Discontinuity Meshing and Hierarchical MultiWavelet Radiosity

S. N. Pattanaik and Kadi Bouatouch
IRISA, Campus de Beaulieu, 35042 Rennes Cedex, France

Abstract

In this paper we propose an extension to the work of Lischinski *et al* [?] on combining hierarchical radiosity method and discontinuity meshing. The extension proposed here supports higher order radiosity function at the estimation step itself. We have carried out this extension using Multi-Wavelet basis functions. Unlike the multi-wavelet radiosity work of Gortler *et al* [?] which uses regular quadtree subdivision, we carry out the subdivision along the discontinuity boundary. In other words, instead of deriving the finer basis functions by uniform parametric dilations of a mother basis function, we derive the finer basis by mapping the mother basis to arbitrary subdomains created during the subdivision along the discontinuity boundary. To carry out *push* and *pull* operations, which are crucial to the hierarchical algorithm, we compute the necessary filter functions. Our method combines the advantages of wavelet radiosity and discontinuity meshing. The preliminary result shows significant computational improvement.

1 Introduction

As of today, hierachical method is the state of the art method of solving radiosity in any general environment. This method was first introduced in [?, ?] to estimate radiosity as piece-wise constant functions over the environment. Subsequently by using multi-wavelet basis functions with higher vanishing moments (> 1) the method was extended [?] to directly estimate radiosity as piecewise higher order polynomial functions. Directly computing higher order functions implies lesser amount of discretisation and hence lesser overall effort in the illumination computation. These methods require a regular discretisation of surfaces. The surfaces needing subdivisions were subdivided by half in each of their parametric dimensions. This type of subdivision gives the best average performance when it is not at all possible to derive the complexity of the underlying radiosity function. However, in certain cases it may be possible to derive some information on the complexity of the function. In such case, instead of carrying out the regular subdivision to finally localise on the complex region it will be most appropriate to use the derived complexity information to directly localise the complexity. We shall take an example of a shadow discontinuity shown in figure 1. Localising this discontinuity by regular subdivision will require a very large number of subdivisions. However, if the discontinuity can be known by some other extraneous method (say discontinuity meshing) then a single subdivision of the surface will be sufficient to capture the complexity. This example tends to suggest that at the pre-processing stage we carry out the subdivision of the environment using a discontinuity mesher. However, such preprocessing approach are extremely expensive because,

- in the absence of the knowledge of the actual illumination distribution in the environment, the pre-discretisation step is likely to create far more number of discrete surfaces than actually required, and
- pre-discretisation would mean an increase in the number of individual interacting elements in the environment, hence a quadratic increase in the illumination computation time.

Figure 1: Discontinuity and Wavelet Dilation Functions.

To overcome this problem Lischinski *et al* [?] introduced the discontinuity driven hierarchical radiosity method in which they combined the advantages of hierarchical radiosity and discontinuity meshing. They carried out the normal hierachical radiosity algorithm to compute interaction between surfaces in the environment, but at the subdivision step of the hierarchical algorithm, instead of using the regular quadtree subdivision, they used irregular subdivision of surfaces along discontinuity boundaries. However, they restricted their method to the estimation of piece-wise constant radiosity functions. For better visual reconstruction of the estimated radiosity function, only in a post-processing step they carried out higher order (quadratic) interpolation.

We extend this work of Lischinski *et al* by directly estimating piecewise higher order radiosity function. Thus work combines the advantages of wavelet-radiosity and discontinuity meshing. As in Gortler *et al* [?] we have used Multi-wavelet basis function to carry out the extension.

The organisation of the paper is as follows. We briefly introduce the function approximation and hierarchical radiosity. Then we derive the filters to carry out the crucial push/pull operation. Finally we demonstrate its successful application to a simple test environment.

2 Radiosity Function Approximation and Hierarchical Algorithm

In an environment with diffusely reflecting and/or emitting surfaces the radiosity of any surface p can be expressed using the following equation:

$$B_p(\bar{x}) = E_p(\bar{x}) + \sum_{q=1}^{N} \int_{A_q} \kappa_{q\to p}(\bar{x},\bar{y})\ B_q(\bar{y}) d\bar{y} \tag{1}$$

where N is the number of surfaces in the environment, $\bar{x}$ and $\bar{y}$ are points and $B(\bar{x})$, $B(\bar{y})$ the radiosity functions over the surfaces p and q respectively, $E_p(\bar{x})$ is the emittance function over the surface p and $\kappa_{q\to p}$, the *kernel* of the integral operator, represents the interaction between surfaces q and p and can be expressed as

$$\kappa_{q\to p}(\bar{x},\bar{y}) = \rho_p(\bar{x}) \frac{\cos\theta_x \cos\theta_y}{\pi r^2} V(\bar{x},\bar{y}) \tag{2}$$

where ρ is the diffuse reflectivity at a point, r the distance between $\bar{x}$, $\bar{y}$ and V the visibility between $\bar{x}$, $\bar{y}$, and θ_x, θ_y respectively are the angles between the line joining $\bar{x}$ and $\bar{y}$, and the surface normals at those points.

