Computers & Graphics 37 (2013) 214–224

Contents lists available at SciVerse ScienceDirect

Computers & Graphics

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

Technical Section

Smooth shadow boundaries with exponentially warped Gaussian filtering

Jesús Gumbau [a,*], Mateu Sbert [b], László Szirmay-Kalos [c], Miguel Chover [a], Carlos González [a]

[a] *Universitat Jaume I, Castellón, Spain*
[b] *Universitat de Girona, Girona, Spain*
[c] *BME, Budapest, Hungary*

ARTICLE INFO

Article history:
Received 30 March 2012
Received in revised form
13 December 2012
Accepted 14 December 2012
Available online 25 January 2013

Keywords:
Shadows
Real-time
Gaussian
Filtering
Exponential

ABSTRACT

Shadow mapping is widely used in computer graphics for efficiently rendering shadows in real-time applications. Shadow maps cannot be filtered as regular textures, thus their limited resolution can cause severe shadow map discretization artifacts in the rendered images. To solve this problem, several techniques have been proposed, including *variance shadow maps* (VSM) and *exponential shadow maps* (ESM). However, these techniques introduce different kinds of “light leaking” artifacts, which are clearly visible in moderately complex scenes. In this paper we propose a new statistical filtering method that approximates the cumulative distribution function (CDF) of depth values by a Gaussian CDF instead of bounding it with Chebyshev Inequality. This approximation significantly reduces “light leaks” and has similar performance and storage requirements compared to the original variance shadow map method. We also show that the combination of this technique with an exponential warp allows us to further reduce the remaining shadowing artifacts from the rendered image.

1. Introduction

Shadow mapping is a popular and efficient technique for solving the shadowing problem in interactive applications. A shadow map approximately represents the light occluding geometry as depth or distance values sampled on a two-dimensional grid. This sampled depth function is then used to reconstruct the distance between the light source and a point to be shaded in order to decide whether or not the point is behind the occluding geometry, i.e. it is in shadow.

The shadow test can also be imagined as a step like *visibility function* $v(z_r) = \epsilon(z_o - z_r)$ that is 0 if the occluder distance z_o from the light source to the occluder surface is smaller than receiver distance z_r measured between the light source and the point to be shaded, and 1 otherwise. However, the light occluder distance is known only in the centers of the shadow map texels, which leads to shadow map sampling artifacts during the visualization process. In order to avoid this issue, the occluding surface (or the visibility function) must be reconstructed and turned into a continuous function. However, as the geometry of the occluder 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 phenomenon of shadow aliasing.

Linear signal theory has a solution for the aliasing problem, which is based on the usage of low-pass filters in order to eliminate high frequencies that are above the Nyquist limit. However, shadow mapping is a non-linear operation as it contains a comparison operation represented by the *step function*. As a consequence, filtering the depth values directly before performing the comparison operation would result in averaged filtered depth values, which would lead to incorrectly calculated shadow boundaries. The problem of the non-linear step function is solved by simply delaying the filtering operation after performing the comparison, which is the basic idea of *percentage closer filtering* (PCF). This approach, which is able to present artifact-free correct anti-aliased shadows, has two major drawbacks. First, the filtering can only be performed when the distance of the shaded point is available, which requires it to be executed separately for every shaded point. Second, many shadow map texture fetch operations must be performed in order to achieve smooth anti-aliased shadow boundaries.

The final goal for solving this problem would be the ability of merging the filtering operations that are executed separately for

---

[*] Corresponding author at: Universitat Jaume I, Castellón, Spain. Tel.: +34 964729062.
E-mail address: jgumbau@uji.es (J. Gumbau).

0097-8493/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.cag.2012.12.004

each shaded point by PCF into a single texture filtering step. We wish to perform filtering as a pre-process before the rendering stage. This can be solved by storing additional information into the shadow maps so that the distribution of depths can be effectively approximated. The information stored in the shadow map must be GPU-friendly, which means that the visibility function can be evaluated from the filtered data with texture fetching modes supported by the graphics hardware, which are restricted to the linear interpolation operation.

Existing techniques aiming at pre-filtering the shadow map can be classified into two main categories, those that transform the problem into a domain where linear signal processing becomes feasible, and those that use non-linear filtering operations based on statistical analysis. Both families of methods allow for efficiently filtering the shadow map on the graphics hardware, but they also introduce artifacts commonly known as “light leaking”, which incorrectly alter the shadow intensity.

The method presented in this paper, which is an extended version of [1], belongs to the category of statistical filtering. It is based on approximating 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 Gaussian CDF instead of bounding it with the Chebyshev Inequality (as happens in VSM) or instead of approximating the step function with an exponential function (as in ESM). The two moments of the depth’s distribution are used to construct a Gaussian CDF. This approach is capable of highly reducing the “light leaking” artifacts present in existing techniques, or even eliminating it for moderately complex scenes, with neither penalty of performance nor increased storage cost. In addition to [1], we present an extension to our technique that builds upon the pure Gaussian approach and allows us to further minimize the appearance “light leaking” artifacts from the rendered image.

