



HAL
open science

Volumetric Vector-based Representation for Indirect Illumination Caching

Romain Pacanowski, Xavier Granier, Christophe Schlick, Poulin Pierre

► **To cite this version:**

Romain Pacanowski, Xavier Granier, Christophe Schlick, Poulin Pierre. Volumetric Vector-based Representation for Indirect Illumination Caching. *Journal of Computer Science and Technology*, 2010, 25 (5), pp.925-932. 10.1007/s11390-010-9377-2 . inria-00505132v2

HAL Id: inria-00505132

<https://inria.hal.science/inria-00505132v2>

Submitted on 5 Feb 2011

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.

Volumetric Vector-Based Representation for Indirect Illumination Caching

Romain Pacanowski¹, *Member, ACM*, Xavier Granier^{2,3}, *Member, ACM*, Eurographics, Christophe Schlick², and Pierre Poulin⁴

¹ CEA-CESTA, BP 2 Le Barp, 33114, France

² INRIA Bordeaux Sud-Ouest / LaBRI - Bordeaux University, Talence, 33405, France

³ State Key Lab of CAD&CG, Zhejiang University, Hangzhou 310027, China

⁴ LIGUM, Dept. I.R.O., Université de Montréal, Montréal QC, H3C 3J7, Canada

E-mail: romain.pacanowski@labri.fr, xavier.granier@inria.fr, schlick@labri.fr, poulin@iro.umontreal.ca

Received 15/12/2009; Revised 01/04/2010.

Abstract This paper introduces a caching technique based on a volumetric representation that captures low-frequency indirect illumination. This structure is intended for efficient storage and manipulation of illumination. It is based on a 3D grid that stores a fixed set of irradiance vectors. During preprocessing, this representation can be built using almost any existing global illumination software. During rendering, the indirect illumination within a voxel is interpolated from its associated irradiance vectors, and is used as additional local light sources.

Compared with other techniques, the 3D vector-based representation of our technique offers increased robustness against local geometric variations of a scene. We thus demonstrate that it may be employed as an efficient and high-quality caching data structure for bidirectional rendering techniques such as particle tracing or Photon Mapping.

Keywords Caching for Global Illumination, photon mapping, Particle Tracing.

1 Motivation

Designing an appropriate representation for global illumination effects is a very difficult task since the range of possible phenomena widely varies between low spatial frequencies (diffuse reflections, soft shadows, etc.) to extremely high spatial frequencies (specular reflections, sharp shadows, caustics).

The cost for an accurate estimation of indirect illumination in a given 3D scene is orders of magnitude more expensive than the computation of direct illumination generated by the usually small number of light sources in that scene. A natural trend is thus to perform this expensive computation only at a limited number of locations in the scene during a preprocessing stage, to store the results in a data structure, and to interpolate between these stored values during final rendering. This general principle has been used since the early years of global illumination computation.

1.1 Indirect Illumination Caching

According to the radiometric quantities that are stored (e.g. radiosity, radiance, irradiance) and the way these quantities are later used in the final rendering step, different flavors of this basic principle have been developed over the

last twenty years. In this paper, we will use *indirect illumination caching* as a generic name for them. Much work has been performed on that research topic to design efficient data structures and devise accurate mathematical representations for this caching process. The recent introduction of programmable graphics hardware has renewed the interest on that topic, which has led to the introduction of GPU-friendly techniques such as precomputed radiance transfer [21] and fully interactive global illumination [26].

However, although caching of indirect illumination may be envisaged to capture low-frequency effects, its memory requirements explode when extended to high-frequency effects. For such effects, stochastic approaches (see [6] for an overview) have proven their accuracy. Some of them, such as Path Tracing [12], Bidirectional Path Tracing [15], and Metropolis Light Transport [24], are view-dependent and require little intermediate storage. They can produce high-quality images, but are generally time consuming.

