

Supplementary Document of “Beyond Hard Shadows: Moment Shadow Maps for Single Scattering, Soft Shadows and Translucent Occluders”

Christoph Peters* Cedrick Münstermann* Nico Wetzstein* Reinhard Klein*
University of Bonn, Germany

This supplementary document provides some relevant mathematical derivations and technical details which cannot be part of the paper due to space constraints. Readers interested in further implementation details are advised to take a look at the provided shader code. Additional results can be found in the supplementary video.

1 Bounds for Rectified Coordinates

In the paper we state that we compute bounds for r , φ and θ such that the entire view frustum is covered. This procedure is not entirely trivial and in the following we provide some additional details. A reference implementation in HLSL is provided as part of the supplementary code.

Single scattering should be accumulated over the entire view ray. Thus, we do not consider the near clipping plane. Let $q_1, \dots, q_4 \in \mathbb{R}^3$ be the coordinates of the four vertices of the far clipping plane of the camera used for main scene rendering in light view space. Without loss of generality let the camera position be in the origin of the coordinate system. Then the maximal value for r is given by

$$r_{\max} := \max_{i \in \{1, \dots, 4\}} \sqrt{(q_i)_1^2 + (q_i)_2^2}.$$

Note that $(q_i)_j$ denotes the j -th entry of the vector q_i for $j \in \{1, 2, 3\}$. Since we ignore the near plane $r_{\min} := 0$.

To compute φ_{\min} and φ_{\max} we compute the maximal pairwise angle enclosed by the vectors $((q_i)_1, (q_i)_2)^T \in \mathbb{R}^2$ for $i \in \{1, \dots, 4\}$. The azimuth of one of the two involved vectors provides φ_{\min} , the other provides φ_{\max} . A special case arises if the light direction or flipped light direction lie within the view frustum (i.e. none of the side clipping planes clips them). In both cases we have to set $\varphi_{\min} = 0$ and $\varphi_{\max} = 2 \cdot \pi$ to indicate that the boundary of the far clipping plane surrounds the camera position in the shadow map.

The computation of θ_{\min} and θ_{\max} is more complicated because the extremal inclination may be realized at the vertices, on the edges or within the area of the far clipping plane. It is convenient to exploit that inclinations depend monotonously on the z -coordinate of normalized vectors, i.e.

$$\theta_i := \arccos \frac{(q_i)_3}{\|q_i\|_2}.$$

Taking the minimal and maximal values of $\theta_1, \dots, \theta_4$ yields the extrema at vertices. The inclination is extremal within the area of the far clipping plane if and only if the light direction or flipped light direction lie within the view frustum (see above). In this case an extremal inclination is $\theta_{\min} = 0$ or $\theta_{\max} = \pi$, respectively.

To find extrema on an edge of the far plane connecting corner points $i, j \in \{1, \dots, 4\}$ we take the derivative of the z -coordinate of normalized vectors on the edge to find critical points:

$$\frac{\partial}{\partial t} \frac{(q_i + t \cdot (q_j - q_i))_3}{\|q_i + t \cdot (q_j - q_i)\|_2} = 0$$

This equation has the unique solution

$$t = \frac{(q_i)_3 \cdot q_i^T \cdot (q_j - q_i) - (q_j - q_i)_3 \cdot \|q_i\|_2^2}{(q_j - q_i)_3 \cdot q_i^T \cdot (q_j - q_i) - (q_i)_3 \cdot \|q_j - q_i\|_2^2}.$$

If $t \in [0, 1]$ we may need to adapt $[\theta_{\min}, \theta_{\max}]$ to include the inclination at this point on the ray.

Note that this whole algorithm only needs to be executed once per frame to get single scattering for one directional light.

2 Sparse Representations of Moments

To filter moments during resampling the paper uses a novel algorithm. This algorithm takes in three moments and represents them by a combination of two Dirac- δ distributions. In the following we prove correctness of this algorithm.

