

Accurate Visibility and Meshing Calculations for Hierarchical Radiosity

George Drettakis, François X. Sillion

► **To cite this version:**

George Drettakis, François X. Sillion. Accurate Visibility and Meshing Calculations for Hierarchical Radiosity. Proceedings of 7th Eurographics Workshop on Rendering (Rendering Techniques '96), 1996, Porto, Portugal. pp.269–279. inria-00510111

HAL Id: inria-00510111

<https://hal.inria.fr/inria-00510111>

Submitted on 17 Aug 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Accurate Visibility and Meshing Calculations for Hierarchical Radiosity

George Drettakis, François Sillion

iMAGIS * Laboratoire GRAVIR/IMAG-INRIA

Abstract: Precise quality control for hierarchical lighting simulations is still a hard problem, due in part to the difficulty of analysing the source of error and to the close interactions between different components of the algorithm. In this paper we attempt to address this issue by examining two of the most central components of these algorithms: *visibility* computation and the *mesh*. We first present an investigation tool in the form of a new hierarchical algorithm: this algorithmic extension encapsulates exact visibility information with respect to the light source in the form of the *backprojection* data structure, and allows the use of *discontinuity meshes* in the solution hierarchy. This tool permits us to study separately the effects of visibility and meshing error on image quality, computational expense as well as solution convergence. Initial experimental results are presented by comparing standard quadtree-based hierarchical radiosity with point-sampling visibility to the approaches incorporating backprojections, discontinuity meshes or both.

1 Introduction

Hierarchical simulation techniques have received a lot of attention in research environments, but their practical use remains impaired by the difficulty of controlling the speed/accuracy tradeoff on which they are based. Error control and solution accuracy issues have been studied to a certain extent for global illumination algorithms [12, 1]. These studies provided a useful categorization of possible error sources, and offered a general framework for error-driven hierarchical refinement. Nonetheless, little has been done in terms of investigating the different causes of error in Hierarchical Radiosity (HR) in particular, and very little is currently known about the quantitative effects on error of different algorithmic choices used during the lighting simulation.

In this paper we attempt to address this problem by providing recommendations, based on theoretical discussion and initial experimental results. We will concentrate our efforts on meshing and visibility computation strategies. We begin by presenting a non-exhaustive list of important algorithmic components in HR and we mention the algorithms that have been proposed to improve these aspects of the simulation. For most of these factors, the precise impact on image or solution quality, as well as possible interactions between them, has not been thoroughly studied.

Important components of the HR simulation algorithm Broadly speaking, two main categories of factors affecting simulation can be identified (following [1]): *discretisation*, concerning mainly issues of mesh construction and data structures, and *computation* which involves the aspects of the algorithm related to form-factor and visibility computation as well as refinement strategy and convergence.

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

Discretisation

Meshing strategy HR relies on the ability to evaluate interactions (energy transfers) at different levels of a hierarchy in the description of the scene. Previous algorithms typically used simple recursive subdivision structures such as the quadtree to represent hierarchical meshes. Another approach consists of computing a *discontinuity mesh* (DM) for much improved representation of direct illumination.

Mixed meshes Simple hierarchical structures such as quadtrees are easy to implement and compact to store, because they rely on implicit information. However they are not suited to some geometrical or topological situations such as the representation of shadow boundaries. Mixed meshes with both triangles and quadrilaterals can therefore be used, and must provide access to connectivity information.

Computation

Visibility calculation Many radiosity implementations to date use point sampling to evaluate visibility factors. Discontinuity meshes, when equipped with the associated visibility information (backprojection) can provide exact visibility computation for direct illumination.

Refinement Strategy Refinement criteria (sometimes called “oracles”) are the core of the HR formulation. Many different criteria have been devised, using varying amounts of information. Possible variables for the refinement decisions are form factor estimates, visibility status, estimate of form factor variance, estimate of radiosity transfer, etc.

