Combining Hierarchical Radiosity and LODs

Reynald Dumont†‡, Kadi Bouatouch†

†IRISA, Campus de Beaulieu
35042 Rennes Cedex, France

‡CRIL Ingénierie, Square Chene Germain
35510 Cesson-Sévigné, France

E-Mail: rdumont@irisa.fr, kadi@irisa.fr

Abstract

Clustering and partitioning based methods have improved drastically Hierarchical Radiosity (HR). The first method reduces the number of links but lacks precision since it does not account precisely for the orientations of the patches, within the emitting and receiving clusters, during the illumination calculation. As for the second approach, it is efficient since it partitions large environments into 3D cells and employs ordering strategies that need only a subset of cells for computing illumination at each step of the resolution process. Nevertheless, the computation complexity still remains high when the cells contain many objects of highly detailed geometry. The new HR, described in this paper, brings solutions to these problems by making use of Levels Of Details (LOD) that approximate the geometry of the object surfaces in the scene.

1 Motivation

Hierarchical radiosity [1], called HR from now on, is widely used for a large range of applications such as: thermal engineering, lighting simulation, radiative transfer within canopies, acoustics, etc. Even for scenes of moderate complexity, HR still is a demanding process in terms of memory and computing resources. This is due to the need of meshing surfaces into elements (or patches), creating links and computing visibility relationships between these elements. For complex building interiors, solutions have been proposed to cope with the memory and computation problems. These solutions rely on a preprocessing step [2, 3, 4, 5] consisting in partitioning the scene into 3D cells and computing a visibility graph expressing visiblity relationships between these cells. Only a subset of cells, concerned with the illumination computation (at the current iteration of the resolution process), is maintained in memory. Once this preprocessing has been effected, ordering strategies [6] are employed for radiosity computation. These strategies need downloading the cells involved in the computation process, removing unnecessary ones from memory and back to the disk. The ordering of these cell transfers makes it possible to reduce expensive read and write operations (to/from the disk or memory). Even though their efficiency has been proved, these strategies do not reduce the complexity of scenes where cells are composed of a large number of small input polygons that finely describe surfaces and objects of highly detailed geometry. Indeed, in a cell, if the number of polygons is high then so do the number of patches, the number of links as well as the memory needed.

As the flux of a patch is proportional to its surface area, when a small patch is selected in isolation as an emitting patch, its flux (to be emitted) may be insignificant. Consequently, its effect on the illumination of the environement may be unnoticeable. On the other hand, if a set of small neighboring patches are together selected as emitting patches, the associated total flux may contribute efficiently to the scene illumination, which will speed up the convergence of the resolution process. One way of carrying out this idea is to make use of levels of details (LOD). More clearly, the surface of an object may be represented by different LOD approximations organized as a hierarchy, called from now on *LOD hierarchy*. The highest level corresponds to the more detailed geometry of the surface. In this way two distant surfaces may interact (in the sense of hierarchical radiosity) at coarse LODs while close surfaces at finer LODs. The number of links is consequently reduced drastically. As seen later on, accounting for LODs in HR raises some problems that have to be solved: energy conservation, occlusion, internal reflection, energy distribution similar to push-pull operations. To our knowledge the use of LOD in HR has not been addressed so far.

This paper is organized as follows. Our approach to LODs is outlined in section 2. Then the different steps involved are in turn described. These steps are: energy emission (section 3.1), distribution of the energy gathered by a receiving patch to its descendants in the associated LOD hierarchy (section 3.2), and internal emission (section 3.3). Before concluding, some results are given in section 4.

2 Outline

2.1 Definitions

We suppose all the objects within the scene are modeled with polygonal surfaces, each one being represented by a hierarchy of LOD approximations. The root of a hierarchy corresponds to the coarsest polygonal approximation of the surface while a leaf to the finest one (see figure 1). A non terminal node and a leaf node are polygons that are called *macrofacet* and *facet* respectively. In a LOD hierarchy, the children of a non terminal node (macrofacet) may be macrofacets and/or facets not coplanar. Note that an input polygon without LOD approximations is also represented by a LOD hierarchy but composed of only one node. A facet may be subdivided into another hierarchy of patches (HP) by the refinement process employed in any classical HR. In this case, an internal node and a leaf node of a HP are referred to as *patch* and *surface element* respectively.

2.2 Creating LOD hierarchies

The LOD hierarchy associated with a surface can be created with the help of any surface simplification method. There are broadly three simplification approaches: vertex decimation, edge contraction and surface refinement. The decimation approach [7, 8] removes recursively one or several vertices from the exact geometry while the contraction approach [9, 10, 11] removes edges and replace them with new vertices. Conversely, the refinement approach [12] starts with a coarse approximation of the surface then refines it progressively by inserting new vertices.

To make this process easier, we have developed an interactive software tool using the MOTIF and OpenGL libraries. With the help of the mouse, the user can select a surface and start the simplification process giving finer and finer approximations.

2.3 Problem statement

In classical HR, only facets are used and refined during the resolution process. A link associated with a pair of mutually visible facets is called initial link. Let N_{IL} be the number of initial links. The number of links created by a HR algorithm has a lower bound equal to N_{IL}. When the scene contains objects of highly detailed geometry, N_{IL} may be high and so does the final number of links.

To overcome this difficulty, one solution is to make use of HR with clustering techniques [13, 14]. However, these techniques do not account precisely for the orientations of the facets within the emitting and receiving clusters during the energy exchanges. In [15] Gibson and Hubbold brought some improvements to account for surface orientation but their method still make assumptions. In addition, they make use of the clustering method by Sillion [14] which assumes a cluster as an isotropic participating medium (which is not generally the case) characterized by an extinction factor (difficult to estimate accurately).

