Simplifying the Representation of Radiance from Multiple Emitters

George Drettakis

iMAGIS / IMAG ⋆

In recent work radiance function properties and discontinuity meshing have been used to construct high quality interpolants representing radiance. Such approaches do not consider the combined effect of multiple sources and thus perform unnecessary discontinuity meshing calculations and often construct interpolants with too fine subdivision. In this research we present an extended structured sampling algorithm that treats scenes with shadows and multiple sources. We then introduce an algorithm which simplifies the mesh based on the interaction of multiple sources. For unoccluded regions an a posteriori simplification technique is used. For regions in shadow, we first compute the maximal umbral/penumbral and penumbral/light boundaries. This construction facilitates the determination of whether full discontinuity meshing is required or whether it can be avoided due to the illumination from another source. An estimate of the error caused by potential simplification is used for this decision. Thus full discontinuity mesh calculation is only incurred in regions where it is necessary resulting in a more compact representation of radiance.

1 Sampling Illumination from Multiple Sources

To accurately render scenes illuminated by area light sources, it is necessary to represent the illumination on surfaces by a simpler, approximating function, even when considering only direct illumination. Piecewise polynomial interpolants are often chosen for this purpose. Such representations are an essential requirement for global illumination computation, in particular for the finite-element style approaches (e.g. [Zatz93, GSCH93]), which extend the radiosity-based method [CoGr85].

In the interpolant construction algorithms presented to date, much effort has been devoted to correctly treating shadow boundaries and identifying the behaviour of radiance. These methods have thus achieved high quality representation of illumination using simple functions. However, despite the significant advances in the field, little has been done to actually compensate for the cumulative effects of illumination from multiple emitters, be they light sources or secondary reflectors.

The importance of identifying these interactions is easy to see: when a single source is present, it may cast a detailed shadow which may require significant computation to

⋆ iMAGIS is a joint research project of CNRS/INRIA/UJF/INPG. Postal address: B.P. 53, F-38041 Grenoble Cedex 9, France. E-mail: `George.Drettakis@imag.fr`.

represent correctly. However, if a second source illuminates the same region in an unobstructed fashion, the shadow will be "washed out" leaving little need for the detailed representation. This phenomenon is illustrated in Fig. 7(a) and (b) (see colour section), in which one and two sources illuminate the environment respectively.

In this paper we propose a solution to this problem, by extending the techniques developed for discontinuity meshing and structured sampling [DrFi93, DrFi94, Dret94]. Throughout we consider only environments of diffusely reflecting surfaces lit by diffusely emitting sources. In the following section we present relevant previous work; we then present the extended structured sampling approach. We then briefly describe the two-pass discontinuity meshing algorithm which incurs the cost of full discontinuity meshing only in the regions required. In the sections that follow, we describe the simplification criteria for two sources, first for the intersection of unoccluded regions and then for the intersection of penumbral/unoccluded regions. For both cases we present first results of a prototype implementation. We then present the extension to multiple sources and summarise the results of the paper.

2 Previous Work

In previous work, the approximations used to represent radiance or radiosity have generally been guided by the requirements of the global illumination calculations. A simpler approach to constructing radiance representations is to examine illumination from a single emitter. The first such approach, in which the nature or structure of radiance for unoccluded regions is examined, was presented by Campbell and Fussell [Camp91]. They observed that radiance for these environments displays a single maximum. This idea was extended by Drettakis and Fiume [DrFi93], who constructed quadratic or linear interpolants tensor-product interpolants which can be shown to satisfy tight error bounds.

It has recently been shown that the computation of shadow boundaries, which are subsequently used to guide interpolant construction, is fundamental for high quality approximation of illumination. The first such work was performed in [Camp91] in which the boundary between penumbral and unoccluded regions was computed. The resulting mesh was then used to build an approximation of radiance of constant-radiance triangular elements. Similar work was performed by Chin and Feiner [ChFe92]. Lischinski et al. [LiTG92] were the first to consider discontinuity surfaces interior to the penumbra, that signify a change in the topological view of the light source (e.g. the appearance or disappearance of a vertex or an edge in the visible portion of the source). They subsequently built a triangulation of the receiver surfaces based on the subdivision of this mesh, and constructed quadratic interpolants over these triangles. A different algorithm was presented by Heckbert [Heck92], in which a similar mesh is computed. Lischinski et al. [LiTG92] also merged the meshes from multiple sources, but no simplification was attempted. In this paper we extend the approach developed in [DrFi94, Dret94]. In this approach the *complete discontinuity mesh* is calculated: the environment is segmented into regions (mesh *faces*), in which the topological structure of the visible region of the source does not change. An abstract representation of the visible part of source, called a *backprojection*, is stored with each face. An example of such a mesh is shown

