

GPU RENDERING



Kevin Margo is a talented VFX/CG supervisor at Blur Studio. He's also working on some very impressive new technology with Chaos Group, maker of the popular V-Ray rendering engine.

Kevin's short film, CONSTRUCT, is being rendered entirely on the GPU. Known for its speed advantages over traditional CPU rendering, the GPU is gaining momentum as the renderer of choice for final frame production.

In addition to final frame rendering, Kevin is pioneering a new method of pre-visualization that incorporates real time ray tracing with live motion capture data using V-Ray for Motion Builder. This new technique allows a director of photography to participate in virtual lighting decisions in real-time. It also allows the director to peer through his camera into the virtual world and see the CG characters interact with virtual sets under full global illumination and ray tracing

Kevin relies on BOXX GPU edition workstations for all aspects of this project. We caught up with Kevin and asked him a few questions regarding this groundbreaking production.

BOXX

Accelerating the rendering process is an advantage, but as you know, faster render times often lead to more creative iterations of your work. Can you talk about how the GPU has helped in this regard?

Kevin

When I think of iterations, the idea of "discreet steps" come to mind: 1. make an adjustment. 2. render/wait a while. 3. review 4. go back to step 1. However it didn't really feel like that when working on the GPU. Anytime I was doing shader or lighting work, I had V-Ray RT running on the side. It was so dynamic and responsive and most often it had progressively rendered something good enough to review and proceed within a single breath. It kept pace with my creative flow. Coffee breaks became about caffeine withdrawal and not about waiting for renders.

BOXX

Perhaps the most impressive piece in all of this is what you have done with V-Ray RT for Motion Builder. How do you see that changing things?

Kevin

Thanks! Yes, V-Ray RT for Motion Builder was/is an amazing collaborative experience with the folks at Chaos Group. They were extremely supportive when I approached them with the idea. Pairing the development of CONSTRUCT with the development of V-Ray for Motion Builder proved very fruitful, with a series of concrete production demands quickly highlighting issues/needs.

I'm so excited about what this workflow could mean for the industry. The scalability of path tracing at interactive frame rates, lighting and composing performances in that context, and the ability of progressive path tracing resolving to a final quality production frame is HUGE. So much creativity is now unleashed in the motion capture volume. Lighting TD's relighting the set (using CG representations of live action lighting kits) while cinematographers establish focal distance, f-stop values and shutter speeds, all while the decision maker/director is in attendance to sign off on realistic lighting decisions at this early stage in production is HUGE. It could alleviate a compositing pipeline burdened to support infinite flexibility, unnecessary if the directors/DPs are on site during performance capture establishing desired lighting.



As a result this simplifies the final rendering/mastering pipeline tremendously. Clients could see a product much nearer the final product potentially weeks/months earlier than previously possible. Camera operators now can compose to color and light, expanding the breadth of creative considerations available. I watch live action camera operators always re-framing as a response to the shapes of color that enter/exit frame. THAT is hugely beneficial. Often a lighting team is forced to re-engineer compositions to locked cameras/performance to achieve good compositions. That feels stale and forced. Now there can be that amazing tension that exists when all elements are in play simultaneously. Everything influencing everything else. Happy accidents and spontaneity ensues.

Additionally, there's potential for this to rethink how VFX-dependent live action films are made. If the VFX industry is enabled to establish lighting and camera direction prior to live action shoots, that is a position to dictate to a live action cinematographer how that set should be lit. Currently, the live action set responds to the direction of the VFX industry, instead of the historically opposite. Gravity and it's lighting methodology is a great example of how a workflow like this could mature. With this creative control could come greater control over the broader filmmaking process and resulting costs.

“Given the cutting edge hardware we’re putting through heavy production, I’m extremely impressed at its stability, which is VERY important when trying to render a one minute teaser in five consecutive nights :) Whatever cooling rig is inside does an impressive job given how intense these cards are.”

-Kevin Margo | Blur Studio

With greater access to these kinds of tools earlier in production, concurrently with the elevated/cherished acting aspect of filmmaking, the VFX industry is further enabled to express their own stories/visions. That's what I'm REALLY excited to see. With a democratizing of the filmmaking process, a tremendous RANGE of creative expression can flourish. I would LOVE to see what a mature virtual production workflow would do for the VFX industry. Similar to what the DSLR industry did for indie filmmaking in the last decade. So much content, so much interesting material previously not cost effective or achievable could be possible.



The role of V-Ray RT for GPUs in the rendering process has evolved from a fast preview renderer into a full-fledged final frame production tool. As Kevin demonstrates, the GPU's role will also play a bigger role in the VFX pipeline. BOXX GPU edition workstations are perfect solutions for GPU centric work and ideal for GPU centric tasks like rendering. They are available with a single CPU, so you can be sure your money is being spent where it matters most. BOXX performance specialists are standing by to help you find the ideal configuration for your needs and budget.

