

WELCOME!

LIVE with ROBERT GREEN:

Demystifying GPUs for CAD Workstations



Robert Green

CAD Management Expert
Cadalyst Contributing Editor

March 21, 2018


cadalyst

TODAY'S PRESENTATION

What is a GPU?

What does a GPU do for CAD?

What are the key terms to know?



How much memory is needed?

What are integrated graphics?

How to specify the right GPU?

WHAT IS A GPU?

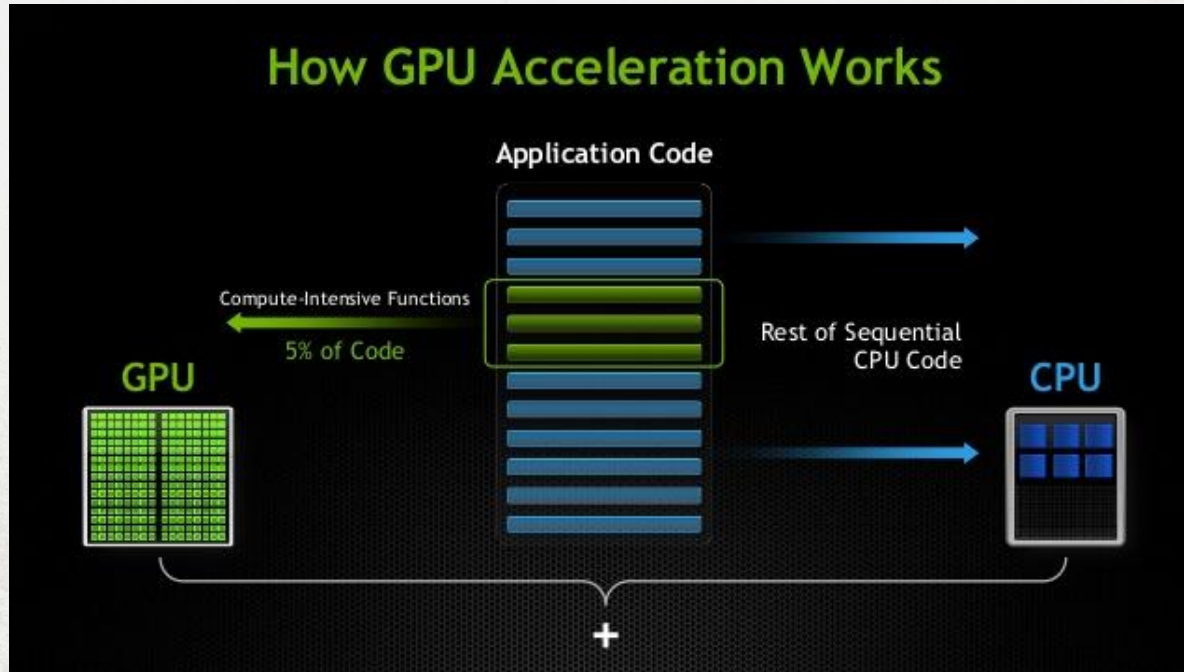
GPUS OFFLOAD GRAPHICS TASKS FROM THE CPU

This unburdens the CPU so it can handle operating system and core processing functions – like CAD modelling and analysis.

The architecture of a GPU is optimized to process video as fast as possible.


CPUs are versatile while GPUs have a narrow task load.

The processor (CPU) takes any instructions that must be processed for rendering to the screen and throws that computing load to the GPU.



Source:
NVIDIA

GPUs use a parallel processing architecture with hundreds to thousands of cores to process video data rapidly to support fast video playback.






CONCLUSION:

Your software use and need for high-resolution, motion-based graphics largely determine what sort of GPU you will need.

**INTEGRATED
GRAPHICS**

VS.

DISCRETE GPUs



WHAT ARE INTEGRATED GRAPHICS?

- ▶▶ A GPU that is a part of the processor
- ▶▶ A lower-cost, less-powerful option than a discrete GPU
- ▶▶ Shares memory with the CPU
- ▶▶ Not upgradable other than adding a discrete GPU
- ▶▶ Typically ideal for 2D CAD and light 3D modeling tasks
- ▶▶ In CAD workstations, CPUs with integrated graphics are typically Intel® Core™ i7 or Intel® Xeon™ processors.



WHAT IS A DISCRETE GPU?

- ▶ Sometimes called a “graphics card,” a discrete GPU is a standalone, add-in board that installs in an expansion slot in the workstation
- ▶ Varying amounts of cores/dedicated memory included
- ▶ Varying resolutions and frame rates supported
- ▶ Varying connecting ports included
- ▶ Priced from a few hundred to several thousand dollars
- ▶ Manufacturers include NVIDIA and AMD



SAVING VIA INTEGRATED GRAPHICS

If integrated graphics support your needs today, save money with an integrated GPU now and add a discrete GPU later as your needs become better defined.

