

PERFORMANT REFLECTIVE BEAUTY: HYBRID RAY TRACED REFLECTIONS IN FAR CRY 6

STEPHANIE BRENHAM, UBISOFT TORONTO

IHOR SZLACHTYCZ, AMD

GDC



REFLECTIONS BACK ON THE UBISOFT TEAM

Stephanie Brenham

- 3D Team Lead Programmer
- Far Cry 6
- The Games Awards Future Class 2021
- Maya & mental ray for Maya

Anton Remezenko

- Senior 3D programmer
- Far Cry 6
- Far Cry 5

Aleksei Shevchenko – 3D programmer Mikhail Shostak – 3D programmer

AMD PUBLIC





FAR CRY 6

- FPS set in a fictional tropical island
- Humid tropical setting
 - See "Simulating Tropical Weather in Far Cry 6"
- Lots of wet and highly reflective surfaces







REFLECTIONS IN FAR CRY 6



- 1. Generate Ray Buffer
- 2. SSLR
- 3. HW Ray trace
- 4. Particles HW Ray trace
- 5. Integrate results

*SSLR based on Tomasz Stachowiak "Stochastic Screen Space Reflections" at SIGGRAPH 2015



IMPORTANCE SAMPLING





SOLUTION: REUSE NEIGHBOURS





RAY BUFFER – TO BE USED LATER



GPUOpen

SCREEN SPACE LOCAL REFLECTIONS

Generate Rays

SSLR

Ray Trace & Lighting

Particles Trace

Integrate



- SS trace faster than HW ray tracing
- SSLR is fallback while BVH builds
- Green shows where we use SSLR

Problem:

SSLR found to be unstable with the limited trace steps



SOLUTION: LINEAR TRACE WITH REFINE





- Replaced Hierarchical Z with Linear Trace
- 64 large steps, 8 refine steps
- Stable with camera rotation
- Same performance as Hierarchical Z

Problem: Edges were being missed



SSLR LINEAR TRACE EDGE CASES





SOLUTION: RANDOM RAY OFFSETS





GPUOpen

SS TILE CLASSIFICATIONS

Generate RaysSSLRRay Trace & LightingParticles TraceIntegrateImage: Solar of the surface of the surfa

SS FAST CONE TRACE



GPUOpen

HW RAY TRACE ACCELERATION STRUCTURES





HW RAY TRACING ONLY IN NECESSARY PLACES

Generate Rays

SSLR

Ray Trace & Lighting



Integrate



- Far distances: Only SSLR
 - BVH stops far from player
- Close distances: Only HW RT
 - Likely to be offscreen reflections
- Mid distances: it depends...

REFLECTING WITH CONFIDENCE: SSLR





REFLECTING WITH CONFIDENCE: HW RAY TRACING



REFLECTING WITH CONFIDENCE: HYBRID

Generate Rays

SSLR

Ray Trace & Lighting

Particles Trace

Integrate



See "Simulating Tropical Weather in Far Cry 6" for more details on environment maps

- Green is full confidence SSLR
 - Use only SSLR
- Teal is mid confidence SSLR
 - Hybrid: use trace with highest confidence
- Blue is no confidence SSLR
 - Use only HW RT
- Other reflections: no confidence in SSLR nor HW RT
 - Use environment map reflections



PEEK BEHIND THE CURTAIN

AVISO

RAYTRACE REFLECTION GRAPH



BVH TREE



- One BLAS per Object (or Material)
 - Tip: Per object is preferred
- Static BLASes are built once
- Skinned BLASes need endless rebuild loop
- First person player character rebuilds each frame
 - Player weapon & attachments as well
- 16 BLASes rebuild per frame (static/skinned)
 - Max 12 skinned



HW RAY TRACING & LIGHTING TOGETHER



|--|

- Lighting shader and ray trace shader in one
- Parameters for lighting unified and stored per vertex
- Each material has specific shader to map its lighting to the unified lighting parameters
- Object material parsed on BLAS creation stage
 - Performance hit of shader variations is a one-time cost during BLAS creation



