



Vulkan Ray Tracing – 5 New Shader Types!



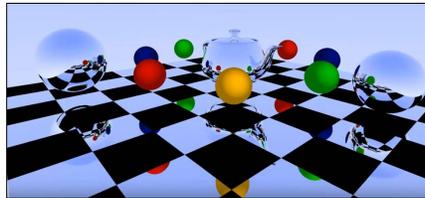
Oregon State University

Mike Bailey

mjb@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/)



VulkanRayTracing.pptx

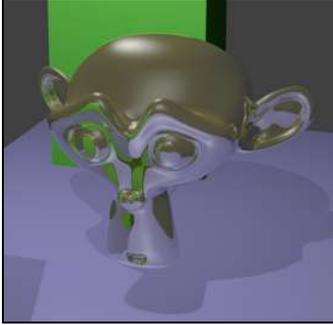
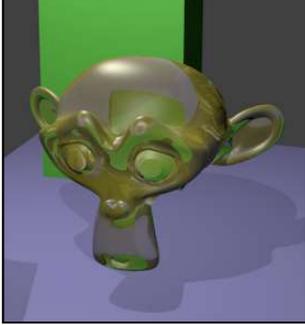
mjb – February 1, 2022

Analog Ray Tracing Example ☺

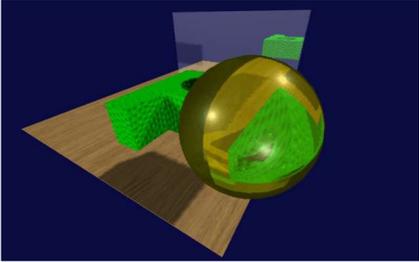


mjb – February 1, 2022

Digital Ray Tracing Examples



Blender



IronCad

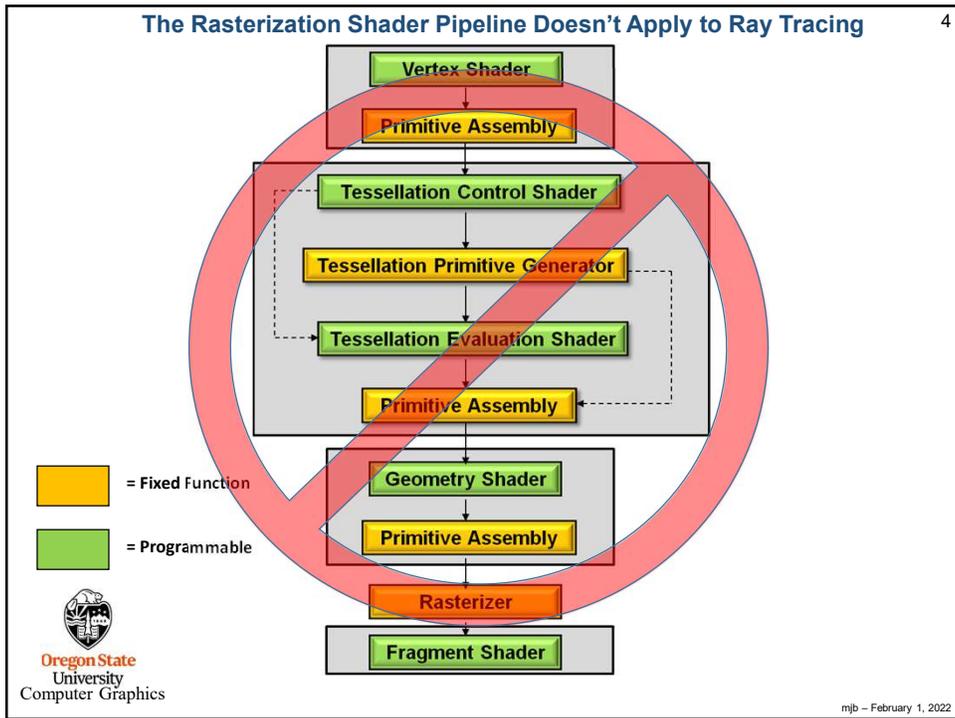


Oregon State University
Computer Graphics

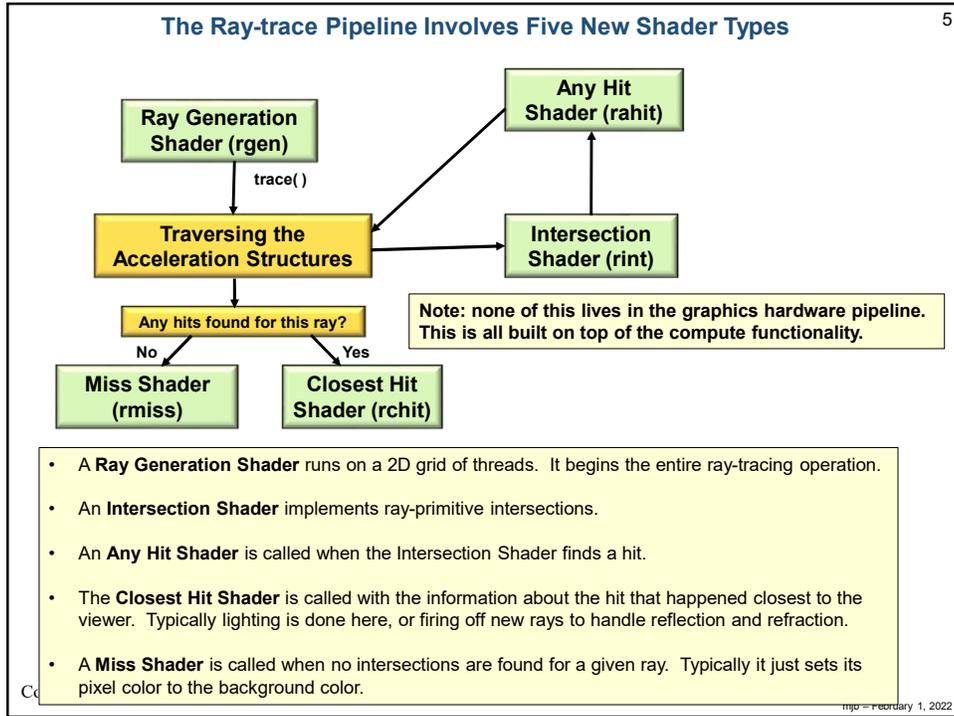
mjb - February 1, 2022

3

3



4



5

The Ray Intersection Process for a Sphere

6

1. Sphere equation: $(x-x_c)^2 + (y-y_c)^2 + (z-z_c)^2 = R^2$
2. Ray equation: $(x,y,z) = (x_0,y_0,z_0) + t*(dx,dy,dz)$

Plugging (x,y,z) from the second equation into the first equation and multiplying-through and simplifying gives:

$$At^2 + Bt + C = 0$$

Solve for t_1, t_2

- A. If both t_1 and t_2 are complex, then the ray missed the sphere.
