

All-Frequency Shadows Using Non-linear Wavelet Lighting Approximation

Ren Ng
Stanford University

Ravi Ramamoorthi
Columbia University

Pat Hanrahan
Stanford University

Abstract

We present a method, based on pre-computed light transport, for real-time rendering of objects under all-frequency, time-varying illumination represented as a high-resolution environment map. Current techniques are limited to small area lights, with sharp shadows, or large low-frequency lights, with very soft shadows. Our main contribution is to approximate the environment map in a wavelet basis, keeping only the largest terms (this is known as a *non-linear approximation*). We obtain further compression by encoding the light transport matrix sparsely but accurately in the same basis. Rendering is performed by multiplying a sparse light vector by a sparse transport matrix, which is very fast. For accurate rendering, using non-linear wavelets is an order of magnitude faster than using linear spherical harmonics, the current best technique.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism; G.1.2 [Numerical Analysis]: Approximation—Nonlinear Approximation, Linear Approximation; G.1.4 [Numerical Analysis]: Quadrature—Error Analysis

Keywords: Shadow Algorithms, Relighting, Image-Based Rendering, Wavelets, Non-linear Approximation, Spherical Harmonics

1 Introduction

In story-telling applications such as movies and games, shadows serve not only as a practical tool for establishing depth and shape, but also as an artistic device for refining mood and character. Subtle lighting design requires the ability to work with lights of different frequencies, or put another way, lights of different sizes. This paper describes real-time relighting of rigid scenes with complex shadowing, under high-resolution illumination that may contain all frequencies. For lighting design from a fixed viewpoint, our technique supports arbitrary reflection models and light transport effects. For applications that require changing viewpoint, we currently handle only transport effects that terminate on diffuse surfaces.

Our goals and general approach are similar to the pre-computed radiance transfer method of Sloan et al. [2002] for low-frequency environments, which extends earlier techniques such as irradiance environment maps [Ramamoorthi and Hanrahan 2001]. Sloan et al. pre-compute light transport in a spherical harmonic basis, and render by calculating a dot product between each vertex's pre-computed transport vector and the harmonic coefficients of the incoming lighting. By using a fixed set of approximating functions for the lighting signal, Sloan et al. employ what is known as a *linear approximation* [Mallat 1998] in a spherical harmonic basis. Their




Figure 1: Comparison of a linear spherical harmonic lighting approximation (left) [Sloan et al. 2002] and our non-linear wavelet approximation. We scale the linear technique to use as many harmonics as wavelets (100). Notice the superior definition of wavelet shadows, which closely match a reference image. The environment depicts the interior of St Peter's Basilica.

approach provides compact storage because high-frequency information is discarded from the transfer matrix, but it is only accurate for very low-frequency lighting.

Our approach differs in that we do not choose our set of approximating functions a priori, but rather choose from a space of tens of thousands based on the given illumination. This process is known as *non-linear approximation*. DeVore [1998] has written a survey of such techniques.

One might consider the use of a spherical harmonic basis for nonlinear approximation, but harmonics do not localize well in space. A small area light would require many harmonics for accurate approximation. Conversely, one might consider using a basis of point-like lights; in fact, techniques such as layered attenuation maps [Agrawala et al. 2000] and the method of Heckbert and Herf [1997] take this approach. However, point lights do not localize well in frequency, and these techniques are limited to small area sources. In addition, nobody has yet found a principled way to represent an arbitrary environment map as a combination of point sources and spherical harmonics.

For these reasons, we choose to use a non-linear *wavelet* approximation. A wavelet basis contains area lights that vary from the size of a cubemap pixel to essentially the size of the entire sky. Consequently, a non-linear wavelet approximation can efficiently represent lighting with features at all frequencies. We show that just 0.5–1% of the wavelet basis functions accurately approximate detailed photographed illumination [Debevec and Malik 1997], and that 0.1–0.2% suffice for simpler environments or synthetic area lights (see our video on the Full Conference DVD-ROM). As an

aside, another important factor in choosing wavelets over spherical harmonics is the speed of their respective transforms: linear time for wavelets, and practically $O(n^{3/2})$ for harmonics, since the asymptotically faster $O(n \log^2 n)$ algorithm is actually slower at these relatively low resolutions.

