

OpenGL Graphics Pipeline

Graphics Pipeline is an abstract model that describes sequence of steps needed to render a 3D scene.

asynchronous operation - The CPU sends rendering commands to the GPU, which then perform rendering operations while the CPU continues with other tasks

VRAM - Memory core on the GPU which stores buffers

- **image buffer** - Has back image buffer, where the GPU renders the scene, and the front image buffer, where the exact pixel data this is visible to the viewport. Buffer Swap is performed when the back image buffer is done rendering and ready to be displayed by the front image buffer.
- **Depth Buffer / Z- Buffer** - Stores values how far the pixel lies in every pixel in the image buffer. Used to perform hidden surface elimination, by allowing pixel drawn that if its depth is less than the pixel already in the image buffer
- **Stencil Buffer** -
- **Texture Maps** - Images applied to surface of an object. Can include other maps such as bump map. Takes up largest amount of VRAM

0. Core Concepts and Vocabulary

rendering - Generate two-dimensional images of 3D scenes

shading - The darkness of an object not in direct light

shadows - the silhouette of one object's shape on the surface of another object

frustum - Region contained within the truncated pyramid shape outlined in white indicates the space visible to the camera.

pixel - specify colors using triples of floating-point numbers between 0 and 1 to represent the amount of red, green, and blue light present in a color; a value of 0 represents no amount of that color is present, while a value of 1 represents that color at full intensity

- **raster** - rendered scene via an array of pixels (picture elements) which will be displayed on a screen, arranged in a 2D grid
- **resolution** - the number of pixels in the raster, the more it is the higher the quality
- **precision** - the number of bits used for each pixel as each bit has two possible values (0 or 1), the number of colors that can be displayed

buffer (data buffer/buffer memory) is a part of a computer's memory that serves as temporary storage for data while it is being moved from one location to another.

- **frame buffer** - Pixel data is stored in a region of memory. A framebuffer may contain multiple buffers that store different types of data for each pixel.
 - **color buffer** - located in frame buffer which stores RGB values. Need this at minimum. Alpha value can also be stored
 - **depth buffer** - located in frame buffer, which stores distances from points on scene objects to the virtual camera. Depth values are used to determine whether the various points on each object are in front of or behind other objects (from the camera's perspective), and thus whether they will be visible when the scene is rendered.
 - **stencil buffer** - store values used in generating advanced effects, such as shadows, refractions, or portal rendering.

1. Application Stage

Initializing the window where the rendered graphics will be displayed.

- Reading data required for the rendering process and sending to the GPU, such as
 - **vertex attributes**, describes appearance of geometric shapes rendered, stored as in **vertex buffer objects** (VBO)
 - images to be applied to surfaces, stored in **texture buffers**
 - source code for **vertex shader** and **fragment shader** programs, sent to GPU to be compiled and loaded.
- Loop that re-renders the scene repeatedly, like 60 fps
- Monitoring hardware for user inputs, handled by the CPU
- **Vertex Array Objects**, manages the associations and whether they are turned on and off, between attributes data stored in VBOs and attribute variables in the vertex shader program

2. Geometry Processing

Determining the final position of each vertex of the geometric shapes to be rendered, implemented by a program called the vertex shader

mesh - a collection of points (vertices) that are grouped into lines or triangles to make a shape of a geometric object

- **vertex** - a point with a data structure holding properties or attributes that are specific to rendering.
 - 3D position of the corresponding point. Mandatory
 - color to be used when rendering the point. Optional

- **texture coordinates** (or UV coordinates) - indicates a point in an image that is mapped to the vertex. Optional
- **normal vector** - indicates the direction perpendicular to a surface, used for lighting calculations. Optional

Vertex shader is applied to each of the vertices to determine the final position each point being rendered, which is typically calculated from a series of transformations:

- **model transformation** - the collection of points defining the intrinsic shape of an object may be translated, rotated, and scaled so that the object appears to have a particular location, orientation, and size with respect to a 3D world.
 - **world space** - coordinates expressed from this frame of reference are said to be in world space
 - **virtual camera** - camera with its own position and orientation in the virtual world.
 - **view transformation** - In order to render the world from the virtual camera's point of view, the coordinates of each object in the world must be converted to a frame of reference relative to the camera itself.
 - **view space** (camera/eye space) - coordinates after view transformation
 - **projection transformation** - clipping points outside the view space
 - **clip space** - points outside the view space
 - perspective projection
 - orthographic projection

In addition to these transformation calculations, the vertex shader may perform additional calculations and send additional information to the **fragment shader** as needed.

The end of geometry processing

3. Rasterization

geometric primitive - How the vertices is connected to produce a shape. OpenGL supports ten of these types, including Points, lines, or triangles, which consist of sets of 1, 2, or 3 points.

Rasterization - Process of filling in the horizontal spans of pixels belonging to a geometric primitive.

primitive assembly - process of grouping points to geometric primitives

Once the geometric primitives have been assembled, the next step is to determine which pixels correspond to the interior of each geometric primitive. Since pixels are discrete units, they will typically only approximate the continuous nature of a geometric shape, and a criterion must be given to clarify which pixels are in the interior.

Three simple criteria could be

1. The entire pixel area is contained within the shape
2. The center point of the pixel is contained within the shape
3. Any part of the pixel is contained within the shape

fragment - For each pixel corresponding to the interior of a shape

- raster position / pixel position - data stored in a fragment
 - **depth** - stored in fragment when in 3D. Needed when points on different geometric objects would overlap from the perspective of the viewer. When this happens, the associated fragments would correspond to the same pixel, and the depth value determines which fragment's data should be used when rendering this pixel
 - Additional data may be assigned to each vertex, such as a color, and passed along from the vertex shader to the fragment shader. In this case, a new data field is added to each fragment.
 - **interpolated** - from the values at the vertices: calculated using a weighted average, depending on the distance from the interior point to each vertex. Optional

4. Pixel Processing

This stage determines the final color of each pixel, storing this data in the color buffer within the frame buffer.

During the first part, a program called the fragment shader is applied to each of the fragments to calculate their final color. This calculation involves various data stored in each fragment, in combination with global data available during rendering, such as

- base color applied to the entire shape
- colors stored in each fragment (interpolated from vertex colors)
- textures, where colors are sampled from locations specified by UV coordinates
- light sources, whose relative position and/or orientation may lighten or darken the color, depending on the direction the surface is facing at a point, specified by normal vector

The GPU handles the following:

- Depth values stored in each fragment are used to resolve visibility issues in a 3D scene, determining which parts of objects are blocked from view by other objects. After the color of a fragment has been calculated, the fragment's depth value will be compared to the value currently stored in the depth buffer at the corresponding pixel coordinates. If the fragment's depth value is smaller than the depth buffer value, then the corresponding point is closer to the viewer than any that were previously processed, and the fragment's color will be used to overwrite the data currently stored in the color buffer at the corresponding pixel coordinates.
- Transparency using alpha values stored in the color of each fragment, determining how much color blend with another color. All opaque objects must be rendered first (in any order), followed by transparent objects ordered from farthest to closest with respect to the

virtual camera.

Revision #15

Created 8 February 2024 09:07:59 by victor

Updated 9 August 2024 22:49:52 by victor