INSTITUTIONEN FÖR DATAVETENSKAP | LUNDS TEKNISKA HÖGSKOLA | PRESENTATIONSDAG 2015-03-20

EXAMENSARBETE Multi-coefficient Parallel Adaptive Wavelet Rendering
STUDENT Robin Somers
HANDLEDARE Michael Doggett (LTH)
EXAMINATOR Tomas Akenine-Möller (LTH)

Photorealistic rendering using signal processing techniques

POPULÄRVETENSKAPLIG SAMMANFATTNING **Robin Somers**

It is a notoriously slow process to create complex photorealistic images. However, by using the power of multiple cores in modern processors, one already efficient method has now been altered to become even faster.

Photorealistic rendering is used today for many different applications, such as for movies, art, advertising and games. It differs from real-time rendering, used for interactive graphics, in that it gives more physically accurate results. In other words it looks more real but it comes at the cost of time. In fact, one such image can take minutes, even hours, to render depending on the complexity of the scene it shows. Consider what this means for a movie containing many visual effects. Every second typically contains at least 24 still images. If each image takes roughly 20 minutes to render and there are 60 minutes containing visual effects in the movie, this means that the total rendering time will be close to 30 000 hours! In practice, to make the time reasonable, giant clusters of computers are used to render several images simultaneously. Nevertheless, any decrease in rendering time can have a profound effect, shaving off thousands of hours. It is no big surprise that much research has gone into doing just that.

The so-called Adaptive Wavelet Rendering method was proposed a couple of years ago and works relatively well. However, in the thesis it is shown how it can be improved by letting the computer perform multiple tasks in parallel rather than just one at a time. This is achieved by using multi-core processing where each processor core can work independently but still in collaboration with the others. The result is that the method works two to three times faster! Not only that, but by slightly modifying the method to work on several regions of an image at the same time it can become up to 30 % faster still! In fact, with the modification, the image below is rendered in only a quarter of the time compared to when using a single processor core. Not every scene gets such a substantial improvement but it never gets worse either. The Adaptive Wavelet Rendering method is just one among many but by parallelizing it, it becomes faster and better suited to todays highly parallel compute systems.

Adaptive Wavelet Rendering

The method uses the concept of wavelets, used in signal processing, to redistribute the pixel information in a not yet completed image. This is done multiple times where each creates a new layer of information. The trick is that these layers are constructed to easily show which parts of an image that need more work and which that are good as they are. Therefore areas with greater complexity can be identified and worked on to a greater extent whereas unnecessary work is avoided for simpler parts of a scene. By being able to target where work is needed much time is saved. The top image to the right shows a rendered image and the bottom image shows how the method has interpreted the complexity, brighter meaning more complex.

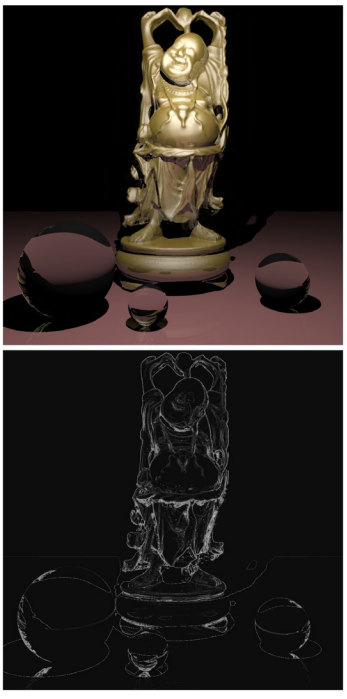

Rendering example