Multipass Rendering in mental ray

Haarm-Pieter Duiker*
ESC Entertainment

Thomas Driemeyer*
mental images

Introduction

Rendering resources are by definition constrained and yet the task of rendering many objects at high quality is frequently encountered in the effects industry. Breaking up scenes into renderable pieces, or passes, is a common approach to rendering large scenes. To achieve the quantity of photo-real effects in frame that were required by the scripts of the Matrix Reloaded and the Matrix Revolutions, mental images and ESC Entertainment developed a system for rendering and automatically compositing many passes. The system we developed has the key advantage that it takes as its compositing primitive not the pixel but the sample.

The System

One difficulty with the standard, pixel-based, approach to depth compositing is that automatically assembling passes can often show artifacts because the depth and color values present in the pixels of each pass are the filtered average of possibly many samples taken from the rendered scene. Each pixel being a filtered average means that in an area of partial occlusion or transparency, meaningful values for transparency or depth may be obscured almost completely if not lost. This can make all but impossible the job of producing a composite that is the same as if all passes had been rendered at once. Taking the scene sample as the primitive for compositing operations means that unnecessary filtering operations can be cut out of the process, removing many standard depth-compositing artifacts.

The system as implemented in mental ray consists of three possible stages. The first stage includes the actual rendering of each pass and the creation of files to hold the original scene samples. The second, an optional stage, allows for arbitrary processing operations on individual samples files. In the third stage, the standard depth compositing operation is applied to the all passes at sample level before the resulting samples are filtered down for pixel resolution output.

The Multipass Rendering system afforded us two main benefits. First and foremost, the system allowed us to render scenes with large amounts of geometry as well as finely detailed lighting and shading calculations. We didn't have to sacrifice one or other, which given the demands of the movies, was not an option. Large amounts of low resolution geometry were used for shadowing purposes while rendering other pieces of high resolution geometry in full detail. The second benefit it afforded us was the use of a two-dimensional motion-blurring algorithm with a minimum of the traditional depth and transparency related artifacts inherent in such a system. Using multipass and two-dimensional motion blurring, we were able to achieve a two and a half dimensional motion blurring system with all the speed and parallel processing benefits of a multipass depth compositing system with a minimum of the artifacts that traditionally show up with either approach.

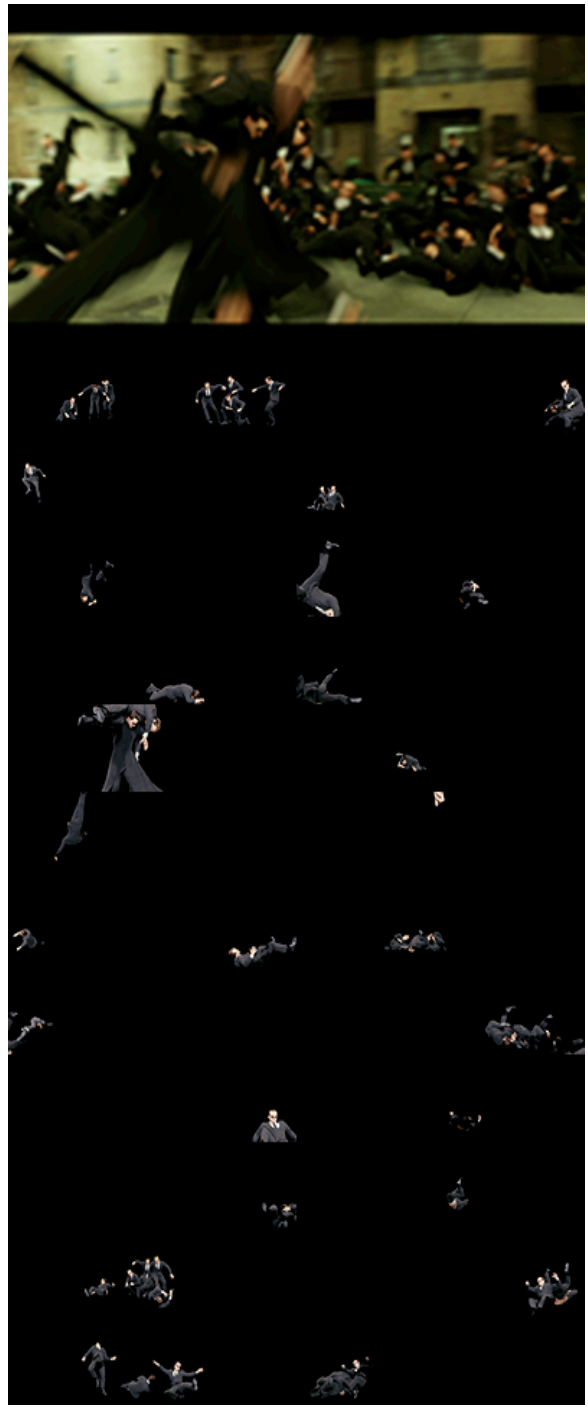

Figure 1: The top image is the full composited result of the combination of the all the passes below.

Conclusion

We have developed and implemented a system for rendering large scenes in multiple passes and compositing the resulting passes in a manner that closely approximates the original scene.

----

*e-mail: hp@escfx.com
thomas@mental.com

Acknowledgements

The team at mental images deserves credit for the implementation of this system. Dev Mannemela has been instrumental in completing a set of tools for the system. Ray Haleblian, Scott Liedtka, and Andrew Harris braved the dangers of using the system in production. Andy Lomas, Olivier Maury, and Scot Shinderman's algorithms helped, and the advice of John Schlag, George Borshukov, and Kim Libreri proved invaluable.