Volumetric Percentage Closer Soft Shadows

Andreas Klein, Björn Tappert, Alfred Nischwitz
Munich University of Applied Sciences
Lothstrasse 64

80335 Munich, Germany
andreas.klein@hm.edu, bjoern.tappert@gmx.de,
nischwitz@cs.hm.edu

Paul Obermeier
MBDA Deutschland GmbH
Hagenauer Forst 27

86529 Schrobenhausen, Germany
paul.obermeier@mbda-systems.de

ABSTRACT

Percentage Closer Soft Shadows is a popular technique to generate contact hardening soft shadows with shadow mapping. Recent research in shadow generation for translucent objects makes it possible to realize shadows for translucent objects in real-time environments. However, for multiple translucent blockers it is unclear how an appropriate blocker depth can be calculated. In this paper, we propose a method to calculate a blocker depth for multiple translucent blockers and therefore, enabling physically plausible soft shadows for opaque and translucent objects in a single approach.

Keywords

Soft Shadows, Contact Hardening Soft Shadows, PCSS, Volumetric Shadows, AVSM

1 INTRODUCTION

Shadow Mapping is a popular method to generate shadows for opaque objects in real-time rendering. The idea is to realize a visibility test by comparing the depth value as seen from the camera with the depth value stored in a depth map.

As shadow mapping assumes point light sources, only hard shadows will be produced. However, it is possible to simulate soft shadows of area light sources by making multiple shadow tests within a filter window and averaging the result. However, these shadows are not physically plausible as the shadows are uniformly soft. In order to achieve physically plausible shadows, the size of the penumbra must be adapted according to the distance between a light blocker and a shadow receiver. These shadows are called contact hardening soft shadows, as the shadow softness increases with the blocker-receiver distance. Contact hardening soft shadows can be realized with shadow mapping by adapting the filter window based on the blocker - receiver distance.

Recent work in shadow generation for translucent objects makes it possible to integrate shadows for translucent objects in real-time environments, such as games. In contrast to shadow maps, which only store the nearest depth value as seen from the light, approaches for translucent shadows store a transmittance function per pixel (Figure 1). A transmittance function encodes the light visibility at each depth value.

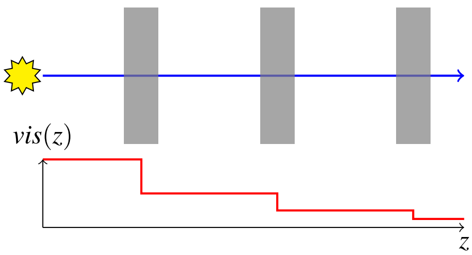

Figure 1: The light intensity is reduced as it passes through a set of translucent blockers. A transmittance function encodes the transmittance at each given z value.

Analogous to shadow mapping, a transmittance function can be used to generate shadows. A shadow exists, if the light intensity of a point light source is reduced by one or multiple translucent blockers. However, these shadows have hard boundaries, despite they appear to be soft compared to opaque blockers since these shadows have a reduced darkness according to the blocker's translucency. As with shadow mapping, it is possible to simulate soft shadows of area light sources by making multiple comparisons within a filter window. However, contact hardening soft shadows requires the distance between a blocker and a receiver to estimate the penumbra width which is used to adapt the filter window. For multiple translucent blockers, it is unclear how an appropriate blocker depth can be calculated.

