Spatio-Temporal Photon Density Estimation Using Bilateral Filtering

Markus Weber, Marco Milch, Karol Myszkowski,
Kirill Dmitriev, Przemyslaw Rokita, and Hans-Peter Seidel
MPI Informatik, Saarbrüecken, Germany*

Abstract

Photon tracing and density estimation are well established techniques in global illumination computation and rendering of high-quality animation sequences. Using traditional density estimation techniques it is difficult to remove stochastic noise inherent for photon-based methods while avoiding overblurring lighting details. In this paper we investigate the use of bilateral filtering for lighting reconstruction based on the local density of photon hit points. Bilateral filtering is applied in spatio-temporal domain and provides control over the level-of-details in reconstructed lighting. All changes of lighting below this level are treated as stochastic noise and are suppressed. Bilateral filtering proves to be efficient in preserving sharp features in lighting which is in particular important for high-quality caustic reconstruction. Also, flickering between subsequent animation frames is substantially reduced due to extending bilateral filtering into temporal domain.

1. Introduction

The rendering of high-quality CG movies featuring global illumination effects is considered very expensive. The main reason of such cost is the design of global illumination algorithms which are optimized for static scenes only. In practice this means that when such algorithms are used for a dynamic scene, all computations have to be repeated from scratch even when only minor changes take place.

An important issue in terms of frame quality is temporal aliasing in animations. Many typical errors in lighting computation and reconstruction cannot be perceived by the human observer when they are coherent in temporal domain. However, they may cause unpleasant flickering and shimmering effects when this coherence is lost.

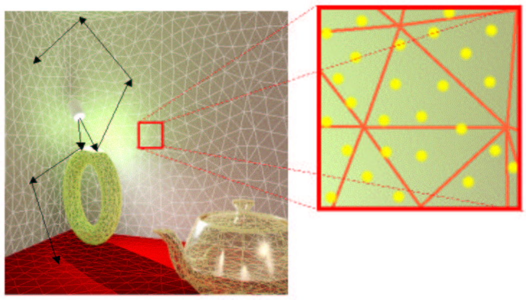

Figure 1: Density estimation photon tracing algorithm [18]: the lighting function is known implicitly as the density of photon hit points for each mesh element in the scene.

In recent years, several global illumination algorithms have been proposed that exploit temporal coherence between frames to improve the computation efficiency and reduce temporal aliasing (refer to [3] for a recent survey on this topic). In this work we build upon an algorithm proposed by Myszkowski et al. [14], in which stochastic photon tracing from light sources towards meshed surfaces in the scene was performed (refer to Figure 1). To reconstruct indirect illumination photons hitting each mesh element were collected in temporal domain for the previous and subsequent frames until a significant change of illumination due to moving objects was detected or a sufficient number of photons was collected. The perception-based Animation Quality Metric was used to decide upon the stopping condition for the photon tracing depending on the perceptibility of stochastic noise (resulting from collecting too few photons) in the reconstructed illumination. The paper by Myszkowski et al. emphasized on the perceptual aspects of steering the global illumination computation and the processing of photons

*email:{mweber,milch,karol,kirill,rokita,hpseidel}@mpi-sb.mpg.de

was limited to simple sharing those photons between frames.

In this paper, we focus on spatio-temporal photon processing whose main goal is the improvement of quality in reconstructed surface illumination by reducing the amount of spatial and temporal blur while keeping the stochastic noise below the perceivable level.

The temporal blur is a result of collecting photons in temporal domain in those scene regions in which lighting changes quickly. In such regions some photon paths computed for previous frames can be invalid for the currently processed frame due to dynamic changes in the scene. To reduce the temporal blur such invalid photons should not be considered and if too few photons are collected in temporal domain to satisfy noise error criteria, the remaining photons must be collected in spatial domain. This in turn may lead to spatial blur in the reconstructed illumination. To reduce this effect the missing photons must be collected only in the neighboring scene regions with coherent (similar) illumination. Collecting photons across the edges of high illumination contrast inherently leads to lighting patterns blurring. Similarly, in temporal domain collecting photons for abruptly changing lighting leads to temporal blur. Clearly, a photon density estimation method capable of reducing stochastic noise while detecting such abrupt spatio-temporal changes of lighting over meshed surfaces is needed.

We propose to extend traditional photon density estimation methods by using spatio-temporal bilateral filtering to reduce stochastic noise while preventing excessive blurring in reconstructed lighting. For lighting estimation of each mesh element photons collected for neighboring mesh elements in space-time domain are considered but the influence of elements with significantly different illumination is suppressed by bilateral filtering. This means that extremely noisy illumination estimates (outliers) as well as estimates resulting from abrupt changes of illumination distribution both in space and time domains are filtered out. As a result better space-time resolution of lighting patterns is obtained while at the same time spatio-temporal noise is significantly reduced. Also, the computation efficiency is significantly improved by reducing the number of photons required to achieve high-quality animations in respect to traditional frame-by-frame rendering approaches.

