Product Integral Binding Coefficients for High-order Wavelets

Nick Michiels, Jeroen Put and Philippe Bekaert
Hasselt University - tUL - iMinds, Expertise Centre for Digital Media
Wetenschapspark 2, 3590 Diepenbeek, Belgium
{*nick.michiels, jeroen.put, philippe.bekaert*}*@uhasselt.be*

Keywords: wavelets, signal processing, product integrals, tensor, rendering

Abstract: This paper provides an efficient algorithm to calculate product integral binding coefficients for a heterogeneous mix of wavelet bases. These product integrals are ubiquitous in multiple applications such as signal processing and rendering. Previous work has focused on simple Haar wavelets. Haar wavelets excel at encoding piecewise constant signals, but are inadequate for compactly representing smooth signals for which high-order wavelets are ideal. Our algorithm provides an efficient way to calculate the tensor of these binding coefficients. The algorithm exploits both the hierarchical nature and vanishing moments of the wavelet bases, as well as the sparsity and symmetry of the tensor. We demonstrate the effectiveness of high-order wavelets with a rendering application. The smoother wavelets represent the signals more effectively and with less blockiness than the Haar wavelets of previous techniques.

1 INTRODUCTION

Efficient rendering of complex scenes with detailed lighting effects is an important application. To achieve this goal, we need to solve the rendering equation at every point in the scene (Kajiya, 1986). The rendering equation calculates the radiance that reaches the observer, depending on the incoming light, the objects and the materials in the scene. Naive sampling to solve this equation is impractical. To make the calculations more tractable, light transport is often precalculated in a factored form. The rendering equation is often evaluated at each point x, viewed from direction ω_o, by calculating the triple product integral over the hemisphere Ω (Ng et al., 2004):

$$B(x,\omega_o) = \sum_i \sum_j \sum_k V_i \, \rho_j(\omega_o) \, L_k \, C_{ijk} \qquad (1)$$

with visibility V, reflectance ρ, environment lighting L and binding coefficient C_{ijk} defined as:

$$C_{ijk} = \int_\Omega \Psi_i(\omega) \, \Psi_j(\omega) \, \Psi_k(\omega) \, d\omega \qquad (2)$$

The factorization process in our rendering application is demonstrated in Figure 7.

This paper focuses on calculating the binding coefficients C_{ijk} for an arbitrary mixture of wavelet basis functions Ψ_i, Ψ_j and Ψ_k. The position and dilation factor are controlled by the basis function numbers i, j and k. Choosing an appropriate basis to represent the various factors becomes critical to achieve high quality results with minimal computation time. This paper recognizes that each factor in the triple product has different signal characteristics. Therefore, it would be efficient to encode each factor with a basis specifically tailored to it. While the Haar wavelet basis Ψ_H excels at encoding piecewise constant signals, like the visibility factor V, it fails at representing smooth signals compactly such as reflectance ρ and lighting L. For this task, smooth high-order wavelets, e.g. the Daubechies-4 basis Ψ_D, are generally considered more appropriate (the number 4 indicates the support size of the basis functions with the smallest dilation). These smooth bases often require an order of magnitude less coefficients for smooth signals. This paper devises a method to calculate product integrals for arbitrary wavelet bases, focusing specifically on triple product integrals for rendering applications. Michiels et al. (Michiels et al., 2013) already showed the applicative advantages of using smooth wavelets for inverse rendering. Figure 1 shows two wavelet bases: Ψ_H on the left and Ψ_D on the right[1]. These bases are orthonormal:

$$\int_\Omega \Psi_{Hi}\Psi_{Hj} \, d\omega = \delta_{ij} \; and \; \int_\Omega \Psi_{Di}\Psi_{Dj} \, d\omega = \delta_{ij}, \quad (3)$$

[1]Without loss of generality, the illustrations in this paper show 1D wavelet functions. Separable wavelets in higher dimensions can easily be obtained by combining two lower dimensional basis functions.

