Adaptive Matrix Column Sampling and Completion for Rendering Participating Media Supplementary Document

Yuchi Huo Rui Wang* Tianlei Hu Wei Hua Hujun Bao*

State Key Lab of CAD&CG, Zhejiang University

1 Discussion of the Work [Huo et al. 2015]

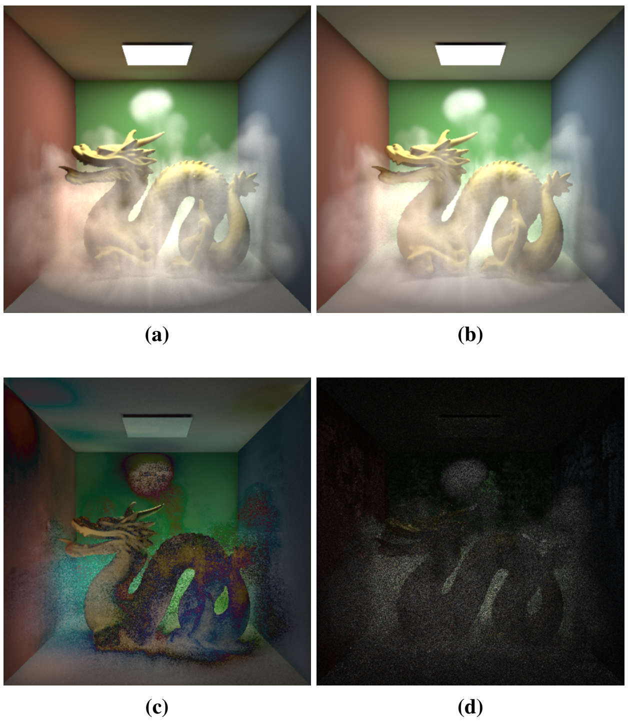

Figure 1: *(a) Image rendered by matrix separation [Huo et al. 2015]. (b) Reference image. (c) Difference image between matrix separation and the reference. (d) Difference image between our method and the reference.*

Huo et al. [2015] proposed a matrix sampling and recovery approach to accelerate the many-light rendering, of which our method is inspired by the sparse sampling and completion idea. In their method, the shading values are densely computed for all matrix elements, and only the visibility terms of the matrix are sparsely sampled and recovered. In particular, three predictors are proposed to predict unknown visibilities from a few visibility samples, and a matrix separation technique is used to separate prediction errors and recover the visibility matrix.

However, such a specific matrix recovery technique used in their method does not consider the scattering light transport, thereby cannot be directly applied to render participating media for two reasons. First, these predictors can only predict binary visibilities, whereas rendering participating media requires both visibility and transmittance. Second, the errors of transmittance are not sparse, therefore they cannot be removed from the matrix separation technique. Figure 1 shows the result of applying their method in rendering participating media. As shown in the difference image, it has serious bias because of the inaccurate recovery of media transmittance. Compared with the matrix separation, the matrix completion used in our method is a more general technique that can recover low-rank lighting matrices of scenes with participating media.

*Corresponding authors: Rui Wang (rwang@cad.zju.edu.cn), Hujun Bao (bao@cad.zju.edu.cn)

2 The Number of Sampled Columns

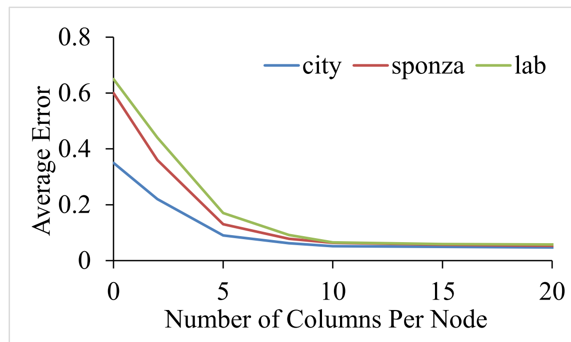

Figure 2: *Average errors rendered with different number of columns per node.*

We test different numbers of sampled column per node in our scenes. As shown in Figure 2, the average errors of the scenes decreases as the number of sampled columns increases, and converges finally. Because the adaptive sampling strategy automatically refines the cut on tree nodes to compute illuminations of nodes, we do not need to sample many columns per node. If the error is too large at one node, the algorithm will adaptively split the node into children to further sample it. Therefore, in our practice, we pick 10 columns per node as default.

References

HUO, Y., WANG, R., JIN, S., LIU, X., AND BAO, H. 2015. A matrix sampling-and-recovery approach for many-lights rendering. *ACM Trans. Graph. 34*, 6 (Oct.), 210:1–210:12.