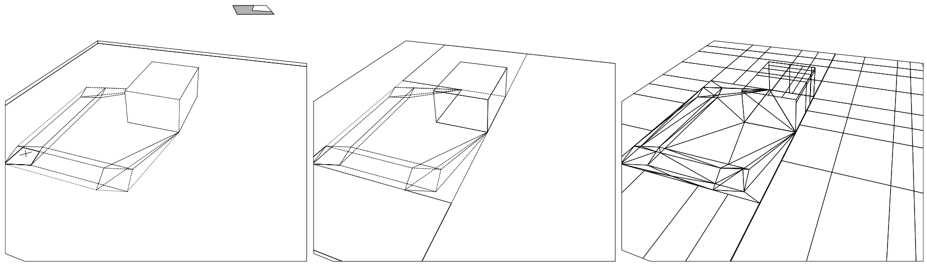

Fig. 1. (a) Mesh and Backprojection (b) Segmentation into Light and Penumbra and (c)Triangular/Tensor Product Interpolants

in Fig. 1(a), where the backprojection of the point marked with a cross is shaded on the source. In [Dret94], the complete mesh is used to construct linear and quadratic interpolants representing radiance in the penumbra. In addition, the structured sampling approach of [DrFi93] was extended ([Dret94]) to handle environments with shadows in the following way. First all regions of shadow are identified and enclosed in a bounding box. Such a regular region enclosing a region of penumbral and umbral faces is called a *penumbral group*. The remaining parts of the receivers (which are unoccluded) are segmented into parallelograms (Fig. 1(b)) on which the structured sampling algorithm is used to create tensor-product interpolants as in [DrFi93] (Fig. 1(c)). Notice in Fig. 1(c) how in the regions of penumbra triangular interpolants are used, while in the unoccluded regions sparse tensor product representations suffice.

3 Extending Structured Sampling for Multiple Emitters

For the purpose of computing reference images in scenes with multiple sources, the discontinuity mesh from each source can be computed independently, and stored with the surface. When rendering using ray-casting, each mesh is queried, the backprojection retrieved for each mesh corresponding to each source and the exact radiance value computed.

To obtain the merged mesh due to several sources, the meshes corresponding to each source are combined. This is performed simply by adding the faces of one mesh into the other. If two light regions with tensor-products are combined, the merged region will contain tensor-product interpolants, while in ever other case (penumbra or umbra combined with penumbra, penumbra or umbra combined with unoccluded) the resulting mesh faces will be triangulated and a combined interpolant built.

By using the structured algorithm in [DrFi93] the radiance function in unoccluded regions for each source is split into regions in which the radiance is well behaved. The algorithm then creates quadratic interpolants and guarantees that the interpolants satisfies tight error bounds. Thus, the combined illumination function over the intersection of two light regions will continue to satisfy these error bounds. Similarly, for the other regions the combination of triangular or tensor-product interpolants is also guaranteed to give high quality results, since the regions have been segmented based on the complete discontinuity mesh.

Because of the guaranteed error bounds for the interpolants representing unoccluded illumination, we can safely use these approximations in our calculations for simplification (see below), instead of the more expensive direct illumination calculation.

4 Two-Pass "On Demand" Discontinuity Meshing

The main cost of the complete discontinuity meshing algorithm is due to the relatively large number of discontinuity surfaces that must be traced in the environment. In addition, it is necessary to search for the existence of discontinuity surfaces (either edge-vertex wedges (EV) or triple-edge quadric surfaces (EEE)), formed by edges and vertices not on the source. To reduce the cost of this computation, we must reduce the number of surfaces traced into the environment.

To do this we separate the mesh computation into two phases: first, the computation of the boundary between light and penumbra, and an estimate of the region between umbra and penumbra, and second the full computation of all discontinuity surfaces interior to the penumbra *only* when required. We call the boundary between penumbra and light the maximal boundary, and the boundary between umbra and penumbra the minimal boundary. The combined maximal and minimal boundary is called the *extremal* boundary.

