The 9th International Symposium on Virtual Reality, Archaeology and Cultural Heritage VAST (2008)
M. Ashley, S. Hermon, A. Proenca, and K. Rodriguez-Echavarria (Editors)

Workshops,
Short and Project Papers

Compact Structures for Interactive Global Illumination on Large Cultural Objects

Romain Pacanowski†[a,c] Mickaël Raynaud†[a] Julien Lacoste‡[b]
Xavier Granier†[a] Patrick Reuter†[a] Christophe Schlick†[a] Pierre Poulin§[c]

[a] INRIA - Bordeaux University (France) [b] LIUPPA - University of Pau (France)
[c] LIGUM, Dept. I.R.O., Université de Montréal (Canada)

Abstract

Cultural Heritage scenes usually consist of very large and detailed 3D objects with high geometric complexity. Even the raw visualization of such 3D objects already involves a large amount of memory and computation time. When trying to improve the sense of immersion and realism by using, global illumination techniques the demand on these resources becomes prohibitive.

Our approach uses regular grids combined with a vector-based representation to efficiently capture low-frequency indirect illumination. A fixed set of irradiance vectors is stored in 3D textures (for complex objects) and in 2D textures (for mostly planar objects). The vector-based representation offers additional robustness against local variations of the geometry. Consequently, the grid resolution can be set independently of geometric complexity, and the memory footprint can therefore be reduced. The irradiance vectors can be precomputed on a simplified geometry. For interactive rendering, we use an appearance-preserving simplification of the geometry. The indirect illumination within a grid cell is interpolated from its associated irradiance vectors, resulting in an everywhere-smooth reconstruction.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and RealismShading; Radiosity; RayTracing;

---

1. Motivation

Global illumination can prove particularly useful for archeology and cultural heritage applications. For instance, the validation of different archaeological hypotheses requires physical plausibility for their virtual reconstructions, with sufficient accuracy in the computed solutions to give confidence to the archaeologists. Another example of application consists in virtual tours for public presentations of cultural heritage sites or scenes. There, visual plausibility under scientific control and validation is sufficient for an improved sense of immersion and a more natural lighting. Our project focuses on this second scenario.

Thanks to nowadays acquisition systems, archaeologists can work with highly detailed models [LPC*00]. However, such extensive models can be difficult to render interactively. In applications where global illumination has to be taken into account, this task is even more difficult. While classical real-time applications can rely on simplified geometry, it is very important in the context of cultural heritage to preserve as many information as possible because they can convey important details about its origin or history, such as weathering, aging signs, or tools marks produced during its creation. Appearance preserving simplification techniques have been proposed to reduce geometric complexity while preserving visual details, but currently, there is no global illumination method well adapted to deal with such representations.

In this paper, we introduce two regular grid structures that are well suited for using global illumination on such detailed objects with a low memory footprint. More precisely, we in-

---

† {pacanows | raynaud | granier | reuter | schlick}@labri.fr
‡ julien.lacoste@pau.fr
§ poulin@iro.umontreal.ca

troduce a 3D regular grid for most objects, and a special 2D regular grid for quasi-planar surfaces. These grids and their associated interpolation schemes guaranty a smooth reconstruction of indirect lighting everywhere in the 3D scene. Furthermore, by using a vectorial representation based on *irradiance vectors* [Arv94], our approach is robust against local variations of both photometric properties (diffuse component of reflectance) and geometric properties (surface normals). These combined properties lead to a more geometry-independent approach and thus, to an overall structure that has low memory requirements and increased scalability.

2. Structures for Indirect Lighting

For reducing the dependency on the geometry, we base our rendering system on the precomputation of indirect lighting in grids of directional irradiance vectors [PRG*08], denoted $\boldsymbol{I_n}(\boldsymbol{p})$. These values are precomputed on a set of grids associated with the different objects of the scene: 3D grids for large objects, 2D grids for quasi-planar ones.

For a diffuse surface of albedo ρ_D, the reflected radiance $L_r(\boldsymbol{p} \rightarrow \boldsymbol{\omega}_o)$ at a point $\boldsymbol{p}$ with normal $\boldsymbol{n}$ is computed as:

$$L_r(\boldsymbol{p} \rightarrow \boldsymbol{\omega}_o) = \frac{\rho_D}{\pi} \left(\boldsymbol{I_n}(\boldsymbol{p}) \cdot \boldsymbol{n}\right) \tag{1}$$

where $\cdot$ denotes the dot product. With this vectorial representation and for a local variation of a surface normal, the reflected radiance can be adjusted, making this representation more **geometrically robust**. By precomputing this value $I_{\boldsymbol{\delta}}$ for the six main directions $\boldsymbol{\delta} = \pm\boldsymbol{x}| \pm \boldsymbol{y}| \pm \boldsymbol{z}$, and with a directional interpolation, we can evaluate an irradiance vector for any normal $\boldsymbol{n} = (n_x, n_y, n_z)$:

$$\boldsymbol{I_n}(\boldsymbol{p}) = \boldsymbol{I_x}(\boldsymbol{p})\, n_x^2 + \boldsymbol{I_y}(\boldsymbol{p})\, n_y^2 + \boldsymbol{I_z}(\boldsymbol{p})\, n_z^2\,, \tag{2}$$

where the choice between $\pm\boldsymbol{x}$ (resp. $\pm\boldsymbol{y}$ and $\pm\boldsymbol{z}$) is done according to the sign of n_x (resp. n_y and n_z).

By precomputing theses values on a regular grid, the data structure is directly supported by the GPU as 2D or 3D textures. We also directly benefit from the spatial interpolation in a cell, increasing the geometric robustness.

As a cultural heritage scene can be composed of many objects with a lot of details, its final size can simply be too large to be stored into core memory. Thanks to their geometric robustness, the directional irradiance vectors can be precomputed with a standard path tracer on a simplified geometry. This process can be parallelized easily and does not require any additional intermediate structure.

For 2D grids, if the 3D position of a grid vertex lies on the associated surface, only half of the directions are accessible. A small translation in the direction of the normal is thus applied. Unfortunately, these displaced vertices, or the vertices of the 3D grids can also reside inside objects: this will result in null irradiance vectors and therefore unwanted shadowed areas. To solve this problem, in a first step, we simply ignore those vertices at the precomputation stage and tag them as invalid. Then, in a second step, we assign to the invalid vertices the average value of their neighboring valid vertices. The unassigned vertices will never be used for the linear interpolation since they are too deep inside the 3D objects.

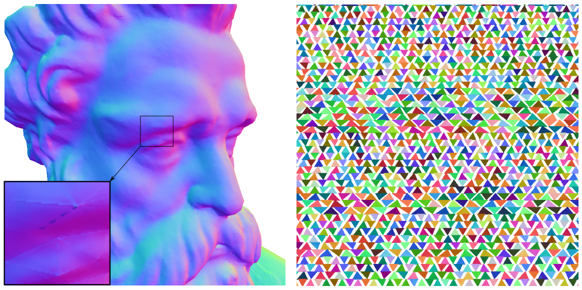
Neptune's head Packed normal maps

Figure 1: *Unparameterized texture atlas. Without global parameterization, the normal maps packed into a texture atlas (left). The white regions are un-defined surface normals. These resulting discontinuities are thus introduced on surface normals (right).*

3. Structures for Geometric Details

Once scanned, archaeological data contains a large amount of details that any visualization algorithm has to deal with. This **original geometry** can be too large, to offer interesting performance when precomputing the global illumination. However, since indirect illumination is a slow varying function [WRC88], it does not depend on all the original details. A **simplified geometry** is thus largely sufficient and it will greatly speed up the preprocess. This coarse level can also be used for distant objects. However, for closer viewpoints, all tiny details can convey important visual information. These details are re-introduced in the intermediate level of details for interactive rendering.

Most of the original details can be represented by a normal map with sufficient visual accuracy. We use **simplified meshes with normal maps** for the final rendering step. This representation enables real-time performance while preserving the visual richness of an object. Since global parameterization of a large object can be difficult to compute, we rely on a volumetric and parameterization-free representation: the appearance-preserving octree (APO) [LBJS07]. An APO encodes a normal map in an octree texture that can be transformed in an atlas of 2D textures. Each triangle of the simplified geometry is projected on the texture plane, and the same fragment shader associates a color to each triangle fragment. Similar to the triangle packing [CMSR98], we do not *shear* the triangles to fit them properly in the atlas. We then lose some small space between triangles in the texture, but distortion while mapping are avoided. With a texture atlas, discontinuities appear at the edges of triangles (cf. Figure 1).

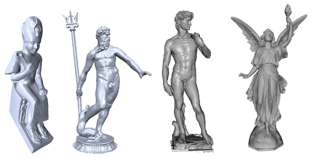

	Ramses	Neptune	David	Lucy
Full/Simplified	1.6M/80K	4M/100K	7.2M/100K	15M/150K
APO	4 min	6 min	4 min	7 min
Texture	6 min	5 min	10 min	10 min

Figure 2: *Size (in polygons) of the full and simplified objects with their preprocessing time for APO and texture atlas generation (done on an AMD 3500+ processor with 2 GB, and on a GeForce 8800 GTS). The preprocessing uses up to 4 GB of memory so the swapping time has been removed from the precomputation.*