LIGHTING IN THE HW RAY TRACING PIPELINE

Generate Rays

SSLR

Ray Trace & Lighting

Particles Trace

Integrate



- **Graphics pipeline**: interpolated values provided by the Rasterizer to Pixel Shaders
- **Ray tracing pipeline**: interpolation done in the lighting shader
 - Unified lighting parameters stored in DXR primitives



PRIMITIVE DATA

Generate Rays

SSLR

Ray Trace & Lighting

Particles Trace

Integrate

DXR Prim, 24byte

DXR Vertex 0

- albedo, RGB8
- smoothness, A8
- normal(oct), RG8
- metallic, B8
- reflectance, A8

DXR Vertex1

DXR Vertex 2

- Unwrap Vertex and Index buffers to Primitives before BLAS created
 - Causes duplication but ensures better coherency for interpolation
- Index buffer not needed for lighting computations
 - Removes an additional indirection
- Vertex data is packed



PRIMITIVE POOL

Generate	Rays
----------	------

SSLR

Ray Trace & Lighting

Particles Trace

Integrate

- Primitive Buffer holds all primitive information from every BLAS
- Primitive Offset Buffer redirection to primitives for BLAS instances in TLAS
 - Required to locate primitive data using the PrimitiveIndex supplied by DXR
- Fixed Geometry count per BLAS max 32



LIGHTING THE PRIMITIVE



HW RAY TRACE BARYCENTRIC COORDINATES

HW RAY TRACE PRIMITIVES LIGHTING

PARTICLES IN FAR CRY 6

Generate Rays

۲ ٠ ۲

SSLR

- Far Cry games are known for
 - explosions

Particles Trace

- dynamic systemic fire
- New for Far Cry 6: Poison
- Particle reflections important for Far Cry 6

Integrate

Problem: Particles are just 2D billboards

Ray Trace & Lighting

PARTICLES HARDWARE RAY TRACE





PARTICLES: THE CASE OF THE MISSING SPRITES



SOLUTION: EXTRA QUADS



SOLUTION: MASK QUADS BY VIEW SPACE AXIS



SSLR PARTICLES

RAY TRACED PARTICLES





GPUOpen

INTEGRATE


REFLECTION MOTION



REFLECTION MOTION : REPROJECTION



REFLECTION MOTION: PREVIOUS FRAME



```
[AMD Official Use Only]
```

REFLECTION MOTION: BACK IN SCREEN SPACE



TEMPORAL ACCUMULATION







FI

SSLR/HW RAY TRACE MASK

HYBRID RAY TRACED REFLECTIONS



OPTIMIZING RAYTRACING IN FAR CRY 6

GDC



REFLECTIONS AT THE AMD TEAM



- Ihor Szlachtycz (presenter)
 - AMD GPU Dev Tech for Far Cry 6
- Zhuo Chen
 - AMD GPU Dev Tech for Far Cry 6
- John Hartwig
 - AMD CPU Dev Tech for Far Cry 6



OPTIMIZATION OVERVIEW

- We will be going over how Hybrid Reflections raytracing was optimized for Far Cry 6
 - Shader table management
 - Hit shader design
 - BVH building
- All performance captures for Far Cry 6 were done on a Radeon 6800XT GPU with a Ryzen 3970X Threadripper
- But first, we will give a quick overview of Radeon GPU Profiler (RGP)

















GENERAL RAYTRACING PERFORMANCE

• Here we have an overview of what a RGP trace looks like with Hybrid RayTraced Reflections





RAY TRACING DIVERGENCE

- Ray tracing is a latency heavy operation, very affected by divergence
- We look at below forms of divergence and how Far Cry 6 tackles them
 - Ray Divergence
 - Shader Table Divergence
 - Resource Divergence





RAY DIVERGENCE

- Different rays can hit different BLAS's with number of levels that the ray has to traverse
- Executing shader will have to traverse for the worst case



SHADER TABLE DIVERGENCE

- Occurs when different rays evaluate different shader table entries for hit/miss
- We have to evaluate all the hit materials in our thread group sequentially, like a branch