To achieve the best of both worlds, state-of-the-art techniques employ a hybrid combination, where view-independent caching is used for low-frequency phenomena and view-dependent computation is performed for high-frequency ones. Among these techniques, *Photon Mapping* [11] has become very popular, thanks to its optimized data

structures and its ability to deal in an integrated manner with various complex illumination effects. *Photon Mapping* is also quite simple to implement and to extend in order to simulate additional phenomena. Based on density estimation [25], photons are stored in a kd-tree where different maps are used for different illumination phenomena (direct illumination, indirect illumination, caustics, etc.). During final gathering, a local search is performed in the kd-tree for each intersection point originating from the camera in order to gather neighboring photons and estimate the local photon density.

1.2 Required Properties

Basically, two main features are desirable for caching strategies. The first feature involves **photometric robustness**, which means that the cached values should be able to capture local variations of radiometric quantities, such as the reflectance of the surface. The second feature involves **geometric robustness**, as the cached values should also be able to capture local variations of geometric quantities, such as the position or the normal of the surface.

As pointed out by Christensen *et al.* [4] and Tabellion *et al.* [23], the efficiency of indirect illumination caching decreases when the geometric complexity of the scene increases. Small scale geometric details, especially high-frequency variations of surface normals, require a densely sampled cache structure in order to capture the subtle changes of illumination. Ten years ago, with the usually lower geometric complexity of 3D scenes, photometric robustness was obviously the main concern when designing cache structures, but nowadays, when even a moderately complex scene contains several hundred thousands of polygons, geometric robustness has become a major concern for caching structures.

To address this issue, this paper proposes to cache indirect illumination with a volumetric vector-based data structure, based on the irradiance vector [2]. This caching structure, called **Irradiance Vector Grid** [18, 19], is only intended for low-frequency indirect illumination, as we assume that direct illumination and high-frequency indirect illumination are dealt with more specific techniques (e.g. soft shadow maps, specular or glossy environment maps, view-dependent ray casting, etc.). The Irradiance Vector Grid offers the following features:

1. robustness against local variations of diffuse reflectance,
2. robustness against local variations of surface normals,
3. smoothness everywhere in the 3D scene,
4. low memory requirements.

Features 1 and 2 increase the efficiency of the cache structure, as diffuse reflectance and normal vector variations are the most important elements involved in low-frequency indirect phenomena. Feature 3 guarantees smooth reconstruction of the indirect illumination from the cache, and Feature 4 increases scalability of the technique.

1.3 Paper Organization

This paper is organized as follows. After recalling some related work in Section 2, we detail how to construct the irradiance vector grid with a special focus on photon tracing in Section 3, and how to interpolate local irradiance at any point in the scene in Section 4. Finally we present and analyze some results, and conclude with potential improvements.

2 Existing Caching Techniques

Caching precomputes a chosen radiometric quantity at different locations in the scene and interpolates between these values everywhere. Three main concerns with caching include: (1) photometric robustness (the type of radiometric quantity); (2) geometric robustness (how this quantity and the interpolation scheme scale when the geometric details increase); and (3) the smoothness of the reconstruction.

2.1 Photometric Independence

Classical caching structures [27, 3, 4] store irradiance, as it allows modifications of the diffuse albedo of materials without having to recompute the cached values. To overcome the limitation to diffuse reflectance or at best, low-frequency BRDFs imposed by irradiance caching, new schemes based on incident radiance caching have been introduced. Encoding incident radiance with spherical harmonics [14, 1] or wavelets [22] is more accurate than constant bases [9], but with both representations, the number of coefficients quickly explodes for high-frequency BRDFs, thus still limiting the method to moderately glossy BRDFs.

2.2 Geometric Independence

Although irradiance is a geometrically dependent quantity, radiance is not. Thus radiance caching theoretically offers geometric robustness. Unfortunately, the usual strategy used to place the precomputed samples [27] depends on the underlying geometry. More generally, the interpolation scheme depends on the distance between sample positions and their corresponding normals (e.g. [23, 26]), leading to high density of samples when the normal field of a surface is strongly varying. Therefore, radiance caching cannot be applied in the case of highly detailed surfaces,

because the required number of samples quickly becomes huge.

