



SQUARE ENIX.

"SHADOWS" OF THE TOMB RAIDER -A RAY TRACING DEEP DIVE

Holger Gruen (NVIDIA), Jon Story (NVIDIA), Michiel Roza (Nixxes) 03/19/2019

www.nvidia.com/GTC

AGENDA

Shadow of the Tomb Raider Why ray traced shadows? DXR shaders (ray generation) GameWorks spatial denoiser DXR acceleration structure Integration in render pipeline Results Future work



AGENDA

Shadow of the Tomb Raider

Why ray traced shadows? DXR shaders (ray generation) GameWorks spatial denoiser DXR acceleration structure Integration in render pipeline

Results

Future work



SHADOW OF THE TOMBRADER

AGENDA

Shadow of the Tomb Raider Why ray traced shadows? DXR shaders (ray generation) GameWorks spatial denoiser DXR acceleration structure Integration in render pipeline Results Future work





Shadow Mapped





WHY RAY TRACED SHADOWS?

There's so much that shadow mapping can't do!

- Pixel perfect shadows
- Translucent shadows
- Point lights
 - Currently faked by using two spot lights
- Area lights





AGENDA

Shadow of the Tomb Raider Why ray traced shadows? DXR shaders (ray generation) GameWorks spatial denoiser DXR acceleration structure Integration in render pipeline Results Future work



DXR SHADERS Mini Agenda

- Noise / random number generation
- Ray generation
- Hit shaders
- Adaptive raytracing
- Translucency
- TAA and jittering



DXR Mini intro





Mini Agenda

- Noise / random number generation
- Ray generation
- Hit shaders
- Adaptive raytracing
- Translucency
- TAA and jittering



Noise / random number generation



Ideally we want to trace <u>many</u> rays to find out how much of the light source a point can see

Full light Penumbra Umbra



- For great performance we want to shoot only 1 ray per pixel
- So instead of one pixel shooting many rays, a neighborhood of pixels samples 'enough' random positions on the light source





- For great performance we want to shoot only 1 ray per pixel
- So instead of one pixel shooting many rays, a neighborhood of pixels samples 'enough' random positions on the light source





- For great performance we want to shoot only 1 ray per pixel
- So instead of one pixel shooting many rays, a neighborhood of pixels samples 'enough' random positions on the light source





- For great performance we want to shoot only 1 ray per pixel
- So instead of one pixel shooting many rays, a neighborhood of pixels samples 'enough' random positions on the light source





- For great performance we want to shoot only 1 ray per pixel
- So instead of one pixel shooting many rays, a neighborhood of pixels samples 'enough' random positions on the light source





- For great performance we want to shoot only 1 ray per pixel
- So instead of one pixel shooting many rays, a neighborhood of pixels samples 'enough' random positions on the light source





- For great performance we want to shoot only 1 ray per pixel
- So instead of one pixel shooting many rays, a neighborhood of pixels samples 'enough' random positions on the light source





- For great performance we want to shoot only 1 ray per pixel
- So instead of one pixel shooting many rays, a neighborhood of pixels samples 'enough' random positions on the light source





- For great performance we want to shoot only 1 ray per pixel
- So instead of one pixel shooting many rays, a neighborhood of pixels samples 'enough' random positions on the light source





DXR SHADERS Noise / random number generation

• Random positions on the light are based on generating pseudo random numbers

http://www.reedbeta.com/blog/quick-and-easy-gpu-random-numbers-in-d3d11/

- The trick is to choose the right random seed for the generator
- We use a seed that is based on the 2D position of the pixel

Seed((pixel_2d_pos) % TILE_SIZE_2D);



Noisy shadows





DXR SHADERS Mini Agenda

- Noise / random number generation
- Ray generation
- Hit shaders
- Adaptive raytracing
- Translucency
- TAA and jittering



DXR SHADERS Ray generation

We use specialized raygen shaders for each light type for optimal performance









Directional light source with an angular extent

Area Cone light

Point light with a spherical area

Rectangular area cone light

NVIDIA



Move some small distance along the normal to prevent self-shadowing!

Ray generation



All light types



No rays for pixels that face away from the current light!