4.1 Extremal Boundary Approximation

The computation of the maximal boundary can be performed exactly, since it is formed exclusively by EV surfaces [Camp91]. Thus these events can be identified in constant time for each object, and subsequently propagated into the environment. The minimal boundary can include EEE events [Tell92], which can be treated by the method described in [DrFi94].

As an example consider the scene shown in Fig. 2. On the left we see the full discontinuity mesh, and on the right the extremal boundary.

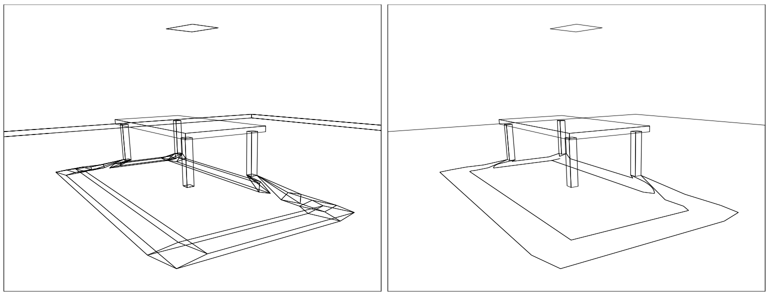

Fig. 2. Complete Mesh vs. Extremal Boundary

The number of discontinuity surfaces traced through the environment is thus reduced significantly. In addition, since no internal detail of the mesh is computed, all non-emitter events are ignored, and the search time for such events is eliminated.

The computation time for the extremal boundary is significantly reduced compared to the computation of the full mesh. In Table 1 we compare the cost of complete discontinuity meshing to the cost of the extremal boundary for the scenes shown in Fig. 1 and Fig. 2, as well as two other more complex scenes. As we can see, the cost of the complete mesh computation is three to four times higher than the just the extremal boundary. It is thus evident that large gains can be achieved if the complete mesh need be computed only when required.

Table 1. Computation time for Complete Mesh and Extremal Boundary

Scene	Polygons	Complete Mesh	Extremal Boundary	Ratio Complete/ Boundary
Box Scene	14	0.74 sec	0.16 sec	4.6
Table Scene	36	1.01 sec	0.31 sec	3.2
Desk Scene	182	17.20 sec	4.42 sec	3.8
Desk & Chair Scene	288	35.20 sec	9.20 sec	3.8

4.2 Local Complete Mesh Construction

As discussed earlier, one of the goals of the approach presented here is to compute portions of the discontinuity mesh only when necessary. The discontinuity meshing algorithm presented in ([Dret94, DrFi94]) is particularly well suited to such an extension.

Given a convex region defined on a receiver for which the complete mesh is desired, a convex volume defined by the source and that region is defined. Using the same spatial subdivision structure as in [DrFi94], the objects contained in this volume can be found efficiently.

To create the full mesh *locally* in the desired region of the receiver, the discontinuity meshing algorithm of [DrFi94] is applied using only the objects within the volume. In this manner, a much smaller number of discontinuity surfaces are traced (only those corresponding to edges and vertices of the selected objects), and the number and expense of searches for non-emitter events is also limited.

5 Simplification Criteria

In this section we discuss the simplification criteria used when two meshes are combined. In Section 6 we show how this applies to an arbitrary number of sources.

Consider a source S_1 and a source S_2. We assume that we have computed the extremal boundaries for the discontinuity mesh for each source, that the environment has been segmented into parallelogram regions of light and penumbra/umbra. We also assume that the structured sampling algorithm has been applied, subdividing the regions of light. In each such unoccluded region a biquadratic tensor product interpolant has been built, which represents the radiance function accurately within strict error bounds. We call each such mesh the simplified mesh for source S. In Fig. 3 we show the simplified mesh for each source for the scene of Fig. 7 (see colour section).

Given the two meshes, M_1 and M_2 respectively, we proceed to "add" M_2 into M_1. Merging is performed this way purely for reasons of algorithmic simplicity. There are three cases that must be treated:

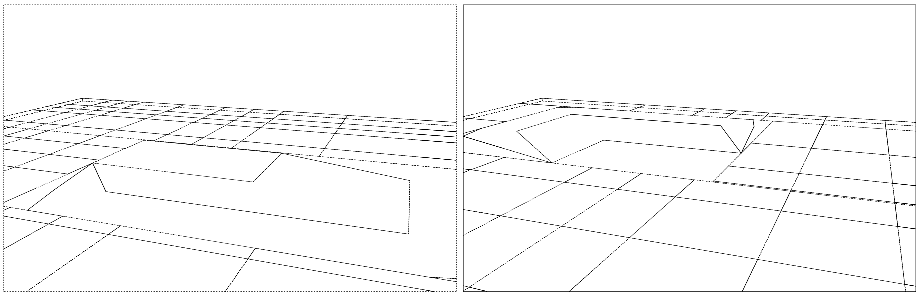

Fig. 3. Simplified Meshes for Box Scene

1. Merging light faces of M_1 with light faces of M_2.
2. Merging light faces of one mesh with penumbral faces of the other.
3. Merging penumbral faces of one mesh with penumbral faces of the other.

