Volume xx (200y), Number z, pp. 1–13

Adaptive Records for Irradiance Caching

Mickaël Ribardière[1,2] , Samuel Carré[1] , Kadi Bouatouch[3]

[1]Centre Scientifique et Technique du Bâtiment (CSTB), France
[2]Université de Rennes 1, France
[3]IRISA - INRIA Rennes, France

Abstract

Irradiance Caching is one of the most widely used algorithms to speed up global illumination. In this paper, we propose an algorithm based on the Irradiance Caching scheme that allows us (1) to adjust the density of cached records according to illumination changes and (2) to store both direct and indirect contributions in the records. To achieve this, a new record footprint is presented. While the original method uses records having circular footprints depending only on geometrical features, our record footprints have a more complex shape which accounts for both geometry and irradiance variations. Irradiance values are computed using a classical Monte Carlo ray tracing method that simplifies the determination of nearby objects and the pre-computation of the shape of the influence zone of the current record. By gathering irradiance due to all the incident rays, illumination changes are evaluated to adjust the footprint's records. As a consequence, the record footprints are smaller where illumination gradients are high. With this technique, the record density will no longer be uniform. Strong variations of irradiance due to direct contributions (for example, due to shadows) can be integrated in the record's data structure. Caching direct illumination is of high importance, especially in the case of scenes having many light sources with complex geometry as well as wide surfaces exposed to daylight. Recomputing direct illumination for the whole image can be very time-consuming, especially for animation rendering or for high resolution pictures. Storing such contributions in the irradiance cache seems to be an appropriate solution to accelerate the final rendering pass.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Rendering, Global Illumination

1. Introduction

Computing global illumination in a reasonable time while handling different effects (glossy reflection, refraction, caustics, subsurface scattering) still is a challenge in computer graphics. To take up this challenge, Irradiance/Radiance caching is one of the most widely used algorithms [WRC88, WH92, TL04, KBPv06, KGW*08]. It is a ray tracing-based method that computes the indirect diffuse illumination component. It relies on the fact that indirect irradiance changes slowly over the object surfaces. Indirect illumination is computed for a sparse set of points, called records, and stored in a cache. Direct illumination is evaluated using classical techniques (shadow map for point light sources, Monte Carlo for area light sources), while indirect illumination is computed either by a Monte Carlo method for the records or using interpolation for the other points.

In some applications, direct illumination can be very costly, for instance in architectural design or in lighting projects requiring accurate and physically-based results (spectral quantities with many wavelets not only 3 as usual). The scene may have many area light sources of complex geometry and specific photometric characteristics. Daylight (sun and sky) is often present as well. In such cases, direct illumination computation can be very time demanding. Storing the direct irradiance together with the original cached irradiance is not possible. Indeed, direct illumination changes quickly over surfaces because of occlusions (shadows). The original Irradiance Caching method [WRC88] and the original gradients [WH92] cannot compensate for high illumina-

tion changes because the record density must be proportional to the rate of illumination change. In addition, gradients are used for computing smooth changes over surfaces and cannot detect occlusions in case of translation. Furthermore, the record density is only controlled by geometry and not by illumination results. This is not sufficient in complex indirect illumination cases in which a significant indirect illumination source illuminates the scene and produces shadows.

This paper presents an adaptive Irradiance Caching algorithm which adapts the record density depending on the surface geometry and the illumination changes. This density control reduces the number of cached records along edges and corners. The record density has to be proportional to the change rate of direct and indirect illumination when high gradients are detected. A novel gradient is proposed to detect these significant illumination changes to adaptively determine the shape and size of the record footprints. This method straightforwardly includes both direct and indirect illumination in the same cache. It is then efficient in terms of direct illumination from complex light sources. Furthermore, storing all kinds of illumination in records allows to easily reuse the cache for interpolation of global illumination at points different from records.

The paper is organized as follows: Section 2 summarizes the related works. Section 3 introduces the new method. Then section 4 describes adaptive record footprints. Section 5 explains how gradients including direct illumination are computed. Section 6 shows how to adapt the size of the influence zone of a record depending on the interpolation accuracy needed. Finally, section 7 presents and discusses some results before concluding.

2. Related Work

The technique proposed in this paper is based on the classical Irradiance Caching algorithm. Irradiance Caching was introduced by Ward *et al.* [WRC88] to speed up global illumination computations based on Monte Carlo ray tracing. It has been widely used ever since [KGW*08]. The algorithm takes advantage of spatial coherence. Irradiance is computed for only a few points over the surfaces of the scene and then interpolated to reconstruct the radiance for each other point. In [WH92], Ward and Heckbert improved this interpolation scheme by using irradiance gradients. As said before, the original algorithm only caches irradiance due to indirect illumination.

In [SM02], Smyk *et al.* increased the density of records according to the gradient magnitude to better reconstruct high changes of indirect illumination. This density is controlled by modifying the records radii. However, this method adds a considerable number of records around the corners and the edges of the scene. Tabellion and Lamorlette [TL04] proposed several improvements of the original method such as using a minimum distance to the near objects instead of the harmonic mean to determine a record radius. As in [SM02], an excessive record density may be assigned to concave objects. In the same way, Křivánek *et al.* [KBPv06] (see [KGW*08] too) proposed other practical modifications to the original Ward's method *et al.* to make it more practical. With *Adaptive Caching* the spatial density of the cached values gets variable. The change rate of indirect illumination controls this adaptive density in order to avoid visible interpolation artifacts. The overlapping area of a record which include a discontinuity in the interpolation will then be reduced. If the change rate in indirect illumination is too high then the radius of a record is smaller, which increases the record density. Unfortunately, with this adaptive caching method, the record which provides the highest discontinuity is often the one which is closest to the considered point and which provides the best interpolation. Such records may be very important for a good interpolation and reducing its zone may add a lot of records in the cache.

Křivánek *et al.* presented in the same paper another improvement referred to as *neighbor clamping*. This technique minimizes the ray leaking problem caused by a poor detection of nearby objects. The radius of a record is controlled as well as all its nearby records. Thus there are more chances to detect small sources of indirect illumination. Arikan *et al.* [AFO05] partitions the incoming radiance into two terms: distant term and near term for distant and nearby emitters. These components are approximated separately. The far-field component does not change rapidly over surfaces. The number of directional samples and the density of cache records (spherical harmonics are used to store and interpolate the incoming radiance) can be less compact. The near-field component is sampled more densely to approximate the rapid changes of the radiance.

One of the strengths of Irradiance Caching is its easy use in animation rendering. Although the algorithm is view-dependent, interpolation requires cached values that can be used irrespective of the point of view in successive frames. A new record is stored when a new part of the scene is viewed by the observer. Caution is needed to avoid memory saturation. The first idea is to maintain a time stamp indicating the last frame the record contributed to. In [SiKDM05], Smyk *et al.* proposed a method to perform a more efficient Irradiance Caching in animation rendering. Their aim was to eliminate unpleasant flickering which can appear in classical implementation in animation rendering. They introduced a data structure called *anchor* to exploit temporal coherence and to permanently link cache locations to points intersected by the final gathering rays. Another strategy was proposed by Brouillat *et al.* [BGB08]. They combine Irradiance Caching and photon mapping as in [Jen96, Jen01]. As photon mapping has the advantage of being fully view-independent, they propose to construct an irradiance cache from a photon map. The cache is advantageously computed for the entire scene which allows interactive walkthrough.

