

# Real-time Selective Rendering

## CSAW-2012 Presentation Abstract

Steven Galea

*Department of Computer Science*

*University of Malta*

*Msida, Malta*

*sgal0012@um.edu.mt*

Abstract - Traditional physically-based renderers can produce highly realistic imagery; however, suffer from lengthy execution times, which make them impractical for use in interactive applications. Selective rendering exploits limitations in the human visual system to render images that are perceptually similar to high-fidelity renderings in a fraction of the time. This paper outlines current research being carried out by the author to tackle this problem, using a combination of ray-tracing acceleration techniques, GPU-based processing, and selective rendering methods. The research will also seek to confirm results published in literature, which indicate that users fail to notice any quality degradation between high-fidelity imagery and a corresponding selective rendering.

## 1 Introduction

At its essence, computer graphics is concerned with synthesizing a pictorial representation of objects and environments from their computer based models [1], through a process known as rendering.

Traditional rendering methods make use of rasterization to project 3D models to a 2D image plane, together with an array of specialized methods to approximate real-world phenomena like shadows, reflections and indirect lighting [2]. These techniques have evolved into mature API's (e.g. OpenGL, DirectX) and can achieve real-time performance at the cost of limited realism and increased application complexity.

In contrast, physically-based renderers seek to replicate the actual properties of lighting and material. For instance, ray-tracing simulates the path taken by light through a virtual environment, and can generate images that not only look realistic, but are also physically plausible. However, rendering in this manner is very computationally intensive, and until recently, physically-based rendering was largely limited to off-line rendering [3].

This paper outlines the authors research in selective rendering with the aim of generating high-fidelity imagery at interactive rates and confirming experimental results published in the literature. Section 2 gives a short description of the problem to be tackled. Section 3 reviews the relevant literature. The proposed solution is discussed in Section 4. Section 5 provides an overview of the current project as well as possible directions for further research. Finally, Section 6 concludes the paper.

## 2 Problem Description

According to Chalmers and Ferko [4], rendering high-fidelity images in real-time has long been the "holy grail" of the computer graphics community. To achieve this, research on physically-based renderers has tried to tackle the problem from different points of view. On one front, researchers have focused on finding effective data structures that reduce the number of operations carried out for each ray of light [3]. Other works have focused on implementing massively-parallel ray tracers on GPUs [5], and distributing rendering load over a computing cluster [6] or supercomputer [7]. However, simulating real-world light interaction of complex scenes at interactive rates on a desktop PC still remains out of reach, even using modern GPUs [8].

More recent attempts have adopted principles from psychology and perception to target computation to parts of the scene that matter most to a human observer [9]. Selective rendering allows for substantial reductions in computation while creating an image that is perceptually equivalent to rendering a high-fidelity image of the whole scene.

## 3 Literature Review

Physically-based rendering algorithms, such as ray-tracing, generate high-fidelity images by simulating the physics of light and its interaction with surface materials. In particular, global illumination effects, such as surface inter-reflections and caustics, play a major role in enhancing the realism of rendered images [10]. However, the huge amount of computation required to produce physically-correct images prevents these methods from being utilized in interactive applications on desktop computers [4]. Several techniques have been developed to improve ray tracing performance:

Acceleration data structures, such as kd-trees and BVH trees, reduce geometric complexity by testing for ray-object intersections against groups of geometric primitives [11].

Instant global illumination (IGI) pre computes the distribution of lighting in the scene by following light paths from light sources in the scene and depositing virtual point lights (VPL) at each vertex [12]. Rendering indirect light is achieved by casting shadow rays from the shaded point to each VPL to determine visibility.

Interleaved sampling exploits the slowly changing nature of the indirect lighting function to include the contribution from a given number of VPLs in less time [13]. This is achieved by modifying the IGI pre-processing stage to generate  $k$  subsets of VPLs. The image is then partitioned into tiles of  $n$  by  $m$  ( $n*m=k$ , usually  $n=m=3$  or  $m=n=5$ ) pixels, with each pixel sampling indirect light from the relative VPL subset. Finally, a discontinuity buffer is used to calculate the contribution of all VPLs at each pixel by interpolating values from neighboring pixels, provided that they do not cross geometric discontinuities [12].

Selective rendering seeks to produce images that are perceived to be of an equivalent quality to a physically-correct image, but using a fraction of the computation. This is possible since the human visual system does not process an image in raster-like fashion, but rapidly directs attention to salient features in a scene, such as sudden movement or brightly coloured objects [14]. This knowledge can be used to construct saliency maps, i.e. images highlighting the most perceptually interesting parts of the scene, which are then used to direct computation appropriately.

## 4 Proposed Solution

Saliency maps generated using image processing techniques are usually expensive to compute and do not give the region of interest where the user is currently looking at, but only where the user is likely to direct attention to. This MSc will look into various methods of achieving real-time performance for rendering using selective methods. This includes the possibility of using an eye tracker to get an accurate measurement of the user's current region of interest, with all the associated computation being carried out by the eye tracking device. This should significantly reduce the time spent on computing the saliency map, making more resources available for rendering. To our knowledge, the use of an eye tracking device for real-time selective rendering has never been explored before.

Additionally, this project will also seek to confirm results published in literature which indicate that users fail to notice any differences in quality between a high-fidelity rendering and a selective rendering of the same scene. Previous psychophysical experiments have used pre-rendered images and animations for the evaluation. If the renderer is successful at achieving real-time performance, the project can also look into carrying out a psychophysical experiment with interactive imagery, where the participants are able to freely navigate the scene.

