Computers & Graphics xxx (2017) xxx–xxx

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag



Special Section on Graphics Interface 2017

# Improved anti-aliasing for Euclidean distance transform shadow mapping

Q1 Márcio Macedo*, Antônio Apolinário

*Federal University of Bahia, Av. Ademar de Barros, 40.170-110 Bahia, Brazil*

ARTICLE INFO

*Article history:*
Received 30 August 2017
Revised 18 November 2017
Accepted 18 November 2017
Available online xxx

*Keywords:*
Rendering
Real time
Shadows
Anti-aliasing

ABSTRACT

High-quality, real-time penumbra rendering remains a challenging problem in computer graphics. Existing techniques for real-time fixed-size penumbra simulation generate aliasing, banding or leaking artifacts that diminish the realism of shadow rendering. Euclidean distance transform shadow mapping aims to solve that by using a normalized Euclidean distance transform to simulate penumbra on the basis of anti-aliased hard shadows generated by revectorization-based shadow mapping. Despite the high visual quality obtained with such a technique, the anti-aliasing provided by shadow revectorization comes at the cost of shadow overestimation artifacts that are introduced in the scene. In this paper, we propose an improved algorithm for Euclidean distance transform shadow mapping by reformulating the visibility function of revectorization-based shadow mapping. Through an additional detailed analysis of the results, we show that we are able to reduce shadow overestimation artifacts for penumbra simulation, generating shadows with higher visual quality than previous fixed-size penumbra shadowing methods, while keeping real-time performance for shadow rendering.



## 1. Introduction

Shadows are essential in several computer graphics applications, such as games and augmented reality, because they add a compelling effect, increasing the visual perception of the user with respect to the rendering of virtual scenes [1]. As pointed in [2], users usually prefer realistic shadows over fake ones when looking into virtual scenes. Unfortunately, accurate shadow rendering is still not feasible for real-time applications, mainly because a high number of samples must be taken from an area light source to approximate the direct illumination term of the rendering equation [3,4], making the shadowing process costly.

One of the most traditional ways to compute shadows in real time is shadow mapping [5]. By approximating the area light source by a single point light source, this technique discretizes the 3D virtual scene, as seen from the point light source viewpoint, into a depth buffer named *shadow map* that is used to aid the real-time shadow computation. However, shadows generated on the basis of a shadow map are prone to aliasing artifacts and temporal incoherence due to the finite resolution of the shadow map. Moreover, differently from an area light source, a point light source is not able, in essence, to cast penumbra in the scene because this type of light source is infinitesimal, such that it cannot be partially occluded in the scene. Therefore, shadow mapping is only able to simulate hard shadows (i.e., shadows without the penumbra effect) in the scene. Unfortunately, such hard shadows are unrealistic, because they are not much present in the real world.

Aliasing artifacts are commonly suppressed by the use of texture linear filtering techniques, such as mip-mapping [6] and anisotropic filtering [7]. However, these strategies cannot be directly applied in the shadow map, because shadow mapping uses a non-linear shadow test to determine the visibility condition of a given fragment [8]. Then, several techniques have been proposed to allow shadow map filtering. Existing techniques either realize shadow filtering after the shadow test [9,10] or filter the shadow map (as done in [11,12]), such that the shadows produced by a modified version of the shadow test are already filtered and anti-aliased. While these techniques minimize aliasing artifacts and simulate *fixed-size* penumbra, they introduce new artifacts in shadow rendering because of the filtering strategy used. Banding artifacts may appear in the final rendering if low-order filter sizes are used to keep real-time performance [9] (Fig. 1(a)). Q3
Techniques that filter the shadow map before the shadow test are prone to light leaking artifacts (in which a shadowed region is erroneously assumed as a lit region) because the filtering may incorrectly affect the shadow test result [13,14] (green closeup in Fig 1(b)). Techniques that filter the shadow map after the shadow test are prone to shadow overestimation artifacts because, during

Q2 * Corresponding author.
*E-mail address:* antonio.apolinario@ufba.br (A. Apolinário).

https://doi.org/10.1016/j.cag.2017.11.006
0097-8493/

Please cite this article as: M. Macedo, A. Apolinário, Improved anti-aliasing for Euclidean distance transform shadow mapping, Computers & Graphics (2017), https://doi.org/10.1016/j.cag.2017.11.006

