



Vulkan.


Textures



Oregon State University
Mike Bailey
mb@cs.oregonstate.edu



This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).



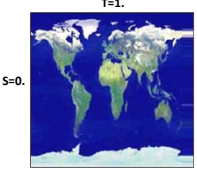
Oregon State University
Computer Graphics

Textures.pptx mjb - September 18, 2018


The Basic Idea

Texture mapping is a computer graphics operation in which a separate image, referred to as the **texture**, is stretched onto a piece of 3D geometry and follows it however it is transformed. This image is also known as a **texture map**. This can be most any image. At one time, some graphics hardware required the image's pixel dimensions to be a **power of two**. This restriction has been lifted on most (all?) graphics cards, but just to be safe... The X and Y dimensions did not need to be the same power of two, just a power of two. So, a 128x512 image would have been OK; a 128x511 image might not have.

Also, to prevent confusion, the texture pixels are not called **pixels**. A pixel is a dot in the final screen image. A dot in the texture image is called a **texture element** or **texel**. Similarly, to avoid terminology confusion, a texture's width and height dimensions are not called X and Y. They are called **S** and **T**. A texture map is not generally indexed by its actual resolution coordinates. Instead, it is indexed by a coordinate system that is resolution-independent. The left side is always **S=0.**, the right side is **S=1.**, the bottom is **T=0.**, and the top is **T=1.** Thus, you do not need to be aware of the texture's resolution when you are specifying coordinates that point into it. Think of S and T as a measure of what fraction of the way you are into the texture.



S=0 T=1
S=1 T=0

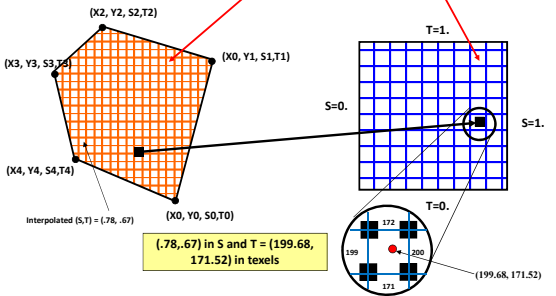


Oregon State University
Computer Graphics


mjb - September 18, 2018

The Basic Idea

The mapping between the geometry of the 3D object and the S and T of the texture image works like this:



You specify an (s,t) pair at each vertex, along with the vertex coordinate. At the same time that the rasterizer is interpolating the coordinates, colors, etc. inside the polygon, it is also interpolating the (s,t) coordinates. Then, when it goes to draw each pixel, it uses that pixel's interpolated (s,t) to look up a color in the texture image.



Oregon State University
Computer Graphics

mjb - September 18, 2018

In OpenGL terms: assigning an (s,t) to each vertex

Enable texture mapping:

```
glEnable( GL_TEXTURE_2D );
```

Draw your polygons, specifying s and t at each vertex:


```
glBegin( GL_POLYGON );
  glTexCoord2f( s0, t0 );
  glNormal3f( nx0, ny0, nz0 );
  glVertex3f( x0, y0, z0 );

  glTexCoord2f( s1, t1 );
  glNormal3f( nx1, ny1, nz1 );
  glVertex3f( x1, y1, z1 );

  ...
glEnd();
```

Disable texture mapping:

```
glDisable( GL_TEXTURE_2D );
```



Oregon State University
Computer Graphics

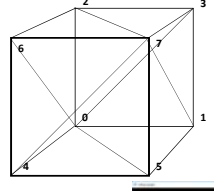

mjb - September 18, 2018

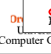
Triangles in an Array of Structures

```

struct vertex
{
  glm::vec3  position;
  glm::vec3  normal;
  glm::vec3  color;
  glm::vec2  texCoord;
};

struct vertex VertexData[] =
{
  // triangle 0-2-3:
  // vertex #0:
  { -1, -1, -1, },
  { 0, 0, -1, },
  { 0, 0, 0, },
  { 1, 0, 0, },
  // vertex #2:
  { -1, 1, -1, },
  { 0, 0, -1, },
  { 0, 0, 0, },
  { 1, 1, 0, },
  // vertex #3:
  { 1, 1, -1, },
  { 0, 0, -1, },
  { 1, 1, 0, },
  { 0, 1, 0, },
};
    
```

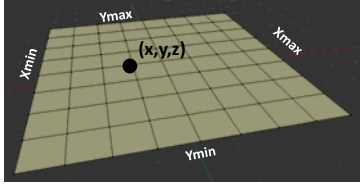


Oregon State University
Computer Graphics


mjb - September 18, 2018

Using a Texture: How do you know what (s,t) to assign to each vertex?