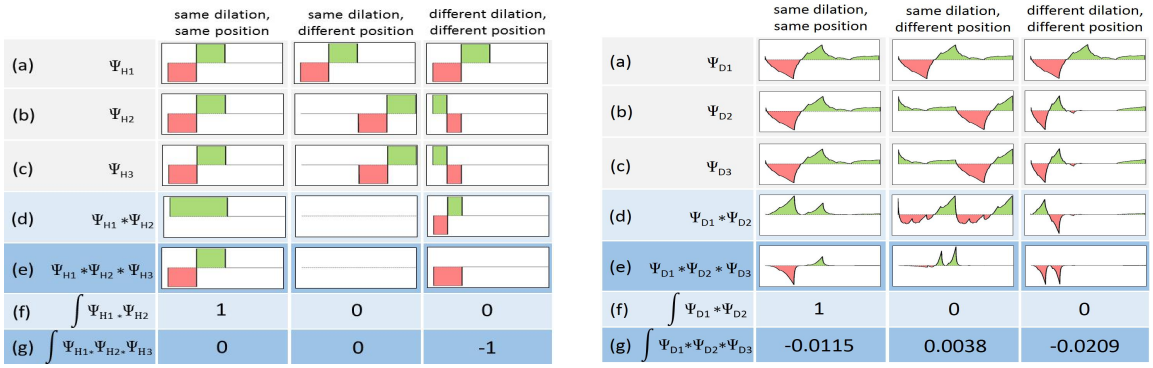

Figure 1: Double and triple product of Haar wavelets (left) and Daubechies wavelets (right). (a), (b) and (c) are the different basis functions Ψ_i, Ψ_j and Ψ_k. Each column represents a permutation of i, j and k that create basis functions with different dilation and/or position. The double product function is depicted in (d) with integral (f). The triple product functions is shown in (e) with integral (g). Green represents the positive part of the basis functions, red the negative part.

where δ_{ij} is the Kronecker delta (Clapham and Nicholson, 2009). The double product integral of such orthonormal bases reduces to a convenient dot product. This can be verified in Figure 1(f). Only basis functions with the same position and dilation result in a non-zero double product integral.

Unfortunately, the triple product integral is substantially more difficult to calculate. An analytical approach was developed for the Haar product integral, which iterates only over non-zero coefficients (Ng et al., 2004). This is possible due to the strongly regular and hierarchical structure of Haar wavelets. Children with a smaller dilation factor fall completely under a constant part of their parent. Figure 1(g) (left) shows that only specific and predictable combinations result in non-zero product integrals.

High-order wavelets, on the other hand, have three properties that prevent them from having simple binding coefficients:

1. Their larger support results in more overlap and circular wrapping.
2. Children are no longer entirely contained by the support of their parent.
3. The subset of the parent support that a child overlaps is not necessarily constant.

It can be seen in Figure 1(g) (right) that the triple product integral is more complex. If we iterate over all combinations of i, j and k and put them in one large tensor we get Figure 2(a) for Haar and Figure 2(b) for Daubechies. With this tensor precalculated, only the sparse non-zero coefficients need to be evaluated at runtime.

The specific contributions of this paper are:

- An efficient algorithm to calculate the tensor with binding coefficients of n-product integrals of a wide range of wavelet basis functions. Mixing of various basis types is supported. The paper focuses on double and triple coefficients, but the approach can easily be extended to quadruple or higher product integrals.
- Analysis of the tensor characteristics such as sparseness and symmetry. We also study the effect of mixing Haar and higher-order bases.
- We demonstrate the effectiveness of our approach with a rendering application.

2 PREVIOUS WORK

Spherical harmonic basis functions were one of the first functions proposed to approximate the factors in the triple product integral (Kautz et al., 2002). Spherical harmonics are an extension of the Fourier transformation to the spherical domain. They have the advantage of being a well studied mathematical tool that can succinctly approximate low-frequency signals. Expressing detailed high-frequency effects with spherical harmonics, however, requires exponentially more coefficients. Another disadvantage is that no efficient analogue of the Fast Fourier Transform (Cooley and Tukey, 1965) exists for harmonic analysis in the spherical domain. A relatively fast alternative algorithm is known, but its time complexity is still superquadratic (Mohlenkamp, 1997). A triple product integral theorem can also be formulated for Legendre polynomials, but due to the global support of this basis, it suffers from the same shortcomings as spherical harmonics (Gupta and Narasimhan, 2007). Their graphical representation of tensor slices inspired the computational approach taken in this paper.

To accurately represent high frequency lighting details, wavelet bases can be utilized (Daubechies,

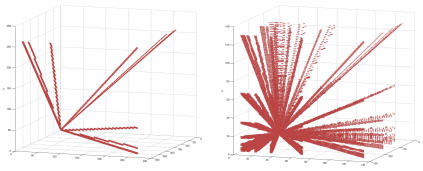

(a) (b)

Figure 2: Tensor of tripling binding coefficients for (a) Haar wavelets and (b) high-order wavelets. The Haar tensor clearly shows the three cases identified by Ng et al. (Ng et al., 2004). High-order wavelets have a broader support, resulting in more non-zero binding coefficients. However, the tensor still remains sparse.

1992). An alternative representation for spherical harmonics in the wavelet domain are spherical wavelets (Sweldens, 1998), which have demonstrated to be a compact and efficient representation. While the possibility to construct an orthogonal spherical wavelet basis with compact support and symmetry has been demonstrated (Lessig and Fiume, 2008), these wavelets are considerably more difficult to construct than their 2D counterparts. Therefore, functions are often parameterized with an area-preserving parameterization (Praun and Hoppe, 2003), so that conventional 2D wavelet analysis can be leveraged.