For the first case, since we have the structured representation in the form of tensor product interpolants for both meshes, we use an a posteriori error estimation to determine whether simplification can be performed. For the second case, we determine the regions of the penumbral group for which complete meshing is necessary. For the third case we currently perform no simplification.

5.1 Light-Light Simplification

The simplest case is the insertion of an unoccluded (light) face F_2 of M_2, into the mesh M_1. The mesh M_1 is searched to find all faces contained inside the boundary of the face F_2 being added. Call these faces $\{f_1, f_2, ..., f_n\}$. Within each such light face f_i of mesh M_1, a (structured) biquadratic interpolant $s_i^1(x, y)$ has been defined. Correspondingly, the structured interpolant in F_2 is $s_2^2(x, y)$.

To determine whether simplification is possible, we proceed to construct two biquadratic interpolants: first a high quality representation of the combined radiance with the region of F_2, denoted $s_\cap(x, y)$ and second a simplified representation, $\tilde{s}(x, y)$. The error incurred by the simpler interpolant (compared to the high-quality interpolant) is used to determine whether simplification can be achieved.

The high-quality interpolant $s_\cap(x, y)$ is defined as follows, in a piecewise fashion over each f_i (this is the interpolant created when combining the meshes as in Section 3):

$$s_\cap = s_2^2(x, y) + s_i^1(x, y), (x, y) \in f_i \tag{1}$$

Since the interpolants $s_i^j(x, y), j = 1, 2$ already constructed are good approximations of the actual radiance function, $s_\cap(x, y)$ is considered to be an accurate approximation of the combined function over the entire domain $F_2 = \cup f_i$.

The second interpolant $\tilde{s}(x, y)$ is defined over F_2 as a simple 9-point biquadratic tensor product, for which the midpoints are used as internal defining nodes. The nodal values are found by querying $s_\cap(x, y)$.

To determine whether the combined illumination from two sources can be represented accurately by the simplified interpolant $\tilde{s}(x, y)$, we use standard approximation theoretic error estimate [Pren89]. As a first approach we compute the L_2-norm of the difference of the simplified and the accurate interpolants.

For the L_2-norm the following quantity is computed:

$$L_2 = \sqrt{\int \int_{F_2} (\tilde{s}(x, y) - s_{\cap}(x, y))^2 dxdy} \quad (2)$$

This integral is computed in a piecewise fashion over each tensor product domain f_i. Since both $\tilde{s}(x, y)$ and $s_{\cap}(x, y)$ are quadratic functions, the integral of Eq.(2) can be computed analytically. In practice, the analytic expression is large and numerically unstable, so a two-dimensional Gauss-Legendre quadrature rule is used. In many cases, the quadrature can give exact results.

If the L_2-norm is less than a user-specified tolerance, the edges of the faces f_i are removed, and radiance in the domain of F_2 is represented by the simplified interpolant $\tilde{s}(x, y)$.

In Fig. 4, we show the result of the simplification criteria applied to a scene of two sources with no shadows. In Fig. 4(a) the original mesh is shown. From Fig. 4(b) it can be seen that T-vertices have been introduced into the mesh. To ensure C^0 continuity, T-vertices are treated as "slave-nodes". First all interpolants of simplified faces are constructed. For each T-vertex, the corresponding value of the neighbouring simplified interpolant replaces the previously assigned nodal value. In Fig. 8(a) (in the colour

