

Photon Differentials in Space and Time

L. Schjøth¹, J.R. Frisvad², K. Erleben¹, and J. Sporring¹

¹ eScience Center, Department of Computer Science, University of Copenhagen
Copenhagen, Denmark

² Department of Informatics and Mathematical Modelling, Technical University of Denmark
Copenhagen, Denmark

{schj,kenny,sporring}@diku.dk, jrff@imm.dtu.dk

Abstract. We present a novel photon mapping algorithm for animations. We extend our previous work on photon differentials [12] with time differentials. The result is a first order model of photon cones in space and time that effectively reduces the number of required photons per frame as well as efficiently reduces temporal aliasing without any need for in-between-frame photon maps.

1 Introduction

Rendering animated scenes with global illumination methods produces some interesting problems, where the perhaps most prominent problem is aliasing caused by the fact that the temporal domain is discretized at a often very low resolution.

Feature films are usually shot at a frame rate of around 25 frames per second (fps). Despite the fact that the human eye easily perceives more than 25 fps, this frame rate still gives a seemingly fluid motion when a film is produced with a traditional movie camera. However, were we to produce a feature film of an animated virtual scene using an unmodified global illumination method at a frame rate of 25 fps, fluid motion is not guaranteed. A typical unmodified global illumination method produces images at instant time in the temporal domain. This can induce temporal aliasing, which is seen as an adverse stroboscopic effect, where the illumination changes rapidly over time. A traditional-camera produced feature film will avoid this problem because the camera has non-zero exposure time. This means that camera-perceived illumination is averaged over the exposure time. In effect, high frequency motion is blurred and therefore seems fluid. This effect is often called motion blur. A temporal aliasing defect not solved by this blurring is the wagon-wheel effect, which is seen as a backwards turning of the spokes of a wheel due to temporal under sampling. However, as our method does not address this particular problem, it will not be discussed further here.

A survey paper for global illumination and rendering solutions exploiting temporal coherence is given in [13].

Brute force methods, such as the accumulation buffer [7], average together in-between frames in order to achieve motion blur. These methods can achieve arbitrary high accuracy but are often prohibitively expensive as full renderings typically have to be made of a large number of in-between frames.

Different global illumination methods also address temporal aliasing by simulating motion blur. Distribution ray tracing [3] achieves motion blur by stochastically sampling the temporal domain as well as the spatial. In [9] the authors adaptively controlled




Fig. 1. A frame from an animation of caustics on the bottom of the sea. Image (a) shows the results using Time dependent photon mapping while image (b) shows our method, Temporal photon differentials.

the temporal and spatial sampling resolution by examining local variations of indirect illumination over time and space in a pilot estimate. In [5] motion blur is modeled as a wedge in the frequency domain. Ray tracing methods using ray differentials have been studied in [8,24,6].

With *time dependent photon mapping* [1], photon mapping was extended such that indirect illumination was estimated using a four dimensional photon map that expanded into the temporal domain. In this manner photons were filtered not only based on their spatial position but also their temporal. A non-linear approach was proposed in [14].

A problem common to the discussed methods is that they all rely on information from in-between frames in order to achieve motion blur. Having this information available places certain restrictions on the animated scene; because a scene description is needed at arbitrary time steps, movement of scene elements either needs to be described as an analytic function, or movement has to be interpolated between frames. The analytical approach is by far the best but demands a certain complexity of the animation application, as well as a tight link to the rendering software. The interpolative approach is error prone, since the animation curve might not be linear. Furthermore, some accelerator for ray-object intersection (such as a bsp-tree) typically needs to be either rebuilt, or at least updated an extra number of times equal to the number in-between frames needed.

In this paper, we propose the use of *temporal photon differentials*. In Figure 1 is shown a frame from a movie of the sea bottom, where the motion of waves causes the caustics to move in the direction of the arrow. As illustrated, our method produces motion blur, and which neither needs in-between frames, nor to over-smooth indirect illumination with high temporal frequency. Our proposed method is an extension of photon differentials [10]. It takes advantage of ray differentials [8] and their extension into the temporal domain [12], and since velocity is a relative concept, our methods handles camera as well as object motion. An earlier version of this paper has appeared in [11].