2. Related work

Shadow mapping [2] was introduced by Lance Williams in 1978 as a fast and efficient method for computing shadows in synthetic images. This method first captures the distances of the shadow caster points from the viewpoint of the light source and stores them into a depth map. Then, the scene is rendered from the camera. In the second step, the point visible in the given pixel is transformed to light space and its distance from the light source is compared to the value stored in the shadow map in its direction. Using this method, it is possible to decide whether the point to be shaded is farther from the light source than the nearest shadow caster. Although this very efficient method scales well on complex scenes since the shadow map complexity is independent of the number of objects, it produces discretization artifacts due to the process of sampling the depth information stored in the shadow map.

Shadow map discretization artifacts can be reduced by two orthogonal approaches: projection optimization and shadow filtering [3]. Projection optimization deals with how shadow caster objects are projected onto the plane of the shadow map in order to optimally allocate texture space for important parts of the scene [4–6]. Shadow filtering considers how data associated with the depth information are filtered in order to reduce aliasing. An important difference between them is that projection optimization is used during shadow map creation, while shadow filtering is executed when rendering the scene deciding how to interpret the data. These approaches are orthogonal and complementary as they can be used together to provide even more efficient anti-aliasing. Our proposed method can be classified as a shadow map filtering approach.

There are several works in the literature for dealing with the discretization problem of shadow maps, including *percentage-closer filtering* [7], which is one of the first methods introduced to alleviate this problem. This method is able to remove aliasing caused by shadow maps by applying a filter that averages the outcomes of depth comparisons against the shadow map. Percentage-closer filtering can be classified as a *post-filtering* approach since it applies the averaging after the non-linear depth test. One of the main drawbacks of this method is that the filtering operation can be executed only when the distance of the shaded point is available, so it should be repeated for every shaded point. Thus, the computational cost of percentage closer filtering becomes prohibitive when filters of large support are used. Unfortunately, relatively large kernels are needed in order to produce smooth anti-aliased shadows. To attack this performance limitation, we should move the filtering operation before the depth comparison and execute it once and globally for all shaded points. Such *pre-filtering* methods belong to two main branches, to those that apply depth transformation [8,9], or to those that are based on statistical analysis [10,11].

One of the most important representative of statistics-based shadow filtering is the method of *variance shadow maps* (VSM) [10]. VSM is based on using the Chebyshev’s Inequality for approximating an upper bound of the light visibility test. For each texel, VSM stores the depth and the squared depth of the shadow casters. These values can be filtered just like regular color textures resulting in the first two moments M_1 and M_2 of the depth values over the shadow filter region. When a point is shaded, the one-tailed verions of the Chebyshev’s Inequality allows us to upper bound the probability that depth z_0 of distribution with mean $\tilde{z}_o$ and variance σ^2 is greater than the receiver depth z_r

$$P(z_o \geq z_r) \leq \frac{\sigma^2}{\sigma^2 + (z_r - \tilde{z}_o)^2},$$

where $\tilde{z}_o = M_1$ is the average of the depth values and $\sigma^2 = M_2 - M_1^2$ is their variance. If receiver depth z_r is greater than the *mean depth* $\tilde{z}_o$, then the variance shadow map method approximates the visibility function by this upper bound for probability $P(z_o \geq z_r)$

$$\nu_{\text{VSM}}(z_r) = \frac{\sigma^2}{\sigma^2 + (z_r - \tilde{z}_o)^2}.$$

If the receiver depth is smaller than the mean depth, we assume that the surface is fully lit and thus $\nu(z_r) = 1$. VSM is an efficient and hardware friendly method. Its performance scales well with the screen resolution. Note that using an upperbound for the probability, the VSM may significantly overestimate the visibility function, making shadows lighter than they should be. This is noticeable as *light leaks* when shadow casters overlap from the light’s point of view. A typical problem occurs on parts of objects that are completely occluded but some amount of light still leaks inside the shadows (see Fig. 1).

Actually, light leaking artifacts show up because the first two moments of the depth do not provide enough information to disambiguate all the possible cases correctly. However, as stated in [12], storing more moments in the shadow map to obtain a sharper upperbound for the visibility function would not solve the problem, because higher-order moments are numerically unstable.

The *layered variance shadow maps* (LVSM) [12] method is an evolution of VSM developed to solve the “light leaking” problem. LVSM divides the light’s depth space into multiple layers, which allows for a better filtering of the two channels of the VSM. Using this technique we can obtain different upper bounds for $P(z_o \geq z_r)$, some tighter than others. When rendering the shadows, multiple