Over the different representations of illumination, vector quantities seem to be prone to geometric robustness. An *Irradiance Vector* [2] is a vector quantity whose norm represents the incident energy over a whole hemisphere and whose direction represents the main incident illumination direction. This quantity results in an increased robustness to geometric variations as shown in the works of Lecot et al. [16], Willmott et al. [28] and Gobbetti et al. [8] for radiosity techniques, and as pointed out by Lehtinen et al. [17]. In general, vector representations are more independent from geometry for caching, as shown by the development of *Light Vectors* [29, 20]. Unfortunately, these Light Vectors are not photometric independent, and therefore are less interesting to our application.

2.3 Continuous Reconstruction

Another issue with caching involves the interpolation scheme used during rendering. Irradiance and radiance caching schemes need to store their samples in an efficient structure (e.g. kd-tree or octree) in order to quickly retrieve them when interpolation is needed. Recent sampling strategies for caching (e.g. [14, 7, 26]) have improved on the smoothness of the reconstruction. However, due to the combined facts that these samples can be placed anywhere and that only interpolation based on local neighborhood can be performed for performance issues, these schemes cannot guarantee a continuous reconstruction of the stored radiometric quantity. On the contrary, with volumetric representations such as irradiance volumes [9], continuous interpolation becomes easier to perform. Unfortunately as the irradiance volumes cache incident radiance, an integration computation is required at rendering.

2.4 Overview

Considering those three issues and the qualities of vectorial quantities, we propose a volumetric data structure based on irradiance vectors, which offers improved geometric robustness and photometric robustness similar to other comparable methods. To provide a smooth reconstruction of indirect illumination, we use a continuous interpolation scheme that does not depend on surface geometry. Furthermore, our representation has lower memory consumption.

3 Building the Irradiance Vector Grid

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 \mathbf{v}^{ijk} 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 ($\pm \mathbf{x} | \pm \mathbf{y} | \pm \mathbf{z}$). Actually, we store an irradiance matrix, as one vector is used for each of the three color channels. In the remainder of this paper, we will note \mathbf{I}_δ^{ijk} the irradiance vector stored at vertex \mathbf{v}^{ijk} in direction δ , where $\delta = \pm \mathbf{x} | \pm \mathbf{y} | \pm \mathbf{z}$.

3.1 Irradiance Vector

For a given wavelength, the *irradiance vector* [2] $\mathbf{I}_n(\mathbf{p})$ is defined for point \mathbf{p} with normal \mathbf{n} as

$$\mathbf{I}_n(\mathbf{p}) = \int_{\Omega_n} L(\mathbf{p} \leftarrow \boldsymbol{\omega}_i) \boldsymbol{\omega}_i d\boldsymbol{\omega}_i$$

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

$$L_r(\mathbf{p} \rightarrow \boldsymbol{\omega}_o) = \frac{\rho_D(\mathbf{p})}{\pi} \langle \mathbf{I}_n(\mathbf{p}), \mathbf{n} \rangle \quad (1)$$

where ρ_D is a diffuse BRDF and $\langle \cdot, \cdot \rangle$ denotes a dot product. The main benefits of irradiance vectors, compared to irradiance, is that for a local variation of normals, the reflected radiance can be adjusted, making this representation more geometrically robust (e.g. [17]). However, this quantity is still dependent on the normal at the sample position. We decided to precompute irradiance vectors for six predefined directions $\pm \mathbf{x} | \pm \mathbf{y} | \pm \mathbf{z}$ associated to normals. This solution improves geometric robustness, while providing an estimate for the irradiance vector at any point in 3D space.

3.2 Estimating the Irradiance Vector

Any global illumination algorithm can be used to estimate the irradiance vectors stored in the grid.

In our implementation, we use Photon Tracing and propagate photons from the light sources through the grid. Every time a photon traverses a voxel face, its contribution is added to the irradiance vector \mathbf{I}_δ^{ijk} associated with the nearest voxel vertex \mathbf{v}^{ijk} and the direction δ provided by the normal of the face (see Figure 1). When a photon hits scene geometry, a classical stochastic reflection is applied according to the local BRDF, and the photon is propagated in the resulting direction.