In the following section we briefly discuss previous work on space-time solutions for global illumination and density estimation. Then, we discuss basic characteristics of bilateral filtering in the context of our space-time density estimation problem. In Section 4 we overview our rendering algorithm which takes advantage of spatio-temporal coherence in lighting for subsequent animation frames. In Section 5 we present our algorithm of lighting reconstruction based on bilateral filtering. In Section 6 we show the results obtained using our techniques and we conclude this work in Section 7.

2. Previous Work

In the following section we discuss only those global illumination techniques which exploit temporal coherence in the computation and can handle efficiently both camera and object motion in the scene. In Section 2.2 we remind basic photon density estimation algorithms that lead to a direct rendering of photon maps.

2.1. Global Illumination for Animations

Temporal coherence in global illumination computations [3] can be considered at various levels ranging from ready to display shaded pixels to simple visibility samples shared between frames. Making use of the coherence at a higher level is generally the approach chosen for interactive applications where fast response time is a crucial factor. For example in the Render Cache technique [19] shaded pixels are re-projected from the previous frames to the current frame, and in the Shading Cache technique [16] working in the object space illumination samples are reused for mesh vertices.

Considering temporal coherence at lower levels, e.g., at the level of single photon paths, usually results in more flexibility in sharing information for many frames at once. Myszkowski et al. [14] and Dmitriev et al. [5] reuse photon hit points in their view-independent stochastic photon tracing techniques. The major drawback of those techniques is over-blurring lighting details due to simple histogram-based photon density reconstruction. Havran et al. [9] reuse complete photon paths from the light source to the eye by extending the view-dependent bi-directional path tracing algorithm into temporal domain. The method significantly improves the animation rendering efficiency in respect to its frame-by-frame counterpart. However, it still requires extensive computation (the cost depends directly on the resolution of frames) to reduce stochastic noise to non-perceivable level.

The temporal coherence can be also considered in the computation of form factors and light propagation in the framework of radiosity methods [6, 4]. However, those methods usually involve high memory costs and are inefficient in handling non-diffuse environments.

In this paper we consider the algorithm proposed in [14] and we include it into the framework of a traditional frame-by-frame animation production software Insight which is available as a plug-in for the 3D Studio Max system. Therefore, for practical reasons, we consider only previous frames while in the original algorithm [14] the following frames are also temporally processed. For the cost reasons we abandon the application of Animation Quality Metric for determining the stopping condition. Instead we use an energy-based error metric for stopping condition, which is perhaps too conservative in terms of the number of traced photons but works very robustly. Our main focus in this paper is an improvement of photon density estimation which leads to a better reconstruction of lighting details and therefore improves the overall animation quality. In the following section we briefly discuss previous work on photon density estimation.

2.2. Density Estimation

The classic density estimation problem can be stated as the construction of an estimate of the probability density function from the observed data points [15]. Thus, the reconstruction of the lighting function based on the location of photon-hitting points, which are obtained as the result of photon tracing computation, is also the density estimation problem [10].

One of the most popular density estimation techniques applied in the global illumination problem is the nearest neighbor method, which is used in the photon mapping algorithm [11]. Usually the number of nearest photons k needed to reconstruct lighting at a given point is set globally for the whole scene which may lead to excessive noise or smoothing in reconstructed lighting. Clearly, k should be decided locally, and some rules for its automatic selection must be provided. In photon mapping this problem is overcome by avoiding a direct display of lighting reconstructed using the photon map but instead a costly final gathering rendering is performed.

Myszkowski [13] attempted to display directly photon maps using a modified nearest neighbor method in which k is locally adjusted based on local lighting complexity. The region at which photons are collected is iteratively expanding until the resulting illumination estimate is significantly different than the ones obtained in the previous steps or a predefined maximum number of photon $k_{\max}$ is collected. This approach significantly reduced blur in the proximity of high-contrast discontinuities in the lighting function at expense of adding some noise in those regions.

Walter [20] considers kernel-based density estimation methods with adaptively controlled kernel width which adjusts to local illumination complexity using perceptually-motivated criteria. Walter applies statistical bias detection techniques to automatically enhance under-resolved illumination features and eliminates boundary bias using a local polynomial density estimation method.

The space-time approaches to density estimation are considered only in two works. Cammarano and Jensen [1] proposed an object space solution for motion blur in which photons are stored with time stamps and the nearest photons in space-time are considered for lighting reconstruction. Myszkowski et al. [14] uses a simple histogram method for space-time photon processing. We extend the latter work by using the bilateral filtering in order to eliminate photons with significantly different space-time distributions from local illumination estimates. We experiment also with other space-time photon processing schemes such as simple Gaussian filters. Our current framework does not require photon hit points storage, simply the number of hits is counted for each mesh element. This reduces significantly the memory requirements in respect to the framework used in [14].