As an extension to glossy surfaces, *Radiance Caching* has been proposed. It consists in projecting the incoming radiances stored in records into a spherical or hemispherical harmonics basis [Gre03, GKPB04]. The Radiance Caching scheme was first proposed in [KŮ05, KGPB05] by Křivánek *et al.* Illumination of diffuse and glossy surfaces are stored in the cache and high frequency BRDF and specular reflections are computed in an additional pass. In [KGBP05], the authors improved radiance gradient computations. The radiance caching scheme has known several improvements. Recently, Herzog *et al.* in [HMS09] used an (ir)radiance caching algorithm using the *lightcuts* method, an adaptive and hierarchical instant-radiosity based algorithm [WFA*05, WABG06]. As in [KGPB05], incident radiances are projected into a hemispherical harmonics basis. An interesting feature in this work is the two-level radiance caching: the first one for indirect irradiance and the second for direct irradiance computed in parallel. Two caches are maintained as well as two light trees. The two caches are tested while direct and indirect irradiance computations are performed individually. Another interesting aspect is the multi-pass adaptive caching proposal. It is an extension of Křivánek *et al.*'s adaptive caching method [KBPv06]. It consists of in reducing the record footprints in one dimension resulting in an ellipsoidal footprint only.

Other techniques allow to interpolate direct and indirect illumination from sparse samples to achieve interactive rendering. One of these methods is the *Render Cache* [WDP99, WDG02]. Only few pixels of the image are computed using ray tracing or path tracing. The values of these pixels are stored in a cache and reused for the subsequent frames. The other pixels of the current frame are computed using interpolation/smoothing filters. The result is an estimate of the current image. The *Render Cache* can also be reused across frames in a walkthrough context by reprojecting shaded pixels according to the new viewpoint. Another technique is the *Edge-and-Point Rendering* proposed by Bala *et al.* [BWG03] and extended by [VALBW06]. It is based on the *Render Cache*. *Edge-and-Point Rendering* improves the rendering of high discontinuity regions such as silhouettes and shadow boundaries by using an *Edge-and-Point Image* which stores the discontinuities. The goal of the above techniques are well suited for interactive rendering but not for a complete and physically based lighting simulation.

In this paper, we focus only on diffuse reflections, glossy reflections could be computed in a second pass.

3. Overview

The Irradiance Caching algorithm exploits spatial coherence by sparsely sampling and interpolating indirect irradiance. Each record R stores the following information:

- x_R, position of the record,
- $\vec{n_R}$, normal at x_R,
- $E(R)$, irradiance computed at x_R,
- d_R, harmonic mean distance ([WRC88,KBPv06]) or minimum distance ([TL04]) to objects visible from x_R.

Let $a \cdot d_R$ be the radius of the influence zone of record R, where a is a user-defined constant related to the maximum error. If the influence zone of a record R covers a given point p it is not necessary to compute an irradiance value at p from scratch, but from the already computed irradiance value of R. The contribution of a record R to a point p is weighted by w_R whose expression is:

$$w_R(p) = \left(\frac{\|x_p - x_R\|}{d_R} + \sqrt{1 - \vec{n_p} \cdot \vec{n_R}} \right)^{-1} \quad (1)$$

Let S be the set of records surrounding point p:

$$S = \left\{ R : w_R(p) > \frac{1}{a} \right\}. \quad (2)$$

If S is not empty, the irradiance at p can be estimated as:

$$E(p) = \frac{\sum_{R \in S} E(R) w_R(p)}{\sum_{R \in S} w_R(p)} \quad (3)$$

The original weighting function (equation 1) has some undesirable properties. This function is not continuous at the border of the influence zone of a record as it does not tend toward zero. As proposed in [KGW*08], we use a slightly modified function:

$$w_R(p) = \left(\frac{\|x_p - x_R\|}{d_R} + \sqrt{1 - \vec{n_p} \cdot \vec{n_R}} \right)^{-1} - \frac{1}{a} \quad (4)$$

A record covers p if $w_R(p) > 0$. Otherwise, a new irradiance record is computed and added to the cache and p becomes the record position. Gradients proposed in [WH92,KGBP05] are designed for improving interpolation. Irradiance at x can be expressed by:

$$E(x) = \int_{\Omega} L(x, \vec{\omega})(\vec{n_x} \cdot \vec{\omega}) d\vec{\omega} \quad (5)$$

In the Irradiance Caching method, this value is computed using Monte Carlo ray tracing with an uniform distribution:

$$E(x) = \frac{2\pi}{N} \sum_{i=1}^{N} L_i(x, \vec{\omega_i})(\vec{n_x} \cdot \vec{\omega_i}) \quad (6)$$

where N is the number of ray samples, L_i the incoming radiance for the i-th sample in the direction $\vec{\omega_i}$ and $\vec{n_x}$ the normal at x. Consider $E_i(x) = 2\pi . L_i(x, \vec{\omega_i})(\vec{n_x} \cdot \vec{\omega_i})$ as an estimate of the irradiance at x (up to a scaling factor 2π) due to the i-th ray sample. In other words, it is an estimation of the irradiance due to the i-th equivalent point source. The determination of equivalent point sources has been used in the *Bidirectional Path Tracing* method [VG94, Vea98, LW93] (see figure 1). An equivalent point light source corresponds to a point light or an eye path which is directly connected to x_R, the position of a new record R (x_R is the first vertex in eye

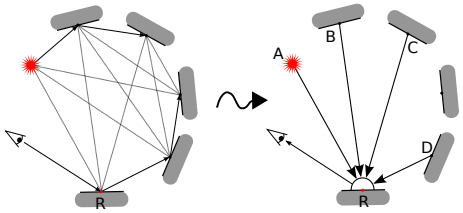

Figure 1: *A sample with a Bidirectional Path Tracing: this sample allows us to know 3 indirect contributions (B,C and D) and 1 direct contribution (A) for R corresponding to 4 equivalent point sources. More tracing would give full coverage of the hemisphere above R.*

path). Now we can substitute $L_i(x,\overrightarrow{\omega_i})(\overrightarrow{n_x}\cdot\overrightarrow{\omega_i})$ in equation 6:

$$E(x)=\frac{1}{N}\sum_{i=1}^{N}E_i(x) \tag{7}$$

Knowing E_i and using ray tracing, an equivalent intensity $I_i^{eq}(\overrightarrow{\omega_i})$ characterizing the i-th equivalent point light source can be expressed:

$$I_i^{eq}(\overrightarrow{\omega_i})=\frac{E_i(x)\cdot d_i^2}{\overrightarrow{n_x}\cdot\overrightarrow{\omega_i}}, \tag{8}$$

where d_i is the distance to the i-th equivalent point source. For the new record R at position x, the irradiance becomes:

$$E(x)=\frac{1}{N}\sum_{i=1}^{N}\frac{I_i^{eq}(\overrightarrow{\omega_i})\cdot(\overrightarrow{n_x}\cdot\overrightarrow{\omega_i})}{d_i^2} \tag{9}$$

As previously said, the original Irradiance Caching scheme does not store direct irradiance. Our method proposes improvements to this technique: use of a new shape of influence zone, a new gradient formulation, an adaptation to strong illumination changes, and the storage of both direct and direct irradiances in the records. These contributions will be detailed later on.

4. Adaptive Records

Adapting the size of the influence zone of a record [KBPv06, HMS09] to the rate of illumination changes increases the number of records while providing better results. The cache density is then adapted to this rate of illumination changes. However, the number of records increases at the edges and the corners too. The same problem appears when using the minimum distance [TL04]. To avoid this problem, Tabellion and Lamorlette control the cache density by using the projected area of a pixel. This results in undersampling far objects and adding a lot of records in case of animation rendering. Though it is necessary to increase the number of records to better estimate the rate of illumination changes, it is important not to increase the density near edges and corners.

We think that the oversampling issue is due to the circular shape of the record's influence zones whose radii are defined by the minimum distance or the harmonic mean of distance to nearby objects. In this section, we propose a new record footprint which better adapts to the geometry. Figure 2 describes the main idea.