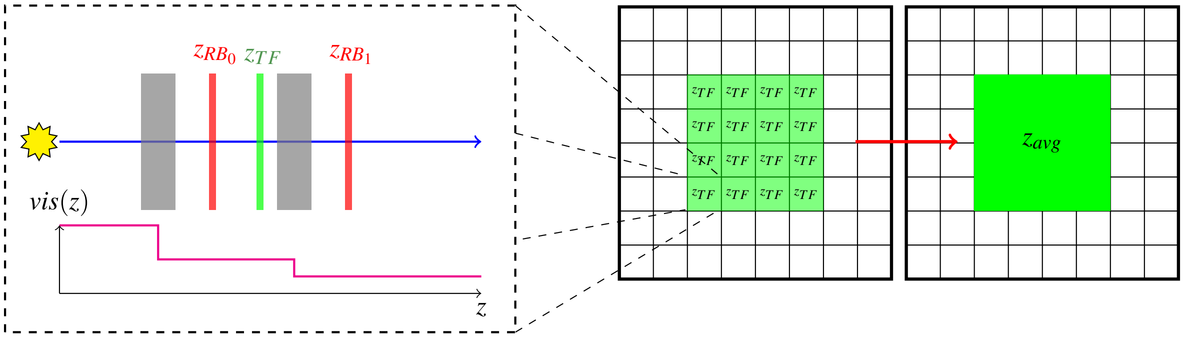

Figure 2: Our algorithm proceeds as follows. First, we sample a single transmittance function and calculate a replacement blocker z_{RB} (red) for each translucent blocker. Second, we compute a weighted average z_{TF} (green) for all replacement blockers of a transmittance function. We repeat these steps for each transmittance function in a filter window and calculate a total replacement blocker z_{avg} for the filter window. Finally, we use the total replacement blocker to estimate the penumbra size.

In this paper, we propose a method to calculate a blocker depth for multiple translucent blockers and therefore, enabling physically plausible soft shadows for opaque and translucent objects in a single approach.

2 RELATED WORK

We focus our review on publications closely related to our work. See Eisemann et al. [Eis11a] for a comprehensive survey on other shadow algorithms.

Shadows from Opaque Blockers

Percentage-closer filtering (PCF) [Ree87a] is a popular method for generating soft shadows. The idea is to build a shadow factor by making multiple shadow comparisons within a user defined filter window. Fernando [Fer05a] proposed percentage-closer soft shadows (PCSS) in order to generate contact hardening soft shadows with PCF. The idea is to first search for blockers within a given filter window and calculate an average blocker depth. A penumbra width can then be estimated using a parallel planes approximation. The penumbra width is used to scale the PCF window.

Shadows from Translucent Blockers

Deep Shadow Maps [Lok00a] stores nodes of a transmittance function per pixel and compresses them to guarantee a fixed absolute error. Salvi et al. [Sal10a] uses in their Adaptive Volumetric Shadow Maps (AVSM) an area based metric to compress a transmittance function with a fixed number of nodes. Several approaches use a basis transformation to compress a transmittance function, e.g. [Jan10a, Del11a].

Depth peeling [Eve01a, Liu06a, Bav08a] uses multiple render passes with dual depth comparison to extract the depth layers of geometry. However, this approach suffers from an unbounded number of rendering passes for complex geometry. Stochastic Transparency [End11a, McG11a] uses a randomized sub-pixel stipple pattern to realize screen door transparency with a fixed set of Multi Sample Anti Aliasing (MSAA) samples.

3 ALGORITHM

Based on the parallel planes approximation of Fernando [Fer05a], we derive an average blocker depth for a list of translucent blockers. We assume that the transmittance functions have already been generated in each frame using a translucent shadow technique, such as AVSM [Sal10a].

The algorithm proceeds as follows (Figure 2). First, we sample a single transmittance function in order to receive two pairs of transmittance - depth values, which represent an extended translucent blocker (Figure 3). From these values we calculate an infinitesimal thin replacement blocker described by a single transmittance - depth pair. Second, we integrate over all samples of a single transmittance function and thus over all blockers along a ray from the light source by computing a weighted average of all replacement blockers. Third, we process step one and two for a set of transmittance functions within a filter window to estimate a total replacement blocker and calculate its depth value. Fourth, we derivate a penumbra width with a parallel planes approximation [Fer05a]. Finally, we generate the shadows by making multiple shadow tests and averaging the result.

Shadows from Translucent Blockers

In order to calculate a shadow factor for translucent blockers with shadow mapping, a modification to the binary shadow test function is necessary.

