An Image-Space Refinement Criterion for Linear Hierarchical Radiosity

Ignacio Martin Xavier Pueyo
Departament d'Informatica i Matematica Aplicada
Universitat de Girona
C/ LLuis Santalo s/n
17003 Girona
Spain

Phone: +34 72 41.87.61 Fax: +34 72 41.83.99
e-mail : {imartin,xavier}@ima.udg.es

Dani Tost
Departament de Llenguatges i Sistemes Informatics
Universitat Politecnica de Catalunya
Av. Diagonal 647, Planta 8
08028 Barcelona
Spain

Phone: +34 3 401.66.67 Fax: +34 3 401.60.90
e-mail: dani@turing.upc.es

Abstract

A Linear Hierarchical Radiosity method using point collocation and triangle meshes is proposed that allows C^0 continuity and performs energy exchange at any level both in the shooter and the receiver; this method leads to an exact representation of the Gouraud shading interpolation that will be used for rendering. A new refinement criterion is presented which tries to improve image quality taking into account: the smoothness of the solution based on pixel intensity values instead of energy ones, and visibility changes along the surfaces for high gradient detection (sharp shadows). In order to perform a efficient refinement a data structure is proposed which isolates every shooter contribution over each receiver and allows to only refine high error interactions.

Keywords: hierarchical radiosity, higher order elements, tone operator, refinement criteria.

1 Introduction

Hierarchical Radiosity (HR) [6] is considered as a major step towards realistic lighting simulations systems. At the heart of HR is the oracle function which predicts the error that will be introduced by linking two patches at a given level of the hierarchy. On the basis of the predicted error, either the link is established, or it is refined and a series of links are computed at lower levels of the hierarchy. The acceptability of a predicted error is related to the level of precision desired in the rendering, which, in turn, depends on the quality of the produced images. However, most existing oracle functions are based on the computation of *object-space* magnitudes, such as form-factors and energy values, and do not take directly into account *image-space* values such as the intensities of the images.

The view-independency of radiosity is often outlined as a major advantage of the global illumination techniques which compute it. However, as stated in [1] it is not always desirable to generate many different views of a single scene. Based on this assumption, several view-dependent radiosity techniques have been proposed such as the Importance-Driven Radiosity [17] and Two-Pass (Multi-Pass) methods [14, 2].

We herein propose an view dependent Linear Hierarchical Radiosity algorithm using point collocation, based on triangular meshes and a quadtree subdivision scheme. A refinement criterion is introduced that takes into account pixel intensities mapped from the radiosities in order to achieve a given degree of smoothness for a fixed point of view. This criterion along with the use of linear elements allow an efficient representation of the radiosity function and a reduction of the number of links computed for a given image quality.

The paper is divided into 7 sections. In section 2, a summary of the previous related work is presented. Next, the basic linear hierarchical method is described. In section 4, a new refinement criterion is proposed and discussed. In section five interaction graphs are presented. The results of the implementation are discussed in section 6 and the conclusions along with the future direction of research are presented in section 7.

2 Previous work

2.1 View Dependent Radiosity Methods

Although most radiosity methods are inherently view-independent, some techniques incorporate view-dependent information at different radiosity processing stages: Two-Pass [14, 13, 9], Multi-Pass methods [2] and Importance-based radiosity. Since this paper is restricted to diffuse interreflections, the two former ones will not be addressed here.

Importance-based radiosity [17] aims at computing as exactly as possible the radiosities of the elements that will have a quantifiable impact on the final image. Importance is defined as the fraction of radiosity emitted by a patch that ultimately reaches the observer. It is a concept dual to radiosity and it is distributed along with it during the HR process. In [3] a different formulation of importance has been proposed, allowing it to be transported similarly to the radiosity. With this approach, importance and radiosity are transported using the same data structures and algorithms. It should be noted that although Importance Driven Radiosity is view-dependent, it is still based on energy values (object space) rather than in mapped intensities (image space).

2.2 Oracles in Hierarchical Radiosity

Oracle functions in HR were first introduced in [6]. The brightness-weighted refinement criterion is based on an upper-bound approximated value of the form-factors F_{ij} and F_{ji} between the involved patches i and j, and the energy transferred. Both are directly compared with the given thresholds ϵ_F and ϵ_{BF}.