An important aspect of our approach is that we pre-compute very high-resolution transport matrices that encode exact transport from any $6 \times 64 \times 64$ cubical environment map. We approximate rendering accurately and efficiently as a giant, sparse matrix multiplication where sparsity derives from two sources: non-linear approximation of the lighting vector, and zero-coefficients within the encoded matrix. The latter is a secondary but important aspect of our technique. Since we encode our transport matrix in a wavelet lighting basis, coherence in the transport operator results in a matrix where only a fraction of the elements are non-zero.

In summary, our main contributions consist of:

- Non-linear approximation of the lighting using a wavelet basis. Because wavelets provide locality in both angular and frequency space, the encoded light vector and light transport matrix are accurately approximated with a reasonable number of terms.
- Interactive rendering of scenes under time-varying, all-frequency lighting. Hard and soft shadows are both captured well (see Figures 1, 3 and 4). Real-time performance is achieved through efficient sparse matrix multiplication.

2 Previous Work

Image relighting and shadowing techniques have a long history in computer graphics. Recently, a number of general soft shadow methods have been proposed, of which perhaps the most relevant are convolution textures [Soler and Sillion 1998], layered attenuation maps [Agrawala et al. 2000] and Heckbert and Herf’s method [1997]. Our method may be considered a generalization of these techniques. The methods developed by Agrawala et al. and Heckbert and Herf use roughly 100 samples to resolve shadows from small area lights; we use roughly 100 wavelet lights to accurately relight from detailed environments. Furthermore, our technique allows real-time, dynamic relighting, which these techniques do not, and naturally supports complex, self-shadowing geometry (see Figure 3), which is problematic for convolution textures.

In image-relighting, most existing techniques make use of linear lighting approximations. Dorsey et al. [1991], Nimeroff et al. [1994], Teo et al. [1997], Ashikhmin and Shirley [2002] and Debevec et al. [2000] all use pre-determined sets of lighting approximation functions, such as points, steerable functions and compressed principal component bases. As with Sloan et al. [2002], commitment to a linear approximation enables compact representation, but limits accuracy.

Debevec et al. [2000] use the highest environment map resolution — roughly 2000 directions, but they focus on largely convex objects without complex shadowing. Their compression scheme is similar to ours; they use JPEG compression to achieve a compression of 20:1. In contrast, we compress the relighting computation 100:1 or 1000:1 by using wavelets to approximate the incoming lighting as well as encoding the transport matrix sparsely. We report real-time results for 24,576 ($6 \times 64 \times 64$) lighting directions.

Our method may bear similarities to a number of light field compression techniques, such as wavelet light fields [Lalonde and Fournier 1999], surface light fields [Wood et al. 2000], and factored forms such as those used in light field mapping [Chen et al. 2002]. The fundamental difference is that these methods deal with static lighting, considering fast *look-up* of a single viewpoint from a table of views, and we focus on real-time relighting, requiring

fast *summation* of illumination. It is not clear that explicit tabular representations or previous compression schemes give real-time performance for this application.

Finally, our method makes use of wavelets for spherical illumination functions. For simplicity, we use a 2D Haar basis over each face of the cubemap, but it would likely be profitable to use a more sophisticated filter and a more elegant tessellation scheme, such as spherical wavelets [Schröder and Sweldens 1995].

3 Algorithms and Implementation

In this section we treat two different cases: *Geometry Relighting* (Figure 1), where the viewpoint changes but we assume diffuse surfaces, and *Image Relighting* (Figures 3 and 4), where the viewpoint is fixed and we support arbitrary reflection models and light transport effects.