2 Temporal Photon Differentials

In our method each photon represents a beam of light that expands, contracts and re-shapes in space and time as it propagates through the scene. We keep track of a photon's coherence by deriving the first order structure of its direction and position with respect to both time and space as it traverse the scene.

Representing a photon as a parameterized ray with origin in \mathbf{x} and the direction $\boldsymbol{\omega}$, we describe the derivatives of a photon with two Jacobian matrices; one for the positional derivatives and one for the directional derivatives. The positional derivatives are then given by

$$D\mathbf{x} = \begin{bmatrix} \frac{\partial x}{\partial u} & \frac{\partial x}{\partial v} & \frac{\partial x}{\partial t} \\ \frac{\partial y}{\partial u} & \frac{\partial y}{\partial v} & \frac{\partial y}{\partial t} \\ \frac{\partial z}{\partial u} & \frac{\partial z}{\partial v} & \frac{\partial z}{\partial t} \end{bmatrix} = [D_u\mathbf{x} \ D_v\mathbf{x} \ D_t\mathbf{x}], \quad (1)$$

where $D\mathbf{x}$ is the Jacobian of the positional derivatives, and $D_u\mathbf{x}$, $D_v\mathbf{x}$ and $D_t\mathbf{x}$ are column vectors that describe the positional derivatives with respect to the scalar variables u , v , and t . The spatial variables u and v relate to light sources, from which the photon differential was emitted.

We write the directional derivatives of the photon as

$$D\boldsymbol{\omega} = [D_u\boldsymbol{\omega} \ D_v\boldsymbol{\omega} \ D_t\boldsymbol{\omega}]. \quad (2)$$

Exactly as with photon differentials, $D_u\boldsymbol{\omega}$, $D_v\boldsymbol{\omega}$, $D_u\mathbf{x}$ and $D_v\mathbf{x}$ are spatially dependent differential vectors. The directional and positional differential vectors with respect to time are new to photon differentials; they are denoted $D_t\boldsymbol{\omega}$ and $D_t\mathbf{x}$ in the above equations. For a scene, in which the light sources are static, the time dependent differential vectors will, initially, be zero. If the scene, additionally, is completely static, they will remain zero-vectors through out the photon's traversal of the scene. In this specific case, our method will behave exactly as ordinary photon differentials: the photons will expand and contract depending on the reflections and refraction encountered during tracing, and their spatially dependent positional differential vectors will form a footprint, which is used in the reconstruction of the indirect illumination.

On the other hand, if we have a dynamic scene, then photon differentials interacting with a non-static scene element will attain non-zero time dependent differential vectors. In this case, the derivatives of a dynamic scene-element's surface positions or normals with respect to time will be non-zero:

$$D_t\mathbf{n} \neq \mathbf{0}, \quad (3)$$

or

$$D_t\mathbf{q} \neq \mathbf{0}, \quad (4)$$

where \mathbf{n} is a surface normal to the element and \mathbf{q} is a position on the element's surface. This again will affect the time dependent derivatives of a photon interacting with the scene element.




Fig. 2. Temporal filter kernel shaped by a spatial kernels translation along the time dependent differential vector

Sporring *et al.* [12] evaluates the full differentials for a parameterized ray. This allows for an extension of parameters such that the derivatives of a ray can be considered with respect to time. From the equations for transfer, reflection and refraction, we observe that non-zero time-dependent element differentials (eg. $D_t \mathbf{q}$) propagate through these interactions to the differentials of the interacting photon. We exploit this behavior such that a footprint from a photon differential traveling in a dynamic scene not only describes the spatial coherence of the ray but also the temporal coherence of the ray.

When a photon differential hits a surface, its positional differential vectors are projected onto the surface's tangent plane at the intersection point. The spatial footprint of the photon differential is the area on the tangent plane of a parallelogram spanned by the positional differential vectors. The spatial footprint can be used to shape an anisotropic filter kernel as illustrated in Figure 3.