Point or Area-based form-factor computation As shown by the work of Wallace *et al.* [20] for progressive refinement radiosity, higher-quality solutions can be obtained when computing radiosity directly at mesh vertices. Area-to-area form factors require an extrapolation/interpolation step which effectively smoothes out some of the defects in the solution but also “blurs” the computed solution in ways which are difficult to quantify.

Convergence The benefits of HR really become apparent when a global solution is sought, i.e. with all interreflection effects. HR convergence is an issue that has received limited attention and it is not known whether some of the choices mentioned above have a significant impact on convergence.

2 Previous Work

2.1 Discontinuity meshing and backprojections

Discontinuity meshing ([9, 12, 5]) has been used to a certain extent for global illumination and HR calculations, but the meshes used have always been *partial* since they do not capture EEE and other important discontinuity surfaces. A result of this simplification is that *backprojections* (defined in e.g. [3, 16, 18]) cannot be computed by these algorithms. Backprojections are data structures permitting efficient determination of the visible part of the source anywhere in the penumbra, thus effectively eliminating all visibility calculation error with respect to the light sources. Backprojections can *only* be computed together with the complete discontinuity mesh (i.e. including EEE and degenerate events), and thus have never been used so far in the context of HR.

2.2 Combination of HR and DM

The most relevant approach to our work is that of Lischinski et al. [12]. In their work a global solution is computed using a BSP tree. When a node is split, an appropriate discontinuity line is chosen. A local pass is subsequently used for display, during which analytic visibility (mainly for primary sources) is computed using an expensive polygon clipping operation (as in [10, 17]). The main difference with the method we present here is the fact that exact visibility was not taken into account during the solution process and thus the different factors (visibility, meshing) affecting error could not be isolated or analysed. Gatenby and Hewitt [5] also developed a hierarchical solution for progressive refinement, but little was presented in terms of solution quality evaluation.

2.3 Error estimation and control

A detailed theoretical presentation of error analysis was developed by Arvo et al [1]. According to their classification, we will be dealing with *discretisation error* (meshing) and *computational error* (visibility calculations). For both types of error, little is known in practice or in quantitative terms. Lischinski et al. [11] have also developed an approach based on error bounds for HR. The approach we present here could be integrated into a system of this type, providing tighter upper and lower bounds in the most difficult cases, those of partial visibility.

3 Efficient combination of HR and backprojections

The first step to allow experimental comparisons of the effect of meshing and visibility calculations on the solution, is the introduction of a new algorithm incorporating backprojections (i.e. exact visibility calculations) and the complete discontinuity mesh in HR. Other than backprojections, this new algorithm constructs a full hierarchy on the input surfaces by clustering the elements of the discontinuity mesh. This allows the use of quadtrees in unoccluded regions as opposed to the use of BSP trees everywhere as in previous partial DM solutions ([12]), while using irregular triangles in shadow regions.

3.1 Accurate visibility computation using backprojections

Given a three-dimensional scene we construct the full discontinuity mesh (with EEE and degenerate events) which has the following property: each cell or face of the mesh contains a data structure called a backprojection which fully describes the partial visibility of the source at each point within the mesh face [3, 16]. An example of such a mesh is shown in Fig. 1(a).

The availability of the backprojections allows us to determine at a low cost the visible part of the source and thus the analytical value of the differential form-factor $dF_{x,S}$, from any point x in the penumbra to the source S . Because no visibility calculation is needed, an accurate estimate of the form factor value is available cheaply during refinement. As we will see below, the exact value is also used to determine irradiance values at vertices within the mesh.

3.2 Mixed triangle-quadrilateral meshes

The discontinuity meshing approach presented in [3, 4] constructs a mixed, non-hierarchical mesh containing triangles in penumbra and in irregular regions of light, and quadrilaterals in large regions of light. The goal here is to create a hierarchy suitable for HR solutions, starting with the DM.

To this end, quadtrees should be used where possible (in particular in unoccluded regions) due to their simplicity and ease of use, while irregular triangular meshes should be used around shadow boundaries.

