CS 33400, Fall 2012
Voicu Popescu

Due: Tuesday Dcemeber 4th at 6:00am

A10—Your choice

In a nutshell

Fast ray tracing, or shadow mapping, depth from stereo, or panorama stitching.

Fast Ray Tracing

Add an acceleration scheme (e.g. 3-D grid, BSP tree, octree, kd-tree) to your ray tracer. Report performance (i.e. average and minimum fps) with and without the acceleration scheme over a path of 100 frames.

Shadow mapping

Implement shadow mapping on the GPU. Hint: you need to pass to the fragment shader the projection-modelview matrix L corresponding to the camera modeling the light. For each fragment (i.e. pixel), project the 3-D surface point (x, y, z, 1) by multiplying by L to obtain (x', y', z', w'). Then divide x', y', and z' by w' to obtain image plane coordinates and "z" values that go from -1 to 1. Normalize to [0, 1], then look up your shadow map.

Depth from stereo

Implement an algorithm that given two overlapping images, their respective planar pinhole cameras and a point on the first image, will search for the corresponding point in the second image by marching on the epipolar segment. The epipolar segment is the projection with the second camera of the first camera ray through the given point onto the image plane of the second camera. If the correspondence is found, the 3-D point should be inferred by triangulation. To test your algorithm, render two images of the auditorium using your HW renderer. The two views should be identical except for a lateral translation. Then find 3-D points using your correspondence finding and triangulation algorithm. Visualize the 3-D points found.

Panorama stitching

Write an algorithm that takes two overlapping images with the same center of projection and merges them into an image with a larger field of view. The algorithm only knows the field of view of the input images and it does not know the relative rotation between the two images.

Turn in

- Source code
- Path used to render demo frames
- 30s 30Hz video showing your result