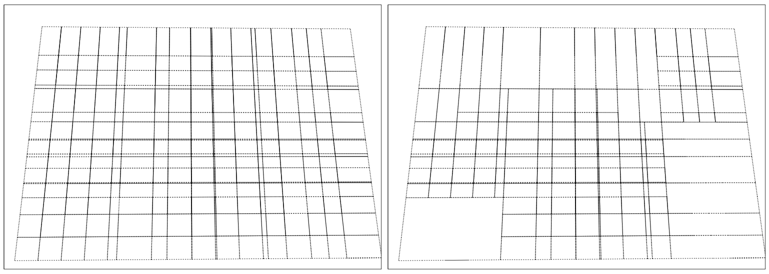

Fig. 4. (a) Original Unoccluded Mesh, and (b) Simplified Mesh

section) we show the image rendered using the original full mesh interpolant. In Fig. 8(b) the result of the construction of the continuous interpolants for the simplified mesh is shown. As can be seen, the resulting images show little difference. However, a more graded variation between simplified and unsimplified regions would be beneficial, using a form of restricted meshing.

5.2 Light-Penumbra Simplification

Consider a penumbral group of a mesh M_P caused by source S_P and a set of light faces of the mesh M_L caused by source S_L, which are contained or cut the penumbral group. We wish to add the light faces into the mesh M_P, and to determine the regions of the penumbral group for which the complete discontinuity mesh must be computed. In contrast to the light-light case, we do not have an accurate representation of the radiance in the penumbra.

To determine whether detailed mesh computation is required, we first construct a medium quality approximation $\hat{s}_\cap(x, y)$ to the radiance in the penumbra, using the extremal boundary, within each light face of M_L. This piecewise approximation takes into account the extremal boundaries of the various sources, and its use is equivalent to the accurate interpolant $s_\cap(x, y)$ for the light-light case. We then construct the simplified interpolant by defining a single biquadratic tensor product $\tilde{s}(x, y)$. The simplification criteria used are the same as in the light-light case.

The construction of $\hat{s}_\cap(x, y)$ proceeds as follows. We first construct an independent mesh defined by the bounding box of the penumbral group. We then add in the extremal boundary of the group of mesh M_P. We show this construction for the box scene and the penumbral group of one source in Fig. 5(a) (refer to Fig. 3(a) and Fig. 7 (colour section) to understand the geometry). In this way, a coarse segmentation of the penumbral group into regions of light, penumbra and umbra has been achieved.

For each vertex inserted into the independent mesh the appropriate illumination value due to source S_P is assigned. For the vertices on the maximal boundary or in the unnoccluded regions this is the direct unoccluded illumination from S_P and for the points on the minimal boundary the value is 0. We then insert all the light faces of M_L

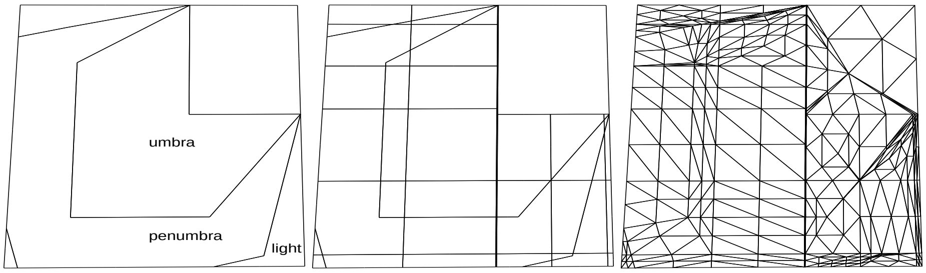

Fig. 5. Mesh for Error Testing: (a) Maximal/Minimal Boundary of penumbral group of M_P, (b) Light faces of M_L added, (c) Triangulation (domain of $\hat{s}_\cap(x, y)$)

that intersect or are contained in the penumbral group boundary (Fig. 5(b)). For the resulting vertices the value of unoccluded illumination is retrieved from the appropriate interpolants of M_L, but it is then necessary to add the appropriate (penumbral) value due to the source S_P. For regions of umbra and light (due to source S_P) this can be found simply. For vertices in regions within the penumbra however it is necessary to retrieve an estimate of the radiance value. This can be achieved by estimating the derivative value of radiance (see below).

The resulting combined mesh is then triangulated (Fig. 5(c)), and the piecewise elements of the interpolant $\hat{s}_\cap(x, y)$ are built. Interior nodal values are computed either directly (if in a region of light or umbra) from the appropriate interpolants in M_L and

M_P, or are averages of the neighbouring nodes if the node is within the penumbra.