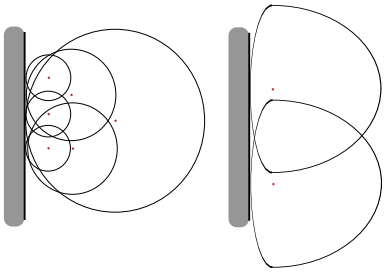

Figure 2: *With a circular influence zone of a record (left) more records are needed. With an influence zone that is small towards close objects and large towards far objects (right), a record covers a better shaped influence zone and finally less records are needed.*

The ideal influence zone of a record can be represented by a planar surface with a curved boundary. This surface cannot be computed exactly but it is approximated through a discretization scheme as explained hereafter. This planar surface lies on the tangent plane associated with a record position and is described by an angular decomposition in 8 pseudo-elliptic zones as seen in figure 3. The local coordinate system of the current influence zone $(\overrightarrow{u},\overrightarrow{v},\overrightarrow{n})$ is used to determine eight axis around the record position R on its tangent plane. The angle between two successive axes k_i and k_{i+1} is equal to $\frac{\Pi}{4}$. An influence zone is then divided into 8 sub-zones, each one being defined by 2 edges and a pseudo-arc. These two edges have the record position R as endpoint and are aligned with 2 successive axes. The lengths of these edges are computed as follows. When computing the irradiance at a record, rays are shot from the record covering an hemisphere placed above it. For each ray close to a footprint axis k_i, we compute the distance to the closest object (called minimum distance from now on) and assign it to the axis. A pseudo-elliptic arc representing the boundary of a sub-zone is not determined exactly. Rather, it is approximated using a linear interpolation between d_1 and d_2 (the lengths of the edges) if k_1 and k_2 are supposed to be the two successive axes supporting the edges. As described in the figure 4, all points P'' lying between k_1 and k_2 on the curved boundary of the surface describe approximatively this pseudo-arc:

$$|\overrightarrow{RP''}|=(1-t)\cdot d_1+t\cdot d_2, \tag{10}$$

with P' the intersection point between $|\overrightarrow{RP''}|$ and $|\overrightarrow{P_1P_2}|$, $t=|\overrightarrow{P1P'}|/|\overrightarrow{P1P2}|$ and $P_i=R+\overrightarrow{k_i}.d_i$.

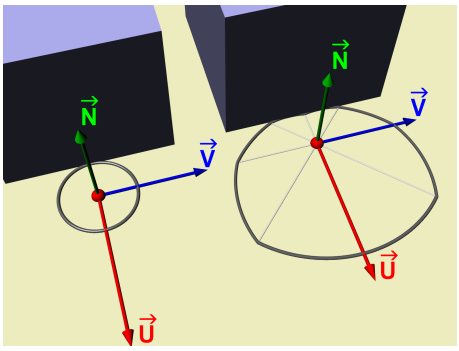

Figure 3: 8 *pseudo-elliptic zones define the influence zone of a record (red balls). Here, the record is near a cube; a circular zone (on the left) is compared to an adaptive record zone (on the right).*

To detect if a point p is covered by the influence zone of a record R, we first test if $|x_p - x_R| < d_{k_{max}}$ with k_{max} being the longest edge of the 8 sub-zones making up the influence zone. Otherwise, p lies outside the zone. The projection of p onto the tangent plane placed at R is necessarily located between two of the eight axes associated with the record. Then, as seen in figure 4, it is possible to find P'' and using equation 10. A function similar to the original Ward *et al.* [WRC88] weighting function $w_R(P)$ can be used:

$$w_R(P) = \left(\frac{|\overrightarrow{RP}|}{|\overrightarrow{RP''}|} + \sqrt{1 - \vec{n}_p \cdot \vec{n}_R} \right)^{-1} - 1 \qquad (11)$$

The accuracy parameter a is set to 1 in equation 11 because the detection of the closest object is more accurate as the length of a footprint axis is a minimum distance. As in the Irradiance Caching method, the record R is rejected if $w_R(P) \leq 0$.

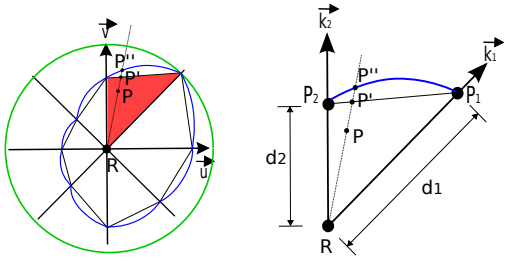

Figure 4: *p is necessarily between 2 axes k_1 and k_2 of the zone assigned to R. The green circle of radius d_{max} (the longest axis) represents the first test used to reject a record.*

In the Cornell Box scene shown in figure 5, our new record's influence zones are compared to circular zones. Both methods use the minimum distance from close objects and exactly the same criterion (same number of paths for computing irradiance and $a = 1$). Direct contributions are not considered. We observe a large gain in terms of number of records, which results in a smaller computation time since each record requires the same computation time for both methods. As described before, with circular zones, records are concentrated on edges and corners. Adaptive zones better adapt to the geometry and result in a smaller total number of records (see grey scale color pictures in figure 5 for the records distribution). The density is lower in the corners and edges for a similar result. However, some artifacts may still exist at the corners because of the ray leaking phenomenon (see [KGW*08]). In what follows, a solution to this issue is proposed allowing the integration of direct contributions in the records' data structure as well.

5. Gradients Computation

The use of new adaptive influence zones requires to reconsider the way a gradient has to be computed. To this end, this section presents a novel method of computing translational and rotational gradients for each of the 8 record axes.

5.1. Computing Irradiance along Axes

Gradients help determining how the irradiance changes at each point within the record zone. Computing gradients can be achieved by making use of equivalent point light sources described in section 3 and of equation 9. The irradiance due to these point light sources can then be computed at every point lying on each axis of the record's influence zone. For point p at position p_x on an axis k of record R, the irradiance is expressed using equation 12:

$$E(p_x) = \frac{1}{N} \sum_{i=1}^{N} \frac{I_i^{eq}(\overrightarrow{\omega_i'}) \cdot (\overrightarrow{n_{p_x}} \cdot \overrightarrow{\omega_i'})}{d_i'^2} V(x_i, p_x), \qquad (12)$$

where $\overrightarrow{\omega_i'}$ is the direction from equivalent point source i to p_x, $I_i^{eq}(\overrightarrow{\omega_i'})$ is its equivalent intensity, d_i is the distance between p_x and the equivalent point light source i. $V(x_i, p_x)$ represents the visibility between x_i, the position of the equivalent point light source i, and p_x. As shown in figure 6, two types of contribution are considered: direct and indirect contributions from light sources and objects respectively. Notations introduced in figure 6 will be used here (*i.e.* $\overrightarrow{\omega_i'} = \overrightarrow{SR}$ and $\overrightarrow{\omega_i'} = \overrightarrow{SR_k'}$). For indirect contribution, if a diffuse material is considered for S then the radiance is uniform. Consequently, the reflected intensity is a function of the cosine of α and the unit solid angle. Thus, for all point R_k' lying on the axis k, the equivalent intensity can be estimated depending on the cosines ratio of α' to α:

$$I_i^{eq}(\overrightarrow{SR_k'}) = I_i^{eq}(\overrightarrow{SR}) \cdot \frac{cos(\alpha')}{cos(\alpha)} \qquad (13)$$

where α is the angle between $\overrightarrow{SR}$ and the sender normal.