Ng. et al. were the first to solve a triple product integral of three factors approximated in Haar bases (Ng et al., 2004). They noticed that the product integral tensor of wavelet functions is very sparse and the calculations can be categorized in a small number of cases, which they exhaustively list in their Haar tripling coefficient theorem. They manually studied the different possible wavelet combinations and their outcome. Only a fraction of the wavelet combinations on different levels resulted in a non-zero integral, so that only these particular cases need to be treated to evaluate the triple product integral.

Previous work has also developed a general technique for importance sampling products of complex functions using wavelets (Clarberg et al., 2005). They perform on-the-fly stochastic sampling of the wavelet scaling coefficients to evaluate a double product integral of the BRDF and an environment map. They base their sampling scheme on the characteristics of the Haar wavelet basis and only double product integrals are demonstrated. In addition, they preprocess the 4D BRDFs, so that only sample points need to be evaluated at runtime, but this means the sample pattern does not adapt when either the environment map or the BRDF is dynamic. Also, the random sampling patterns they use introduce considerable noise in the resulting images.

A generalized Haar integral coefficient theorem was proposed for evaluating arbitrary dimensional Haar product integral coefficients (Sun and Mukherjee, 2006). They extend the approach of Ng. et al. and create an efficient sublinear algorithm to evaluate these N-product integrals. As in previous work however, they are limited to simple Haar wavelet bases.

There exists also a geometry-dependent basis for diffuse precomputed radiance transfer (Nowrouzezahrai et al., 2007). Their basis is derived from Principal Component Analysis of the sampled transport functions at each vertex. They only demonstrate double product integral capabilities and the rendering results are diffuse only. Interpolation artifacts also arise due to the dependency of the basis on the geometric representation of the scene.

An affine double and triple product integral theory was developed, enabling one of the product functions to be scaled and translated (Sun and Ramamoorthi, 2009). They demonstrate that these operations are very sparse and scale with linear complexity. This sparsity enables them to add some of the first near-field lighting effects. In their disposition, they give specific attention to the common Haar wavelets and rely on its non-overlapping property. They state that an implementation for non-Haar wavelets is more expensive but that their general approach can be similar applied to general wavelets. They do not, however, provide a solution to this problem.

Previous work has tried to exploit the fact that in areas where the visibility factor is constant, the triple product integral reduces to a double product integral (Inger et al., 2013). Nevertheless, in the areas where a full triple product evaluation is needed, a fall back to an expensive pixel domain integral is still required. In addition, mixing of arbitrary and high-order wavelet bases is not supported.

This paper differs from previous work in that a computational approach is taken to calculate the tensor coefficients for triple product integrals. This allows the use and mixing of a wide range of high-order wavelets, as opposed to simple Haar wavelets. High-order wavelets can compactly approximate the smooth factors in the product. In the next section, the use of these high-order wavelets will be motivated. The rest of this paper will explain our computational approach and evaluate the results.

3 WAVELET PRODUCT INTEGRAL

Many functions are naturally expressed in the spherical domain. Solving the rendering equation (see Equation 2) at each point in space can be considered a spherical convolution operation in signal processing.

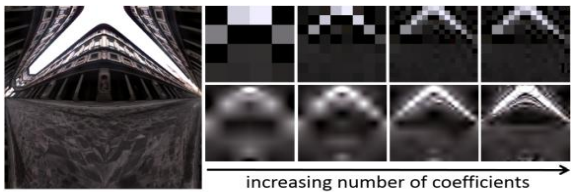

Figure 3: Reconstruction of an environment map (left) with a small number of coefficients. A Haar wavelet basis is used in the top row on the right. The result is much more blocky in comparison to the smoother Daubechies-6 reconstruction in the bottom row.

Each term in the convolution, in the rendering case V, ρ and L, is expanded in an appropriate basis Ψ:

$$V(\omega) = \sum_i V_i \Psi_i(\omega) \tag{4}$$

$$\rho(\omega) = \sum_j \rho_j \Psi_j(\omega) \tag{5}$$

$$L(\omega) = \sum_k L_k \Psi_k(\omega) \tag{6}$$

As described before, previous work either used harmonic analysis to model the functions compactly on the sphere, spherical wavelets or regular 2D Haar wavelets. Harmonic analysis has the advantage of being a natural choice of representation for spherical functions, but requires a high number of coefficients to represent high-frequency detail. Haar wavelets on the other hand are better suited to compactly represent these local details, because the basis functions have a small support. Haar functions lack smoothness, however, which makes them less suitable for the approximation of smooth signals. On the other hand, high-order wavelet bases (e.g. Daubechies) are better tailored to representing smooth signals. Figure 3 shows a comparison between the Haar wavelet basis and the Daubechies-4 wavelet basis. A reconstruction with few coefficients leads to a noticeably more blocky result when using the Haar basis. In this particular example, the original signal requires roughly 5 times less coefficients in the Daubechies basis. This motivates our approach to perform the product integral calculation on this sparser representation.

