




RADEON




# A REVIEW OF GPUOPEN EFFECTS

TAKAHIRO HARADA & JASON LACROIX




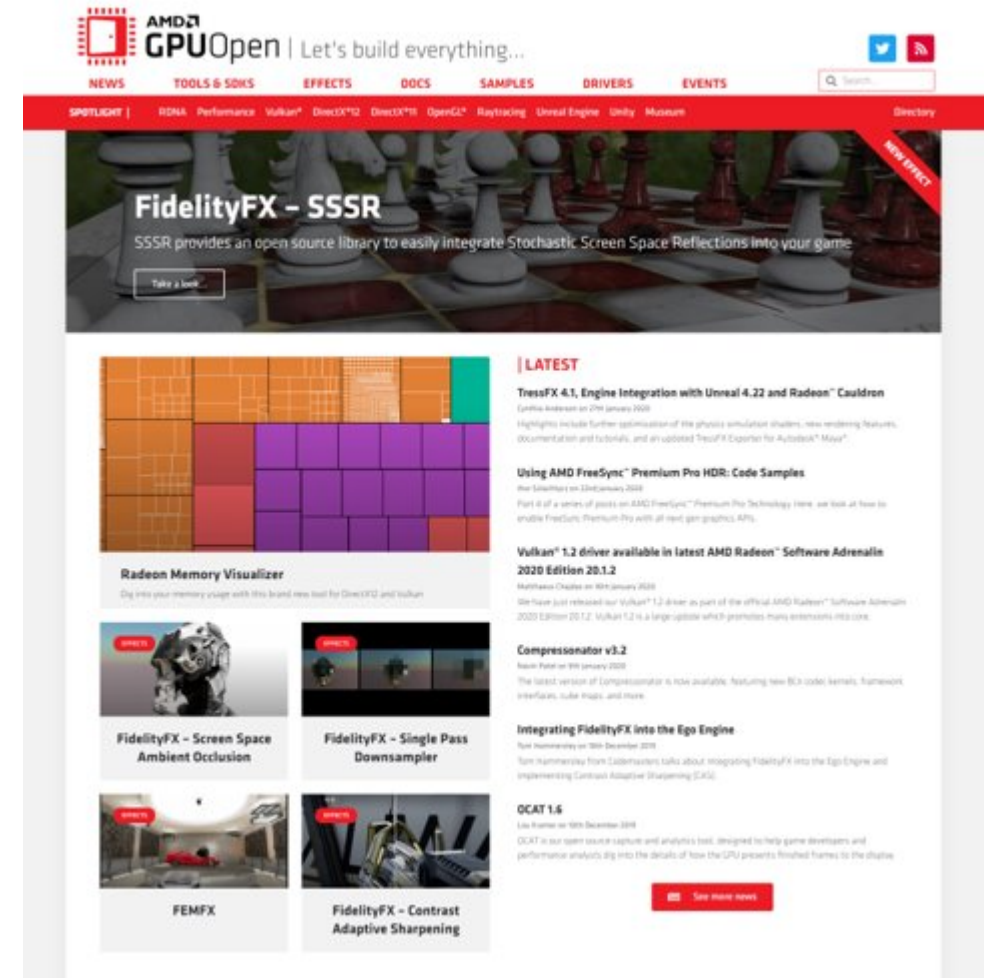
AMD  
GPUOpen



- An initiative designed to help developers make better content by “opening up” the GPU
- Contains a variety of software modules across various GPU needs:
  - Effects and render features
  - Tools, SDKs, and libraries
  - Patches and drivers
- Software hosted on GitHub with no “black box” implementations or licensing fees
- Website provides:
  - The latest news and information on all GPUOpen software
  - Tutorials and samples to help you optimise your game
  - A central location for up-to-date GPU and CPU documentation
  - Information about upcoming events and previous presentations