In Importance-Driven Radiosity, the oracle [17] takes into account the importance of an interaction by computing the error as the product of the estimated importance with the estimated radiosity, the reflectance and an error bound ΔF_{ij} of the form-factor: $\|\tilde{I}_i \rho_i \Delta F_{ij} \tilde{B}_j\|$, where $\tilde{I}_i$ is the estimated importance of patch i, ρ_i is the reflectivity of patch i and $\tilde{B}_j$ is the estimated radiosity of patch j. The error bound ΔF_{ij} is set to the difference between the upper $(\overline{F_{ij}})$ and lower bound $(\underline{F_{ij}})$ of the form-factor which are computed as the maximum and minimum values respectively of point-to-disk form-factors at given samples on the patches.

In [10] a new strategy is proposed which refines all links that have a significant effect on the total error bound, independently of the magnitude of their energy transfer. The error is: $\underline{I_i} \Delta K_{ij} \bar{B}_j$ where $\underline{I_i}$ is the lower bound of the importance, ΔK_{ij} is the difference between the upper and lower bounds of $\rho_i F_{ij}$ and $\bar{B}_j$ is the upper bound of the radiosity.

Finally, Sillion and Drettakis [16] propose a feature-based oracle meeting a user-defined criterion. This metric takes into account the image quality by measuring how well illumination features are represented in it.

2.3 Radiosity with Linear Elements

The radiosity equation, which describes the radiosity B for all points x in an environment, can be expressed as [7]:

$$B(x) = E(x) + \rho(x) \int_S B(x')G(x,x')dA'$$

Being a continuous dimension function, it is generally approximated as a weighted sum of a finite number of simple textitbasis functions $N_i(x)$, defined over a subregion of the function domain:

$$B(x) \approx \tilde{B}(x) = \sum_{i=1}^{n} B_i N_i(x)$$

Point collocation and Galerkin Radiosity are the two most referenced methods proposed in order to minimize the error of the approximation of B [4].

In [12], a linear basis approximation with point collocation requiring a correct energy balance at the vertices of the patches is proposed, leading to the expression of a Vertex-to-Vertex Form-Factor.

3 Linear Hierarchical Radiosity

As stated above, point collocation and Galerkin Radiosity are the two most referenced methods proposed in order to minimize the error of the approximation of the radiosity. Both methods lead to a system of linear equations:

$$B_i^k = E_i^k + \rho_i^k \sum_j \sum_k F_{ij}^{kl} B_j^l$$

where B_i^k, E_i^k, ρ_i^k respectively represent radiosity coefficient, emissivity and reflectivity for basis function k of patch i; F_{ij}^{kl} is the transport coefficient from basis function l of patch j towards basis function k of patch i; finally B_j^l is the radiosity coefficient for basis function l of patch j.

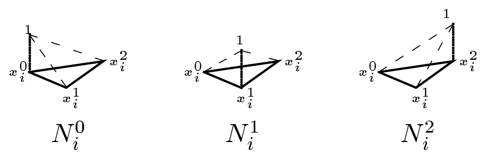

Figure 1: *The three basis functions $N_i^0, N_i^1 and N_i^2$ for triangle i*

Our algorithm uses a radiosity method with linear basis function and the point collocation technique using only triangles so we can use simple basis functions. For every patch we use three basis functions as depicted in Figure 1, so $l, k = 1..3$. Given those assumptions, the formula for the transport coefficient between two basis functions is as follows:

$$F_{ij}^{kl} = \int_{A_j} N_j^l G(x_i^k, x) dA \tag{1}$$

where x_i^k is the point on the support domain of basis function N_i^k chosen for integration on triangle i; this point is called the node of the basis function; finally, $G(x_i^k, x)$ is the geometric term that includes visibility. We choose this point to be at the maximum of the basis function, i.e. at the corresponding vertex.

If we want a hierarchical representation of the radiosity function we must be able to allow interaction at different levels; that means subdividing the patches and perform transport with finer basis functions. This refinement procedure is depicted in Figure 2. Refinement can be done at the receiver or at the shooter.

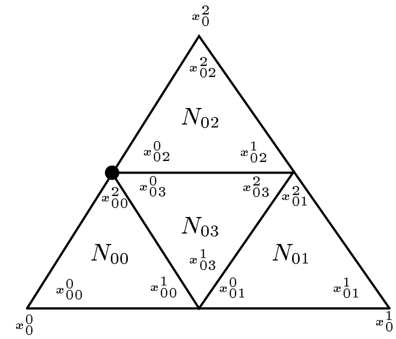

Figure 2: *Placement of basis function in one refinement*

To represent the interaction between two patches nine transport coefficients are needed, the square of the number of basis functions per patch. If we refine the receiver we perform a recursive subdivision following a quadtree scheme (regular subdivision in four subtriangles) and create the new transport coefficients for the new basis functions and destroy the old ones; if we consider one refinement of a receiver, 12 new basis functions are considered (4 patches, 3 basis per patch) and 36 transport coefficients must be computed, thus x_i becomes $x_{i0}...x_{i3}$. However, if we look more closely at equation 1 we realize that two basis functions with different domains but with the same node have the same transport coefficients for the same shooter; for example, following Figure 2, basis functions N_{00}^2, N_{03}^0 and N_{02}^0 share the same node so $F_{00,j}^{2l} = F_{03,j}^{0l} = F_{02,j}^{0l}$, as shown by the black dot. Also, transport coefficients from the upper level can be reused; for example, basis functions N_0^0 and N_{00}^0 share the same node and thus $F_{0,j}^{0l} = F_{00,j}^{0l}$. Following these rules there are only 9 new transport coefficients to be computed.

