Pulsar Houdini Volume Modeling Plug-ins Suite

Presented by Rony Edde



Intro

I. Voxel containers and CFDs paradigm. II. In house proprietary solutions. III. Geometry based volume modeling advantages. IV.Common in house tools limitations and overhead. V. Proposed solution. VI.Non CFD volume modeling techniques and solutions. VII.Real time demos. VIII.Improvements and TODOs.

I. Voxel containers and CFDs

- 1. Storing volumetric data in voxels is expensive.
- 2. A 10x10x10 volume doubled in resolution results in 8 times the initial resolution.

2x 2y 2z => 8 xyz

1000v => 8000

- 3. Solving Navier-Stokes equations across all voxels exponentially increases simulation time and in many cases alters the results between iterations due to the unpredictable nature of fluids.
- 4. Smooth look prevents high level of detail at reasonable resolutions.
- 5. Directors who demand full control when fluid sims are tricky to control.

II. In house proprietary solutions

- 1. Proprietary volumetric tools aren't necessarily an emulation of fluids. In many cases these techniques are used to generate unrealistic elements as well. Here are a few techniques:
 - 1. Use points to generate clouds.
 - 2- Use splines/curves to generate volumetric tubes.
 - **3-** Use points to generate wisps.
 - **4** Use surfaces to generate wisps.
- 2. These methods offer a good array of results that "emulate" a CFD look. They also generate a very high detail volume at increased resolutions without the overhead of running fluid simulations.

III. Geometry based volume modeling advantages

- 1. Time independent. This gives the ability to farm out all frames simultaneously effectively reducing the time taken to generate a full frame range.
- 2. Dealing with relatively simple equations to solve, we can increase the volume size or regenerate a volume without sacrificing too much time.
- 3. No need to resim upon manipulating input.
- 4. Using input geometry to generate a volume gives the user complete control.

IV. Common in house tools limitations and overhead

- 1. Inability To visualize volume in real time without conversion overhead.
- 2. Need to render with lights to see the volume which is expensive.
- 3. New hires need to be trained in order to use the scripting language required to create/fill/manipulate buffers.
- 4. Overhead of converting the input geometry, then sending the data to a proprietary tool which generates the volume which could be converted to a Houdini volume for viewing. This of course is a waste of resources.

IV. Common in house tools limitations and overhead



V. Proposed solution

In order to bypass the overhead of data conversion and scripting, I developed multi-threaded plug-ins for Houdini that run transparent to the user without leaving the Houdini geo engine. Here are the advantages to using such a system:

- 1. Real time feedback in OpenGL.
- 2. Uses custom Houdini nodes that keep the work flow active for post processing.
- 3. Speed gain due to the lack of data conversion and file system overhead.
- 4. User familiarity with the Houdini interface.

V. Proposed solution



1. Fill with clouds plug-in.

This plug-in generates cloud volumes from input points using a custom Perlin noise implementation.



Input Point

Result

Video



P=input point R=radius $\rho=density$ V=voxel at index(i, j, k)

2. Fill with curves plug-in.

This plug-in generates a noised tube volume from input curves using a custom Perlin noise.





3. Fill with wisps plug-in.

This plug-in generates wisps from input points using custom Perlin noise distribution.





4. Fill with surface wisps plug-in.

This plug-in generates wisps from a mesh using custom Perlin noise distribution based on UVW coordinates.





Inefficient vector interpolation example:







UV coordinates interpolation for noise displacement:



Using frustum buffers reduces the memory footprint and optimizes the buffer without sacrificing resolution.

The benefits:

- 1. More resolution close to camera where it really matters.
- 2. Volume generation in optimized by processing geometry that is only inside the camera frustum.





Velocity Paradigm:

When transferring velocity vectors from points to neighboring voxels there are 2 possible options for each voxel.

- 1. Only one point is close enough to directly affects the voxel value.
- 2. More than one point is affecting the voxel and this case causes problems because we can only store one velocity component value per voxel.

Two points with different velocity vectors inside the same voxel.



The combined or average value is incorrect. The result is incorrect velocities.



The solution is to bake the velocity into the density and ignore the velocity field.



Baked velocity example:



VII. Real time demos



1. Generalizing object types with a generic attribute class.

When dealing with several types of primitives and data, it's often difficult to store information without previous knowledge of the type. This lead to the development of a simple but efficient AttributeContainer class.

The AttributeContainer class stores multiple arbitrary data types in one single object. This means that we can store an entire scene with miscellaneous object types which facilitates API updates and changes.

Sample code usage:

// using a test class we declared for testing
TEST test;
test.a=l; test.b=2; test.x = 3.0; test.y = 4.0;
Attr<TEST>* tt = new Attr<TEST>("TEST");
tt->setVal(test):

vector<float> a; a.push_back(1.23); a.push_back(5.954); a.push_back(12.256);

AttribContainer A; A.addAttrib(10.Of); A.addAttrib(12); A.addAttrib(a); A.addAttrib(a);

// retrieving values test

cout << "Attrib A t yal: " << A.getAttrib<float>(0)->getVal() << endl; cout << "Attrib A u yal: " << A.getAttrib<int>(1)->getVal() << endl; cout << "Attrib A z0 yal: " << A.getAttrib< vector<float> >(2)->getVal()[0] << endl; cout << "Attrib A z1 yal: " << A.getAttrib< vector<float> >(2)->getVal()[1] << endl;</pre>

//retrieving values from custom class

cout << "Attrib A test yal a: " << A.getAttrib<TEST>(3)->getVal().a << endl; cout << "Attrib A test yal b: " << A.getAttrib<TEST>(3)->getVal().b << endl; cout << "Attrib A test yal c: " << A.getAttrib<TEST>(3)->getVal().x << endl; cout << "Attrib A test yal d: " << A.getAttrib<TEST>(3)->getVal().y << endl; cout << "Attrib A test type: " << A.getAttrib<float>(0)->getType().c_str() << endl; cout << "Attrib A test type: " << A.getAttrib<float>(1)->getType().c_str() << endl; cout << "Attrib A test type: " << A.getAttrib<float>(1)->getType().c_str() << endl; cout << "Attrib A test type: " << A.getAttrib<float>(2)->getType().c_str() << endl; cout << "Attrib A test type: " << A.getAttrib<float>(2)->getType().c_str() << endl;</pre>

Output:

devel/houdini_plugins% ./AttribContainer Attrib A t val: 10 Attrib A u val: 12 Attrib A z0 val: 1.23 Attrib A z1 val: 5.954 Attrib A test val a: 1 Attrib A test val b: 2 Attrib A test val c: 3 Attrib A test val d: 4 Attrib A test type: f Attrib A test type: i Attrib A test type: i Attrib A test type: St6vectorIfSaIfEE Attrib A test type: 4TEST

2. GPU Optimization.

Optimizing with the GPU requires 2 main tests to pass.

- 1. The sub containers should not be too small otherwise time would be wasted updating between GPU and CPU memory.
- 2. The sub voxels need to be bigger than a certain size but not bigger than what the GPU memory can handle.

Conclusion

Modeling volumes has been a life saver on many feature films and they make directors happy.

The speed, control and detail they offer is unprecedented. Being able to integrate them within a Houdini pipeline is a good advantage for real time visualization and post processing.

The result is happy artists, supervisors and Directors.

Thank you so much for your time!