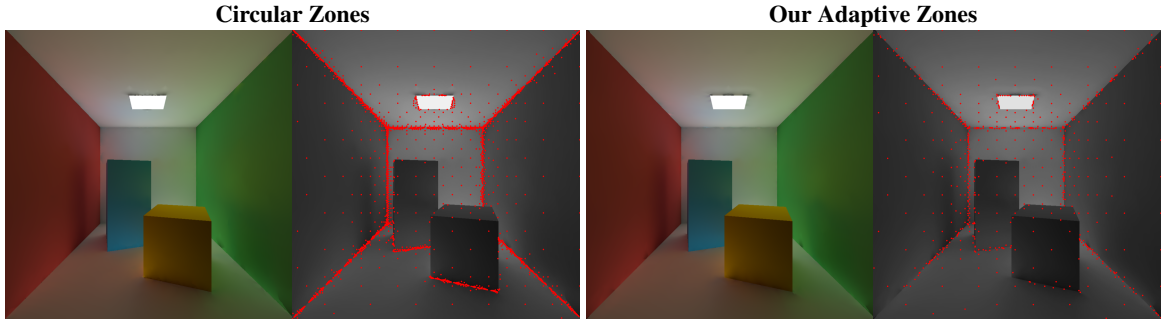

Figure 5: *Circular zones (left) and our adaptive zones (right) with indirect illumination only; both methods use the minimum distance from close objects and exactly the same criteria. Computation time for one record is the same for both methods.*

When direct contribution is also stored in a record, it is necessary to consider real point light sources. For a real point light source (radiance is not uniform), the equivalent point light source intensity is proportional to the real light source intensity $I(\overrightarrow{SR})$ of source i. Thus the equivalent intensity at R'_k on the axis k is approximated by the ratio of $I(\overrightarrow{SR'_k})$ to $I(\overrightarrow{SR})$:

$$I_i^{eq}(\overrightarrow{SR'_k}) = I_i^{eq}(\overrightarrow{SR}) \cdot \frac{I(\overrightarrow{SR'_k})}{I(\overrightarrow{SR})} \tag{14}$$

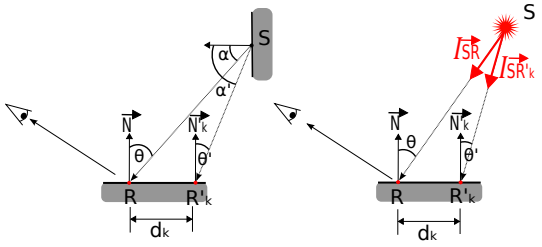

Figure 6: *2 types of equivalent points sources (direct on the right and indirect on the left) are detected for computing irradiance in R. They are used to determine the irradiance at R'_k on the axis k.*

Using these two formulations, irradiance can now be computed at each point of each axis.

5.2. Translational gradients

5.2.1. Computation

A new method is used to compute gradients in order to integrate the direct illumination contributions: a second-order Taylor's expansion is introduced for interpolation in the adaptive influence zones. It requires solving an equation expressing the irradiance $S_k(x)$ at a point x on the axis k of the record R:

$$\tilde{E}_k(x) = E(R) + G_k^1 \Delta_k + G_k^2 \frac{\Delta_k^2}{2} \tag{15}$$

where $\Delta_k = |x - R|_k$ and G_k^1 and G_k^2 are useful interpolation gradients stored in the record and so $\tilde{E}_k(x) = E(x)$. As mentioned before, it is possible to know the value of $\tilde{E}_k(x)$ (the irradiance at point x) for all points x on an axis k thanks to equation 12. If we choose x' at a distance $\Delta' = d_k/2$ from R with d_k the length of the axis k (minimum distance to the closest object along axis k) and x at a distance $\Delta = d_k$, we obtain the following set of equations:

$$\begin{cases} E(x') &= E(R) + G_k^1 \Delta' + G_k^2 \Delta'^2/2 \\ E(x) &= E(R) + G_k^1 \Delta + G_k^2 \Delta^2/2 \end{cases} \tag{16}$$

Applying equation 12 to express $E(x')$ and $E(x)$ results in a system whose solutions are G_k^1 and G_k^2. The approximated irradiance can now be quickly determined at every point on each axis k of each record using equation 15.

5.2.2. Interpolation

As stated in section 4, a point P belongs to the influence zone of a record R if its projection onto the tangent plane of R lies within a curved sub-zone defined by 2 axes k_1 and k_2 (see figure 4). Equation 15 is expressed for each axis k_n as:

$$\begin{aligned} \tilde{E}_{k_n}(P) \;=\; & E(R) + G_k^1 \frac{|\overrightarrow{RP}|}{|\overrightarrow{RP'}|} \cdot d_{k_n} \\ & + \frac{1}{2} G_k^2 \left(\frac{|\overrightarrow{RP}|}{|\overrightarrow{RP'}|} \cdot d_{k_n} \right)^2 \end{aligned} \tag{17}$$

We assume that the lighting change is linear, so the irradiance at P due to R is evaluated using a linear interpolation of the 2 values $\tilde{E}_{k_1}(P)$ and $\tilde{E}_{k_2}(P)$:

$$E(P) = w_R(P) \cdot \left[\tilde{E}_{k_1}(P) \cdot (1-\tau) + \tau \cdot \tilde{E}_{k_2}(P)\right], \qquad (18)$$

where after Thales theorem $\tau = |\overrightarrow{P_1P'}|/|\overrightarrow{P_1P_2}|$. We define $Tr(P) = \left[\tilde{E}_{k_1}(P) \cdot (1-\tau) + \tau \cdot \tilde{E}_{k_2}(P)\right]$ as the translational term for the point P. The above assumption provides good results in practice.

5.3. Rotational gradients

5.3.1. Computation

Rotational gradients help determining the irradiance variation due to normal perturbation. With adaptive record zones, it is possible to use the classical Ward and Heckbert's rotational gradients [WH92]. However, these gradients are not efficient when direct illumination is stored in the records, because direct illumination is a high frequency signal. Consequently, a new method of rotational gradient computation is proposed in this section. Our approach computes eight Rotational gradients, each corresponding to a normal perturbation of maximum angle α_{max} along one axis associated with the influence zone of a record. Using equation 7, irradiance can be computed for a normal vector $\overrightarrow{n'_k}$, a perturbation of the normal vector $\overrightarrow{n_R}$ of the record R along the axis k:

$$E(\overrightarrow{n'_k}) = \frac{1}{N} \sum_{i,(\overrightarrow{n'_k}\cdot\overrightarrow{\omega_i})>0}^{N} \frac{I_i^{eq}(\overrightarrow{\omega_i}) \cdot (\overrightarrow{n'_k} \cdot \overrightarrow{\omega_i})}{d_i^2} \qquad (19)$$

A second-order Taylor expansion is a function of Δ_θ, the angular distance between the normal vector $\overrightarrow{n_R}$ of the record R, and $\overrightarrow{n'_k}$, the perturbed normal vector by an angle α_{max}:

$$\tilde{E}(\alpha_{max}) = E(R) + G_k^{rot1}\Delta_\theta + \frac{1}{2}G_k^{rot2}\Delta_\theta^2 \qquad (20)$$

Similarly to translational gradients and equation 16, we get:

$$\begin{cases} E\left(\frac{\alpha_{max}}{2}\right) &= E(R) + G_k^{rot1}\frac{\Delta_\theta}{2} + \frac{1}{2}G_k^{rot2}\left(\frac{\Delta_\theta}{2}\right)^2 \\ E(\alpha_{max}) &= E(R) + G_k^{rot1}\Delta_\theta + \frac{1}{2}G_k^{rot2}(\Delta_\theta)^2 \end{cases} \qquad (21)$$

5.3.2. Interpolation

As in the case of translation interpolation, when a point P lies within the influence zone of the record R, then its orthogonal projection onto the tangent plane of R lies necessarily between two successive axes. Let us suppose that these axes are $\overrightarrow{k_1}$ and $\overrightarrow{k_2}$. The normal vector $\overrightarrow{n_P}$ at P also