The photon propagation is accomplished in two steps. First, using a ray tracing acceleration structure, we find the closest intersection with the scene. Then, we propagate the photon into the Irradiance Vector Grid without any intersection test. Once all photons have been treated, a normalization step is performed on the irradiance vectors

$$\mathbf{I}_\delta^{ijk} = \frac{1}{A_\delta^{ijk}} \sum_{n=1}^{N_\delta^{ijk}} \phi_n \boldsymbol{\omega}_n$$




Figure 1. Irradiance Vector computation illustrated in 2D. When a photon hits a face of a voxel, its contribution is added to the irradiance vector associated with the nearest voxel vertex. Photon contribution added to (Left) $I_{-x}^{i,j}$ and (Right) $I_{+y}^{i,j+1}$.

where N_{δ}^{ijk} is the number of photons that have contributed to the irradiance vector at vertex v^{ijk} in direction δ , A_{δ}^{ijk} is the area of the rectangular cell centered at v^{ijk} in direction δ , and ϕ_n is the energy of the particle n . As we use a uniform rectangular 3D grid, the area of such cell is simply the area of the voxel face oriented in the same direction, except for grid boundary vertices, where the area is divided by two, and for grid corner vertices, where it is divided by four.

Our density estimation correctly accounts for the intersection of the photons with the grid. Therefore, our approach does not suffer from the classical boundary bias of Photon Mapping, where gathering spheres with large radii are used to collect photons, such as in room corners and along contours of flat surfaces. Also, unlike the strategy of Havran et al. [10], our approach does not need to store all the rays generated from the photon propagation.

In fact we compute the colored irradiance vector, defined by three primary colors (R, G, B). We thus need, for each δ , three irradiance vectors stored in a 3×3 matrix \mathbf{M}_{δ}

$$\mathbf{M}_{\delta} = [\mathbf{I}_{R,\delta} \ \mathbf{I}_{G,\delta} \ \mathbf{I}_{B,\delta}] = \frac{1}{A_{\delta}^{ijk}} \sum_{n=1}^{N_{\delta}^{ijk}} \omega_n \phi_n^T,$$

where ϕ_n is the colored energy of the particle n .

4 Using the Irradiance Vertex Grid

4.1 Interpolation of Irradiance Vectors

In order to compute smooth indirect illumination, we interpolate an irradiance vector for each point \mathbf{p} with normal \mathbf{n} that needs to be shaded. This interpolation is computed by a spatial interpolation according to \mathbf{p} , followed by a directional interpolation according to \mathbf{n} .

In the first step, the irradiance vector $\mathbf{I}_{\delta}(\mathbf{p})$ is obtained by spatial interpolation of the irradiance vectors $\mathbf{I}_{\delta}^{ijk}$ stored

at the grid vertices surrounding point \mathbf{p} . The interpolation is only computed for three out of the six possible directions of δ . The choice between $\pm x$ (resp. $\pm y$ and $\pm z$) is based according to the sign of n_x (resp. n_y and n_z). Trilinear or tricubic interpolation provides satisfactorily smooth results for spatial interpolation. In the second step, the final interpolated irradiance vector $\mathbf{I}_n(\mathbf{p})$ is obtained by remapping the three spatially interpolated irradiance vectors according to the normal direction \mathbf{n} at point \mathbf{p}

$$\mathbf{I}_n(\mathbf{p}) = \mathbf{I}_x(\mathbf{p}) n_x^2 + \mathbf{I}_y(\mathbf{p}) n_y^2 + \mathbf{I}_z(\mathbf{p}) n_z^2.$$

This interpolation scheme does not introduce any error when \mathbf{n} is colinear with one of the main directions ($\pm x | \pm y | \pm z$).

4.2 Caching Strategy

Even if it has been previously used for real-time GPU rendering [19], the Irradiance Vector Grid is also appropriate in the context of off-line global illumination rendering. Because it provides spatial and directional smooth reconstruction of irradiance vectors, it can be used directly to estimate diffuse indirect illumination, without propagating secondary rays for final gathering. Direct use of cached values could also be done using existing techniques [27, 14], but they do not guarantee a continuous reconstruction of the indirect illumination. Moreover, as they do not offer geometric robustness, these techniques would require a large number of samples to provide an accurate estimate for highly detailed geometry.