The mixed mesh structure poses certain problems of connectivity, since inhomogeneous mesh elements co-exist. In particular, neighbour-finding is handled by adding simple adjacency information at the vertices of the original mesh. When a quadrilateral or a triangle are subsequently subdivided, regular quadtrees or regular triangle hierarchies are created. An example of an initial mesh (before subdivision) for our test scene is shown in Figure 1(b) (see next section for the construction algorithm); the same mesh subdivided after iterating is shown in Figure 4(b). When searching for a neighbour within a regular mesh (triangular or quadrilateral) implicit neighbourhood relationships are maintained, and when crossing a shared edge, the neighbourhood information stored at the vertex is used to find the appropriate quadtree or triangle mesh. We climb up the hierarchy until the parent maintaining the appropriate information is found.

3.3 Constructing a true hierarchy from the discontinuity mesh

For each receiver containing a penumbral or umbral zone, after the discontinuity meshing and triangulation steps, we have a set of triangles corresponding to this partially lit or occluded region. To construct a hierarchy we attach these triangles at appropriate levels of a standard quadtree, such that in unoccluded regions illumination is represented with the regular quadtree structure.

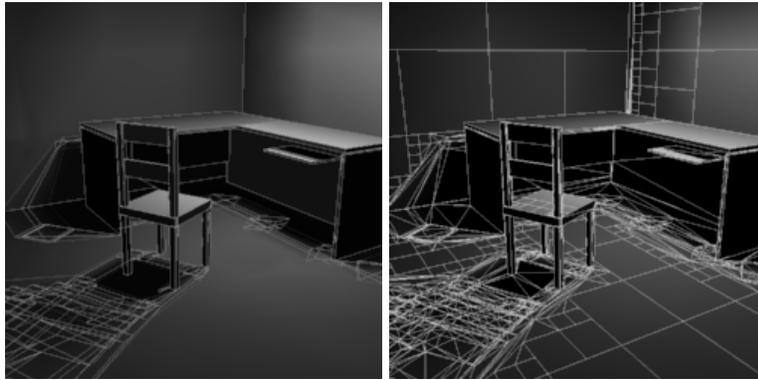


Fig. 1. (a) Discontinuity mesh and (b) HR/DM mesh for a test scene

Quadtree Subdivision We start by recursively subdividing the receiver using a standard quadtree. If a child of the quadtree contains no partially lit or occluded region, initial subdivision (i.e. the subdivision performed before BF-refinement) terminates. The unoccluded quadtree leaf elements are also inserted into a temporary face-edge-vertex data structure mesh. If on the other hand a child contains part of the penumbra or umbra, subdivision continues until a predefined maximal depth is reached.

When no more subdivision is possible, the penumbral/lit boundary edges are inserted into mesh. When this process is complete, mesh contains a set of faces corresponding

to the completely unoccluded quadtree leaves, and a (usually highly irregular) face between the leaves and the penumbra/lit boundary. This face is triangulated, and the resulting triangles are then added to the list of penumbral/umbral triangles.

The final step requires “clustering” of these triangles (both original penumbra triangles from DM as well as the new triangles in lit regions) from mesh so that they can be correctly attached to an appropriate level of the quadtree.

“Clustering” for Penumbral and boundary regions To perform the clustering step, a 3D clustering bottom-up construction ([15]) is adapted to 2D. A multi-level grid is constructed, such that the smallest grid cell has the size of a maximal depth quadtree leaf. Each level of the grid is visited, and triangles entirely contained in a cell at a given level of the grid are attached to the corresponding quadtree inner node (if it exists). The triangles contained in a given cell at a given level are grouped to form an internal node. This node is then inserted at the appropriate level higher in the multi-level grid, if it is contained in a cell at that level. The contents of this grid cell will in turn be attached to the appropriate level of the quadtree.

In this manner we have a mixed hierarchy which starts at the root as a normal quadtree, and has children which may be regular quadtree subdivisions, or agglomerations of triangles or individual triangles (in the case of elongated triangles which can occur in the context of discontinuity meshing). An example is shown in Figure 1(b).

