Deep Compositing

Daniel Heckenberg* Johannes Saam† Colin Doncaster‡ Chris Cooper§
Animal Logic

Figure 1: *Integration of separately rendered cloud volume and intersecting objects can be achieved without holdout mattes or any manual intervention (left). Deep images for each render are combined in 3d (right) and may be futher manipulated before flattening.*

Abstract

At Animal Logic, we exploit the potential of images produced by 3d rendering rather than physical cameras to improve pipeline and workflow processes in compositing, in particular facilitating the integration of stereoscopic and volumetric elements. Whilst traditional compositing happens between flat images, many operations depend on the depth relationship of image elements. We render and composite with "deep images", which contain information about coverage and depth along each pixel. This enables artists to achieve standard compositing operations without manual, explicit image segmentation or layering as well as decoupling rendering and compositing by obviating holdout mattes. Post-render application of depth-of-field may also be approximated with fewer artefacts.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Bitmap and framebuffer operations

Keywords: compositing, deep opacity, deep shadow maps

1 Deep images

In contrast to a simple image that stores only the final appearance at each pixel, a deep image records the 3d renderer's list of all points in depth that contribute to the pixel. Every surface or volume sample making up the final pixel is stored with its appearance information (colour and opacity) as well as its distance to the camera. As transparent objects may have boundary and interior samples, a depth extent can be represented also as shown in figure 1.

2 Workflow and Pipeline Advantages

Depth ordering requires no manual layer creation or stacking. Objects can also be rendered separately and combined via compositing without requiring holdout mattes. These advantages are particularly significant for stereoscopic elements where layer orders may have to differ for each camera and for volumetric elements where often no simple layer decomposition is possible.

Merging of scene elements doesn't require explicit layering or holdout mattes, so rendering of separate objects may be completely decoupled and is independent of compositing requirements. Multiple samples over pixel depth allow higher quality depth-of-field approximation allowing this to be tuned without re-rendering.

3 Deep Compositing Operations

Typical compositing operations on layers [Porter and Duff 1984] map directly to operations between deep pixel samples. Image processing operations are trivally operations on all of the pixel samples, with an added dimension of control for parameter variation over depth. For example, colour grading may be applied with variation though all three dimensions of scene space. We have implemented node-based operations for deep images within our compositing software to allow artists to work as they are accustomed. The extra data in deep images allows a number of novel compositing operations: 3d clipping using primitves or planes, accurate fog and volumetric effects and high quality post-render depth-of-field.

4 Implementation

Deep images are used in several rendering and shading approaches for shadowing and occlusion calculation. Most prominent is Pixar's RenderMan deep shadow technique[Lokovic and Veach 2000] in which deep textures are rendered through a light camera and evaluated at points visible to the render camera. We use the same rendering approach but store deep data from the render camera, still using Pixar's deep shadow map format which records an opacity function over the distance to the camera for each pixel. When working with deep images in compositing, we convert that data using an un-over operation to discrete samples with depth and extent.

References

LOKOVIC, T., AND VEACH, E. 2000. Deep shadow maps. *SIGGRAPH Comput. Graph.*, 385–392.

PORTER, T., AND DUFF, T. 1984. Compositing digital images. *SIGGRAPH Comput. Graph. 18*, 3, 253–259.

*e-mail: danielh@al.com.au
†Now at Dr.D Studios, e-mail: johannes.saam@gmail.com
‡Now at Storm Studios, e-mail: colin@colindoncaster.com
§Now at Dr.D Studios, e-mail: chris.cooper@drdstudios.com