Managing Geometry Complexity for Illumination Computation of Cultural Heritage Scenes

Xavier Granier, Romain Pacanowski, Christophe Schlick

Tiré-à-part des Actes du colloque
Virtual Retrospect 2007

Pessac (France) 14, 15 et 16 novembre 2007
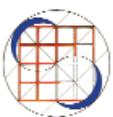
X. Granier, R. Pacanowski, C. Schlick,

AUSONIUS

archéovision

Managing Geometry Complexity for Illumination Computation of Cultural Heritage Scenes

Xavier Granier granier@labri.fr
Romain Pacanowski pacanows@labri.fr
Christophe Schlick schlick@labri.fr

INRIA Bordeaux Sud-Ouest – LaBRI – Université de Bordeaux
Domaine Universitaire - 351, cours de la Libération
33405 Talence Cedex - France

http://iparla.labri.fr/members/

Résumé : Pour l'étude du patrimoine, de plus en plus d'objets 3D sont acquis par le biais de scanners 3D [Levoy 2000]. Les objets ainsi acquis contiennent de nombreux détails et fournissent une très grande richesse visuelle. Mais pour les afficher, leur très grande complexité géométrique nécessite l'utilisation d'algorithmes spécifiques. Nous montrons ici comment simplifier ces objets pa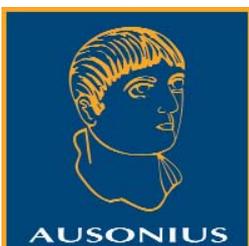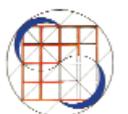r un maillage de faible résolution et une collection de cartes de normales [Boubekeur 2005] pour préserver les détails. Avec cette représentation, nous montrons comment il est possible de calculer un éclairement réaliste à l'aide d'une grille et de données vectorielles [Pacanowski 2005]. Cette grille permet de capturer efficacement les basses fréquences d'un éclairement indirect. Nous utilisons des textures 3D (pour des gros objets) et potentiellement des textures 2D (pour les objets quasi-plans) afin de stocker un nombre prédéterminé de vecteurs d'irradiance. Ces grilles sont calculées au cours d'un pré-calcul à l'aide de n'importe quelle méthode stochastique de calcul d'éclairement global. Pour l'affichage, l'éclairement indirect dû à la grille est interpolé au sein de la cellule associée à la position courante, fournissant ainsi une représentation continue. De plus, cette approche vectorielle permet une plus grande robustesse aux variations locales des propriétés géométriques de la scène.

Mots-clés : scanner 3D, maillage, texture.

Abstract : For cultural heritage, more and more 3D objects are acquired using 3D scanners [Levoy 2000]. The resulting objects are very detailed with a large visual richness but their geometric complexity requires specific methods to render them. We first show how to simplify those objects using a low-resolution mesh with its associated normal maps [Boubekeur 2005] which encode details. Using this representation, we show how to add global illumination with a grid-based and vector-based representation [Pacanowski 2005]. This grid captures efficiently low-frequency indirect illumination. We use 3D textures (for large objects) and 2D textures (for quasi-planar objects) for storing a fixed set of irradiance vectors. These grids are built during a preprocessing step by using almost any existing stochastic global illumination approach. During the rendering step, the indirect illumination within a grid cell is interpolated from its associated irradiance vectors, resulting in a smooth everywhere representation. Furthermore, the vector-based representation offers additional robustness against local variations of geometric properties of the scene.

Keywords : 3D scanner, mesh, texture.

INTRODUCTION

For archeology and cultural heritage, global illumination can be used in two main contexts. The first one is the virtual reconstruction and the validation of different hypothesis. In this context of physical plausibility, the algorithms have to be sufficiently accurate in order to be confident in the computed solutions. The second one is the public presentation of cultural heritage scenes, like for a virtual tour. In this context, the main goal is to reach visual plausibility in order to get convincing results. In this paper, we focus on this second context.