# LET'S BUILD A NEW GPUOPEN...

- Brand new, modern, dynamic website
- Easy to find the information you need quickly
- Read the latest news and see what's popular
- Learn new tips and techniques from our engineers
- Looks good on mobile platforms too!
- New social media presence
  - @GPUOpen 



# EFFECTS

A look at recently released samples

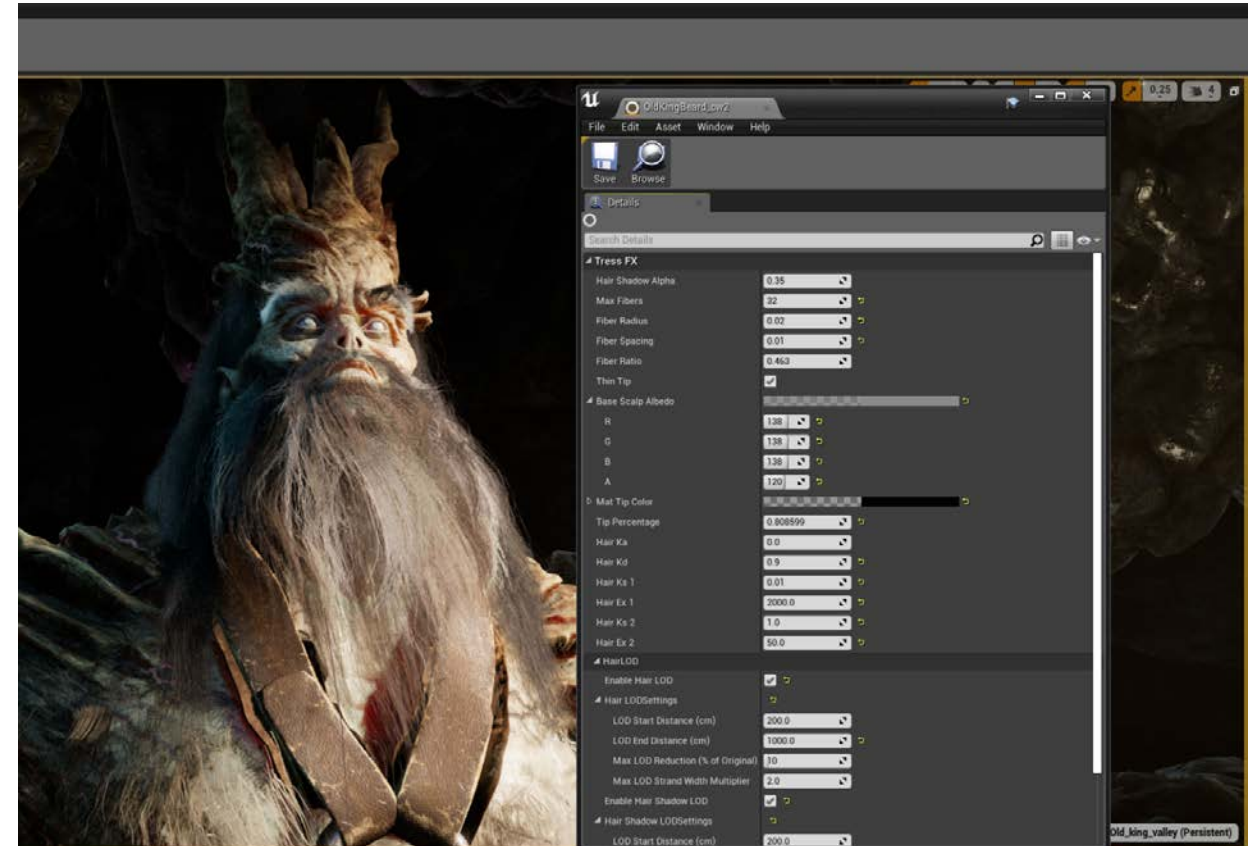
# TRESSFX 4.1

- Self-contained solution for hair simulation
- Implementation into Radeon® Cauldron framework
  - DirectX® 12 and Vulkan® with full source
- Optimized physics simulation
  - Faster velocity shock propagation
  - Simplified local shape constraints
  - Reorganization of dispatches
- StrandUV support
- New LOD system
- New and improved Autodesk® Maya® exporter