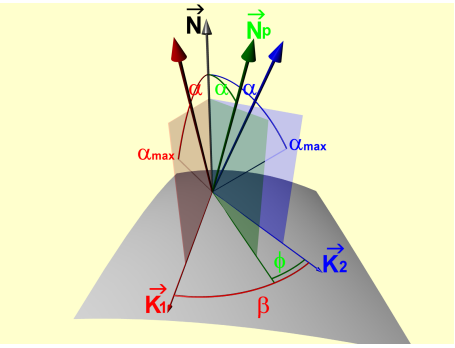

Figure 7: *Rotational interpolation. ϕ is the angle between the ($\vec{N}_R$,$\vec{N}_P$) and the ($\vec{k}_2$,$\vec{N}_R$) planes.*

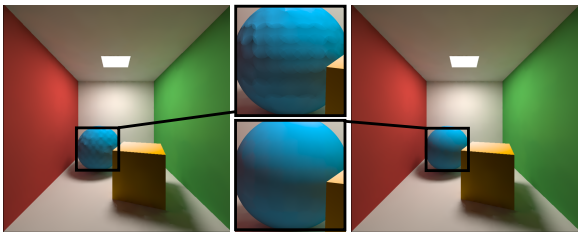

Figure 8: *Ward and Heckbert's rotational gradients (left) and our rotational gradients (right) with direct and indirect illumination; both methods give the same number of records (around 150 records on the sphere). In our gradients, α_{max} is fixed to 20 degrees.*

ranges between two normal perturbation vectors on each axis (see figure 7). As in the case of translation interpolation, we make the same empirical assumption regarding irradiance changes: we consider them as linear between two successive axes. The rotational term $Rot(P)$ may be written as:

$$Rot(P) = \left[\frac{\phi}{\beta} \cdot \tilde{E}_{k_1}(\alpha) + \left(1 - \frac{\phi}{\beta}\right) \cdot \tilde{E}_{k_2}(\alpha)\right], \qquad (22)$$

where ϕ is the angle between the projection plane of normal $\overrightarrow{n_R}$ and $\overrightarrow{k_2}$, and $\beta = \pi/4$ is the angle between $\overrightarrow{k_1}$ and $\overrightarrow{k_2}$ (see figure 7). The interpolation of irradiance at a point P within the influence zone of the record R using rotational gradients is expressed as:

$$E(P) = w_R(P) \cdot (Tr(P) + Rot(P)) \qquad (23)$$

Doing a slight modification of the weighting function which must tend toward 0 when α tends toward α_{max} to avoid a discontinuity at the border of the influence area (see [KGW*08] and section 3), we get:

$$w_R(P) = \left(\frac{|\overrightarrow{RP'}|}{|\overrightarrow{RP}|} - 1 \right) \cdot \left(1 - \frac{\alpha}{\alpha_{max}} \right) \qquad (24)$$

In fact, we replace equation 11 by equation 24 because we allow a rotational perturbation of an angle α ranging between 0 and α_{max} and assume that the weights vary linearly with the angle α.

Figure 8 shows an example where Ward and Heckbert's rotational gradients are not accurate enough when direct contributions are stored in the records. Both methods use adaptive records described in section 4 as well as translational gradients described in the previous section. Our rotational gradients provide smoother results on curved surfaces.

6. Size of Zones Based on Interpolation Accuracy

Storing direct lighting into records saves a lot of time in case of complex lighting conditions. However, despite the quadratic gradients interpolation, some lighting situations, such as the presence of shadows, can be sources of artifacts. To cope with these cases, it is possible to adjust the size of a record's influence zone depending on the irradiance change over this zone. The higher the irradiance change along an axis k_i of the zone, the smaller the length of this axis. In this paper, the interpolation quality is controlled during the computation of a record (figure 9). Ray tracing is used to compute the irradiance of a record R together with its initial length axes (size of the influence zone for each axis on the tangent plane associated with R). For each axis k_i, the translational gradients are computed at a point R' at a distance d_i (the detected distance) from R and at a point R'' at distance $d_i/2$. Two other intermediate points are used: R_1 at distance $d_i/4$ and R_2 at distance $3d_i/4$. The actual values for these 2 points are determined using equation 12 and the interpolated values $E_{interp}(R_1)$ and $E_{interp}(R_2)$ using the current translational gradients computed using the irradiances evaluated at points R' and R''. These values are then compared: if the error $\frac{E(R_1)-E_{interp}(R_1)}{E(R_1)}$ (or $\frac{E(R_2)-E_{interp}(R_2)}{E(R_2)}$ for point R_2) exceeds a threshold ρ set by the user, then the interpolation quality is not satisfactory. In this case, the process starts again with a new smaller distance d_i. We reuse the values previously calculated (see figure 9) to speed up computation. This technique has the effect of increasing the concentration of records in areas where irradiance gradients are not valid for interpolation purpose.

7. Results and Discussion

General Comments on Results

In physics-based lighting design, light/material interaction must be simulated as accurately as possible. Results must of course be consistent with the physical reality. All results presented in this paper are obtained without any restriction on the number of indirect bounces. The physical quantities, radiance and irradiance, are represented by spectra of 40 wavelengths. The light sources and their photometric features are defined by *IES* (Illuminating Engineering Society of North America) and stored in specific files (see standard [IES02]).

The *Adaptive Records for Irradiance Caching* method presented in this paper (called *ARIC* from now on) has been compared with several other techniques. The classical BPT (*Bidirectional Path Tracing* [VG94, Vea98] and [LW93]) was first considered. Then, two versions of the Irradiance Caching algorithm have been used. A first method (called *MDIC*) relies on circular zones and uses a minimum distance for computing the weighting coefficients. A second method (called *HMDIC*) uses circular influence zones, harmonic mean distance (rather than minimum distance), *neighbor clamping* and the *adaptive caching* test proposed in [KBPv06]. When applied to a record R, being part of a set S of records whose influence zones cover a visible point P, this test does not consider R as a contributor (amongst the records in S) to the irradiance at P if it causes a discontinuity when evaluating the irradiance at R through interpolation. It is of course possible to include this optimization in our method. Nevertheless, the results presented in this section have been obtained without this optimization. For all the other Irradiance Caching methods, direct illumination is not stored in the records, rather it is recomputed in the rendering pass.

All the methods that are compared to ours are integrated into a same renderer. Results have been obtained on an Intel Core 2 Q9550 (2.83 GHz) with 4 GByte of RAM (using a single core) running on a 64 bits version of linux operating system.

Cornell Box scene results

Figure 10 shows renderings of the well-known Cornell Box scene with the four above mentioned methods. The rendering resolution is 800×800 pixels. Rendering statistics are given in table 1. BPT has been calibrated to take approximately the same time as our new method (12 samples per pixel). The Irradiance Caching methods have been run with the objective of providing the same perceptive results. To evaluate irradiance similar convergence criteria for Monte Carlo sampling are used for all methods. Direct contributions are recomputed for all pixels (for the *MDIC* and *HMDIC* methods), 30 shadow rays are cast to sample the area light source. Parameter a is set to 0.4 for the *HMDIC* method while it is equal to 1 for the *MDIC* method. Note that with the *ARIC* method, the records are concentrated around the shadow to better capture the abrupt illumination changes. For the same image quality, our adaptive approach outperforms the other methods in terms of rendering time.