It must be noted that due to the bottom-up construction, we have the ability to insert the *entire* DM into the hierarchy (as is done here). In many cases simplification should probably be performed, but the generality of the method permits maximal flexibility.

4 Impact of algorithmic choices on solution and image quality

In this section we revisit the different algorithmic components of HR mentioned in the introduction, and discuss their relevance. This discussion serves both as a first attempt to investigate the influence these factors have on the solution as well as on each other, and as motivation for the experimental approach developed in the following section.

4.1 Meshing strategy

The use of quadtrees has many advantages: the structure is simple to handle and manipulate, it allows implicit neighbour finding operations, and provides well shaped elements which is important in the context of any numerical approximation ([13]). Interpolation and extrapolation operations are also readily performed in these structures since elements respect regular ratios. Nonetheless, the very regularity of the quadtree structure hides its inadequacy in representing high-frequency irregular information such as shadow boundaries.

Discontinuity meshing provides an appropriate solution to the problems of visual representation of shadow boundaries. The problems with such meshes are however numerous. Other than the issues related to numerical accuracy in construction [19], these meshes tend to contain far too many elements ([4]), and they result in badly formed triangles which pose problems for interpolation/extrapolation operations as well as being formally unadaptable to finite element approaches ([13]). It is difficult to determine a priori when the use of such meshes is advisable in the context of HR.

4.2 Visibility calculation

In traditional HR approaches, visibility is computed by ray-casting between two patches p and q . A form-factor disc approximation is then multiplied by the fraction of rays

blocked, which is used as a visibility estimate. This approximation influences the form-factor estimation as well as refinement, and has repercussions which are difficult to isolate.

Given the mesh and backprojections, two important changes can be incorporated into the treatment of the direct illumination links: characterisation of links as partial, occluded or unoccluded (in the spirit of [19]) can be performed accurately immediately after the discontinuity meshing step and the calculation of irradiance values at the vertices during the solution is exact. The estimate of area-to-area form-factors can also be significantly improved since each sample of the kernel function is calculated with the exact visible portion of the source.

4.3 Direct Illumination at Vertices for HR

In standard HR (using a piecewise constant approximation), irradiance is "pushed" down the hierarchy to the leaf nodes ([6, 14]). This irradiance is then converted to radiosity which is subsequently extrapolated to the vertices of the leaves and "pulled" up the hierarchy. If the backprojections are available, we can compute exact irradiance values due to light sources at all vertices very cheaply, since no visibility computation is required (points are either in a penumbral (or umbral) mesh face or in light). It is thus only natural to skip the "push" step for the light source, and simply evaluate the exact irradiance at the vertices of the mesh. These values are then averaged, resulting in a radiosity value assigned to the leaf, and then "pulled" in the normal manner up the hierarchy. We note that this direct shading should only be performed at the vertices originally in the discontinuity mesh, or for vertices on hierarchy leaves with a link to the source (in the current implementation it is performed at all vertices).

This approach has a double advantage of producing visually accurate results (see Figure 5(iv) in Colour Section) while simultaneously providing a highly accurate, hierarchical representation of direct illumination, which will, hopefully, result in an overall higher quality global illumination simulation.

4.4 Refinement criteria

In traditional HR approaches, "BF" refinement has been used, which essentially requires a link between two surface elements to be refined if their mutual form-factor multiplied by the power on the link is larger than a threshold [8]. The philosophy of this approach is to reduce the respective size of the elements on the two sides of a link, thus reducing visibility error since all interactions tend to be either occluded or visible ([19]), and reducing "integration error" of the area form-factor integral since the kernel varies. Another approach involves a "smoothness" criterion for the kernel used in Wavelet-based HR ([7]).

When exact visibility is available, neither of these criteria is entirely satisfactory. Since we can cheaply determine the visible part of the source, we could apply a (point-to-area) form-factor variation criterion (this is similar in spirit to the "smoothness" criterion). Subdivision will be thus better adapted to the variation of irradiance on the receiver, since the source is never subdivided.