The intensity of the light is reduced by the transmittance value of a blocker. For a single translucent blocker with depth z and an alpha value α the translucent shadow test function is given by:

$$S_T(z) = (1-\alpha) + \alpha S(z_L, z_S)$$

where z_L is the blocker depth transformed into light's coordinate system (light space), z_S is the value in the depth map and $S(z_L, z_S)$ is the binary shadow test function:

$$S(z_L, z_S) = \begin{cases} 0, & \text{if } z_L > z_S \\ 1, & \text{if } z_L \leq z_S \end{cases}$$

For multiple overlapping blockers, the translucent shadow test function can be applied using the over operator [Por84a]:

$$S_T(z_i) = (1-\alpha_i) S_T(z_{i-1}) + \alpha S(z_{i_L}, z_{i_S}) =$$
$$= \sum_{m=0}^{i} \left[\alpha_{i-m} S(z_{i-m_L}, z_{i-m_S}) \prod_{n=i-m+1}^{i} (1-\alpha_n) \right]$$

Estimate a Replacement Blocker for a Single Transmittance Function

In contrast to standard shadow mapping approaches, multiple translucent blockers are available for each pixel and each would produce a different penumbra. In order to calculate a contact hardening soft shadow with a parallel planes approximation, such as in PCSS, an average blocker depth must be estimated. Our idea is to replace multiple blockers with one appropriate replacement blocker. First, we show how a replacement blocker can be calculated for a single translucent blocker and then how they can be averaged for multiple blockers.

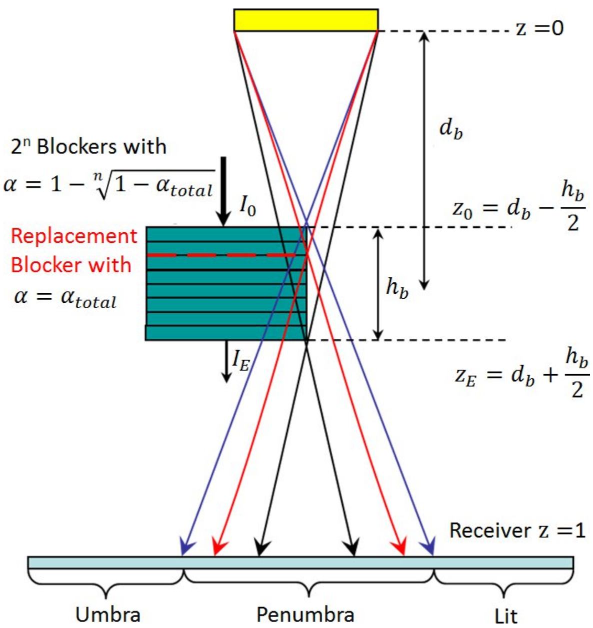

Figure 3: Our idea is to replace a blocker of dimension h_b with n thin layers and compute a replacement blocker (depicted as a red dashed line) with the total transmittance $(1-\alpha_{total})$.

We assume that a single blocker has a constant absorption with the coefficient α which reduces the light intensity I_0 (Figure 3):

$$I_E = (1-\alpha) I_0 \tag{1}$$

In a homogenous material, the reduction of the light intensity can be expressed using the entry depth of a blocker z_0 and an absorption coefficient b:

$$I(z) = I_0 \cdot e^{-b(z-z_0)} \tag{2}$$

Due to the exponential reduction, we place the replacement blocker at the position where the light intensity is reduced by half of the difference value:

$$\left(\frac{I_E - I_0}{2} \right) \tag{3}$$

We choose this position for the replacement blocker, as the mean squared error to a reference solution is smaller compared to a medium position (see Section 3).

