

SOHO: Orthogonal and Symmetric Haar Wavelets on the Sphere

CHRISTIAN LESSIG and EUGENE FIUME
University of Toronto

We propose the SOHO wavelet basis—the first spherical Haar wavelet basis that is both orthogonal and symmetric, making it particularly well suited for the approximation and processing of all-frequency signals on the sphere. We obtain the basis with a novel spherical subdivision scheme that defines a partition acting as the domain of the basis functions. Our construction refutes earlier claims doubting the existence of a basis that is both orthogonal and symmetric. Experimental results for the representation of spherical signals verify that the superior theoretical properties of the SOHO wavelet basis are also relevant in practice.

Categories and Subject Descriptors: I.3.0 [Computing Methodologies]: Computer Graphics—General; G.1.0 [Numerical Analysis]: General—Numerical algorithms; G.1.2 [Numerical Analysis]: Approximation—Nonlinear approximation

General Terms: Theory, Measurement

Additional Key Words and Phrases: Wavelet transform, spherical signals

ACM Reference Format:

Lessig, C. and Fiume, E. 2008. SOHO: Orthogonal and symmetric Haar wavelets on the sphere. *ACM Trans. Graph.* 27, 1, Article 4 (March 2008), 11 pages. DOI = 10.1145/1330511.1330515 <http://doi.acm.org/10.1145/1330511.1330515>

1. INTRODUCTION

Many signals are naturally parametrized over the sphere S^2 . Examples from computer graphics include bidirectional reflectance distribution functions (BRDFs), radiance, and visibility. Spherically parametrized signals can also be found in many other fields, including astronomy, physics, climate modeling, and medical imaging. An efficient and distortion free representation of spherical signals is therefore of importance. Of particular interest are the ability to approximate a wide range of signals accurately with a small number of basis function coefficients, and the possibility of obtaining computationally efficient algorithms to process a signal in its basis representation.

A variety of representations for spherical signals has been proposed in the literature. Spherical Harmonics (SH) have been popular for the representation of low-frequency signals. The global support of the basis functions makes the SH basis however inefficient for representing high-frequency data, and for processing signals directly in the basis representation. Spherical Radial Basis Functions (SRBF) are localized in both space and frequency and thus efficiently represent all-frequency signals. Projecting a signal into an SRBF basis is, however, prohibitively expensive, and efficiently processing a signal is difficult due to the unstructured nature of the basis. Wavelets, in contrast, are both localized in space and frequency and hierarchically structured leading to fast algorithms for basis projection and processing a signal in its basis representation. Additionally, wavelet

bases represent dissimilarities in a signal, making them well suited for the compact representation and approximation of real-world signals.

In the past, planar wavelets have often been used to represent spherical signals. These representations suffer however from parametrization artifacts which are unavoidable when the sphere is mapped onto a planar domain. Spherical wavelets are free of distortion but the representations proposed in the literature are limited in their efficiency for approximating and processing all-frequency signals. We argue that the following three properties are important for an efficient representation.

Orthogonality. An orthogonal basis has a variety of theoretical and practical advantages. Most notably, the optimal approximation in the ℓ_2 norm can be found efficiently. In many cases orthogonality also leads to more efficient algorithms, for example for computing product integrals, and establishing properties of a representation is often easier for orthogonal bases.

(Very) Compact Support. The costs of basis projection and processing a signal in its basis representation depend heavily on the support size of the basis functions. Haar-like bases have minimal support and computations in the representation are thus very efficient. In the literature it has often been argued that Haar-like bases are well suited only for the representation of piecewise constant functions. However, both theoretical and practical results have shown that they

The authors wish to acknowledge the Natural Sciences and Engineering Research Council of Canada for funding this project.

Author's address: Dynamics Graphics Project, Department of Computer Science, University of Toronto, 40 St. George St., Toronto, ON M5S 2E4, Canada; email: {lessig,elf}@dgp.toronto.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2008 ACM 0730-0301/2008/03-ART4 \$5.00 DOI 10.1145/1330511.1330515 <http://doi.acm.org/10.1145/1330511.1330515>

also efficiently represent functions of bounded variation, that is, all-frequency signals as found in many applications [Donoho 1993; Ng et al. 2003].

Symmetry. Local symmetry of the basis functions guarantees an orientation-free representation of features in the basis, preventing distortion when a signal is approximated in the basis. The high sensitivity of the human visual system to asymmetric artifacts makes this particularly important for the visual quality of approximated signals.

We conclude that an orthogonal and symmetric spherical Haar wavelet basis is particularly well suited for the efficient approximation and processing of all-frequency signals defined on the sphere. However, none of the bases proposed in the literature satisfies all of the above properties. We therefore developed the *SOHO wavelet basis*. To our knowledge this is the first spherical Haar wavelet basis that is both orthogonal and symmetric. The key to the derivation of the basis is a novel subdivision scheme of the sphere that defines the partition acting as the domain of the basis functions. The derivation of the SOHO wavelets refutes earlier claims doubting the existence of such a basis [Bonneau 1999].

The practical relevance of the superior theoretical properties of the SOHO wavelet basis has been validated experimentally. Results for the representation of different spherical signals show that the SOHO wavelet basis provides competitive or lower error rates than other spherical Haar wavelet bases when signals are approximated in the basis representation. The visual quality of reconstructed signals affirms these results.

Many applications such as environment map rendering in computer graphics, molecular electronics structure calculations in physics, and data set alignment in medical imaging require the rotation of signals. Efficient and accurate algorithms to rotate a signals in its basis representation are therefore of high practical importance. For the SOHO wavelet basis, we developed basis transformation matrices to rotate signals represented in our new basis. The elements of these matrices can be computed analytically and, in contrast to planar representation of the sphere such as cubemaps or the octahedral map used by Wang et al. [2006], the rotation is mathematically well-defined. An analysis of the structure of the rotation matrices shows that they are very sparse and quasi-block symmetric. This is important for the efficient computation of rotations and reduces otherwise significant matrix storage costs [Wang et al. 2006]. We verified experimentally that the rotation in a spherically parametrized representation is significantly more efficient than in a planar representation of the sphere. A more detailed discussion of rotation in spherical Haar wavelet bases is beyond the scope of this paper. For further details see Lessig [2007].¹

2. RELATED WORK

Various representations for spherical signals have been proposed in the literature. *Spherical Harmonics* [MacRobert 1948] have been popular in physics and chemistry. The SH basis has also been employed in other fields such as geoscience and medical imaging [Clarke et al. 2004; Katsuyuki et al. 2001]. In computer graphics, Spherical Harmonics have been used for example by Westin et al. [1992] to represent BRDFs, and Sillion et al. [1991] employed the basis for the representation of exitant radiance in offline radiosity computations. Cabral et al. [1987], and later Ramamoorthi and Hanrahan [2002], Kautz et al. [2002] and Sloan et al. [2002] used

the SH basis for environment map rendering and Precomputed Radiance Transfer (PRT). In the past decade, Spherical Harmonics have been complemented by different spherical and hemispherical harmonic bases [Makhotkin 1996; Koenderink et al. 1996; Alfeld et al. 1996a; 1996b; Gautron et al. 2004; Sloan et al. 2005]. However, the global support of the basis functions prevents harmonic bases from efficiently representing high-frequency signals. Wavelets, in contrast, are localized in both space and frequency and therefore efficient for the representation of all-frequency signals.