# TRESSFX 4.1 – UE4 ®

- External code module requiring minimal hooks
- Reference renderer for hair
- Provides a starting point for developers to integrate TressFX into their Unreal ® -based games
- Features
  - Unreal ® shadow system support\*
    - \*Dynamic lights only
  - Parameter blending via blend maps
  - TressFX component
  - .uasset support for TressFX data files
    - Can be drag and dropped into Unreal ®



# DEMO

Tressfx 4.1 UE4

# FIDELITYFX – SSSR


- Stochastic screen space reflections (SSSR)
- Based on industry leading algorithm
- Hierarchical depth buffer traversal kernel
- Glossy reflections via Ray Jittering
- High performant spatio-temporal denoisers
- Tile classifier that allows skipping non-reflective areas
- HLSL SM 6.0 with wave-level optimizations
- D3D12
- Vulkan (coming soon)





# FIDELITYFX – SSSR

- Surface requirements
  - Depth Hierarchy (2x2 minimum)
  - Per pixel motion vectors
  - Resolved scene
  - Normals for current and last frame
  - Roughness for current and last frame





# FIDELITYFX – SSSR

- Tile Classification Pass
  - Which pixel needs a ray
  - Which tile needs a denoiser
- Intersection Pass
  - Hierarchical trace through the depth buffer to find first intersection
  - Looks up environment map as fallback
  - Variable rate traversal
    - 1 ray per pixel for mirror surfaces
    - 1 ray per 4 pixels for rough surfaces
- Denoiser Passes



# FIDELITYFX – SSSR

- Spatial Denoiser Pass
  - 16 samples of Halton Sequence (2,3) discretized to 7x7 region
  - Edge stopping weight based on normal
  - Gaussian weight based on depth difference
  - Accumulate all radiance values
  - Normalize using accumulated weights
- Temporal Denoiser Pass
  - Temporal reprojection based on surface position
  - Temporal reprojection based on hit position
  - Picks the best result between the two
- Edge-Aware À-Trous Wavelet Filter (EAW)
  - Get rid of remaining spatial noise



# DEMO

FidelityFX: SSSR

# FIDELITYFX – CACAO

- Combined Adaptive Compute Ambient Occlusion (CACAO)
- A highly optimized implementation of SSAO
- Multiple settings which balance performance and quality
- Requirements
  - Depth buffer
  - Normal buffer (optional)
- Supports
  - DX12
  - Vulkan (coming soon)



# DEMO

FidelityFX: CACAO

# FIDELITYFX – LPM

- Luma Preserving Mapper
  - Help balance content for HDR and wide gamut
  - Tonemap(Luma(RGB)) rather than Luma(Tonemap(RGB))
- Key Takeaways:
  - Works with any positive RGB input
  - Targets any RGB output
    - Presets for available for common colorspace
      - REC709, P3, REC2020, HDR10, FS2
  - Completely ALU based solution without LUT
    - Can be easily decoupled to run to async compute queue
  - Supports FP16
    - Better perf on Vulkan
- Available for Vulkan and D3D12