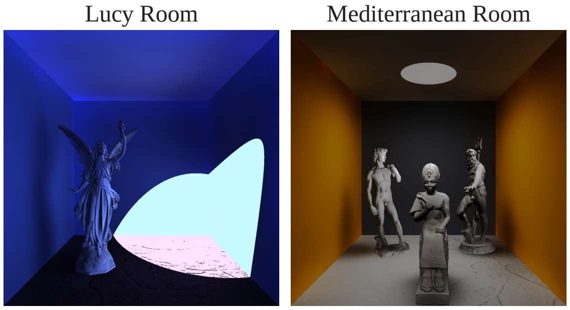

Figure 3: *The test scenes. The first scene is dedicated to Lucy statue. The floor consists in a normal mapped quad and a spot-light is directed toward the back-right corner. The second scene is dedicated to statues gathered from around the Mediterranean sea from different places and periods. It consists in three statues, a normal-mapped floor and a spot light directed toward the ceiling.*

4. Results

Unless mentioned otherwise, all precomputation of illumination grids and rendering have been done on an Intel Q6600 (4 GB of RAM) with an Nvidia GeForce 8800 GTX. We experiment our grids structure for indirect illumination with two scenes (cf. Figure 3) where the illumination on the main objects is thus mostly indirect. In all scenes, we use 2D grids to capture the indirect illumination and shadows for the walls, floor, and ceiling, while the 3D grids capture them for the archeological objects. For each scene and as described in previous sections, we use a simplified mesh to precompute the indirect illumination, and this mesh with normal mapping for the final interactive rendering.

This approach allows to speed up precomputation without loosing too many features in the illumination. To validate this point, we compare (cf. Figure 4) the indirect irradiance obtained when using grids precomputed with the simplified (cf. Figure 4(b)) and full geometries (cf. Figure 4(c)). As shown in Figure 4(d) the average difference in perceptually uniform *Lab* color space is very low. The maximum possible

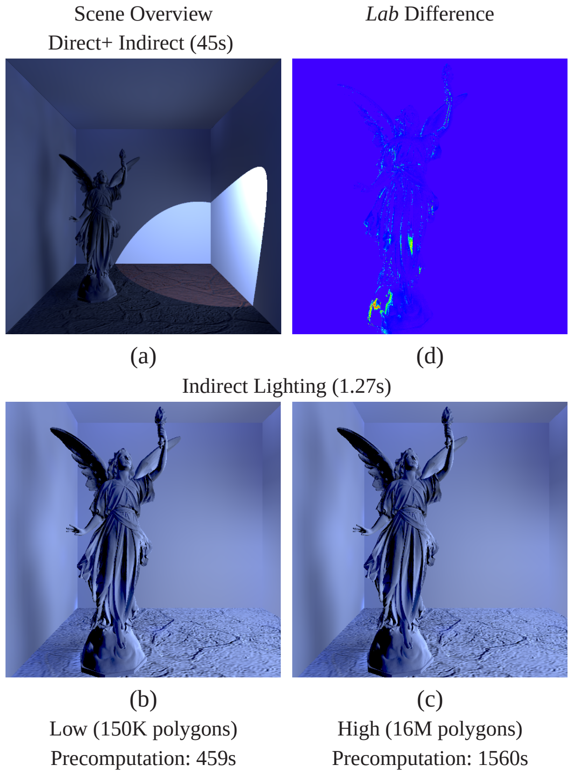

Figure 4: *Precomputation on the full versus simplified geometry. The resolution of the 3D grid used to capture the indirect illumination around this statue is* $8 \times 16 \times 8$*. In (b) and (c), we use software rendering to show the indirect irradiance from the grid precomputed on respectively low (b) and high (c) geometry. Computing the difference in Lab color space (d) between (b) and (c) shows that the maximum error is 69 (15% of maximum possible error) and is very localized. The average error is 0.4 (0.09%).*

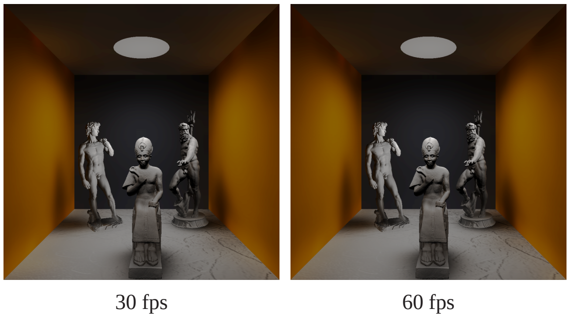

Figure 5: *APO (left) versus Texture Atlas (right). No noticeable differences are visible on statues. Note the indirect reflection of the wall on the left and right statues and on the left and right part of the floor. The indirect lighting correctly follows the varation of normal vectors.*