- B. If $t_1 = t_2$, then the ray brushed the sphere at a tangent point.
- C. If both t_1 and t_2 are real and different, then the ray entered and exited the sphere.

In Vulkan terms:
 $gl_WorldRayOrigin = (x_0,y_0,z_0)$
 $gl_Hit = t$
 $gl_WorldRayDirection = (dx,dy,dz)$

mjb - February 1, 2022

6

The Ray Intersection Process for a Cube 7

1. Plane equation: $Ax + By + Cz + D = 0$
2. Ray equation: $(x,y,z) = (x_0,y_0,z_0) + t^*(dx,dy,dz)$

Plugging (x,y,z) from the second equation into the first equation and multiplying-through and simplifying gives:

$At + B = 0$
Solve for t

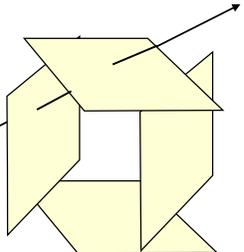
A cube is actually the intersection of 6 half-space planes (just 4 are shown here). Each of these will produce its own t intersection value. Treat them as pairs: (t_{x1}, t_{x2}) , (t_{y1}, t_{y2}) , (t_{z1}, t_{z2})

The ultimate entry and exit values are:

$t_{min} = \max(\min(t_{x1}, t_{x2}), \min(t_{y1}, t_{y2}), \min(t_{z1}, t_{z2}))$
 $t_{max} = \min(\max(t_{x1}, t_{x2}), \max(t_{y1}, t_{y2}), \max(t_{z1}, t_{z2}))$