3. Bilateral Filtering

Bilateral filtering was introduced in image processing and computer vision literature as an efficient tool to combat noise [17]. It is a non-linear filter which performs feature preserving smoothing in a single pass. Those two seemingly contradictory goals are achieved using the Gaussian filter $w_S(x) = e^{-x^2/2\sigma_S^2}$ in the image space with simultaneous penalizing large variations in pixel intensities using another Gaussian filter $w_I(x) = e^{-x^2/2\sigma_I^2}$. σ_S denotes the spatial filter support whose extent determines the set $\Omega(q)$ of pixels p which are centered around a pixel q. σ_I controls the influence of p on the filtered pixel intensity $\hat{I}(q)$ as a function of the difference in intensities $|I(q) - I(p)|$ in the input image for pixels p and q. The bilateral filter can be formulated as follows:

$$\hat{I}(q)=\frac{\sum\limits_{p\in\Omega(q)} w_S(\parallel q-p\parallel)w_I(I(q)-I(p))I(p)}{\sum\limits_{p\in\Omega(q)} w_S(\parallel q-p\parallel)w_I(I(q)-I(p))} \tag{1}$$

Bilateral filtering was successfully used in a number of different applications in computer graphics. Durand and Dorsey [7] proposed a novel tone mapping operator in which high-dynamic range images are decomposed into two layers (a base layer encoding large scale variations and a detail layer) and then the contrast is compressed only in the former layer while leaving the detail untouched. The base layer is obtained by means of the bilateral filter which preserves high contrast edges and removes high-spatial frequency details of lower contrast. Fleishman et al. [8] and Jones et al. [12] applied bilateral filtering to mesh denoising while preserving important features of input geometrical models. This required the extension of bilateral filter formulation from two-dimensions to manifolds in three dimensions. Choudhury and Tumblin [2] introduced trilateral filtering which additionally smoothes signals towards a sharply bounded, piecewise linear approximation. Trilateral filtering performs slightly better than its bilateral counterpart for signals with sharp changes in gradients and quickly changing signals for which too many $I(p)$ samples is treated as outliers. Those improvements are achieved at expense of significant increase in the computation cost.

In this work we consider bilateral filtering in the context of global illumination computation which is performed for mesh elements using photon density estimation. The basic bilateral filtering formulation in Equation 1 remains valid. The only difference is that instead of neighboring pixels we consider neighboring mesh elements. Thus, the distance $\parallel q-p\parallel$ is simply measured between the gravity centers of mesh elements. Since in our application the photon density estimation is performed in space-time domain we need to extend the bilateral filter by introducing the time stamps t_i for estimated illumination $I_{t_i}(q)$ and $I_{t_i}(p)$ at mesh elements q and p. Also, an additional Gaussian filter w_T with the support σ_T is added to control the temporal extent of filtering:

$$\hat{I}(q) \quad = \quad \frac{\sum\limits_{t_i=t_0}^{t_{\max}-1}\sum\limits_{p\in\Omega(q)} \omega_S\cdot\omega_I\cdot\omega_T\cdot I_{t_i}(p)}{\sum\limits_{t_i=t_0}^{t_{\max}-1}\sum\limits_{p\in\Omega(q)} \omega_S\cdot\omega_I\cdot\omega_T} \tag{2}$$

$$\begin{aligned}\omega_S &= w_S(\parallel q-p\parallel)\\ \omega_I &= w_I(I_{t_0}(q)-I_{t_i}(p))\\ \omega_T &= w_T(t_i-t_0)\end{aligned}$$

where $t_0=0$ denotes the currently processed frame and $t_{\max}$ denotes the maximum number of frames processed in the temporal domain.

In the following section we describe our lighting reconstruction algorithm utilizing space-time bilateral filtering as specified in Equation 2. We start from the outline of overall photon tracing and animation rendering procedure.

4. Rendering Overview

After preparing animation paths for all moving objects and camera using 3D Studio Max the rendering plug-in Insight is activated to perform the global illumination computation and frame rendering. The computation is initialized by tracing a pilot number of photons N_{pilot} for the first frame in the animation and estimating the resulting global accuracy of lighting simulation E_{pilot} (refer to [18] for details how to compute E). Based on E_{pilot}, the requested maximum error $E_{\max}$ and $t_{\max}$ (those two parameters are decided by the user) the number of photons per frame N_{frame} is found as:

$$N_{\text{frame}}=\frac{E^2_{\text{pilot}}}{E^2_{\max}}\cdot\frac{N_{\text{pilot}}}{t_{\max}} \tag{3}$$

This number of photons is then traced for each frame and the number of hit points for each mesh element are recorded.

Since the temporal processing requires immediately $t_{\max}$ frames, photon tracing is repeated $t_{\max}$ times for the very first frame in the animation. For each instance of the first frame N_{frame} different photons are traced and the resulting number of hit points are counted separately for all mesh elements. As the subsequent frames are computed the photon hit point counters for each instance of the first frame are one-by-one overwritten. After this initialization the processing of each subsequent frame proceeds as follows:

1. Update object positions accordingly to their animation paths.
2. Update ray tracing acceleration data.
3. Trace N_{frame} photons and count hit points for each mesh element.