4.5 Convergence

The issue of convergence is rarely addressed in HR research. Following the original definition and structure of the algorithm [8], we define convergence as a sequence of iterations involving a Refine step followed by an (optional) loop of Gather-Push/Pull operations until the result is no longer modified. The iteration terminates when no new links are created ("convergence").

One of the issues we wish to investigate is the effect of meshing strategy and visibility on the rate of convergence. Intuitively it seems that a good representation of direct illumination and accurate visibility computations should improve the convergence rate.

5 Comparison of some strategies and initial experimental results

The new algorithm presented in Section 3 provides an environment that allows a first investigation of the relative effect of meshing and visibility error on image and solution quality.

5.1 Experimental configurations

Four different algorithmic configurations are considered: (i) HR, using regular quadtree subdivision and ray-cast visibility calculation (QT/RT), (ii) HR regular quadtree subdivision and using backprojections (QT/BP), (iii) HR with full discontinuity meshing and ray-cast visibility (DM/RT) and finally (iv) HR with full discontinuity meshing and backprojections (DM/BP).

Comparing (i) and (ii) quantifies the effect of visibility error in form-factor computation, and the resulting effect on the solution. Comparing (i) and (iii) demonstrates the importance of the use of the discontinuity mesh as a basis for the subdivision in HR. Finally, the combined effect of the mesh and accurate visibility becomes evident by comparing configuration (iv) to the others.

Two test environments “Desk+Chair” (Fig. 1) and “Books” (Figure 5 in the colour section) with specific points of view have been chosen. The scenes are lit by large light sources giving rise to large regions of penumbra (which favours the use of backprojections). In addition, the “Books” scene contains many regions of small fine shadow, for which discontinuity meshing is advantageous. There are 133 polygons (268 distinct edges) in “Desk+Chair” and 241 polygons (484 distinct edges) in “Books”.

A reference solution is computed for both scenes, using a standard quadtree subdivided very finely, and run to convergence with a very small tolerance value. We compute an L_1 error on the pixel RGB values.

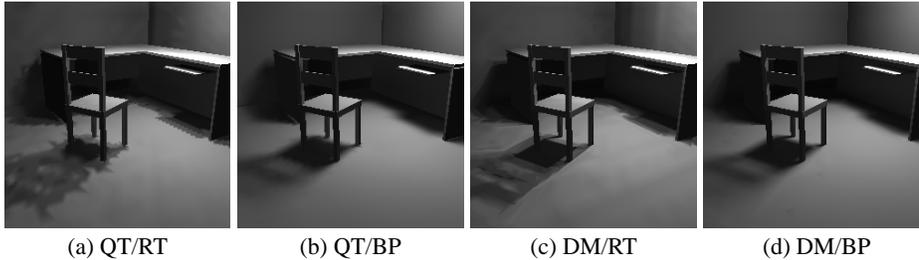


Fig. 2. “Desk+Chair” images

5.2 Quality evaluation and test suites

For each run we present the total computation time t_t to convergence, the final number of leaves at convergence l_c and the L_1 image error e after convergence for the given tolerance. We also report (where applicable) DM construction time t_{dm} ((ii)-(iv)) triangle

clustering time t_c ((iii) and (iv)) and the number l_{dm} of leaf elements after the mixed hierarchy construction but before subdivision ((iii) and (iv)). All reported timings are on an SGI R4400 Indigo 2 at 150 Mhz.

For the two scenes, we have attempted to maintain approximately the same number of elements, to provide a “fair” comparison. This requires the judicious choice of parameters $BF-\epsilon$, minimum area size and the visibility factor as defined in [8]. The resulting images for “Books” are shown in colour in Figure 5, while small versions of “Desk+Chair” are shown in Figure 2. Meshes for “Desks+Chair” are shown for illustration in Fig. 4 for cases QT/BP and DM/BP. The numerical results are summarised in Table 1.