BOXX

Many people are interested in the potential render acceleration of GPU rendering compared to traditional CPU rendering. You mentioned your HD frames are rendering in 5-10 minutes. Can you provide any insight into how long your frames would take on the CPU?

Kevin

I just did some benchmarks. Keep in mind the GPUs are stacked, so the results are impressive. Image attached demonstrating the render time difference (See below). I capped paths per pixel to 39 since the CPU was going to take FOREVER to render.

So by my calculations that's almost a 70x performance boost with the GPUs. That's kinda crazy now that I've done the numbers. Another way to look at this, with RT GPU I was rendering between 800 and 1024 paths per pixel, which in most cases were acceptable final production quality, arriving at 4-14 minute frame times. Now look at what 16 minutes on CPU took and how much noise there is.

**V-Ray RT 3.0 CPU - i7 4960K/4.5ghz/6 cores,
12 threads (overclocked)
39 paths per pixel
RenderTime: 16min 4 seconds**

**V-Ray RT 3.0 GPU - 1x K6000. 2x K40's
39 paths per pixel
RenderTime: 13.9 seconds**

Based on Kevin's results, a single Tesla K40 or single GTX Titan would be approximately 20X faster than the single i7 CPU in the scenario described below.



BOXX

Given the fact that V-Ray RT 3.0 now supports render elements, do you think GPU final frame rendering is a viable option for production?

Kevin

Render Elements are very appealing from a production standpoint. Having access to a variety of AOV's enables a robust and flexible compositing pipeline that mirrors traditional CPU based VFX workflows. However with CONSTRUCT being an entirely CG project I had ZERO need for render elements. The inherent beauty of V-Ray RT required no need for any compositing. A single beauty render was used for every shot in the teaser, with a bit of color grading to sweeten the image. I love the purity and elegance of path tracing. If photorealism is my goal, there's nothing I can do additionally to a render other than screw it up in comp—so I didn't comp.

BOXX

In your GTC talk, you discussed several methods you employed that enable you to stay within the constraints of the GPUs memory. Can you talk about those techniques? Did you even approach the memory limits?

Kevin

Memory is the primary limitation with GPUs at the moment. However, with each generation of cards that becomes less of an issue. The K6000's and K40's with 12 GB memory entered the realm of what high quality production demands. Even still, frequently I encounter scene files using 36+ GB of CPU RAM, so I realized that adopting a 'rationing' mentality to make CONSTRUCT a viable GPU rendered project was necessary. Reducing the unique geometry footprint by using instancing as much as possible helped. The house under construction is essentially thousands of instanced boxes in the form of repeating wood planks. Unique grass patches and trees were kept to a minimum, again instancing those as much as possible.

Relying on shader-based color variations were extremely useful in creating visual complexity with minimal assets. The robots again were all the same instanced geometry with shader color variations.

So with those geometry efficient approaches, next came texture maps. This, by far, proved the largest memory consumer. Very quickly, a scene full of a few dozen 2k or 4k texture maps would demand 6-10 gigs of memory. This was where the most focused optimizations needed to occur. Our character modelers implemented pixel-saving strategies like storing multiple data channels (reflect,gloss,bump) inside the 3 RGB color channels to be accessed individually in the shader for their respective data usage. We avoided color info storage in diffuse maps, instead grayscale maps used to mix various colors together. Collapsing complex shader trees to single bitmaps per component also was needed at times. These are all tricks game engines have used for years.

Lastly, I had various resolutions for every bitmap in the scene, from 256 ^2 up to 4k. For each shot I started with a low end resolution like 256 or 512...did some test renders with DOF and motion blur enabled, and evaluated what scene elements needed higher resolutions and increased those accordingly. In the end, most scenes hovered between 4 and 6 GB of memory, with the most complex scenes touching 7 or 8 GB, which was actually quite comfortable since there's always a little more overhead on the main card handling windows/app operations. My K6000 had this responsibility, so it was always about 1 GB higher than the two compute only K40s.



BOXX

You've been using your BOXX workstations in a production environment for some time now. Can you talk about how your experience has been so far?

Kevin

Given the cutting edge hardware we're putting through heavy production, I'm extremely impressed at it's stability, which is VERY important when trying to render a one minute teaser in

five consecutive nights. Whatever cooling rig is inside does an impressive job given how intense these cards are.

The role of V-Ray RT for GPUs in the rendering process has evolved from a fast preview renderer, into a full-fledged final frame production tool. As Kevin demonstrates, the GPU's role will also play a bigger role in in the VFX pipeline. BOXX GPU edition workstations are perfect solutions for GPU-centric work.