Furthermore, the acquisition systems nowadays provide very large 3D models with a lot of geometric details [Levoy 2000], increasing the resulting visual richness but also, the required computation time of global illumination solutions. Traditionally, caching schemes are used to reduce the complexity of indirect illumination. Unfortunately, their efficiency generally diminishes with the increasing geometric complexity. Small surface details require a densely sampled cache in order to capture the subtle illumination variations induced by local variations of the normal (e.g., normal maps, micro-polygons).

For complex models, Tabellion and Lamorlette [Lamorlette 2004] use a simplified geometry for the caching step. For the final rendering, the full resolution is used in order to provide visually rich solutions. Boubekeur *et al.* [Boubekeur 2005] have shown that a two-level representation can be sufficient to preserve the rich appearance of an object : a simplified geometry for the preservation of the global shape, and a

normal-map for the local details. If the geometry is simplified, the appearance is still highly detailed, leading to high visual quality. As previously said, such models still require a densely sampled cache with current approaches.

We present some grid structures well-suited for such detailed objects, and more precisely a volumetric grid for very complex ones and a 2D grid for quasi-planar surfaces. From the grids and their associated interpolation schemes presented in this paper, we guarantee a smooth reconstruction everywhere in the 3D scene, making the cache representation directly accessible for the final rendering. Our approach is based on irradiance vector [Arvo 1994]. From the irradiance vector, our representation inherits robustness against local variations of both photometric properties (diffuse component of reflectance) and geometric properties (surface normal vectors). Finally, our overall structure has low memory requirements and thus increases the scalability of the method.

GEOMETRIC SIMPLIFICATION

Before any lighting computation, and in order to reduce the computational complexity to fit the 3D model with available memory, we need to reduce the number of polygons without losing the visual richness of original models. For this purpose, we use a mixed representation, based on a simplified geometry for the definition of the overall shape and topology, and normal maps for preserving their local variations on the surface. Since the lighting mostly depends on the normal, this approach preserves the objects' appearance.

The large size of 3D models is not only a problem for the lighting simulation. We also have to deal with the fact that the main memory would not be sufficient to load the complete model for the simplification process. We need to use out-of-core solutions. We thus decomposed our algorithm in four steps (more detailed presentation of this approach can be found in the paper [Boubekeur 2005]) :

• We perform an out-of-core simplification of the huge model using a uniform resampling [Lindstrom 2000, Boubekeur 2005]. A non-uniform and adaptive solution [Boubekeur 2006] could be also used.

• The resulting simplified point set is quickly converted into a triangle mesh [Boubekeur 2005, 2006] and organized in a bounding box tree-hierarchy.

• Each leaf of this tree hierarchy is associated with a quadrilateral texture. All the points of the original model are thus streamed through this tree and distributed to their corresponding leaves, where the point normal is projected into the associated texture. This streaming process is the key step of our technique, as it allows us to handle large models with limited memory. At the end of this step, each triangle of the low-resolution mesh is associated with a high-resolution (and possible sparse) normal map.

• Since there is no guaranty that each texel of the normal map corresponds to an existing point and thus an existing

110

normal, the resulting map can contains some holes. We thus use a diffusion algorithm, to get a continuous normal field, interpolating the original normals of the huge model.

Thanks to this out of core simplification, the memory requirement is largely reduced while preserving most of the appearance details (fig. 1). Lighting can be computed on such a representation, but an adapted approach has to be developed.

ILLUMINATION COMPUTATION FOR COMPLEX GEOMETRY

In computer graphics, illumination computation has been extendedly studied. Although direct illumination of an object is easily computed by summing the contribution of every light source seen from the object, the indirect lighting computation is a more challenging task. Indeed, when considering indirect illumination, every object in the scene plays potentially the role of a light source. Therefore, many algorithms (cf. [Dutré 2006]) have been developed over the last decade to perform indirect illumination computation. Among these algorithms, some of them use specific precomputed structures that cache the indirect illumination. As pointed out by Tabellion and Lamorlette [Tabellion 2004], caching efficiency generally diminishes with the increasing geometric complexity of a 3D scene. Previous approaches [Wilmott 1999] based on Hierarchical Radiosity methods use simplification methods to reduce the size of the models. However, these methods are difficult to implement and are not robust to strong geometric variations.