The easiest way to figure out what s and t are at a particular vertex is to figure out what fraction across the object the vertex is living at. For a plane,



$$s = \frac{x - X_{min}}{X_{max} - X_{min}} \quad t = \frac{y - Y_{min}}{Y_{max} - Y_{min}}$$

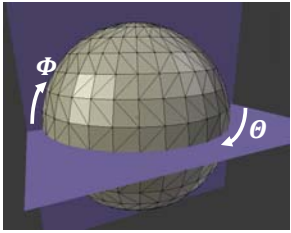


Oregon State University
Computer Graphics

mjb - September 18, 2018

Using a Texture: How do you know what (s,t) to assign to each vertex? 7


Or, for a sphere,



$$s = \frac{\theta - (-\pi)}{2\pi} \quad t = \frac{\phi - (-\frac{\pi}{2})}{\pi}$$

From the Sphere code:

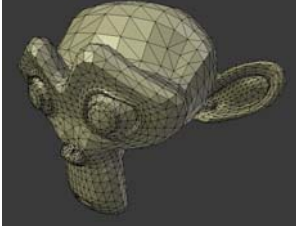
$$s = (\text{In}g + M_PI) / (2 * M_PI);$$

$$t = (\text{lat} + M_PI/2.) / M_PI;$$


Oregon State University Computer Graphics

Using a Texture: How do you know what (s,t) to assign to each vertex? 8

Uh-oh. Now what? Here's where it gets tougher...




$s = ?$ $t = ?$

Oregon State University Computer Graphics

mjb - September 18, 2018

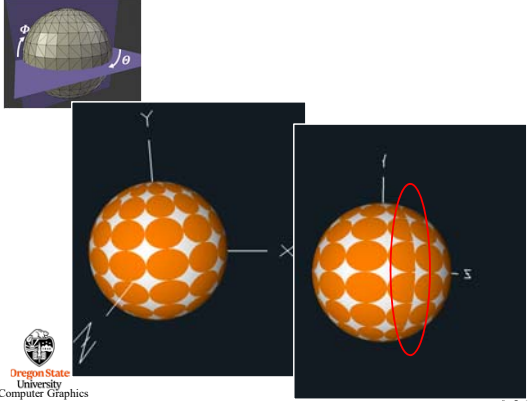
You really are at the mercy of whoever did the modeling... 9



Oregon State University Computer Graphics

mjb - September 18, 2018

Be careful where s abruptly transitions from 1. back to 0. 10



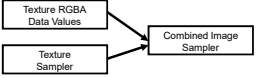
Oregon State University Computer Graphics

mjb - September 18, 2018

11

```

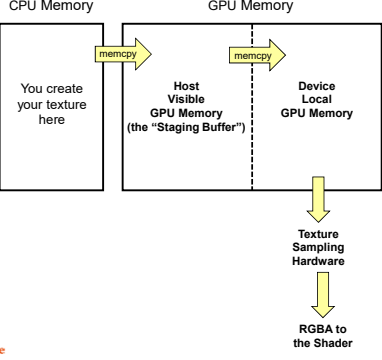
VkDescriptorSetLayoutBinding TexSamplerSet(1);
TexSamplerSet(0).binding = 0;
TexSamplerSet(0).descriptorType = VK_DESCRIPTOR_TYPE_COMBINED_IMAGE_SAMPLER;
// uniform sampler2D uSampler
// vec4 rgba = texture( uSampler, vST );
TexSamplerSet(0).descriptorCount = 1;
TexSamplerSet(0).stageFlags = VK_SHADER_STAGE_FRAGMENT_BIT;
TexSamplerSet(0).pImmutableSamplers = (VkSampler *)nullptr;
...
VkDescriptorImageInfo vdi0;
vdi0.sampler = MyPuppyTexture.texSampler;
vdi0.imageView = MyPuppyTexture.texImageView;
vdi0.imageLayout = VK_IMAGE_LAYOUT_SHADER_READ_ONLY_OPTIMAL;
...
VkWriteDescriptorSet wds3;
wds3.sType = VK_STRUCTURE_TYPE_WRITE_DESCRIPTOR_SET;
wds3.pNext = nullptr;
wds3.dstSet = DescriptorSets[3];
wds3.dstBinding = 0;
wds3.dstArrayElement = 0;
wds3.descriptorCount = 1;
wds3.descriptorType = VK_DESCRIPTOR_TYPE_COMBINED_IMAGE_SAMPLER;
wds3.pBufferInfo = (VkDescriptorBufferInfo *)nullptr;
wds3.pImageInfo = &vdi0;
wds3.pTexelBufferView = (VkBufferView *)nullptr;
    
```