3.1 The Haar tripling coefficient theorem

The Haar tripling coefficient theorem binds appropriate weighting factors to each of the wavelet coefficients that combine to form the product integral. They use an analytical approach to formulate the very limited set of cases that arise when calculating the triple product integral of simple Haar basis functions. A recursive algorithm is given with sublinear complexity, which calculates only the relevant coefficients. Their manual approach does not scale well when using smoother wavelets. This paper therefore takes a computational approach to the problem and calculates the tripling coefficients for a wide range of wavelets automatically.

Ng et al. realized that the different binding coefficients are very redundant for a simple Haar basis (Ng et al., 2004). Their final tripling coefficient theorem has three cases, with variations for scaling and wavelet combinations. Furthermore, Haar wavelets with smaller dilation factors are always fully contained in a subset of the parents support with constant value. These characteristics allow for a fast sublinear time algorithm, that iterates over all non-zero coefficients.

3.2 General tripling coefficient theorem

In contrast to the Haar wavelets, high-order wavelets are not piecewise constant functions anymore. Since high-order wavelets have a larger support, the number of non-zero combinations will increase rapidly. It is no longer feasible to manually indentify and enumerate all the various cases. Figure 4 visually compares the overlapping properties for Haar and the high-order Daubechies wavelets. It can be seen that the convolution of Haar wavelets yields a much sparser tensor than the convolution of high-order wavelets. Haar wavelets with smaller dilation never overlap more than one coarser Haar wavelet. This is not true for high-order wavelets. The support of the high-order wavelets remains compact, however, only being enlarged by a constant factor. Because of this, the tensor of binding coefficients remains sparse.

3.2.1 Naive approach

It is easy to envision a naive approach to calculate the tensor, based on the permutation of all i, j and k and calculating the binding coefficients C_{ijk} (Equation 2). This quickly becomes intractable, due to its time complexity $O(n^D * r^2)$ with D the dimensionality of the product integral (double, triple, quadruple, ...), n the total number of translations and dilations of the basis function and r^2 the resolution of the 2D wavelet functions that are integrated. If the original signal contains many high-frequency perturbations, n can grow profoundly large. As a result the computation time increases drastically. For only a resolution of 512x512, a triple product integral for 2D wavelets requires an order of 9.5×10^{14} calculations. Clearly more sophisticated methods are required.

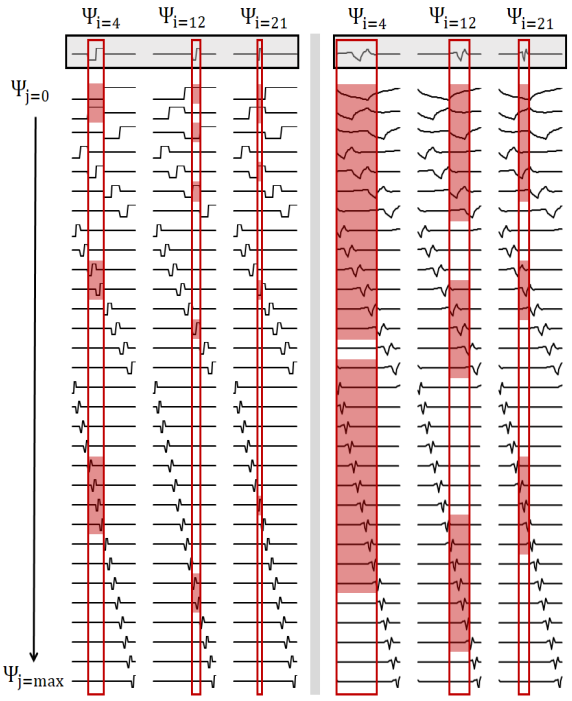

Figure 4: Hierarchically overlapping wavelet basis functions. Each wavelet basis function Ψ_i will have a fixed overlap with a set of other wavelet functions Ψ_j. The three columns on the left represent the overlap in case of Haar wavelets, specifically for $\Psi_{i=4,12,21}$. The three rightmost columns show the analogous cases for the Daubechies-4 wavelets.