Visibility between clusters has been improved in [16]. In addition, clustering techniques have also been applied to environments containing specular surfaces [17, 18]. These techniques assume that clusters are approximated as point light sources since they rely on intensity. Another clustering approach accounting for specular surfaces has been proposed in [19]. The method makes use of two kinds of hierarchy: line space hierarchy and line space hierarchy with clusters. Moreover, the authors do not give any detail about the way they account for the orientation of the facets within the receiving and emitting clusters.

As our objective is lighting simulation we cannot content ourself with these strong assumptions.

Our solution is to combine HR and LODs for reducing N_{IL} significantly, which entails a smaller number of links, a lower visibility calculation cost and less memory for storing the different data structures involved. Indeed, two surfaces may interact at a coarse LOD if they are sufficiently distant. In this case, these are macrofacets that interact mutually instead of facets or patches or surface elements. Unlike clustering-based HR, our method accounts precisely for the orientations of the emitting and receiving macrofacets (and facets) whenever a macrofacet shoots its energy or distributes it to its descendants.

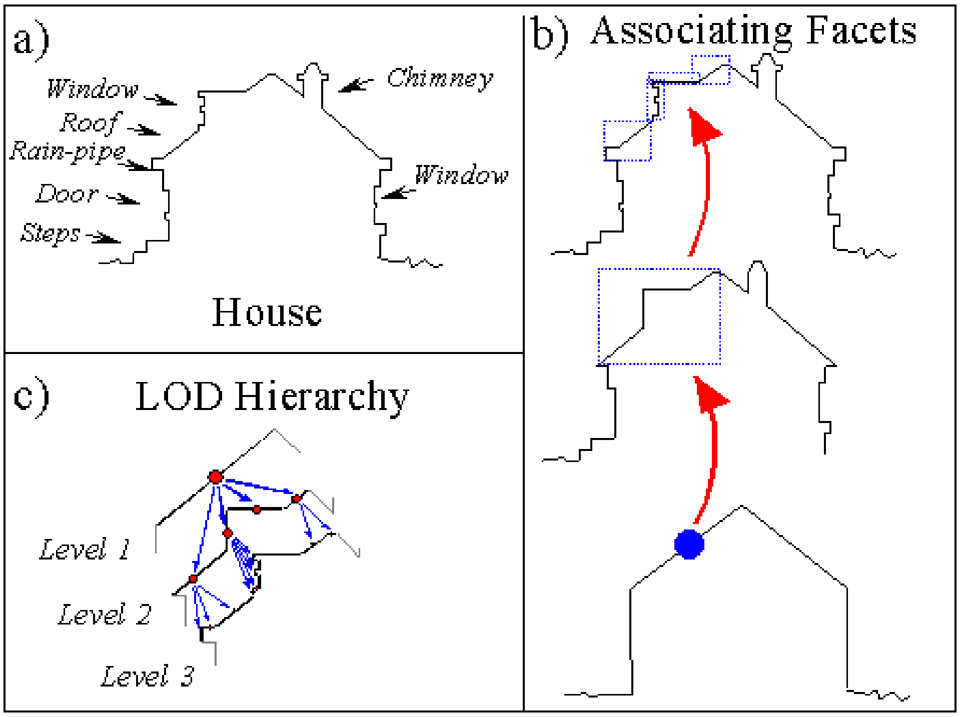

Figure 1: LOD hierarchy

2.4 Algorithm overview

In this section we describe only the basic ideas of the algorithm, details are given in the following sections.

Our HR uses a *shooting* technique and removes all the created links after each iteration for a reason of memory saving. Let us call *root macrofacet* a macrofacet at the root of a LOD hierarchy. The algorithm starts by selecting the emitting root macrofacet E of maximum flux. E has to shoot its flux to each other visible root macrofacet R. In fact, E and R may not interact directly but through nodes of their respective LOD hierarchies. More precisely, the two LOD hierarchies are top-down traversed and two estimated form factors F_{E-R} and F_{R-E} are calculated for the two current nodes N_E and N_R, belonging to E and R respectively. If F_{E-R} and F_{R-E} are below a certain threshold T_{LOD} then these two nodes are linked and N_E shoots its energy to N_R using F_{E-R} weighted by a visibility factor. Before shooting the flux of N_E, all the descendants (finer LOD) of N_E exchange one to another their flux according to one iteration of the Jacobi resolution method. This operation will be referred to as *internal shooting*. Then, for each N_R a new radiosity B_c of N_E is computed (depending on the location of N_R in the scene). This radiosity accounts for the internal exchanges within N_E. B_c will be called *contextual radiosity* because it depends on the context of N_R. Next, N_R distributes the gathered energy to its descendants. This process is illustrated by figure 2.

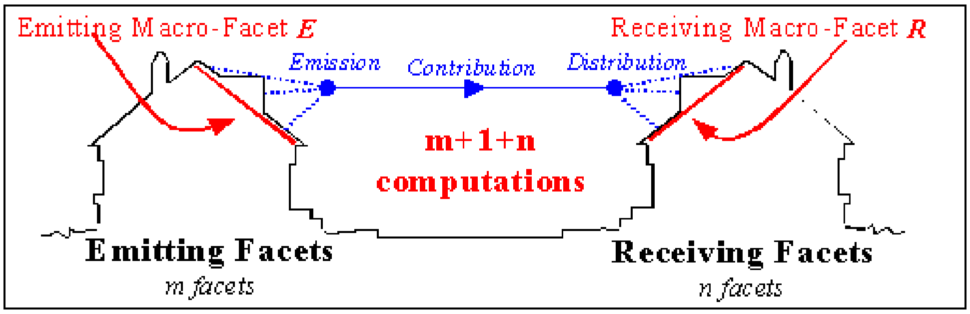

Figure 2: LOD hierarchy

Note that if the traversal of R's LOD hierarchy leads to a leaf P while F_{E-R} still remains higher than T_{LOD} then P is subdivided into a hierarchy of patches as in classical radiosity but with another refinement criterion T_P using a more precise form factor computation method. In this case E is linked to a patch or a surface element P_R (belonging to P) if the computed form factor $F_{N_E-P_R}$ is below T_P. Then N_E shoots its flux to P_R using the precise form factor. This gathered flux is then pushed/pulled to the ascendants/descendants of P_R through the hierarchy of patches associated with P.