Ray generation





Ray generation









DXR SHADERS Mini Agenda

- Noise / random number generation
- Ray generation
- Hit shaders
- Adaptive raytracing
- Translucency
- TAA and jittering


DXR SHADERS Hit shaders

- We don't use the closest hit along the ray
- Instead we use the first opaque one that gets reported to anyhit()
 - Translucency for DXR ultra quality is an exception
- We fetch texture coords for alpha tested prims to carry out the alpha test



DXR SHADERS Hit Shaders

{

• Opaque Geometry

```
void OpaqueClosestHit(...)
{
```

```
payload.hitT = RayTCurrent();
payload.visibility = 0.0f;
```

```
• Alpha-tested Geometry
```

```
void AlphaClosestHit(...)
{
    payload.hitT = RayTCurrent();
}
```

```
void OpaqueAnyHit(...)
```

```
AcceptHitAndEndSearch();
```

```
void AlphaAnyHit(...)
```

```
float alpha = GetHitAlpha(bary);
```

```
if( alpha < g_fAlphaThreshold )
        IgnoreHit();
else {</pre>
```

```
payload.visibility = 0.0f;
AcceptHitAndEndSearch();
```



DXR SHADERS Mini Agenda

- Noise / random number generation
- Ray generation
- Hit shaders
- Adaptive raytracing
- Translucency
- TAA and jittering



DXR SHADERS Adaptive raytracing

- Adaptive raytracing is only used for ultra DXR quality
- In this mode we cast more than 1 ray for some pixels
- Let's dive into the details ...



Adaptive raytracing

Phase 1 - Cast 1 ray per pixel

Phase 2 - Cast up to 2 more rays ,where necessary'



Neighborhood of pixel

Visibility:

- Black => can't see the light
- White => can see the light

hitT:

- Distance to blocker along ray
- White means 'infinity'
- Dark means close-by blocker

Adaptive raytracing

- Phase 1 Cast 1 ray per pixel
- Phase 2 Cast up to 2 more rays ,where necessary'

Visibility



hitT





Adaptive raytracing

Phase 1 - Cast 1 ray per pixel

Phase 2 - Cast up to 2 more rays ,where necessary'



Visibility





Adaptive raytracing

Phase 1 - Cast 1 ray per pixel

Phase 2 - Cast up to 2 more rays ,where necessary'



Visibility

hitT



Visibility = (Visibility0+Visibility1)/2

Adaptive raytracing

- Phase 1 Cast 1 ray per pixel
- Phase 2 Cast up to 2 more rays ,where necessary'

Visibility







Visibility = (Visibility0+Visibility1+Visibility2)/3

DXR SHADERS Mini Agenda

- Noise / random number generation
- Ray generation
- Hit shaders
- Adaptive raytracing
- Translucency
- TAA and jittering



DXR SHADERS Translucency

- Ultra DXR quality also features translucent shadows
- We support up to 3 layers of translucency
 - Mainly to keep performance at acceptable levels
- Translucency should be straightforward with DXR right?
 - Want to keep on using anyhit() instead of iterated closesthit()
 - Let's look at the details ...















DXR SHADERS Translucency

```
void TranslucentAnyHit(...)
{
      float alpha = GetHitAlpha(bary, PrimID);
      if( alpha >= g fAlphaThreshold )
            payload.visibility -= ( 1.0f / 3.0f );
      if( payload.visibility < 0.01f ) {</pre>
            payload.visibility = 0.0f;
            AcceptHitAndEndSearch();
      else
             IqnoreHit();
```



Opaque raytraced shadow

Translucent raytraced shadows

DXR SHADERS Mini Agenda

- Noise / random number generation
- Ray generation
- Hit shaders
- Adaptive raytracing
- Translucency
- TAA and jittering



DXR SHADERS TAA + Jittering

• Like many games SotTR uses jittered TAA

• Each frame adds a 'random' subpixel offset to all geometry

• Surprisingly this creates problems with flickering shadows!



TAA + Jittering



a 2x2 pixel grid

- The red dots are the pixel centers
- This is where rasterized geometry is sampled



TAA + Jittering





The intermediate positions and the grid are shown to help understand how 3D positions change across the quad



TAA + Jittering



Jitter somewhat ...