The Irradiance Vector Grid could also be used as an efficient caching structure for stochastic approaches. For instance, similarly to Photon Mapping, our grid may be accessed indirectly by shooting secondary rays from points being shaded. The main advantages are that high-frequency details such as indirect soft shadows would be well preserved, and that reconstruction errors would be masked

since a diffuse or low-glossy reflection is similar to a low-pass filter [5]. Therefore, a simple trilinear spatial interpolation would provide a good compromise between speed and quality.

5 Results

All the results have been computed on an AMD 64bit 3500+ processor with 2GB of memory. Images in Figures 3 and 4 have a 640×480 resolution.

5.1 Geometric Robustness

Due to its robustness against geometric details, it is possible with our technique to precompute the volumetric irradiance vectors on a low resolution version of the objects, similarly to the approach developed by Tabellion et al. [23]. In Figure 2, the same grid resolution ($20 \times 22 \times 20$) and the same number of photons (8 M) are used during precomputation. The resulting recomputed illuminations in the two images are extremely similar (less than 1% maximum pixel difference) for a precomputation time divided by three. Refining the geometry modifies object curvatures, and therefore photon reflection directions. This explains the small differences located on the two objects and on the room ceiling.

5.2 Off-Line Caching

To evaluate our approach, we present two test scenes with complex lighting and many geometric details (approx. 232 MB of triangles). The first configuration (Figure 3) presents a scene mostly directly illuminated from 11 light sources. The second configuration (Figure 4) is a classical two-room scene where one room is indirectly illuminated by one light source located in the other room. Both scenes have more than 8 M polygons due the highly detailed objects. They illustrate the geometric robustness of our technique.

We compared our technique to Photon Mapping combined with Christensen’s [3] precomputed samples. Precomputation time for our technique involves shooting photons and accumulating their contribution in the irradiance vector grid, and for Christensen’s, it involves shooting photons, balancing the kd-tree, and precomputing irradiance. For all images generated with Christensen’s technique, we fixed the number of precomputed samples to one fourth of the total number of photon hits, as originally suggested [3].

Figure 3 compares the results obtained with our technique and Christensen’s technique for the direct illumination configuration. For a fair comparison, a reference solution has been computed with a high density (1600 rays per pixel) Path-Tracing algorithm. For equivalent precomputation times, our technique allows a much faster computation (51s vs. 5 527s) of the indirect illumination component due

to its direct access to the cached values. In our technique, the final gathering step spends most of its time computing ray-geometry intersections and direct illumination, whereas Christensen’s cache requires to cast a large number of additional gathering rays, since direct access produces very objectionable illumination patterns with zones of constant irradiance values similar to Voronoi diagrams. This is a well-known artefact [3] preventing direct access to cached irradiance.

For the indirect illumination configuration presented in Figure 4, we have also tried a direct access to our structure, but to capture finer shadows cast by the columns, a higher-resolution grid ($30 \times 20 \times 92$) was built and consequently a larger number of photons needed to be shot. This resulted in a much larger precomputation time, but our method with direct access is still much faster (2 311s vs. 47 971s) than Christensen’s caching technique for a similar quality. We have also tested our technique with an indirect access (cf. Figure 4) to the cached values by using a lower resolution grid ($12 \times 8 \times 20$) traversed by the same number of photons as with Christensen’s method. Our technique reduces both precomputation time (150s vs. 466s) and reconstruction time of the indirect illumination (2 376s vs. 10 050s), while retaining similar illumination features. Access time to our grid is constant in the number of the cached samples, instead of being logarithmic as with a kd-tree [3, 4], which explains these gains.




Figure 5. Construction time as a function of grid resolutions and number of photons shot. The construction time grows linearly with the number of photons shot.

5.3 Discussion

The construction time of our Irradiance Vector Grid structure grows mainly linearly with the number of photons shot (see Figure 5), albeit a small overhead exists when increasing the size of the grid. Note that in Figure 4




Figure 2. Geometric robustness with direct access. (Left) The precomputation is performed on a low-detailed geometry, and (Center) on a high-detailed geometry. (Right) The difference color-coded RGB image (maximum pixel difference $\varepsilon = 1.5\%$).