Theorem 1. *Let $b_1, b_2, b_3 \in \mathbb{R}$ be moments with positive variance $\sigma^2 := b_2 - b_1^2 > 0$. Let*

$$c := \begin{pmatrix} c_0 \\ c_1 \\ c_2 \end{pmatrix} := \begin{pmatrix} b_1 \cdot b_3 - b_2^2 \\ b_1 \cdot b_2 - b_3 \\ b_2 - b_1^2 \end{pmatrix} \in \mathbb{R}^3.$$

Then the polynomial $c_2 \cdot z^2 + c_1 \cdot z + c_0 = 0$ has two distinct roots $z_2, z_3 \in \mathbb{R}$. Let $w_3 := \frac{b_1 - z_2}{z_3 - z_2}$ and $w_2 := 1 - w_3$. Then we know

$$\forall j \in \{1, 2, 3\} : b_j = w_2 \cdot z_2^j + w_3 \cdot z_3^j.$$

Proof. The proof builds upon the fact that sequences of moments inducing a singular Hankel matrix have exactly one representation by a sparse distribution with half as many points of support as moments are given [Peters and Klein 2015].

Let $b_4 \in \mathbb{R}$ be a fourth moment and consider the Hankel matrix

$$B := \begin{pmatrix} 1 & b_1 & b_2 \\ b_1 & b_2 & b_3 \\ b_2 & b_3 & b_4 \end{pmatrix}.$$

The determinant of this matrix depends linearly upon b_4 :

$$\det B = b_4 \cdot \det \begin{pmatrix} 1 & b_1 & 0 \\ b_1 & b_2 & 0 \\ b_2 & b_3 & 1 \end{pmatrix} + \det \begin{pmatrix} 1 & b_1 & b_2 \\ b_1 & b_2 & b_3 \\ b_2 & b_3 & 0 \end{pmatrix}$$

Note that

$$\det \begin{pmatrix} 1 & b_1 & 0 \\ b_1 & b_2 & 0 \\ b_2 & b_3 & 1 \end{pmatrix} = \det \begin{pmatrix} 1 & b_1 \\ b_1 & b_2 \end{pmatrix} = b_2 - b_1^2 = \sigma^2 > 0.$$

Thus, we can choose b_4 such that B is a positive semi-definite, singular matrix.

* {peters,muenste,wetzstei,rk}@cs.uni-bonn.de

Now we observe that c is a non-zero vector in the kernel of B :

$$\begin{aligned} B \cdot c &= \begin{pmatrix} 1 & b_1 & b_2 \\ b_1 & b_2 & b_3 \\ b_2 & b_3 & b_4 \end{pmatrix} \cdot \begin{pmatrix} b_1 \cdot b_3 - b_2^2 \\ b_1 \cdot b_2 - b_3 \\ b_2 - b_1^2 \end{pmatrix} \\ &= \begin{pmatrix} b_1 \cdot b_3 - b_2^2 + b_1 \cdot (b_1 \cdot b_2 - b_3) + b_2 \cdot (b_2 - b_1^2) \\ b_1 \cdot (b_1 \cdot b_3 - b_2^2) + b_2 \cdot (b_1 \cdot b_2 - b_3) + b_3 \cdot (b_2 - b_1^2) \\ b_2 \cdot (b_1 \cdot b_3 - b_2^2) + b_3 \cdot (b_1 \cdot b_2 - b_3) + b_4 \cdot (b_2 - b_1^2) \end{pmatrix} \\ &= \begin{pmatrix} 0 \\ 0 \\ \det B \end{pmatrix} = 0 \end{aligned}$$

Therefore, it can be used for construction of the unique sparse representation of the moments $1, b_1, b_2, b_3, b_4$ [Peters and Klein 2015, Section 7.3]. From this we immediately see that distinct roots z_2, z_3 must exist because they are points of support of this sparse distribution. Furthermore, the corresponding weights w_2 and w_3 are uniquely determined by the system of linear equations

$$\begin{pmatrix} 1 & 1 \\ z_2 & z_3 \end{pmatrix} \cdot \begin{pmatrix} w_2 \\ w_3 \end{pmatrix} = \begin{pmatrix} 1 \\ b_1 \end{pmatrix}$$

and this is solved by setting w_2, w_3 as specified. \square

Note that the algorithm presented in the paper computes c_0 from c_1 and c_2 . This scheme has been derived via a Cholesky decomposition and gives a crucial advantage with regard to numerical stability.