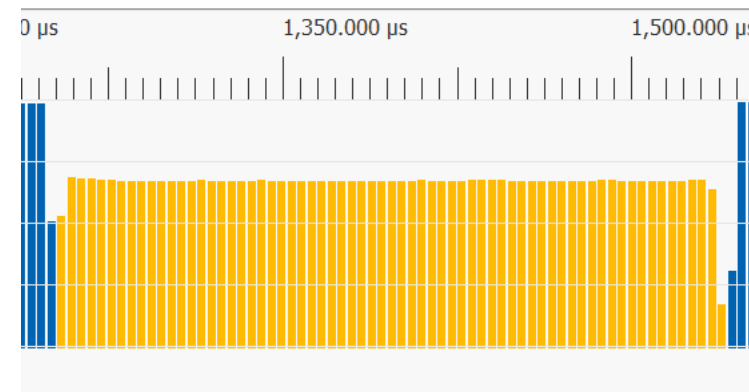
# DEMO

FidelityFX: Luma Preserving Mapper

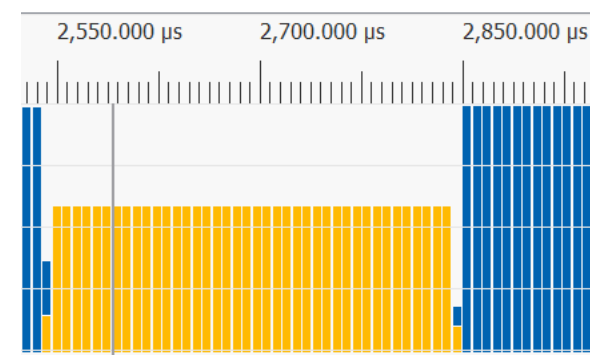


# FIDELITYFX – CAS

- Contrast Adaptive Sharpening (CAS)
- First effect in the FidelityFX library
- For **sharpening** and optional upsampling
  - Enhances sharpness and local high-frequency contrast lost due to TAA
  - Created to provide natural sharpness without artifacts with low overhead
  - Upsampling designed for Dynamic Resolution Scaling (DRS)
- Single compute dispatch call



CAS Upsample & Sharpen 1440p to 4K: 0.287ms RX 5700 XT




CAS Sharpen Only 4K: 0.301ms RX 5700 XT

# FIDELITYFX – CAS



# FIDELITYFX – SPD

- Single Pass Downsampler (SPD)
- Downsamples texture from 4kx4k down to 1kx1k
- User configurable 2x2 reduction operation
- GLSL and HLSL code available
- Compute shader based solution
- Suitable for async compute
- 2 modes of execution
  - Using wave operations and LDS
  - Using LDS without wave operations
- FP16 support

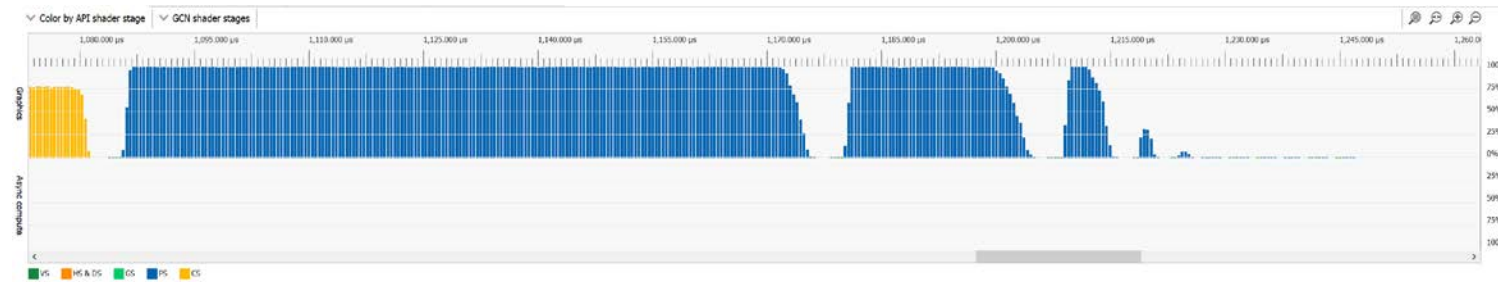


# FIDELITYFX – SPD

- SPD vs. traditional multipass pixel shader
- 4k source image
- 11 mip levels generated
- SPD: 0.127 ms
- Multi-pass Pixel Shader: 0.166 ms



SPD - 0.127 ms on Radeon RX 5700 XT





Multi-pass Pixel Shader - 0.166 ms on Radeon RX 5700 XT

# FEMFX

# FEM FUNDAMENTALS

- Physics method for deformable materials
- Models a solid as a mesh of elements (tetrahedra)
- Each element has material parameters
  - Stiffness
  - Volume preservation
  - Stress limit before permanent (plastic) deformation
  - Stress limit before fracture