Figure 3. Scene with mostly direct illumination. 16 rays per pixel were used and 5 M photons shot. (Left) Our technique with a $40 \times 50 \times 40$ grid (24.5 MB) is used directly with a tricubic interpolation scheme. (Center) Photon Mapping with Christensen's cache: 50 photons of the 5 M (124 MB) were used to precompute each irradiance sample. (Right) Reference solution computed by Path Tracing with 1 600 rays per pixel. The general illumination patterns are similar in all three techniques, however our much faster indirect illumination computation (51s vs. 5 527s) strongly reduces the total rendering time (1 341s vs. 6 618s).

(Left), precomputation time may be reduced by lowering the grid size and shooting less photons. This will lose high-frequency details but indirect effects of the light near geometrically complex objects will still be captured. Noise due to undersampling of the illumination during the photon shooting pass appears mostly on larger surfaces and on higher frequency details such as shadows. The quality of the reconstruction directly scales with the number of photons shot, as the illumination details scale with the size of the grid. However, even if unbiased, our estimation can lead

to light leakages and under-estimation of irradiance values, depending on the relative geometric configuration with the grid. These artefacts are reduced by increasing the resolution of the grid. Some of these artefacts may also be reduced by using an adaptive volumetric structure such as an octree, or by an improved estimation strategy [13]. Finally, it should be mentioned that our technique cannot reach the reconstruction quality of Ray maps [10], but as counterpart, less storage is required.




Figure 4. Scene with mostly indirect illumination. (Left) Our technique using directly a $30 \times 20 \times 92$ uncompressed grid (11.4 MB) constructed with 80 M photons. (Center) Christensen’s method reference image with 3 200 rays to sample the hemisphere and 5 M photons stored (124 MB). (Right) Our technique using indirectly a $12 \times 8 \times 20$ uncompressed grid (590 KB) with the same number of photons shot and the same number of rays to sample the hemisphere.

6 Conclusion

In this paper, we have presented a new representation for indirect illumination, based on a 3D grid of irradiance vectors. This representation allows a smooth-everywhere reconstruction of the irradiance. Thanks to the irradiance vectors, the resulting solution is more robust to local variations of geometry, as shown with the presented results.

We have implemented this 3D structure as indirect illumination caching scheme for Photon Mapping. The results show that diffuse inter-reflections are well captured. Compared to existing solutions, our approach requires fewer cached samples for higher-quality cached indirect illumination. Additionally, it does not require any photon storage. Furthermore, our irradiance cache can be accessed directly during the final gathering pass.

Future Work

We propose mainly two directions for improving our technique. Firstly, the estimation of irradiance vectors is based on a regular grid. Identifying proper resolutions for the grid and the number of photons to shoot requires a good understanding of the illumination effects. A more automatic estimate would facilitate the use of our technique. For larger scenes, we would like to use a multiresolution and adaptive structure in order to reduce the construction cost. This would also allow better capture of indirect illumination near surfaces and maybe find better automatic estimates.

Secondly, since the scheme is based on linear interpolation, some artefacts may appear for complex variations of illumination. The introduction of gradients would improve the smoothness of the reconstruction where incoming illu-

mination varies quickly. Improved visibility estimation will also provide higher-quality reconstruction.

Acknowledgement

Xavier Granier is supported by the Open Project Program of the State Key Lab of CAD&CG (Grant No. A1007), Zhejiang University.

References

- [1] Okan Arıkan, David A. Forsyth, and James F. O’Brien. Fast and Detailed Approximate Global Illumination by Irradiance Decomposition. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 24(3):1108 – 1114, 2005.
- [2] James R. Arvo. The irradiance Jacobian for partially occluded polyhedral sources. In *Proc. SIGGRAPH ’94*, annual conference on computer graphics and interactive techniques, pages 343–350, Orlando, Florida, July 1994. ACM.
- [3] Per H. Christensen. Faster Global Photon Map Global Illumination. *J. of Graphics Tools*, 4(3):1–10, 1999.
- [4] Per H. Christensen and Dana Batali. An Irradiance Atlas for Global Illumination in Complex Production Scenes. In *Proc. Eurographics Symposium on Rendering 2004*, pages 133–141, Norrköping, Sweden, June 2004. Eurographics.
- [5] Frédo Durand, Nicolas Holzschuch, Cyril Soler, Eric Chan, and François Sillion. A Frequency Analysis