The time dependent positional differential vector, $D_t \mathbf{x}$, tells us either how the photon's footprint is going to behave over consecutive frames, or how the footprint has behaved in former frames. In the former case, the direction of $D_t \mathbf{x}$ predicts the direction on the surface that the footprint will move, and the magnitude of the vector predicts how far the footprint is likely to move. Basically, the magnitude and the direction of $D_t \mathbf{x}$ depends on the estimation method used to calculate the time derivatives of an element, which again depends on the geometry representation. In the present method, we simply use finite differences and triangle meshes. Except for the last frame, in which we use backward differences, we estimate the time dependent differentials using forward differences. When we want to predict how a footprint is going to behave, having intersected a moving element, we estimate the element's positional time derivatives by

$$D_t \mathbf{q}_f = e(\mathbf{q}_{f+1} - \mathbf{q}_f), \quad (5)$$

where $D_t \mathbf{q}_f$ is the derivative of the vertex \mathbf{q}_f with respect to time at frame step f , and e is the exposure. The exposure is a parameter for how much we trust our prediction.




Fig. 3. Spatial filter kernel shaped by the positional differential vectors, $D_u \mathbf{x}$ and $D_v \mathbf{x}$




Fig. 4. Temporal filter kernel shaped by a spatial kernels translation along the time dependent differential vector

Generally, it works as a smoothing parameter for the time dependent footprint that decides how much motion blur we induce. Its unit is given in frames as it depends on the movement of the scene elements between frames. The exposure is related to the exposure time by the frame rate such that the exposure time is equal to the exposure divided by the frame rate.

The time dependent footprint constitutes an integration of the spatial footprint over the time dependent differential vector such that the spatial footprint is elongated along the vector. We achieve this by translating the spatial footprint along the time dependent differential vector. As in the spatial case, the time dependent footprint describes a filter kernel. In Figure 4(a), $D_t \mathbf{x}_{pd}$ is the time dependent differential vector, \mathbf{x}_{pd} is the center of the spatial kernel, and \mathbf{x} is the estimation point, for which the kernel weight is estimated.

The kernel is translated along $D_t \mathbf{x}_{pd}$ to the point, \mathbf{x}'_{pd} , on the line segment, $(\mathbf{x}_{pd} \rightarrow \mathbf{x}_{pd} + D_t \mathbf{x}_{pd})$, where \mathbf{x}'_{pd} is the point on the segment having the shortest distance to the estimation point, \mathbf{x} . Using \mathbf{x}'_{pd} as center for the spatial kernel, the resulting time dependent kernel will achieve an elongated shape as illustrated in Figure 4(b).

The irradiance of the time dependent photon differential is estimated as

$$E_{pd} = \Phi_{pd} / A_{pd}, \tag{6}$$




Fig. 5. First order approximation of the sweeping area of the kernel on the surface of an object induced by relative motion of the light source and the object. The approximate area swept is the sum of the area of the kernel and the rectangle spanned by $D_t \mathbf{x}$ and l , where \mathbf{x}_t is the initial central point of projection, $D_t \mathbf{x}$ is the vector of change of \mathbf{x}_t by time, and l is the spread of the kernel in the direction perpendicular to $D_t \mathbf{x}$

where Φ_{pd} is the radiant flux carried by the photon, and A_{pd} is the surface area, to which the radiant flux is incident. For the time dependent photon differential, this area is the area of the time dependent kernel. Referring to Figure 5 this area is calculated as

$$A_{pd} = \frac{1}{4}\pi|D_u \mathbf{x} \times D_v \mathbf{x}| + l|D_t \mathbf{x}|, \quad (7)$$

where the first term is the area of the spatial kernel and the second term is the area of a rectangle. One side of the rectangle is the length of the time dependent differential vector and the other is the length of the spatial kernel in a direction perpendicular to the time dependent differential vector.

Having defined the time dependent kernel as well as the irradiance of the photon differential, we can now formulate a radiance estimate for temporal photon differentials.

