

**Vulkan.**  
Ray-Tracing: Acceleration Structures

Oregon State University

Mike Bailey  
mjb@cs.oregonstate.edu

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License

Oregon State University Computer Graphics

AccelerationStructures.pptx  
nb - March 3, 2023

1

**Acceleration Structures**

- A Bottom-level Acceleration Structure (BLAS) reads the vertex data from vertex (and possibly index VkBuffers) to determine Axis-Aligned Bounding Boxes (AABBs).
- You can also supply your own AABB information to a BLAS.
- A single Top-level Acceleration Structure (TLAS) holds Instances, which are transformations and pointers to (potentially) multiple BLASes.
- Each BLAS is essentially used as a Model Coordinate bounding box, while the single TLAS is used as a World Coordinate bounding box.

Top Level Acceleration Structure  
Instance  
Instance  
Instance  
Instance  
Bottom Level Acceleration Structure  
Bottom Level Acceleration Structure  
Bottom Level Acceleration Structure  
Oregon State University Computer Graphics  
nb - March 3, 2023

2

**Creating the Bottom Level Acceleration Structures**

```

VkAccelerationStructure
    BottomLevelAccelerationStructure:
        vasgt; // VulkanAccelerationStructureGeometryTrianglesData
        vasgt.sType = VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_GEOMETRY_TRIANGLES_DATA;
        vasgt.pNext = nullptr;
        vasgt.vertexFormat = VK_FORMAT_VEC3; // = VK_FORMAT_R32G32B32_SFLOAT
        vasgt.vertexData.deviceAddress = MyVertexDataBuffer.vdm; // device address of the array of vertex structs
        vasgt.vertexData.size = sizeof(Vertex) * 3; // size of vertex data
        vasgt.indexType = VK_INDEX_TYPE_UINT_32; // we're not using index data
        vasgt.indexData = VK_NULL_HANDLE; // but if we were, they would be 32-bit unsigned in
        vasgt.transformData = 0; // no transform here

        vasgd; // VulkanAccelerationStructureGeometryData
        vasgd.sType = VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_GEOMETRY; // this is a union, not a struct
        vasgd.pNext = nullptr;
        vasgd.geometryType = VK_GEOMETRY_TYPE_TRIANGLES; // VK_GEOMETRY_TYPE_TRIANGLES or VK_GEOMETRY_TYPE_AABS or VK_GEOMETRY_TYPE_INSTANCE
        vasgd.geometry = vasg; // VulkanAccelerationStructureGeometry
        vasgd.flags = VK_GEOMETRY_OPAQUE_BIT; // members of VkGeometryFlagBits
        vasgd.dstAccelerationStructure = BottomLevelAccelerationStructure;
    
```

C:\

3

**Creating the Bottom Level Acceleration Structures**

```

VkAccelerationStructureBuildRangeInfo
    vasbri; // # triangles, # bounding boxes, or # instances
    vasbri.primitiveCount = sizeof(Vertices) / sizeof(Vertex) / 3;
    vasbri.primitiveOffset = 0;
    vasbri.firstVertex = 0;
    vasbri.transformOffset = 0;

VkAccelerationStructureBuildGeometryInfo
    vasbgi; // VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_BUILD_GEOMETRY_INFO
    vasbgi.sType = VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_TYPE_BOTTOM_LEVEL;
    vasbgi.pNext = nullptr;
    vasbgi.flags = VK_BUILD_ACCELERATION_STRUCTURE_PREFER_FAST_TRACE_BIT;
    vasbgi.mode = VK_BUILD_ACCELERATION_STRUCTURE_MODE_BUILD; // not re-building
    vasbgi.scratchData.deviceAddress = vkDeviceOrHostAddress; // ptr to array of VkAccelerationStructureGeometry structs
    vasbgi.geometryCount = 1;
    vasbgi.pGeometries = vasg; // VulkanAccelerationStructureGeometry
    vasbgi.pPGeometries = nullptr;
    vasbgi.dsAccelerationStructure = VK_NULL_HANDLE; // will be set later
    vasbgi.scratchData.deviceAddress = <<vkDeviceOrHostAddress >>; // will be set later

