

Improving Shadow Map Filtering with Statistical Analysis

J. Gumbau¹ and L. Szirmay-Kalos² and M. Sbert³ and M. Chover¹

¹Universitat Jaume I & Castellón, Spain

²BUTE & Budapest, Hungary

³Universitat de Girona & Girona, Spain

Abstract

Shadow maps are widely used in real-time applications. Shadow maps cannot be filtered linearly as regular textures, thus undersampling leads to severe aliasing. This problem has been attacked by methods that transform the depth values to allow approximate linear filtering and to approaches based on statistical analysis, which suffer from “light bleeding” artifacts. In this paper we propose a new statistical filtering method for shadow maps, which approximates the cumulative distribution function (CDF) of depths with a power function. This approximation significantly reduces “light bleeding” artifacts, maintaining performance and spatial costs. Like existing techniques, the algorithm is easy to implement on the graphics hardware and is fairly scalable.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Line and curve generation

1. Introduction

Shadow mapping is a popular and effective way of solving the shadowing problem. The depth buffer represents the occluder geometry in a discretized form. From this information, we need to reconstruct the distance from the light source and the occluder in the direction of the shaded point in order to decide whether or not occlusion happens. The shadow test can also be imagined as a step like *visibility function* $v(z_r) = \mathcal{E}(z_o - z_r)$ that is 0 if the occluder’s distance from the light source, z_o , is smaller than the receiver distance z_r , and 1 otherwise. Unfortunately, the occluder distance is known only in the centers of the texels, so this function must be reconstructed in other points. As the occluder geometry may involve high frequency variations, the occluder distance function can only be approximately reconstructed, and high frequency components may distort the reconstructed signal even at low frequencies, which leads to the well known shadow aliasing. Linear signal theory has a solution for the aliasing problem, which low-pass filters the signal to eliminate frequencies above the Nyquist limit, but shadow mapping is a non-linear operation as it contains a comparison operation represented by the *step function*.

Our method approximates the probability that the shaded point passes the depth test. This probability is obtained from

the approximation of the cumulative distribution function of depths with a power function. The mean, minimum and the maximum of the distribution are needed to fit the power function. Our approach is capable of highly reducing “light bleeding” artifacts, or even eliminating it for moderately complex scenes, with no penalty of performance or storage costs. Moreover, for very complex scenes, it can be converted to a layered approach (in the same way as layered variance shadow maps) for completely eliminating these artifacts.

2. Related work

Although shadow mapping is a widely used hardware friendly method for computing shadows in real-time scenes it suffers from aliasing, which can be reduced by two orthogonal complementary approaches: projection optimization and shadow filtering [SWP10]. The former deals with how shadow caster objects are projected over the shadow map in order to optimize texture space for important parts of the scene [WSP04] [MT04] [SDD03]. The latter deals with how the shadow map is filtered in order to reduce aliasing.

One of the first methods introduced to alleviate aliasing in shadow mapping is *percentage-closer filtering* (PCF)

[RSC87] which averages the results of depth comparisons instead of the depth values themselves. Unfortunately, this becomes expensive when large filter sizes are used. To solve this problem *Convolution shadow maps* [AMB*07] apply depth transformation and approximately express the visibility function in a product form $\varepsilon(z_o - z_r) \approx \sum_i g_i(z_r) \cdot h_i(z_o)$ and then linearly filter the $h_i(z_o)$ factors. *Exponential shadow maps* (ESM) [Sal07] approximate the shadow test using a single exponential function. Variance Shadow Maps (VSM) [DL06] store the depth and the squared depth of the shadow casters. These values are filtered as regular textures. The filtered values are used at rendering time to calculate the first two moments M_1 and M_2 of the depth values over the shadow filter region. Then, the one-tailed version of the Chebyshev's Inequality allows us to approximate an upper bound of the probability of shadowing if receiver. However, these techniques introduce "light bleeding" artifacts. *Layered variance shadow maps* (LVSM) [LM08] solves the "light bleeding" artifacts by dividing the light's depth space into multiple layers which allows for a correct filtering of the shadow map but introduces the problem of selecting the number and the optimal placement of the warps, decreasing performance and increasing storage cost as more layers are used.