Figure 11 shows the difference between an image generated with our method and a reference image. The image with a 800×800 pixel resolution is computed with 1300 records. The average relative difference on all non-zero points is

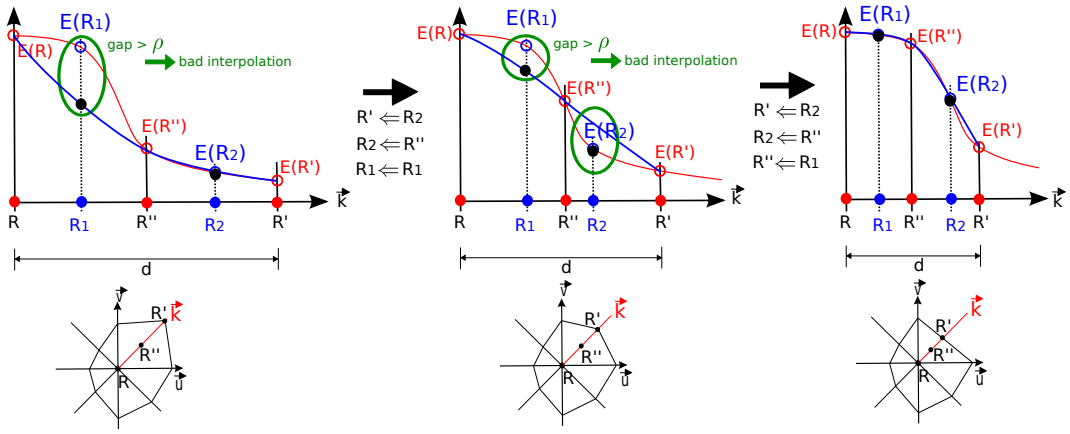

Figure 9: *Intermediate points (R_1 and R_2) are used to determine if the interpolation (the blue curve) is accurate with control points R' and R''. Interpolation is inaccurate if the relative difference between interpolated values for R_1 and R_2 (represented by black dots) and the actual values $E(R_1)$ and $E(R_2)$ (the blues circles) exceeds the threshold ρ fixed by user.*

1%. Most errors are located on the edges of the shadows: a smaller threshold ρ_{mini} could solve the problem. Other errors located on surfaces can be corrected by a better Monte Carlo sampling.

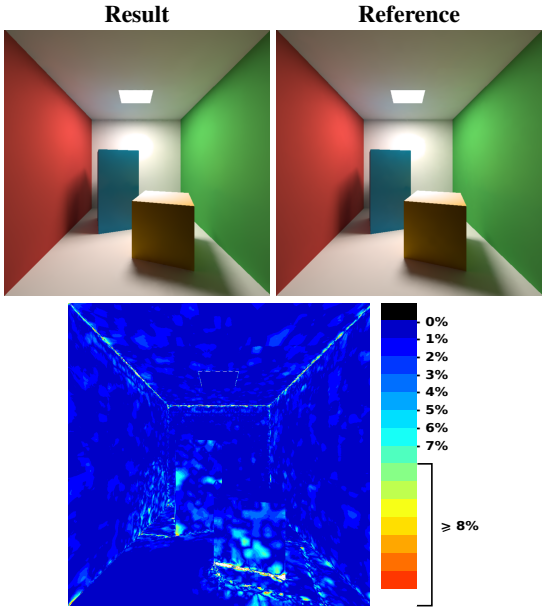

Figure 11: *Relative difference between the image calculated with our method and a reference picture: the mean relative difference is* 1%.

Llama, torus and grid scene results

The scene (Llama, torus and grid) shown in figure 12 is composed of curved surfaces and exhibits high-frequency illumination. High-frequency illumination due to small area light sources is difficult to compute with Irradiance Caching because some effects can be missed. Monte Carlo sampling seems to be better because few shadow rays are needed to capture direct illumination. Storing direct illumination in the records could be less interesting for this kind of scene. Fig-

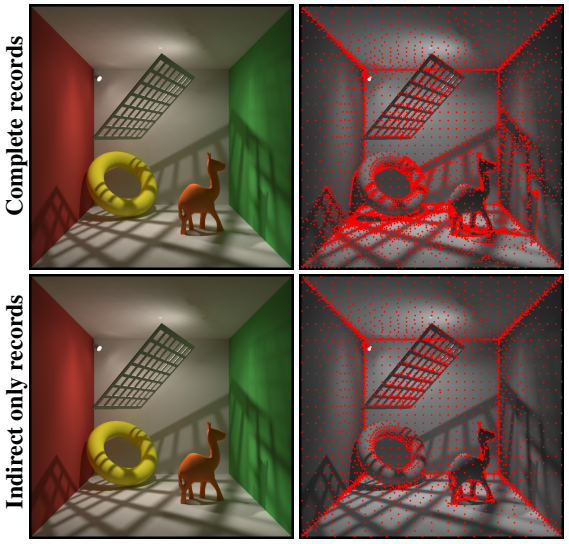

Figure 12: *Scene composed of curved surfaces and exhibiting high-frequency illumination. The first row is computed with adaptive records storing direct and indirect illuminations while the second row with adaptive records storing only indirect illumination (direct illumination is recomputed in a rendering pass). The images of the second row are of better quality but requires twice the computation time*

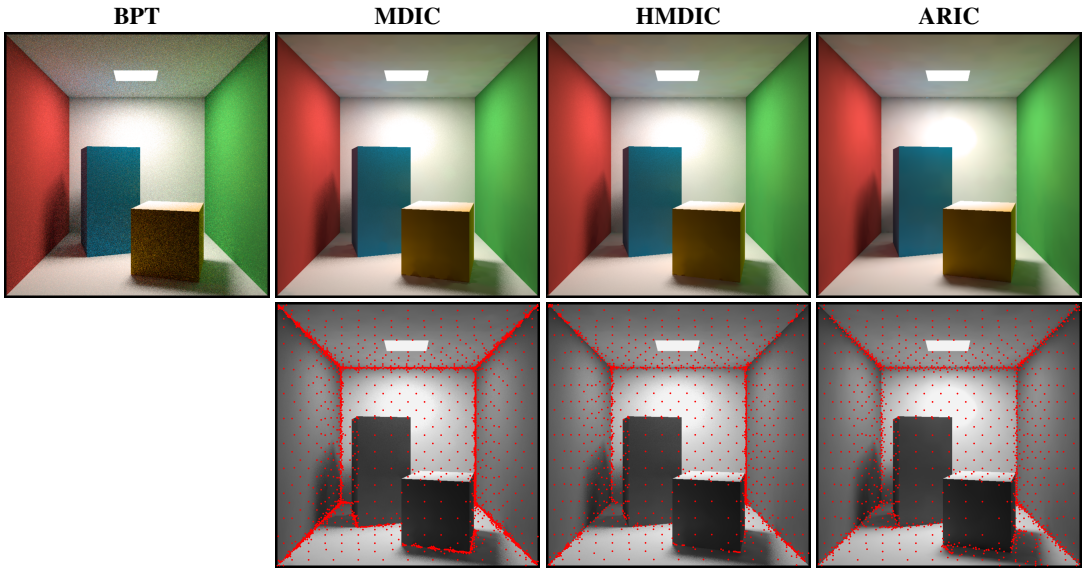

Figure 10: *Four views of the Cornell Box scene rendered with four different methods. With our adaptive records method (on the right), the record density is higher in areas of high gradients (here along the shadow edges)*

.

	# records created	Cache filling	Rendering pass	Total time
BPT	-	-	81	81
MDIC	2798	200	47	247
HMDIC	1443	102	47	149
ARIC	**1371**	**74**	**6**	**80**

Table 1: *Rendering statistics for the Cornell Box scene: cache filling is the time (expressed in seconds) spent for creating records and storing them in the cache. Rendering pass is the time spent for creating the final pictures (interpolation from the cached records and evaluation of the direct contributions for the MDIC and HMDIC methods). BPT needs only a Rendering Pass.*

ure 12 compares two version of our *ARIC* : (1) direct and indirect illuminations are stored in the records, (2) only indirect contributions are stored. The same parameters are used for both versions. The direct contributions are recomputed for version 2. As our method adapts the cache density in regions of high illumination change, the results of figure 12 show two different records distributions for the two versions. For version 1, the records are concentrated around shadow edges. In this example, more records are needed for version 1 (6016 records) than for version 2 (3241 records) because of the high frequency illumination changes. However, version 1 takes twice less time than version 2 for exactly the same parameters. Consequently, version 1 of our method performs well in case of strong indirect or direct illumination changes.