The depth of the replacement blocker is then given by

$$z_{RB} = z_0 - \frac{1}{b} \cdot ln \left(\frac{I_E + I_0}{2 \cdot I_0} \right) \tag{4}$$

with:

$$b = -\frac{1}{z_E - z_0} \cdot ln(1-\alpha) \tag{5}$$

We express the depth of the replacement blocker for a single translucent blocker using the alpha value α:

$$z_{RB} = z_0 - (z_E - z_0) \cdot \frac{ln\left(1 - \frac{\alpha}{2}\right)}{ln(1-\alpha)} \tag{6}$$

In order to obtain a single replacement blocker depth z_{TF} for a transmittance function, which may consists of multiple translucent blockers, we calculate an weighted average of all replacement blockers:

$$z_{TF} = \frac{\sum_{i=1}^{n} \Delta_i \cdot z_{RB_i}}{\sum_{i=1}^{n} \Delta_i} \tag{7}$$

where Δ_i is a measure for the part of the light intensity which is blocked by the i-th replacement blocker:

$$\Delta_i = \alpha_i \prod_{j=1}^{i-1} (1-\alpha_j) \tag{8}$$

Analysis of the Approximation

One way to calculate PCSS for multiple translucent blockers is to evaluate the PCSS function for each sample on a transmittance function T:

$$S_f(z_R) \approx \int_{i \in T} \alpha_i PCSS(z_R, z_i) vis(z_i) dz \tag{9}$$

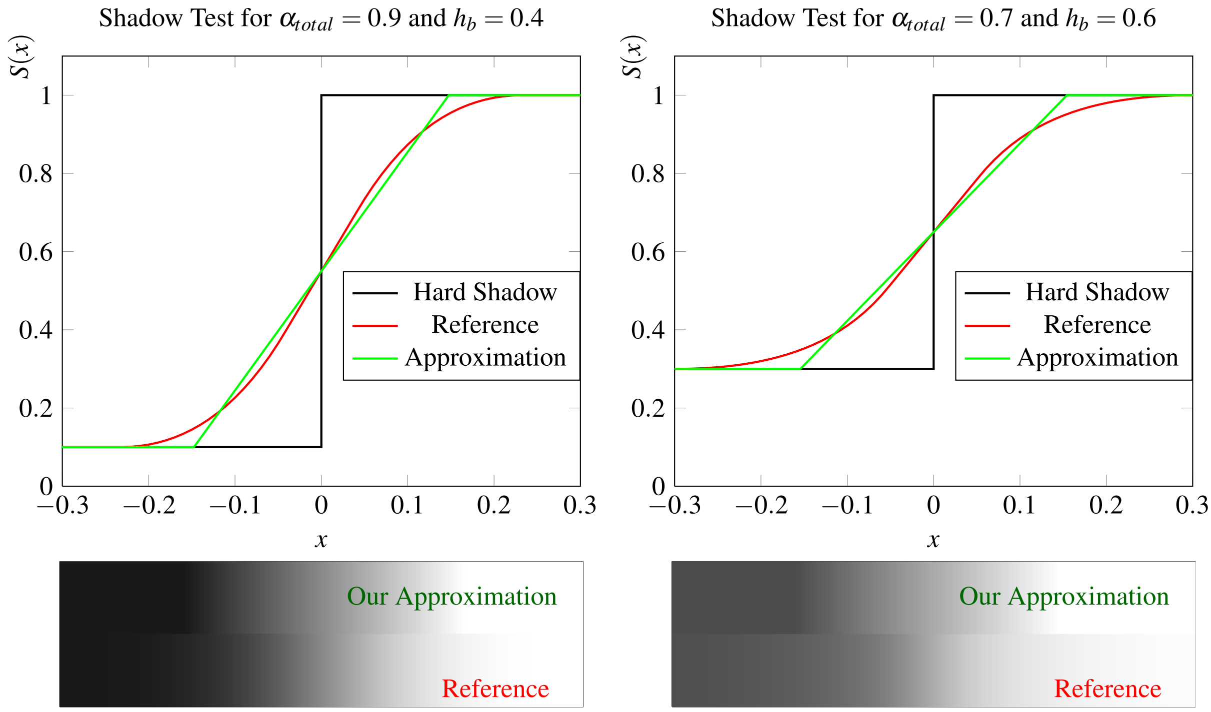

Figure 4: Comparison of the shadow test of our approximation that uses a single replacement blocker against a reference solution with n infinitesimal blockers. The plots at the bottom displays the resulting shadow factors at a shadow boundary.

