The Deep Magic

an industry op-ed by Steve Wright

the new deep compositing in visual effects

I was recently fondling a shiny new beta version of the new Nuke6.3 release and among the exciting new goodies like a planar tracker, an integral particle system, and a new displacement shader for stereo conversion; I discovered Nuke's deep magic - integral support for the new deep image technology. With deep images you can do deep compositing, and this is a real game-changer for the world of compositing visual effects.

The basic idea behind deep images is that each pixel contains not just RGB, alpha, and depth Z information, but multiple layers of transparency and depth that reach deep into the scene. The number of samples depends on the compositing requirements, but can be 10 to 20 or even more. Each pixel of the rendered layer "knows" where it belongs in depth and can be composited with other deep image layers without holdout mattes. The pixels literally sort themselves out back to front. If holdout mattes are needed for any reason they can be trivially generated from the deep image information.

The inspiration for deep images was surely from Pixar's Renderman which has supported deep shadows for some time now using their proprietary dtex file format. Deep shadows were first used for shadowing the fur in Pixar's "Monster's Inc." (released 2001) and have since been accepted as an expensive but essential photorealistic component of cg renders today. They are larger and slower than traditional shadow maps, but look fantastic for hair, fur (think "King Kong"), and shadow casting from semi-transparent objects. Not to be used lightly for every cg object, just for those that warrant the extra expense.

True deep images were developed at Weta Digital in New Zealand for James Cameron's blockbuster "Avatar". The jungle scenes were so complex that multiple re-renders were out of the question. With traditional render layers, as the cg characters walk through the cg jungle the holdout maps for the foreground and background would have to be re-rendered whenever a character was re-animated, which happened constantly. Using deep image renders only the characters had to be re-rendered when they were re-animated. They could then be re-composited with the existing deep forest renders because the depth information required to layer the forest and characters together was carried in the rendered image rather than being baked into a long list of render layers.

Deep compositing also really pays off for complex scenes that have rapid changes or great complexity in the layering order of the foreground and background. Imagine a flock of birds flying through trees, or a character walking through high grass, or (to take an appropriate Avatar example), furious air battles with banshees flying between the trees.

Avatar game screen capture courtesy of Ubisoft

There are also huge advantages over traditional depth Z compositing which typically creates edge artifacts like the example shown below. Since each pixel only has a single transparency value at a single depth, so edge issues such as motion blur and depth of field are not accurately represented resulting in aliasing and hard edges. Compositing with deep images, however, cleans up all those edge artifacts because the great increase in information about each pixel allows for visually correct blending of the edges.

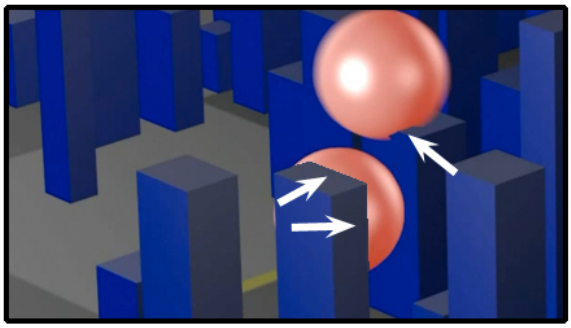

Edge artifacts from typical depth Z composite

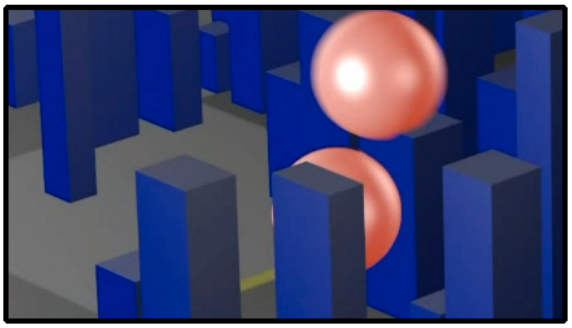

Clean edges with deep compositing *(images courtesy of The Foundry)*

One concern about deep images is the huge increase in file size for the additional information. An efficient compression scheme will be an essential component of a future industry standard file format for deep images. I'm sure that the steady increase in disc size, network bandwidth, and processing speed will minimize this problem in the future.

The main obstacle to industry acceptance of this powerful new form of compositing is an established file format for storing images. Currently, each vfx studio implements it in its own internal proprietary format or adapts Pixar's dtex format to their needs. The Foundry, makers of Nuke, (my personal favorite compositing software) are releasing a beta version with support for deep images and deep compositing in the upcoming Nuke6.3. The Foundry is also working towards the establishment of industry wide standards for deep images in OpenEXR 2.0 and will offer a free maintenance release when OpenEXR 2.0 is available. I can hardly wait.

Steve Wright