	Lucy Room	Mediterranean Room
Full/Simpl. geom.	15M/150K	12.8M/280 K
APO	21 MB	94 MB
Texture Atlas	45 MB	108 MB
3D grids	1 - 8×16×8	3 - 16×16×16
	144 KB	1.7 MB
2D grids	6 - 32×32	6 - 32×32
	264 KB	264 KB
Lighting Precomp.	55 min	260 min

Table 1: *Size of the different structures. The geometry size corresponds to the number of polygons. We present in parentheses the number of 2D and 3D grids used in each scene, their resolution, and the total amount of memory. Note how small the latter is compared to the size of the simplified geometry. We use 16K paths for each grid vertex.*

	APO	Texture Atlas	Full geom.
Lucy Room	35fps	60fps	3-8fps
Mediterranean Room	30fps	60fps	5-8fps

Table 2: *Average rendering frame-rate for an 800×800 pixels viewport. For a given scene, the same grids are used for indirect illumination with the different representation of the geometry.*

error reaches 15%, but mostly in darker areas where it is less noticeable.

Archeological objects are represented either with an APO or an atlas version to preserve the important details. As shown in Figure 2, this precomputation step is quite time consuming minutes but is only done once, and allows thereafter interactive rendering. As shown in Table 2 rendering frame-rates are higher with texture atlases than APOs because the octree traversal needs more GPU instructions than a single texture lookup used with atlases. However, as mentioned earlier, the quality of the geometric details are better represented with APOs, and accurate texture atlases would require difficult global parameterizations of the unsimplified objects.

Nevertheless, our illumination grids may be used with any geometric representation or level of details. We demonstrate this in Figure 5. The left image of Figure 5 shows the direct and indirect illumination with the APO structure while the right image with the texture atlas. The difference between the two images is barely noticeable on statues. Notice that one may use other representations such as displacement or relief mappings with our illumination grids. This approach allows users to work independently on the representations for the geometry and for the indirect illumination.

Moreover, bare in mind that once the illumination has been precomputed our rendering system allows archeologists to navigate in 3D scenes and visualize the objects interactively whereas this would be unattainable with full geometries (cf. Table 2). Note also that the overhead due to our representation of the indirection illumination is low compare to the required memory for the geometry (cf. Table 1).

5. Conclusion and Future Work

In this paper, we have presented a representation for indirect illumination. Based on geometry-independent grids (2D for quasi-planar surfaces and 3D for generic surfaces) and a vectorial representation of the irradiance, this approach is well suited for rendering of cultural heritage objects with global illumination. Our structure is robust to simplification techniques with appearance preservation of effects due to aging and tool marks.

Based on this first validation, improvements can be done on both geometry and illumination. About geometry, we can improve rendering speed of APOs by reducing the required number of texture lookups. For the texture atlas of normal maps, a parameterization of the simplified object could improve both their compactness and their continuity.

About illumination, we have to accelerate the precomputation, since we do not take currently into account the high coherency of the sample paths. We also have to work on the integration of more complex direct illumination. Furthermore, we are currently working a more robust directional basis for vectorial representations which would improve the interpolation scheme and provide a more accurate solution.

References

[Arv94] ARVO J.: The irradiance Jacobian for partially occluded polyhedral sources. In *Proc. ACM SIGGRAPH'94* (1994), pp. 343–350.

[CMSR98] CIGNONI P., MONTANI C., SCOPIGNO R., ROCCHINI C.: A general method for preserving attribute values on simplified meshes. In *VIS '98: Proc. conference on Visualization '98* (1998), IEEE Computer Society Press, pp. 59–66.

[LBJS07] LACOSTE J., BOUBEKEUR T., JOBARD B., SCHLICK C.: Appearance preserving octree-textures. In *GRAPHITE '07: Proc. international conference on Computer graphics and interactive techniques in Australia and Southeast Asia* (2007), ACM, pp. 87–93.

[LPC*00] LEVOY M., PULLI K., CURLESS B., RUSINKIEWICZ S., KOLLER D., PEREIRA L., GINZTON M., ANDERSON S., DAVIS J., GINSBERG J., SHADE J., FULK D.: The digital Michelangelo project: 3D scanning of large statues. In *ACM SIGGRAPH '00* (2000), pp. 131–144.

[PRG*08] PACANOWSKI R., RAYNAUD M., GRANIER X., REUTER P., SCHLICK C., POULIN P.: Efficient streaming of 3d scenes with complex geometry and complex lighting. In *Web3D 2008: Proc. International Conference on 3D Web Technology* (2008), ACM Press.

[WRC88] WARD G., RUBINSTEIN F., CLEAR R.: A ray tracing solution for diffuse interreflection. In *Proc. ACM SIGGRAPH'88* (1988), pp. 85–92.