with:

$$vis(z) = \prod_{z_j < z} (1 - \alpha_j) \tag{10}$$

$$PCSS(z_R, z_B) = \omega_L \frac{z_R - z_B}{z_B} \sum_{n=0}^{k} f(z_n - z_r) S(z_n, z_R) \tag{11}$$

where f is a filter function, ω_L the diameter of the light source, z_R the depth of the receiver and z_B the depth of the blocker. However, this solution requires a shadow factor computation for each single blocker. Instead, we approximate this integral by using a single replacement blocker with total alpha:

$$S_f(z_R) \approx \left(\sum_i \alpha_i \right) PCSS(z_R, z_{TF}) \tag{12}$$

In Figure 4, we compare this approximation of the shadow test function against a reference solution (Eq. 9) for different blocker sizes and alpha values. It can be observed, that our approximation calculates a linear transition between penumbra and lit regions and vice versa where a smooth transition is correct.

Figure 5 compares the shadow test function when choosing different depths for a replacement blocker. By using a replacement blocker at half of the absorption, the mean squared error to the reference solution is smaller compared to a shadow test function that uses a medium depth for a replacement blocker.

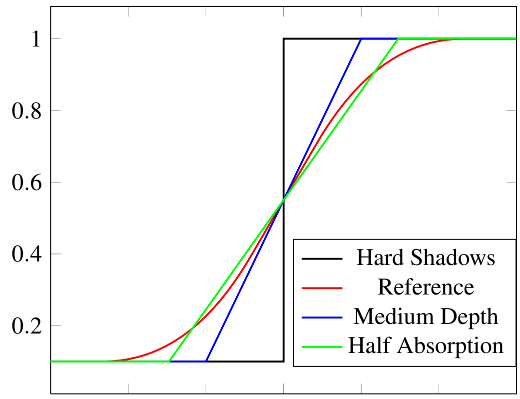

Figure 5: We choose the position where the light intensity is reduced by half of the difference value $\frac{I_E - I_0}{2}$ (half absorption), as it approximates the reference solution more accurate.

As our approximation is based on PCSS, we introduce the same limitations, such as overestimating the penumbra area.

4 IMPLEMENTATION

We implemented our algorithm using Adaptive Volumetric Shadow Maps [Sal10a].

Our implementation proceeds as follows. First, we render an AVSM in each frame. Second, we sample the AVSM within the given search radius to calculate an average blocker depth. We terminate the algorithm, if there are no pixels with a transmittance smaller than one. Otherwise, we calculate a replacement blocker for each transmittance function in the given search radius and calculate the average depth as follows:

$$z_{avg} = \frac{\sum_{i=1}^{n} \Delta_i \cdot z_{TF_i}}{\sum_{i=1}^{n} \Delta_i} \tag{13}$$

In the next step, we estimate the penumbra width using the parallel planes approximation [Fer05a]:

$$\omega_{Penumbra} = \frac{(z_R - z_{avg})}{z_{avg}} \omega_L \tag{14}$$

Finally, we generate the shadow factor by comparing multiple samples and averaging the result.

5 RESULTS

We used 8 nodes for the transmittance function in the AVSM implementation. The shadows of our algorithm were generated using 25 Poisson samples for the blocker search as well as in the final shadow factor computation. The AVSM size was 1024 x 1024 and the screen resolution was 1680 x 1050. Table 1 shows the performance results and Figure 7 compares the visual results. The reference solution is realized by replacing the area light source with 128 point light sources. Figure 6 shows an example scene with opaque and translucent blockers. The performance results were obtained on an Intel Xeon E5620 CPU with 2.4 GHz, 8 GB RAM and a NVIDIA GeForce GTX 680 graphics card with 2048 MB memory.