Spherical Radial Basis Functions have been used widely, for example, in astronomy and geoscience [Fisher et al. 1993; Narcowich and Ward 1996; Freedden et al. 1998; Freedden 1999]. Recently, these bases have also been employed in computer graphics [Green et al. 2006; Tsai and Shih 2006]. SRBFs are localized in both space and frequency and can thus efficiently represent all-frequency signals. In contrast to wavelets, however, obtaining the basis representation of a signal is prohibitively expensive [Green et al. 2006; Tsai and Shih 2006]. The unstructured nature of SRBF bases makes it furthermore difficult to develop efficient algorithms to process a signal in its basis representation. Such optimizations have shown to be efficient for wavelets [Ng et al. 2004; Sun and Mukherjee 2006].

Different wavelet representations for spherical signals have been proposed in the literature. Lounsbery et al. [1997] defined wavelet bases over subdivision surfaces that can represent sphere-like shapes. Their construction can employ a wide range of subdivision schemes but none of the resulting bases is orthogonal. To make the bases practical it was additionally necessary to truncate the globally supported basis functions. Truncated basis functions, however, no longer form true wavelet bases.

Girardi and Sweldens [1997] developed orthogonal Haar wavelet bases over general measure spaces L_p . The scaling functions employed in their work are identical to those of the SOHO wavelet basis but their wavelet construction does not yield a symmetric basis on the sphere.

In their seminal work, Schröder and Sweldens [1995] proposed different vertex- and face-based spherical wavelets. Based on the work by Girardi and Sweldens [1997], the authors developed the *Bio-Haar* wavelets, a semi-orthogonal and symmetric spherical Haar wavelet basis. *Lifting* was used to obtain smooth, spherical wavelet bases. Schröder and Sweldens verified experimentally that the bases developed in their work are well suited for the representation of common spherical signals from computer graphics. An interesting side result of their experiments is that Haar-like wavelets are as efficient as smoother bases for the representation of image-like signals. This confirmed earlier theoretical results by Donoho [1993] that showed that Haar-like bases are close to optimal for the representation of functions of bounded variation. Ng et al. [2003] later provided similar results demonstrating that Haar-like wavelets efficiently represent natural, all-frequency signals, and that these bases clearly outperform Spherical Harmonics.

Based on Schröder and Sweldens [1995], Nielson et al. [1997] and later Bonneau [1999] developed semi-orthogonal, symmetric spherical Haar wavelet bases that are *nearly orthogonal*, in the sense that they become orthogonal in the limit as the subdivision level of the wavelet domain goes to infinity and the area of the domains goes to zero. Recently, Roşca [2005] likewise developed a family of nearly orthogonal spherical Haar wavelet bases. However, none of these works provided detailed experimental results comparing their newly developed bases to incumbent spherical Haar wavelet bases.

Ma et al. [2006] used a Haar-like pseudo wavelet basis over the sphere for PRT. The basis is identical to that proposed by Bonneau [1999], but the authors assumed the subdivision of a

¹An extended version of this work can be found in Lessig's thesis, available online at <http://www.dgp.toronto.edu/people/lessig/soho>

partition yields child domains of equal area. This is in general not true; indeed the *pseudo Haar wavelets* are not a basis of $L_2(\mathbb{S}^2, d\omega)$ [Lessig 2007].

Wavelets parametrized over planar domains have also been used to represent spherical signals [Ng et al. 2003, 2004; Zhou et al. 2005; Wang et al. 2006; Sun and Mukherjee 2006]. These techniques are limited in that a planar parametrization of the sphere unavoidably leads to distortion.

The SOHO wavelets are inspired by the seminal bases developed by Schröder and Sweldens [1995] and by Bonneau [1999]. In contrast to these representations, our basis is both orthogonal and symmetric. Unlike Bonneau, we are also able to show that our new basis is an unconditional basis of the space $L_2(\mathbb{S}^2, d\omega)$ [Lessig 2007].

3. SECOND-GENERATION WAVELETS

We shall present an overview of second-generation wavelets, which provides the necessary background for the derivation of the SOHO wavelet basis in Section 4. A more comprehensive introduction may be found, for example, in Sweldens [1996] or Lessig [2007].

Let $L_2(X)$ be the space of functions with finite energy defined over a domain $X \subseteq \mathbb{R}^n$, and let $\langle \cdot, \cdot \rangle$ be an inner product on X . A *multiresolution analysis* of X is a sequence of nested spaces $V_j \subset V_{j+1}$ on different levels j whose union is dense in $L_2(X)$. Bases of the spaces V_j are formed by sets of *scaling basis functions* $\{\varphi_{j,k} \mid k \in \mathcal{K}(j)\}$, where $\mathcal{K}(j)$ is an index set defined over all scaling basis functions on level j . The strictly nested structure of the V_j implies the existence of *difference spaces* W_j such that $V_j \oplus W_j = V_{j+1}$. The W_j are spanned by sets of *wavelet basis functions* $\{\psi_{j,m} \mid m \in \mathcal{M}(j)\}$ with $\mathcal{M}(j)$ being an index set defined over all $\psi_{j,m}$ on level j . For all levels $j \in \mathcal{J}$, V_j and W_j are subspaces of V_{j+1} implying the existence of refinement relationships

$$\varphi_{j,k} = \sum_{l \in \mathcal{L}(j,k)} h_{j,k,l} \varphi_{j+1,l} \quad \text{and} \quad \psi_{j,m} = \sum_{l \in \mathcal{L}(j,m)} g_{j,m,l} \varphi_{j+1,l} \quad (1)$$

with *scaling filter coefficients* $h_{j,k,l}$ and *wavelet filter coefficients* $g_{j,m,l}$. The index sets $\mathcal{L}(j,k) \subseteq \mathcal{K}(j+1)$ and $\mathcal{L}(j,m) \subseteq \mathcal{K}(j+1)$ are defined exclusively over the nonzero filter coefficients $h_{j,k,l}$ and $g_{j,m,l}$, respectively. If not mentioned otherwise, in the following l is assumed to run over $\mathcal{L}(j,k)$ or $\mathcal{L}(j,m)$.