# BENEFITS

- A different kind of physics for different effects
- Simulates deformable materials
  - Bending metal
  - Wood that flexes and breaks
  - Compressing, bouncing tires
  - Melting objects
  - “Alien” materials
- Improves fidelity
  - Realistic flexing
  - Storing and release of energy
  - Snap when breaking
  - Less brittle-looking fracture
- New interaction possibilities



# DEMO





# FEMFX LIBRARY

- Open source with GPUOpen license
- Multithreaded CPU implementation
  - For interaction with gameplay and other systems
  - Following trend of increasing CPU cores
  - SIMD optimizations
- Release includes
  - Houdini authoring tools
  - UE4 Plugin as example engine integration

# IMPLEMENTATION HIGHLIGHTS

- Implicit integration for stiff materials
  - Benefits stability and convergence
- Constraint-based contact between objects
  - Supports stacking
  - Compatibility with rigid-body solvers
- Continuous collision detection
  - Reduces missed collisions, tunneling
  - Using approximate method
- Multiplicative plasticity
  - Found behavior better than additive
  - Intuitive limits
- Sleeping feature to freeze slow-moving objects
- Disabling feature to remove object from simulation
- Non-fracture groups



# THREADING APPROACH

- Library uses all async threading
  - Dispatched tasks detect completion of work
  - One of these tasks submits a task to continue
  - Handles all the task synchronization
- High level threading features implemented in the library
  - Parallel for
  - Task graph
- Has a callback interface to support external task systems
  - Mostly just async task submit

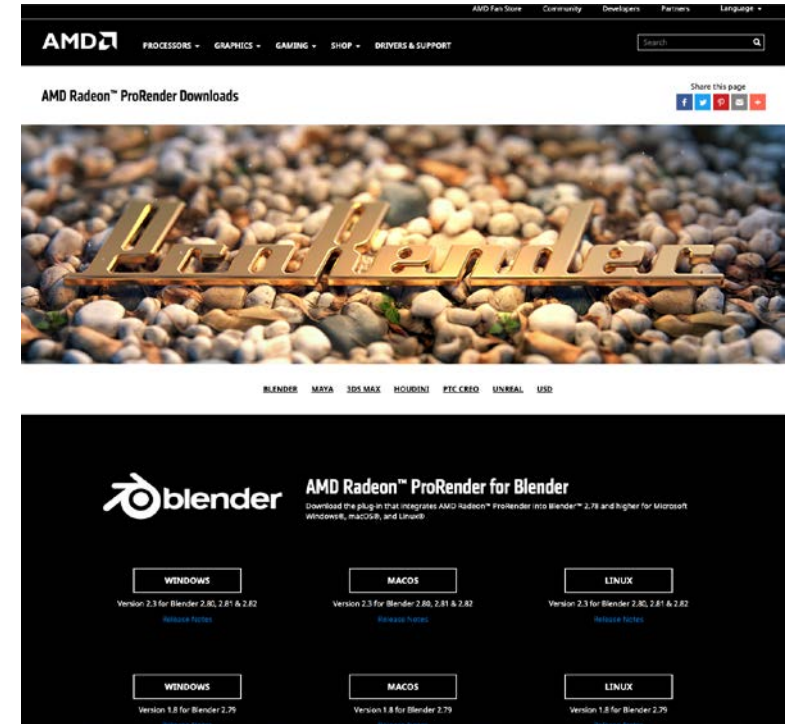
# RENDERING

- Demonstrated in UE4 Plugin
- Render mesh is separate from tetrahedral mesh
  - Allows more visual detail than simulation detail
- Preprocessing finds correspondence between them
  - Attaches each render vertex to closest tetrahedron
  - Finds barycentric coords of vertex in tetrahedron (four weights)
- At run-time vertices driven by tetrahedral mesh and weights
  - Like skinning
  - Can be done in a vertex shader
- Render meshes authored in pieces
  - Which break up with fractures in the tetrahedral mesh
- After fracture, may re-assign vertex to another tetrahedron
  - To keep vertices of the same render mesh piece together