2.1 The Temporal Radiance Estimate

Reflected radiance from temporal photon differentials can be estimated by

$$\hat{L}_r(\mathbf{x}, \boldsymbol{\omega}) = \sum_{pd=1}^n f_r(\mathbf{x}, \boldsymbol{\omega}_{pd}, \boldsymbol{\omega}) E_{pd}(\mathbf{x}, \boldsymbol{\omega}_{pd}) K((\mathbf{x} - \mathbf{x}'_{pd})^T \mathbf{M}_{pd}^T \mathbf{M}_{pd} (\mathbf{x} - \mathbf{x}'_{pd})), \quad (8)$$

where \mathbf{x} is a position on an illuminated surface, $\boldsymbol{\omega}$ is reflection direction considered, f_r is the bi-directional reflectance function, $\boldsymbol{\omega}_{pd}$ is the incident ray direction, \mathbf{x}'_{pd} is the translated center of spatial kernel, E_{pd} is the irradiance of the temporal photon differential, K is a scalar kernel functions examples of which are shown in Table 1 and \mathbf{M}_{pd} is a matrix that transforms from world coordinates to the filter space of the spatial kernel, as illustrated in Figure 6.

The temporal radiance estimate can be extended as to include filtering in time. One intuitive approach is to weight the part of the differential, which is closest in time the highest, where the time is estimated from the photon hit point. This can be achieved by again using a simple scalar kernel as those presented in Table 1. To the kernel, we

Table 1. Some rotational symmetric kernel functions. The normalization constants k_* are found as the solution to $1 = \int_{|\mathbf{y}|^2 < 1} K(|\mathbf{y}|^2) d\mathbf{y}$ and depend on the dimensionality of \mathbf{y} .

| Kernel | $K(\mathbf{y} ^2)$ | |
|--------------|--|---------------------------|
| Uniform | $k_{\text{uni}} 1$ | if $ \mathbf{y} ^2 < 1$, |
| | 0 | otherwise |
| Cone | $k_{\text{cone}}(1 - \sqrt{ \mathbf{y} ^2})$ | if $ \mathbf{y} ^2 < 1$, |
| | 0 | otherwise |
| Epanechnikov | $k_{\text{epan}}(1 - \mathbf{y} ^2)$ | if $ \mathbf{y} ^2 < 1$, |
| | 0 | otherwise |
| Quartic | $k_{\text{quart}}(1 - \mathbf{y} ^2)^2$ | if $ \mathbf{y} ^2 < 1$, |
| | 0 | otherwise |
| Gaussian | $k_{\text{gauss}} \exp(-\frac{1}{2} \mathbf{y} ^2)$ | |




Fig. 6. Transformation from geometry space to filter space by the matrix M_{pd} . The ellipse inside the parallelogram is the footprint of the photon differential. When transformed into filter space the ellipse becomes a unit circle.

input the squared distance along the time dependent differential vector, $D_t \mathbf{x}_{pd}$, relative to furthest point of the kernel along negative $D_t \mathbf{x}_{pd}$. This is illustrated in Figure 5.

With time filtering the temporal radiance estimate is formulated as

$$\hat{L}_r(\mathbf{x}, \omega) = \sum_{pd=1}^n f_r(\mathbf{x}, \omega_{pd}, \omega) E_{pd}(\mathbf{x}, \omega_{pd}) K \left(\frac{(\mathbf{x}_t - \mathbf{x}'_{pd})^T (\mathbf{x}_t - \mathbf{x}'_{pd})}{h_t^2} \right), \quad (9)$$

where K is a kernel function (See Table 1), h_t is the length of the temporal kernel along $D_t \mathbf{x}_{pd}$, and \mathbf{x}_t is the furthest point of the kernel in the direction $-D_t \mathbf{x}_{pd}$. With the formulation of the temporal radiance estimate, we now have a method, which reconstructs indirect illumination based on a virtual scenes dynamics. This allows for motion blur. In the following we will make a simple analysis of the method.

3 Results

We first test our proposed method using a case study. The case study is a simple animated scene, in which a sinusoidal wave moves horizontally in a direction perpendicular to the wave crests. The wave is illuminated from above by collimated light, which it refracts such that the light form caustics on a plane beneath the wave. A virtual camera is placed such that the caustics are clearly visible.