In the same way, if the traversal of E's LOD hierarchy leads to a leaf P while F_{R-E} still remains higher than T_{LOD}, the same process is applied to P.

F_{E-R} is a point-surface form factor whose appex is the center of the emitting macrofacet (or facet) E and subtented by the receiving macrofacet (or facet) R. F_{R-E} is a point-surface form factor whose appex is the center of the receiving macrofacet (or facet) R and subtented by the emitting macrofacet (or facet) E (figure 4). When trying to link two primitives, these two form factors do not account for the visibility factor so as to not slow down our HR algorithm. On the other hand, this factor is taken into account once the linking has been performed. In order to be more vigilant regarding the choice of the LOD approximation, we choose a small value of the threshold T_{LOD}. On the other hand, T_P may be chosen larger or equal to T_{LOD}.

The algorithms describing our HR are given by the figures 3, 6 and 7. Let us explain the role of some functions invoked in figures 6 and 7. As in classical HR, the function *Refine_And_Link()* recursively subdivides either the emitting facet only (leaf of a LOD hierarchy) or the receiving facet only or both depending on the value of the last parameter of this function which is S_E, S_R and $S_{E/R}$ respectively. It also establishes links when the refinement criterion is met. This criterion compares the precise form factor with the threshold T_P.

Shooting() is a function which shoots the energy of the emitter through the links determined by the function *Refine_And_Link()*, as in classical HR.

Shoot_Energy() computes the radiosity B_R gathered by a macrofacet using the form factor F_{E-R} (figure 4) weighted by a visibility factor.

Choose_Node(N_E,N_R) selects a node (N_E if $F_{R-E} > F_{E-R}$ else N_R), the children of which are passed in turn to the function *Link_And_Shoot()* together with the other node. If one of the nodes (N_E or N_R) is a leaf of a LOD hierarchy, the other is systematically selected (in this case, links are established between a macrofacet and patches).

```
HR_With_LOD()
    While not convergence() {;
        E = Root_Macrofacet_Of_Maximum_Flux() ;
        Internal_Shooting(E);
        for each root macrofacet R {
            /* An input polygon without LODs is also a root macrofacet */
            Link_And_Shoot(E, R, T_LOD, T_P) ;
            /* T_LOD and T_P are form factor thresholds */
            Remove_Links(E);
            for each leaf P of R's LOD hierarchy {
                Push_Pull(P) ;
            }
        }
        ΔΦ_E = 0 ; /* Reset of the unshot flux of the emitting macrofacet */
    }
```

Figure 3: New Hierarchical Radiosity

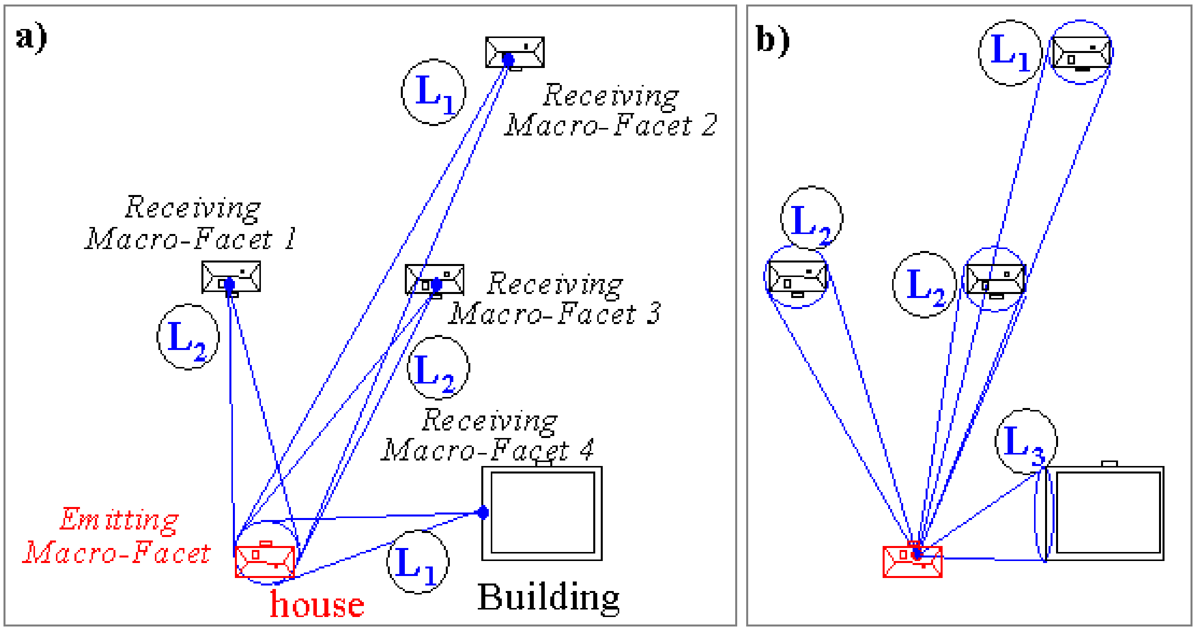

Figure 4: Form factors. L_1, L_2, L_3 mean: *level* 1, 2 and 3 respectively. a) means that we are selecting the LOD level of the emitting macrofacet while b) that of the receiving macrofacet.

3 Emission and Distribution

Let E and R be the emitting and receiving macrofacets respectively. Let us call e_k and r_k the children nodes of E and R in their respective LOD hierarchies. e_k and r_k may be facets or macrofacets.