# RADEON™ PRORENDER

# RADEON™ PRORENDER

- A powerful physically-based path traced rendering engine
- Supports GPUs and CPUs (OpenCL™ or Apple® Metal®)
- Works across Windows®, Linux®, and macOS®
  
- Plugins, integrations for artists
  - Maya, Max, Blender plugins
  - Integrations from our partners
  - USD Hydra delegate (Houdini)
- <https://www.amd.com/en/technologies/radeon-prorender-downloads>



# RADEON™ PRORENDER

- SDK for developers
- Open access to the SDK at GDC2020
- Radeon ProRender 2.0
  - Re architected for better experience
  - Improvements
    - Rendering speed
    - GPU memory reduction
    - MaterialX support



# RADEON™ PRORENDER

## FULL SPECTRUM RENDERING

- Approximated faster render modes
- Cross-platform, cross-vendor
- Implemented using
  - Vulkan, RadeonRays 3.0
  - DX12, RadeonRays 4.0
- Render modes:
  - Raster
  - Hybrid
  - Biased path tracer
- Dynamic lightmap GI
- Voxel GI



Rasterized



Hybrid



Biased



Full path traced



# RADEON RAYS

- Radeon Rays 4.0
  - Releasing at GDC 2020
- Radeon Rays 3.0
  - Implemented in Vulkan®
  - Supports both AMD GPUs and CPUs
  - Feature highlight
    - GPU-accelerated Bounding Volume Hierarchy (BVH) building
    - Half-precision (FP16) computation support
  - Works across Windows® and Linux®
- Radeon Rays 2.0
  - Implemented in OpenCL™
  - Supports both GPUs and CPUs
  - Works across Windows® and Linux®

# RADEON RAYS 4.0

- Unifies the zoo of SW/HW raytracing APIs
- Today, we have A LOT of APIs
  - Windows: DX12 Compute/DXR, Vulkan Compute/RT
  - Linux: Vulkan Compute/RT
  - MacOS: Metal Performance Shader
- Makes it platform / API independent
- Different level of intersection APIs depending on your needs
- Supports more intersection primitives


# RADEON RAYS 4.0

- Batch intersection
  - Good for pro apps
  - No graphics / compute API interop required
  - Cross-vendor
- DXR1.1 compute shader intrinsic (via interop)
  - AMD only
  - For performance critical workflows / gaming
- Fully customizable BVH traversal from compute shader
  - AMD only
  - For ultimate performance on AMD
  - Enables algorithms not available with DXR
    - Flexible LOD, early traversal stop, mixed nodes, on-the-fly BVH generation, etc.

# RADEON™ IMAGE FILTER LIBRARY

- Collection of image post processing filters
- Started to provide post processing effect for Radeon™ ProRender
  - Tone mapping
  - Denoising
  - Bloom
  - etc...
- Now includes our AI effects for denoising and upsampling.
- Based on OpenCL™, Metal, **DX12** (new for GDC) and **DML for AI**
  - Interop with OpenCL™, Metal, DX12 but also OpenGL & DX11
- Part of RPR SDK

<https://www.amd.com/en/technologies/radeon-prorender-developers>  
<https://radeon-pro.github.io/RadeonProRenderDocs/rif/about.html>