However radiosity coefficients cannot be shared because they represent radiosity at different levels and are computed after pushing and pulling radiosity over the hierarchy. In Figure 3 an example of this procedure is shown.

The push-pull procedure is illustrated in Figure 3 for triangle P_{01}. It is performed in three steps after gathering energy through the links:

1. Add the contribution from above. It's the contribution gathered in higher levels corresponding to the patch P_{01}. This contribution is added to the coefficients B_{01}^k, resulting in the new ones

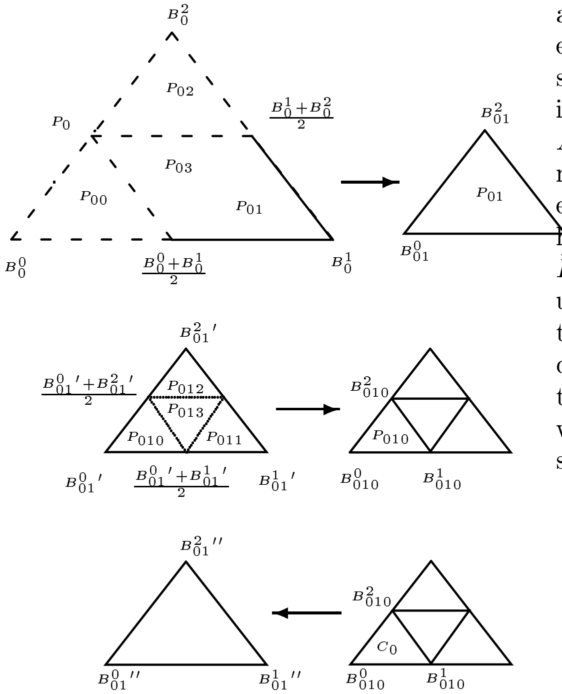

Figure 3: *Push-pull procedure in three steps*

$B_{01}^k{}'$.

2. Once we have radiosity gathered at this level we compute the function corresponding to each child at next level down by linear interpolation and then call recursively the push-pull procedure for those children. For exemple, following Figure 3, child P_{010} will receive $(\frac{B_{01}^0{}'+B_{01}^1{}'}{2}, B_{01}^1{}', \frac{B_{01}^1{}'+B_{01}^2{}'}{2}$.

3. When all the children return from the recursive call we set the radiosity coefficients B_{01}^k combining the children functions by averaging the coefficients of triangles P_{01i}. Then we return from the call.

Step 2 can be modified if patch at level k has neighbours of level $k+1$. In this case when pushing radiosity from patch at level k linear interpolation for calculating child functions may be incorrect. In Figure 4.a,b two shooters A and B contribute to a given mesh. Shooter A has contributed at level of patches $P_0, P_{10}, P_{12}, P_{13}$ and shooter B to $P_1, P_{00}, P_{01}, P_{02}, P_{03}$. After the gathering step we start pushing from patch P_0 to patches $P_{0i}, i = 0..3$, and we use linear interpolation to calculate the coefficients to be passed to the children; but is possible that neighbours have gathered more accurate information. If we consider radiosity from shooter A, in Figure 4.c the value passed to node at the middle of the diagonal should be computed by linear interpolation, $(B_0^0 + B_0^1)/2$, but the exact value has been computed at the neighbours and stored at $B_{10}^2 = B_{12}^0 = B_{13}^0$. So we do not interpolate and use that value to compute the functions to be passed to the children. We will need to isolate each shooter contribution to be able to identify different contributions at different levels and use the right values; we will also need to use a restricted quadtree in order to simplify neighbour identification.