3.1 Emission

The goal now is to compute the flux of E that has to be actually considered for the energy transfer $E \rightarrow R$. This flux divided by the surface area of E will be called *contextual radiosity* (denoted B_c) from now on. Its particularity is that it accounts for the orientations of the nodes e_k, which makes more accurate the energy transfer between two nodes of LOD hierarchies.

The total flux of E is :

$$\Phi_E = A_E \cdot B_E, \tag{1}$$

where A_E and B_E are the surface area and the radiosity of E respectively.

The flux $\Phi_{E\rightarrow R}$ emitted by E toward R is:

$$\Phi_{E\rightarrow R} = A_E \cdot B_c \cdot F_{E-R}, \tag{2}$$

F_{E-R} being the form factor between E and R.

Taking into account the orientations of the nodes e_k (facets or macrofacets), yields:

$$\Phi_{E\rightarrow R} = \sum_k \Phi_{e_k\rightarrow R} = \sum_k A_{e_k} \cdot B_{e_k} \cdot F_{e_k-R}, \tag{3}$$

As mentioned before, two nodes E and R of LOD hierarchies are linked only if the form factors (not including the visibility factor) F_{E-R} and F_{R-E} are below a threshold T_{LOD}. If T_{LOD} is small enough the F_{e_k-R} may be approximated as differential form factors. Consequently:

$$F_{e_k-R} = \frac{\cos\Theta_{e_kR} \cdot \cos\Theta_{Re_k} \cdot A_R}{\pi \cdot d^2_{e_k-R}}, \tag{4}$$

d_{e_k-R} being the distance separating the centers of e_k and R.

Similarly we can write:

$$F_{E-R} = \frac{\cos\Theta_{ER} \cdot \cos\Theta_{RE} \cdot A_R}{\pi \cdot d^2_{E-R}}, \tag{5}$$

From equations 2, 3, 4 and 5 we get:

$$B_c = \sum_k \frac{A_{e_k} \cdot B_{e_k} \cdot F_{e_k-R}}{A_E \cdot F_{E-R}} \approx \sum_k \frac{A_{e_k}}{A_E} \cdot \frac{\cos\Theta_{e_kR}}{\cos\Theta_{ER}} \cdot B_{e_k}, \tag{6}$$

assuming $d_{e_k-R} \approx d_{E-R}$ and $\Theta_{Re_k} \approx \Theta_{RE}$ $\forall k$ (figure 5). B_c has to be multiplied by a visibility factor V_{e_kR} between e_k and R.

The radiosity due to E of the receiver node R (of a LOD hierarchy) is then:

$$B_R = \rho_R \cdot F_{R-E} \cdot B_c, \tag{7}$$

where ρ_R, the average reflectivity of R, is given by : $\rho_R = \sum_k \frac{A_{r_k}}{A_R} \cdot \rho_{r_k}$

Here F_{R-E} includes the visibility factor evaluated by tracing visibility rays between sample points on E and R.

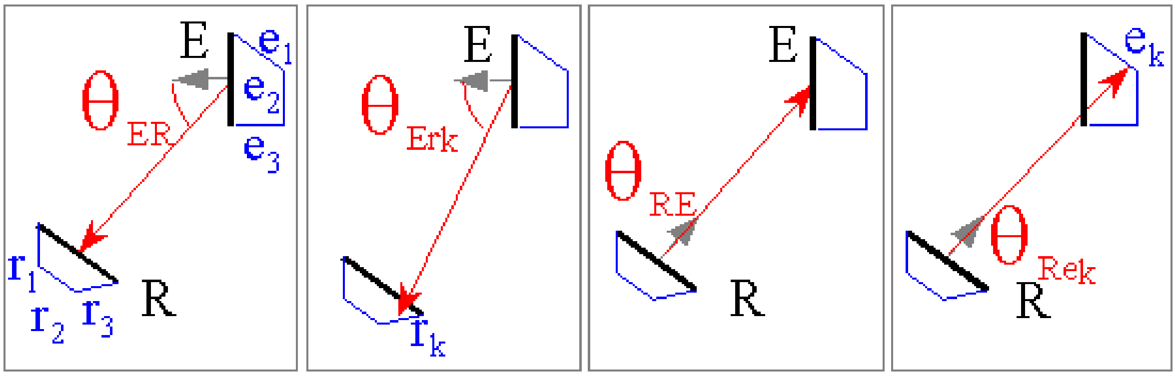

Figure 5: Angles

3.2 Distribution

The radiosity gathered by a facet (leaf of a LOD hierarchy) is pushed/pulled through the associated patch hierarchy (PH). As for a macrofacet R, whenever it gathers energy from E it distributes it to its descendants (facets and macrofacets) as explained hereafter. This energy distribution must account for the orientations of E and the children nodes r_k of R.

Recall that B_R is the radiosity gathered by R and due to the emitter E. Let us denote B_{r_k} the radiosity that has to be distributed to each r_k by its parent macrofacet R. We can write:

$$A_R \cdot B_R = \rho_R \cdot F_{R-E} \cdot B_c \cdot A_R = \rho_R \cdot F_{E-R} \cdot B_c \cdot A_E \tag{8}$$

$$A_{r_k} \cdot B_{r_k} = \rho_{r_k} \cdot F_{r_k-E} \cdot B_c \cdot A_{r_k} = \rho_{r_k} \cdot F_{E-r_k} \cdot B_c \cdot A_E \tag{9}$$

As before we suppose T_{LOD} small enough for using the differential form of F_{E-r_k}. We get then:

$$F_{E-r_k} = \frac{\cos\Theta_{Er_k} \cdot \cos\Theta_{r_kE} \cdot A_{r_k}}{\pi \cdot d^2_{E-r_k}}, \tag{10}$$