A wavelet basis Ψ is formed by the scaling basis function at the topmost level and the wavelet basis functions on all levels $j \in \mathcal{J}$

$$\Psi \equiv \{\psi_i \mid i \in \mathcal{I}\} \equiv \{\varphi_{0,0}, \psi_{j,m} \mid j \in \mathcal{J}, m \in \mathcal{M}(j)\},$$

with \mathcal{I} being an index set defined over *all* basis functions of Ψ . A wavelet basis is *orthogonal* iff

$$\langle \varphi_{j,k}, \varphi_{j,k'} \rangle = \delta_{k,k'} \quad \langle \psi_{j,m}, \psi_{j',m'} \rangle = \delta_{j,j'} \delta_{m,m'} \quad \langle \varphi_{j,k}, \psi_{j',m'} \rangle = 0, \quad (2)$$

where $\delta_{i,j}$ is the Kronecker delta. For a *semi-orthogonal* wavelet basis the orthogonality of the wavelet basis functions in Equation (2) is no longer satisfied, and for *biorthogonal* wavelets also the scaling basis functions are not orthogonal. If not mentioned otherwise, in the following we will mean “orthogonal wavelet” when we use “wavelet.”

Instead of working with the basis functions it is often more convenient to employ the filter coefficients. The conditions in Equation (2)

can for example be written as

$$\sum_l h_{j,k,l} h_{j,k',l} = \delta_{k,k'} \quad \sum_l g_{j,m,l} g_{j',m',l} = \delta_{j,j'} \delta_{m,m'}$$

$$\sum_l h_{j,k,l} g_{j,m,l} = 0.$$

With an orthogonal wavelet basis, a function $f \in L_2(X)$ can be represented as

$$f = \sum_{i \in \mathcal{I}} \langle f, \psi_i \rangle \psi_i = \sum_{i \in \mathcal{I}} \gamma_i \psi_i$$

where the γ_i are the basis function coefficients. Computing inner products to obtain the γ_i would be expensive and limit the practicality of wavelets. The *fast wavelet transform* projects a function into its wavelet basis representation and reconstructs it in linear time. This is accomplished by using the scaling and wavelet filter coefficients instead of the basis functions. An analysis step of the fast wavelet transform takes the form

$$\lambda_{j,k} = \sum_l h_{j,k,l} \lambda_{j+1,l} \quad \text{and} \quad \gamma_{j,m} = \sum_l g_{j,m,l} \lambda_{j+1,l}, \quad (3)$$

computing the basis function coefficients at level j as a linear combination of the scaling function coefficients at level $j+1$. A *synthesis step* takes the form

$$\lambda_{j+1,l} = \sum_k h_{j,k,l} \lambda_{j,k} + \sum_m g_{j,m,l} \gamma_{j,m}, \quad (4)$$

reconstructing the scaling function coefficients at level $j+1$ from the basis function coefficients at level j . Here, k and m run only over the nonzero filter coefficients $h_{j,k,l}$ and $g_{j,m,l}$.

A *partition* $\{S_{j,n} \mid j \in \mathcal{J}, n \in \mathcal{N}(j)\}$ formed by measurable subsets $S_{j,n}$ of X is used to construct the basis functions. For every level j , the $S_{j,n}$ form a simple cover of X . Domains on different levels are strictly nested and partitions satisfying $S_{j+1,n'} \subset S_{j,n}$ are called the *children* $S_{j+1,l}^n$ of $S_{j,n}$. In the limit when $j \rightarrow \infty$ every $S_{j,n}$ converges to a single point. The index set $\mathcal{K}(j)$ defined over the scaling basis functions satisfies $\mathcal{K}(j) \subseteq \mathcal{N}(j)$. Subdivision schemes for embeddings $X \subset \mathbb{R}^3$ are well-known examples for partitions.

Some of the most popular wavelet bases are *Haar-like* wavelets. For these bases the scaling basis functions are defined as $\varphi_{j,k} = \eta_{j,k} \chi_{j,k}$, with $\mathcal{K}(j) = \mathcal{N}(j)$, where $\chi_{j,k}$ is the characteristic (or inclusion) function of the partition $S_{j,k}$, and $\eta_{j,k} \in \mathbb{R}$ is a scaling factor that is usually chosen such that the basis functions are normalized. It follows from the definition of the scaling basis functions that the Haar wavelet basis functions associated with a domain $S_{j,k}$ are defined exclusively over the children $S_{j,l}^k$.

In this paper we develop a Haar-like wavelet basis that is *symmetric* in the sense that the basis function coefficients are locally invariant with respect to the labeling of all but one child partition of a partition $S_{j,n}$ [Nielsen et al. 1997]. More formally, let $S_{j+1,l}^n$ be the child partitions of $S_{j,n}$, and let $\mathcal{P}(j,n)$ be an index set defined over all but one (fixed) child partition. For a Haar-like wavelet basis, $\mathcal{P}(j,n)$ can also be used to index the basis function coefficients $\Gamma = (\gamma_{j,1}, \dots, \gamma_{j,m})$, with $m = |\mathcal{P}(j,n)|$, of the wavelet basis functions associated with $S_{j,n}$. A basis is then *symmetric* iff an arbitrary permutation $\sigma(\mathcal{P}(j,n))$ implies $\sigma(\Gamma)$, that is altering the labeling of the partitions changes the order of the basis function coefficients accordingly, but does not affect the value of the coefficients γ_i . As example consider a four-fold subdivision of a triangular domain $T_{j,n}$ where the children $T_{j+1,l}^n$ are obtained by subdividing the sides of the parent triangle (cf. Figure 1). If we choose the central child, $T_{j+1,0}^n$, as the fixed domain then a basis is symmetric if we can change the




Fig. 1. Subdivision of a spherical triangle. The labeling of the entities of a spherical triangle $T_{j,k}$ is shown in (a), the 4-fold subdivision yielding the child triangles in (b).

labeling of the three outer child triangles without altering the basis function coefficients associated with $T_{j,n}$.

4. SOHO WAVELETS

We now derive the SOHO wavelet basis. The basis spans the space $L_2(\mathbb{S}^2)$ of functions with finite energy on the sphere \mathbb{S}^2 . We employ the standard inner product

$$\langle f, g \rangle = \int_{\mathbb{S}^2} fg \, d\omega,$$

on \mathbb{S}^2 . The measure $d\omega$ is defined as $d\omega \equiv d\omega(\theta, \phi) = \sin \theta \, d\theta \, d\phi$.

4.1 The Partition Scheme

The partition \mathcal{T} , over which the basis functions of the SOHO wavelet basis are defined, is formed by a set of spherical triangles $\mathcal{T} = \{T_{j,k} \mid j \in \mathcal{J}, k \in \mathcal{K}(j)\}$. The domains at the coarsest level $T_{0,k}$ are obtained by projecting a platonic solid with triangular faces such as the octahedron or the icosahedron onto the sphere. The domains at finer levels are formed by recursively subdividing every spherical triangle $T_{j,k}$ into four child triangles $T_{j+1,l}^k$. As shown in Figure 1(b), these are obtained by inserting one new vertex $v_{j,k}^l$ on each of the arcs forming the sides of the $T_{j,k}$. The partition \mathcal{T} is thus a forest of partition trees and the domains $T_{0,k}$ at the coarsest level are the root nodes of these trees.