In order to reduce this problem, this paper introduces a volumetric representation for indirect illumination based on irradiance vectors. From the irradiance vector, our representation inherits robustness against local variations of both photometric properties (diffuse component of reflectance) and geometric properties (surface normal vectors). From the volumetric structure and the associated interpolation scheme presented in this paper, we guarantee a smooth reconstruction everywhere in the 3D scene, making the cache representation directly accessible for the final rendering. Finally, our overall structure has low memory requirements and thus increases the scalability of the method. Furthermore, these qualities make this approach suitable for interactive hardware rendering.

Irradiance Vector Grid (IVG)

Our structure is based on an axis-aligned uniform rectangular 3D grid, divided into $N_i \times N_j \times N_k$ voxels. At each vertex $v^{i,j,k}$ of the grid (where $i \in [0, N_i]$ $j \in [0, N_j]$, $k \in [0, N_k]$), six irradiance vectors are stored, one for each main direction ($\pm x | \pm y | \pm z$). Actually we store an irradiance matrix, as one vector is used for each color channel. In the remaining of this paper, we will note $I_\delta^{i,j,k}$ the irradiance vector stored at $v^{i,j,k}$ in the direction δ where $\delta = \pm x | \pm y | \pm z$.

Fig. 1. Original models and their simplified versions.
Notice the largely reduce size and the preserved appearance.

Irradiance Vector

For a given wavelength, the irradiance vector $I_n(p)$ is defined for a point p with normal $n(n_x, n_y, n_z)$ as :

$$I_n(p) = \int_{\Omega_n} L(p \leftarrow \omega_i)\, \omega_i\, d\omega_i \, ,$$

where $L(p \leftarrow \omega_i)$ represents the incident radiance at p from direction ω_i, $d\omega_i$ the differential solid angle sustained by ω_i and Ω_n the hemisphere centered at p oriented toward n. The irradiance vector stores radiometric and geometric information and is directly related to the diffusely reflected radiance :

$$L_r(p \rightarrow \omega_o) = \frac{\rho_d}{\pi}\, \langle I_n(p), n \rangle \, ,$$

where ρ_d is the diffuse BRDF and $\langle \ , \ \rangle$ denotes a dot product. The main benefits of irradiance vectors compared to irradiance is that for a local variation of the normal, the reflected radiance can be adjusted, making this representation more geometrically robust. Finally, irradiance vectors are precomputed by using any classical algorithm such as Photon-Tracing or Monte-Carlo sampling.

Irradiance Vector Interpolation

In order to compute smooth indirect illumination, we interpolate an irradiance vector for each point p with normal n that needs to be shaded. This interpolation is performed in two successive steps : a spatial interpolation according to p and then a directional interpolation according to n. In the first step, the irradiance vector $I_\delta(p)$ is obtained by spatial interpolation of the irradiance vectors $I_\delta^{i,j,k}$ stored at the grid vertices surrounding point p. The interpolation is only done for three out of the six possible directions of δ. The choice between $\pm x$ (resp. $\pm y$ and $\pm z$) is done according to the sign of n_x (resp. n_y and n_z). Trilinear or tricubic interpolation approximate satisfactory smooth results for spatial interpolation. In the second step, the final interpolated irradiance vector $I_n(p)$ is obtained by remapping the three spatially interpolated irradiance vectors according to the normal direction n at point p :

$$I_n(p) = I_x(p)\, n_x^2 + I_y(p)\, n_y^2 + I_z(p)\, n_z^2 \, .$$

Hardware Implementation

The great benefit of using a 3D regular grid is that the data structure can be straightforwardly uploaded on GPU as a 3D texture. In our case, the interpolated irradiance vectors are simply used by the fragment shader as additional light sources that are meant to encode indirect illumination. Remember that each grid vertex holds one irradiance vector I_λ per color channel for each of the six δ directions, i.e., $3 \times 3 \times 6 = 54$ floating point numbers. To reduce the number of texture fetches which may be costly in current graphics hardware, we compress the 3 irradiance vectors as one (r, g, b) color and one direction d :

$$r = |I_R|, \qquad g = |I_g|, \qquad b = |I_b|, \qquad d = \frac{I_R + I_G + I_B}{|I_R + I_G + I_B|} \, .$$