## 5 Project Status

A prototype to demonstrate proof of concept is being developed as an extension to the GPU-based ray-tracing framework NVIDIA OptiX [15]. The framework provides a high-performance ray-tracing core comprising a small set of building blocks that are used in most ray-tracing algorithms.

Several basic components such as model-loading and rendering of direct lighting have already been developed. Rendering of global illumination is achieved by a GPU-based implementation of the IGI algorithm, together with interleaved sampling and discontinuity buffer to improve execution time.

Interleaved sampling provides good speedups at the cost of some degradation in image quality, especially near geometric discontinuities, where the averaging across nearby samples cannot be performed. The loss in image quality arises mainly from the fixed 3x3 or 5x5 tiling in interleaved sampling. Current research is investigating whether it is possible to obtain a better quality image by smartly positioning the samples across the image according to saliency. This is done by considering the full set of VPLs at each sample, but positioning the samples densely in salient regions (e.g. at pixel intervals), and more sparsely (e.g. sampling once per 3x3 or 5x5 tile) in less salient regions. Missing illumination values can then be interpolated from nearby samples.

Preliminary results show that this technique gives a good quality indirect lighting image. However, a more in depth analysis still needs to be carried out to compare the output of the proposed method with that of traditional interleaved sampling.

The saliency map currently being used is a GPU implementation of the method by Hou et al. [16]. Although fast, the method is quite basic and several enhancements have been proposed in literature to improve accuracy. Future work will look into substituting the current saliency map with an improved version, thus improving the quality of the indirect lighting image.

Once the selective rendering component is ready, the project may also look into using an eye tracker to provide the saliency information instead of the saliency map.

## 6 Conclusion

This paper started with a review of the existing literature, outlining the performance issues of rendering high-fidelity imagery at interactive rates. An overview of the current research was given, proposing a combination of ray-tracing acceleration techniques, GPU-based processing, and selective rendering methods as a way to tackle this problem. The current status of the project was discussed, together with preliminary results and directions for future research.

## References

- [1] J. D. Foley, A. v. Dam, S. K. Feiner and J. F. Hughes, *Computer Graphics: Principles and Practice in C (2nd Edition)*, Addison-Wesley Professional, 1995.
- [2] J. Schmittler, S. Woop, D. Wagner, W. J. Paul and P. Slusallek, "Realtime ray tracing of dynamic scenes on an FPGA chip," in *ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware*, New York, USA, 2004.
- [3] I. Wald, W. R. Mark, J. Gnther, S. Boulos, T. Ize, W. Hunt, S. G. Parker and P. Shirley, "State of the Art in Ray Tracing Animated Scenes," *Computer Graphics Forum*, vol. 28, no. 6, pp. 1691-1722, 2009.
- [4] A. Chalmers and A. Ferko, "Levels of Realism: From Virtual Reality to Real Virtuality," in *24th Spring Conference on Computer Graphics*, New York, USA, 2008.
- [5] J. Gunther, S. Popov, H.-P. Seidel and P. Slusallek, "Real-time Ray Tracing on GPU with BVH-based Packet Traversal," in *IEEE/Eurographics Symposium on Interactive Ray Tracing*, 2007.
- [6] I. Wald, P. Slusallek, C. Benthin and M. Wagner, "Interactive Distributed Ray Tracing of Highly Complex Models," in *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, London, UK, 2001.
- [7] I. Wald and P. Slusallek, "State of the Art in Interactive Ray Tracing," 2001.
- [8] A. Chalmers, K. Debattista, G. Mastoropoulou and L. P. d. Santos, "There-Reality: Selective Rendering in High Fidelity Virtual Environments," *The International Journal of Virtual Reality*, 2007.
- [9] D. Bartz, D. Cunningham, J. Fischer and C. Wallraven, "The Role of Perception for Computer Graphics," in *Eurographics State-of-the-Art-Reports*, 2008.
- [10] K. Myszkowski, T. Tawara, H. Akamine and H.-P. Seidel, "Perception-Guided Global Illumination Solution for Animation Rendering," in *In Proceedings of SIGGRAPH 2001*, Los Angeles, California, USA, 2001.
- [11] E. Reinhard, A. Chalmers and F. W. Jansen, "Overview of Parallel Photo-realistic Graphics," in *Eurographics '98 State of the Art Reports*, 1998.
- [12] I. Wald, T. Kollig, C. Benthin, A. Keller and P. Slusallek, "Interactive global illumination using fast ray tracing," in *EGRW '02 Proceedings of the 13th Eurographics workshop on Rendering*, 2002.
- [13] A. Keller and W. Heidrich, "Interleaved Sampling," in *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, London, UK, 2001.
- [14] S. Lee, G. J. Kim and S. Choi, "Real-time tracking of visually attended objects in interactive virtual environments," in *Proceedings of the 2007 ACM symposium on Virtual reality software and technology*, New York, USA, 2007.
- [15] NVIDIA, "NVIDIA OptiX ray tracing engine," NVIDIA, [Online]. Available: <http://developer.nvidia.com/optix>. [Accessed 1 May 2012].
- [16] X. Hou and L. Zhang, "Saliency Detection: A Spectral Residual Approach," in *IEEE Conference on Computer Vision and Pattern Recognition*, Minneapolis, Minnesota, USA, 2007.