

Using Levels of Detail to Speedup Radiosity Computation

Reynald Dumont, Kadi Bouatouch

► **To cite this version:**

Reynald Dumont, Kadi Bouatouch. Using Levels of Detail to Speedup Radiosity Computation. [Research Report] RR-3602, INRIA. 1999. inria-00073077

HAL Id: inria-00073077

<https://hal.inria.fr/inria-00073077>

Submitted on 24 May 2006

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.

*Using Levels of Detail to speedup Radiosity
Computation*

Reynald Dumont, Kadi Bouatouch

N° 3602

Janvier 1999

_____ THÈME 3 _____



*Rapport
de recherche*



Using Levels of Detail to speedup Radiosity Computation

Reynald Dumont, Kadi Bouatouch

Thème 3 — Interaction homme-machine,
images, données, connaissances
Projet SIAMES

Rapport de recherche n3602 — Janvier 1999 — 20 pages

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.

Key-words: Computer graphics, Global illumination, Hierarchical Radiosity, Levels of Detail

(Résumé : tsvp)

Utilisation de Niveaux de Détail pour l'Accélération du calcul de Radiosité

Résumé : Les méthodes de regroupement (Clustering en anglais) et de structuration sont des optimisations importantes de la radiosité hiérarchique. La première méthode réduit le nombre de liens générés mais est imprécise car les orientations des carreaux au sein des groupes de carreaux récepteur et émetteur sont mal prises en compte lors du calcul d'illumination. Quant à la seconde approche est très utile car elle divise de très grandes scènes en plusieurs cellules 3D et emploie des stratégies d'ordonnancement qui ne nécessite qu'un sous-ensemble de cellules à chaque étape du calcul d'illumination. Néanmoins, la complexité numérique reste toujours très élevée lorsque les cellules contiennent un nombre important d'objets de géométrie complexe. Le nouvel algorithme de radiosité hiérarchique présenté ici apporte des solutions à ces problèmes en utilisant les niveaux de détail (NDD) pour approcher la géométrie des surfaces des objets de la scène.

Mots-clé : Synthèse d'images, Illumination globale, Radiosité hiérarchique, Niveaux de détail

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 visibility 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, 7] 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 environment 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) and internal emission (section 3.3), distribution of the energy gathered by a receiving patch to its descendants in the associated LOD hierarchy (section 3.2). 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 (HLOD). 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 [8, 9] removes recursively one or several vertices from the exact geometry while the contraction approach [10, 11, 12] removes edges and replace them with new vertices. Conversely, the refinement approach [13] 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 [14, 15]. However, these techniques do not account precisely for the orientations of the facets within the emitting and receiving clusters during the energy exchanges.

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

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.

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 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 apex 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 apex is the center of the receiving macrofacet (or facet) R and subtented by the emitting macrofacet (or facet) E (figure 5). 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.

The algorithms describing our HR are given by the figures 2, 3 and 4.

```

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);
      }
    }
    E = 0;
  }

```

Figure 2: New Hierarchical Radiosity

```

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 ( $F_{R-E} < T_{LOD}$ ) {
   $B_c$  = Compute_Contextual_Radiosity( $N_E$ ,  $N_R$ );
  Link( $N_E$ ,  $N_R$ );
   $B_R$  = Shoot_energy( $N_E$ ,  $N_R$ );
  Distribute_Energy_To_Descendants( $N_E$ ,  $N_R$ ,  $B_R$ );
}
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 {
    if  $N_E$  is not a leaf of a LOD hierarchy {
       $B_c$  = Compute_Contextual_Radiosity( $N_E$ ,  $N_R$ );
    }
    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 3: Shooting process: first part, continued in figure 4

Let us explain the role of some functions invoked in figures 3 and 4. 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 .

```

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 {
    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( $N_E, N_R$ );
      Distribute_Energy_To_Descendants( $N_E, N_R, B_R$ ) ;
    } 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 4: Shooting process: continued

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 5) 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.




Figure 5: Form factors.

If one of the nodes (N_E or N_R) is a leaf of a LOD hierarchy, the other is systematically selected.

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 :

$$\Phi_E = A_E \cdot B_E, \quad (1)$$

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:

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

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:

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

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:

$$F_{e_k-R} = \frac{\cos \Theta_{e_k R} \cdot \cos \Theta_{R e_k} \cdot A_R}{\pi \cdot d_{e_k-R}^2}, \quad (4)$$

d_{e_k-R} being the distance separating the centers of e_k and R . The angles $\Theta_{e_k R}$ and $\Theta_{R e_k}$ are shown in figure 6.

Similarly we can write:

$$F_{E-R} = \frac{\cos \Theta_{ER} \cdot \cos \Theta_{RE} \cdot A_R}{\pi \cdot d_{E-R}^2}, \quad (5)$$




Figure 6: Angles

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

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

assuming $d_{e_k-R} \approx d_{E-R}$ and $\Theta_{R e_k} \approx \Theta_{RE} \forall k$ (figure 6). B_c has to be multiplied by a visibility factor $V_{e_k R}$ between e_k and R .

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

$$B_R = \rho_R \cdot F_{R-E} \cdot B_c, \quad (7)$$

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 .