First, consider direct illumination from an environment map:

$$B(\mathbf{x}, \omega_o) = \iint_{\Omega} L(\omega) S(\mathbf{x}, \omega) f_r(\mathbf{x}, \omega \rightarrow \omega_o) (\omega \cdot n(\mathbf{x})) d\omega, \quad (1)$$

where \mathbf{x} is the sample location of a vertex, or a pixel in the image, ω_o is the viewing direction, ω indexes incident directions, L is the environment map, S is the (binary) visibility function indicating if a ray from \mathbf{x} in direction ω is shadowed, f_r is the reflection function at location \mathbf{x} , and $(\omega \cdot n)$ is the cosine of the incident angle.

Geometry Relighting In this case, f_r depends only on surface location (diffuse assumption), so we may define a combined transport function:

$$T(\mathbf{x}, \omega) = S(\mathbf{x}, \omega) f_r(\mathbf{x}) (\omega \cdot n(\mathbf{x})). \quad (2)$$

Image Relighting In this case ω_o depends only on \mathbf{x} , so the combined transport is given by

$$T(\mathbf{x}, \omega) = S(\mathbf{x}, \omega) f_r(\mathbf{x}, \omega \rightarrow \omega_o(\mathbf{x})) (\omega \cdot n(\mathbf{x})). \quad (3)$$

In both cases, Equation 1 no longer depends on ω_o , and using numerical cubature on the integral (with appropriate normalizing weights that we omit here),

$$B(\mathbf{x}_i) = \sum_j T(\mathbf{x}_i, \omega_j) L(\omega_j), \quad (4)$$

which is easier to write simply in matrix notation,

$$B = TL, \quad (5)$$

where T is the light transport matrix, and B and L are column vectors for computed radiance and incident illumination.

It is important to note that this form applies much more generally than originally introduced in Equation 1. First, the illumination may be parameterized over any manifold, not just a sphere, so we can manipulate not just environment maps, which represent distant lighting, but also a local light on a ceiling or from a window, for instance. Second, the function T can include global light transport effects, such as interreflection, subsurface scattering and caustics; the only restriction is in the case of geometry relighting, where these effects must terminate on a diffuse surface. Finally, Equation 5 holds not just for an explicit tabular representation of L , but when it is expressed in any orthonormal basis, like spherical harmonics or wavelets.

3.1 Pre-computation

Our method requires a significant amount of pre-computation, which we accelerate in some cases by using graphics hardware. Storage costs for T may be large if we desire accurate light transport from high resolution and high frequency environment maps. However, if the lighting is truly low-frequency, then low-resolution maps suffice and our representation is comparably compact to linear spherical harmonic techniques [Sloan et al. 2002].

Rendering Raw Transport Matrix We first compute transport matrix T in the raw lighting basis of cubemap pixel lights. We use two different methods which provide different benefits.

The first method is to ray trace the columns of the matrix, which are simply images of the scene illuminated under a single pixel of the cubemap environment. This method is simple, and clearly allows any light transport effect supported by the ray-tracer to be encoded into the matrix. However it works only for image-relighting.

The second method is to compute the matrix rows rather than columns. In this case we handle only direct illumination, which means that each row is simply the visibility cubemap for a particular \mathbf{x} , weighted by the reflection function and cosine term at that location. This simplification allows use of graphics hardware. We rasterize a high-resolution visibility hemi-cube [Cohen and Greenberg 1985] at each \mathbf{x} , read back the pixels, apply the weighting due to the reflection function and down-sample to the desired environment resolution.

Wavelet Transform We project each row of the raw matrix (corresponding to the dot-product transfer of light from L onto a single output location \mathbf{x}) onto a wavelet basis. Our transform is 2D Haar over each cubemap face. It is worth noting that after the transformation is complete the representation is still lossless, and a full matrix multiplication corresponds to exact re-lighting.