- of Light Transport. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 24(3), 2005.
- [6] Philip Dutré, Kavita Bala, and Philippe Bekaert. *Advanced Global Illumination (Second Edition)*. A. K. Peters, Ltd., 2006.
- [7] Václav Gassenbauer, Jaroslav Křivánek, and Kadi Bouatouch. Spatial Directional Radiance Caching. *Comput. Graph. Forum (Proc. Eurographics Symposium on Rendering)*, 28(4):1189–1198, 2009.
- [8] Enrico Gobbetti, Leonardo Spanò, and Marco Agus. Hierarchical higher order face cluster radiosity for global illumination walkthroughs of complex non-diffuse environments. *Comput. Graph. Forum (Proc. Eurographics)*, 22(3):563–572, 2003.
- [9] Gene Greger, Peter Shirley, Philip M. Hubbard, and Donald P. Greenberg. Irradiance Volume. *IEEE Comput. Graph. Appl.*, 18(2):32–43, 1992.
- [10] Vlastimil Havran, Jiri Bittner, Robert Herzog, and Hans-Peter Seidel. Ray maps for global illumination. In *Proc. Eurographics Symposium on Rendering 2005*, pages 43–54, Konstanz, Germany, June 2005. Eurographics.
- [11] Henrik Wann Jensen. *Realistic Image Synthesis using Photon Mapping*. A.K. Peters, 2001.
- [12] Jim T. Kajiya. The rendering equation. In *Proc. SIGGRAPH '86*, annual conference on computer graphics and interactive techniques, pages 143–150, Dallas, Texas, August 1986. ACM.
- [13] Janne Kontkanen and Samuli Laine. Sampling Pre-computed Volumetric Lighting. *J. Graph. Tools*, 11(3):1–16, 2006.
- [14] Jaroslav Křivánek, Kadi Bouatouch, Sumanta N. Patanaik, and Jiří Žára. Making Radiance and Irradiance Caching Practical: Adaptive Caching and Neighbor Clamping. In *Proc. Eurographics Symposium on Rendering 2006*, pages 127–138, Nicosia, Cyprus, June 2006. Eurographics.
- [15] Eric P. F. Lafortune and Yves D. Willems. Bidirectional Path Tracing. In *Proc. International Conference on Computational Graphics and Visualization Techniques (Compugraphics '93)*, pages 145–153, Alvor, Portugal, December 1993.
- [16] Grégory Lecot, Bruno Lévy, Laurent Alonso, and Jean-Claude Paul. Master-element vector irradiance for large tessellated models. In *GRAPHITE '05: Proc. international conference on computer graphics and interactive techniques in Australasia and South East Asia*, pages 315–322, Dunedin, New Zealand, November 2005. ACM.
- [17] Jaakko Lehtinen, Matthias Zwicker, Emmanuel Turquin, Janne Kontkanen, Frédo Durand, François Sillion, and Timo Aila. A Meshless Hierarchical Representation for Light Transport. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 27(3), 2008.
- [18] Romain Pacanowski, Mickaël Raynaud, Xavier Granier, Patrick Reuter, Christophe Schlick, and Pierre Poulin. Efficient Streaming of 3D Scenes with Complex Geometry and Complex Lighting. In *Web3D '08: Proc. international symposium on 3D web technology*, pages 11–17, Los Angeles, California, August 2008. ACM.
- [19] Romain Pacanowski, Mickaël Raynaud, Julien Lacoste, Xavier Granier, Patrick Reuter, Christophe Schlick, and Pierre Poulin. Compact Structures for Interactive Global Illumination on Large Cultural Objects. In *International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST 2008): Shorts and Projects*, Braga, Portugal, December 2008. Eurographics.
- [20] Xavier Serpaggi and Bernard Péroche. An Adaptive Method for Indirect Illumination Using Light Vectors. *Comput. Graph. Forum (Proc. Eurographics 2001)*, 20(3):278–287, 2001.
- [21] Peter-Pike Sloan, Jan Kautz, and John Snyder. Pre-computed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 21(3):527–536, 2002.
- [22] Marc Stamminger, Annette Scheel, Xavier Granier, Frédéric Perez-Carzorla, George Drettakis, and François Sillion. Efficient Glossy Global Illumination with Interactive Viewing. *Comput. Graph. Forum*, 19(1):13–25, 2000.
- [23] Eric Tabellion and Arnauld Lamorlette. An Approximate Global Illumination System for Computer Generated Films. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 23(3):469–476, 2004.
- [24] Eric Veach and Leonidas J. Guibas. Metropolis light transport. In *Proc. SIGGRAPH '97*, annual conference on computer graphics and interactive techniques, pages 65–76, Los Angeles, California, August 1997. ACM/Addison-Wesley Publishing Co.
- [25] Bruce Walter, Philip M. Hubbard, Peter Shirley, and Donald P. Greenberg. Global illumination using local