Solution of this type of equations is often carried out by using funcion approximation technique. This method seeks for the exact solution by projecting the involved functions onto a finite dimensional space, *i.e.* the radiosity function B is approximated by a linear combination of basis functions:

$$B(u,v) = \sum_{k=1}^{n} b_k \mathcal{N}_k(u,v)$$

where n is the number of basis functions, b_k's are unknown approximation coefficients and $\mathcal{N}_k$'s are the chosen basis functions.
A set of linear equations [?] of the type

$$b_{p,i} = \sum_{q=1}^{N} \sum_{j=1}^{n_q} K_{q\to p,ij} b_{q,j} \tag{3}$$

where $i = 1..n_p$ and n_p, n_q are the number of approximation terms for $B_p(u,v)$ and $B_q(s,t)$,

$$b_{q,j} = \int_{u=0}^{1}\int_{v=0}^{1} \mathcal{N}_j(u,v) B_p(u,v) du dv,\ b_{q,i} = \int_{s=0}^{1}\int_{t=0}^{1} \mathcal{N}_i(s,t) B_q(s,t) ds dt$$

and $K_{q\to p,ij}$ $\int_{u=0}^{1}\int_{v=0}^{1}\int_{s=0}^{1}\int_{t=0}^{1} \kappa_{q\to p}(u,v,s,t)\mathcal{N}_i(u,v)\mathcal{N}_j(s,t) \left\| \frac{\delta x}{\delta s} \times \frac{\delta x}{\delta t} \right\| ds dt du dv$

Figure 2: Push-Pull operation.

are then derived by using orthogonal basis functions and Galerkin error minimisation techqnique, and the solution is carried out by following an iterative method such as Gauss-Seidel method or Southwelll relaxation method till convergence. The complexity of solving this system for the whole environment is $\mathcal{O}(\sum n_p)^2$. The accuracy of the resulting solution depends on how correctly the radiosity function on each surface has been approximated to arrive at the linear equation. The brute force method of improving the accuracty is to increase the number of basis function n_p for the approximation of radiosity function of each surface p. However, as the complexity of the technique indicates the quadratic increase in computation time with the increase in the number of basis function, one must make optimal use of the number of basis function for any appromixation.

Hierarchical algorithm [**?**] addresses this problem by making some crucial observations. They are :

- To write an expression for the unknown radiosity function of surface p of the type given in equation (3), we need the expression of radiosity function of all other surfaces q which are fully or partially visible to p.
- It is possible to decide on the maximum discretisation of each surface for any given accuracy in the solution of the radiosity function.
- If a particular accuracy in computation of radiosity of surface p requires a finer discretisation of surface q, then it is not necessary that for the same computational accuracy in radiosity of another surface r will require the discretisation of q to the same fineness. Probably a coarser discretisation or much finer discretisation may be required.

The last observation is very crucial because if we can use different levels of discretisation of a particular surface to set up the system of equations of the type (3) then the overall complexity of the solution is bound to be much less than $\mathcal{O}(\sum n_p)^2$ where n_p is the maximum discretisation of the surface p. Making use of this observation demands that one must

- associate a hierarchy of subdivisions with each surface and
- at anytime during the iterative solution of radiosity, maintain with each level of the hierarchy the radiosity information commensurate to the level.

Maintaining radiosity function at different levels during the iteration requires 2 operation known as push and *pull* operation. In the next section we describe the basic ingradients necessary for these operations.