3.2.2 Hierarchical approach

A second approach exploits the characteristics of higher-order wavelet bases. By design, each basis function has a fixed support. For example, a Daubechies-4 wavelet starts with a dilation factor of size 4 and will further dilate accordingly (4, 10, 22, 46, ...). As a result, a wavelet basis function Ψ_i will overlap only with a fixed amount of other wavelet basis functions Ψ_j, each having a different position and dilation factor. Figure 4 illustrates the overlap of several basis functions Ψ_i with other basis functions Ψ_j. The figure displays that a basis function with a narrow dilation factor overlaps only with a small number of other functions. Only these specific combinations will possibly result in non-zero binding coefficients. All other combinations can safely be skipped. The time complexity is now reduced to $O(n * C * log^{D-1}(n) * r^2)$ with C a constant related to the enlargement of support for high-order wavelets (in case of Daubechies-4, $C \simeq 3$). The reconstruction of a 512×512 resolution will now take approximately 6.5×10^{10} permutations, which is several orders of magnitude less than the naive approach.

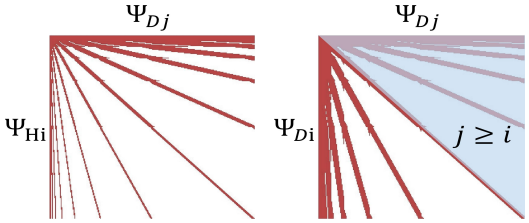

Figure 5: Tensor mirroring. (a) shows the overlap regions for a double product of Daubechies-8 wavelets Ψ_{Di} and Haar wavelets Ψ_{Hj}. (b) shows the overlap regions for a double product where both Ψ_{Di} and Ψ_{Dj} are Daubechies-8 wavelets functions. In the case of two identical wavelet bases, the tensor is fully symmetric. The values in the upper triangle, indicated in blue, can be mirrored into the lower triangle. Note that in all cases the tensor is sparse.

3.2.3 Tensor Mirroring

When the mix of wavelet bases in a product integral is homogeneous, symmetry can be observed in the plotted tensor. For the double product case, the tensor is depicted in Figure 5(b)). Since $\int \Psi_i \, \Psi_j = \int \Psi_j \, \Psi_i$ is true, each combination of basis functions where $j \geq i$ can be mirrored around the diagonal, eliminating half the work. In the case of a triple product integral with i, j and k or a general product integral, taking advantage of the symmetry is done in an analogous manner.

3.2.4 Wavelet sliding

Previous sections have discussed improvements in computational complexity by reducing the number of wavelet permutations. The bottleneck for the overall tensor calculation is not the amount of wavelet permutations (e.g. 10^5 for 512×512), but rather the actual work performed in each permutation (r^2 or 512^2 multiplications for integration). A lot of these calculations are redundant. This section will explain how they can be avoided.

Let us observe two basis functions $\Psi_{i=x}$ and $\Psi_{i=y}$ and their respective sets of overlapping basis functions $O_{i=x}$ and $O_{i=y}$ for which holds:

$$O_i = \{\Psi_j : (\int (\Psi_i \Psi_j) \neq 0, j \geq i\} \qquad (7)$$

In Figure 6, $\Psi_{i=x}$ is circled in red and $O_{i=x}$ is marked in red. $\Psi_{i=y}$ and $O_{i=y}$ are marked green analogously. In the case of $\Psi_{i=x}$ and $\Psi_{i=y}$ having the same dilation factor, the amount of overlapping functions will be equal ($\#O_{i=x} = \#O_{i=y}$). However, the overlapping functions of $\Psi_{i=y}$ are shifted with a sliding factor s:

$$s = (y - x) \times support(\Psi_{i=x}) \qquad (8)$$

Using the observation above, only the overlapping functions of one specific translation need to be cal-

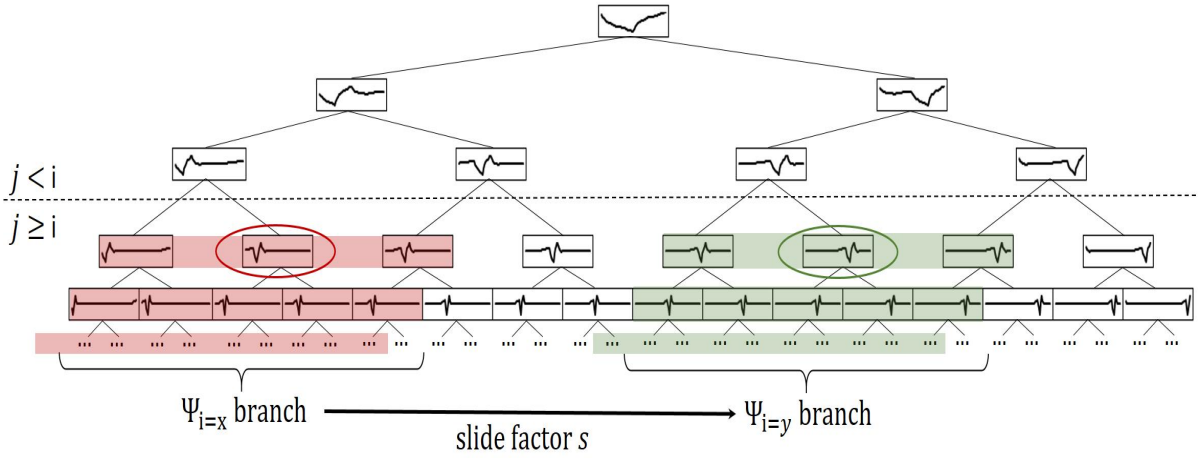

Figure 6: Wavelet sliding. Given a wavelet basis function $\Psi_{i=x}$ (circled in red), there is a branch in the wavelet tree of overlapping basis functions (red). The precalculated binding coefficient values of this branch can be reused for every translated $\Psi_{i=y}$ (circled in green) with the same dilation. The branch is translated with slide factor s (see Equation 8).