Quantization and Storage We dither and quantize the matrix elements to 6, 7 or 8 bits, depending on the desired level of fidelity, and discard all zero coefficients. Because the wavelet transform translates coherence into small coefficients, the resulting matrix is sparse. To prevent using memory bandwidth for lights that are unused in relighting a frame, we store the coefficients for each light contiguously in memory. For improved cache coherence in accumulating the outgoing radiance, we segment the matrix into blocks of 256 output samples. We store each coefficient as an 8-bit value and an 8-bit block index.

3.2 Rendering

We use a very small number of matrix columns to relight each frame. In this section we describe how relighting works as a sparse matrix multiplication.

Wavelet Transform We weight L by the normalization factors required by numerical cubature, then perform a fast wavelet transform. We use a 2D Haar transform here as well because Haar is an orthonormal basis. In general one would project the light onto the dual wavelet basis.

Non-linear Lighting Approximation We choose a subset of wavelet basis lights to use for the current frame. We have experimented with three methods for selecting the lights. In the first method, *unweighted selection*, each light’s priority is simply the magnitude of its wavelet coefficient. This is easily shown to be the optimal choice for minimizing L^2 error in the illumination [Mallat 1998]. However, this is not the same as minimizing error of the output images, which is our ultimate goal.

The second method, *transport-weighted selection*, scales the priority of the i^{th} light by the energy of the i^{th} column of T . For image relighting the weight is the image energy of the scene lit by the i^{th} light; for geometry relighting, the weight is the irradiance energy of the mesh lit by that light. These column energies may be pre-computed and the weighting applied with negligible cost. The net effect is to reduce the priority of lights which contribute little or no variation in the output vector. An example of a zero-priority light is the one directly beneath the floor in Figures 1, 3 and 4.

The third method, *area-weighted selection*, scales the priority of each wavelet light by its area. Giving preference to larger lights helps to resolve the diffuse color of the scene more quickly. Figures 1, 3 and 4 are generated with this selection scheme.

Sparse Block Matrix Multiplication We relight the scene by performing the matrix multiplication formulated in Equation 5 where L is sparse because of non-linear approximation, and T is sparse because of wavelet encoding. We have implemented this multiplication entirely on the CPU. Graphics hardware has evolved to support some forms of sparse matrix-vector multiplication [Bolz et al. 2003], but it is not clearly general enough to support our algorithm efficiently because we exploit sparsity in both L and T .

4 Results

In this section we first present an error analysis in which we compare our technique against linear spherical harmonics. We then present data for our compression rates in the lighting vector and transport matrix, and for the speed of our relighting system.

4.1 Error Analysis

We compare the error induced by approximations using non-linear wavelets and linear spherical harmonics. We make comparisons in two spaces: in the cubemap pixels of the lighting approximation, and the image pixels of the final renderings.

To present our findings we use two acquired environments: St. Peter’s Basilica and Grace Cathedral (see cubemap insets, Figures 1 and 4) [Debevec and Malik 1997]. These environments contain a dynamic range of 5 orders of magnitude and carry energy across essentially all frequencies, as illustrated by the plots of spherical harmonic energy in Section 4.1.1. They contain area light sources in a range of sizes, some of which approach point lights in St. Peter’s. Additional environments are shown on the video.

4.1.1 Error in Lighting Approximation

The process of choosing a small number of terms to approximate the lighting induces error by neglecting the contributions of all basis functions that are discarded. We visualize the approximate lighting directly (Figure 2) by projecting the wavelet and spherical harmonic lights back onto cubemap pixels.

Figure 2 shows that 100 (area-weighted) non-linear wavelets resolve the high-energy lights in the environment, but the first 100 spherical harmonics do not. 4096 wavelets capture even the low energy detail, but even 10,000 harmonics blur the brightest lights





Figure 2: Comparison of cubemap lighting approximations for St. Peter’s Basilica, with linear spherical harmonics (SH) and area-weighted, non-linear wavelets (W) of different orders. Two views of the same data are presented because the luminance values have a dynamic range of 10^5 . High-energy lights (more than 10^4) are displayed in the left half, and low-energy lights (less than 10^2) in the right half. Notice that 4096 wavelets provide very high accuracy (0.6% error), but even 10,000 harmonics do not (48% error). Also note that 100 wavelets resolve high-energy details well.

and exhibit sharp ringing in the low energy details. The graphs below compare the L^2 lighting error for linear spherical harmonic and non-linear wavelet lighting approximations, as a function of the number of terms used. They show that the linear approximation is fundamentally less accurate at all approximation orders, with an error decay that is 2-3 orders of magnitude slower than for non-linear wavelets. Note that the horizontal axis is plotted on a log scale.