Oregon State University Computer Graphics

mjb - September 18, 2018

Memory Types 12



Oregon State University Computer Graphics

mjb - September 18, 2018


```

// =====
// translate the texture buffer layout a second time
// =====
{
    VkImageSubresourceRange visr;
    visr.aspectMask = VK_IMAGE_ASPECT_COLOR_BIT;
    visr.baseMipLevel = 0;
    visr.levelCount = 1;
    visr.baseArrayLayer = 0;
    visr.layerCount = 1;

    VkImageMemoryBarrier vimb;
    vimb.sType = VK_STRUCTURE_TYPE_IMAGE_MEMORY_BARRIER;
    vimb.pNext = nullptr;
    vimb.oldLayout = VK_IMAGE_LAYOUT_TRANSFER_DST_OPTIMAL;
    vimb.newLayout = VK_IMAGE_LAYOUT_SHADER_READ_ONLY_OPTIMAL;
    vimb.srcQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED;
    vimb.dstQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED;
    vimb.image = textureImage;
    vimb.srcAccessMask = 0;
    vimb.dstAccessMask = VK_ACCESS_SHADER_READ_BIT;
    vimb.subresourceRange = visr;

    vkCmdPipelineBarrier(TextureCommandBuffer,
        VK_PIPELINE_STAGE_TRANSFER_BIT, VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT, 0,
        0, (VkMemoryBarrier*)nullptr,
        0, (VkBufferMemoryBarrier*)nullptr,
        1, (Vimb*)&vimb);
}

result = vkEndCommandBuffer( TextureCommandBuffer );

VkSubmitInfo vsti;
vsti.sType = VK_STRUCTURE_TYPE_SUBMIT_INFO;
vsti.pNext = nullptr;
vsti.commandBufferCount = 1;
vsti.pCommandBuffers = &TextureCommandBuffer;
vsti.waitSemaphoreCount = 0;
vsti.pWaitSemaphores = (VkSemaphore*)nullptr;
vsti.signalSemaphores = (VkSemaphore*)nullptr;
vsti.pFinishedSemaphore = (VkSemaphore*)nullptr;
vsti.pFinishedSemaphore = (VkSemaphore*)nullptr;

result = vkQueueSubmit( Queue, 1, IN &vsti, VK_NULL_HANDLE );
result = vkQueueWaitIdle( Queue );
    
```



```

// create an image view for the texture image:

VkImageSubresourceRange visr;
visr.aspectMask = VK_IMAGE_ASPECT_COLOR_BIT;
visr.baseMipLevel = 0;
visr.levelCount = 1;
visr.baseArrayLayer = 0;
visr.layerCount = 1;

VkImageViewCreateInfo vici;
vici.sType = VK_STRUCTURE_TYPE_IMAGE_VIEW_CREATE_INFO;
vici.pNext = nullptr;
vici.flags = 0;
vici.image = textureImage;
vici.viewType = VK_IMAGE_VIEW_TYPE_2D;
vici.format = VK_FORMAT_R8G8B8A8_UNORM;
vici.components.r = VK_COMPONENT_SWIZZLE_R;
vici.components.g = VK_COMPONENT_SWIZZLE_G;
vici.components.b = VK_COMPONENT_SWIZZLE_B;
vici.components.a = VK_COMPONENT_SWIZZLE_A;
vici.subresourceRange = visr;

result = vkCreateImageView( LogicalDevice, IN &vici, PALLOCATOR_OUT &pMyTexture->textureImageView);

return result;
    
```



Note that, at this point, the CPU buffer and the GPU Staging Buffer are no longer needed, and can be destroyed.

Reading in a Texture from a BMP File

```

typedef struct MyTexture
{
    uint32_t width;
    uint32_t height;
    VkImage texture;
    VkImageView textureView;
    VkSampler texSampler;
    VkDeviceMemory vdm;
} MyTexture;

...

MyTexture MyPuppyTexture;
    
```

```

result = Init06TextureBufferAndFillFromBmpFile ( "puppy.bmp", &MyTexturePuppy);
Init06TextureSampler( &MyPuppyTexture.texSampler );
    
```

This function can be found in the **sample.cpp** file. The BMP file needs to be created by something that writes uncompressed 24-bit color BMP files, or was converted to the uncompressed BMP format by a tool such as ImageMagick's *convert*, Adobe *Photoshop*, or GNU's *GIMP*.



Anisotropic Texture Filtering

https://en.wikipedia.org/wiki/Anisotropic_filtering