Although \mathcal{T} is defined similarly to the partition used by Schröder and Sweldens [1995] we do not employ the geodesic bisector to obtain the positions of the new vertices $v_{j,k}^l$. For the SOHO wavelet basis, the vertex positions are chosen so that the *areas* of the three outer child triangles $T_{j+1,1}^k$, $T_{j+1,2}^k$, and $T_{j+1,3}^k$ are equal. This is the key to the derivation of a basis that is both orthogonal and symmetric. We will detail the novel subdivision scheme employed in our work in Section 4.4.

The area of a spherical triangle $T_{j,k}$ will be denoted $\alpha_{j,k}$, and we define $\tau_{j,k} \equiv \tau_{j,k}(\omega)$ to be the characteristic function of $T_{j,k}$.

4.2 Scaling Basis Functions

For a Haar-like basis, the scaling basis functions are defined as $\varphi_{j,k} = \eta_{j,k} \tau_{j,k}$. With $\eta_{j,k} = 1/\sqrt{\alpha_{j,k}}$, it follows immediately from the disjoint nature of the $T_{j,k}$ for fixed j that the $\varphi_{j,k}$ on the same level are orthonormal, as required in Equation (2). The scaling functions

for the SOHO wavelet basis are thus

$$\varphi_{j,k} = \frac{\tau_{j,k}}{\sqrt{\alpha_{j,k}}}.$$

Given the $\varphi_{j,k}$, the filter coefficients $h_{j,k,l}$ must be chosen to satisfy Equation (1). It follows from the partition that $|\mathcal{L}(j, k)| = 4$ and that the union of the child domains $\tau_{j+1,l}^k$ is again $\tau_{j,k}$. The filter coefficients are therefore

$$h_{j,k,l} = \frac{\sqrt{\alpha_{j+1,l}^k}}{\sqrt{\alpha_{j,k}}}.$$

For the partition \mathcal{T} and the filter coefficients $h_{j,k,l}$ in Equation (4.2), the *cascade algorithm* converges to the scaling functions in Equation (4.2) [Lessig 2007]. This is a necessary condition for the existence of a wavelet basis [Sweldens 1996].

4.3 Wavelet Basis Functions

We use a custom two-step approach to derive the wavelet basis functions: In the first step a semi-orthogonal basis is developed, and in the second step orthogonality and symmetry are enforced.

For a Haar-like basis, the wavelet basis functions $\psi_{j,k}^l$ with $l = 0, 1, 2$, associated with a partition $T_{j,k}$ are exclusively defined over the child partitions $T_{j+1,l}^k$. Wavelet basis functions $\psi_{j,k_1}^{l_1}$ and $\psi_{j,k_2}^{l_2}$ defined on the same level j but over different partitions are thus trivially orthogonal. For $\langle \psi_{j_1,k_1}^{l_1}, \psi_{j_2,k_2}^{l_2} \rangle = 0$ with $j_1 \neq j_2$ to be true, we require that the wavelet basis functions have a vanishing integral. It is easy to show that for a Haar-like basis this in fact implies that the wavelet basis functions on different levels are orthogonal [Lessig 2007]. In the following it is therefore sufficient to consider only one spherical triangle $T_{j,k}$, with fixed but arbitrary j and k , together with its child triangles $T_{j+1,l}^k$, for the derivation of the wavelet basis functions [Bonneau 1999]. The areas of interest will be abbreviated $\alpha_l \equiv \alpha_{j+1,l}^k$ and $\alpha_p \equiv \alpha_{j,k}$; analogous notation will be used for the characteristic functions $\tau_{j,k}$. In some cases we will also omit the indices j and k . It is then understood that these are the fixed j and k of $T_{j,k}$.

Considering only one partition $T_{j,k}$, the analysis and synthesis steps in Equations (3) and (4) can be expressed as compact matrix-vector products. Perfect reconstruction requires that the analysis and synthesis matrices $A_{j,k}$ and $S_{j,k}$, respectively, satisfy $A_{j,k} = S_{j,k}^{-1}$; for an orthonormal basis this simplifies to $A_{j,k} = S_{j,k}^T$. A synthesis step in matrix-vector notation is of the form

$$\begin{pmatrix} \lambda_{j+1,0} \\ \lambda_{j+1,1} \\ \lambda_{j+1,2} \\ \lambda_{j+1,3} \end{pmatrix} = \begin{pmatrix} h_0 & g_0^0 & g_0^1 & g_0^2 \\ h_1 & g_1^0 & g_1^1 & g_1^2 \\ h_2 & g_2^0 & g_2^1 & g_2^2 \\ h_3 & g_3^0 & g_3^1 & g_3^2 \end{pmatrix} \begin{pmatrix} \lambda_j \\ \gamma_j^0 \\ \gamma_j^1 \\ \gamma_j^2 \end{pmatrix}, \quad (5)$$

where g_l^i denotes the l^{th} filter coefficient associated with the i^{th} wavelet basis function $\psi_{j,k}^i$ defined over $T_{j,k}$, and the h_l are the filter coefficients derived in Section 4.2.

In the following we will derive a semi-orthogonal spherical Haar wavelet basis. For fixed $T_{j,k}$, the basis functions of such a basis have to satisfy

$$\langle \psi_{j,k}^0, \varphi_{j,k} \rangle = \langle \psi_{j,k}^1, \varphi_{j,k} \rangle = \langle \psi_{j,k}^2, \varphi_{j,k} \rangle = 0. \quad (6)$$

Equation (6) can be written in Dirac bra-ket notation:

$$[(\Phi_j \mid \Psi_j)] = 0, \quad (7)$$

where $[(X \mid Y)]$ denotes the matrix of inner products of the two function sets X and Y , and Φ_j and Ψ_j are matrices containing the




Fig. 2. The scaling basis function and the two possible sets of wavelet basis functions defined over a partition at level 0. The top row shows the basis functions where the positive sign has been employed in the computation of parameter a in Equation (9), the bottom row shows the basis functions where the negative sign has been used. The basis is defined over a partition derived from an octahedron. Reddish hues indicate positive filter coefficients, bluish hues negative ones.

scaling and wavelet basis functions defined over $T_{j,k}$, respectively. The wavelet basis functions in Equation (7) can be expanded with the refinement relationship in Equation (1) [Stollnitz et al. 1996; Finkelstein and Salesin 1994], yielding

$$[(\Phi_j | \Phi_{j+1})]G_j = 0.$$

The matrix Φ_j is degenerate, containing only $\varphi_{j,k}$; Φ_{j+1} is formed by the four scaling functions $\varphi_{j+1,l}^k$ defined over the child domains $T_{j+1,l}^k$; and G_j contains the desired wavelet basis function filter coefficients g_l^j . The matrix G_j spans the nullspace of $[(\Phi_j | \Phi_{j+1})]$ and is obtained using existing techniques, giving

$$G_j = \begin{bmatrix} -\frac{\sqrt{\alpha_1}}{\sqrt{\alpha_0}} & -\frac{\sqrt{\alpha_2}}{\sqrt{\alpha_0}} & -\frac{\sqrt{\alpha_3}}{\sqrt{\alpha_0}} \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (8)$$

It is easy to show that the wavelet basis functions defined in Equation (8) have a vanishing integral [Lessig 2007].