3. Reconstruction of the cumulative distribution with the Power function

Our algorithm is a statistical method that allows for eliminating light bleeding artifacts. Compared to VSM, our algorithm replaces the Chebyshev's Inequality by a power function approximation that is able to approximate the depth distribution more accurately, highly reducing "light bleeding" artifacts. It can also be combined with a layered approach (just like LVSM) to completely eliminate artifacts for complex scenes. Reconstruction of CDF was also proposed in [Gru08] and for volumetric ambient occlusion [RSKU*10].

The application of probability theory techniques in shadow mapping is possible because we can make a few fundamental assumptions on the unknown visibility function:

- At $z = 0$, that is when the shaded point is at the light source, the visibility function is 1.
- At $z = \infty$, that is when the shaded point is very far from the light source, the visibility function is 0.
- The visibility function is monotonically decreasing.

From the point of view of statistics, *uncertainty* is involved in the depth buffer, since the depth values are known only in lexel centers. The goal is to guess the visibility function at arbitrary point with minimizing this inherent uncertainty. Unlike in signal processing, we are not constrained by linear operations thus by the selection of a proper estimation, the depth testing can be made more robust.

In order to propose a practically useful statistical approach, we need to consider two additional requirements.

- Unoccluded planar objects are expected to be fully lit, so the visibility function must give value 1 for the average depth value.
- As variables are obtained from the depth buffer, we have to define them in a way which allows separable filtering.

In order to compute the visibility function, we start with the probability of no occlusion $P(z_o \geq z_r)$ and bias it to avoid self occlusions. Self occlusions are eliminated for planar surfaces if average depth \bar{z}_o is associated with visibility 1, thus our proposed visibility approximation is:

$$v(z_r) = \frac{P(z_o \geq z_r)}{P(z_o \geq \bar{z}_o)}, \text{ if } z_r > \bar{z}_o \text{ and } 1 \text{ otherwise.}$$

The probabilities are computed from the cumulative probability distribution $F(z) = P(z_o < z)$ of random variable z_o :

$$v(z_r) = \frac{1 - F(z_r)}{1 - F(\bar{z}_o)}, \text{ if } z_r > \bar{z}_o \text{ and } 1 \text{ otherwise.} \quad (1)$$

Our approach is based on reconstructing the cumulative distribution from three different values that are obtained by filtering an area of the depth buffer. The larger the area of interest, the more blurred the shadows. This filtering operation is configured as a separable kernel that calculates the minimum value z_{\min} , the maximum value z_{\max} , and the mean \bar{z}_o of the depth values.

Our cumulative distribution $F(z)$ must be zero if $z \leq z_{\min}$, equal to 1 if $z \geq z_{\max}$, and non-decreasing in between. The following normalized depth parameter is introduced for notational simplicity:

$$t = \frac{z - z_{\min}}{z_{\max} - z_{\min}} = \frac{z - z_{\min}}{\Delta z},$$

where $\Delta z = z_{\max} - z_{\min}$. Using the normalized parameter, the cumulative distribution function must be zero if $t \leq 0$, and 1 if $t \geq 1$, and non-decreasing in the $[0, 1]$ interval. Taking into account the distribution of all possible shadow map values, the cumulative distribution function may be a step function at t_{\min} and t_{\max} at the two extremes, respectively. Therefore, our goal is to find a function that increases from 0 to 1 and has the flexibility to adapt to all options between the two extreme cases. Considering these, we propose to use the function t^β where β is the parameter of data fitting (see Figure 1). Thus, the cumulative distribution is

$$F(z) = t^\beta \text{ where } z(t) = t\Delta z + z_{\min}. \quad (2)$$

Let us consider the constraint on the mean:

$$\begin{aligned} \bar{z}_o &= \int_{z_{\min}}^{z_{\max}} z dF = \int_0^1 z(t) \frac{dF}{dt} dt = \\ &= \int_0^1 (t\Delta z + z_{\min}) \beta t^{\beta-1} dt = \frac{\beta \Delta z}{\beta + 1} + z_{\min} \end{aligned}$$