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:

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

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

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

$$F_{E-r_k} = \frac{\cos \Theta_{E-r_k} \cdot \cos \Theta_{r_k E} \cdot A_{r_k}}{\pi \cdot d_{E-r_k}^2}, \quad (10)$$

By dividing equations 8 by 9, substituting equation 10, assuming $\Theta_{E-r_k} \approx \Theta_{E-R}$, $d_{E r_k} \approx d_{ER}$, $\forall k$ (figure 6) and adding the visibility factor $V_{E r_k}$, we obtain:

$$B_{r_k} \approx B_R \cdot \frac{\cos \Theta_{r_k E}}{\cos \Theta_{RE}} \cdot \frac{\rho_{r_k}}{\rho_R} \cdot V_{E r_k} \quad (11)$$

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 8. Suppose a root macrofacet I has been chosen for shooting its energy (see figure 2). 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*. Let us explain the internal shooting operation with the help of figure 8. 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 7. To make the *Internal Shooting* faster, two nodes of different parents (such as i_j and i_l in figure 8) 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 8.

4 Results

This section provides some results for two test scenes *Scene1* and *Scene2*. *Scene2* contains more objects than *Scene1*. They are made up of different kinds of objects given in table 1. In this table, when the name of an object is followed by *LOD*, that means the surface of the object is represented by a LOD hierarchy. For these scenes, each LOD hierarchy has a maximum depth equal to 2. In other words, at the coarsest level (depth level equal to 1) the surface of an object is represented by a certain number of root macrofacets, each of them is in turn represented at the finest level (depth level equal to 2) by facets (see table 2).

The figures 9 and 10 are wireframe images of *Scene1* for the two depth levels.

For each scene, the tables 3 and 4 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.

Even for scenes of moderate complexity, the speed-ups obtained are encouraging.

```

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  $E$  ( $C_E \neq C'_E$ ) {
      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 7: Internal Shooting

In the figures 11 and 12, the upper image has been generated *without LOD* while the middle one has been obtained with our method. The lower image is the difference between these two images. This difference image is given in grey scale, the larger the intensity of a pixel, the lower the difference. The figure 13 shows a closer view of *Scene1*. Look at the borders of the armrests of the sofas, they still are round even though we used two level LOD hierarchies.

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




Figure 8: Internal Shooting

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. Even for the simple test scenes used in this paper, the obtained speed-ups are significant. For more complex scenes we expect higher speed-ups.

	Number of input polygons (facets) at the finest LOD	Number of objects in the scene	Kinds of objects in the scene
Scene1	1011	13	1 light source (LOD) 4 sofas (LOD), 1 table 1 frame (LOD), 1 tray and 5 cups (LOD)
Scene2	1505	21	1 light source (LOD) 4 sofas (LOD), 1 table 4 frames (LOD), 2 trays (LOD) and 10 cups (LOD)

Table 1: Description of the scene tests.

	Number of root macrofacets, depth 1	Number of facets, depth 2
Light source	6	37
Sofa	21	138
Tray	1	30
Cup	6	64
Frame	1	30

Table 2: Two level LOD hierarchies

Scene1	Without LOD	With LOD	Speed up
Time for direct lighting	140 s	16 s	8.75
Total time	81 mn	31 mn	2.7
$P\Phi_{unshot}$	18.8 %	18.3 %	-
Memory	11 Mb	11 Mb	-

Table 3: Results for *Scene1*

Scene2	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	11 Mb	11 Mb	-

Table 4: Results for *Scene2*




Figure 9: Scene1: Coarse LOD




Figure 10: Scene1: Finer LOD






Figure 11: Scene1



INRIA

Figure 12: Scene2



RR n3602

Figure 13: Close view of Scene1

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*. PhD thesis, University of north Carolina at Chapel hill, 1990.
- [3] C. Séquin T. Funkhouser, S. Teller and D. Khorramabadi. The uc berkeley system for interactive visualization 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 architectural environments. *To appear in Visualization and Computer Animation*, 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] D. Meneveaux, K. Bouatouch, and E. Maisel. Memory management schemes for radiosity computation in complex environments. In *CGI'98 Hannover*, June 1998.
- [8] 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.
- [9] Taylor D.C. and Barrett W.A. An algorithm for continuous resolution polygonalizations of a discrete surface. In *Graphics Interface*, 1994.
- [10] Garland M. and Heckbert P. Surface simplification using quadric error metrics. In *Siggraph*, 1997.
- [11] Hoppe H. Progressive meshes. In *Siggraph*, pages 99–108, 1996.
- [12] Ronfard R. and Rossignac J. Full range approximation of triangulated polyhedra. In *Eurographics*, 1996.
- [13] Schaffner 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.
- [14] 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.
- [15] F. Sillion. Clustering and volume scattering for hierarchical radiosity calculations. *5th Eurographics Workshop on Rendering*, 1994.



Unit e de recherche INRIA Lorraine, Technop le de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS L ES NANCY
Unit e de recherche INRIA Rennes, Irisa, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unit e de recherche INRIA Rh ne-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unit e de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unit e de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

 diteur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399