DXR SHADERS TAA + Jittering

- Jittering changes the WS position that is sampled at pixel centers
 - It also changes the depth values at the pixel centers
- Jittering changes the reconstructed world space positions



Shadow ray origins jitter as well



DXR SHADERS TAA + Jittering

Jittered ray positions are not problematic in general, but:

- We typically shoot only one ray per pixel
 - Which is equivalent to 'point sampling' of the visibility signal
- Large areas of flat ground are problematic
 - Vertical jittering leads to large differences in WS positions
 - Also visible with shadow maps but less because of SM filtering







Solutions:

- 1. Currently we render an extra depth pass without jittering
 - Use non-jittered depth to reconstruct WS ray origins
- 2. Future: Render ddx/ddy(1/z_buffer_depth) with depth pass
 - Reconstruct non-jittered depth
 - Use non-jittered depth to construct WS ray origins
 - 1/z-buffer-depth is linear in screen space



AGENDA

Shadow of the Tomb Raider Why ray traced shadows? DXR shaders (ray generation) GameWorks spatial denoiser DXR acceleration structure Integration in render pipeline Results Future work



INPUT / OUTPUT



ISOTROPIC KERNEL



5 NIXXES 📀 NVIDIA.

ANISOTROPIC KERNEL



OVERLAPPING PENUMBRA #1



nixxes

OVERLAPPING PENUMBRA #2



nixxes



BLEEDING ARTIFACTS





CUSTOMIZED BOUNDARY DETECTION





COULD WE DO LESS WORK?

200 FPS (5.02 ms/frame)	
Keyboard Shortcuts (?)	
V Global Controls	
462.864 Time	
1.000 Time Scale	
Reset Faue Stor	
Terror Canture Video Canture	
Enable FXAA	
GFSDK_ShadowLib Technology Type	
Raster Techniques	
RT Denoiser Technques	
V Display type	
Display Scene	
Display Ray Trace Buffers	
1 Lightindex	
V Spot Light 1	
-3.645 2.243 -1.488 World Position	
0.220 -0.975 0.000 Direction	
1.413 Opening Angle	
8:255 G:198 8:102 Color	
19.686 Intensity	
V. Shadow Type	
RT Denoise A Technique	
Ray Trace	
Ground Truth Mode (196 Rays Per Pixel)	
4 Rays Per Pixel	
7. Kernal Radius	
Half Res Filter	
Aniso Filter	
0.001 Depth Epsilon Min	
0.019 Depth Epsilon Percent	
0.001 Depth Epsilon Planar	
Debug Display: Depth Boundaries	
Debug Display: Penumbra Mask	
0.001 Hit Epsilon Min	
200.000 Hit Epsilon Range	
1.179 Hit Epsilon Power.	
0.070 Spot Light Source Radius	
And the second s	
The set of	
and the state of the	
a state and a second with the	
FILL HAR BOLLEN THE THE	
and a second second	
ALL	



PENUMBRA MASK




IMPORTANT FEATURES

- Half resolution denoising
 - Drastically improves performance
 - SOTTR uses this mode for ALL light types
- MSAA input Depth & Normal buffers supported
 - Still only requires single sample Visibility & HitT buffers
 - Produces MSAA shadow mask
- Sub-viewports supported for local light sources
 - Just need to figure out screen area affected



AGENDA

Shadow of the Tomb Raider Why ray traced shadows? DXR shaders (ray generation) GameWorks spatial denoiser **DXR** acceleration structure Integration in render pipeline Results Future work



BLAS "a mesh"

- Vertex and index buffers for each geometry
- Straightforward for static geometry

```
struct D3D12_RAYTRACING_GEOMETRY_TRIANGLES_DESC
{
    DXGI_FORMAT IndexFormat;
    DXGI_FORMAT VertexFormat;
    UINT IndexCount;
    UINT VertexCount;
    D3D12_GPU_VIRTUAL_ADDRESS IndexBuffer;
    D3D12_GPU_VIRTUAL_ADDRESS VertexBuffer;
}
```

BLAS

What about skinned objects and vertex animations?