Solving this equation for β , we get:

$$\beta = \frac{\tilde{z}_o - z_{\min}}{z_{\max} - \tilde{z}_o}. \quad (3)$$

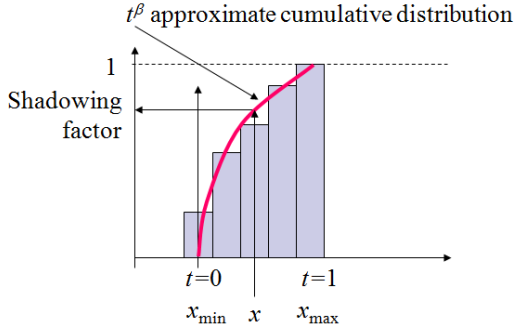


Figure 1: Approximation of the cumulative distribution function with t^β .

The visibility formula needs the computation of the cumulative distribution for the expected depth, which corresponds to

$$\tilde{t} = \frac{\tilde{z}_o - z_{\min}}{\Delta z} = \frac{\beta}{\beta + 1}.$$

The visibility function is obtained from equation 1:

$$v(t) = \frac{1 - t^\beta}{1 - \left(\frac{\beta}{\beta+1}\right)^\beta}, \text{ if } t > \frac{\beta}{\beta+1} \text{ and } 1 \text{ otherwise.} \quad (4)$$

4. Results

This section presents visual quality and performance tests to compare our new filtering approach with VSM and ESM. All quality and performance tests were generated on an Intel Core2 Quad Q9550 CPU @ 2.83Ghz with a NVIDIA GeForce 280GTX using Direct3D 10.

Figure 3 shows a visual comparison between our approach, VSM and ESM. It can be seen how our approach is able of eliminating “light bleeding” artifacts in these scenes. Our Power CDF method needs at least three channels for storing the mean, minimum and maximum values. As shown in Figure 2 VSM and ESM are very sensitive to floating point precision. However, the Power CDF reconstruction is able to render artifact-free soft shadows even with 16 bits per channel. This enables for improving performance, as shown in Table 1.

Table 1 compares performance obtained with our approaches and with the Chebyshev’s Inequality approximation. All times were measured using the Car scene at a screen resolution of 1920×1080 . We used two different shadow map precisions. It is important to note that, as it can be seen




Figure 2: Visual quality comparison using different shadow map floating point precisions.

	512x512	1024x1024	2048x2048
Chebyshev	481 fps	342 fps	142 fps
Exponential	495 fps	392 fps	165 fps
CDF recons. 32bits	420 fps	243 fps	82 fps
CDF recons. 16bits	468 fps	381 fps	160 fps

Table 1: Performance comparison table at 1920×1080 . VSM and ESM used 32 bits per channel in the shadow maps.

in Figure 3, our Power CDF method is able of performing sharper antialiased shadows compared to other methods using the same amount of texture samples and the same filtering kernel size. Figure 4 shows an analysis of the shadow sharpness of our Power CDF function compared to percentage closer filtering. This property of our method allows us to use lower resolution shadow maps while providing similar antialiasing sharpness (see Figure 4 for an example). Compared to the exponential shadow maps, our algorithm is not sensitive the light leaking artifacts introduced by the multiple distant occluders.

5. Conclusion

We have developed a new shadow filtering approach based on replacing the Chebyshev’s Inequality of the traditional VSM by evaluating a power function. This allows for enhancing the visual results of shadows by reducing “light bleeding” artifacts. Compared to ESM our technique is able of avoiding artifacts in contact shadows as shown in Figure 3. Performance is even improved by using a 16 bit floating point precision shadow map, which is sufficient for high quality results. Moreover, the Power CDF reconstruction generates sharper anti-aliased shadows, which allows the use of lower resolution shadow maps.