By dividing equations 8 by 9, substituting equation 10, assuming $\Theta_{E-r_k} \approx \Theta_{E-R}, d_{Er_k} \approx d_{ER}, \forall k$ (figure 5) and adding the visibility factor V_{Er_k}, we obtain:

$$B_{r_k} \approx B_R \cdot \frac{\cos\Theta_{r_kE}}{\cos\Theta_{RE}} \cdot \frac{\rho_{r_k}}{\rho_R} \cdot V_{Er_k} \tag{11}$$

We see that equation 10 accounts for the orientations of the nodes r_k and the spatial location of the emitter, which makes more precise the radiosity distribution. Note that the equations 6, 7 and 11, provide us with simple and fast expressions giving the contextual and distributed radiosities. Indeed, suppose that E and R have respectively m and n children. Only $(m+n)$ form factors F_{e_k-R} and F_{E-r_k} are computed (including visibility), avoiding then the evaluation of $m*n$ form factors $F_{e_k-r_k}$.

3.3 Internal Shooting

Consider figure 9. Suppose a root macrofacet I has been chosen for shooting its energy (see figure 3). Before invoking the function *Link_And_Shoot()*, all the I's descendants must shoot their energy one to another according to one iteration of Jacobi type. To this end, the LOD hierarchy associated with I is traversed from the leaves to the root. At each depth level of this hierarchy:

- the nodes having the same parent shoot their energy one to another,
- move to the upper depth level and repeat.

This energy transfer between the children nodes of the same parent will be called *Internal Shooting*. Its role is to avoid sub-estimating the flux emitted by a macrofacet. Let us explain the internal shooting operation with the help of figure 9. First, the facets i_j and i_k shoot their energy one to another and so do i_l and i_m, using classical HR. Then we compute the contextual radiosity of I_1 and shoot its energy to I_2 using the updated radiosities of i_j and i_k. Then I_2 distributes its gathered energy to its descendants (say i_l and i_m). We proceed in the same way for shooting the energy of I_2 toward I_1. This operation is repeated for the other upper depth levels (i.e. coarser approximations) in the LOD hierarchy as explained by the algorithm of figure 8. To make the *Internal Shooting* faster, two nodes of different parents (such as i_j and i_l in figure 9) do not exchange energy directly but through their respective parents (I_1 and I_2). Note that, at the first call, the argument of *Internal_Shooting*() is a root macrofacet, say I in figure 9.

```
Link_And_Shoot(LOD_HR_Node N_E, LOD_HR_Node N_R, T_LOD, T_P)
    if N_E and N_R are leaves of LOD hierarchies {
        Refine_And_Link(N_E, N_R, T_P, S_E/R) ;
        Shooting(N_E, N_R);
        }
    if (F_E-R < T_LOD) and (F_R-E < T_LOD) {
        Link(N_E, N_R) ;
        B_c = Compute_Contextual_Radiosity(N_E,N_R); /* Equation 6 */
        B_R = Shoot_energy(N_E, N_R) ; /* Equation 7 */
        Distribute_Energy_To_Descendants(N_E,N_R,B_R) ; /* Equation 11 */
        }
    if (F_E-R > T_LOD) and (F_R-E < T_LOD) {
        if N_R is not a leaf of a LOD hierarchy {
            for each child C_NR of N_R
                Link_And_Shoot(N_E,C_NR,T_LOD, T_P) ;
            }
            else { /* N_R is a leaf */
                if N_E is not a leaf of a LOD hierarchy {
                    B_c = Compute_Contextual_Radiosity(N_E,N_R); /* Equation 6 */
                    }
                if (F_E-R < T_P) {
                    Link(N_E, N_R) ;
                    Shooting(N_E, N_R);
                    }
                    else {
                        Refine_And_Link(N_E, N_R, T_P,S_R) ;
                        Shooting(N_E, N_R);
                        }
                }
        }
```

Figure 6: Shooting process: first part, continued in figure 7

```
if (F_{E-R} < T_{LOD}) and (F_{R-E} > T_{LOD}) {
    if N_E is not a leaf of a LOD hierarchy {
        for each child C_{NE} of N_E {
            Link_And_Shoot(C_{NE}, N_R, T_{LOD}, T_P) ;
        }
        else { /* N_E is a leaf */
            if (F_{R-E} < T_P {
                Link(N_E, N_R) ;
                }
                else {
                    Refine_And_Link(N_E, N_R, T_P, S_E ) ;
                    }
            if N_R is not a leaf of a LOD hierarchy {
                B_R = Shoot_Energy(N_E, N_R); /* Equation 7 */
                Distribute_Energy_To_Descendants(N_E, N_R, B_R) ; /* Equation 11 */
                } else {
                    Shooting(N_E, N_R);
                    }
            }
    }
if (F_{E-R} > T_{LOD}) and (F_{R-E} > T_{LOD}) {
    Which = Choose_Node(N_E, N_R) ;
    for each child C_W of Which {
        if (Which = N_E)    {
            Link_And_Shoot(C_W, N_R, T_{LOD}, T_P) ;
            }
            else {
                Link_And_Shoot(N_E, C_W, T_{LOD}, T_P) ;
                }
        }
    }
```

Figure 7: Shooting process: continued

```
Internal_Shooting(E)
    /* E is the emitting macrofacet */
    for each child C_E of E {
        if C_E is not a leaf of LOD hierarchy {
            Internal_Shooting(C_E);
            }
    }
    for each child C_E of E {
        for each child C'_E of E (C_E ≠ C'_E) {
            if C_E and C'_E are leaves of LOD hierarchy {
                /* Classical HR */
                Refine_And_Link(C_E, C'_E, T_P, S_E/R) ;
                /* S_E/R tells us that the emitting and receiving facets may be refined */
                Shooting(C_E, C'_E) ;
                Remove_Links(C_E);
                Push_Pull(C_E);
                }
                else { /* Energy exchange with LODs */
                    B_c = Compute_Contextual_Radiosity(C_E, C'_E) ; /* equation 6 */
                    B_R = Shoot_Energy(C_E, C'_E) ; /* equation 7 */
                    Distribute_Energy_To_Descendants(C_E, C'_E, B_R) ; /* equation 11 */
                    }
        }
    }
```