We omit the transport-weighted approximation from these graphs because it depends on T and therefore changes with the scene. Notice that area-weighted selection converges slightly less quickly than the unweighted scheme, which is optimal in this error metric.

It might be surprising that spherical harmonics exhibit a 50% L^2 error even with thousands of terms, but this is in fact common for high-dynamic range imagery involving localized bright light sources. For an analytic directional light source or delta function, the spherical harmonic coefficients do not converge, but increase as the square root of the spherical harmonic order. In analyzing the statistics of natural illumination, Dror et al. [2001] have noted that the spherical harmonic coefficients remain almost flat until fairly large numbers of terms (1000-3000), corroborating our results.

4.1.2 Error in Computed Radiance

The error of the computed radiance is of primary importance since we visualize this error directly. We focus in this section on the radiance computed at the output image pixels of the plant (Figure 3) and teapot (Figure 4) scenes, but the same results hold for the irradiance computed at the vertices of the buddha geometry.

We make use of two error metrics. The first is the standard L^2 norm, which measures convergence in image energy. However, since shadow details occur at boundaries that are essentially one dimensional and carry little energy, blurred shadows – which are perceptually striking – register low L^2 error. For this reason we also compute the Sobolev H^1 norm, which is given by


$$\|f\|_{H^1} = \left(\int f^2(x) + |\nabla f(x)|^2 dx \right)^{1/2}$$

The H^1 norm measures error in boundary details better, because it weights the energy of the derivative. Note that the L^2 energy is included in the H^1 norm, and we subtract it out in reporting the relative H^1 error.

The graphs in the next column provide a numerical comparison of the error in Figures 3 and 4. These graphs illustrate two important points. First, when using fewer than 1000 terms, area-weighted wavelet approximation performs better than the transport-weighted or unweighted methods described in Section 3.2. Second, area-weighted wavelet approximation converges exponentially faster than linear harmonic approximation. Note that the horizontal axes are plotted on a log scale.

Figures 3 and 4 provide a visual comparison between the best wavelet approximation and spherical harmonic approximation. These images show that a 25 term [Sloan et al. 2002] linear harmonic approximation completely blurs shadow details, and even an impractically large 2,000 term expansion cannot resolve the highest frequencies. In contrast, a non-linear wavelet approximation requires 1-2 orders fewer terms. 200 wavelets resolve essentially

all shadow details cast by these environments. The error values below each picture show that very low relative L^2 error is needed for perceptually accurate shadows: 1% for the plant, and 0.1% for the teapot. At these levels of accuracy, the error graphs show that area-weighted wavelet approximation requires only 1–10% as many terms as linear spherical harmonics.



For truly low-frequency lighting, linear harmonics and non-linear wavelets perform comparably, both visually and numerically. For example, the Uffizi Gallery lighting environment is well approximated by a single large area light source for the visible sky. When the plant is illuminated under this environment, using 25 linear harmonics produces 7% error, and 25 non-linear wavelets 5% error.

4.2 Performance

4.2.1 Lighting Compression

We compress the lighting by 2-3 orders of magnitude. In practice, this means that out of 24,576 possible wavelet lights, we use between 20 and 200 lights. We find that for St Peter's and Grace, using 100 terms produces visually pleasing animations. For smooth outdoor environments (with no direct illumination from the sun) as few as 20-30 terms suffice. Furthermore, area lights of any size, including illumination from a single cubemap pixel, are well approximated by 30-40 wavelet lights.