3 Push/Pull Operation

Let us consider the simplest case of an one-level surface discretisation shown in figure 2, where we have only 2 levels of discretisation associated with a surface. Given the appoximation of a function, $B(\bar{x})$, defined over the level 0, the *push* operation computes the approximation of the the function at level 1. And similarly, given the approximation at level 1, the *pull* operation computes the approximation of the function at level 0. Let us define three sets of orthonormal basis functions $N_0^{(m,n)}(\bar{x})$, $N_1^{(m,n)}(\bar{x})$ and $N_2^{(m,n)}(\bar{x})$ such that they are zero outside the domains Ω_0, Ω_1 and Ω_2, and if we assume that a unit square in parametric domains (u, v), (u', v') and (u'', v'') span Ω_0, Ω_1 and Ω_2 respectively then

$$\begin{aligned} N_0^{(m,n)}\left(\bar{x}(u,v)\right) &= \mathcal{M}^{(m,n)}(u,v) \\ N_1^{(m,n)}\left(\bar{x}(u',v')\right) &= \mathcal{M}^{(m,n)}(u',v') \\ N_2^{(m,n)}\left(\bar{x}(u'',v'')\right) &= \mathcal{M}^{(m,n)}(u'',v'') \end{aligned}$$

where $\mathcal{M}^{(m,n)}$'s are the scaling function of a multiwavelet basis of vanishing moment M.

With these definitions we can write the approximation of the radiosity function at the given two levels as follows:

Approximation at Level 0 :

$$\begin{aligned} B(\bar{x}) &\approx \sum_{m=1}^{M}\sum_{n=1}^{M} b_0^{(m,n)} N_0^{(m,n)}(\bar{x}) \qquad (4)\\ b_0^{(m,n)} &= \int_{u=0}^{1}\int_{v=0}^{1} B(\bar{x}) N_0^{(m,n)}(\bar{x}) du dv \end{aligned}$$

Approximation at Level 1 :

$$\begin{aligned} B(\bar{x}) &\approx \sum_{m=1}^{M}\sum_{n=1}^{M} b_1^{(m,n)} N_1^{(m,n)}(\bar{x}) + \sum_{m=1}^{M}\sum_{n=1}^{M} b_2^{(m,n)} N_2^{(m,n)}(\bar{x}) \qquad (5)\\ \text{where } b_1^{(m,n)} &= \int_{u'=0}^{1}\int_{v'=0}^{1} B(\bar{x}) N_1^{(m,n)}(\bar{x}) du' dv' \text{ and } b_2^{(m,n)} = \int_{u''=0}^{1}\int_{v''=0}^{1} B(\bar{x}) N_2^{(m,n)}(\bar{x}) du'' dv'' \end{aligned}$$

It must be noted that $\bar{x}$ in each of the above equations has the appropriate parametric expression.

Push Filters : As explained above, *push*ing amounts to computing the approximation coefficients $b_1^{(m,n)}$'s and $b_2^{(m,n)}$'s from the approximation coefficients $b_0^{(m,n)}$'s.

$$\begin{aligned} b_1^{(m,n)} &= \int_{u'=0}^{1}\int_{v'=0}^{1} B(\bar{x}) N_1^{(m,n)}(\bar{x}) du' dv' = \int_{u'=0}^{1}\int_{v'=0}^{1} \left[\sum_{p=1}^{M}\sum_{q=1}^{M} b_0^{(p,q)} N_0^{(p,q)}(\bar{x})\right] N_1^{(m,n)}(\bar{x}) du' dv'\\ &= \sum_{p=1}^{M}\sum_{q=1}^{M}\left[\int_{u'=0}^{1}\int_{v'=0}^{1} N_0^{(p,q)}(\bar{x}) N_1^{(m,n)}(\bar{x}) du' dv'\right] b_0^{(p,q)} = \sum_{p=1}^{M}\sum_{q=1}^{M} c_{0\to 1}^{(m,n,p,q)} b_0^{(p,q)} \end{aligned}$$

and similarly

$$\begin{aligned} b_2^{(m,n)} &= \sum_{p=1}^{M}\sum_{q=1}^{M} c_{0\to 2}^{(m,n,p,q)} b_0^{(p,q)}\\ \text{where}&\\ c_{0\to 1}^{(m,n,p,q)} &= \int_{u'=0}^{1}\int_{v'=0}^{1} N_0^{(p,q)}(\bar{x}) N_1^{(m,n)}(\bar{x}) du' dv' \text{ and } c_{0\to 2}^{(m,n,p,q)} = \int_{u''=0}^{1}\int_{v''=0}^{1} N_0^{(p,q)}(\bar{x}) N_2^{(m,n)}(\bar{x}) du'' dv'' \end{aligned}$$

Pull Filters : These filters are responsible for computing $b_0^{(m,n)}$'s from $b_1^{(m,n)}$'s and $b_2^{(m,n)}$'s.