Figure 8: Internal Shooting

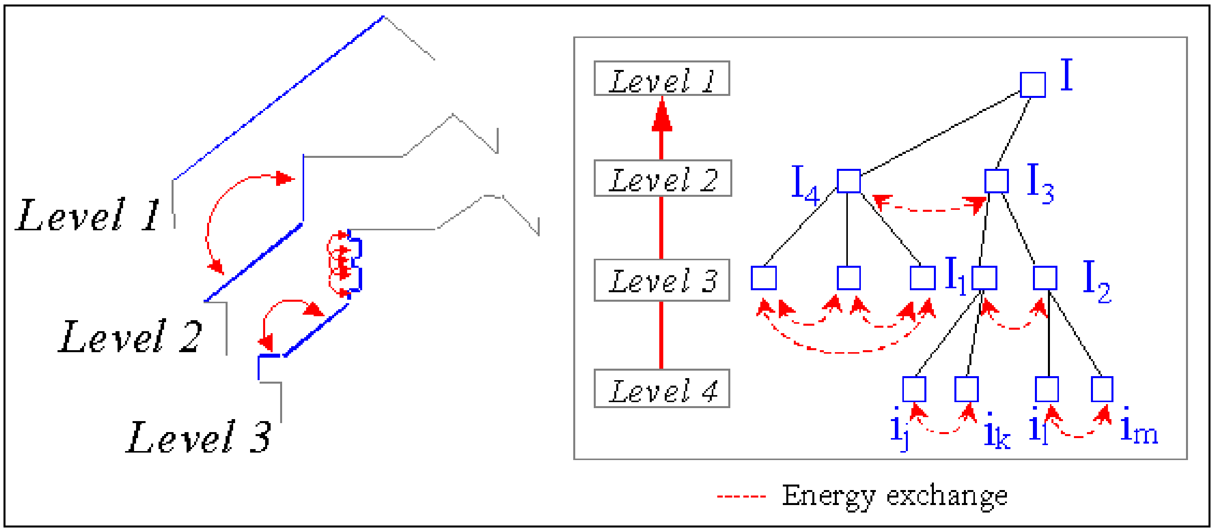

Figure 9: Internal Shooting

4 Results

This section provides some results for two test scenes *Scene1* and *Scene2*. It is made up of different kinds of objects approximated with different LODs. *Scene1* contains 1505 input polygons at the finest level and 192 at the coarsest level. *Scene2* is modeled with 10031 input polygons at the finest level and 2922 at the coarsest level. This scene contains 144 different objects (among which 116 are approximated with LODs) and 23 light sources (also approximated with LODs).

For each scene, the tables 1 and 2 give the run time for direct lighting (only the light source emits its energy), the total run time, the proportion of the total remaining unshot flux $P\Phi_{unshot}$ and the amount of memory required. In these tables *Without LOD* means that the finest LOD approximation is used for each object surface and classical HR is employed, while *With LOD* means that HR and LOD are combined as explained in this paper. Note that :

$$P\Phi_{unshot} = \frac{\sum_i \Phi_i^{unshot}}{\Phi_{lightSources}}$$

where i is the index of a root macrofacet.

From these tables we can see that the unshot flux $P\Phi_{unshot}$ is not affected by the use of LODs. We mean that the internal shooting and the energy distribution as implemented in our algorithm prevent from over or sub-estimating the energy transfer between geometric primitives such as macrofacet or facets.

The speed-ups obtained are encouraging.

In the figure 10, the left image has been generated *without LOD* while the right one has been obtained with our method. A top view of *Scene2* is given in figure 11 while figure 12 shows *Scene2* viewed from different viewpoints.

5 Conclusion

We have proposed a global illumination method combining hierarchical radiosity and LOD approximations of object surfaces. To our knowledge, it is the first method proposed so far. Unlike clustering techniques our approach accounts precisely for the orientations of the children nodes of the emitting and receiving macrofacets. When selecting a macrofacet for shooting, internal shooting followed by the computation of the associated contextual radiosity make the energy emission more precise. A macrofacet may be linked to another macrofacet, a facet, a patch or a surface element. The used shooting process depends on the linked primitives. The obtained speed-ups are significant and encouraging. We think that our approach can easily be extended for handling specular surfaces.

Scene1	Without LOD	With LOD	Speed up
Time for direct lighting	190 s	24 s	8
Total time	112 mn	36 mn	3.1
$P\Phi_{unshot}$	17.88 %	18.3 %	-
Memory	13 Mb	13 Mb	-

Table 1: Results for *Scene1*

Scene2	Without LOD	With LOD	Speed up
Time for direct lighting	356 mn	17 mn	20.9
Total time	633 mn	91 mn	7
$P\Phi_{unshot}$	19.76 %	19.83 %	-
Memory	31 Mb	31 Mb	-