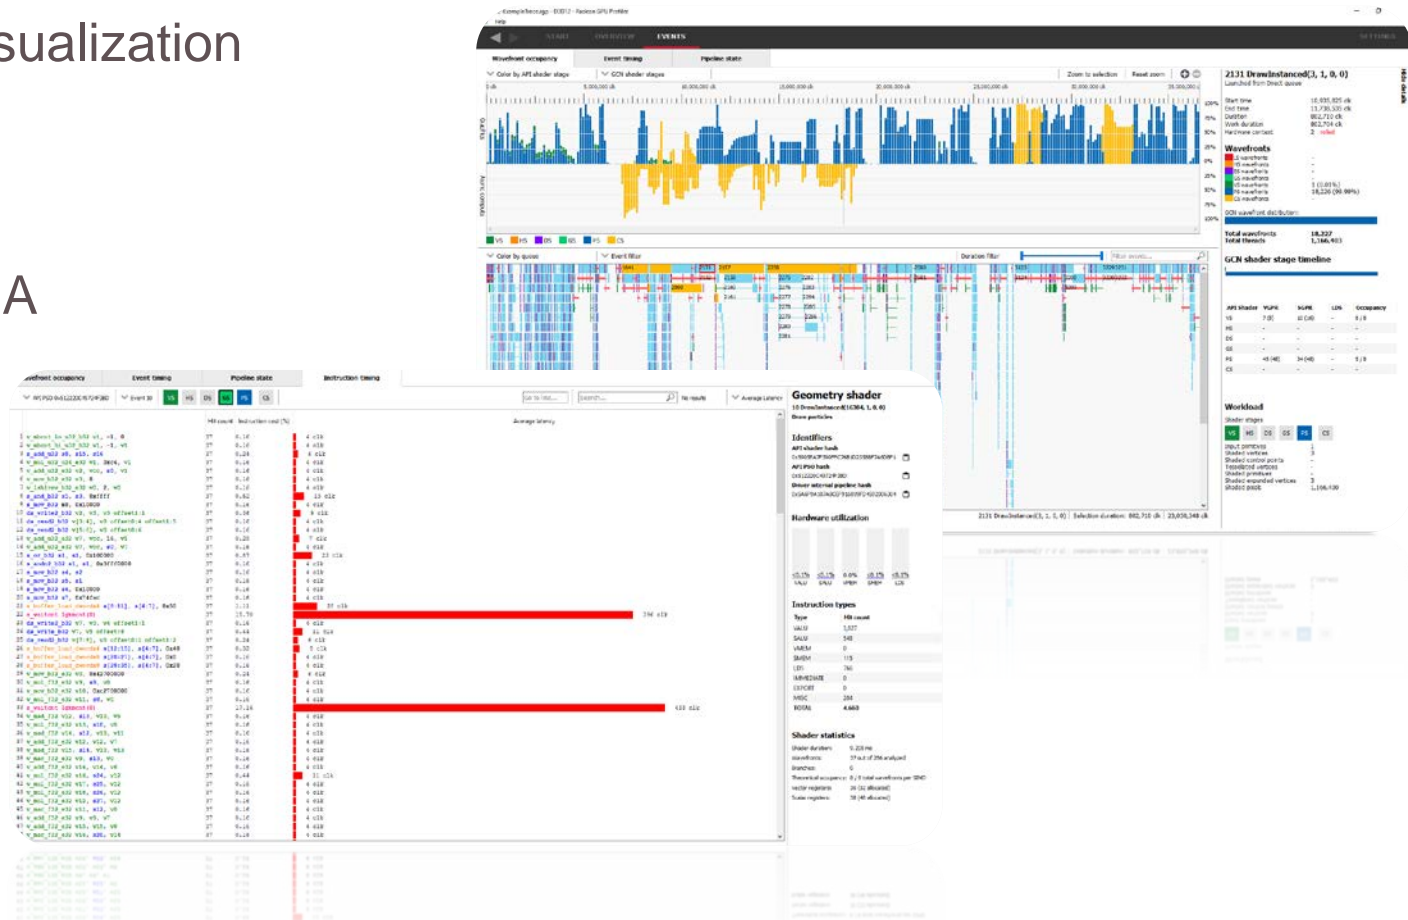


# TOOLS

The right tools for building great effects


# RADEON GPU PROFILER

- Draw, compute, and barrier visualization
- Pipeline state inspection
- WGP occupancy
- Instruction timings & RDNA ISA



