CHVOSCENON

From Production to Real-Time Ray Tracing with V-Ray GPU and Project Lavina

Vladimir Koylazov, CTO, ChaosGroup Phillip Miller, VP of Product Management, ChaosGroup Alexander Soklev, Lead Developer, V-Ray GPU, ChaosGroup

CHVOSCHOD

creating powerful rendering and simulation technology to help you visualize anything imaginable

- global leader in computer graphics, with +20 years experience
- shaping the future of creative storytelling, VR, and digital design
- V-Ray is the industry standard within top design centers, architectural firms, advertising agencies, and visual effects studios
- 92 out of the world's top 100 architecture firms depend upon V-Ray
- used in +200 feature films, such as Captain Marvel and Black Panther and cinematic TV like Game of Thrones



SELECTED CLIENT LIST







INDEPENDENCE ONY

REE STAT

ACADEMY OF MOTION PICTURE ARTS AND SCIENCES

ZODIAC



CHAOZGROUP our mission

Solving hard problems so you don't have to

- Includes constantly looking into new technologies that may help our customers (and we love to play with new toys!)
- And looking for problems that need solving...



GPU Ray Tracing

Started in 2008 with "V-Ray RT"

 Using whatever it takes: GP-GPU, CUDA, OpenCL, OpitX, DXR, Vulkan, etc.

Why?

- We believe increasing **artist efficiency** is equal in importance to image quality
- Faster Rendering
 - = faster feedback
 - = faster decisions
 - = more design iterations
 - = better results & better designs

V-Ray GPU

+10 years of development delivers **the world's most fully-featured** GPU production renderer

- Uses same data and workflow as V-Ray
- Projects can start in V-Ray GPU and easily transition to V-Ray (if required)
- Maximizes performance by utilizing all GPUs and/or CPUs in the system
- Delivers identical results across hardware configurations for complete flexibility
- First to include NVLink support for pooling GPU memory (shipping since 2016)

and/or

V-Ray GPU

- Also enabling new possibilities
- Real-time motion capture rendering
- In 2014 using a cluster of VCAs housing scores of GPUs
- Looking forward to doing this now with better quality on a single RTX card...

V-Ray Solutions

() v.ray

V-Ray Next - since GTC '18

Faster, More Responsive, Smarter, Easier to Use, Focused on Production

All features presented last year delivered in:

- V-Ray Next for 3ds Max May 2018
- V-Ray Next for Maya November 2018
- V-Ray Next for Unreal November 2018
- V-Ray Next for Sketchup February 2019
- V-Ray Next for Rhino
 March 2019
- in Beta: Houdini, Katana, Nuke, Modo, Blender

FASTER

- 2X Faster compared to previous V-Ray 3.6 across all GPU generations
 - New rendering architecture
 - Designed for the future like RTX

- Adaptative Dome Light up to 7X faster
- Volumetrics & Hair now accelerated, and many times faster

- New hair shader
 - Glints and glitter controls

- Environment Fog and Volumetrics
 - Supports textured volumetrics
 - Improved sampling strategies for homogeneous volumes

Edges texture

Bercon noise

VRscans

More Responsive

- Support for the NVidia AI Denoiser (supplied within OptiX) – great for interaction
- Reworked interactive behavior for far faster start-ups, updates and iterations

 this alone makes it seem twice faster
- Interactive modes (IPR etc.) now share the same settings as production
- Your interactive results match your production results every time

Easier to Use

- New public API queries feature support status
- Allows GUI to reflect what is supported by V-Ray GPU
- V-Ray Next products make V-Ray GPU simple to use – if you can see it, it's supported

Smarter

- Auto Exposure and Auto White Balance
- Exposed as part of Camera settings
- Great results with no effort delivers smooth animations through changing lighting.

Focused on Production

- GPU Bucket rendering
 - Automatic configuration of best settings
- Better scaling with distributed rendering (DR)
- Cryptomatte support
- Resumable Rendering (start where you left off)
- Improved Chaos denoiser
 - Consistency across render elements allowing faster renderers to be used in post

And More:

- Improved displacement quality & speed for a fraction of the memory
- Improved lens effects
- Refraction dispersion
- On-demand textures now also in IPR
- Expanded NVLink support for the new bridge types on RTX cards

V-Ray Next – some user feedback

Tomasz Wyszolmirski 🗘 Here's an list of things that made our experience better, since V-Ray Next and update 1 was released: 1. We've seen great improvement in user experience. Active Shade is gone, IPR is working smoother than ever. UI is simplified and easier to use. Update 1 greatly improved IPR stability.

Volume grid support with bitmaps is kickass, allows for some very interesting setups on top of ability to render simulated fire and smoke.

3. Fast environment fog is impressive. This heavily reduced rendering times.

4. New hair shader is simply the best there is. Period.

 Up to 2x faster rendering in general as compared to 3.6. Because of that we didn't have to expand our render farm with many more GPUs (Sorry NVIDIA).

 You should brag about vrscene, It's great underappreciated feature. And by nature of V-Ray GPU, it works flawlessly. I've seen it was recently improved in interactivity as well, but didn't test it yet.
 Viewport IPR is super fun to work with and also since it's introduction, the IPR performance also improved.

8. New lens effects and dispersion makes V-Ray GPU great for rendering jewelry.

 Bucket rendering, that introduced cryptomatte and better scaling on multi-GPU machines. Though still could be improved it's great start.

It's not ordered list of improvements that I feel they are equally important.