GPU

TERMINOLOGY



SPECIFICATIONS

GPU Memory	8 GB GDDR5
Memory Interface	256-bit
Memory Bandwidth	Up to 243 GB/s
NVIDIA CUDA® Cores	1792
System Interface	PCI Express 3.0 x16
Max Power Consumption	105 W
Thermal Solution	Active
Form Factor	4.4" H x 9.5" L, Single Slot, Full Height
Display Connectors	4x DP 1.4
Max Simultaneous Displays	4 direct, 4 DP 1.4 Multi-Stream
Display Resolution	4x 4096x2160 @ 120Hz 4x 5120x2880 @ 60Hz
Graphics APIs	Shader Model 5.1, OpenGL 4.5 ⁺ , DirectX 12.0 ⁺ , Vulkan 1.0 ⁺
Compute APIs	CUDA, DirectCompute, OpenCL™

GPU CORES & MEMORY

- ▶▶ Expressed in core number and GB
- ▶▶ More cores and memory mean more images can be processed independently of the CPU
- ▶▶ Memory interface determines the number of memory paths and bandwidth from CPU to GPU
- ▶▶ Most GPU manufacturers harmonize core number and amount of memory for optimal performance ...



HOW MUCH MEMORY TO BUY?

Memory supports core performance, so the real question is, How much performance do you need? Memory follows cores and bandwidth, which together dictate speed.



HOW DO I GAUGE MY NEEDS?

The more graphics and video you need — especially if running at the same time — the higher your performance needs will be. In these cases, “buying up” can protect against slowdowns during demanding session work.



SPECIFICATIONS

GPU Memory	8 GB GDDR5
Memory Interface	256-bit
Memory Bandwidth	Up to 243 GB/s
NVIDIA CUDA® Cores	1792
System Interface	PCI Express 3.0 x16
Max Power Consumption	105 W



SPECIFICATIONS

GPU Memory	16 GB HBM2
Memory Interface	4096-bit
Memory Bandwidth	Up to 717 GB/s
NVIDIA CUDA® Cores	3584
NVIDIA NVLINK™	2 GP100 boards supported
System Interface	PCI Express 3.0 x16
Max Power Consumption	235 W

INTERFACE & POWER

- ▶▶ Typically a PCI Express expansion slot (or two)
- ▶▶ Power consumption in watts
- ▶▶ Workstation case size and power supply are a consideration



CASE & POWER SUPPLY?

In general, as GPU power increases, so does GPU size and power consumption. You'll need more expansion slots, and larger cases/power supplies to assure proper operation.



SPECIFICATIONS

GPU Memory	8 GB GDDR5
Memory Interface	256-bit
Memory Bandwidth	Up to 243 GB/s
NVIDIA CUDA® Cores	1792
System Interface	PCI Express 3.0 x16
Max Power Consumption	105 W
Thermal Solution	Active
Form Factor	4.4" H x 9.5" L, Single Slot, Full Height
Display Connectors	4x DP 1.4
Max Simultaneous Displays	4 direct, 4 DP 1.4 Multi-Stream
Display Resolution	4x 4096x2160 @ 120Hz 4x 5120x2880 @ 60Hz
Graphics APIs	Shader Model 5.1, OpenGL 4.5 ⁺ , DirectX 12.0 ⁺ , Vulkan 1.0 ⁺
Compute APIs	CUDA, DirectCompute, OpenCL™

VIDEO RESOLUTION

- ▶▶ Typically 4096x2160 (4K) maximum supported
- ▶▶ Sometimes 5120x2880 (5K) maximum supported
- ▶▶ Refresh rate of 120 Hz desirable
- ▶▶ Lower resolutions like 1920x1080 HD supported



SPECIFICATIONS

GPU Memory	8 GB GDDR5
Memory Interface	256-bit
Memory Bandwidth	Up to 243 GB/s
NVIDIA CUDA® Cores	1792
System Interface	PCI Express 3.0 x16
Max Power Consumption	105 W
Thermal Solution	Active
Form Factor	4.4" H x 9.5" L, Single Slot, Full Height
Display Connectors	4x DP 1.4
Max Simultaneous Displays	4 direct, 4 DP 1.4 Multi-Stream
Display Resolution	4x 4096x2160 @ 120Hz 4x 5120x2880 @ 60Hz
Graphics APIs	Shader Model 5.1, OpenGL 4.5 ⁺ , DirectX 12.0 ⁵ , Vulkan 1.0 ⁴
Compute APIs	CUDA, DirectCompute, OpenCL™

CONNECTOR TECHNOLOGY

- ▶▶ Typically Display Port (DP) 1.4
- ▶▶ Sometimes HDMI or DP w/Audio supported
- ▶▶ Typically more than two monitors supported
- ▶▶ Typically 120-Hz refresh rates supported



CONCLUSION:

Resolution and monitor connectors predict whether existing monitors can be used and dictate specs for new monitor purchases.