	AVSM Hard Shadows	VPCSS Soft Shadows
Dragon (Fig. 7a) (871K tris)	5.1 ms	10.1 ms
Particle (Fig. 7b) (5K particles)	3.8 ms	9.4 ms
Hairball (Fig. 7c) (2.88M tris)	26.3 ms	33.4 ms
Tank (Fig. 6) (136K tris 5K particles)	4.3 ms	10.9 ms

Table 1: Performance results in milliseconds with a screen resolution of 1680 x 1050. Note that these timings also include the generation of the transmittance functions with AVSM.

Figure 6: Example scene with soft shadows resulting from opaque and translucent blockers.

6 CONCLUSIONS AND FUTURE WORK

We presented a method to calculate an average blocker depth for multiple translucent blockers. Analog to standard shadow mapping, a contact hardening soft shadow can now be generated by using a constant amount of shadow tests. The results show that our algorithm creates physically plausible soft shadows for opaque and translucent blockers.

For future work we wish to investigate how a replacement blocker can be calculated for wavelength-dependent transmissive blockers.

7 ACKNOWLEDGMENTS

We thank the anonymous reviewers for their useful comments. The work of A. Klein is funded by MBDA Deutschland GmbH.

8 REFERENCES

[Bav08a] Bavoil, L. and Meyers, K. Order independent transparency with dual depth peeling. NVIDIA technical report, 2008.

[Del11a] Delalendre, C., Gautron, P., Marvie, J.E. and Francois, G. Transmittance Function Mapping. In Conf.proc of Interactive 3D Graphics and Games 2011. 2011.

[Eis11a] Eisemann E., Schwarz M., Assarsson U., Wimmer M. Real-Time Shadows, Taylor & Francis, 2011.

[End11a] Enderton, E., Sintorn, E., Shirley, P. and Luebke, D. Stochastic Transparency. IEEE Transactions on Visualization and Computer Graphics August 2011. 2011.

[Eve01a] Everitt, C. Interactive order-independent transparency. NVIDIA white paper, 2001.

[Fer05a] Fernando R. Percentage-Closer Soft Shadows. ACM SIGGRAPH 2005 Sketches, 2005.

[Jan10a] Jansen, J. and Bavoil, L. Fourier Opacity Mapping. In Conf.proc of Interactive 3D Graphics and Games 2010. 2010.

[Liu06a] Liu, B., Wei, L.-Y, and Xu, Y.-Q. Multi-layer depth peeling via fragment sort. Microsoft Research technical report. 2006

[Lok00a] Lokovic, T and Veach, E. Deep Shadow Maps. In Conf.proc of SIGGRAPH 00. 2000.

[McG11a] McGuire, M. and Enderton, E. Colored Stochastic Shadow Maps. In Conf.proc of Interactive 3D Graphics and Games 2011. 2011.

[Por84a] Porter, T and Duff, T. Composing digital images. In Conf.proc. SIGGRAPH 84, 1984.

[Ree87a] Reeves W. T., Salesin D. H., Cook R. L. Rendering antialiased shadows with depth maps. In Conf.proc SIGGRAPH 87, ACM, 283-291, 1987.

[Sal10a] Salvi, M., Vidimce, K., Lauritzen, A. and Lefohn A. Adaptive Volumetric Shadow Maps. In Conf.proc. EGSR 2010. 2010

[Wil78a] Williams L. Casting curved shadows on curved surfaces. In Conf.proc. SIGGRAPH 78, ACM, 270-274, 1978.

(a)

(b)

(c)

Figure 7: Resulting shadows from the dragon (a), particle (b) and hairball (c) datasets. From left to right: Hard shadows, our algorithm and reference solution. (a) The dragon was rendered with $\alpha = 0.3$. (b) In the particle dataset, the shadows resulting from a hard shadow test look already soft. However, our algorithm as well as the reference solution softens the shadows further, as the blocker-receiver distance increases. (c) The hairball dataset was rendered with $\alpha = 1.0$.