4. Reconstruct the indirect lighting for the current frame using photon hit counters for $t_{\max}$ recent frames.

5. Ray trace the current frame using lighting reconstructed in the previous step.

6. Reset the number of photon hit points in the data structures for the oldest frame.

In our current implementation it is assumed that the number of mesh elements does not change for all frames in a temporally processed animation sequence, which is a reasonable assumption for rigid body animation. For animations involving changes in mesh connectivity and topology some tracking of mesh parents would be required to process them in temporal domain.

Because of fixed number of mesh elements the data structures for storing hit point counters are static and very simple. For each mesh element $t_{\max}$ counters of photon hit points are used to store the data for each frame separately. Those counters are cyclically updated when the subsequent animation frames are processed.

During ray tracing the direct illumination is computed from scratch for each frame and the indirect illumination is estimated using our spatio-temporal photon density estimation algorithm.

In the following section we focus on the lighting reconstruction using space-time bilateral filtering.

5. Lighting Reconstruction

Our density estimation algorithm is based on Equation 2. The choice of σ values used in this equation we describe in Section 5.1. This choice is very important because it decides upon the accuracy of reconstructed lighting and overall animation quality. In Section 5.2 we provide details concerning mesh illumination estimates $I_{t_0}(q)$ and $I_{t_i}(p)$ which are also used in Equation 2. Finally, in Section 5.3 we summarize our bilateral filtering algorithm and we consider its simplifications towards the traditional Gaussian space-time filtering.

5.1. Bilateral Filtering Parameters

A proper application of space-time bilateral filtering for mesh illumination processing requires reasonable values of σ for all three Gaussian filters in Equation 2. Increasing σ_T and σ_S leads to increasing the number of photons used for illumination estimate for a given mesh element. This results in the reduction of noise but may also lead to increase of the spatial and temporal blur as discussed in the Introduction section. The level of blur is controlled by the value of σ_I. The smaller σ_I the more details in reconstructed lighting possibly at expense of increasing noise. In the following sections we discuss the choice of values for σ_T, σ_S, and σ_I.

5.1.1. Temporal Domain: σ_T

The value of σ_T is decided by the user by fixing the maximum number of temporally processed frames $t_{\max}$. We set $\sigma_T = 0.5 \cdot t_{\max}$, which results in the smallest temporal weight $w_T(t_{\max -1}) \approx 0.135$. The weight values w_T are precomputed for all frames $[t_0, \ldots, t_{\max -1}]$ and stored in the lookup table.

5.1.2. Spatial Domain: σ_S

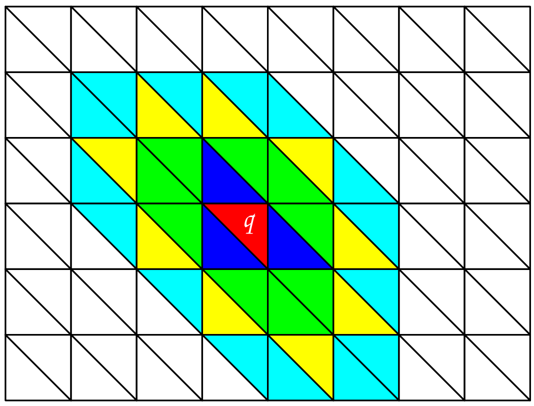

Figure 2: Ring neighborhoods for the mesh element q. In the neighborhood expansion for the mesh element q three edge-connected triangles are considered as the 1-ring neighborhood (marked in blue). The 2-ring neighborhood (marked in green) contains all remaining triangles which are vertex connected to q (usually there are 9 such triangles for a regular mesh). The 3- and 4-ring neighborhoods are composed of triangles that are edge- and vertex-connected (marked in yellow and cyan, respectively) to all triangles in the 2-ring neighborhood. Further ring neighborhoods are found in the same way by searching for edge- and vertex-connected triangles to the previous even ring neighborhoods.

The value of σ_S is selected adaptively for each processed mesh element q as a function of lighting in local neighborhood. The main objective here is to keep the stochastic noise in filtered lighting on a similar level for all mesh elements.

Let us recall that for the static scene a local measure of stochastic noise in illumination estimate can be computed as a relative mean deviation $\sigma_{\%}$ (expressed in percents) [18]:

$$\sigma_{\%} = \frac{100\%}{\sqrt{n}} \tag{4}$$

where n is the number of photons hitting the filter support. For example the suppression of stochastic noise to the level of $\sigma_{\%} = 5\%$ requires that 400 photons hit the filter support. In the derivation of Equation 4 (refer to [18]) it is assumed that the probability to hit the filter support area by a particular photon traced in the scene is sufficiently small, which usually holds well in practice.

We use the illumination accuracy in Equation 4 to select adaptively σ_S for each mesh element q. At first, the number of photons n_q hitting q for all temporally processed frames $[t_0, \ldots, t_{\max-1}]$ is computed. In order to achieve the desired illumination accuracy $\sigma_{\%}$ the condition $n < n_q$ must be satisfied. If this is not the case then subsequent neighboring mesh elements p must be processed until the total number of collected photons is bigger than n. For each processed p all photons for frames $[t_0, \ldots, t_{\max-1}]$ are always considered.