For each region corresponding to a light face F_L, the interpolant $\tilde{s}(x, y)$ is constructed. This interpolant is a simple 9-point bi-quadratic Lagrange interpolant. The values for nodes corresponding to vertices in the combined mesh have already been assigned and those that remain are found by querying the interpolant $\hat{s}_\cap(x, y)$. We then compute the L_2-norm error in the same manner as for the unoccluded case for the triangles of $\tilde{s}(x, y)$ which lie in umbra or penumbra. The integral is computed over each triangle included in the domain of the light face F_L. If the L_2 error is less than the predefined tolerance, the edges of F_L are inserted into M_P, the extremal boundary edges contained in F_L are removed from the mesh M_P and radiance within this region is represented by the simplified interpolant.

If the error is greater than the user specified tolerance, the region of the original light face is marked as requiring further meshing. After processing all light faces, the complete mesh is locally computed only for the regions required.

5.3 First Implementation and Discussion

To verify the algorithm, we have implemented the light-penumbra simplification by first computing the complete mesh, and then simplifying the mesh where appropriate. The full construction of the extremal boundary and the simplification algorithm have been implemented as described above, with the exception of the local backprojection estimate. Instead, for the light-face vertices within penumbra, the exact penumbral radiance is retrieved from the (complete) mesh of source S_P.

As mentioned above, for the penumbral regions only the portions of the simplified mesh in penumbra or umbra are taken into consideration for the L_2-norm computation. As noted in [Dret94], edges leading to a singular vertex display a particularly rapid variation. To correctly account for this, in faces for which singular edges exist the light faces are also considered in the L_2-norm calculation.

The results of the implementation are shown in Fig. 6. We first show the unsimplified combined mesh (a), and then the simplified mesh for tolerance values 0.005 and 0.001 respectively (b) and (c). The corresponding shaded images are shown in Fig.

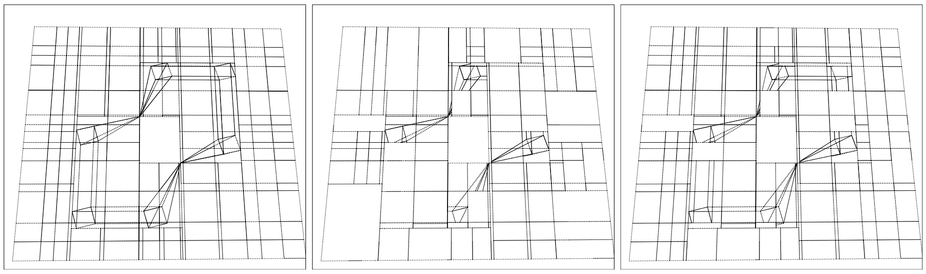

Fig. 6. (a) Unsimplified Combined Mesh and Simplified Mesh for (b) Tolerance 0.005 and (c) 0.001

9(a),(b) in the colour section. The results of a more complicated test are shown in Fig.

?? to **??**. The complete mesh and resulting image are shown in Fig. **??**, and the reduced meshes and images in Fig. **??** and **??** for tolerances equal to 0.1 and 0.005 respectively.

Overall the method shows promising first results. Little difference can be seen in the simplified images compared to complete mesh image for the simple scene (Fig. 7 (b)), and the simplification appears to occur in desirable regions of the mesh as the tolerance grows. Similarly the simplified images for the table scene (Fig. **??**, **??**) appear to maintain relatively high quality, since simplification occurs in the regions in which the detail of the penumbra is not very important.

In the tests performed it can be seen that the use of the L_2-norm can sometimes cause undesirable simplification (e.g., the shadow boundary of the front leg in Fig. **??**(a)). A possible solution is to maintain the extremal boundary instead of substituting with a tensor product.

5.4 Penumbral Radiance Estimates

Given the maximal and minimal boundary we propose here an estimate of the radiance at any point in the penumbra using local backprojection information. By construction, the minimal or maximal edges of the discontinuity mesh include information about the local change of the backprojection. Thus a good estimate of the radiance at a point P, known to be in penumbra, can be found by approximation.

To perform this approximation we first find the edge on the minimal boundary for which the two endpoints are closest to P. We then calculate the backprojection into the penumbra locally in a direction defined by the midpoint of the minimal edge and the point P. Given the backprojection, we estimate the radiance derivative, then build a Hermite cubic from the values and the derivative estimates, and determine the radiance value at P using the cubic. Experimental verification will determine the quality of this approach.

