DEPARTMENT OF COMPUTER SCIENCE | LUNDS TEKNISKA HÖGSKOLA | PRESENTATION DAY 2016-06-03

MASTER'S THESIS Multi-layered G-Buffers for Real-Time Reflections
STUDENT Mattias Simonsson
SUPERVISOR Michael Doggett (LTH)
EXAMINER Flavius Gruian (LTH)

Screen space reflections with multiple layers

POPULAR SCIENCE SUMMARY **Mattias Simonsson**

Real-time reflections can be rendered quickly with decent quality in screen space. We investigate and evaluate the performance of a method that improves the quality of the reflections

Real-time graphics is all about approximating the look of the world in a limited amount of time. However, as computers grow more and more powerful, the amount of computations that can be done in the same amount of time increases.

Reproducing the behavior of reflective surfaces, like mirrors and water, is expensive but still possible to do thanks to the power of modern hardware. To render, or draw, reflections on a surface, we need to find out where the light that bounced on the surface into our eyes came from. Because of the complexity of scenes in many modern applications, this takes a lot of time to do accurately.

By reducing the complexity of the scene in some way, we can render the reflections faster. A way to do this is to render the reflections in screen space. This can be thought of as taking a photograph of the scene and then drawing reflections on top of this photograph later.

Reducing the amount of available information before rendering reflections improves performance but reduces accuracy of the reflections. As an example, if a player is standing in front of a mirror in a third-person game, the mirror will not be able to reflect the face of the player. The information is simply not available.

In my thesis I combine screen space reflections with a technique known as depth peeling. With depth peeling, the scene is rendered into multiple 2D images, or layers, instead of just one, with each layer containing unique information. This will improve the accuracy of the reflections but also reduce performance because the scene needs to be rendered multiple times. I evaluate and compare the performance cost relative to the increase in image quality in my thesis.

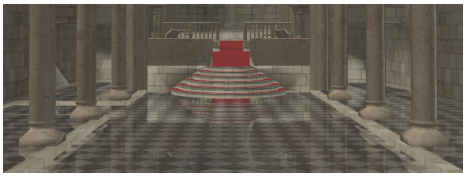

Figure 1: Example of screen space reflections

The results suggest that while using multiple layers does improve image quality, this improvement is in most cases too small to justify the very large performance cost of depth peeling. The results also suggest that if multiple layers are used, the amount of layers should be two or maybe three. In general, each successive layer adds less information than the previous.