We have rendered the scene using temporal photon differentials, and Cammarano and Jensen's time dependent photon mapping. The images in Figure 7 are renderings of the same frame but at different exposures, e . They were rendered using temporal photon differentials and a photon map containing only 1000 photons. From the images we notice that the temporal photon differentials assume the expected behavior. As the exposure increases the caustics are blurred acquiring a comets tail away from the direction of movement. This is the behavior chosen at implementation time. We could just as well have placed the time dependent kernel centered over the photon intersection point and likewise have centered the time filtering or we could just have centered the filtering. As it is, the time differential is 'trailing' after the photon both in respect to placement and filtering. As we shall see, though it is hardly visible, the same strategy has been implemented for time dependent photon mapping.

The renderings in Figure 8 have all except (a) been created with time dependent photon mapping. Additionally, all images were rendered using the same exposure, $e = 1.0$ frame. Figure 8(a) has been included for comparison, it was rendered using temporal photon differentials and is a copy of the image in Figure 7 with an exposure 1.0 frame. First of all, what we see from image 8(b) is that the bias versus variance trade-off provided by time dependent photon mapping is too poor to produce palpable caustics. For this reason a much higher number photons have been used to render the images in the three rightmost columns. Of these, the top row is based on a photon map containing as much as 480 000 photons while the bottom row is based on a photon map contain 40 000 photons. From left to right the temporal resolution increases from 0 to 2 to 10 in-between frames. The spatial bandwidth for the renderings was chosen as to decrease noise to an acceptable level. This leads to the perhaps most important observation, namely that a low temporal resolution produces visible bands that can only be removed by filtering beyond what removes normal noise. This complicates matter, as an increase of photons no longer is a guarantee for high quality illumination.

Photon differentials are free of this concern as the blurring is based on the first order derivatives of object movement and not on finite animation steps. In the implementation presented here one additional frame is need in order to estimate the derivatives.

The images in Figure 9 present a more complex—all though very artificial—scene, in which a cylinder is rotating counterclockwise around its one end. This rotation means that the speed of the cylinder will increase as we move from the turning point down the length of the cylinder. As a result the produced caustic becomes more blurred when refracted from the high-speed end of the cylinder. Figure 9(a) gives the solution provided by temporal photon differentials while the images in Figure 9(b) and 9(c) were produced with time dependent photon mapping. All images were rendered with the same number of photons contained in the photon map. However, the two latter images were rendered with different bandwidths. From these two images we see that at this obviously




Fig. 7. Renderings of the case study scene using temporal photon differentials. The number under the renderings indicate exposure, e . All images are rendered at the same frame step using a photon map containing 1 000 photons.







Fig. 8. Comparing temporal photon differentials with time dependent photon differentials. Rendering in (a) shows temporal photon differentials using 1 000 photons, (b) shows time dependent photon differentials using 1 000 photons, (c)-(e) shows time dependent photon differentials using 480 000 photons with 0, 2, and 10 in-between frames respectively, and (f)-(g) shows time dependent photon differentials using 40 000 photons with 0, 2, and 10 in-between frames respectively. The exposure for all images is 1.0 frame.



(a)



(b)



(c)

Fig. 9. Rendering of a rotating cylinder. All images were rendered with an exposure of 1.0 frames, using a photon map containing 5000 photons. Figure 9(a) was rendered using temporal photon differentials. Figure 9(b) and 9(c) were rendered using time dependent photon mapping the former using 100 photons per radiance estimate and the latter using 250 photons in the radiance estimate.




Fig. 10. Projected time differentials from a rotating glass cylinder

low temporal resolution an increase in bandwidth can help remove the temporal bands that time dependent photon mapping is prone to. The price, however, is an unwanted blurring of the front of the caustic.

Finally, Figure 10 solely depicts the photons' time differentials as they are projected down on the plane beneath the cylinder. A high exposure time has been used as to facilitate the illustration. The image confirms that the time differential vectors become longer when refracted from the high-speed end of the cylinder, thus elongating the time dependent kernel used in the temporal radiance estimate.