culated, e.g. the red area in Figure 6. We call this a branch. The precalculated branch of binding coefficients can then be reused for all other translated basis functions with that specific dilation factor. This will reduce the amount of work drastically to $O(C * log^D(n) * r^2)$. The 512×512 example will now take approximately 3.1×10^8 calculations.

3.2.5 Vanishing moments

An additional advantage of high-order Daubechies wavelets is that they provide the maximum amount of vanishing moments for wavelets of that specific support size (Daubechies, 1992). In general, smoother wavelets generate more vanishing moments. The interesting property of wavelets with vanishing moments is that their product integral is not only zero when the basis functions do not overlap, but even for certain translations within their support. This yields extra sparsity in the tensor. The position of these vanishing moments are predictable and their determination is incorporated in the wavelet sliding algorithm of the previous section. Afterwards, the data entries of vanishing moments can be removed from the precalculated branch.

4 APPLICATIONS

Our main motivation for using high-order wavelets is that they provide a compact and high-quality approximation of smooth functions. This is particularly interesting for triple product integration to evaluate the rendering equation (see Figure 7). The visibility factor is a piecewise constant function, for

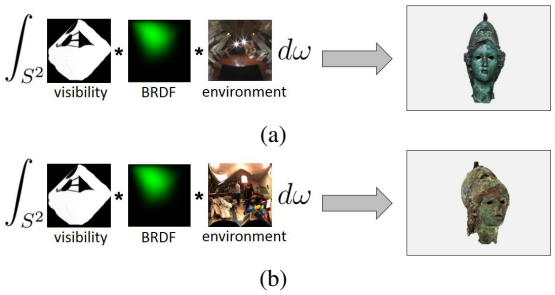

Figure 7: Triple product integration in our rendering application. Solving the rendering equation requires taking an integral over the hemisphere at each point in the scene. The three factors of the product integral are visibility, BRDF (bidirectional reflectance distribution function) and lighting. The output of all the integrals for the entire image results in a rendered viewpoint. (a) and (b) show an example of two rendered views, both with different lighting conditions.

which Haar wavelets are ideally suited. We argue that the lighting environment map and certainly the BRDF (bidirectional reflectance distribution function) exhibit, in general, profoundly more smoothness. These factors are better represented with a smoother wavelet, for example Daubechies. In contrast to previous methods, our algorithm is able to calculate the triple product integral with a heterogeneous mix of wavelet bases, where each factor is coded in a basis specifically tailored to the signal characteristics.

Figures 8 and 9 compare the visual quality of blocky Haar wavelets and smoother Daubechies-6 wavelets for various models and compression rates. The ground truth rendering at full quality is included for reference, alongside zoomed pictures on areas with the largest differences. The L_2 norm is pro-