```

Top Level Acceleration Structure  
Instance  
Instance  
Instance  
Instance  
Bottom Level Acceleration Structure  
Bottom Level Acceleration Structure  
Bottom Level Acceleration Structure  
Oregon State University Computer Graphics  
nb - March 3, 2023

4

**Creating the Bottom Level Acceleration Structures**

```

VkAccelerationStructureBuildSizesInfo
    vasbsi; // VulkanAccelerationStructureBuildSizesInfo
    vasbsi.sType = VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_BUILD_SIZES_INFO;

vkGetAccelerationStructureBuildSizes( LogicalDevice, VK_ACCELERATION_STRUCTURE_BUILD_TYPE_DEVICE,
    IN &vasbgi, IN &vasbri.primitiveCount, OUT &vasbsi );
//<< vasbsi.accelerationStructureSize is how big the buffer should be >>
//<< use VK_BUFFER_USAGE_[ACCELERATION_STRUCTURE_STORAGE | SHADER_DEVICE_ADDRESS | STORAGE_BUFFER]_BIT when creating that buffer >>

VkAccelerationStructureCreateInfo
    vasci; // VulkanAccelerationStructureCreateInfo
    vasci.sType = VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_CREATE_INFO;
    vasci.pNext = nullptr;
    vasci.createFlags = 0;
    vasci.buffer = MyBottomLevelBuffer.buffer; // where BLAS will be stored
    vasci.offset = 0;
    vasci.size = vasbsi.accelerationStructureSize;
    vasci.type = VK_ACCELERATION_STRUCTURE_TYPE_BOTTOM_LEVEL;
    vasci.deviceAddress = nullptr;

VkAccelerationStructure
    BottomLevelAccelerationStructure:
        result = vkCreateAccelerationStructure( LogicalDevice, IN &vasci, PALLOCATOR, OUT &BottomLevelAccelerationStructure );
        vasbgi.dsAccelerationStructure = BottomLevelAccelerationStructure;
        vasbgi.scratchData.deviceAddress = MyBottomLevelBuffer.vdm;

```

At this point, BottomLevelAccelerationStructure is just a handle. We need to call vkCmdBuildAccelerationStructure() to populate it.

5

**Building the Bottom Level Acceleration Structures**

```

VkAccelerationStructureBuildGeometryInfo
    vasbgi; // already created...
    ...
VkAccelerationStructureBuildRangeInfo
    vasbri; // already created...
    ...
vkCmdBuildAccelerationStructure( CommandBuffer, 1, IN &vasbgi, IN &vasbri );

```

The BLAS vkCmdBuildAccelerationStructure command must be submitted right away. It must complete before attempting to build a TLAS.

Top Level Acceleration Structure  
Instance  
Instance  
Instance  
Instance  
Bottom Level Acceleration Structure  
Bottom Level Acceleration Structure  
Bottom Level Acceleration Structure  
Oregon State University Computer Graphics  
nb - March 3, 2023

6

Submitting the BLAS `vkCmdBuildAccelerationStructure` Command

```

VkCommandBufferBeginInfo
    vcbi.sType = VK_STRUCTURE_TYPE_COMMAND_BUFFER_BEGIN_INFO;
    vcbi.pNext = nullptr;
    vcbi.flags = VK_COMMAND_BUFFER_USAGE_ONE_TIME_SUBMIT_BIT;
    vcbi.pInheritanceInfo = (VkCommandBufferInheritanceInfo *)nullptr;
result = vkBeginCommandBuffer( AccelerationStructureCommandBuffer, IN &vcbi );

vkCmdBuildAccelerationStructure( AccelerationStructureCommandBuffer, 1, IN &vasbgi, IN &vasbri );

result = vkEndCommandBuffer( AccelerationStructureCommandBuffer );