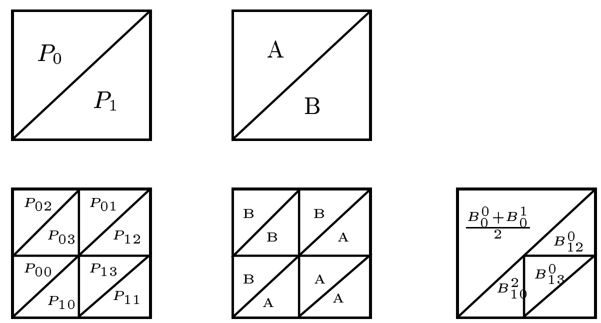

(a) *(b)* *(c)*

Figure 4: *T-vertex problem:* a) *hierarchy of patches,* b) *contribution of shooter* A *and* B*, c) vertex where interpolation is not used*

We have described refinement on the receiver and how energy is distributed over the hierarchy, but refinement can also be done at the shooter for accuracy on the transport coefficient. We will call a link the group of three transport coefficients associated to the three basis functions of a single shooter triangle to a certain receiver basis function. If a given refinement criterion decides that this link is not accurate enough then the shooter triangle is subdivided (if necessary) and four new links are established from the children to the receiver basis function. If the shooter was already subdivided then its radiosity coefficient is used directly; if not, the shooter is subdivided and the radiosities of the children are computed as in step 2 of the push-pull procedure.

We have described above a linear hierarchical method for radiosity but a refinement criterion is needed that decides when to perform the refinement of both shooters and receivers. In the next section we depict this criterion that consists of two oracles: one that tells us if a receiver must be refined, and other one that tells us if a shooter must be refined.

4 Refinement criterion

As said in previous sections, most refinement citeria [6] take into account the energy transported in each interaction. However, the relationship between the energy of a given patch and the intensity of the pixels on which this patch is mapped is not linear [18]. Therefore some refinements in the energy transfer will have a very little noticeable effect on the image. We herein propose a refinement criterion that progressively improves image quality working directly on the Gouraud shaded image to be rendered.

4.1 Receiver oracle

The oracle for receiver refinement takes into account the smoothness of the apoximation and the possible visibility artifacts (sharp shadows).

We assume that, at the rendering stage, a hardware implemented Z-Buffer visibility method along with Gouraud shading will be used. In addition, the scene surfaces are meshed in triangular patches as ststed before, avoiding Gouraud shading orientation dependence. The intensity value E at a given pixel X is therefore obtained by linearly interpolating, at image-precision, the intensity values of the three vertices $V_l^i, 1 \leq l \leq 3$ of the patch i that is visible at the pixel:

$$E(X) = \mathrm{LinearInterpolation}(E(V_i^0), E(V_i^1), E(V_i^2), X)$$

The vertices intensities are computed by applying a mapping function to the radiosity values B_i^l at the vertices:

$$E(V_i^l) = \mathrm{Map}(B_i^l), \quad 1 \leq l \leq 3$$

In this paper the mapping function proposed in [18] is used. This mapping uses an empiric model for the human eye response and takes into account the maximum intensity of the image and some parameters related to the display where it will be visualized.However it can be substituted by any other one.

A transition in the Gouraud shading at the shared edge between two neighbouring patches of a same surface is expected to be continuous. In other terms, the change in the intensity along a scan-line accross two coplanar adjacent patches is expected to be the same at the right side and the left side of the separating edge. The angle between the two linear intensity distributions at any point on the separating edge is a mesure of the shading discontinuity. Figure 5 illustrates this concept.

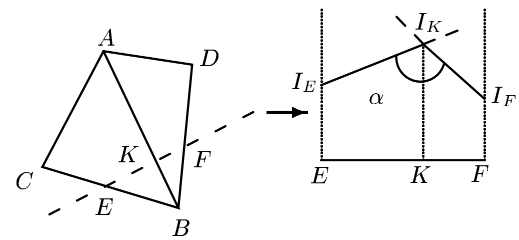

Figure 5: *Gouraud shading discontinuity; angle α mesures it*

Obviously, if the surface is curved, the angle value depends also on the geometrical angle formed by the patches which are no more coplanar.

A discontinuity at the separating edge indicates either an inaccurate radiosity estimation or a poor sampling ratio on the surface [15]. Assuming that the radiosity estimation is acceptable at the vertices, a discontinuity should lead to a subdivision of the two patches and a refinement of their vertices radiosity. However, the discontinuity will be significant only if the vertices radiosities are computed directly and not reconstructed by averaging the radiosities of the neighbouring patches, as with constant basis functions radiosity. With the linear hierarchical radiosity method depicted in previous section we acomplish this goal.

The question that arises is if the inverse statement is valid, i.e. may a continuity in the Gouraud shading accross two adjacent patches fail at identifying a poor sampling ratio?