Nawras Ryhan The new IPR is DOPE! Period. It is like my most wanted feature comes true. It was crazy moment when I fired the IPR in huge city scene (expecting Max to crash) but I was able to move around and play with lights and composition in almost realtime (using only one 1080ti)

Faking Fog in post days is gone.. The speed of rendering tog now is great that I would start adding fog to every scene loool

Muhammed The GPU IPR uses less VRAM for geometry and displacement compared to before...Overall Vray GPU has been very effecient using memory, haven't ran out of memory on my 1080tis for awhile Muhammed Top feature has to be between the reworked GPU IPR and Adaptive domelight..

The GPU IPR in Maya is by far the best in the industry by miles, -Debug shading, offers isolating lights/textures/objects interactively in the VFB or Maya's Viewport. And it has other useful modes for look dev (lighting/AO/normals/wirframe/UVs)

-runs directly in Maya's IPR, and moving geometry around has been improved quite a lot which is huge deal for what we do. interactivity for the IPR overall is way better than what we had in 3.6 -Viewport denoising with Optix for instant feedback, which works in Maya's viewport or the VFB

About adaptive domelight, HDRI lighting for interiors is one click setup really, no need for portals. And renders has been 2x or 3x faster compared to using portals..Haven't seen this kind of improvement in any other engine really

These 2 features has been "game changing" in our workflow. There are more cool features like VRscans support, the new Bloom/Glare, new volumetrics rendering, Bucket mode, autoexposure/white balance, the new hair material and metalness for PBR shading ,denoised render elements and the new UI for materials/lights and render settings.

Many of these features are not avialble in any other GPU engine, same for Nvlink support. Working nicely here on 2x 2080s I need to create a new thread and upload few more example scenes. Having a single slider for render settings is big deal compared to other engines. GGX tail falloff and Vray dirt are massive features.

V-Ray GPU – in the works

Currently being Worked On:

- Full out-of-core rendering
- RT Core Support
- alSurface shader
- Texture nodes (Curvature & Distance Tex)
- Deep Output
- OSL procedural textures (patterns)
- Faster startup times
- Faster rendering on x86

V-Ray GPU – in the works

RT Core Support

- About 2X faster than our CUDA path on the same RTX hardware
- Already supporting:
 - +95% of V-Ray Shaders
 - All light types
 - Adaptive Lights
 - Light Cache
 - Particles
 - Hair
 - Motion Blur
 - SSS
 - IPR

See it in action at the GTC NVIDIA booth

V-Ray GPU – in the works

RT Core Support

- In testing with select customers today, in broader Beta this summer
- RT Core support will have complete parity when it ships

 just as capabilities & results for x86 and CUDA match now
- Product goal to simply take advantage of RT Cores whenever suitable RTX cards are found
- If you want access, monitor the Chaos Group forums @ https://forums.chaosgroup.com/
- Preliminary performance results from RT cores

Scene

RT Core acceleration within Turing GPUs

- Ray traced rendering comprised of two major parts: Ray Casting vs. Shading
 - Intersecting rays with geometry vs.
 - processing materials, textures, lighting, and effects
- RT Cores within Turing GPUs accelerate just the ray casting portion

Project Lavina

Research project to explore RTX

What if RT Cores were always there?

Goals:

- > 24 FPS
- Pure ray tracing no rasterization
- Work directly with V-Ray data
- Continue increasing capabilities while staying within real-time budget
- Also support large data sets (with out-of-core for geometry)

Research in Action

Standalone viewer of vrscenes

- Using Microsoft's DXR
- Achieving +5X over CUDA on the same hardware
- Loads vrscenes exported from any V-Ray product
- No adjustment to materials, lights or geometry
- Supports out-of-core geometry for large scenes
- Includes collision detection
 when rays are free, why not?!

🜒 GPU mem : 2.31 Gbytes, Frame : 19617, fps : 34.01, denoise : 4.90ms, loaded blases : 931, loaded last frame : 0, unloaded last frame : 0, total : 0, cam passes : 0, shadow passes : 0 🥚 – 👘 🛛

Project Lavina

Possibilities with V-Ray GPU

Lavina focuses on real-time speeds

impressive quality without pressure for final frame

V-Ray GPU focuses on quality & capability

very interactive, without pressure to be real-time

Cloud Rendering

Ideal for when you have a "time problem"

- Can scale for the most demanding of jobs.
- Pay only for what you need, when you need it.
- A "super computer" from any machine.
- Cloud "Economy Magic"

Cloud Rendering

Economies

- Cloud providers charge linearly for computing (e.g., 128 core/hours costs 8x 16 core/hours)
 - If scaling is linear, then it costs no more to produce the same image more quickly
- Big discounts for "preemptible" machines
 - Plan to stop & resume and you can save 80%
- Windows costs a premium
 - Run on Linux
- Both V-Ray and V-Ray GPU scale very linearly within system, are resumable, and support Linux

GPU Cloud Rendering

Economies

Speed often doesn't matter if it costs more

speeds are GPU-only

- In 2018, cloud GPU prices reduced
 especially for preemptable machines
- GPUs cost more/hour

	Rental Ratio
64 Core	1.0
1 P100	1.7
1 V100	2.2
4 V100	5.6

GPU Cloud Rendering

What if?

- Built-in to what you already use
- Nothing new to learn
- Nothing to manage
- As easy as rendering locally
- Really efficient
- Cost effective

Chaos Cloud

Available Now

- Proven after 6 months of Beta with 20k users sending 500k jobs producing nearly 6 million images
- A simple "Cloud" button in the likes of 3ds Max, Maya, C4D, SketchUp, Rhino, Revit, etc.
- Same settings as normal with a job name
- Everything is included all cloud aspects, software licensing, etc.
- Never upload the same asset twice
- Monitor and manage from any web browser
- Only using CPU instances today. GPU instances to be provided later this year after we make them easy

