

David Mazzocco

CS599 Scribe Notes

GPU Programming

Shading Languages/Tools

Offline

Renderman (Pixar)

Houdini

Gelato

Real Time

Architecture Review Board (ARB) assembly (low level)

Cg (NVIDIA 2005, high level)

GLSL (OpenGL Shader Language, high level)

HLSL (Direct3D, high level)

GPU-CPU Comparison

GPUs are more powerful than CPUs in terms of computation speed and bandwidth, as well as a faster growth rate and lower cost. This is due to GPUs being designed to compute a single task, versus the general nature of CPUs. This comes at a cost for GPUs by taking a large performance penalty for flow control/branching.

The CPU-GPU memory bus is bidirectional, but is biased towards moving memory to the GPU. One way to avoid a bottleneck on the bus is to use display lists, which caches memory on the GPU.

Graphics Pipeline

Application

The application determines what graphical data to send to the GPU to be rasterized, such as primitives and texture data.

Transform

The transform stage involves both the camera transform and perspective projection to generate screenspace coordinates for all vertices, plus a z-value for depth information. Lighting is also applied here (up to 8 lights in OpenGL), and uses the Phong model as default. This stage can be overwritten using a vertex shader.

Assemble Primitives

Triangles and other primitives are assembled from the vertices according to the input data. This stage can be overwritten by a geometry shader, which is a relatively new type of shader. Geometry shaders can allow for different methods of tessellation, as well as generate new geometry.

Rasterization

Pixels are superimposed over the image. Texture coordinates and vertex colors are interpolated through each pixel to determine the final color. The default behavior will use vertex colors to determine the pixel color, or will also use texture information if enabled. Textures can be used to either replace or modulate with the underlying vertex colors. This stage can be modified using a pixel/fragment shader.