RESOURCE DIVERGENCE

- Occurs when a shader issues a resource read, but different threads in the wave can be accessing different resources
- We need to loop through all unique resources in our wave and have each issue the read for the threads using that resource
 - In HLSL, this is done via NonUniformResourceIndex



Textures[textureId].Sample



Texture.Sample
Read Texture



RAY TRACING HIT EXECUTION

```
if (ray intersects)
{
    uint hit_table_index = calc_hit_table_index();
    switch (hit_table_index)
    {
        case 0:
            ray_result = shader_table_entry_0(ray_hit);
            break;
        case 1:
            ray_result = shader_table_entry_1(ray_hit);
            break;
        ...
        case n:
        ray_result = shader_table_entry_n(ray_hit);
        break;
        case n:
        ray_result = shader_table_entry_n(ray_hit);
        ray_result = shader_table_entry_n(ray_hit);
        break;
        case n:
        ray_result = shader_table_entry_n(ray
```

- Small shader table:
 - Everything gets inlined and the code behaves like a series of if statements
- Large shader table:
 - Too many shaders to inline, compiler will emulate a function calling convention with loading/storing function arguments/return values in LDS
- hit_table_index does not have to be uniform for a thread group



SHADER TABLES (NON INLINED)

- Example of what you will see in RGP when your shader table isn't inlined
- Lots of LDS usage, instruction cache misses, wasted instructions for writing/loading function arguments
- On a Radeon[™] 6800XT at 4K, forcing the shader table to not be inlined is about 2-3x slower



Shader Table

Shader table

ISA

Information

(1) Allocated resources are driven by the shader table entries with the highest pressure on each resource, regardless of the call count.

Export name	Туре	VGPRs	Stack size (B)	LDS (B)	Scratch memory (B)	Total latency (clks)	Instruction cost (%)	Instruction hit count	Call count *	API shader hash	Internal pipel
RayGen	Ray generation	126	240	0	0	N/A	N/A	N/A	N/A	0x2193817871F6D4F9	0xD4BE70F6FF7
TraceRaysAmdInternal	Traversal	88	64	4,096	0	N/A	N/A	N/A	N/A	0xAAA1CCC20000000FBD5CBD0CA959ED5	0xF790F03A9669
ClosestHit	Closest hit	80	224	640	0	N/A	N/A	N/A	N/A	0x8C16B2C8D34C7892	0x4AA2A0E81E4
Miss	Miss	19	16	0	0	N/A	N/A	N/A	N/A	0xF78FE40F86CE1DC9	0x4CEF6A572CE
_amdgpu_cs_main	N/A	0	0	0	0	N/A	N/A	N/A	N/A	0x61DB175FE66864C9	0x301A80234BEI



v writelane b32 v24, s52, 21 v writelane b32 v24, s51, 22 v writelane b32 v24, s50, 23

GPUOpen

AMD PUBLIC

PERFORMANT REFLECTIVE BEAUTY IN FC 6

GDC - MARCH 2022 61

FAR CRY 6 SHADER TABLE

- Far Cry 6 has a single hit shader which interpolates vertex attributes that store all material properties
- Avoids lots of shader table divergence with only 1 hit and 1 empty miss shader
- Avoids resource divergence since all DXR Prims are in global buffers
 - No per material textures get accessed
- Ray divergence is partially mitigated by BVH building
 - We will get to this in next section

uint geometryOffset = instanceIndex * gs_DxrGeometriesPerInstance + geometryIndex; uint primOffset = RTPrimOffsetBuffer[geometryOffset];

DXRPrim prim = DXRPrimUnpack(RTPrimBuffer[primOffset + primitiveIndex]);
float3 color = COMPUTE_BARYCENTRIC_POINT(barycentrics, prim.vertices[0].color, prim.vertices[1].color, prim.vertices[2].color);
// ...
float reflectance = COMPUTE_BARYCENTRIC_POINT(barycentrics, prim.vertices[0].reflectance, prim.vertices[1].reflectance, prim.vertices[2].reflectance);

payload.hitColor = ApplyLighting(WorldRayOrigin(), WorldRayDirection(), RayTCurrent(), color, ... reflectance);



AMD RAYTRACING PERFORMANCE GUIDE

- You can find our performance recommendations at https://gpuopen.com/performance/
- We have raytracing and BVH recommendations in the Ray Tracing section



Optimizing a modern real-time renderer can be a somewhat daunting task. Explicit APIs hand you more control over how you craft your frame than ever before, allowing you to achieve higher frame rates, prettier pixels, and ultimately, better use of the hardware.