First of all, many different patches acting as shooters contribute to a given radiosity value. These contributions may compensate each other when they are added up. Thus non linear contributions of each patch separately may give a globally linear radiosity distribution. This is illustrated in Figure 6.a, where two shooters 1 and 2 are contributing over adjacent receiver patches i and j; the radiosity function resulting from adding the two shooter contributions seem continuous across the edge separating the two receivers. However each isolated shooter contribution shows a high slope change across the edge, but opposite in sign so they cancel when they are added. Therefore, the criterion must test interpolation continuity for each shooter contribution separately. As a consequence, not all shooters interacting with a given patch may cause discontinuity and thus the receiver patch refinement should be realized only for the shooting patches which cause it, which is consistent with the HR strategy.

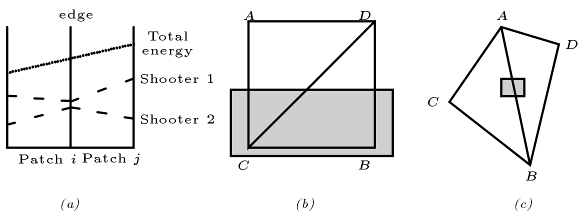

Figure 6: *Gouraud Continuity failing at refining*

A second aspect to be considered is the visibility of the separating edge vertices towards a shooting surface. Figure 6.b illustrates this case: a shadow is projected over two adjacent patches in such a way that vertices A and D are outside the shadow area whereas B and C are inside. If radiosity at A and D is fairly the same and radiosity at B and C also happens to be nearly the same, then Gouraud shading of the patches gives a smooth transition between them and fails at detecting that they should be subdivided. However this case can be easily detected because it corresponds to a discontinuity in the visibility along the edge. Therefore the Gouraud shading smoothness criterion has to take into account the difference of visibility of the edge vertices towards the shooting patch leading eventually to a subdvision.

Finally, the initial meshing must have a minimum density in order to avoid precision cases such as illustrated in Figure 6.c where a small shadow is projected on the center of one or two of the patches. As point collocation integrates over the shooter surface but at the receiver this integration is performed on a sigle point the initial meshing must thus be al least such that the maximum patch size is less than the minimum projection size of the smallest object in any scene surface. However, in practice, the initial meshing size can be fixed to a uniform value without loose of accuracy.

Summarizing, for each patch contribution, the proposed oracle function tests Gouraud shading continuity at the edges between adjacent receiving patches and checks the visibility of the edge vertices towards the shooter. The oracle thus decides whether to subdivide or not and, in the former case, new links are established. As the radiosity is linear, the new links are created between the shooter and the new sampling points. The refinement of the new links is exposed below.

In addition, the oracle takes into account importance expressed according to a linear formulation proposed in [11]. Importance is trasmited throught the same links as radiosity and is treated in the same way. Initial importance is computed in terms of the point of view used for the final rendering.

The oracle therefore searches all the edges over each mesh and perform three steps for every shooter contributing at the level of the adjacent patches:

1. Compute the visibility variation between the two vertices V_1^k and V_2^k of the edge corresponding to the shooter k:

$$\Delta V_{12}^k = \|V_1^k - V_2^k\|$$

2. This value is tested against a given threshold in order to decide if visibility change is small enough to consider the change of slope as a valid measure of smoothnes:

    - If ΔV_{12}^k is below a given threshold then we can use the change in the slope as a measure of smoothnes. As this oracle is applied many times during the algorithm we will use an easy but non conservative estimation of this change of slope.

      Since we are considering image values we must compute the slope change using pixel intensity values instead of radiosity ones. Since this map is not linear we also must take into account the contribution of the other shooters in order to estimate the real change of slope produced by shooter k; the greater the energy of the other shooters the smaller the slope change.

      Using geometry shown in Figure 7, we compute the slope change as:

$$\Delta E_{12}^k =$$

$$\frac{1}{2}||\mathrm{Map}(B_1 + B_{\min}) + \mathrm{Map}(B_2 + B_{\min})$$

$$-\mathrm{Map}(B_3 + B_{\min}) + \mathrm{Map}(B_4 + B_{\min})||$$

where, $B_{\min}$ is the minimum radiosity of the four vertices, without considering the contribution of the shooting patch k , and B_l, $1 \leq l \leq 4$ are the radiosities at the vertices of the two adjacent patches being $l = 1$ and $l = 2$ the separating edge vertices.