An appropriate number of terms is easily chosen per-frame to provide a desired level of accuracy in the lighting approximation. Temporal artifacts result when so few terms are used that discarded coefficients have significant non-zero values. In this case, artifacts may occur because non-zero lights turn on and off discontinuously as they are alternately retained and discarded from frame to frame.

4.2.2 Transport Matrix Compression

The plant and teapot scenes have fixed viewpoint and contain 512×512 pixels. The buddha scene contains 160,000 vertices over the statue and floor, and may be rendered with good fidelity at 1280×1024 screen resolution. We pre-compute light transport assuming up to $6 \times 64 \times 64$ cubemaps for St Peter and Grace, and as little as $6 \times 4 \times 4$ cubemaps for low-frequency environments.

Pre-computation for an 80,000 vertex dataset (buddha without floor), with 400,000 visibility samples takes under 3 hours. Pre-computation times for image-relighting depend on the ray-tracer.

The pre-computed wavelet matrices are sparse at high resolution. With 7-bit quantization, as few as 10% of the elements are non-zero. This decreases to as few as 3% with 6-bit quantization, at the expense of additional dithering noise in matrix images.

The following table reports the sparsity (Sp. columns) as the fraction of non-zero coefficients compared to the full number of matrix elements, and the size of our encoded matrices using 6-bit quantization. Matrices are 60% larger with 7-bit quantization.




Figure 3: Image relighting of a plant scene in St Peter’s Basilica displays sharp shadows from small windows. We compare shadow fidelity using linear spherical harmonics (top row) and non-linear wavelets selected with area-weighted prioritization (bottom). The number of terms is given in parentheses below each picture, as are the L^2 and H^1 error. Notice that for high quality shadows wavelets require two orders of magnitude fewer coefficients for comparable error (compare W (200) and SH (20,000)). Model courtesy of O. Deussen, P. Hanrahan, B. Lintermann, R. Mech, M. Pharr, and P. Prusinkiewicz.




Figure 4: Image relighting of a glossy teapot in Grace Cathedral (inset cubemap) epitomizes all-frequency cast shadows from area lights of different sizes. We compare fidelity of the shadows and the specular highlight, using linear spherical harmonics (top row) and area-weighted, non-linear wavelets (bottom). Note that harmonics blur the specular highlight at 200 terms and the tip of the spout’s shadow at 2000. Both features are well resolved with 200 wavelets.

Light Res.	Buddha		Teapot		Plant	
	Sp.	Size	Sp.	Size	Sp.	Size
$6 \times 4 \times 4$	32%	8.5 MB	19%	9.5 MB	17%	8.4 MB
$6 \times 8 \times 8$	32%	34 MB	16%	32 MB	15%	31 MB
$6 \times 16 \times 16$	26%	113 MB	14%	111 MB	14%	115 MB
$6 \times 32 \times 32$	10%	185 MB	10%	331 MB	12%	394 MB
$6 \times 64 \times 64$	3.8%	314 MB	3.1%	476 MB	7.4%	1.0 GB

The matrices are large but manageable for high-resolution lighting. In the cases where Sloan et al. provide a perceptually valid approximation, the lighting is very low-resolution, and a comparably compact wavelet matrix suffices. The video shows environments for which these low-resolution approximations are valid.

It should also be noted that scaling the methods of Debevec et al. [2000] or Sloan et al. [2002] to our resolutions would result in similarly sized datasets, but their systems would not work in real-time because of the lack of sparsification provided by lighting compression. In the case of linear spherical harmonics the resulting matrices would also not provide the sparsity exhibited by wavelet matrices, because harmonics are globally supported basis functions.

4.2.3 Total Compression and Rendering Speed

The combined lighting and matrix compression ratio is on the order of $10^3:1$ or $10^4:1$. The associated reduction in computation allows interactive relighting from environments that are at least ten times greater in resolution than previously demonstrated.

We report performance on a commodity 2.8 GHz Pentium 4 computer with an nVIDIA GeForce4 graphics card. We present end-to-end frame rates (“Full” columns below), as well as refresh rates for just relighting (“Relight” columns below). Relighting includes everything in Section 3.2 except rasterization and display.