Our AMD RDNA[®] 2 Performance Guide will help guide you through the optimization process with a collection of tidbits, tips, and tricks which aim to support you in your performance quest.





BVH OVERVIEW

GDC



BLAS VERTEX GENERATION

- Far Cry 6 uses compute shader to generate vertices for BLAS building instead of a VS/GS/HS/DS
 - Can't do per vertex/object culling if done in a geometry pass
 - Avoids redundancy with vertex cache miss
- Compute shader in Far Cry 6 stores material data per vertex which are calculated via reading material textures of the simplified models at the vertices
- Vertices generated by CS for BLAS can be reused for shadows and potentially other passes

	VS Write	GS Write	HS/DS Write	CS Write
Culled Geometry	No	No	No	N/A
Vertex Cache Miss	Yes	Yes	Yes	No



BVH BUILDING

• Looking at the RGP trace again, you can see that BVH building is executed on async compute





ASYNC BVH BUILDING

- BLAS/TLAS generation doesn't usually depend on other parts of the frame until you start ray tracing
 - Can overlap with Gbuffer pass/shadows/early frame workloads
- Below is an image of non overlapped BVH generation (0.4-0.8ms slower):



GPUOpen

ASYNC BVH BUILDING

· · · · · · · · · · · · · · · · · · ·							•	_
Frame 1909	49 750 000 up	52 500 000 uc	Frame 1910	60,000,000 us	62 750 000 up	67 500 000 uc	•	l ow an
μο το,000.000 μο	το, / 50.000 μS	52,500.000 μs	50,250.000 µS	ου,ουο.ουο με	03,730.000 μs	07,500.000 μs		an
								needeo
Graphics queue	••••••••••••••••••••••••••••••••••••••							
								queue
								•
	ID3D12	Ç		100	ID3 ID3D12Queue::E			• On
	1D3D12Qu	e ID		ID3	D12 ID3D12Queue::E ID	03D12		011
	1D3D12Qu			1030	12Q ID3D12QueueE IL			qu
	ID3D12Qu	L ID ID3D12Queue::Ex		ID3D12Oueu	e::E> ID3D12Queue::E ID	3D12 ID ID3D12Queue::E		
ID3I	ID: ID3D12Queue::	E ID ID3D12Queue::Ex I	D3D12Queue::ExecuteC	ID3D12Queue:	:Exe ID3D12Queue::E ID	03D12 ID ID3D12Queue::E		
ID3I ID3D12Queue::ExecuteCommand	Lists(1) ID3D12Queue::	E ID ID3D12Queue::Ex I	D3D12Queue::ExecuteC ID3D12Qu	eue::ExecuteCommandLis	ts(1) ID3D12Queue::E ID	3D12 ID ID3D12Queue::E	7	Fartry
ID3E ID3D12Queue::ExecuteCommand	Lists(4) ID3D12Queue::	E ID ID3D12Queue::Ex I	D3D12Queue::ExecuteC ID3D12Qu	eue::ExecuteCommandLis	ts(4) ID3D12Queue::E ID	03D12 ID ID3D12Queue::E		۔ اندا بیر ا
ID3I ID3D12Queue::ExecuteCommand	Lists(1) ID3D12Queue::	ID ID3D12Queue:: I	D3D12Queue::ExecuteC ID3D12Qu	eue::ExecuteCommandLis	ts(1) ID3D12Queue::E	ID ID3D12Queue::E		IISL WIL
Compute queue								to just
	1020120				1020120	- (ιυ յսςι
	ID3D12Q	ieue::ExecuteCommandL			1D3D12Queue::Execut	eCommandLists(1)		
			100012Quee			/		
Graphics Compute	DMA Wait	Signal Present						
			•		6			
			Queue	submissions	Comma	na Duffers		
CPU								
AMD Ryzen		ODU						
	r 🗕 💻	GPU		0	-			
3970X		AMD Radeo	n 🔤 👘	8		88		
- 32-Core		RX 6800 X1		total	,	rotal		
Brocossor				total				
Processor								
A								
-			Queue	Count Key	Zueue	Count Key		
The a	polication is CDU	bound	Graphics que	Je 7	Graphics queue	37		
The a		bound.	Compute que	ue 1	Compute queue	1		
Frame	duration 15,608.63	39 µs						
F								
		AMI	D PUBLIC PE	RFORMANTR	EFLECTIVE BE	AUTY IN FC 6	STEPH	ANIE AND IHOR