```

VkSubmitInfo

```

    vsi.sType = VK_STRUCTURE_TYPE_SUBMIT_INFO;
    vsi.pNext = nullptr;
    vsi.commandBufferCount = 1;
    vsi.pCommandBuffers = &AccelerationStructureCommandBuffer;
    vsi.waitSemaphoreCount = 0;
    vsi.pWaitSemaphores = (VkSemaphore *)nullptr;
    vsi.signalSemaphoreCount = 0;
    vsi.pSignalSemaphores = (VkSemaphore *)nullptr;
    vsi.pWaitDstStageMask = (VkPipelineStageFlags *)nullptr;
result = vkQueueSubmit( Queue, 1, IN &vsi, 0 );
result = vkQueueWaitIdle( Queue );

```

7

Creating the Top Level Acceleration Structure

```

VkAccelerationStructureDeviceAddressInfo
    vasdai.sType = VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_DEVICE_ADDRESS_INFO;
    vasdai.accelerationStructure = BottomLevelAccelerationStructure;

glm::mat4 instanceRotation = glm::rotate( glm::mat4(1.), rotAngle, axis );
vmtm;
vmtm.matrix = glm::mat3x4( instanceRotation );

```

VkAccelerationStructureInstanceInfo

```

    vasi.transformMatrix = vmtm.matrix;
    vasi.pNext = nullptr;
    vasi.shaderBindingTableRecordOffset = 0;
    vasi.flags = VK_GEOMETRY_INSTANCE_TRIANGLE_FACING_CULL_DISABLE_BIT;
    vasi.accelerationStructureReference = vkGetAccelerationStructureDeviceAddress( LogicalDevice, &vasdai );

```

VkAccelerationStructureBuildRangeInfo

```

    vasbri.primitiveCount = 1;
    vasbri.primitiveOffset = 0;
    vasbri.firstVertex = 0;
    vasbri.transformOffset = 0;

```

8

Creating the Top Level Acceleration Structure

```

VkAccelerationStructureGeometryInstancesData
    vasgid.sType = VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_GEOMETRY_INSTANCES_DATA;
    vasgid.pNext = nullptr;
    vasgid.instances.arrayOfPointers = VK_FALSE;
    vasgid.instances.data.deviceAddress = << VkDeviceOrHostAddress of TopLevelAccelerationStructure >>;

```

VkAccelerationStructureBuildSizesInfo

```

    vasbsi.sType = VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_BUILD_SIZES_INFO;
    Vasbsi.pNext = nullptr;
    Vasbsi.buffer = << where TLAS will be stored >>;
    Vasbsi.primitiveCount = OUT &vasbri.primitiveCount;
    Vasbsi.size = OUT &vasbri.size;

```

VkAccelerationStructureGeometry

```

    vasg.sType = VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_GEOMETRY;
    vasg.geometryType = VK_GEOMETRY_TYPE_INSTANCES;
    vasg.geometry.instances = vasgid;

```

VkAccelerationStructureInfo

```

    vasi.sType = VK_ACCELERATION_STRUCTURE_TYPE_TOP_LEVEL;
    vasi.pNext = nullptr;
    vasi.flags = 0;
    vasi.instanceCount = 0;
    vasi.geometryCount = 1;
    vasi.pGeometries = &vasg;

```

Oregon State University Computer Graphics

mb - March 3, 2023

9

Creating the Top Level Acceleration Structure

```

VkAccelerationStructureCreateInfo
    vasci.sType = VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_CREATE_INFO;
    vasci.pNext = nullptr;
    vasci.createFlags = ???;
    vasci.buffer = << where TLAS will be stored >>;
    vasci.offset = 0;
    vasci.size = vasbri.accelerationStructureSize;
    vasci.type = VK_ACCELERATION_STRUCTURE_TYPE_TOP_LEVEL;
    vasci.deviceAddress = nullptr;

```

VkAccelerationStructure

```

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

```

vasi.dstAccelerationStructure = TopLevelAccelerationStructure;

Create scratch buffer: buildSize, buildScratchSize, VK\_BUFFER\_USAGE\_STORAGE\_BUFFER\_BIT  
vasi.scratchData.deviceAddress << device address of tlas scratch buffer handle >>

At this point, TopLevelAccelerationStructure is just a handle. We need to call `vkCmdBuildAccelerationStructure()` to populate it.