All mesh elements p that are included into the spatial filter support are organized into ring neighborhoods surrounding q (refer to Figure 2). The subsequent ring neighborhoods are considered one-by-one starting from the smallest rings until n photons are collected or the maximum number of rings to be processed is reached. The number is decided by the user to prevent excessive computation costs for very dim regions in the scene, possibly at expense of some visible noise.

The illumination accuracy measure in Equation 4 holds well for photon density estimation in static environments but in practice we extend the spatial filter support by one ring neighborhood to compensate for temporal changes in illumination and the resulting suppression of outliers by w_I. The external ring radius r_{ring} is found as the largest distance between centroids of q and p in this ring. The value of $\sigma_S = 0.5 \cdot r_{\text{ring}}$ is set, which is equivalent to the spatial filter truncation at $2\sigma_S$ as suggested in [7].

5.1.3. Influence: σ_I

The value of σ_I must adapt to the local illumination level in order to be sensitive to local lighting patterns. We set $\sigma_I = l \cdot I_{t_0}(q)$ where l is set globally for the whole scene by the user based upon the desired level of details in reconstructed lighting (limited by the mesh granularity) and the amount of tolerable stochastic noise. In the following section we describe our approach to estimate the value of $I_{t_0}(q)$.

5.2. Mesh Illumination Estimates

At the first step of space-time density estimation for mesh element q and frame t_0 the reference illumination $I_{t_0}(q)$ must be found because it is required to compute σ_I and w_I in Equation 2. Since for each frame only a relatively small number of photons is traced (refer to Equation 3) the estimate of $I_{t_0}(q)$ based only on photons for frame t_0 would be too noisy. To increase the robustness of $I_{t_0}(q)$ estimate the number of photons hitting q is averaged for five frames $[t_0, \ldots, t_4]$ (five frames were chosen experimentally and this is usually a good trade-off between excessive noise and blur). For the same reason illumination $I_{t_i}(p)$ for any mesh element $p \in \Omega(q)$ is estimated based on photons collected for frames $[t_{i-2}, \ldots, t_{i+2}]$ which results in the following formulation:

$$I_{t_i}(p) = \frac{\sum\limits_{-2 \le j \le 2} n_{t_{i+j}}(p)}{5 \cdot p_{\text{Area}}} \cdot e \tag{5}$$

where $n_{t_i}(p)$ denotes the number of photons for the mesh element p and frame t_i, p_{Area} is the surface area of p, and e is the radiant flux carried by each photon (refer to [11] for details how e can be computed). For the $I_{t_i}(p)$ estimate at the beginning or the end of considered animation segment (frames t_0, t_1, $t_{\max-2}$, and $t_{\max-1}$) the five nearest frames in time domain that are available are chosen.

In our estimates $I_{t_0}(q)$ and $I_{t_i}(p)$ we average the number of photons in time domain because for most scene regions local illumination changes slowly and our emphasis is on preserving spatial details in lighting distribution. Figure 3 summarizes our scheme for selecting frames to derive $I_{t_0}(q)$ and $I_{t_i}(p)$ estimates.

5.3. Filtering

When the values of σ_T, σ_S, σ_I, and $I_{t_0}(q)$ are computed for a given mesh element q bilateral filtering as specified in Equation 2 is performed. At first $I_{t_i}(q)$ for the mesh element q and for frames $[t_0, \ldots, t_{\max-1}]$ are processed. In the same way $I_{t_i}(p)$ for all mesh elements p (located within the

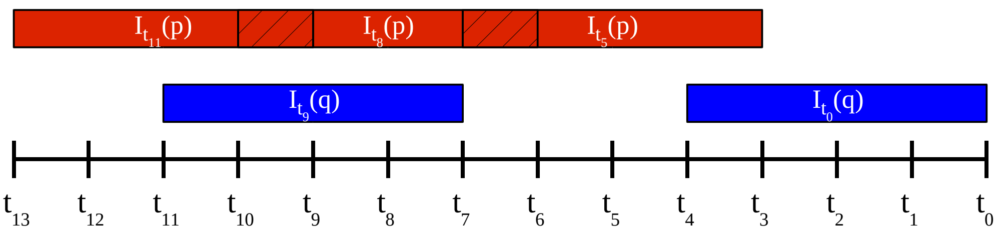

Figure 3: Illumination estimates $I_{t_i}(p)$ for mesh elements p and frames t_i that are used in Equation 2 are computed by averaging the number of photons hitting each p for five neighboring frames in time domain.

ring of radius r_{ring} centered at q) are processed starting from the smallest ring neighborhoods.

Note that the choice of σ_T and σ_S, is valid for other space-time filtering techniques, for example simple Gaussian filtering. In the latter case σ_I must not be estimated and the weights w_I are simply ignored in Equation 2. All the remaining steps of animation rendering remain the same as for bilateral filtering.