6 Treating Multiple Sources

The simplification algorithm begins by computing the extremal boundary for each of the n sources in the scene. The light regions are computed, and the structured algorithm run for each surface. The result is a list of simplified meshes for each surface: $\{M_1, ..M_n\}$. The algorithm proceeds by merging the first two meshes. The combined mesh M_c is then merged with mesh M_3 etc.

For a pair $\{M_c, M_j\}$ we first insert the light faces of M_j into the mesh . If a light face of M_j contains exclusively light faces of M_c, or there is a parallelogram subregion of M_j with this property, the light-light simplification is applied. The penumbral regions of both meshes are then visited, and the simplification algorithm is run for each penumbral group. A list of regions marked as "potentially requiring meshing" is stored, together with a pointer to the appropriate source. In addition, the interpolant $\hat{s}_\cap(x, y)$ is stored and used in subsequent tests for error bound checking. If a subsequent source eliminates the need for the meshing, the corresponding regions are deleted from the list. At the end of this process, there will be a list of regions for which the complete mesh is applied.

7 Conclusions

In this paper we have presented an algorithm which allows more compact representation of radiance due to multiple emitters based on careful error analysis, and allows the cost of discontinuity meshing to be deferred until it is required.

To achieve this goal, the L_2-norm is used to compare an accurate representation of radiance over a domain with a simpler one. When the simpler interpolant satisfies a given error tolerance, it is used. For regions with unobstructed views of all sources, this is performed as an a posteriori step. For regions in penumbra for one source and light for another, a low-quality discontinuity mesh is first computed, and an approximation to radiance built, which is then compared to a simpler interpolant. Results of a first implementation show promising reduction of the mesh, and good quality images when using the simplified interpolant.

For the future, it is extremely interesting to apply these ideas to complex environments with many sources, to determine the savings, both in the representation of unoccluded regions, but more importantly in the computation time for discontinuity meshing. The subsequent step is the usage of these algorithms in a global illumination context, since for secondary reflection the need for complete meshing is highly unlikely.

8 Acknowledgments

The author is a post-doctoral fellow hosted by INRIA, under an ERCIM fellowship. Many of the ideas presented originate in the authors' Ph.D. thesis [Dret94] at the University of Toronto, under the supervision of Prof. Eugene Fiume. The software system used for the implementation includes many software components written by researchers at the Dynamic Graphics Project at Toronto.

References

[Camp91] Campbell, A. T. III, Fussell, D. S.: An Analytic Approach to Illumination with Area Light Sources, Technical Report TR-91-25, Computer Science Dept., University of Texas at Austin, August 1991.

[ChFe92] Chin N., Feiner S.: Fact Object Precision Shadow Generation for Area Light Source using BSP Trees, ACM Computer Graphics (SIGGRAPH Symp. on Interactive 3D Graphics), March 1992.

[CoGr85] Michael F., Greenberg D. P. : The Hemi-cube, A Radiosity Solution For Complex Environments, ACM Computer Graphics (SIGGRAPH '85 Proceedings), Vol. 19, No. 3, July 1985.

[DrFi93] Drettakis G., Fiume E. L. : Accurate and Consistent Reconstruction of Illumination Functions Using Structured Sampling, Proceedings of the Eurographics Conference, Barcelona, Spain, September 1993.

[DrFi94] Drettakis, G., Fiume E. L. : A Fast and Accurate Shadow Algorithm for Area Light Sources Using Backprojection ACM SIGGRAPH Computer Graphics, Annual Conference Series, July 1994.

[Dret94] Drettakis G.: Structured Sampling and Reconstruction of Illumination for Image Synthesis Ph.D. Thesis, Dept. of Computer Science, University of Toronto, also available as CSRI Tech. Report-293 (ftp site ftp.csri.toronto.edu:csri-technical-reports/293), January 1994.

[GSCH93] Gortler S. J., Schroeder P., Cohen M. F., Hanrahan P.: Wavelet Radiosity, ACM SIGGRAPH Computer Graphics, Annual Conference Series, August 1993.

[HaSA91] Hanarahan P., Salzman D., Aupperle L.: A Rapid Hierarchical Radiosity Algorithm, ACM Computer Graphics (SIGGRAPH '91 Proceedings).

[Heck92] Heckbert P.: Discontinuity Meshing for Radiosity, 3rd Eurographics Workshop on Rendering, Bristol, UK May 1992.

[HeWi91] Heckbert P., Winget, J. M.: Finite Element Methods for Global Illumination, EECS/University of California Berkeley, Report No. UCB/CSD 91/643, July 1991.