- *i* amounts of synchronization are ded between graphics and compute ue
 - On PC, sync between graphics and compute queue is higher than console
- Cry 6 has one async compute command with all the BVH building calls and syncs ust before SSLR

GDC - MARCH 2022

68

ASYNC BVH BUILDING

- Some frames have much higher BVH building times vs others, which async helps hide
 - Make sure you launch async early enough to cover the variance $_{\Lambda}$



BVH TRAVERSAL COST



- Rays that miss geometry still pay the traversal cost for the TLAS hierarchy
 - Part of why having particles in separate TLAS gave good perf wins
- Best to avoid putting objects into TLAS if they are far away or unlikely to affect rendering result
- Camera is usually deep in a tree, so all rays can be affected by the higher traversal cost

Ray Divergence



BVH TRAVERSAL COST



- Far Cry 6 limits the objects that are in the TLAS to only a radius around the player
 - Terrain clip mapped is contained with a longer range
- BLAS objects are at a lower LOD, less nodes to traverse to get to triangles
 - More correct rendering for distance objects, fully detailed objects would alias
 - Helps address ray divergence



HYBRID REFLECTIONS BLAS/TLAS TIPS

- Consider using CS for vertex generation for BLAS
 - Far Cry 6 generates BLAS positions and DXR Prim attributes in a single CS pass
- We recommend trying async compute for BLAS/TLAS building
 - Easy to overlap with early/mid parts of your frame, big perf win for Far Cry 6
- Try to limit amount of geometry in your TLAS to avoid cost of traversing large TLASes per frame
 - For Far Cry 6 reflections, only objects within a radius of 100 are put into the BVH
- Consider having lower LODs for reflection BLASes for faster tracing and lower VRAM usage
 - Far Cry 6 uses a constant lower LOD for objects in BLASes




HYBRID REFLECTIONS SAMPLE

GDC



[AMD Official Use Only]

HYBRID REFLECTIONS OVERVIEW



FEEDBACK COUNTERS

- We store a feedback counter for each set of 8x8 pixels (1 tile)
- The counter consists of one U32 with hit/miss counters for 2 frames of statistics



• When deciding if a ray should be Hybrid or DXR, we use reprojected counters from previous frame

1 U32 per 8x8 tile of pixels





FEEDBACK COUNTERS

- Reproject counters by picking a random motion vector in the 8x8 tile
- We can switch pixels from SSLR to DXR, also need a way to switch from DXR to SSLR
 - Problem is all our DXR rays will always hit something with high confidence
- We solved this by randomly converting pixels in tiles from DXR to Hybrid





CLASSIFICATION

Successful screen space intersection



Reproject using motion vectors



CONCLUSION

- Overall, Hybrid Raytraced Reflections is a powerful technique that gives a good quality boost over SSLR
- Keeping raytraced reflections hybrid really helps in making the technique viable for real time
 - Lots of performance saved by using SSLR where applicable
- Plan ahead for asset system integration and handling materials in ray tracing shaders
- Important to keep in mind how ray tracing runs on HW and its performance characteristics
- AMDs Hybrid Reflection sample is currently available at <u>https://gpuopen.com/learn/hybrid-reflections/</u>
 - DX12 implementation with docs and MIT licensed public source code



DISCLAIMER AND ATTRIBUTION

DISCLAIMER

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. GD-18

© 2022 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, Radeon, Ryzen, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective owners.