linear density estimation. In *Proc. SIGGRAPH '97*, annual conference on computer graphics and interactive techniques, pages 217–259, Los Angeles, California, August 1997. ACM/Addison-Wesley Publishing Co.

- [26] Rui Wang, Rui Wang, Kun Zhou, Minghao Pan, and Hujun Bao. An efficient GPU-based approach for interactive global illumination. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 28(3):1–8, 2009.
- [27] Gregory J. Ward, Francis M. Rubinstein, and Robert D. Clear. A ray tracing solution for diffuse interreflection. In *Proc. SIGGRAPH '88*, annual conference on computer graphics and interactive techniques, pages 85–92, Atlanta, Georgia, August 1988. ACM.
- [28] Andrew J. Willmott, Paul S. Heckbert, and Michael Garland. Face Cluster Radiosity. In *Proc. Eurographics Workshop on Rendering 99*, pages 293–304, Granada, Spain, June 1999. Springer-Verlag.
- [29] Jacques Zaninetti, Xavier Serpaggi, and Bernard Péroche. A Vector Approach for Global Illumination in Ray Tracing. *Comput. Graph. Forum (Proc. Eurographics '98)*, 17(3), 1998.



Romain Pacanowski received his Engineer degree from EFREI and his MS degree from the University of Bordeaux. He holds a PhD from the University of Bordeaux (France) and Université de Montréal (Canada). He is currently working for the CEA as postdoctoral fellow. His research interest includes realistic rendering as well as appearance modeling and acquisition. He is a member for ACM.

interest includes realistic rendering as well as appearance modeling and acquisition. He is a member for ACM.



Xavier Granier received his Engineer and MS degrees from the

Grenoble Institute of Technology and his PhD Diploma for University Joseph Fourier in Grenoble (France). He is currently at INRIA Bordeaux Sud-Ouest (France) as a research scientist. His research interest includes realistic lighting and expressive rendering, appearance modeling and acquisition. He is a member for ACM and EUROGRAPHICS.



Christophe Schlick Christophe Schlick is professor in Computer Science at the University of Bordeaux 2 (France) where he has recently headed the Applied Mathematics and Computer Science Department. After having received his PhD in 1992 for his work on BRDF models and Monte Carlo techniques, his research interests have embraced many aspects of computer graphics, including global illumination, procedural texture and geometric synthesis, curves and surfaces, point based modeling and rendering.

Monte Carlo techniques, his research interests have embraced many aspects of computer graphics, including global illumination, procedural texture and geometric synthesis, curves and surfaces, point based modeling and rendering.



Pierre Poulin is a full professor in the Computer Science and Operations Research department of the Université de Montréal. He holds a Ph.D. from the University of British Columbia and a M.Sc. from the University of Toronto, both in Computer Science. He has served on program

committees of more than 35 international conferences. His research interests cover a wide range of topics, including image synthesis, image-based modeling, procedural modeling, natural phenomena, scientific visualization, and computer animation.