- If instead V_{diff}^k is superior or equal to the threshold, then a more conservative estimation of the slope change is used. We use the difference between the maximum and the minimum pixel intensity values of the four vertices, taking into account again the contribution of other shooters:

$$\Delta E_{12}^k =$$

$$\mathrm{MAX}_{l=1..4}\{\mathrm{Map}(B_{\min} + B_l)\} -$$

$$\mathrm{MIN}_{l=1..4}\{\mathrm{Map}(B_{\min} + B_l)\}$$

3. The estimator E_{diff}^k is tested against the slope threshold, weighted by the importance of the two adjacent patches:

$$E_{diff}^k \cdot (I_{left} + I_{right}) < \epsilon_{pixel}$$

where E_{diff}^k is the estimated intensity variation across the edge, due to the shooter patch k, I_{left} and I_{right} are the importance of the adjacent patches $left$ and $right$, and ϵ_{pixel} is the threshold.

If the expression become false then patches $right$ and $left$ are subdivided and new links created from shooter k.

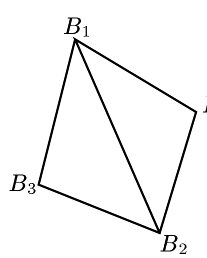

Figure 7: *Vertex naming when computing change of slope estimator*

This formulation has been used assuming that the triangular meshing is regular, that is, the triangles are equilateral. Other expressions could also be used in presence of irregular triangles.

4.2 Shooter oracle

The above discussion has described how a receiver has to be subdivided for increasing accuracy of the solution. But the criterion is based on accurate radiosity values at the vertices. If not enough accuracy is achieved at those vertices then the oracle will fail.

When a link is initially established between a shooter and a vertex (associated to a basis function at a given level) we must decide if transferring energy at that level is accurate enough to guarantee a good aproximation of the radiosity value at the vertex. If not we destroy that link, subdivide the shooter and create four new links from the children to the vertex, and then test those new links recursively for accuracy.

This lead us to a another oracle that decides if a link at a given level is accurate enough. We use an energy based oracle that compares the energy carried by the link with a given threshold, and if higher refines it. This criterion uses the maximum of the three products of shooter coefficients by the transport coefficients, weighted by the importance of the receiver triangle, and test it against a given threshold.

5 Interactions Graphs and Progressive Refinement

As mentioned in the previous section, a requirement of both the hierarchical method and the refinement criterion is to isolate the contribution of each shooting patch to each receiver basis function:

- When performing the push-pull procedure shooter contributions nust be kept isolated in order to push the right energies.
- When applying the refinement citeria on the receiver each shooter contribution is tested separatedly for smoothness.

In order to manage each shooting patch separatedly, a graph data structure is used. For efficiency reasons, patches belonging to a same surface are processed altogether in al single mesh; for example, an sphere will be discretized into a set of patches all of them belonging to the same mesh. Thus, for each shooting mesh, a graph is computed for any mesh receiving energy from it. The graph is a representation of the receiver-mesh at the level of refinement in which the energy transport between the shooter and the receiver is computed. The nodes of the graph are vertices of the receiver and the arcs are the edges between them. Figure 8 shows the two different graph representations of a same mesh, for two different shooters 1 and 2. The numbers in the patches of Figure 8.a show at which level each shooter is contributing. Figures 8.b and 8.c show the two graphs for the receiver mesh corresponding to shooters 1 and 2 respectively. Note how the graphs have a restricted quadtree structure.

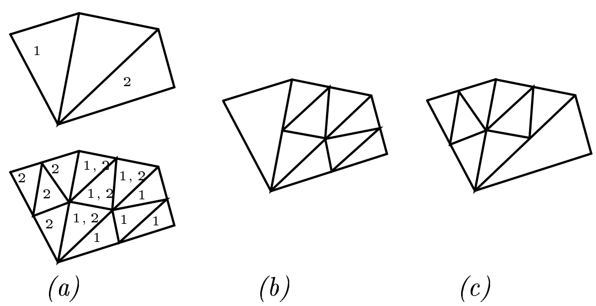

Figure 8: *Two graph corresponding to the same mesh for different shooters* $1, 2$*: a) mesh of triangles showing at which level each shooter contributes, b) graph for shooter* 1*, c) graph for shooter* 2

However, if all meshes are linked against any other one, the memory requirement of nmesh^2 graphs may be a serious drawback of the proposed method. Therefore, as proposed in [8], a progressive refinement strategy is used, allowing us to avoid an initial quadratic linking and therefore reducing the number of computed graphs. At the begining of each iteration, mesh-to-mesh interactions are checked and, if necessary, new interaction graphs are created. The experimental results in section 6 show that the memory requirement overhead due to the interaction graphs has an average value of 30 percent.