SPECIFICATIONS

GPU Memory	8 GB GDDR5
Memory Interface	256-bit
Memory Bandwidth	Up to 243 GB/s
NVIDIA CUDA® Cores	1792
System Interface	PCI Express 3.0 x16
Max Power Consumption	105 W
Thermal Solution	Active
Form Factor	4.4" H x 9.5" L, Single Slot, Full Height
Display Connectors	4x DP 1.4
Max Simultaneous Displays	4 direct, 4 DP 1.4 Multi-Stream
Display Resolution	4x 4096x2160 @ 120Hz 4x 5120x2880 @ 60Hz
Graphics APIs	Shader Model 5.1, OpenGL 4.5 ^a , DirectX 12.0 ^b , Vulkan 1.0 ^c
Compute APIs	CUDA, DirectCompute, OpenCL™

API/DRIVER TECHNOLOGY

- ▶▶ DirectX 11 or higher for Windows CAD platforms
- ▶▶ OpenGL 4.x for iOS, Android, Windows apps
- ▶▶ Check software documentation for versions required

- ▶▶ Note: Sometimes old software doesn't recognize newer graphics driver technology without installing updates to the software itself.



THE GOLDDILOCKS GPU ZONE

Purchase the GPU that meets your monitor and resolution needs as a baseline so you don't overspend. Buy one step up today to avoid obsolescence, or upgrade later as technology evolves.

FRAME RATES **& VIRTUAL** **REALITY (VR)**

WHAT IS FRAME RATE?



- ▶▶ The number of frames per second (fps) the GPU can send to your screen
 - ▶▶ Higher frame rate = smoother video
 - ▶▶ Higher resolution = slower frame rate
 - ▶▶ More monitors/processes = slower frame rate
- ▶▶ For realistic, immersive video (like virtual reality) on screens, 30 fps is the gold standard
- ▶▶ For headset-based VR playback, 90 fps is recommended to avoid nausea and disorientation



WHERE'S THE FRAME RATE?

- ▶▶ Manufacturers won't list it because there are too many variables to consider.
- ▶▶ To find frame rate data, search online based on your graphics card model and software application, for example:

“NVIDIA Quadro 4000 3ds Max frame rate”

- ▶▶ Results will come from a variety of sources, but you'll be able to see in general how other users with your same software find the performance.



CONCLUSION:

If playback and VR frame rates are your principle concern (rather than running CAD applications), then “buying up” into a higher performance “VR Ready” GPU is justified.

PRO vs. GAMING **GPUs**

ARE YOU DESIGNING OR VISUALIZING?

CAD and modeling tools are more likely to use DirectX support – which professional GPUs (NVIDIA Quadro or AMD FirePro) are optimized for and tested extensively with.

Visualization/video tools are more likely to use OpenGL and be judged on high frame rates – parameters for which gaming GPUs (NVIDIA GeForce or AMD Radeon) are optimized.



GAMING GPU_s ARE EXPENSIVE RIGHT NOW

Recent trends in BitCoin and crypto currency mining have made gaming GPUs highly sought after, which has led to price increases.

It's no longer necessarily true that high-performance gaming GPUs are less expensive than professional GPUs.

ISV **CERTIFICATION**

PROFESSIONAL GPU_s HELP AVOID LOCKUPS & DRIVER ISSUES

Independent Software Vendors (ISVs) such as Autodesk, Dassault Systèmes, etc., test new products with professional GPUs (NVIDIA Quadro or AMD FirePro) to certify hardware performance.

Professional GPU drivers are generally updated as CAD software evolves and are thus less likely to experience problems than gaming GPUs.

Ultimately, professional GPUs are about avoiding CAD downtime.



CONCLUSION:

Software developers will recommend GPUs that they've tested to ensure optimal performance. From a compatibility and warranty standpoint, it is advised to go with recommendation for your software.

WRAP-UP **AND QUESTIONS**

FINAL THOUGHTS

- ▶▶ There is much to consider, but knowing your needs is paramount.
 - ▶▶ Buy what you need today, or one step above.
 - ▶▶ For 2D CAD, consider integrated graphics
 - ▶▶ For high-performance CAD, think professional discrete GPUs
 - ▶▶ For pure video playback or VR, think gaming GPUs

Q & A

▶▶ Submit a question using the Questions panel in your GoToWebinar interface.

Panelist: Robert Green
Moderator: Nancy Johnson



THANK YOU!

Watch your inbox for notification of the webinar recording.
For webinar archives, other CAD management resources, or
to subscribe to Robert Green's *CAD Manager's Newsletter*, see
www.cadalyst.com/RGwebinar.

cadalyst