Computers & Graphics 38 (2014) 374–381

Contents lists available at ScienceDirect

Computers & Graphics

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

Special Section on CAD/Graphics 2013

Parallel and adaptive visibility sampling for rendering dynamic scenes with spatially varying reflectance

Rui Wang [a], Minghao Pan [a], Xiang Han [a], Weifeng Chen [b], Hujun Bao [a,*]

[a] State Key Lab. of CAD&CG, Zhejiang University, China
[b] School of Information, Zhejiang University of Finance & Economics, China

ARTICLE INFO

Article history:
Received 3 August 2013
Received in revised form
27 October 2013
Accepted 28 October 2013
Available online 15 November 2013

Keywords:
Adaptive visibility sampling
Environmental lighting
Dynamic scenes

ABSTRACT

Fast rendering of dynamic scenes with natural illumination, all-frequency shadows and spatially varying reflections is important but challenging. One main difficulty brought by moving objects is that the runtime visibility update of dynamic occlusion is usually time-consuming and slow. In this paper, we present a new visibility sampling technique and show that efficient all-frequency rendering of dynamic scenes can be achieved by sampling visibility of dynamic objects in an adaptive and parallel way. First, we propose a two-level adaptive sampling scheme to distribute sample points spatially and compute visibility maps angularly on each sample point. Then, we present a parallel hemispherical distance transform to convert these visibility maps into spherical signed distance fields. Finally, using such a distance-based visibility representation, we integrate our visibility sampling algorithm in the all-frequency rendering framework for scenes with spatially varying BRDFs. With an entire GPU-based implementation, our algorithm enables interactive all-frequency rendering of moderate dynamic scenes with environment lighting and spatially varying reflectance.

1. Introduction

Natural illumination, complex shadows and detailed reflections are all important in realistic image synthesis, but usually require high rendering costs. Precomputed radiance transfer (PRT) [1,2], or precomputation-based rendering, has shown that such rendering costs can be greatly reduced by precomputing the light transport in static scenes. Various basis functions, such as spherical harmonics (SH) basis [1], wavelet basis [2], polynomials [3], and spherical radial basis functions [4,5], have been proposed to approximate complex lighting, all-frequency visibility and spatially varying BRDFs. However, the requirement of static scenes has so far excluded a lot of PRT methods from many important applications with dynamic objects.

One main challenge brought by dynamic scenes is that the runtime visibility update of dynamic occlusion is usually time-consuming. Some work has been done to extend the PRT framework to dynamic scenes [6–9]. However, those methods more or less have some restrictions, such as handling only movements of rigid objects, rendering with low-frequency shadows, or requiring preprocess on scene's geometry. These restrictions make these methods unable to truly support all-frequency rendering of fully dynamic scenes.

In this paper, we present a new visibility sampling technique and show that efficient all-frequency rendering of dynamic scenes can be achieved by sampling visibility of dynamic objects in an adaptive and parallel way. Based on the observation in [10] that the light transportation in a scene is local and low rank, we first propose a two-level adaptive runtime visibility sampling scheme to distribute sample points spatially, and compute visibility maps angularly on these sample points. Then, we convert these sampled visibility maps into hemispherical signed distance fields and interpolate between them for shading points. A parallel hemispherical distance transform algorithm is presented to make this conversion fast. Finally, using such a distance-based visibility representation, we integrate our visibility sampling in the all-frequency rendering framework for scenes with spatially varying BRDFs. With an entire GPU-based implementation, our algorithm enables interactive all-frequency rendering of moderate dynamic scenes.

2. Related work

Our paper focuses on all-frequency rendering of dynamic scenes. Our work is based on a PRT rendering framework and extends the precomputation of visibility to a runtime adaptive visibility sampling. Thus, we start with a brief review of the PRT technique and then introduce some techniques that are most relevant to our work.

* Corresponding author.
E-mail addresses: rwang@cad.zju.edu.cn (R. Wang), bao@cad.zju.edu.cn (H. Bao).

0097-8493/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.cag.2013.10.036

Precomputation-based rendering. PRT technique [1,2] enables real-time rendering under natural lighting, complex shadowing and detailed materials. Various basis functions, such as spherical harmonics (SH) basis [1], wavelet basis [2], polynomials [3], and spherical radial basis functions [4], have been proposed to approximate lighting, visibility and BRDFs. A comprehensive survey of current PRT techniques can be found in [11]. Our work is based on the all-frequency rendering framework proposed in [5] and then extended in [12]. By approximating lighting and BRDFs into spherical Gaussians and representing visibility into signed spherical distance fields, real-time rendering is achieved for scenes with spatially varying BRDFs. In our method, we replace the visibility precomputation by a parallel and adaptive visibility sampling algorithm so as to enable fast rendering of dynamic scenes.

The light transport response of vertices in a scene is local and of low rank [13]. It can be highly compressed by clustered principal component analysis (CPCA) [14] or wavelets [2]. Huang and Ramamoorthi [10] proposed a precomputing method to sample the light transport matrix adaptively and sparsely. Although they demonstrated faster precomputation time than previous methods, their sampling and reconstruction method is still too slow for dynamic scenes. Our work is inspired by their work but focuses on runtime update of visibility.

Rendering dynamic scenes under complex lighting. The PRT technique has been extended to render dynamic scenes. Zhou et al. [6] proposed shadow fields to generate soft shadows in scenes with moveable objects. Sun et al. [15] generalized the idea to wavelet product to render scenes with dynamic glossy objects. However, these methods only handle the movement of rigid objects. Zonal spherical harmonics [7] with spheres approximation [8] enable rendering of deformable objects, but is limited to SH-represented low-frequency shadows. Iwasaki et al. [9] presented an all-frequency rendering method for dynamic scenes. But, their method still relies on a precomputed spheres approximation, which is unable to handle arbitrary movement of objects. Our adaptive visibility sampling method does not have such restrictions on scenes.