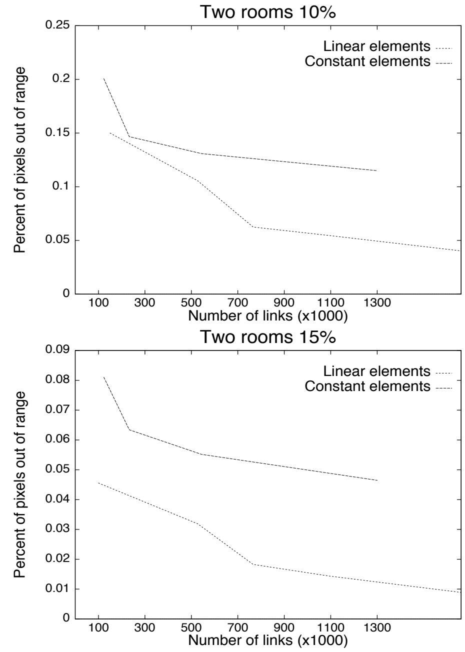

Figure 9: *Results with the two rooms scene.*

6 Implementation and results

In order to evaluate the proposed method we have compared it with the progressive refinement proposed in [8].

Colour plates 1.a and 1.b show the first scene used; it is composed by two rooms with four light sources each one, and some furniture; there are 386 initial surfaces. The second scene in colour plates 2.a and 2.b is composed by four rooms with two light sources each one and one table per room; there are 296 initial surfaces. Colour plates 3.a and 3.b show a detail of the mesh for first and second scene.

Colour plates 1.a and 2.a show the reference images. Colour plates 1.b and 2.b show the images computed with the linear method. Finally, colur plate 3 show a detail of the first scene along with the mesh generated by the linear method; note how the shadows are well recognized and represented.

The scenes have been processed at different levels of image quality, mesured in comparison to a reference image with a metric proposed in [5]. This metric is a pixel-by-pixel absolute difference of the approximate image from the refence image; we have used 10 % and 15% as differences. The reference image has been computed with a classical hierarchical method plus a final gathering pass per pixel [3]. The number of links for each image has been mesured and Figures 9 and 10 show plots for classical and linear methods.

It can be seen that the proposed method converges more quickly to a good image quality in terms of the ten-per-cent norm. Our method does not perform much better for a small number of links but when detailed images are required classical hierarchical method tends to produce good average images but with no well resolved details like sharp shadows; however the refinement criterion presented here allows to converge to detailed images faster. Other measurements like MSE (Minimum Square Error) or PSNR (Peak Signal to Noise Ratio) do not give much better results, but this a common problem when analizing image quality. In fact we believe that better measures for image quality would show better results for this method.

In all executions the memory overhead due to the interaction graphs is constant, with an average value of 30 percent.

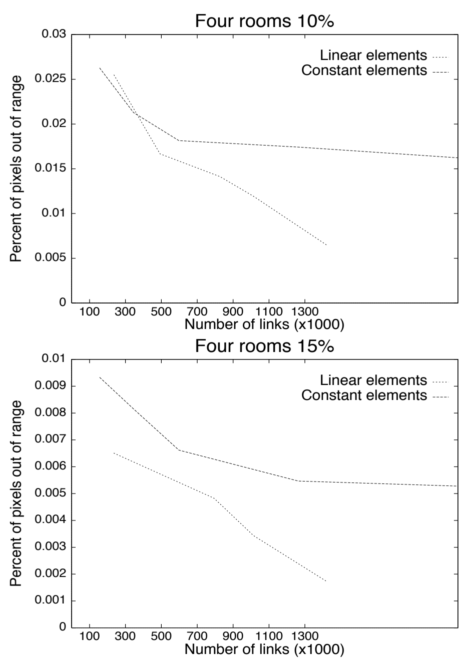

Figure 10: *Results with two the tfour rooms scene.*

7 Conclusions and future work

A new linear HR algorithm for triangular meshes using point collocation has been proposed, based on a refinement criterion that takes into account the expected smoothness of Gouraud shading at the separating edges between adjacent patches of a same surface. The method requires the storage of an interaction graph for any important mesh to any shooter that have a significant impact on it. This data structure alows to refine separately the interactions of a patch with each shooter.

The current implementation results show that this criterion reduces significantly the number of links for a given accuracy while keeping a reasonable memory requirement overhead.

We are currently working on the analysis of a better meshing strategy, adapted to the shadow areas. Another future research line is the use of clusters in order to reduce the number of graphs and thus making the progressive linking strategy unnecessary and avoiding energy loses and memory overhead.

Acknowledgments

This work is partially supported by the grant TIC 95-0630-C05-05 of the CICYT. Ignacio Martin has developed this work with a FI (Formacio de Investigadors) grant from the Generalitat de Catalunya. Thanks to George Drettakis and Frederic Perez for his unvaluable comments and help.