Table 2: Results for *Scene2*

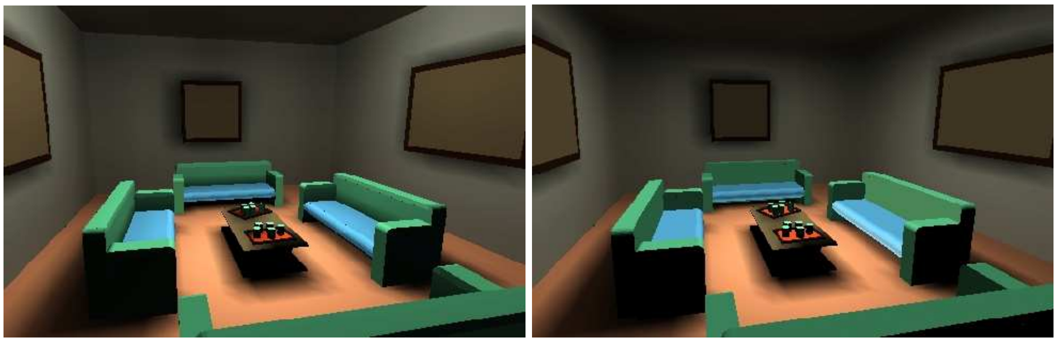

Figure 10: *Scene1*

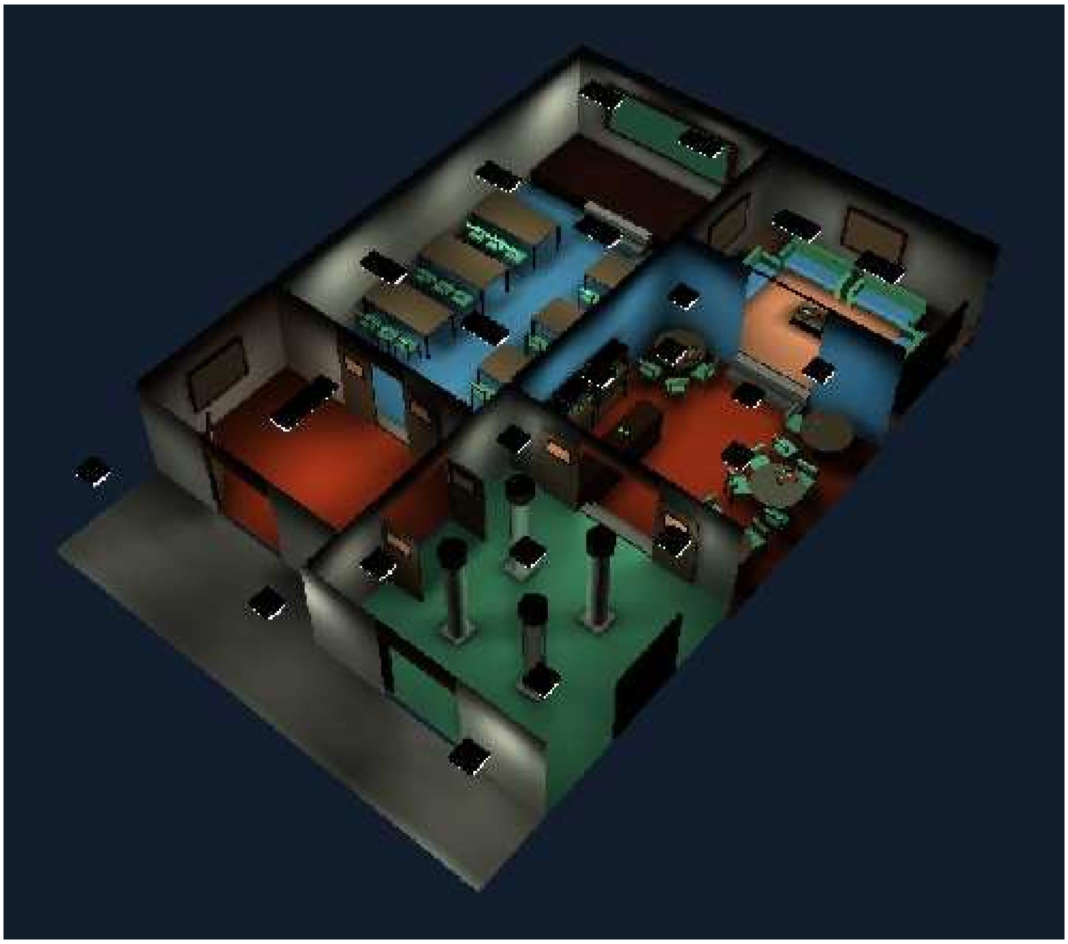

Figure 11: Top view of *Scene2*

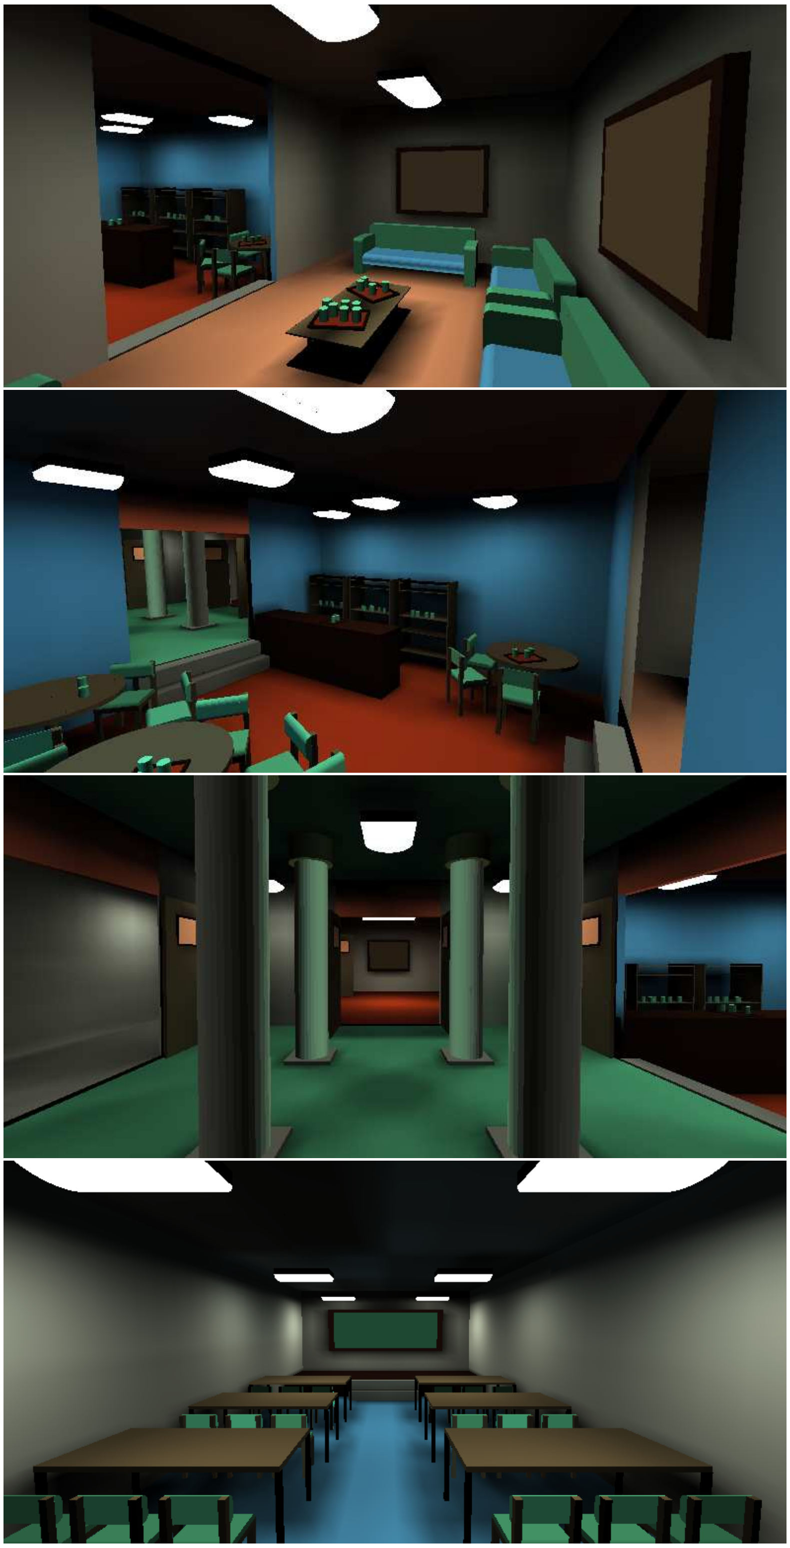

Figure 12: *Scene2*: Views of different rooms

References

[1] Pat Hanrahan & David Salzman & Larry Aupperle. A rapid hierarchical radiosity algorithm for unoccluded environments. *Computer Graphics*, 25(4):197–205, July 1991.

[2] J. M. Airey. Increasing update rates in the building walkthrough system with automatic model-space subdivision and potentially visible set calculation. 1990.

[3] C. Séquin T. Funkhouser, S. Teller and D. Khorramabadi. The uc berkeley system for interactive visulalization of large architectural models. *Presence*, 5(1):13–44, 1996.

[4] Seth Jared Teller. *Visibility Computations in Density Occluded Polyhedral Environments*. PhD thesis, University of California at Berkeley, 1992.

[5] D Meneveaux, K Bouatouch, E Maisel, and R Delmont. A new partitioning method for architectual environments. *Journal of Visualization and Computer Animation*, 9:195–213, 1998.

[6] Seth Teller & Celeste Fowler & Thomas Funkhouser & Pat Hanrahan. Partitioning and ordering large radiosity computations. In *Computer Graphics Proceedings, Annual Conference Series*, pages 443–450, 1994.

[7] De Berg M. and Dobrindt K. An levels of detail in terrains. Technical Report UU-CS-1995-12, University of Utrecht, The Netherlands, April 1995.

[8] Taylor D.C. and Barrett W.A. An algorithm for continuous resolution polygonalizations of a discrete surface. In *Graphics Interface*, 1994.

[9] Garland M. and Heckbert P. Surface simplification using quadric error metrics. In *Siggraph*, 1997.