Given the semi-orthogonal basis derived above, we now have to enforce symmetry and orthogonality of the wavelet basis functions, while taking care to preserve the properties that have already been established. Let $\hat{S}_{j,k}$ be a synthesis matrix formed by the h_l and g_l^j derived previously. We first tried to augment each of the wavelet basis function filter coefficients g_l^j in $\hat{S}_{j,k}$ with a free parameter. The desired properties of the basis could then be formulated as a linear system, and the solution to the system would be the wavelet basis functions. However, we were not able to find such a solution; indeed a solution to the system might not exist.

To obtain a basis having the desired properties we therefore required that the area of the three outer child partitions $T_{j+1,1}^k$, $T_{j+1,2}^k$, and $T_{j+1,3}^k$ be equal (cf. Figure 1). In Section 4.4 it will be shown that the partition \mathcal{T} can be constructed so that this constraint is satisfied. With the area-isometry of the three outer child triangles, symmetry can be guaranteed by the following parametrization of the synthesis

matrix:

$$\hat{S}_{j,k} = \begin{bmatrix} \frac{\sqrt{\alpha_0}}{\sqrt{\alpha_p}} & -c\frac{\sqrt{\alpha_1}}{\sqrt{\alpha_0}} & -c\frac{\sqrt{\alpha_1}}{\sqrt{\alpha_0}} & -c\frac{\sqrt{\alpha_1}}{\sqrt{\alpha_0}} \\ \frac{\sqrt{\alpha_1}}{\sqrt{\alpha_p}} & b & a & a \\ \frac{\sqrt{\alpha_1}}{\sqrt{\alpha_p}} & a & b & a \\ \frac{\sqrt{\alpha_1}}{\sqrt{\alpha_p}} & a & a & b \end{bmatrix},$$

where a , b , and c are the remaining free parameters. Enforcing orthogonality of the basis then yields a simple linear system whose solution are the wavelet basis functions for the SOHO wavelet basis (cf. Appendix A):

$$\begin{aligned} \psi_{j,k}^0 &= \frac{\Lambda_1}{\Lambda_0} \tau_0 + \frac{1}{\Lambda_1} ((-2a+1)\tau_1 + a\tau_2 + a\tau_3) \\ \psi_{j,k}^1 &= \frac{\Lambda_1}{\Lambda_0} \tau_0 + \frac{1}{\Lambda_1} (a\tau_1 + (-2a+1)\tau_2 + a\tau_3) \\ \psi_{j,k}^2 &= \frac{\Lambda_1}{\Lambda_0} \tau_0 + \frac{1}{\Lambda_1} (a\tau_1 + a\tau_2 + (-2a+1)\tau_3), \end{aligned}$$

where

$$a = \frac{\alpha_0 \pm \sqrt{\alpha_0^2 + 3\alpha_0\alpha_1}}{3\alpha_0} \quad (9)$$

and $\Lambda_l \equiv \sqrt{\alpha_{j+1,l}^k}$. The two solutions for parameter a yield two different sets of wavelet basis functions. These are shown in Figure 2. The area measures α_0 and α_1 are always positive and thus a is guaranteed to be real for both solutions.

Given the orthogonal basis derived above, an orthonormal basis can be obtained by normalizing the wavelet basis functions. Lessig further establishes that the SOHO wavelets form an unconditional basis of the space $L_2(\mathbb{S}^2, d\omega)$ [Lessig 2007].

Table I.

Shape distortion for the partition proposed in this work and the geodesic bisector subdivision employed by Schröder and Sweldens [1995]. The minimum internal angle (in degrees) over all partitions was used as distortion measure.

Subdivision Scheme	Octahedron		Icosahedron	
Level	5	7	5	7
Our subdivision	43.0864	43.0417	53.5981	53.5808
Geodesic midpoint	45.0345	45.0022	54.0163	54.0010

4.4 Construction of the Partition

The previous section demonstrated that the SOHO wavelet basis can be constructed provided the three outer child triangles $T_{j+1,1}^k$, $T_{j+1,2}^k$, and $T_{j+1,3}^k$ of $T_{j,k}$ have equal area. The geodesic bisector subdivision employed by Schröder and Sweldens [1995], Nielson et al. [1997], and Bonneau [1999] does not have this property. The partition \mathcal{T} , as defined in Section 4.1, only imposes a topology [Nielson et al. 1997]. The vertices $v_{j,k}^l$ can therefore be positioned so that $\alpha_1 = \alpha_2 = \alpha_3$. Let $v_{j,k}^1$ still be the geodesic bisector. The positions of $v_{j,k}^2$ and $v_{j,k}^3$ can then be obtained with a system of equations:

$$\begin{aligned} \cot\left(\frac{E}{2}\right) &= \cot(C) + \frac{\cot\left(\frac{1}{2}\beta_1\right)\cot\left(\frac{1}{2}\gamma\right)}{\sin(C)} \\ \cot\left(\frac{E}{2}\right) &= \cot(B) + \frac{\cot\left(\frac{1}{2}\beta_2\right)\cot\left(\frac{1}{2}\gamma\right)}{\sin(B)} \\ \cot\left(\frac{E}{2}\right) &= \cot(A) + \frac{\cot\left(\frac{1}{2}b - \frac{1}{2}\beta_1\right)\cot\left(\frac{1}{2}c - \frac{1}{2}\beta_2\right)}{\sin(A)}, \end{aligned}$$

where we employed Equation (36) from Todhunter's book [1901] to define the system. The variables on the right hand side of the equations are given in Figure 1, and E denotes the spherical excess of the three outer child domains. Solving the system for β_1 and β_2 yields the desired vertex positions. The resulting formulae are lengthy and may be found in Appendix B. Lessig [2007] establishes that with a consistent labeling of $T_{j,k}$ exactly one solution for the area equality exists.

For the partition \mathcal{T} it is desirable to yield spherical triangles that are uniform. We have not yet been able to prove bounds on the shape distortion introduced by our new subdivision. Numerical experiments show however that it is not significantly larger than for the geodesic bisector subdivision employed by Schröder and Sweldens [1995]. We used the minimum internal angle over all partitions to measure the shape distortion. The results are given in Table I.