$$\begin{aligned} b_0^{(m,n)} &= \int_{u=0}^{1}\int_{v=0}^{1} B(\bar{x}) N_0^{(m,n)}(\bar{x}) du dv\\ &= \int_{u=0}^{1}\int_{v=0}^{1}\left[\sum_{p=1}^{M}\sum_{q=1}^{M} b_1^{(p,q)} N_1^{(p,q)}(\bar{x}) + \sum_{p=1}^{M}\sum_{q=1}^{M} b_2^{(p,q)} N_2^{(p,q)}(\bar{x})\right] N_0^{(p,q)}(\bar{x}) du dv\\ &= \sum_{p=1}^{M}\sum_{q=1}^{M}\left[b_1^{(p,q)}\int_{u=0}^{1}\int_{v=0}^{1} N_1^{(p,q)}(\bar{x}) N_0^{(m,n)}(\bar{x}) du dv + b_2^{(p,q)}\int_{u=0}^{1}\int_{v=0}^{1} N_2^{(p,q)}(\bar{x}) N_0^{(m,n)}(\bar{x}) du dv\right]\\ &= \sum_{m=1}^{M}\sum_{n=1}^{M}\left[c_{1\to 0}^{(m,n,p,q)} b_1^{(p,q)} + c_{2\to 0}^{(m,n,p,q)} b_2^{(p,q)}\right] \end{aligned}$$

where

$$c_{1\to 0}^{(m,n,p,q)} = \int_{u=0}^{1}\int_{v=0}^{1} N_1^{(p,q)}(\bar{x})N_0^{(m,n)}(\bar{x})dudv \quad \text{and} \quad c_{2\to 0}^{(m,n,p,q)} = \int_{u=0}^{1}\int_{v=0}^{1} N_2^{(p,q)}(\bar{x})N_0^{(m,n)}(\bar{x})dudv$$

Relation Between Push and Pull Filters : As we are using orthogonal basis functions, it may be worthwhile to see if there exists any relationship between the *push/pull* filters, which can reduce the computational effort for evaluating these filters. We shall first try to find this relationship between $c_{1\to 0}^{(m,n,p,q)}$ and $c_{0\to 1}^{(m,n,p,q)}$. As derived above, $c_{1\to 0}^{(m,n,p,q)}$ is an integral over domain Ω_0 whereas $c_{0\to 1}^{(m,n,p,q)}$ is an integral over domain Ω_1. As both are biparametric surfaces, it is possible to find a parametric mapping $\mathcal{F}$ from the Ω_0 to Ω_1, *i.e.*

$$\mathcal{F} : (u, v) \to (u', v')$$

Using this mapping we shall find the relationship as follows:

$$\begin{aligned} c_{1\to 0}^{(m,n,p,q)} &= \int_{u=0}^{1}\int_{v=0}^{1} N_1^{(p,q)}(\bar{x})N_0^{(m,n)}(\bar{x})dudv \\ &= \int_{u'=0}^{U}\int_{v'=0}^{V} N_1^{(p,q)}(\bar{x})N_0^{(m,n)}(\bar{x}) \left\| \frac{\delta u}{\delta u'} \times \frac{\delta v}{\delta v'} \right\| du'dv' \quad \text{where } \mathcal{F} : [0,1]\times[0,1] \to [0,U]\times[0,V] \\ &= \int_{u'=0}^{1}\int_{v'=0}^{1} N_1^{(p,q)}(\bar{x})N_0^{(m,n)}(\bar{x}) \left\| \frac{du}{du'} \times \frac{dv}{dv'} \right\| du'dv' \quad \text{as } \Omega_1 \subset \Omega_0 \end{aligned}$$

This derivation tells us that, when $\|\frac{du}{du'} \times \frac{dv}{dv'}\|$ is a constant function ($\mathcal{C}$), we can find a simplified relationship as follows:

$$c_{1\to 0}^{(m,n,p,q)} = \mathcal{C}\int_{u'=0}^{1}\int_{v'=0}^{1} N_1^{(p,q)}(\bar{x})N_0^{(m,n)}(\bar{x})du'dv' = \mathcal{C}\ c_{0\to 1}^{(m,n,p,q)}$$

For example, when we have subdivision by uniform dilation of a factor 2, as in [?, ?] then we have $\|\frac{du}{du'} \times \frac{dv}{dv'}\|$ is constant and is equal to 0.25. So $c_{1\to 0}^{(m,n,p,q)} = 0.25\ c_{0\to 1}^{(m,n,p,q)}$.
However, for arbitrary subdivisions, it will not be possible to find any such simple relationship. Thus one has to compute the push filters and the pull filters separately.