# RADEON GPU PROFILER

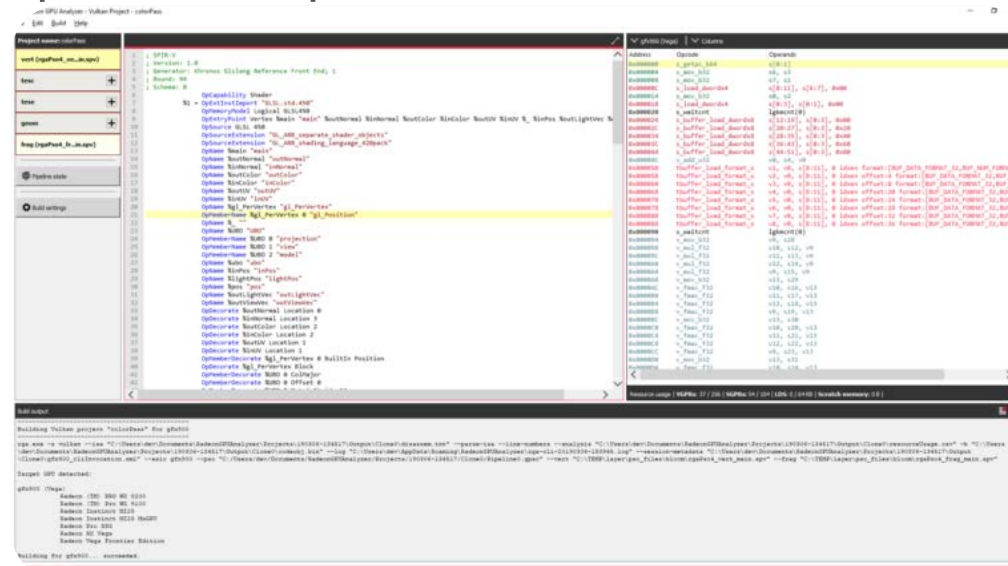
- RGP 1.7
  - Support for Radeon RX 5500 and Radeon RX 5300
  - New pipelines overview pane
  - New flag to show wave32 vs wave64 on RDNA hardware
  - New overlays:
    - User events
    - Hardware contexts
    - Command buffers
    - Render targets



<https://gpuopen.com/gaming-product/radeon-gpu-profiler-rgp/>

# RADEON GPU ANALYZER

- Offline compilation of shaders to AMD GCN ISA / RDNA ISA
- Shader register liveness analysis
- Support for DX11, OpenCL & Vulkan (SPIR-V)
- Full support for D3D12 Pipeline Compilation as of v2.3





# RADEON MEMORY VIEWER

- Create deep-dive snapshots at any point in time (even after the fact)
- Track memory-related events over time
  - Resource creation (render target, buffer, etc.)
  - Full virtual & physical memory map
  - Memory paging
  - Generate full history of any resource
  - Understand when resources were last accessed
- Snapshot comparisons
  - Detect memory leaks



## Disclaimer:

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED 'AS IS.' AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Use of third-party marks / products is for informational purposes only and no endorsement of or by AMD is intended or implied. GD-83

© 2020 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, Radeon and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. Windows and DirectX are registered trademarks of Microsoft Corporation in the US and other jurisdictions. Vulkan and the Vulkan logo are trademarks of Khronos Group Inc. Unreal and its logo are registered trademarks of Epic Games in the US and elsewhere. Autodesk, the Autodesk logo, and Maya are registered trademarks of Autodesk, Inc., and/or its subsidiaries and/or affiliates in the USA and/or other countries. Houdini is a registered trademark of SideFX in the US and elsewhere. Twitter Logo is a registered trademark of Twitter in the US and elsewhere. MacOS is a trademark of Apple Inc. in the US and elsewhere. Other names are for informational purposes only and may be trademarks of their respective owners.