These two vectors are encoded in two 16 bits 3D textures, and therefore the information for the six δ directions requires 12 3D textures.

RESULTS

To illustrate the geometric robustness of our caching structure we compare the indirect illumination obtained when the structure is precomputed with a model at full resolution (fig. 2) vs. at low resolution (fig. 3). For both figures the indirect illumination is stored in a 13x14x10 grid for the dragon model and in a 10x20x10 grid for the buddha model. As illustrated in fig. 4 the perceptive difference is very low with only one 1 unity in the uniform perceptive Lab color space. Furthermore, the precomputation is 3 times faster when using the low resolution model. The same scene rendered with direct and indirect illumination is shown in fig. 5. Furthermore, fig. 6 shows a more complex scene where most illumination is indirect.

Fig. 2. Indirect illumination Visualization on high resolution models (1M triangles).

Fig. 3. Indirect illumination Visualization on low resolution models (100K triangles). The precomputation is 3 times faster than in fig. 2.

Fig. 4. Lab difference between fig. 2 and fig. 3. Brighter color means bigger error.

Fig. 5. Final image computed with direct and indirect illumination. The overhead introduced to include indirect illumination with our grid is very small (5% of the total rendering time).

CONCLUSION AND FUTURE WORK

To summarize, we presented a method to compute indirect illumination in a scene with complex models. Our method proceeds in two steps :

A geometry simplification process that preserves the original normals and store them in a texture (normal map).

An appropriate structure that stores the indirect illumination. This structure is robust to geometric variations of the underlying geometry.

Our coming work is to improve the precomputation of our illumination structure. In fact, it is still difficult to handle some illumination configuration. Therefore, we would like to introduce an adaptive criterion to refine the structure where illumination varies too quickly.

References

Arvo, J. (1994): "The Irradiance Jacobian for Partially Occluded Polyhedral Sources", in: *ACM SIGGRAPH'94*, 343-350.

Boubekeur, T., F. Duguet and C. Schlick (2005): "Volume-Surface Trees", in: *EUROGRAPHICS VAST 2005*, 75-82.

Boubekeur, T., W. Heidrich, X. Granier and C. Schlick (2006): "Rapid Visualization of Large Point-based Surfaces", in: *Computer Graphics Forum (Proc. EUROGRAPHICS 2006)*, 25(3), 399-406.

Dutré P., K. Bala and P. Bekaert (2006): *Advanced Global Illumination*, A.K. Peters, Ltd.

Levoy, M., K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, Ginsb., J. Shade and D. Fulk (2000): "The Digital Michelangelo Project: 3D Scanning of Large Statues", in: *ACM SIGGRAPH '00*, 131–144.

Lindstrom, P. (2000): "Out-of-Core Simplification of Large Polygonal Models", in: *ACM SIGGRAPH '00*, 259-262.

Pacanowski, R., X. Granier and C. Schlick (2005): "Nouvelle représentation directionnelle pour l'éclairage global", in : *Actes des 18ᵉ Journées de l'Association Française d'Informatique Graphique (AFIG)*.

Tabellion, E. and A. Lamorlette (2004): "An Approximate Global Illumination System for Computer Generated Films", in: *ACM SIGGRAPH '04,* 469-476.

Willmott, A. J., P. Heckbert and M. Garland (1999): "Face Cluster Radiosity", in: *Eurographics Rendering Workshop*, 293–304.