5. EXPERIMENTAL EVALUATION

The performance of a basis depends on a wide range of factors and superior theoretical properties are no guarantee of better results in practice. We therefore performed a variety of experiments to assess the practical efficacy of the SOHO wavelet basis. To provide insights for a wide range of applications we focused on experiments which are independent of specific settings.²

5.1 Methodology

In the experiments the SOHO wavelet basis was compared to six previously proposed spherical Haar wavelet bases: the Bio-Haar

²The Matlab code used for the experiments presented in this section can be found online under <http://www.dgp.toronto.edu/people/lessig/soho/matlab.html>.

basis developed by Schröder and Sweldens [1995], the two nearly orthogonal bases proposed by Nielson et al. [1997], the pseudo Haar wavelets used by Ma et al. [2006], and the two nearly orthogonal bases developed by Bonneau [1999]. All bases have been employed to represent three signals: a unimodal function with only low-frequency content, a piecewise constant function with only high-frequency features, and an image-like signal with content in the full frequency spectrum. For convenience we will refer to the test signals as *BRDF*, *visibility map*, and *texture map*, respectively (cf. Figure 3).

When a signal is represented in a wavelet basis, typically a large proportion of the basis function coefficients is very small or zero, and a small number of coefficients is sufficient to obtain reconstructions that closely resemble the original signal. In the experiments we therefore investigated the connection between the error in reconstructed signals and the number of basis function coefficients used to obtain the reconstructions. As error measures we employed the ℓ_1 and the ℓ_2 norm. In the literature it has been argued that for images the ℓ_1 norm better corresponds to the perceived image quality than other numerical error measures [DeVore et al. 1992]. Our test signals can be seen as images on the sphere and we therefore wanted to explore whether or not the ℓ_1 norm is in fact a more appropriate error measure for the signals than the ℓ_2 norm which is the standard norm for the space $L_2(\mathbb{S}^2, d\omega)$. Numerical error measures are valuable in many contexts, but they are limited in their ability to measure the quality of a signal as it is perceived by humans [Pratt 1991]. In Figure 3 we therefore provide plots of reconstructed signals for visual inspection.

The problem of finding the optimal approximation in a basis representation, that is the set of basis function coefficients which minimizes the reconstruction error for a fixed number k of coefficients, is nontrivial [Gross 1996]. Only with the ℓ_2 norm and for an orthogonal basis is it possible to efficiently obtain the optimal approximation: in this case in fact the k largest basis function coefficients yield the minimal reconstruction error. One of the motivations for the development of nearly orthogonal spherical Haar wavelet bases was to obtain representations which show in practice the same behavior as orthogonal bases. Following a similar idea, Ma et al. likewise assumed that the pseudo Haar wavelets provide the same benefits as orthogonal bases. We thus employed the ℓ_2 optimal approximation strategy for orthogonal bases not only for the SOHO wavelet basis but also for the nearly orthogonal Haar wavelets and the pseudo Haar basis. For the semi-orthogonal Bio-Haar basis the approximation given by the k largest basis function coefficients is far from optimal [Lessig 2007]. We therefore employed for the Bio-Haar basis the basis-specific ℓ_2 optimal approximation strategy, requiring to compute inner products $\langle \psi_{j,i}^n, \psi_{j,i'}^n \rangle$ between wavelet basis functions defined over the same domain $T_{j,n}$ [Lessig 2007], thereby disregarding the fact that the computations are significantly more expensive than for the other bases.

As shown in Figure 2, two different SOHO wavelet bases can be obtained by either using the positive or the negative sign in the computation of parameter a in Equation (9). In the experiments the basis in which the negative sign had been employed performed slightly better and in the remainder of the paper we will therefore refer to this basis as the SOHO wavelet basis.

5.2 Evaluation of Approximation Performance

The subdivision scheme proposed in Section 4.4 allows the SOHO wavelet basis to be defined over a partition derived from a tetrahedron, an octahedron, or an icosahedron. Although in our experiments none of the resulting bases provided a clear advantage, we argue that




Fig. 3. Reconstructed signals for the SOHO wavelet basis, the Bio-Haar basis, and the pseudo Haar basis (from left to right).

in general the octahedron is the best choice for the base polyhedron. The basis induced by the octahedron never yielded the highest error rate in any of our experiments and it can directly represent hemispherical signals. Additionally, the alignment of a partition derived from an octahedron with other parametrizations of the sphere is useful in many applications. In the experiments discussed in the following we always employed the octahedron as base polyhedron.

For the spherical Haar wavelet bases employed in the experiments, in Figure 4 the error rates in the ℓ_1 and the ℓ_2 norm resulting from an approximation of the test signals with an increasing number of nonzero basis function coefficients are shown. The basis Bonneau2 performed for all signals very similarly to the SOHO wavelet basis and in the plots both representations are in most cases indistinguishable. For the visibility map the basis Bonneau1 provided higher error rates than the SOHO wavelet basis; both bases achieved however very similar results for the other two signals. The two bases developed by Nielson et al. provided for the texture map

almost the same error rates as the SOHO wavelet basis; for the visibility map however, in particular in the ℓ_1 norm, the error rates were higher. For the BRDF, the basis Nielson2 performed slightly better than the SOHO wavelet basis whereas the basis Nielson1 obtained inferior results. With the exception of the BRDF where the pseudo Haar basis achieved similar results than the SOHO wavelets, the Bio-Haar basis and the pseudo Haar wavelets provided always higher error rates than the other bases. Note that independent of the basis for the visibility map about 5% of the basis function coefficients were sufficient to provide reconstructions with virtually no error. The full information of the signal could therefore be represented with a small fraction of all coefficients.

In Figure 3 reconstructions of the test signals with a subset of all basis function coefficients are shown. The plots for the SOHO wavelet basis look in most cases very similar to those of the nearly orthogonal bases and we therefore provide reconstructions only for the SOHO wavelet basis, the Bio-Haar basis, and the pseudo Haar




Fig. 4. Area-corrected error rates in the ℓ_1 and ℓ_2 norm for approximations of the test signals. The total number of basis function coefficients of the signals was 131,072.

wavelets. For the texture map we employed 8,192, or 6.25%, of the original 131,072 basis function coefficients for reconstruction, for the other two signals 1,024, or 0.78%, of the coefficients were retained. Although the perceived quality is an inherently subjective measure, we argue that the SOHO wavelet basis achieved for all signals visually more pleasing results than the other two bases shown in the comparison. For the texture map the SOHO wavelet basis provides sharper edges and the reconstructed signal appears less noisy. The basis also preserves more detail. This can be seen in the center of Africa where the two lakes south of Lake Victoria are clearly visible only for the SOHO wavelets. For the BRDF, the reconstructed signal for the SOHO wavelet basis appears significantly smoother than the reconstructions for the Bio-Haar basis and the pseudo Haar wavelets, and it therefore better resembles the original Lambertian BRDF. For the visibility map, the SOHO wavelet basis provides again sharper edges than the other two bases.

5.3 Discussion

We believe that the presented results provide valuable insights about the strengths and weaknesses of spherical Haar wavelets and in particular of the SOHO wavelet basis. Some experiments have been omitted as the results have already been presented elsewhere. Ng et al. [2003], for example, showed that Haar-like bases are significantly more efficient than Spherical Harmonics for the representation of all-frequency functions, and Schröder and Sweldens [1995] demonstrated that spherical Haar wavelet bases are as efficient as smoother, spherical wavelets for the representation of image-like signals.