[LiTG92] Lischinski D., Tampieri F., Greenberg D. P.: Discontinuity Meshing for Accurate Radiosity, IEEE Computer Graphics and Applications, November 1992.

[NiNa83] Nishita, T., Nakamae, E. : Half-tone Representation of 3-D Objects Illuminated by Area Sources or Polyhedron Sources, Compsac, Proc. IEEE 7th Intl. Comp. Soft and Applications Conference, 237–242, November 1983.

[Pren89] Prenter, P.M.: Splines and Variational Methods, John Wiley and Sons, New York, 1989.

[SaLD92] Salesin D., Lischinski D., DeRose T. : Reconstructing Illumination Functions with Selected Discontinuities, 3rd Eurographics Workshop on Rendering, Bristol, UK May 1992.

[Tell92] Teller S.: Computing the Antipenumbra of an Area Light Source ACM Computer Graphics (SIGGRAPH 92 Proceedings), July 1992.

[Zatz93] Zatz H. R.: Galerkin Radiosity: A Higher Order Solution Method for Global Illumination, ACM Computer Graphics (SIGGRAPH '93 Proceedings), August 1993.

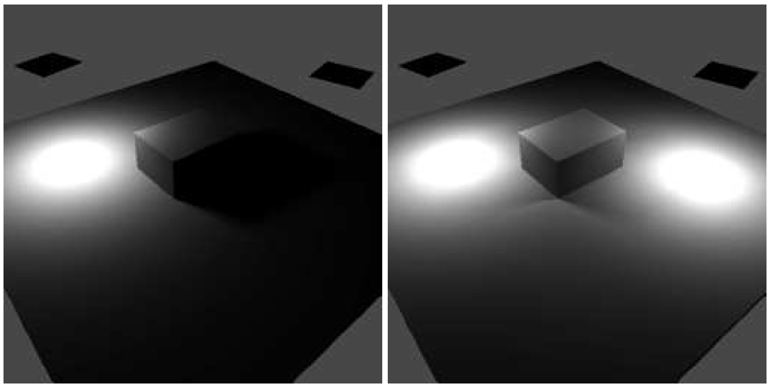

Fig. 7. (a) One Source and (b) Two Source Images

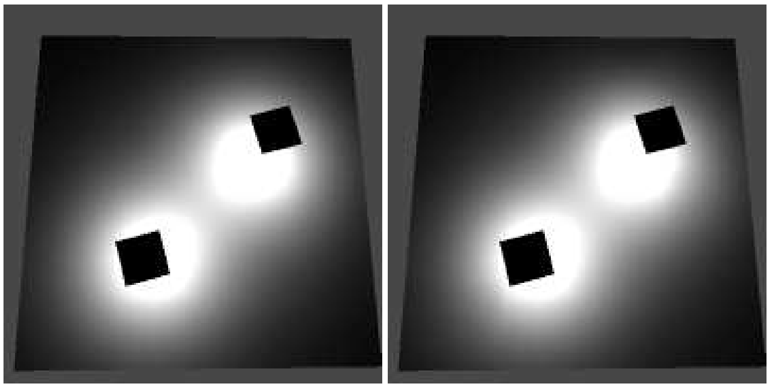

Fig. 8. Images for (a) Original and (b) Simplified Unoccluded Meshes

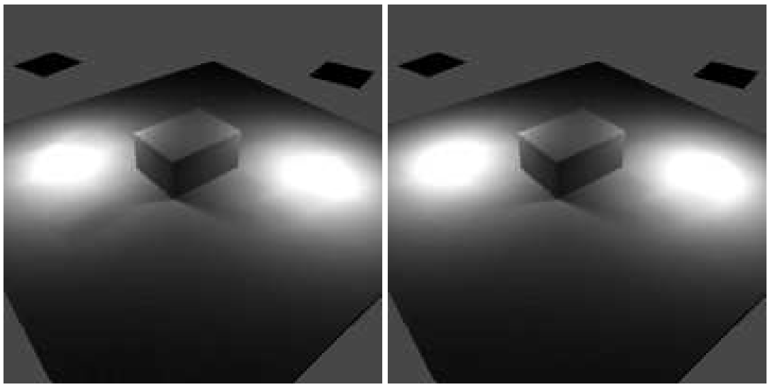

Fig. 9. Images of Simplified Meshes (a) Tolerance = 0.005, (b) Tolerance = 0.001