6. Results

In this section we present the results that we obtained using our techniques for four scenes of various complexity (refer to Table 1 for details).

Scene	Number of Triangles	Number of Lights
GLASS	122,244	1
BOX	44,444	4
SALON	159,252	39
SALA	27,507	12

Table 1: Test scenes complexity.

The goal of the first experiment was to evaluate the quality of reconstructed caustic for the GLASS scene shown in Figure 4d. Only the static case was considered, i.e., time domain was ignored. Figure 4a shows an unfiltered caustic that results from simple bucketing of photon hit points for each mesh element. Figure 4b shows the result of Gaussian filtering with adaptive σ_S for $n = 400$ i.e. $\sigma_{\%} = 5\%$ (refer to Section 5.1.2). The result of bilateral filtering shown in Figure 4c was obtained for the same setting of σ_S. The value $l = 1$ was set for the parameter that decides upon the influence of illumination samples from neighboring triangles to the illumination value at the filtered triangle (refer to Section 5.1.3). As it can be seen bilateral filtering produces the best results. The noise is significantly reduced in respect to the non-filtered image in Figure 4a and excessive blur visible in Figure 4b is avoided. All images shown in Figure 4 were obtained using 557,000 photons.

In the second experiment the performance of Gaussian and bilateral filtering is evaluated for an animated sequence. The number of temporally processed frames $t_{\max} = 30$ were considered and $n = 400$ and $l = 0.5$ were set. Figures 5a and c show the result of Gaussian filtering that leads to temporal blur in reconstructed lighting. Such excessive blurring (perhaps at expense of some visible noise) can be avoided when bilateral filtering is applied (refer to Figures 5b and d). For all frames in Figure 5 454,400 photons per frame were considered.

Figures 6b and 7b show sample animation frames obtained using bilateral filtering. The number of temporally processed frames $t_{\max} = 30$ were considered and $n = 400$ and $l = 0.5$ were set for both scenes. Those scenes contain many very small mesh elements that collect a very small number of photons and therefore exhibit significant noise in the lighting function. This leads to strong fluctuations of illumination between neighboring triangles, which bilateral filtering interprets as important lighting details and tends to preserve. To overcome this problem we adaptively increase σ_I for those mesh elements that collect a very small number of photons. For Figures 6b and 7b such an increase of σ_I is performed when the number of collected photons is below 16 per mesh element. In practice, this means that for such elements filtering characteristics adaptively evolves towards Gaussian filtering when the number of photons decreases. Figures 6a and 7a show the corresponding results for the same number photons without any filtering. In order to obtain a similar quality of frames as in Figures 6b and 7b for traditional frame-by-frame computation and filtering performed only in the spatial domain at least ten times more photons must be traced. However, tem-

