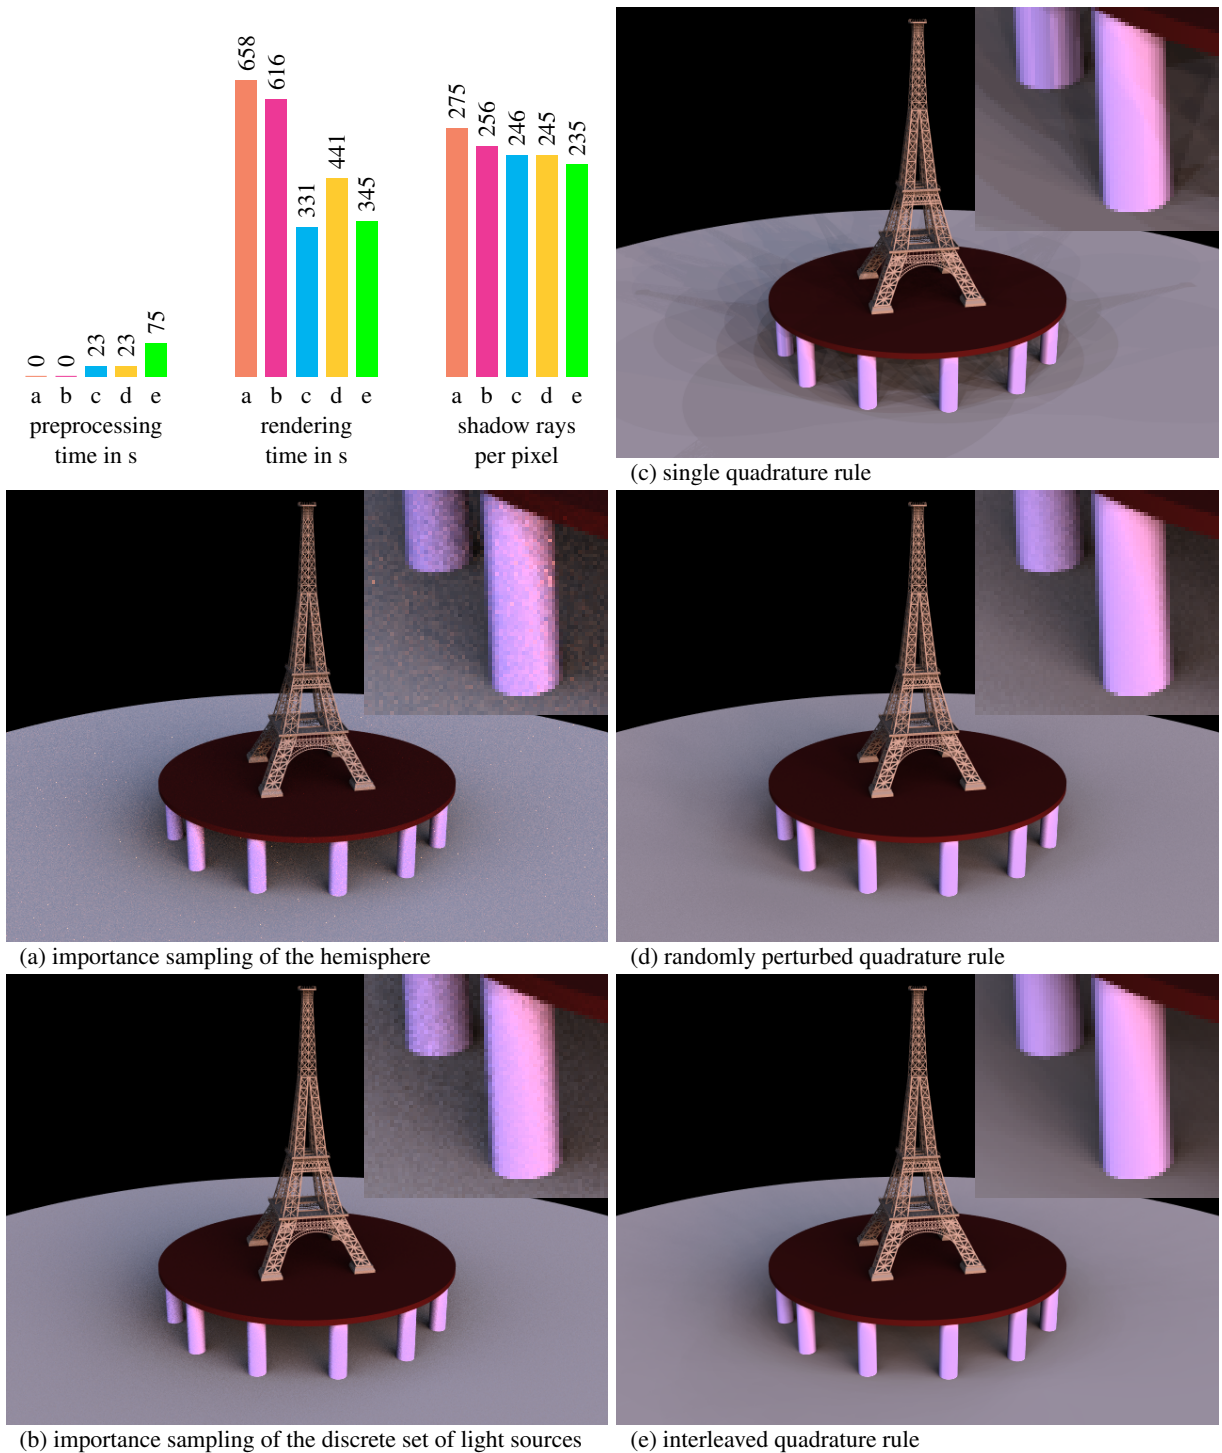


Figure 6: Comparison of direct illumination computed by (a) importance sampling of the hemisphere, (b) importance sampling of the discrete set of light sources, (c) a single, (d) a randomly perturbed, and (e) an interleaved quadrature rule. For anti-aliasing 16 primary rays per pixel were traced. For each primary ray 32 samples were used for the techniques on the left. 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. The diagram displays the preprocessing and rendering times on a 650 MHz Pentium III for the above images at a resolution of 640×480 pixels and the number of shadow rays per pixel.