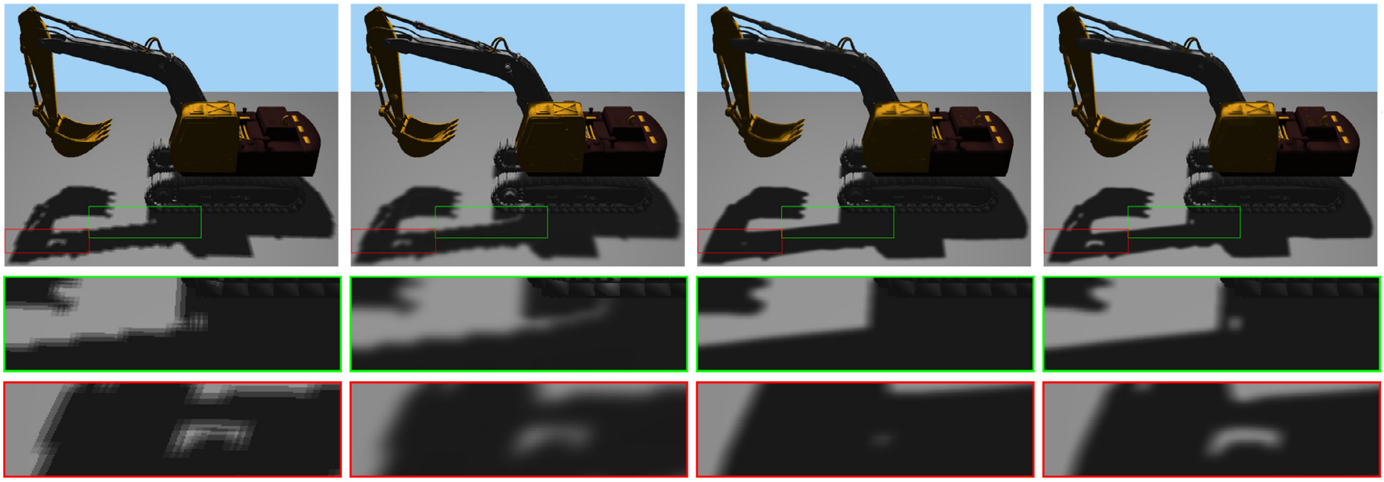

(a) PCF (b) VSM (c) EDTSM (d) EDTSM*

**Fig. 1.** Fixed-size penumbra produced by different techniques. For a low-order filter size, shadow map filtering techniques, such as PCF, generate shadows with aliasing and banding artifacts (a). Shadow map pre-filtering techniques, such as VSM, are prone to light leaking artifacts (green closeup in (b)). EDTSM suffers from shadow overestimation artifacts (c). The proposed approach (here named EDTSM*) is able to minimize those artifacts efficiently (d). Images were generated for the Excavator model using a $512^2$ shadow map resolution. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the shadow anti-aliasing, they can incorrectly merge parts of the shadow boundary that are originally disconnected [10] (Fig. 1(c)). Finally, filter size may directly affect the quality of the penumbra simulated. Small filter sizes may produce penumbra with blurred aliasing artifacts along the shadow boundary (Fig. 1(a)). On the other hand, large filter sizes may suppress fine details of shadows into penumbra.

Recently, Euclidean distance transform shadow mapping (EDTSM) was introduced to solve most of the problems mentioned before [15]. To do so, the technique first computes anti-aliased hard shadows using revectorization-based shadow mapping (RBSM) [10]. Then, an exact normalized EDT is computed from anti-aliased hard shadows using parallel banding algorithm (PBA) [16], which runs on the GPU. Finally, to reduce skeleton artifacts generated by the EDT, a simple mean filter is applied over the shadow boundary. Indeed, EDTSM is able to simulate fixed-size penumbra with less aliasing, banding and leaking artifacts than previous work, while keeping high frame rates. However, by the use of RBSM as hard shadow anti-aliasing technique, EDTSM suffers from shadow overestimation artifacts, which decrease the realism of shadow rendering (Fig. 1(c)).

In this work, which is an invited extension of our Graphics Interface 2017 paper [15], our main contribution is the enhancement of the RBSM visibility function to solve the problem of shadow overestimation, as shown in Fig. 1(d). Doing so, we can improve not only the quality of the hard shadow anti-aliasing provided by RBSM, but also the quality of the fixed-size penumbra simulation generated by EDTSM, keeping the processing time with a marginal overhead (about 1% of additional cost).

## 2. Related work

In this section, we review relevant work related to the proposed solution. We mainly cover techniques which provide real-time fixed-size penumbra simulation. For a more complete review of existing shadow mapping techniques, we suggest the reader to see the following books [17,18].

Several strategies have already been proposed to solve the aliasing problem of shadow mapping by warping [19,20], partitioning [21,22], traversing [10,23] or incorporating additional geometric information into the shadow map [24–26]. Unfortunately, none of these strategies are able to simulate penumbra, focusing only on the anti-aliasing of hard shadows.

The most traditional algorithm for fixed-size penumbra simulation is the percentage-closer filtering (PCF) [9]. As an extension of shadow mapping, PCF takes the results of shadow tests performed over a filter region and averages them to determine the final shadow intensity. By filtering the shadow test results, rather than the shadow map itself, PCF is not prone to light leaking artifacts, but provides real-time performance, while keeping low memory consumption for penumbra simulation. However, PCF does not support texture pre-filtering, does not provide scalability in terms of filter size, and requires a high number of samples to solve banding artifacts.

To make the shadow filtering scalable, variance shadow mapping (VSM) [11] uses Chebyshev's inequality, depth and squared depth stored in the shadow map to determine the shadow intensity of a surface point by means of a probability of whether the point is in shadow. VSM supports shadow map pre-filtering and is scalable for the filter size, but generates light leaking artifacts in shadows.

To reduce the light leaking artifacts of VSM, convolution shadow mapping (CSM) [27] uses Fourier series to approximate and linearize the shadow test. In CSM, the shadow map is converted into filtered basis textures that are used to determine the final shadow intensity as a weighted sum of basis functions stored in basis textures. CSM supports pre-filtering and reduces light leaking artifacts as compared to VSM, at the cost of more memory consumption and processing time than VSM.

To minimize the processing time required by CSM, exponential shadow mapping (ESM) [12,13] approximates the shadow test by an exponential function, rather than Fourier series. ESM stores exponent-transformed depth values into the shadow map, which are later used for penumbra simulation. ESM is faster and requires less memory footprint than CSM, while generating visual results similar to the ones obtained with VSM.

To improve the visual quality of both VSM and ESM, exponential variance shadow mapping (EVSM) [28] merges both ESM and VSM theories to produce high-quality fixed-size penumbra simulation. In EVSM, light leaking only occurs at places where both ESM and VSM techniques generate such an artifact.