Solution	t_t (s)	t_{dm} (s)	t_c (s)	l_{dm}	l_c	e		t_t (s)	t_{dm} (s)	t_c (s)	l_{dm}	l_c	e
(i) QT/RT	413.6	-	-	-	3911	6.7		838.1	-	-	-	8759	9.2
(ii) QT/BP	531.3	71.4	-	-	3650	2.8		1076.6	394.5	-	-	8168	4.2
(iii) DM/RT	604.9	71.4	98.5	3365	4502	6.0		1332.8	394.5	95.0	5929	8320	8.2
(iv) DM/BP	310.1	71.4	98.5	3244	3982	3.0		782.7	394.5	95.0	5775	6969	3.7

Table 1. Test results: (left) “Desk+Chair” and (right) “Books”

5.3 Results of Experimental Study

The test results presented above are by no means definitive or complete. The nature of experimental work is such that it is difficult to come to concrete conclusions from a set of given tests. Nonetheless, we believe that the results presented provide interesting insight into the problems related to visibility and meshing in the context of HR.

For both scenes we see that DM/BP provides the most computationally efficient solution, despite the overhead of mesh creation and hierarchy construction. The image quality is always better (Fig. 2, and Colour Fig. 5), and numerical accuracy is either better or on a par with all available alternatives. In the case of fine shadow features, the discontinuity mesh is particularly advantageous.

Visibility Visibility accuracy is of predominant importance. Solution (ii) QT/BP is numerically the most accurate for “Desk+Chair”, but DM/BP is visually superior (see Fig. 2, 5). DM/BP is however more accurate numerically for “Books”. The use of backprojections enhances numerical and visual quality more than the use of DM alone. The visual quality of QT/BP can be very high, as is the case for “Desk+Chair”.

Convergence In Fig. 3 we compare the L_1 image error at each iteration for the two scenes. The accuracy of the visibility computation appears to directly influence convergence. Solution (ii) has the best behaviour for “Desk+Chair”, but the DM is more important for “Books”, for which the DM/BP solution has the lowest error. In the case where QT/BP is numerically better, the difference is insignificant.

Meshing Discontinuity meshing without exact visibility results in visual artifacts, and is computationally expensive. This is particularly evident for “Desk+Chair” (Fig. 2). Previous algorithms avoided this problem by using “final gather” type approaches [12]

Nonetheless, the irregular meshes produced add a high overhead in the global solution, simply by the sheer number of leaf elements at the outset (Table 1) (a similar observation was made by Lischinski et al. [12]). An obvious remedy is to investigate the use of simplification techniques for meshing (e.g., [2, 4]), while maintaining the original backprojection information for visibility computations.

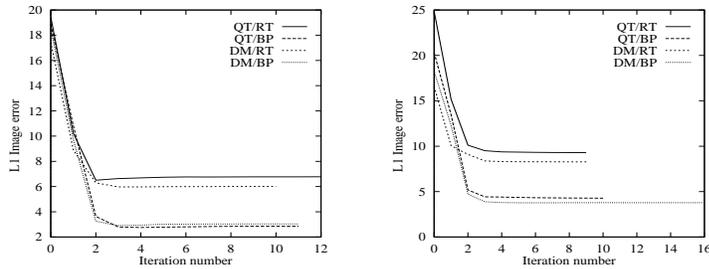


Fig. 3. Convergence for different approaches (a) “Desk+Chair” (b) “Books”