References

[1] Shenchang Eric Chen. Incremental radiosity: An extension of progressive radiosity to an interactive image synthesis system. In *Computer Graphics (ACM SIGGRAPH '90 Proceedings)*, pages 135–144, August 1990. Published as Computer Graphics (ACM SIGGRAPH '90 Proceedings), volume 24, number 4.

[2] Shenchang Eric Chen, Holly E. Rushmeier, Gavin Miller, and Douglass Turner. A progressive multi-pass method for global illumination. In *Computer Graphics (ACM SIGGRAPH '91 Proceedings)*, pages 164–174, July 1991. Published as Computer Graphics (ACM SIGGRAPH '91 Proceedings), volume 25, number 4.

[3] Per H. Christensen, David H. Salesin, and Tony D. DeRose. A continuous adjoint formulation for radiance transport. In *Fourth Eurographics Workshop on Rendering*, pages 95–104, Paris, France, June 1993. Eurographics Technical Report Series EG 93 RW.

[4] Michael F. Cohen and John R. Wallace. *Radiosity and Realistic Image Synthesis.* Academic Press Professional, Boston, MA, 1993. Excellent book on radiosity algorithms.

[5] George Drettakis and Eugene Fiume. Concrete computation of global illumination using structured sampling. In *Third Eurographics Workshop on Rendering*, pages 189–202, Bristol, UK, May 1992.

[6] Pat Hanrahan, David Salzman, and Larry Aupperle. A rapid hierarchical radiosity algorithm. In *Computer Graphics (ACM SIGGRAPH '91 Proceedings)*, pages 197–206, July 1991. Published as Computer Graphics (ACM SIGGRAPH '91 Proceedings), volume 25, number 4.

[7] Paul Heckbert and J. M. Winget. Finite element methods for global illumination. Technical Report UCB/CSD 91/643, University of California at Berkeley, Berkeley, CA, 1991.

[8] Nicolas Holzschuch, Francois Sillion, and George Drettakis. An efficient progressive refinement strategy for hierarchical radiosity. In *Fifth Eurographics Workshop on Rendering*, pages 343–357, Darmstadt, Germany, June 1994. also available in "Photorealistic Rendering Techniques" (eds. G. Sakas, P. Shirley, and S. Muller).

[9] Arjan J. F. Kok and Frederik W. Jansen. Adaptive sampling of area light sources in ray tracing including diffuse interreflection. *Computer Graphics Forum (Eurographics '92)*, 11(3):C289–C298, September 1992.

[10] Dani Lischinski, Brian Smits, and Donald P. Greenberg. Bounds and error estimates for radiosity. In *Computer Graphics Proceedings, Annual Conference Series, 1994 (ACM SIGGRAPH '94 Proceedings)*, pages 67–74, 1994.

[11] I. Martin, D. Tost, and X. Pueyo. An extention of importance to linear hierarchical radiosity. *Report LSI-96-R6*, 1996.

[12] Nelson L. Max and Michael J. Allison. Linear radiosity approximation using vertex-to-vertex form factors. In David Kirk, editor, *Graphics Gems III*, pages 318–323. Academic Press Professional, Boston, MA, 1992.

[13] Mark C. Reichert. A two-pass radiosity method driven by lights and viewer position. Master's thesis, Program of Computer Graphics, Cornell University, January 1992.

[14] Holly E. Rushmeier. Extending the radiosity method to transmitting and specularly reflecting surfaces. Master's thesis, Program of Computer Graphics, Cornell University, 1986. longer version of paper.

[15] Francois Sillion. Detection of shadow boundaries for adaptive meshing in radiosity. In James Arvo, editor, *Graphics Gems II*, pages 311–315. Academic Press Professional, Boston, MA, 1991.

[16] Francois Sillion and George Drettakis. Feature-based control of visibility error: A multiresolution clustering algorithm for global illumination. In *Computer Graphics Proceedings, Annual Conference Series, 1995 (ACM SIGGRAPH '95 Proceedings)*, pages 145–152, 1995.

[17] Brian E. Smits, James R. Arvo, and David H. Salesin. An importance-driven radiosity algorithm. In *Computer Graphics (ACM SIGGRAPH '92 Proceedings)*, pages 273–282, July 1992. Published as Computer Graphics (ACM SIGGRAPH '92 Proceedings), volume 26, number 4.

[18] Jack Tumblin and Holly E. Rushmeier. Tone reproduction for realistic computer generated images. Tech. Report GIT-GVU-91-13, Graphics, Visualization & Usability Center, Coll. of Computing, Georgia Institute of Tech., 1991. a reasonable method to scale radiosity solutions.