4 Conclusions

In this paper we have presented *temporal photon differentials* – a global illumination method for rendering animated scenes. *Temporal photon differentials* elegantly handles time filtering such that frames can be rendered on a one to one basis. Since velocity is a relative concept, our method handles camera as well as object motion, further in contrast to similar dynamic scene renderers, *temporal photon differentials* does not need in-between frames in order to avoid temporal aliasing. Finally, *temporal photon differentials* includes photon differential's ability to efficiently and sharply model caustics in still frames, and it further improves the computational efficiency for animations by reusing photons from earlier frames through the first order spatial-temporal model of photon incidence.

References

1. Cammarano, M., Jensen, H.W.: Time dependent photon mapping. In: EGWR 2002: Proceedings of the 13th Eurographics Workshop on Rendering, pp. 135–144. Eurographics Association, Aire-la-Ville (2002)
2. Christensen, P.H., Laur, D.M., Fong, J., Wooten, W.L., Batali, D.: Ray differentials and multiresolution geometry caching for distribution ray tracing in complex scenes. *Computer Graphics Forum (Proceedings of Eurographics 2003)* 22(3), 543–552 (2003)

3. Cook, R.L., Porter, T., Carpenter, L.: Distributed ray tracing. *Computer Graphics (Proceedings of ACM SIGGRAPH 1984)* 18(3), 137–145 (1984)
4. Durand, F., Holzschuch, N., Soler, C., Chan, E., Sillion, F.X.: A frequency analysis of light transport. In: Marks, J. (ed.) *ACM Transaction on Graphics (Proceedings of ACM SIGGRAPH 2005)*, vol. 24(3), pp. 1115–1126 (2005)
5. Egan, K., Tseng, Y.T., Holzschuch, N., Durand, F., Ramamoorthi, R.: Frequency Analysis and Sheared Reconstruction for Rendering Motion Blur. *ACM Transaction on Graphics (Proceedings of ACM SIGGRAPH 2009)* 28(3), 93:1–93:13 (2009)
6. Gjøøl, M., Larsen, B.D., Christensen, N.J.: Final gathering using ray differentials. In: *Proceedings of WSCG 2008*, pp. 41–46 (2008)
7. Haeberli, P., Akeley, K.: The accumulation buffer: hardware support for high-quality rendering. *Computer Graphics (Proceedings of ACM SIGGRAPH 1990)* 24(4), 309–318 (1990)
8. Igehy, H.: Tracing ray differential. In: Rockwood, A. (ed.) *Proceedings of ACM Siggraph 1999*, pp. 179–186. Addison Wesley Longman, Los Angeles (1999)
9. Myszkowski, K., Tawara, T., Akamine, H., Seidel, H.P.: Perception-guided global illumination solution for animation rendering. In: *Proceedings of ACM SIGGRAPH 2001*, pp. 221–230. ACM, New York (2001)
10. Schjøth, L., Frisvad, J.R., Erleben, K., Sporring, J.: Photon differentials. In: *GRAPHITE 2007: Proceedings of the 5th International Conference on Computer Graphics and Interactive Techniques in Australia and Southeast Asia*, pp. 179–186. ACM, New York (2007)
11. Schjøth, L., Frisvad, J.R., Erleben, K., Sporring, J.: Temporal photon differentials. In: Richard, P., Braz, J., Hilton, A. (eds.) *Proceedings of GRAPP 2010 International Conference on Computer Graphics Theory and Applications (2010)* ISBN: 978-989-674-026-9
12. Sporring, J., Schjøth, L., Erleben, K.: Spatial and Temporal Ray Differentials. Tech. Rep. 09/04, Institute of Computer Science, University of Copenhagen, Copenhagen, Denmark, (September 2009) ISSN: 0107-8283
13. Tawara, T., Myszkowski, K., Dmitriev, K., Havran, V., Damez, C., Seidel, H.P.: Exploiting temporal coherence in global illumination. In: *SCCG 2004: Proceedings of the 20th Spring Conference on Computer Graphics*, pp. 23–33. ACM, New York (2004)
14. Weber, M., Milch, M., Myszkowski, K., Dmitriev, K., Rokita, P., Seidel, H.-P.: Spatio-temporal photon density estimation using bilateral filtering. In: Cohen-Or, D., Jain, L., Magnenat-Thalmann, N. (eds.) *Computer Graphics International (CGI 2004)*, pp. 120–127. IEEE, Crete (2004)