So far we have discussed the operation from the top most level to the next level. The same discussion can be extened to any pair of levels. For every level of discretisation we have to compute the equivalent filter function. It must be noted here that similar filters were also required for the *push/pull* using multi-wavelets with uniform dilation. But this uniform dilation allowed one

- to use pre defined filter coefficients, and
- the filter coefficents were same between any pair of consecutive discretisation levels.

Where as, in the push/pull operation with arbitrary discretisation

- the filter function is likely to vary for every pair of consecutive discretisation level, and
- these must be computed at each level during the discretisation process.

4 Re-parameterisation

Subdivision of a biparametric surface along an arbitrary boundary may lead to patches which are difficult to directly parameterise (example: figure 3(a)). All our above discussion assumes that we are able to map the mother multi-wavelet basis function to the domain of the subdivided patch, we must find a mechanism of reparameterisation. To do this, we use a very simple approach. The approach is demonstrated in figure 3(b) which avoids the parameterisation problem posed int figure 3(a). We make sure that each discretisation step leads to discretisation of only one parametric dimension. If it is not so, we introduce another extraneous step to guarantee this discretisation.

Figure 3: Re-parametrisation Issue.

Figure 4: A typical scene.

5 Results

We show here the advantages of applying higher order approximation with subdivision across discontinuity boundary using a simple test scene given in figure 4. Figures 5, 6, 7 and 8 compile the results obtained with various strategies. The strategies are respectively:
(a) uniform subdivision with constant basis function (?? patches),
(b) subdivision at discontinuity boundary with constant basis function (?? patches),
(c) uniform subdivision with multi-wavelet basis function of vanishing moment 4 (262 patches), and
(d) subdivision at discontinuity boundary with multi-wavelet basis function of vanishing moment 4 (3 patches).
In all these cases the subdivision has been carried out keeping the error-bound same.
Note that working with multi-wavelet basis function of vanishing moment 4 involves about 64 times more work for each patch in the given case. So if we compare uniform subdivision results, even though higher order basis functions results in a smaller number patches, still then for the given error bound the overall effort far exceeds that due to constant basis function. However, the result due to arbitrary subdivision with higher order basis function performs much favourably compared to all other strategy. We are carrying out the actual test with real large scenes.

6 Conclusion

We believe that our proposed technique can be used in the complex environments with substantial benefits. We are carrying out these tests. In the discussion of the paper, we have emphacised on the subdivision along the discontinuity boundary. However, the method is not limited to this. As long as one is able to decide on a best boundary of subdivision one can apply the above method. We are planning to extend the method to the adaptive mesh generation work of Campbell and Fussel [**?**].

References

[1] Dani Lischinski, Filippo Tampieri, and Don Greenberg. Combining hierarchical radiosity and discontinuity meshing. *Computer Graphics (SIGGRAPH '93 Proceedings)*, 27(4):199–208, Aug. 1993.

[2] Steven Gortler, Peter Schroder, Michel F. Cohen, and Pat Hanrahan. Wavelet radiosity. *Computer Graphics (SIGGRAPH '93 Proceedings)*, 27(4):221–230, 1993.

[3] Pat Hanrahan and David Salzman. A rapid hierarchical radiosity algorithm for unoccluded environments. *Proceedings Eurographics Workshop on Photosimulation, Realism and Physics in Computer Graphics*, pages 151–171, June 1990.

[4] Pat Hanrahan, David Salzman, and Larry Aupperle. A rapid hierarchical radiosity algorithm. *Computer Graphics (SIGGRAPH '91 Proceedings)*, 25(4):197–206, July 1991.

[5] Harold R. Zatz. Galerkin radiosity. *Computer Graphics (SIGGRAPH '93 Proceedings)*, 27(4):213–220, 1993.

[6] A.T. Campbell, III, and Donald S. Fussell. Adaptive mesh generation for global diffuse illumination. *Computer Graphics (SIGGRAPH '90 Proceedings)*, 24(4):155–164, Aug. 1990.

Figure 5: Uniform meshing + Constant Basis.

Figure 6: Discontinuity meshing + Constant Basis.

Figure 7: Uniform meshing + Higher Order Basis.

Figure 8: Discontinuity meshing + Higher Order Basis.