Multi-Resolution Depth-of-Field Rendering

Yuna Jeong Kangtae Kim Sungkil Lee
Sungkyunkwan University*

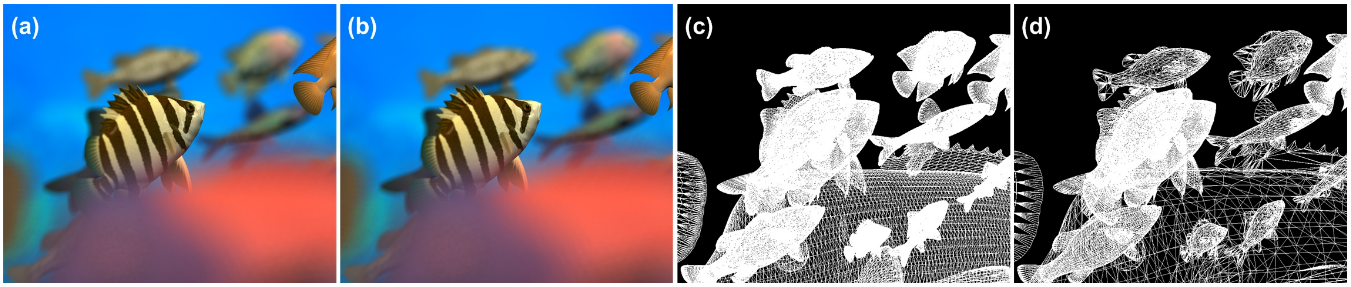

Figure 1: *Example images rendered without LOD (a) and with our LOD management (b) for depth-of-field effects. (c) and (d) are the wireframe representations of (a) and (b), respectively. The output qualities of our LOD management are virtually indistinguishable from those of reference renderings, while improving rendering performance up to 3.5 times.*

1 Introduction

Depth of field (DOF) refers to a distance range in which photographic capture yields acceptable sharp imagery. The DOF effect is crucial in improving the perceptual realism of synthetic images and drawing user's attention. In graphics, object-based approaches served as reference [Cook et al. 1984; Haeberli and Akeley 1990]. However, their low performance, resulting from the repeated rendering for different lens samples, has impeded their real-time use.

Recent advances in real-time DOF rendering have been driven to image-based approaches. Among them, [Lee et al. 2010] attained noticeable improvements in quality using layer decomposition and image-based ray tracing. But, these methods may potentially suffer from artifacts with an insufficient number of depth layers. Furthermore, geometric details projected to less than a pixel cannot be represented, leading to the lack of antialiasing. In contrast, object-based approaches naturally achieve fine geometric details and antialiasing.

This paper presents a new DOF rendering algorithm, based on the distributed rasterization [Haeberli and Akeley 1990] and LOD management. Our solution allows us to maintain the benefit of the object-based approach without spatiotemporal quality loss, while achieving real-time performance. A key idea is that geometric degradation of models is not perceived when they are blurred. Hence, lower details can be used for blurred models, greatly improving performance. Another challenge here is avoiding temporal popping artifacts resulting from transitions between adjacent discrete levels. To avoid this problem, we propose a novel blending method for LOD.

2 Our Approach

We briefly describe the pipeline of our framework. Given 3D models, we generate their discrete LOD representations, applying successive simplification at a step of half size. For every single frame of rendering, we first compute the LOD for each object using its bounding box—this level is not an integer but a real number. Then, the rasterization of each model is repeated for different lens samples and accumulated into a single output buffer.

The LOD of an object is proportional to the degree of blur and the reciprocal of its projection size. The degree of blur can be efficiently approximated using the size of circle of confusion at the center depth of the object's bounding box. Also, the bounding box gives us the rough size of the object's projection.

We use two different levels adjacent to the model's LOD, instead of a single integer level that obviously leads to the popping. The fractional part of the LOD elicits how many times each model needs to be rendered. The smooth change of the fraction along with animation enables to avoid temporal popping. Then, we need to map each lens sample index to one of the two models. This mapping is greatly leveraged by Halton sampling. The lens sample indices less than the same fraction of the sample number correspond to the lower-level model and the others the higher-level model. Meanwhile, the Halton sampling maintains the uniform distribution for each model.

Finally, we report performance comparison of the reference, our method, and [Lee et al. 2010] for the scene of 20 fishes (1.8 M faces); see Figure 1 for quality. An Intel i5 machine with GTX 560Ti was used at 1024×768. The speedup of our method reached up to 3.5 against the reference (Figure 2). Also, our method slightly outperformed [Lee et al. 2010] for 32 samples, while having reductions at more samples. This results from the constant cost of layer decomposition for [Lee et al. 2010]. In the future, we are planning to include sampling LOD and other methods for better performance.

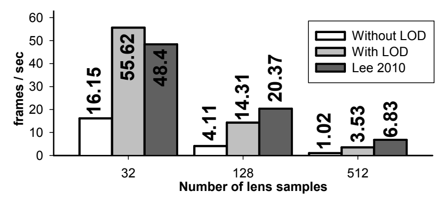

Figure 2: *The summary of performance measurements.*

Acknowledgements. This work was funded by the Basic Science Research Program through the National Research Foundation of Korea (2011-0014015).

References

COOK, R., PORTER, T., AND CARPENTER, L. 1984. Distributed ray tracing. In *Proc. ACM SIGGRAPH*, vol. 18, ACM, 137–145.

HAEBERLI, P., AND AKELEY, K. 1990. The accumulation buffer: Hardware support for high-quality rendering. In *Proc. ACM SIGGRAPH*, vol. 24, ACM, 309–318.

LEE, S., EISEMANN, E., AND SEIDEL, H.-P. 2010. Real-Time Lens Blur Effects and Focus Control. *ACM Trans. Graphics (Proc. ACM SIGGRAPH'10) 29*, 4, 65:1–7.

*e-mail: {jeongyuna, sonata, sungkil}@skku.edu

SIGGRAPH 2012, Los Angeles, California, August 5 – 9, 2012.
ISBN 978-1-4503-1435-0/12/0008