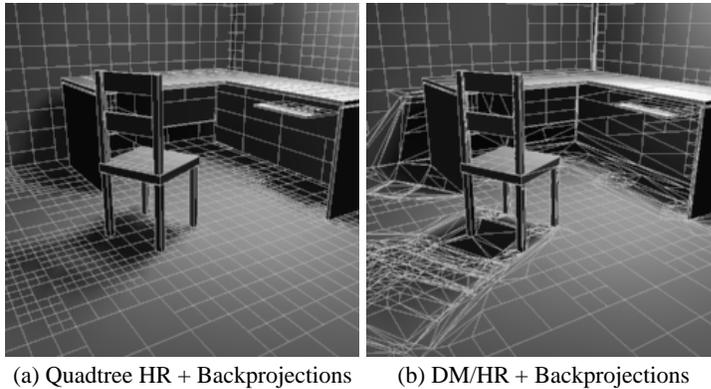


Fig. 4. Test suite mesh images

6 Conclusions

We have presented a first approach to investigating sources of error due to visibility computation and meshing strategies. A list of important factors affecting HR computations was presented and discussed. To facilitate experimental investigation we introduced a new hierarchical radiosity algorithm which incorporates backprojections (and thus exact visibility with respect to the source) and discontinuity meshes.

This approach has permitted the comparison of standard quadtree-based HR using traditional point-sampling for visibility with the case where visibility is calculated with backprojections, HR discontinuity mesh with point sampling and finally HR discontinuity meshing with backprojections.

A number of interesting observations were made from an experimental study relating the different effects of the use of analytic visibility (backprojections) and discontinuity meshes for HR light transport. Overall, it was observed that visibility accuracy is much more important than the use of meshing. Nonetheless, DM with BPs adds to overall visual quality. Many more experimental tests are required to confirm the observations made here as well as to investigate other aspects of the solution process.

Numerical difficulties and robustness problems are inherent in all discontinuity meshing approaches. A comprehensive solution to this problem is being pursued for scenes of moderate complexity. A algorithm to simplify discontinuity meshes for the mixed hierarchy is also currently being investigated.

Acknowledgements The first author thanks Xavier Pueyo for initial discussions on the subject.

References

1. James Arvo, Kenneth Torrance, and Brian Smits. A framework for the analysis of error in global illumination algorithms. In *Computer Graphics Proceedings, Annual Conference Series: SIGGRAPH '94* (Orlando, FL), pages 75–84. ACM, July 1994.
2. George Drettakis. Simplifying the representation of radiance from multiple emitters. In *Proceedings of 5th EG Workshop on Rendering*, Darmstadt, Germany, June 1994.
3. George Drettakis and Eugene Fiume. A fast shadow algorithm for area light sources using back projection. In *Computer Graphics Proceedings, Annual Conference Series: SIGGRAPH '94* (Orlando, FL), pages 223–230. ACM SIGGRAPH, New York, July 1994.
4. George Drettakis and Eugene Fiume. Structured penumbral irradiance computation. 1996. Submitted for publication.
5. Neil Gatenby and W. T. Hewitt. Optimizing discontinuity meshing radiosity. In *Fifth Eurographics Workshop on Rendering*, pages 249–258, Darmstadt, Germany, June 1994.
6. Reid Gershbein, Peter Schröder, and Pat Hanrahan. Textures and radiosity: Controlling emission and reflection with texture maps. In *Computer Graphics Proceedings, Annual Conference Series: SIGGRAPH '94* (Orlando, FL), pages 51–58. ACM, July 1994.
7. Steven J. Gortler, Peter Schröder, Michael F. Cohen, and Pat Hanrahan. Wavelet radiosity. In *Computer Graphics Proceedings, Annual Conference Series: SIGGRAPH '93* (Anaheim, CA, USA), pages 221–230. ACM SIGGRAPH, New York, August 1993.
8. Pat Hanrahan, David Saltzman, and Larry Aupperle. A rapid hierarchical radiosity algorithm. *Computer Graphics*, 25(4):197–206, August 1991. SIGGRAPH '91 Las Vegas.
9. Paul Heckbert. Discontinuity meshing for radiosity. *Third Eurographics Workshop on Rendering*, pages 203–226, May 1992.
10. A. T. Campbell III and Donald S. Fussell. An analytic approach to illumination with area light sources. Technical Report TR-91-25, CS Dpt. U of Texas, Austin, TX, August 1991.
11. Dani Lischinski, Brian Smits, and Donald P. Greenberg. Bounds and error estimates for radiosity. In *Computer Graphics Proceedings, Annual Conference Series: SIGGRAPH '94* (Orlando, FL), pages 67–74. ACM, July 1994.
12. Dani Lischinski, Filippo Tampieri, and Donald P. Greenberg. Combining hierarchical radiosity and discontinuity meshing. In *Computer Graphics Proceedings, Annual Conference Series: SIGGRAPH '93* (Anaheim, CA, USA), pages 199–208. ACM, August 1993.
13. P. M. Prenter. *Splines and Variational Methods*. John Wiley & Sons Inc, New York, 1989.
14. François Sillion. A unified hierarchical algorithm for global illumination with scattering volumes and object clusters. *IEEE Trans. on Vis. and Comp. Graphics*, 1(3), September 1995.
15. François Sillion and George Drettakis. Feature-based control of visibility error: A multiresolution cluste ring algorithm for global illumination. In *Computer Graphics Proceedings, Annual Conference Series: SIGGRAPH '95* (Los Angeles, CA), pages 145–152. ACM, August 1995.
16. A. James Stewart and Sherif Ghali. Fast computation of shadow boundaries using spatial coherence and backprojections. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94* (Orlando, Florida, July 24–29, 1994), Computer Graphics Proceedings, Annual Conference Series, pages 231–238. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.
17. Filippo Tampieri. *Discontinuity Meshing for Radiosity Image Synthesis*. PhD thesis, Department of Computer Science, Cornell University, Ithaca, New York, 1993. PhD Thesis.
18. Seth J. Teller. Computing the antipenumbra of an area light. *Computer Graphics*, 26(4):139–148, July 1992. Proceedings of SIGGRAPH '92 in Chicago (USA).
19. Seth J. Teller and Patrick M. Hanrahan. Global visibility algorithms for illumination computations. In *Computer Graphics Proceedings, Annual Conference Series: SIGGRAPH '93* (Anaheim, CA, USA), pages 239–246. ACM, August 1993.
20. John R. Wallace, Kells A. Elmquist, and Eric A. Haines. A ray tracing algorithm for progressive radiosity. *Computer Graphics*, 23(3):315–324, July 1989. Proceedings SIGGRAPH '89 in Boston.