With the rapid development of computational power of GPU, some methods using GPU-based point lights rendering techniques have been proposed for dynamic scenes. Annen et al. [16] proposed a real-time method for all-frequency shadows in dynamic scenes based on convolution shadow maps. Ritschel et al. [17] presented imperfect shadow maps to compute indirect illuminations from a points approximation of the scene. Besides their work, there is a large body of recent work on GPU-based global illumination, such as [18–20]. These techniques use GPU to achieve fast illumination, thus support dynamic scenes. However, similar to most real-time point-based shadow map techniques, these methods do not support integrating the BRDF across the area light source domain, which limits its usage in all-frequency rendering of non-diffuse BRDFs, especially spatially varying BRDFs.

Euclidean distance transformation. In this paper, we present a GPU-based parallel hemispherical distance transformation algorithm to generate spherical signed distance fields for visibility interpolation. Our work is inspired by the Euclidean distance transform (EDT) algorithms on 2D images. EDT is an important method in computer vision and geometry processing. A comprehensive survey on 2D image can be found in [21]. Generally, fast EDT algorithms are designed in a dimensionality reduction manner that it first computes in one dimension, e.g. for each row, and then computes in the second dimension, e.g. for each column. Such a dimensionality reduction idea was first proposed in [22] and then improved in [23–26]. Recently, Cao et al. [26] have further extended the idea to GPUs and achieved fast 2D EDTs by dividing data into different bands. However, due to a non-uniform spherical parameterization, it is non-intuitive to directly apply these 2D EDT algorithms to the hemispherical domain. Michikawa and Suzuki [27] proposed a sequential algorithm for the spherical EDT but it is too slow for our application. In this paper, we introduce a new parallel distance transformation algorithm, which is designed specifically for the distance transform on the hemisphere.

3. Overview

Our visibility sampling algorithm is based on the all-frequency rendering framework for spatially varying BRDFs, which was first proposed in [5] and then extended in [12] for better visibility approximations. In this section, we first briefly describe the all-frequency rendering framework and then introduce our visibility sampling algorithm.

3.1. Rendering framework

The direct lighting at point $\mathbf{x}$ in view direction ω_o can be computed by the following integral:

$$L_o(\mathbf{x},\omega_o) = \int_{\Omega_{2\pi}} L_i(\mathbf{x},\omega) f(\mathbf{x},\omega,\omega_o) V(\mathbf{x},\omega)(\mathbf{n}\cdot\omega)\, d\omega \tag{1}$$

where L_o is the outgoing radiance, L_i is the incident lighting, $f(\mathbf{x},\omega,\omega_o)$ is the BRDF, $V(\mathbf{x},\omega)$ is the visibility function, and $(\mathbf{n}\cdot\omega)$ is the cosine term. To handle the all-frequency rendering with spatially varying reflectance, Wang et al. [5] approximated the lighting and BRDFs by spherical Gaussians, and represented the visibility into spherical signed distance functions. In this paper, we follow the formation, where the incident lighting and BRDFs are computed as

$$\begin{aligned} L_i(\omega) &\approx \sum_l G_l(\omega;\mathbf{p}_l,\lambda_l,\mu_l), \\ f(\omega,\omega_o) &\approx \sum_m G_m(\omega;\mathbf{p}_m^{\omega_o},\lambda_m^{\omega_o},\mu_m^{\omega_o}). \end{aligned} \tag{2}$$

in which a spherical Gaussian G has the following form:

$$G(\omega;\mathbf{p},\lambda,\mu) = \mu\cdot e^{\lambda(\mathbf{p}\cdot\omega-1)} \tag{3}$$

where $\mathbf{p}$ is the lobe direction, λ is the lobe sharpness and μ is the lobe amplitude. The spherical signed distance function (SSDF) [28] is used as an intermediate visibility representation, where for a direction $\mathbf{v}$, it stores a signed angular distance to the closest visibility boundary. The sign of the function encodes whether the direction $\mathbf{v}$ is occluded or not. It is defined as

$$D(\mathbf{v}) = \begin{cases} +\min\limits_{V(\mathbf{t})=0} \arccos(\mathbf{t}\cdot\mathbf{v}) & \text{if } V(\mathbf{v})=1 \\ -\min\limits_{V(\mathbf{t})=1} \arccos(\mathbf{t}\cdot\mathbf{v}) & \text{if } V(\mathbf{v})=0 \end{cases} \tag{4}$$

where $\mathbf{v}$ and $\mathbf{t}$ are vectors on a unit sphere, $V(\cdot)$ is the visibility function that returns 0 when being occluded and 1 otherwise. To better represent the visibility, in [12], the SSDF, $D(\mathbf{v})$, is extended to adaptively sampled distance field [29], where the domain of the SSDF is adaptively divided into a set of cells, and visibility boundaries are approximated by iso-lines in these cells, Fig. 2(c). In this way, in each cell, the rendering integral is

$$L_o^{ij} = \int_{\varphi_i}^{\varphi_{i+1}} \int_{\theta_j}^{k\varphi+b} \sum_{l,m} G_{l,m}(\omega)(\mathbf{n}\cdot\omega)\ \sin\theta\, d\theta\, d\varphi. \tag{5}$$

where (φ,θ) is the 2D parameterizations of a sphere, Fig. 2(a), $\varphi_i,\varphi_{i+1},\theta_j,\theta_{j+1}$ are upper and lower bounds of a cell, and $k\varphi+b$ is the iso-line function in the cell, Fig. 2(c). The final outgoing radiance of Eq. (1) can be computed by summing radiance from all cells. In this rendering framework, binary visibility maps are