In conclusion, the new technique presented in this paper produces good quality anti-aliased shadows at high performance, and are very scalable for complex scenes.

References

[AMB*07] ANNEN T., MERTENS T., BEKAERT P., SEIDEL H.-P., KAUTZ J.: Convolution shadow maps. In *Rendering Tech-*




Figure 3: Notice the artifacts introduced by VSM creating an effect of overlapping shadows, quite noticeable in the spheres scene. Artifacts introduced by ESM are distinguishable on the shadows at contact points (h) and are caused by the difference of values stored in the exponential shadow map. Our technique is able of eliminating “light bleeding” artifacts introduced in both VSM and ESM.

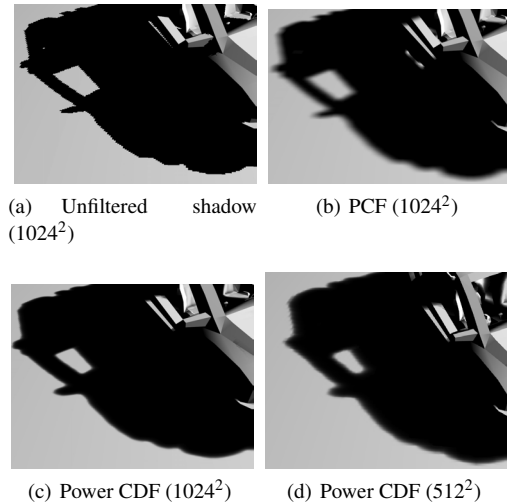


Figure 4: Visual quality comparison of shadows sharpness comparing PCF with the Power CDF with the same number of samples and filtering kernel size.

niques 2007: Eurographics Symposium on Rendering (2007), Eurographics, pp. 51–60. [2](#)

[DL06] DONNELLY W., LAURITZEN A.: Variance shadow maps. In *13D '06: Proceedings of the 2006 symposium on Interactive 3D graphics and games* (New York, NY, USA, 2006), ACM, pp. 161–165. [2](#)

[Gru08] GRUN H.: Approximate cumulative distribution function shadow mapping. In *Shader X 7* (2008), Engel W., (Ed.), Course Technology. [2](#)

[LM08] LAURITZEN A., MCCOOL M.: Layered variance shadow maps. In *GI '08: Proceedings of graphics interface 2008* (Toronto, Ont., Canada, Canada, 2008), Canadian Information Processing Society, pp. 139–146. [2](#)

[MT04] MARTIN T., TAN T.: Anti-aliasing and continuity with trapezoidal shadow maps. In *Rendering Techniques* (2004), pp. 153–160. [1](#)

[RSC87] REEVES W. T., SALESIN D. H., COOK R. L.: Rendering antialiased shadows with depth maps. *SIGGRAPH Comput. Graph.* 21, 4 (1987), 283–291. [1](#)

[RSKU*10] RUIZ M., SZIRMAY-KALOS L., UMENHOFFER T., BOADA I., FEIXAS M., SBERT M.: Volumetric ambient occlusion for volumetric models. In *Proceedings of CGI* (2010). [2](#)

[Sal07] SALVI M.: Rendering filtered shadows with exponential shadow maps. *ShaderX⁶: Advanced rendering techniques* (2007). [2](#)

[SDD03] STAMMINGER M., DRETTAKIS G., DACHSBACHER C.: Perspective shadow maps. In *Game Programming Gems IV*. 2003. [1](#)

[SWP10] SCHERZER D., WIMMER M., PURGATHOFER W.: A survey of real-time hard shadow mapping methods. In *State of the Art Reports Eurographics* (2010). [1](#)

[WSP04] WIMMER M., SCHERZER D., PURGATHOFER W.: Light space perspective shadow maps. In *Rendering Techniques 2004 (Proceedings Eurographics Symposium on Rendering)* (2004), pp. 143–151. [1](#)