Figure 13 shows the impact of the threshold ρ on a close-up view of the scene shown in figure 13 and computed with our *ARIC* method where the records store direct and indirect illuminations. Both results are computed with the same parameters, only the threshold ρ is different. A restrictive threshold (for example 1%) allows to concentrate records density around strong illumination gradients while a permissive threshold (for example 10%) can miss this kind of effects. In case of complete records with a low threshold in a scene composed of high frequency illumination, interpolation artifacts appear around shadow edges. A restrictive threshold solves the problem by adding records around high frequency illumination effects.

Villa Arpel scene results

Figure 14 shows two close-up views of the villa Arpel scene generated with our method and the two other Irradiance Caching methods. The image in Figure 15 represents the complete villa Arpel scene generated with our method. The images resolutions are 800×800 pixels for the close-up views and 1920×1080 pixels (full high definition) for figure 15. The rendering results are given in table 2 for the two close-up views. The results have been obtained with com-

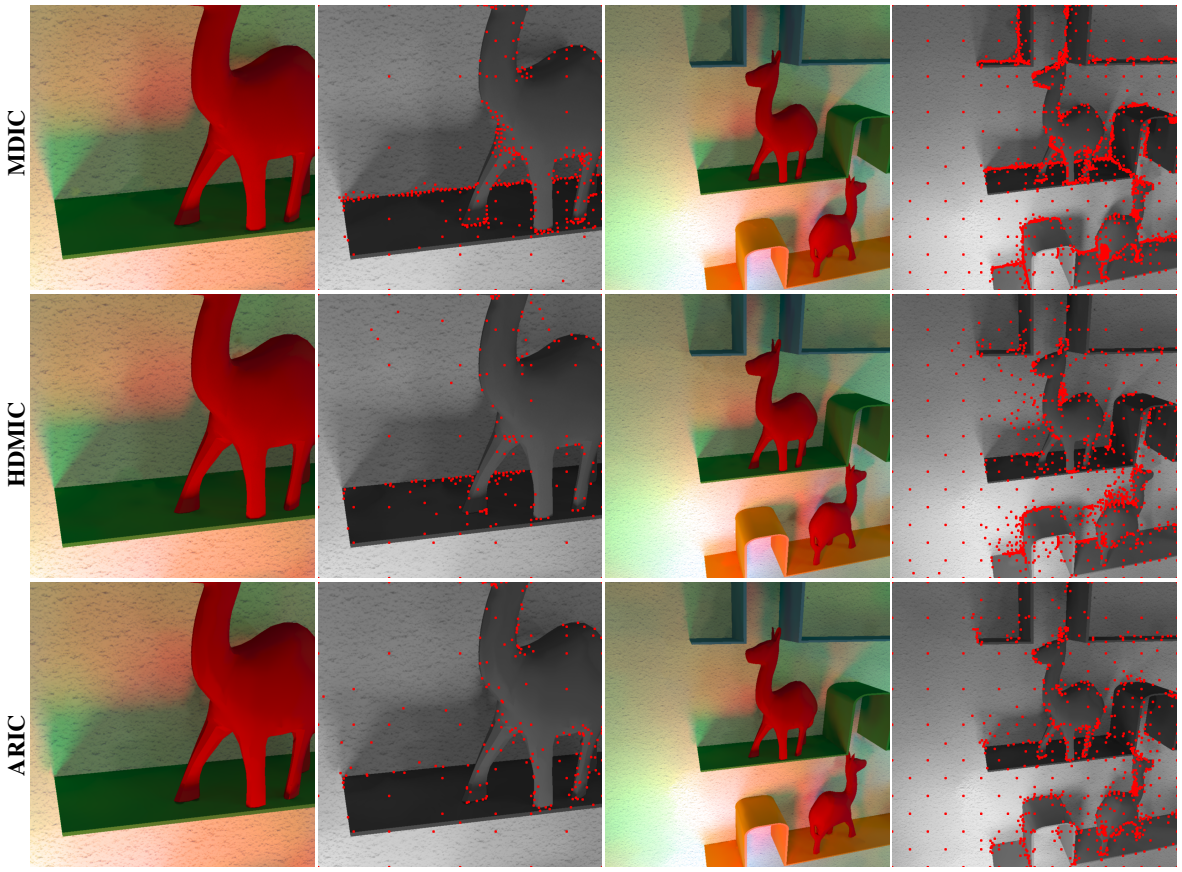

Figure 14: *2 close-up views of the villa Arpel scene rendered with different methods. The cached points are represented in red color.*

		# records created	Cache filling	Total time
first close-up	*MDIC*	380	7m. 59s.	2h. 47m.
	HMDIC	197	3m. 3s.	2h. 40m. 30s.
	ARIC	**209**	**3m. 10s.**	**3m. 16s.**
second close-up	*MDIC*	1247	18m. 15s.	2h. 34m. 6s.
	HMDIC	798	12m. 9s.	2h. 56m. 24s.
	ARIC	**523**	**6m. 50s.**	**6m. 57s.**

Table 2: *Rendering statistics for results seen in figure 14.*

plex artificial lighting conditions (11 light sources made up of 760 polygons with different colored spectra) and overcast sky for daylighting conditions obtained with a standard CIE sky model (see CIE standards in [CIE96]). Given the high number of light sources, it is particularly interesting to store direct contributions in records to save computation time. All the methods have the same parameters. Direct contributions are computed using 4000 shadow rays for all the artificial light sources. Parameter a is set to 0.4 for the *HMDIC* method and 1 for the methods using a minimum distance. The image in Figure 15 has been computed in 1 hour and 6 minutes (and 33 seconds for the rendering pass) with 6962 cached points, while the *HMDIC* method took 1 hour and 30 minutes (and 8 hours for the rendering pass) with 6656 records. As shown in figure 14, our method provides a better record distribution while reducing the total number of records. In addition, for the same image quality, our ARIC method outperforms the other methods.

Figure 15: *Villa Arpel scene rendered with our adaptive records method.*

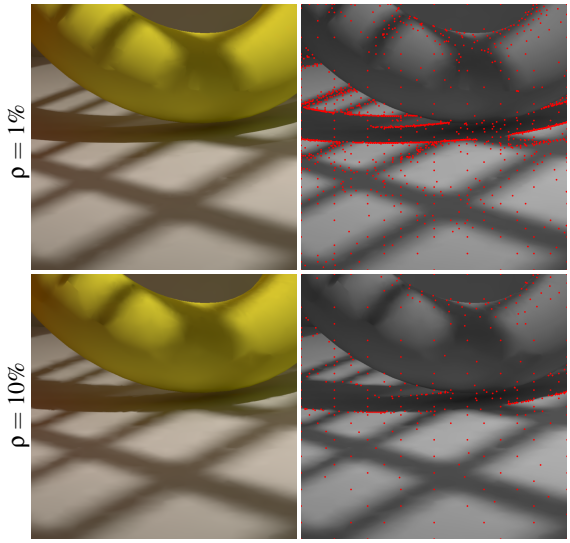

Figure 13: *Impact of the threshold* ρ *on the quality of the high-frequency illumination effects and on the records density.*

8. Conclusion

We have presented a new method to adaptively compute records in an Irradiance Caching algorithm. We have proposed a new record footprint specifically built to account for both geometrical and irradiance changes over surfaces. Our approach prevents the cache from being too dense at the edges and borders of the scene's objects. Including irradiance changes in the translational gradient computations leads to a denser cache in areas subject to large irradiance changes. It thus compensates for any inaccurate interpolation appearing with complex changes of irradiance. Such a feature allowed us to successfully store direct illumination in the records and to significantly speed up the rendering pass. The storage of the direct irradiance is interesting especially in the case of complex lighting conditions such as combined daylighting and artificial lighting. Likewise, large area light sources or animation rendering could also benefit from this characteristic because direct lighting does not have to be computed for successive frames but rather interpolated.