Terms	Buddha		Teapot		Plant	
	Full	Relight	Full	Relight	Full	Relight
40	19 / 9.8	38 / 13	15 / 6.9	25 / 7.4	18 / 8.1	25 / 9.9
100	14 / 6.2	21 / 7.7	8.8 / 3.4	10 / 4.4	12 / 5.1	13 / 5.5
200	9.5 / 3.5	13 / 4.4	5.0 / 2.5	5.3 / 2.4	6.6 / 2.6	6.7 / 3.3

Rates are presented in Hz. Each cell contains rates for both monochromatic / color rendering. In color rendering, each channel is handled separately, with no optimization to use vector operations.

5 Discussion and Conclusion

At the heart of our technique lies the simple and natural idea to compress the relighting computation by taking the lighting data into account. Our most important discovery is the fact that, in a wavelet representation, less than 1% of the lighting information is required to produce accurate renderings under detailed natural environments. Previous research in image-based relighting has not exploited this critical source of data redundancy.

An interesting feature of our results is that area-weighted coefficient selection is superior to unweighted or transport-weighted selection. This is best understood in terms of the work of Ramamoorthi and Hanrahan [2001]. They show that unshadowed, diffuse reflection preserves only the lowest-frequency lighting energy. As a result, area-weighting works well because it prioritizes larger, lower-frequency lights, and because our scenes are mostly diffuse.

The teapot scene contains a glossy term, however, and for this scene transport-weighted selection works almost as well as area-weighting (see the error graphs for the teapot in Section 4.1.2). Area-weighting works less well in this case because glossy surfaces reflect energy from smaller, higher-frequency regions of the sky.

This result invites further study into the design of non-linear coefficient selection schemes. The 100-200 wavelet coefficients used to approximate the St. Peter and Grace environments is large compared to low-order harmonic approximations. This reflects the high information content in these environments, and we believe that accurate all-frequency shadows cannot be generated with many fewer terms. Nevertheless, more intelligent coefficient selection has the potential to minimize the number of terms needed in practice.

We have chosen to focus on shadows in this paper, because shadows are an example of *all-frequency light transport effects*, which produce different images under lights of different frequencies. Shadows are all-frequency because lights of different sizes create shadow boundaries with varying blurriness. Caustics are another all-frequency effect. In contrast, some types of light transport, such as multiple bounces of diffuse reflection, always blur lighting and curtail high frequencies. One might call these *exclusively low-frequency effects*.

We think it would be useful to systematically classify global illumination effects according to which range of lighting frequencies they preserve. Using such a classification, it may be possible to separate a relighting problem into exclusively low-frequency effects, which are well approximated by linear methods such as that of Sloan et al. [2002], and all-frequency effects, which require more costly techniques such as the one proposed in this paper.

The main costs of our non-linear approach are memory and pre-computation time: the transport matrices are very large. However, in this paper we have focussed on simplifying the matrix-vector multiplication, not on compressing the representation of the matrix itself. An area of future work is to apply wavelet transforms on the matrix columns as well as the rows, for a full 4D compression scheme. A more challenging research direction is to support fully dynamic geometry, which probably implies eliminating the dependence on pre-computed data altogether.

Another challenging open problem is to support realistic materials with changing viewpoint. One advantage of low-order harmonics is that they lead to fast methods in this case [Sloan et al. 2002]. However, these only apply when both lighting and reflection are very smooth; otherwise, specular highlights and shadows are drastically blurred (Figure 4). For detailed lighting environments, new methods are still needed. Our recent work suggests that non-linear approximation can be usefully applied in this area also.

In summary, this paper introduces fast and accurate rendering under detailed illumination, demonstrating its viability and potential. We hope that bringing these rich visual effects into the interactive realm will provide a new medium for artistic expression.

