

RealTime Rendering

When we're talking about computer generated images, whatever the individual images are, one of the main concerns is how much time will be needed for each image to be rendered. This is quite true in both animated movies and visual effects given that creating the realistic impression desired requires more and more computing time: photons, radiosity, caustics, and shadows, to name a few, multiplied by the complexity of the geometry in the scene.

When dealing with Video Games, or interactive experiences, this problem is all the more important. When the end user is doing something, be it with a joystick or with a motion tracking sensor, he needs to get immediate feedback - meaning that the image must be rendered in less than 20 miliseconds. As such, research in this field is focused on creating images, as appealing as the ones seen in animated films, rendered significantly faster.

Some of the latest research in Real Time rendering focuses on immersing the user in a credible environment; not necessarily an accurate one. The advances shown by Bungie Studio are quite impressive: they mostly focus their research on rendering atmosphere, paying particular attention to how light and the sun behave with earth's atmosphere. Their tools enable artists to tweak a few sliders to change the look of a scene (without changing the 3d models, nor the position of the sun). Using these sliders, one can easily mimic realistic scenery, change a scene to have night ambience, or give the impression of a walk on the surface of an unknown world.

A lot of work is also being done regarding the use of multiple lights in a real time scene: multiple being up to 400 lights (which is quite a lot, even for precalculated movies). In a classic video game, the number of lights is usually limited to 8 or 16. However, in real time, by combining the most similar light sources and changing the rendering pipeline to be more efficient, it is possible to get very realistic rendering.

Although some real time demos have been presented at previous SIGGRAPH conferences, this year's exposure brought real time to the next level in more than one way.

The biggest adjustments this year were the new Real Time Rendering category added to the Computer Animation Festival, as well as the Real Time Live session held on the Wednesday afternoon. There was always interesting research being done in real time rendering, but the problem was how to show it to the SIGGRAPH audience while maintaining the interesting part about it: the interactivity. Some demos made their way to the Computer Animation Festival in previous years, but they were only video captures of real time graphics. As such, the audience did not always grasp the full extent and achievement of these demos.



Flower (PS3, ThatGameCompany)

This year, in the Real Time screenings, instead of showing recorded video of each of the featured projects, all the authors manipulated their projects in real time, while concurrently commenting on the hows and whys of their demos. Thus, these presentations became even more interesting: one could watch the live demo itself while also getting a glimpse at what really happened behind the scenes.

The demo screenings shown at SIGGRAPH were very rich, both in terms of quality and diversity. Three full feature games were presented, each with different gameplay and rendering styles, but all very impressive. Attendees were privy to **Flower**, with its dreamlike rendering, **Split Second** with its intense and explosive graphics, and of course **Fight Night 4** – complete with special appearance from Will Wright (one of the Keynote Speakers this year, and creator of Spore) in digital form, who managed to beat Mike Tyson in a fair fight.



Split Second (XBox 360 / PS3 / PC, Disney Interactive)

Following the featured game demos, came the technical demos. Three of these were presented by NVIDIA: a fluid dynamic rendering (which gave a very accurate result), a realistic hair simulation (in which both the animation and rendering were of a quality one would expect in a movie), and a demo depicting **Medusa** (complete with shaders, effects and millions of triangles).

AMD also presented a technical demo – it depicted the lives of little creatures called **Froblins**. The rendering in this demo was very impressive: each Froblin could potentially be rendered with up to 1.6 million polygons (depending on the level of detail they were seen at), and the scene featured hundreds of these frog goblins. This demo also intensively used General Purpose computation on GPU (GPGPU). Indeed, why should one limit the use of the GPU to render vertices and pixels, when one can use these computational resources to add AI (as was done in Froblins).

All the aforementioned advances are mostly high level advances, using the core pipeline available and adding new steps as needed. As mentioned in the paper “**GRAMPS**, a programming model for Graphic Pipelines”, to prevent things from getting too complicated a group from Stanford University tried to create a programmable pipeline. The concept they came up with provides the user with a generic tool that lets him create his own pipeline (adapted to his desired rendering specifications), allowing for the use of the GPU for alternate tasks.

On a different note, for the past couple of decades, people have sought to bridge the gap between real time and visual effects/CG animation. This year, a French studio (Studio Delacave) presented their use of a real time game engine to render a feature-length animated movie: **The True Story of Puss'n Boots**. This approach enabled them to render the movie in no time; some portions of which were even recomputed a few weeks before the deadline. In addition,

artists were also able to make changes to the materials, lights and passes, in real time - and get immediate results for feedback.



The True Story of Puss'n Boots (Studio Delacave)

The longest time required to render an image for the movie Puss'n Boots was 20 minutes (with the image including effects such as ambient occlusion, fake global illumination, etc.). This is a first step that shows that real time and precalculated rendering are leaning towards each other. Even though some effects will always need time to be precomputed, we're definitely moving towards movie quality real time interactive experiences, and movies will continue to have more flexibility as the rendering time shortens with the use of real time rendering solutions.

Xavier Gouchet