3 Six Moment Shadow Mapping

As explained in the paper six moment shadow mapping requires an affine transform to optimize the representation of moments for minimal loss of information during quantization. We found this transform by numerical optimization on a cluster of computers and it is given in Table 1.

The volume of the set of meaningful vectors of moments is given by

$$\prod_{j=1}^6 \frac{((j-1)!)^2}{(2 \cdot j - 1)!} = \frac{1}{44008272000}$$

[Peters and Klein 2015, Proposition 19]. This means that an optimal quantization for a uniform distribution of moments could gain

$$\log_2 44008272000 \approx 35.357$$

bits of entropy. This certainly cannot be achieved by an affine transform but the differential entropy gain of 30.5 bits for the transform in Table 1 gets reasonably close. Further optimization in the space of affine transforms could improve on this slightly.

For an implementation of the 4×4 Cholesky decomposition and the solution of the cubic polynomial we refer the reader to the provided shader code.

In the paper we state that we compute the relevant sum of the weights w_1, w_2, w_3, w_4 by constructing an interpolation polynomial. This works as follows. Suppose we want the linear combination of weights $\sum_{i=1}^4 f_i \cdot w_i$ where $f_1, f_2, f_3, f_4 \in \mathbb{R}$. In practice we set $f_1 := \beta$ and for $i \in \{2, 3, 4\}$

$$f_i := \begin{cases} 1 & \text{if } z_i < z_f, \\ 0 & \text{otherwise.} \end{cases}$$

By definition of the moments b_1, b_2, b_3

$$\begin{pmatrix} 1 & 1 & 1 & 1 \\ z_1 & z_2 & z_3 & z_4 \\ z_1^2 & z_2^2 & z_3^2 & z_4^2 \\ z_1^3 & z_2^3 & z_3^3 & z_4^3 \end{pmatrix} \cdot \begin{pmatrix} w_1 \\ w_2 \\ w_3 \\ w_4 \end{pmatrix} = \begin{pmatrix} 1 \\ b_1 \\ b_2 \\ b_3 \end{pmatrix}.$$

It follows that

$$\begin{aligned} \sum_{i=1}^4 f_i \cdot w_i &= \begin{pmatrix} f_1 \\ f_2 \\ f_3 \\ f_4 \end{pmatrix}^\top \cdot \begin{pmatrix} w_1 \\ w_2 \\ w_3 \\ w_4 \end{pmatrix} \\ &= \begin{pmatrix} f_1 \\ f_2 \\ f_3 \\ f_4 \end{pmatrix}^\top \cdot \begin{pmatrix} 1 & 1 & 1 & 1 \\ z_1 & z_2 & z_3 & z_4 \\ z_1^2 & z_2^2 & z_3^2 & z_4^2 \\ z_1^3 & z_2^3 & z_3^3 & z_4^3 \end{pmatrix}^{-1} \cdot \begin{pmatrix} 1 \\ b_1 \\ b_2 \\ b_3 \end{pmatrix} \\ &= \begin{pmatrix} 1 \\ b_1 \\ b_2 \\ b_3 \end{pmatrix}^\top \cdot \left(\begin{pmatrix} 1 & z_1 & z_1^2 & z_1^3 \\ 1 & z_2 & z_2^2 & z_2^3 \\ 1 & z_3 & z_3^2 & z_3^3 \\ 1 & z_4 & z_4^2 & z_4^3 \end{pmatrix} \right)^{-1} \cdot \begin{pmatrix} f_1 \\ f_2 \\ f_3 \\ f_4 \end{pmatrix} \end{aligned}$$

The matrix in this last expression is a Vandermonde matrix. Thus, the vector

$$d := \begin{pmatrix} 1 & z_1 & z_1^2 & z_1^3 \\ 1 & z_2 & z_2^2 & z_2^3 \\ 1 & z_3 & z_3^2 & z_3^3 \\ 1 & z_4 & z_4^2 & z_4^3 \end{pmatrix}^{-1} \cdot \begin{pmatrix} f_1 \\ f_2 \\ f_3 \\ f_4 \end{pmatrix}$$

holds coefficients of the interpolation polynomial taking value f_i at point z_i . Such a polynomial can be constructed efficiently and robustly as Newton polynomial using divided differences. Once the coefficients have been computed, the end result is given by