In a future work, the Radiance Caching scheme will be investigated in order to include in the adaptive record method more directional information such as glossy reflection.

9. Acknoledgements

This work has been carried out with the financial support of ANRT within the framework of the CSTB Institut Carnot 2009 research program.

References

[Adr89] ADRIAN W.: Visibility of targets: Model for calculation. *Lighting Research and Technology 21*, 4 (1989), 181–188.

[AFO05] ARIKAN O., FORSYTH D. A., O'BRIEN J. F.: Fast and detailed approximate global illumination by irradiance decomposition. In *Proceedings of ACM SIGGRAPH 2005* (2005), ACM Press. 2

[BGB08] BROUILLAT J., GAUTRON P., BOUATOUCH K.: Photon-driven irradiance cache. *Computer Graphics Forum 27* (2008), 1971–1978. 2

[BWG03] BALA K., WALTER B., GREENBERG D. P.: Combining edges and points for interactive high-quality rendering. *ACM Trans. Graph. 22*, 3 (2003), 631–640. 3

[CIE96] CIE: Spatial distribution of daylight - cie standard overcast sky and clear sky. In *CIE S 003/E-1996* (1996). 11

[GKPB04] GAUTRON P., KŘIVÁNEK J., PATTANAIK S. N., BOUATOUCH K.: A novel hemispherical basis for accurate and efficient rendering. In *Rendering Techniques 2004, Eurographics Symposium on Rendering* (June 2004), pp. 321–330. 3

[Gre03] GREEN R.: Spherical harmonic lighting: The gritty details. *Archives of the Game Developers Conference* (March 2003). 3

[HMS09] HERZOG R., MYSZKOWSKI K., SEIDEL H.-P.: Anisotropic radiance-cache splatting for efficiently computing high-quality global illumination with lightcuts. In *Computer Graphics Forum (Proc. EUROGRAPHICS)* (München, Germany, 2009), Stamminger M., Dutré P., (Eds.), vol. 28, Wiley-Blackwell, pp. 259–268. 3, 4

[IES02] IESNA: Iesna standard file format for the electronic transfer of photometric data and related information. In *LM-63-02* (September 2002). 8

[Jen96] JENSEN H. W.: Global illumination using photon maps. In *Proceedings of the eurographics workshop on Rendering techniques '96* (London, UK, 1996), Pueyo E. X., Schröder P., (Eds.), Springer-Verlag, pp. 21–30. 2

[Jen01] JENSEN H. W.: *Realistic image synthesis using photon mapping*. A. K. Peters, Ltd., Natick, MA, USA, 2001. 2

[KBPv06] KŘIVÁNEK J., BOUATOUCH K., PATTANAIK S. N., ŽÁRA J.: Making radiance and irradiance caching practical: Adaptive caching and neighbor clamping. In *Rendering Techniques 2006, Eurographics Symposium on Rendering* (Nicosia, Cyprus, June 2006), Akenine-Möller T., Heidrich W., (Eds.), Eurographics Association, Eurographics Association. 1, 2, 3, 4, 8

[KGBP05] KŘIVÁNEK J., GAUTRON P., BOUATOUCH K., PATTANAIK S.: Improved radiance gradient computation. In *SCCG '05: Proceedings of the 21st spring conference on Computer graphics* (New York, NY, USA, 2005), ACM Press, pp. 155–159. 3

[KGPB05] KŘIVÁNEK J., GAUTRON P., PATTANAIK S., BOUATOUCH K.: Radiance caching for efficient global illumination computation. *IEEE Transactions on Visualization and Computer Graphics 11*, 5 (September/October 2005), 550–561. http://graphics.cs.ucf.edu/RCache/index.php. 3

[KGW*08] KŘIVÁNEK J., GAUTRON P., WARD G., JENSEN H. W., TABELLION E., CHRISTENSEN P. H.: Practical global illumination with irradiance caching. *ACM SIGGRAPH '08 Class* (2008). http://www.graphics.cornell.edu/~jaroslav/papers/2008-irradiance_caching_class/index.htm. 1, 2, 3, 5, 7

[KŮ05] KŘIVÁNEK J.: *Radiance Caching for Global Illumination Computation on Glossy Surfaces*. Ph.d. thesis, Université de Rennes 1 and Czech Technical University in Prague, December 2005. 3

[LW93] LAFORTUNE E. P., WILLEMS Y. D.: Bi-directional path tracing. In *Proceedings of Third International Conference on Computational Graphics and Visualization Techniques (Compugraphics '93)* (Alvor, Portugal, 1993), Santo H. P., (Ed.), pp. 145–153. 3, 8

[SiKDM05] SMYK M., ICHI KINUWAKI S., DURIKOVIC R., MYSZKOWSKI K.: Temporally coherent irradiance caching for high quality animation rendering. In *The European Association for Computer Graphics 26th Annual Conference EUROGRAPHICS 2005* (Dublin, Ireland, 2005), vol. 24 of *Computer Graphics Forum*, Blackwell, pp. 401–412. 2

[SM02] SMYK M., MYSZKOWSKI K.: Quality improvements for indirect illumination interpolation. In *Proceedings of the International Conference on Computer Vision and Graphics* (2002). 2

[TL04] TABELLION E., LAMORLETTE A.: An approximate global illumination system for computer generated films. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Papers* (New York, NY, USA, 2004), ACM, pp. 469–476. 1, 2, 3, 4

[VALBW06] VELÁZQUEZ-ARMENDÁRIZ E., LEE E., BALA K., WALTER B.: Implementing the render cache and the edge-and-point image on graphics hardware. In *GI '06: Proceedings of the 2006 conference on Graphics interface* (Toronto, Ont., Canada, Canada, 2006), Canadian Information Processing Society, pp. 211–217. 3

[Vea98] VEACH E.: *Robust monte carlo methods for light transport simulation*. PhD thesis, Stanford University, Stanford, CA, USA, 1998. Adviser-Leonidas J. Guibas. 3, 8

[VG94] VEACH E., GUIBAS L. J.: Bidirectional estimators for light transport. In *Eurographics Rendering Workshop* (Darmstadt, Germany, June 1994), Springer-Verlag, New York, pp. 147–162. 3, 8

[WABG06] WALTER B., ARBREE A., BALA K., GREENBERG D. P.: Multidimensional lightcuts. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers* (New York, NY, USA, 2006), ACM, pp. 1081–1088. 3

[WDG02] WALTER B., DRETTAKIS G., GREENBERG D.: Enhancing and optimizing the render cache. In *Proceedings of the Eurographics Workshop on Rendering* (June 2002), Debevec P., Gibson S., (Eds.), Eurographics, ACM Press. 3

[WDP99] WALTER B., DRETTAKIS G., PARKER S.: Interactive rendering using the render cache. In *Rendering Techniques (Proceedings of the Eurographics Workshop on Rendering)* (New York, NY, Jun 1999), Lischinski D., Larson G., (Eds.), vol. 10, Springer-Verlag/Wien, pp. 235–246. 3

[WFA*05] WALTER B., FERNANDEZ S., ARBREE A., BALA K., DONIKIAN M., GREENBERG D. P.: Lightcuts: a scalable approach to illumination. *ACM Trans. Graph. 24*, 3 (2005), 1098–1107. 3

[WH92] WARD G., HECKBERT P.: Irradiance gradients. In *Eurographics Rendering Workshop* (May 1992), pp. 85–98. 1, 2, 3, 7

[WRC88] WARD G. J., RUBINSTEIN F. M., CLEAR R. D.: A ray tracing solution for diffuse interreflection. In *SIGGRAPH '88: Proceedings of the 15th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1988), ACM, pp. 85–92. 1, 2, 3, 5