}

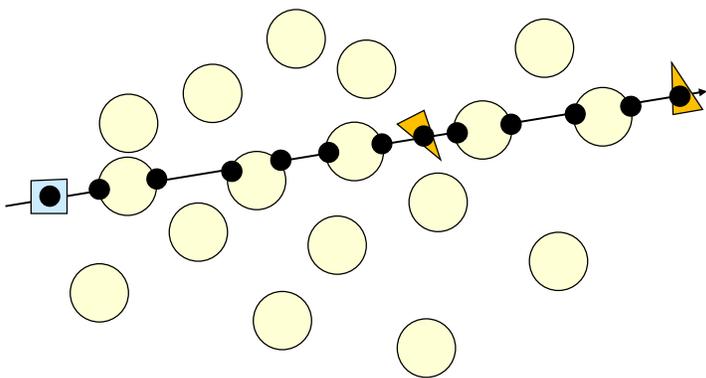
This works for all convex solids




mjb - February 1, 2022

7

In a Raytracing, each ray typically hits a lot of Things 8




mjb - February 1, 2022

8

Acceleration Structures

- Bottom-level Acceleration Structure (BLAS) holds the vertex data and is built from vertex and index VkBuffers
- The BLAS can also hold transformations, but it looks like usually the BLAS holds vertices in the original Model Coordinates.
- Top-level Acceleration Structure (TLAS) holds a pointer to elements of the BLAS and a transformation.
- The BLAS is used as a Model Coordinate bounding box.
- The TLAS is used as a World Coordinate bounding box.
- A TLAS can instance multiple BLAS's.



February 1, 2022

9

Creating Bottom Level Acceleration Structures

```

vkCreateAccelerationStructure( LogicalDevice, IN &vasi, PALLOCATOR, OUT &BottomLevelAccelerationStructure );

VkAccelerationStructureInfo
vasi.sType = VK_ACCELERATION_STRUCTURE_TYPE_BOTTOM_LEVEL;
vasi.flags = 0;
vasi.pNext = nullptr;
vasi.instanceCount = 0;
vasi.geometryCount = << number of vertex buffers >>
vasi.pGeometries = << vertex buffer pointers >>

VkAccelerationStructureCreateInfo
vasci.sType = VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_CREATE_INFO;
vasci.pNext = nullptr;
vasci.info = &vasi;
vasci.compactedSize = 0;
    
```



mjb - February 1, 2022

10

Creating Top Level Acceleration Structures

```

vkCreateAccelerationStructure              TopLevelAccelerationStructure;

VkAccelerationStructureInfo
vasi.sType = VK_ACCELERATION_STRUCTURE_TYPE_TOP_LEVEL;
vasi.flags = 0;
vasi.pNext = nullptr;
vasi.instanceCount = << number of bottom level acceleration structure instances >>;
vasi.geometryCount = 0;
vasi.pGeometries = VK_NULL_HANDLE;

VkAccelerationStructureCreateInfo
vasci.sType = VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_CREATE_INFO;
vasci.pNext = nullptr;
vasci.info = &vasi;
vasci.compactedSize = 0;

result = vkCreateAccelerationStructure( LogicalDevice, &vasci, PALLOCATOR, &TopLevelAccelerationStructure );
    
```

mjb - February 1, 2022

11

Ray Generation Shader

Gets all of the rays going and writes the final color to the pixel

```

layout( location = 1 ) rayPayload myPayload
{
    vec4 color;
};

void
main( )
{
    trace( topLevel, ..., 1 );
    imageStore( framebuffer, gl_GlobalInvocationID.xy, color );
}
    
```

A "payload" is information that keeps getting passed through the process. Different stages can add to it. It is finally consumed at the very end, in this case by writing *color* into the pixel being worked on.

mjb - February 1, 2022

12

A New Built-in Function

```

void trace
(
    accelerationStructure    topLevel,
    uint                     rayFlags,
    uint                     cullMask,
    uint                     sbtRecordOffset,
    uint                     sbtRecordStride,
    uint                     missIndex,
    vec3                     origin,
    float                    tmin,
    vec3                     direction,
    float                    tmax,
    int                      payload
);
        
```

In Vulkan terms:
gl_WorldRayOrigin = (x₀,y₀,z₀)
gl_Hit = t
gl_WorldRayDirection = (dx,dy,dz)