The results shown in Figure 4 demonstrate that the error rates in the ℓ_1 and the ℓ_2 norm are highly correlated, resembling earlier observations by Schröder and Sweldens [1995]. A similar correlation can be observed between both numerical error measures employed in our experiments and the visual quality of reconstructed signals; this would appear to contradict DeVore et al. [1992] who argued that the ℓ_1 error norm better corresponds to the visual quality of image-like signals than other numerical error measures. Our results therefore suggest that it is in applications justifiable to employ the ℓ_2 norm to accurately estimate the error in the ℓ_1 norm or the visual quality of reconstructed signals. This is important as only for the ℓ_2 norm is it possible to efficiently find the optimal approximation; in log-linear time for orthogonal bases and still at moderate costs for many other practically relevant bases.

It is wellknown that Haar-like bases are well suited for the representation of piecewise constant functions such as visibility maps. The experiments presented in this section demonstrate, however, that these bases are also efficient for the representation of smooth signals. For the BRDF for example 0.78% of all basis function coefficients were sufficient to provide visually pleasing reconstructions as shown in Figure 3. The texture map has features in the full frequency spectrum and thus more basis function coefficients had to be retained. Reconstructions which closely resemble the original signal could however still be obtained with not more than 5% of the basis function coefficients.

Prior to the experiments, we had anticipated that the nearly orthogonal spherical Haar wavelet bases would outperform the Bio-Haar and the pseudo Haar wavelets. That these bases provided results very similar to those of the SOHO wavelet basis was however surprising to us. We initially thought that this might result from the high subdivision level on which the signals were defined so that the bases already approached their limit properties. Additional experiments showed however that the good performance of the nearly orthogonal spherical Haar wavelet bases is independent

of the resolution of the input signal, hinting that the design of the basis functions rather than the near orthogonality of the bases causes the results. It will be interesting to explore how the nearly orthogonal bases perform for other applications where orthogonality is of importance.

Considering that the pseudo Haar wavelets are not a basis of $L_2(\mathbb{S}^2, d\omega)$ we were surprised that the representation provided for almost all of the experiments results competitive to those of the true spherical Haar wavelet bases. The assumption that the geodesic bisector subdivision yields child domains with equal area seems therefore in practice to be reasonable. For applications the pseudo Haar wavelets provide the advantage that the filter coefficients are constant and do not depend on the area of the domains of the partition.

5.4 Costs of Wavelet Transform

For spherical Haar wavelet bases it is expensive to compute the area of the partition domains and the filter coefficients which are necessary to perform wavelet transforms. For the SOHO wavelet basis additional costs result from our novel subdivision scheme. Both the partition and the filter coefficients can however be precomputed so that the costs of performing wavelet transforms are independent of the basis and similar to those of wavelets in $2D$.

6. FUTURE WORK

We believe that many applications might benefit from the use of the SOHO wavelet basis. In computer graphics, for example the solution of the rendering equation [Kajiya 1986] is likely to be more efficient with a representation of the factors of the product integral equation in the SOHO wavelet basis. The light transport factors in this case are usually not aligned and have to be rotated before a solution can be obtained efficiently. It would be interesting to explore the possibility of computing basis transformation matrices for these rotations at runtime, thereby avoiding approximations and reducing the otherwise significant storage requirements [Wang et al. 2006].

Beyond computer graphics, applications for example in medical imaging, astrophysics, and geoscience might benefit from the use of the SOHO wavelet basis. In medical imaging in particular the orthonormality of the basis will be of interest; for example the ability to rigorously establish error bounds. For the very large data sets in astrophysics and geoscience the superior approximation performance of the SOHO wavelet basis will be beneficial.

Some theoretical questions also remain. In this work we derived an orthogonal and symmetric spherical Haar wavelet basis by requiring area-isometry of the three outer child partitions. The question if such a basis exists without the area constraint is still unanswered. Another open problem is if smooth, orthogonal and symmetric wavelets on \mathbb{S}^2 exist, and if such a basis provides practical advantages over the SOHO wavelet basis.

7. CONCLUSION

In this work we developed the SOHO wavelet basis, a novel spherical Haar wavelet basis that is both orthogonal and symmetric, clarifying previous work that doubted the existence of such a basis. Experimental results verify that the superior theoretical properties of the SOHO wavelet basis are also of practical relevance.

Combining the findings of this paper, we believe that the SOHO wavelet basis is an attractive representation for the approximation and processing of all-frequency signals on the sphere, and we anticipate that the basis will enable more efficient solutions for many problems in computer graphics and beyond.

ACKNOWLEDGMENTS

We would like to thank Silvio Borac̆ for early feedback, and Michael Daum for help on editing the paper.

A. Appendix

Derivation of the wavelet basis functions for the SOHO wavelet basis (Mathematica document):

<http://www.dgp.toronto.edu/people/lessig/soho/soho-wavelets.nb>

B. Appendix

Development of a spherical subdivision scheme with area-isometry for the three outer child triangles (Mathematica document):