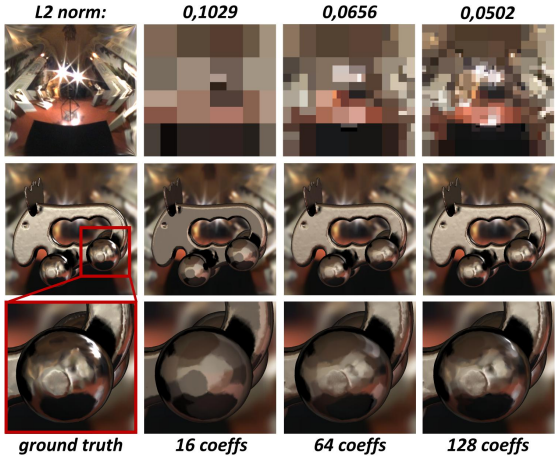

(a) Haar

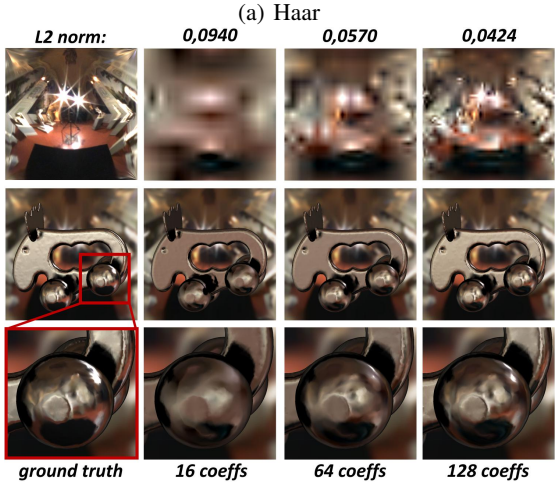

(b) Daubechies-6

Figure 8: Quality comparison of the Elch dataset. The environment map is compressed with: (a) the Haar wavelet basis and (b) the Daubechies-6 wavelet basis. The three right columns show the quality differences with respectively 16, 64 and 128 coefficients retained. In (a), the mesh is rendered with the Haar triple product integral (Ng et al., 2004). (b) is rendered with our tensor calculation. The results in (b) have smaller L_2 norms and converge faster to the ground truth (depicted in the left column).

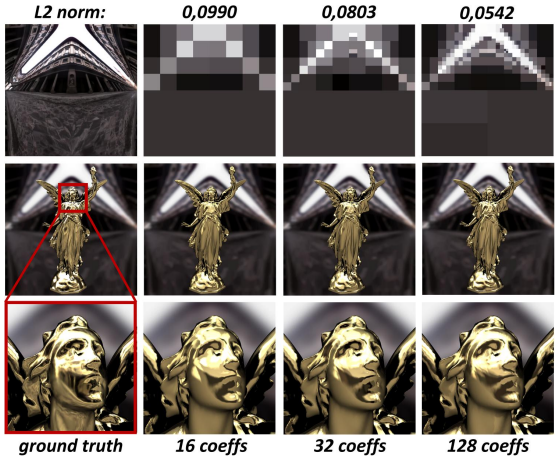

(a) Haar

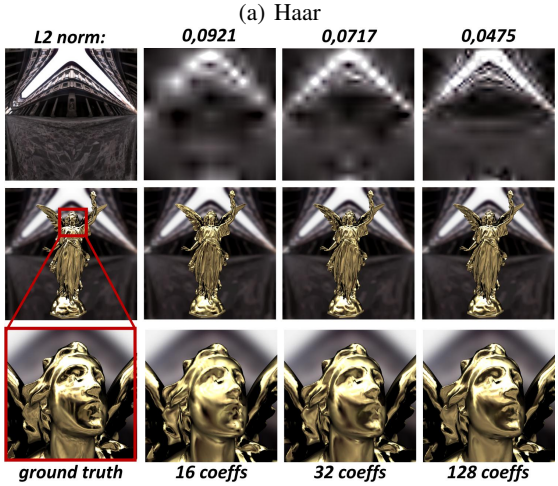

(b) Daubechies-6

Figure 9: Quality comparison of the Lucy dataset. The environment map is compressed with: (a) the Haar wavelet basis and (b) the Daubechies-6 wavelet basis. The three right columns show the quality differences with respectively 16, 32 and 128 coefficients retained. In (a), the mesh is rendered with the Haar triple product integral (Ng et al., 2004). (b) is rendered with our tensor calculation. The results in (b) have smaller L_2 norms and converge faster to the ground truth (depicted in the left column).

vided with each rendering as a quantitative comparison measure. A rendering with a smaller norm is closer to the ground truth. It can be seen that smooth high-order wavelets outperform Haar wavelets both qualitatively and quantitatively.

Figure 10 demonstrates the ability of our application to render with different kinds of wavelets. It is possible to mix and match wavelets to optimally represent the signals of all factors in the product integral calculation. In general, the smoothness characteristics of the wavelet basis should match those of the signal. In that case the wavelets will be able to represent the signal with a minimum of coefficients.

5 CONCLUSION

This paper has provided a method to calculate the binding coefficients of general product integrals. Our method is able to cope with a mix of heterogeneous bases. This allows the representation of factors in the product integral with piecewise constant or smooth basis functions, depending on the signal properties. The algorithm exploits both the hierarchical nature and vanishing moments of the wavelets basis, as well as the sparsity and symmetry of the tensor. Our rendering application demonstrates that the tensor-based product integral leads to less blockiness in the results.