Oregon State University Computer Graphics mjb - February 1, 2022

13

Intersection Shader

```

hitAttribute vec3 attribs
void main()
{
    SpherePrimitive sph = spheres[ gl_PrimitiveID ];
    vec3 orig = gl_WorldRayOrigin;
    vec3 dir = normalize( gl_WorldRayDirection );
    ...
    float discr = b*b - 4.*a*c;
    if( discr < 0. )
        return;

    float tmp = ( -b - sqrt(discr) ) / ( 2.*a );
    if( gl_RayTmin < tmp && tmp < gl_RayTmax )
    {
        vec3 p = orig + tmp * dir;
        attribs = p;
        reportIntersection( tmp, 0 );
        return;
    }
    tmp = ( -b + sqrt(discr) ) / ( 2.*a );
    if( gl_RayTmin < tmp && tmp < gl_RayTmax )
    {
        vec3 p = orig + tmp * dir;
        attribs = p;
        reportIntersection( tmp, 0 );
        return;
    }
}
        
```

Intersect a ray with an arbitrary 3D object.
 Passes data to the Any Hit shader.
 There is a built-in ray-triangle Intersection Shader.

The diagram illustrates the intersection of a ray with a sphere. A yellow sphere is shown with a red ray passing through it. Red arrows point from the code snippets in the 'main' function to the corresponding parts of the sphere and ray. Below this, a flowchart shows the ray-tracing pipeline: 'Ray Generation Shader (rgen)' calls 'trace()', which leads to 'Traversing the Acceleration Structures'. This block then branches based on 'Any hits found for this ray?'. If 'No', it goes to 'Miss Shader (miss)'. If 'Yes', it goes to 'Closest Hit Shader (rhit)'. The 'Closest Hit Shader (rhit)' then calls the 'Intersection Shader (rint)', which in turn calls the 'Any Hit Shader (rahit)'. The 'Any Hit Shader (rahit)' also receives input from 'Traversing the Acceleration Structures'.

Oregon State University Computer Graphics mjb - February 1, 2022

14

Miss Shader

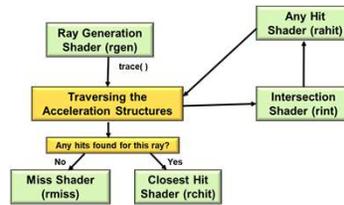
15

Handle a ray that doesn't hit *any* objects

```

rayPayload myPayload
{
    vec4 color;
};

void
main( )
{
    color = vec4( 0., 0., 0., 1. );
}
    
```



mjb - February 1, 2022

15

Any Hit Shader

16

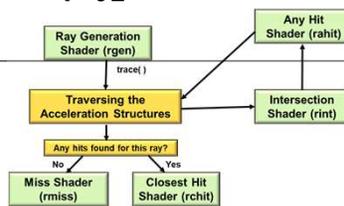
Handle a ray that hits *anything*.
Store information on each hit.
Can reject a hit.

```

layout( binding = 4, set = 0) buffer outputProperties
{
    float outputValues[ ];
} outputData;

layout(location = 0) rayPayloadIn uint outputId;
layout(location = 1) rayPayloadIn uint hitCounter;
hitAttribute vec3 attribs;

void
main( )
{
    outputData.outputValues[ outputId + hitCounter ] = gl_PrimitiveID;
    hitCounter = hitCounter + 1;
}
    
```



mjb - February 1, 2022

16

Closest Hit Shader

17

Handle the intersection closest to the viewer.
Collects data from the Any Hit shader.
Can spawn more rays.

```

rayPayload myPayload
{
    vec4 color;
};

void main()
{
    vec3 stp = gl_WorldRayOrigin + gl_Hit * gl_WorldRayDirection;
    color = texture( MaterialUnit, stp ); // material properties lookup
}
    
```

In Vulkan terms:
gl_WorldRayOrigin = (x₀,y₀,z₀)
gl_Hit = t
gl_WorldRayDirection = (dx,dy,dz)

```

graph TD
    rgen[Ray Generation Shader (rgen)] -- traceNV() --> traver[Traversing the Acceleration Structures]
    traver --> rint[Intersection Shader (rint)]
    traver --> q[Any hits found for this ray?]
    q -- No --> rmiss[Miss Shader (rmiss)]
    q -- Yes --> rchit[Closest Hit Shader (rchit)]
    rchit --> rahit[Any Hit Shader (rahit)]
    rint --> rahit
    
```

Oregon State University Computer Graphics mjb - February 1, 2022

17

Other New Built-in Functions

18

void **terminateRay**();

void **ignoreIntersection**();

void **reportIntersection**(float hit, uint hitKind);