10

Building the Top Level Acceleration Structure

```

...
VkBuidAccelerationStructureBuildGeometryInfo      vasbgi;
...
VkBuidAccelerationStructureBuildRangeInfo        vasbri;
...
vkCmdBuildAccelerationStructure( CommandBuffer, 1, IN &vasbgi, IN &vasbri );

```

11

Submitting the vkCmdBuildAccelerationStructure

```

VkCommandBufferBeginInfo
    vcbi.sType = VK_STRUCTURE_TYPE_COMMAND_BUFFER_BEGIN_INFO;
    vcbi.pNext = nullptr;
    vcbi.flags = VK_COMMAND_BUFFER_USAGE_ONE_TIME_SUBMIT_BIT;
    vcbi.pInheritanceInfo = (VkCommandBufferInheritanceInfo *)nullptr;
result = vkBeginCommandBuffer( AccelerationStructureCommandBuffer, IN &vcbi );

vkCmdBuildAccelerationStructure( AccelerationStructureCommandBuffer, 1, IN &vasbgi, IN &vasbri );

result = vkEndCommandBuffer( AccelerationStructureCommandBuffer );

```

VkSubmitInfo

```

    vsi.sType = VK_STRUCTURE_TYPE_SUBMIT_INFO;
    vsi.pNext = nullptr;
    vsi.commandBufferCount = 1;
    vsi.pCommandBuffers = &AccelerationStructureCommandBuffer;
    vsi.waitSemaphoreCount = 0;
    vsi.pWaitSemaphores = (VkSemaphore *)nullptr;
    vsi.signalSemaphoreCount = 0;
    vsi.pSignalSemaphores = (VkSemaphore *)nullptr;
    vsi.pWaitDstStageMask = (VkPipelineStageFlags *)nullptr;
result = vkQueueSubmit( Queue, 1, IN &vsi, 0 );
result = vkQueueWaitIdle( Queue );

```

12

Other Information for Creating the Top Level Acceleration Structure 13

```

VkAccelerationStructureGeometryAabbsData          vasgad;
vasgad.sType = VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_GEOMETRY_AABBS_DATA;
vasgad.pNext = nullptr;
vasgad.data = << VkDeviceOrHostAddressConst >>;
vasgad.stride = 0;

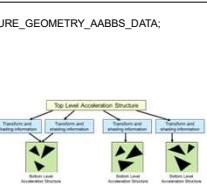
VkAccelerationStructureInstance                  vasi;
vasi.transform = << VkTransformMatrix >>;
vasi.instanceCustomIndex = << uint32_t:24 >>;
vasi.mask = 0xff
instanceShaderBindingTableRecordOffset = << uint32_t:24 >>;
vasi.flags = 0;
vasi.accelerationStructureReference = << uint64_t >>;

VkAabbPositions                                vap;
vap.minX, .minY, .minZ;
vap.maxX, .maxY, .maxZ;

VkTransformMatrix                               vtm;
vtm.matrix = float [3][4];                      // glm::mat3x4

```

Oregon State University Computer Graphics



Why a 3x4 Matrix? 14

glm::mat4

$$\begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdot \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}$$

glm::mat3x4

$$\begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdot \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}$$

Because we are not doing perspective here, we really don't need the bottom row

```

glm::mat4 mat = glm::mat4( 1. );
mat = glm::rotate( mat, rotAngle, zaxis );
vtm.matrix = glm::mat3x4( mat );

```

Oregon State University Computer Graphics

13

14