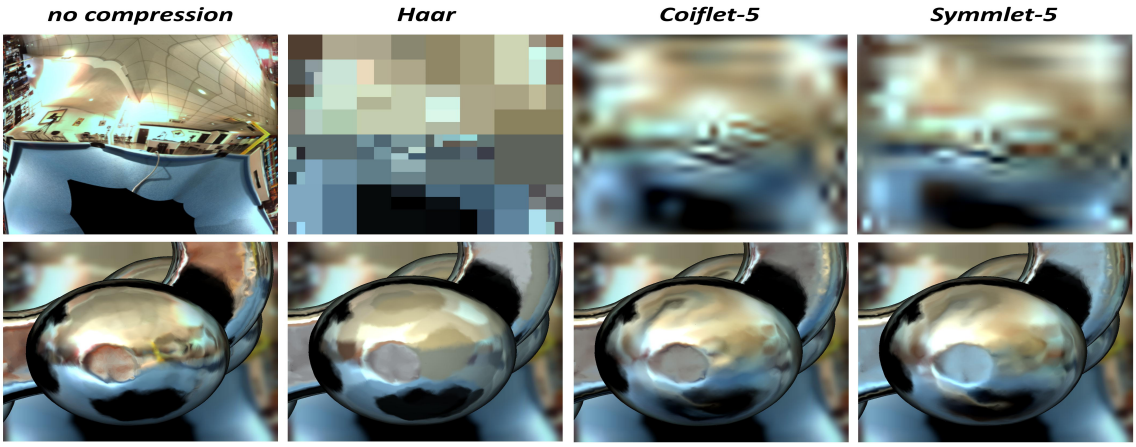

Figure 10: Product integral calculation with different wavelet basis functions. In this case, a compression with Haar, Coiflet-5 and Symmlet-5 is executed on the environment map. The result of the triple product integral on the Elch dataset is shown respectively.

ACKNOWLEDGEMENTS

The authors acknowledge financial support by the European Commission (FP7 IP SCENE), the European Regional Development Fund (ERDF), the Flemish Government, iMinds and IWT.

REFERENCES

Clapham, C. and Nicholson, J. (2009). *The Concise Oxford Dictionary of Mathematics*. Oxford Paperback Reference. OUP Oxford.

Clarberg, P., Jarosz, W., Akenine-Möller, T., and Jensen, H. W. (2005). Wavelet importance sampling: efficiently evaluating products of complex functions. In *ACM SIGGRAPH 2005 Papers*, SIGGRAPH '05, pages 1166–1175, New York, NY, USA. ACM.

Cooley, J. and Tukey, J. (1965). An algorithm for the machine calculation of complex fourier series. *Mathematics of Computation*, 19(90):297–301.

Daubechies, I. (1992). *Ten Lectures on Wavelets*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA.

Gupta, M. and Narasimhan, S. G. (2007). Legendre polynomials triple product integral and lower-degree approximation of polynomials using chebyshev polynomials. Technical Report CMU-RI-TR-07-22, Robotics Institute, Pittsburgh, PA.

Inger, Y., Farbman, Z., and Lischinski, D. (2013). Locally adaptive products for all-frequency relighting. *Computer Graphics Forum (Proceedings of Eurographics 2013)*, 32(2pt1):73–82.

Kajiya, J. T. (1986). The rendering equation. *SIGGRAPH Comput. Graph.*, 20(4):143–150.

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, pages 291–296, Aire-la-Ville, Switzerland, Switzerland. Eurographics Association.

Lessig, C. and Fiume, E. (2008). SOHO: Orthogonal and Symmetric Haar Wavelets on the Sphere. *ACM Trans. Graph.*, 27(1).

Michiels, N., Put, J., Haber, T., Klaudiny, M., and Bekaert, P. (2013). High-order wavelets for hierarchical refinement in inverse rendering. In *ACM SIGGRAPH 2013 Posters*, SIGGRAPH '13, pages 99:1–99:1, New York, NY, USA. ACM.

Mohlenkamp, M. J. (1997). *A Fast Transform for Spherical Harmonics*. PhD thesis, New Haven, CT, USA. UMI Order No. GAX97-33952.

Ng, R., Ramamoorthi, R., and Hanrahan, P. (2004). Triple product wavelet integrals for all-frequency relighting. In *ACM SIGGRAPH 2004 Papers*, SIGGRAPH '04, pages 477–487, New York, NY, USA. ACM.

Nowrouzezahrai, D., Simari, P., Kalogerakis, E., and Fiume, E. (2007). Eigentransport for efficient and accurate all-frequency relighting. In *Proceedings of the 5th international conference on Computer graphics and interactive techniques in Australia and Southeast Asia*, GRAPHITE '07, pages 163–169, New York, NY, USA. ACM.

Praun, E. and Hoppe, H. (2003). Spherical parametrization and remeshing. *ACM TOG*, 22(3):340–349.

Sun, B. and Ramamoorthi, R. (2009). Affine double- and triple-product wavelet integrals for rendering. *ACM Trans. Graph.*, 28(2):14:1–14:17.

Sun, W. and Mukherjee, A. (2006). Generalized wavelet product integral for rendering dynamic glossy objects. In *ACM SIGGRAPH 2006 Papers*, SIGGRAPH '06, pages 955–966, New York, NY, USA. ACM.

Sweldens, W. (1998). The lifting scheme: A construction of second generation wavelets. *SIAM J. Math. Anal.*, 29(2):511–546.