} Loosely equivalent to "discard"

Oregon State University Computer Graphics mjb - February 1, 2022

18

Ray Trace Pipeline Data Structure

19

```

VkPipeline          RaytracePipeline;
VkPipelineLayout   PipelineLayout;

VkPipelineLayoutCreateInfo
  vplci.sType       = VK_STRUCTURE_TYPE_PIPELINE_LAYOUT_CREATE_INFO;
  vplci.pNext       = nullptr;
  vplci.flags       = 0;
  vplci.setLayoutCount = 1;
  vplci.pSetLayouts = &descriptorSetLayout;
  vplci.pushConstantRangeCount = 0;
  vplci.pushConstantRanges = nullptr;

result = vkCreatePipelineLayout( LogicalDevice, IN &vplci, nullptr, OUT &PipelineLayout);

VkRayTracingPipelineCreateInfo
  vrtpci.sType       = VK_STRUCTURE_TYPE_RAY_TRACING_PIPELINE_CREATE_INFO;
  vrtpci.pNext       = nullptr;
  vrtpci.flags       = 0;
  vrtpci.stageCount = << # of shader stages in the ray-trace pipeline >>
  vrtpci.pStages     = << what those shader stages are >>
  vrtpci.groupCount  = << # of shader groups >>
  vrtpci.pGroups     = << pointer to the groups (a group is a combination of shader programs >>
  vrtpci.maxRecursionDepth = << how many recursion layers deep the ray tracing is allowed to go >>;
  vrtpci.layout      = PipelineLayout;
  vrtpci.basePipelineHandle = VK_NULL_HANDLE;
  vrtpci.basePipelineIndex = 0;

result = vkCreateRayTracingPipelines( LogicalDevice, PALLOCATOR, 1, IN &vrtpci, nullptr, OUT &RaytracePipeline);
    
```



Oregon State
University
Computer Graphics

mjb - February 1, 2022

19

The Trigger comes from the Command Buffer: vkCmdBindPipeline() and vkCmdTraceRays()

20

```

vkCmdBindPipeline( CommandBuffer, VK_PIPELINE_BIND_POINT_RAYTRACING, RaytracePipeline );

vkCmdTraceRays( CommandBuffer,
  raygenShaderBindingTableBuffer, raygenShaderBindingOffset,
  missShaderBindingTableBuffer, missShaderBindingOffset, missShaderBindingStride,
  hitShaderBindingTableBuffer, hitShaderBindingOffset, hitShaderBindingStride,
  callableShaderBindingTableBuffer, callableShaderBindingOffset, callableShaderBindingStride
  width, height, depth );,
    
```



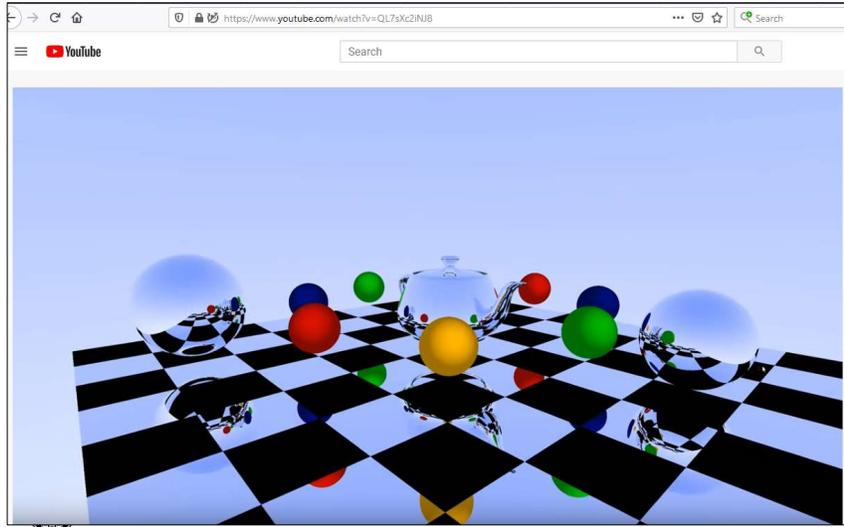
Oregon State
University
Computer Graphics

mjb - February 1, 2022

20

Check This Out!

21



Oregon State
University
Computer Graphics

<https://www.youtube.com/watch?v=QL7sXc2iNJ8>

mjb - February 1, 2022

21