- Each vertex needs to be fully transformed!
- Foundation Engine uses shader graphs
- Added a shader permutation in VS template for exporting a transformed vertex buffer
- Run a pass for all dynamic objects before building

Skinning gone wrong Inner demon ©

BLAS Lara's hair

- PureHair, an evolution of TressFX
 - Simulates control points
 - Renders strands of hair as camera facing quads
- Everything needs to be actual geometry in the AS
 - Make the simplest cylinder possible for every strand









BLAS Rebuild/refit strategy

- Two modes of updating dynamic BLASes in DXR:
 - Rebuild, essentially "replacing" the old one (~100M tris/sec)
 - Refit, for "small" model changes (~1000M tris/sec, 10x as fast!)
 - Catch: ray trace performance might degrade!
 - Top refitting throughput only for large enough workloads
- We chose to always refit BLAS unless # vertices change



TLAS "A scene"

- Static BLASes can be instanced
- Always rebuilding TLAS seems to be fast enough (<1ms)

```
struct
D3D12_BUILD_RAYTRACING_ACCELERATION_STRUCTURE_INPUTS
{
    UINT NumDescs;
    D3D12_GPU_VIRTUAL_ADDRESS InstanceDescs;
}
```

ACCELERATION STRUCTURE

About LODs of meshes

- Every LOD level is stored in a separate BLAS
- Using LOD 0 for everything caused self-shadowing artifacts!
- Just use the same LOD we use for rendering
- What about LOD fading?
 - Use "most visible" LOD





AGENDA

Shadow of the Tomb Raider Why ray traced shadows? DXR shaders (ray generation) GameWorks spatial denoiser DXR acceleration structure Integration in render pipeline Results Future work



Forward+ renderer



Now with ray traced shadows!



Now with ray tracing!





Async compute



WHY DO WE STILL NEED SHADOW MAPPING?

Translucent rendering

- Translucent rendering has no depth write
- Can't use shadow resolve pass!
- We cannot shoot rays from pixel shaders



WHY DO WE STILL NEED SHADOW MAPPING? Performance!

- ► Updating entire scene full of dynamic objects costs up to 20ms of BLAS refits ⊗
- AS culling using existing shadow map culling



WHY DO WE STILL NEED SHADOW MAPPING? Performance!

- For directional lights:
 - Replace only nearest cascades with ray traced shadows
- For local lights:
 - Distance based fade to shadow map
- How do we choose these distances?



ARTIST TOOLS Let lighting artists decide!

E-Ray Tracing	· · · · · · · · · · · · · · · ·
Light is active when Raytracing	
Light is active when Shadowmapping	
Enable Raytraced Shadows	Force raytraced
Light Source Radius	5
Override near plane	
Near Plane	100
Sun Cone Angle	30
Raytracing light priority	50
RT Fade Start Distance	200
RT Cull (fade out) Distance	300
Shadow Cull From Pivot	
Resolution mode for this light	Full Res
- Area Light	
Area Light Width	50
Area Light Height	50
Area Light Up Vector	0 1 0



AGENDA

Shadow of the Tomb Raider Why ray traced shadows? DXR shaders (ray generation) GameWorks spatial denoiser DXR acceleration structure Integration in render pipeline Results Future work



No point light shadows

C.L

A

Raytraced point light shadows

Shadow mapped area light shadows

Raytraced area light shadows

Shadow mapped sun light shadows

Raytraced sun light-shadows

AGENDA

Shadow of the Tomb Raider Why ray traced shadows? DXR shaders (ray generation) GameWorks spatial denoiser DXR acceleration structure Integration in render pipeline Results Future work



FUTURE WORK

- Reconstruct non-jittered depth
- Ray traced shadows on translucent geometry
- Tessellation
- Content authoring with ray tracing in mind
- Use vertex transform pass for rasterization as well
- ► GI / Reflections / AO / ... ?







C R Y S T A L DYNAMICS

QUESTIONS

Holger Gruen (hgruen@nvidia.com)

Jon Story (jons@nvidia.com)



www.nvidia.com/GTC