<http://www.dgp.toronto.edu/people/lessig/soho/soho-area-equality.nb>

REFERENCES

- ALFELD, P., NEAMTU, M., AND SCHUMAKER, L. L. 1996a. Bernstein-Bézier polynomials on spheres and sphere-like surfaces. *Comput. Aid. Geom. Des.* 13, 4, 333–349.
- ALFELD, P., NEAMTU, M., AND SCHUMAKER, L. L. 1996b. Fitting scattered data on sphere-like surfaces using spherical splines. *J. Comput. Appl. Math.* 73, 1-2, 5–43.
- BONNEAU, G.-P. 1999. Optimal triangular Haar bases for spherical data. In *Proceedings of the Conference on Visualization (VIS'99)*. IEEE Computer Society Press, Los Alamitos, CA, 279–284.
- CABRAL, B., MAX, N., AND SPRINGMEYER, R. 1987. Bidirectional reflection functions from surface bump maps. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'87)*. ACM Press, New York, NY, 273–281.
- CLARKE, P. J., LAVALEE, D. A., BLEWITT, G., AND VAN DAM, T. 2004. Choice of basis functions for the representation of seasonal surface loading signals in geodetic time series. *AGU Fall Meeting Abstracts*, A121+.
- DEVORE, R. A., JAWERTH, B., AND LUCIER, B. J. 1992. Image compression through wavelet transform coding. *IEEE Trans. Inform. Theory* 38, 2, 719–746.
- DONOHO, D. L. 1993. Unconditional bases are optimal bases for data compression and statistical estimation. *Appl. Comp. Harm. Anal.* 1, 100–115.
- FINKELSTEIN, A. AND SALESIN, D. H. 1994. Multiresolution curves. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'94)*. ACM Press, New York, NY, 261–268.
- FISHER, N. I., LEWIS, T., AND EMBLETON, B. J. J. 1993. *Statistical Analysis of Spherical Data*. Cambridge University Press, Cambridge, UK.
- FREEDEN, W. 1999. *Multiscale Modelling of Spaceborne Geodata*. B.G. Teubner, Stuttgart, Leipzig.
- FREEDEN, W., GERVEN, T., AND SCHREINER, M. 1998. *Constructive Approximation on the Sphere (With Applications to Geomathematics)*. Oxford Sciences Publication, Clarendon Press, Oxford University.
- GAUTRON, P., KŘIVÁNEK, J., PATTANAIK, S. N., AND BOUATOUCH, K. 2004. A novel hemispherical basis for accurate and efficient rendering. In *Proceedings of the Eurographics Symposium on Rendering*. 321–330.
- GIRARDI, M. AND SWELDENS, W. 1997. A new class of unbalanced Haar wavelets that form an unconditional basis for L_p on general measure spaces. *J. Fourier Anal. Appl.* 3, 4.
- GREEN, P., KAUTZ, J., MATUSIK, W., AND DURAND, F. 2006. View-dependent precomputed light transport using nonlinear Gaussian function approximations. In *Proceedings of the Symposium on Interactive 3D Graphics and Games (SI3D'06)*. ACM Press, New York, NY, 7–14.
- GROSS, M. H. 1996. L2 Optimal Oracles and Compression Strategies for Semiorthogonal Wavelets.
- KAJIYA, J. T. 1986. The rendering equation. In *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'86)*. ACM Press, New York, NY, 143–150.
- KATSUYUKI, T., GENSHENG, L. Z., AND GULLBERG, G. T. 2001. Cone-beam image reconstruction using spherical harmonics. *Phys. Medicine Biol.* 46, N127–N138(1).
- KAUTZ, J., SLOAN, P.-P., AND SNYDER, J. 2002. Fast, arbitrary BRDF shading for low-frequency lighting using spherical harmonics. In *Proceedings of the 13th Eurographics Workshop on Rendering (EGRW'02)*. Eurographics Association, Aire-la-Ville, Switzerland, 291–296.
- KOENDERINK, J. J., VAN DOORN, A. J., AND STAVRIDIS, M. 1996. Bidirectional reflection distribution function expressed in terms of surface scattering modes. In *Proceedings of the 4th European Conference on Computer Vision-Volume II (ECCV'96)*. Springer-Verlag, 28–39.
- LESSIG, C. 2007. Orthogonal and symmetric wavelets on the sphere. MSc thesis. University of Toronto.
- LOUNSBERY, M., DEROSE, T. D., AND WARREN, J. 1997. Multiresolution analysis for surfaces of arbitrary topological type. *ACM Trans. Graph.* 16, 1, 34–73.
- MA, W.-C., HSIAO, C.-T., LEE, K.-Y., CHUANG, Y.-Y., AND CHEN, B.-Y. 2006. Real-time triple product relighting using spherical local-frame parameterization. *Vis. Comput.* 9-11, 682–692. Pacific Graphics 2006 Conference Proceedings.
- MACROBERT, T. M. 1948. *Spherical Harmonics; An Elementary Treatise on Harmonic Functions, with Applications*. Dover Publications.
- MAKHOTKIN, O. A. 1996. Analysis of radiative transfer between surfaces by hemispherical harmonics. *J. Quant. Spectros. Radiat. Tran.* 56, 869–879.
- NARCOWICH, F. J. AND WARD, J. D. 1996. Nonstationary wavelets on the m-sphere for scattered data. *App. Comput. Harm. Anal.* 3, 324–336.
- NG, R., RAMAMOORTHY, R., AND HANRAHAN, P. 2003. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Trans. Graph.* 22, 3, 376–381.
- NG, R., RAMAMOORTHY, R., AND HANRAHAN, P. 2004. Triple product wavelet integrals for all-frequency relighting. *ACM Trans. Graph.* 23, 3, 477–487.
- NIELSON, G. M., JUNG, I.-H., AND SUNG, J. 1997. Haar wavelets over triangular domains with applications to multiresolution models for flow over a sphere. In *Proceedings of the 8th Conference on Visualization (VIS'97)*. IEEE Computer Society Press, Los Alamitos, CA 143–ff.
- PRATT, W. K. 1991. *Digital Image Processing* 2nd Ed. John Wiley & Sons, Inc., New York, NY.
- RAMAMOORTHY, R. AND HANRAHAN, P. 2002. Frequency space environment map rendering. In *ACM SIGGRAPH 2002 Papers*. ACM Press, New York, NY, 517–526.
- ROȘCA, D. 2005. Haar wavelets on spherical triangulations. In *Advances in Multiresolution for Geometric Modelling*, N. A. Dodgson, M. S. Floater, and M. A. Sabin, Eds. Mathematics and Visualization. Springer, 405–417.
- SCHRÖDER, P. AND SWELDENS, W. 1995. Spherical wavelets: Efficiently representing functions on the sphere. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'95)*. ACM Press, New York, NY, 161–172.
- SILLION, F., ARVO, J., WESTIN, S., AND GREENBERG, D. P. 1991. A global illumination solution for general reflectance distributions. In *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'91)*. ACM Press, New York, NY, 187–196.

- SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *ACM SIGGRAPH 2002 Papers*. ACM Press, New York, NY, 527–536.
- SLOAN, P.-P., LUNA, B., AND SNYDER, J. 2005. Local, deformable pre-computed radiance transfer. In *ACM SIGGRAPH 2005 Papers*. ACM Press, New York, NY, 1216–1224.
- STOLLNITZ, E. J., DEROSE, T. D., AND SALESIN, D. H. 1996. *Wavelets for Computer Graphics: Theory and Applications*. Morgan Kaufmann Publishers Inc., San Francisco, CA.
- SUN, W. AND MUKHERJEE, A. 2006. Generalized wavelet product integral for rendering dynamic glossy objects. *ACM Trans. Graph.* 25, 3, 955–966.
- SWELDENS, W. 1996. The lifting scheme: A custom-design construction of biorthogonal wavelets. *Appl. Comput. Harmon. Anal.* 3, 2, 186–200.
- TODHUNTER, I. 1901. *Spherical Trigonometry, for the Use of Colleges and Schools*. Macmillan, New York.
- TSAI, Y.-T. AND SHIH, Z.-C. 2006. All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. In *ACM SIGGRAPH 2006 Papers*. ACM Press, New York, NY, 967–976.
- WANG, R., NG, R., LUEBKE, D., AND HUMPHREYS, G. 2006. Efficient wavelet rotation for environment map rendering. In *Proceedings of the Eurographics Symposium on Rendering*. Springer-Verlag, Vienna. Published as *Rendering Techniques 2006*.
- WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *ACM SIGGRAPH 1992 Papers*. ACM Press, New York, NY, 255–264.
- ZHOU, K., HU, Y., LIN, S., GUO, B., AND SHUM, H.-Y. 2005. Precomputed shadow fields for dynamic scenes. In *ACM SIGGRAPH 2005 Papers*. ACM Press, New York, NY, 1196–1201.

Received May 2007; revised October 2007; accepted November 2007