$$d_1 + \sum_{i=1}^3 d_{i+1} \cdot b_i.$$

4 Biasing for Clamped Depth Values

Moment shadow mapping requires biasing of moments. Otherwise quantization errors can lead to meaningless vectors of moments which cannot be explained by any distribution of depth values. The biasing originally proposed for four moment shadow mapping with 16-bit quantization is

$$\begin{aligned} b' &:= \alpha \cdot b^* + (1 - \alpha) \cdot b & \text{where} \\ b^* &:= \left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2} \right)^\top & \text{and } \alpha := 3 \cdot 10^{-5}. \end{aligned}$$

Here $b \in \mathbb{R}^4$ is the vector of moments from the moment shadow map (i.e. including quantization errors) and b' is the biased vector of moments that is used as input to the reconstruction algorithm.

The reasoning behind this is that this choice is optimal in the average case (under certain assumptions on the distribution of b) [Peters 2013]. The vector b^* corresponds to the distribution $Z^* := \frac{1}{2} \cdot (\delta_0 + \delta_1)$ so we effectively add virtual shadow casters at the extreme depth values zero and one with a weight of $\frac{\alpha}{2}$.

From this representation the worst case becomes apparent. If b corresponds to a distribution of depth values taking only the values zero and/or one, the biased vector b' still corresponds to such a distribution. Such distributions with no more than two different values

$$\theta_6(b) := \begin{pmatrix} -0.39982174417 & 5.372499433 & -19.400931206 & 5.9601506283 & 21.716342727 & -14.2481344 \\ -0.2036458037 & 61.687438614 & -280.26108456 & 480.86777452 & -370.6170286 & 109.52605942 \\ 35.997693025 & -419.96056141 & 1791.8207666 & -3455.6578547 & 3071.7060681 & -1023.906108 \\ 24.253667437 & -187.46243964 & 499.10228252 & -516.52881873 & 144.63179523 & 37.003250859 \\ 14.571662352 & -64.120119623 & 75.292307128 & 8.798154992 & -43.073239823 & 8.5315616209 \\ 2.5464326857 & -78.04982013 & 350.77691236 & -674.56153641 & 598.68902736 & -199.39897152 \end{pmatrix} \cdot b$$

$$+ \begin{pmatrix} 0.9999446933 \\ 0.0003698424 \\ 0.0000000000 \\ 0.0000807802 \\ 0.0000059274 \\ 0.9774602669 \end{pmatrix}$$

Table 1: The affine transform which should be applied to vectors of six moments before quantization to minimize loss of information.

correspond to the boundary of the set of meaningful vectors of four moments. Therefore, the above biasing is not efficient in this case.

This worst case becomes relevant as soon as depth values are being clamped into the interval $[0, 1]$. For prefiltered single scattering we employ such clamping. Therefore, we choose to add a uniform component to the distribution Z^* . This way we ensure that b^* lies in the interior of the (convex) set of meaningful vectors of moments. In consequence sufficiently strong biasing guarantees valid vectors of moments. In our experiments we use $\alpha = 3 \cdot 10^{-5}$ and

$$b^* = \frac{4}{5} \cdot \left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2} \right)^T + \frac{1}{5} \cdot \left(\frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \frac{1}{5} \right)^T$$

for prefiltered single scattering with four moments.

For prefiltered single scattering with six moments we proceed in an analogous fashion but with a stronger weight on the uniform component of the bias. This is necessary because when using six moments distributions with three different depth values still produce vectors of moments on the boundary of the meaningful domain. Thus, the bias is defined by $\alpha := 4 \cdot 10^{-3}$ and

$$b^* = \frac{1}{2} \cdot \left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2} \right)^T + \frac{1}{2} \cdot \left(\frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \frac{1}{5}, \frac{1}{6}, \frac{1}{7} \right)^T.$$

References

- PETERS, C., AND KLEIN, R. 2015. Moment shadow mapping. In *Proceedings of the 19th Symposium on Interactive 3D Graphics and Games*, ACM, 7–14.
- PETERS, C. 2013. *Moment Shadow Mapping*. Master’s thesis, University of Bonn.