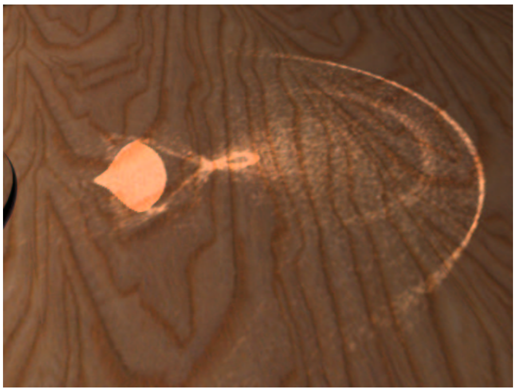

(a) No filtering

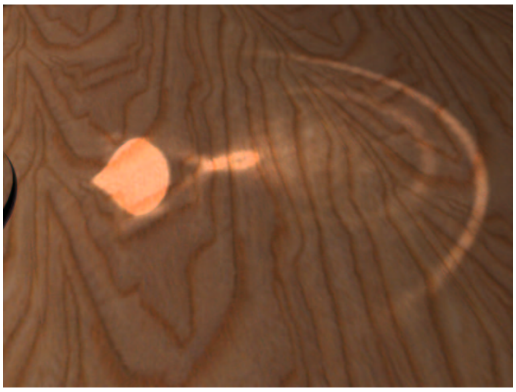

(b) Gaussian filtering

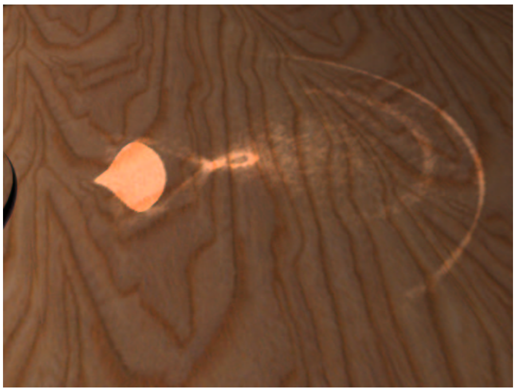

(c) Bilateral filtering

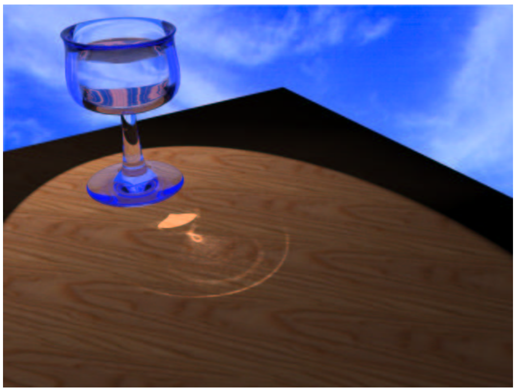

(d) Bilateral filtering

Figure 4: Caustic reconstruction for the GLASS scene. Notice that the shape of caustics is affected by the polygonal structure of the glass.

poral flickering is very difficult to eliminate when the processing of photons is limited only to spatial domain.

Scene	GLASS	BOX	SALON	SALA
Ph. Trac.	10.82	8.19	19.87	5.23
Bilateral	3.57	10.60	29.00	6.74
Rendering	11.34	14.35	68.00	43.26
Total	25.73	33.14	116.87	55.23

Table 2: Timings for various stages of the frame computation. Rendering is performed using ray tracing and the indirect lighting (computed and filtered in the photon tracing stage) is interpolated based on the values stored in mesh vertices. All timings are given in seconds and are measured on a Pentium 4, 3.2 GHz processor.

Table 2 summarizes timings obtained for all four tested scenes. As can be seen the overhead incurred by spatio-temporal bilateral filtering is quite significant and directly depends on the number of polygons. In our current implementation the filtering is performed for all mesh elements in the scene. This could be improved by lazily performed filtering only for those mesh elements that are directly visible.

7. Conclusions

In this paper we investigated the use of bilateral filtering for lighting reconstruction in the framework of histogram-based density estimation method. We found that bilateral filtering has many desirable properties in this task such as efficient removal of stochastic noise while preserving sharp features in the lighting function. Through applying bilateral filtering in spatio-temporal domain a smaller number of photons per frame can be traced and temporal flickering is substantially reduced.

We applied bilateral filtering in the framework

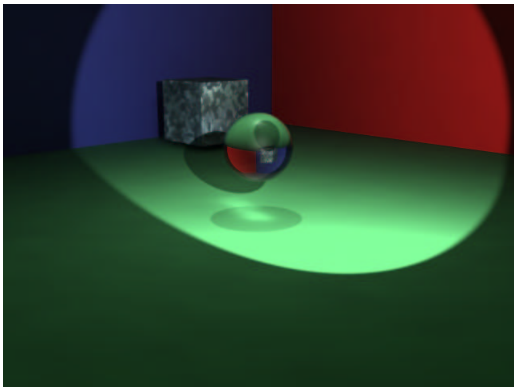

(a) Gaussian filtering

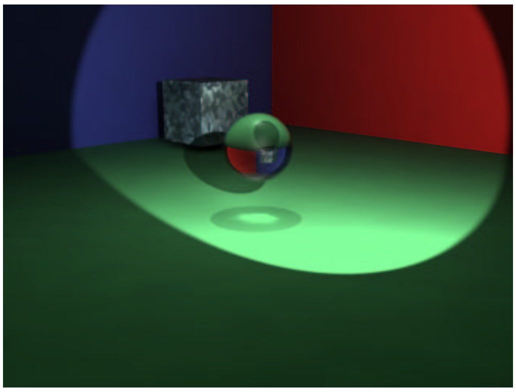

(b) Bilateral filtering

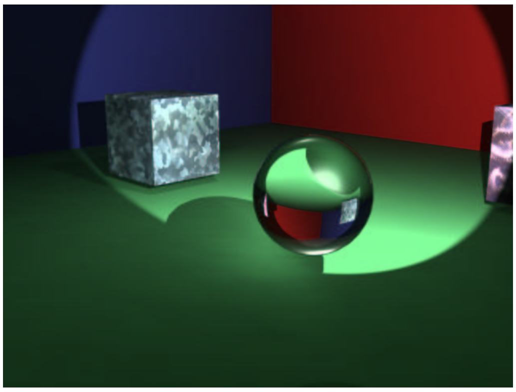

(c) Gaussian filtering

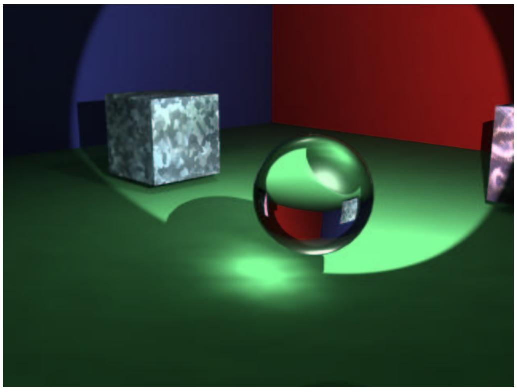

(d) Bilateral filtering

Figure 5: Space-time lighting reconstruction for the BOX scene. Notice that the caustic resulting from the flying glass sphere is blurred along the sphere motion trajectory in a) or is almost completely washed out in c) for Gaussian filtering. Those undesirable effects are significantly suppressed in b) and d) when bilateral filtering is used.

of Insight plug-in for the 3D Studio Max system. As a result an efficient solution for high-quality animation rendering was obtained. Although, we perform lighting reconstruction in spatio-temporal domain, we avoided costly photon storage as it was required in the algorithm proposed by Myszkowski et al. [14].