Acknowledgments

We would like to thank Henrik Wann Jensen for helping compute transport matrices and for discussions; Paul Debevec for the lighting environments (<http://www.debevec.org/probes/>); Mike Houston, Kekoa Proudfoot and Eric Chan for help with writing and code; and Peter Schröder and Gene Golub for advice. We would also like to thank the anonymous reviewers for their valuable suggestions. This work was supported by a Stanford School of Engineering Fellowship, an NSF grant (NSF-ISS-0085864-2), and a Department of Energy grant (Contract LLL-B341491).

References

- AGRAWALA, M., RAMAMOORTHI, R., HEIRICH, A., AND MOLL, L. 2000. Efficient image-based methods for rendering soft shadows. In *Proceedings of SIGGRAPH 2000*, 375–384.
- ASHIKHMIN, M., AND SHIRLEY, P. 2002. Steerable illumination textures. *ACM Transactions on Graphics* 21, 1, 1–19.
- BOLZ, J., FARMER, I., GRINSPUN, E., AND SCHRÖDER, P. 2003. Sparse matrix solvers on the GPU: Conjugate gradients and multigrid. *To Appear in ACM Transactions on Graphics*.
- CHEN, W., BOUGUET, J., CHU, M., AND GRZESZCZUK, R. 2002. Light field mapping: Efficient representation and hardware rendering of surface light fields. *ACM Transactions on Graphics* 21, 3, 447–456.
- COHEN, M. F., AND GREENBERG, D. P. 1985. The hemi-cube: a radiosity solution for complex environments. *Computer Graphics (Proceedings of SIGGRAPH 85)* 19, 3, 31–40.
- DEBEVEC, P., AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. *Proc. of SIGGRAPH 1997*, 369–378.
- DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *Proceedings of SIGGRAPH 2000*.
- DEVORE, R. 1998. Nonlinear approximation. *Acta Numerica* 7, 51–150.
- DORSEY, J., SILLION, F., AND GREENBERG, D. 1991. Design and simulation of opera lighting and projection effects. In *Computer Graphics (Proceedings of SIGGRAPH 91)*, vol. 25, 41–50.
- DROR, R., LEUNG, T., ADELSON, E., AND WILLSKY, A. 2001. Statistics of real-world illumination. In *CVPR 01*, II–164–II–171.
- HECKBERT, P. S., AND HERF, M. 1997. Simulating soft shadows with graphics hardware. Tech. Rep. CMU-CS-97-104, Carnegie Mellon U.
- LALONDE, P., AND FOURNIER, A. 1999. Interactive rendering of wavelet projected light fields. In *Graphics Interface 99*, 107–114.
- MALLAT, S. 1998. *A Wavelet Tour of Signal Processing (Chapter 9)*. Academic Press.
- NIMEROFF, J., SIMONCELLI, E., AND DORSEY, J. 1994. Efficient re-rendering of naturally illuminated environments. In *Fifth Eurographics Workshop on Rendering*, 359–373.
- RAMAMOORTHI, R., AND HANRAHAN, P. 2001. An efficient representation for irradiance environment maps. In *Proceedings of SIGGRAPH 2001*, 497–500.
- SCHRÖDER, P., AND SWELDENS, W. 1995. Spherical wavelets: Efficiently representing functions on the sphere. In *Proceedings of SIGGRAPH 1995*, 161–172.
- SLOAN, P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics* 21, 3, 527–536.
- SOLER, C., AND SILLION, F. 1998. Fast calculation of soft shadow textures using convolution. In *Proceedings of SIGGRAPH 1998*, 321–332.
- TEO, P., SIMONCELLI, E., AND HEEGER, D. 1997. Efficient linear re-rendering for interactive lighting design. Tech. Rep. STAN-CS-TN-97-60, Stanford University.
- WOOD, D., AZUMA, D., ALDINGER, K., CURLESS, B., DUCHAMP, T., SALESIN, D., AND STUETZLE, W. 2000. Surface light fields for 3D photography. In *Proceedings of SIGGRAPH 2000*, 287–296.