TidScene: Pixar's Pipeline Backplane

Arun Rao
Pixar Animation Studios
Emeryville, CA, USA
rao@pixar.com

1. Introduction

The construction of an animated CG scene can be thought of as a series of steps involving the creation of geometric models, placing them in a hierarchical scene graph, and layering different kinds of information (animation, shading) on the graph. The scene thus constructed is then traversed by a renderer to produce the final sequence of images.

All these steps are performed in different applications. Even though the scene itself may be very large and geometrically complex, the amount of data actually authored in each application is usually quite small. It is therefore very desirable to avoid copying data as much as possible, and instead employ a referencing mechanism with sparse overrides.

We present our solution to this problem, a multi-file scene database called *TidScene*.

2. What is TidScene?

TidScene is a file-format, and supporting code libraries. The low-level, on-disk representation is Berkeley DB, constructed with a special schema which is interpreted by the associated code. This schema provides the following facilities:

- Storage of time-indexed data for well-known scenegraph entities: meshes, curves, transforms, shaders, procedurals, cameras etc.
- A transform hierarchy of nodes representing these entities
- Storage of arbitrarily-typed blind data on each node
- Representation of *references* to other files

The TidScene libraries add:

- A method of constructing a lightweight, in-memory representation of the scenegraph embodied in a file, which might itself be an aggregate of referenced files.
- A standard *override* semantic, by which selected attributes in one file may be altered by values for the same attribute in another.
- A convenient API by which the scene can be reasoned about, without regard to referencing structure
- Methods for traversing the graph for rendering and other purposes.

3. Use of TidScene in Pixar's Pipeline

TidScene is now the primary means of scene data interchange between different applications. Applications will typically receive a scene from another application, add new data and/or modify existing data by creating another file that "layers" on top of the incoming file.

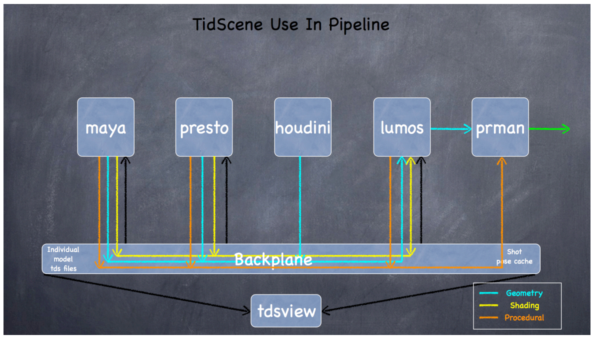

Figure 1. *Data interchange between applications*

For example, our animation system *Presto* assembles a hierarchy of references representing the scene, and deposits data for the animation on top of this aggregate. This "pose-cache" is then passed on to Houdini for further FX augmentation.

New applications may be deployed in the pipeline by writing a TidScene importer and exporter for that application.

4. Benefits

TidScene has now been in use at Pixar for two features (*Brave, Monsters University)* and one short (*The Blue Umbrella)*. It has fundamentally changed the way data flows between applications.

We call out the following significant benefits that have accrued:

- Significantly more data sharing, far less file copying
- "Just-in-time" data access. Scene data can be reasoned about, but not actually read until needed.
- Easy new application deployment
- Uniform, lightweight API for examining and manipulating data at any stage
- Single file-format for geometry representation *and* scene assembly

SIGGRAPH 2013, July 21 – 25, 2013, Anaheim, California.

ACM 978-1-4503-2261-4/13/07