The main limitation of our current approach is the lack of adaptivity in the mesh structure used to collect photons. The mesh granularity limits the spatial resolution of reconstructed lighting details. As future work we plan to remove this limitation. Another interesting research topic is to consider differences in the perceived brightness (perceptual quantity) instead of currently considered illumination (photometric quantity) while computing the influence of local lighting in bilateral filtering. This could result in more precise control over the preservation of perceivable details in reconstructed lighting.

References

[1] M. Cammarano and H.W. Jensen. Time Dependent Photon Mapping. In *Proceedings of the 13th Eurographics Workshop on Rendering*, pages 135–144, 2002.

[2] P. Choudhury and J. Tumblin. The trilateral filter for high contrast images and meshes. In *14th Eurographics Symposium on Rendering*, pages 186–196, 2003.

[3] C. Damez, K. Dmitriev, and K. Myszkowski. State of the art in global illumination for interactive applications and high-quality anima-

(a) No filtering

(b) Bilateral filtering

Figure 6: Space-time lighting reconstruction for the SALA scene using 109,400 photons per frame.

(a) No filtering

(b) Bilateral filtering

Figure 7: Space-time lighting reconstruction for the SALON scene using 622,200 photons per frame.

tions. *Computer Graphics Forum*, 22(1):55–77, March 2003.

[4] C. Damez and F.X. Sillion. Space-Time Hierarchical Radiosity for High-Quality Animations. In *Eurographics Rendering Workshop 1999*, pages 235–246, 1999.

[5] K. Dmitriev, S. Brabec, K. Myszkowski, and H.-P. Seidel. Interactive Global Illumination Using Selective Photon Tracing. In *Proceedings of the 13th Eurographics Workshop on Rendering*, pages 25–36, 2002.

[6] G. Drettakis and F.X. Sillion. Interactive Update of Global Illumination Using a Line-Space Hierarchy. In *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 57–64, 1997.

[7] F. Durand and J. Dorsey. Fast bilateral filtering for the display of high-dynamic-range images. *ACM Transactions on Graphics*, 21(3):257–266, 2002.

[8] S. Fleishman, I. Drori, and D. Cohen-Or. Bilateral mesh denoising. *ACM Transactions on Graphics*, 22(3):950–953, 2003.

[9] V. Havran, C. Damez, K. Myszkowski, and H-P. Seidel. An efficient spatio-temporal architecture for animation rendering. In *14th Eurographics Symposium on Rendering*, pages 106–117, 2003.

[10] P. Heckbert. Adaptive Radiosity Textures for Bidirectional Ray Tracing. In *Computer Graphics (ACM SIGGRAPH '90 Proceedings)*, pages 145–154, 1990.

[11] H.W. Jensen. *Realistic Image Synthesis Using Photon Mapping*. AK, Peters, 2001.

[12] T.R. Jones, F. Durand, and M. Desbrun. Non-iterative, feature-preserving mesh smoothing. *ACM Transactions on Graphics*, 22(3):943–949, 2003.

[13] K. Myszkowski. Lighting reconstruction using fast and adaptive density estimation techniques. In *Eurographics Rendering Workshop 1997*, pages 251–262, 1997.

[14] K. Myszkowski, T. Tawara, H. Akamine, and H-P. Seidel. Perception-Guided Global Illumination Solution for Animation Rendering. In *Proceedings of ACM SIGGRAPH 2001*, pages 221–230, 2001.

[15] B.W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapmann and Hall, London, 1985.

[16] P. Tole, F. Pellacini, B. Walter, and D.P. Greenberg. Interactive Global Illumination in Dynamic Scenes. *ACM Transactions on Graphics*, 21(3):537–546, 2002.

[17] C. Tomasi and Roberto Manduchi. Bilateral Filtering for Gray and Colored Images. In *Proceedings IEEE International Conference on Computer Visiion*, pages 836–846, 1998.

[18] V. Volevich, K. Myszkowski, A. Khodulev, and Kopylov E.A. Using the Visible Differences Predictor to Improve Performance of Progressive Global Illumination Computations. *ACM Transactions on Graphics*, 19(2):122–161, 2000.

[19] B. Walter, G. Drettakis, and S. Parker. Interactive Rendering using the Render Cache. In *Proceedings of the 10th Eurographics Workshop on Rendering*, pages 235–246, 1999.

[20] B.J. Walter. *Density Estimation Techniques for Global Illumination*. Ph.D. thesis, Cornell University, 1998.