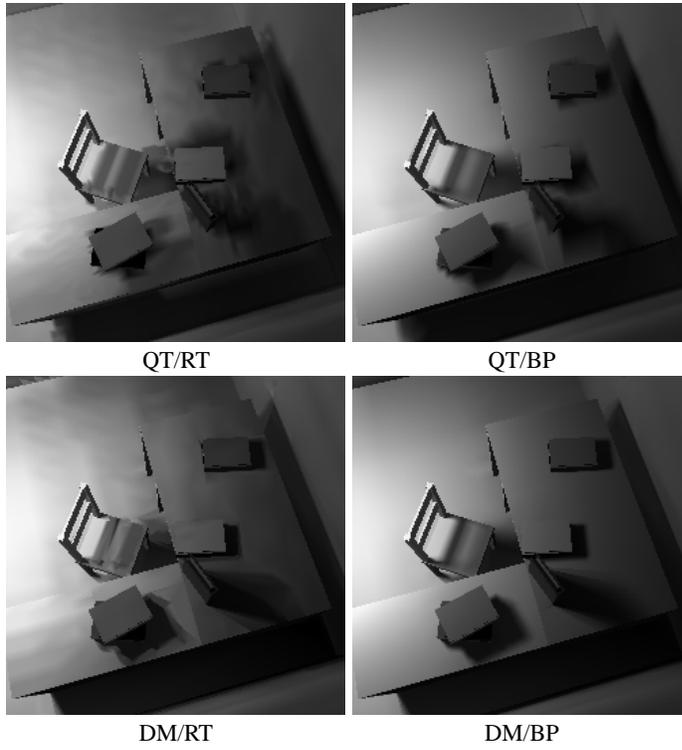


Fig. 5. Results of the four selected combinations for the "Books" image.