[10] Hoppe H. Progressive meshes. In *Siggraph*, pages 99–108, 1996.

[11] Ronfard R. and Rossignac J. Full range approximation of triangulated polyhedra. In *Eurographics*, 1996.

[12] Schauffer G. and Sturzlinger W. Generating multiple levels of details from polygonal geometry models. In *Second Eurographics Workshop on virtual environment , realism and real time*, 1995.

[13] Brian Smits & James Arvo & Donald Greenberg. A clustering algorithm for radiosity in complex environments. In *Computer Graphics Proceedings, Annual Conference Series*, pages 435–442, 1994.

[14] F. Sillion. Clustering and volume scattering for hierarchical radiosity calculations. *5th Eurographics Workshop on Rendering*, 1994.

[15] S. Gibson and R.J. Hubbold. Efficient hierarchical refinement and clustering for radiosity in complex environments. *Computer Graphics Forum*, 15:297–310, 1996.

[16] F. Sillion and G. Drettakis. Feature-based control of visibility: A multiresolution clustering algorithm for global illumination. In *Siggraph'95*. ACM, 1995.

[17] F. Sillion, G. Drettakis, and C. Soler. *A Clustering Algorithm For Radiance Calculation in General*. Springer Verlag, Rendering Techniques (Proceedings of the 6th Eurographics rendering workshop), 1995.

[18] P.H. Christensen, D. Lischinski, E. Stollnitz, and D. Salesin. Clustering for glossy global illumination. *ACM Transaction On Graphics*, 16(1):3–33, January 1997.

[19] M. Stamminger, P. Slusallek, and H.P. Siedel. *Three Point Clustering for Radiance Computaions*. Springer Verlag, Rendering Techniques (Proceedings of the 9th Eurographics Workshop on Rendering), 1998.