



# NVIDIA MAXIMUS TECHNOLOGY FOR ANSYS MECHANICAL

DU-06467-001\_v01 | September 2012

User Guide



# DOCUMENT CHANGE HISTORY

DU-06467-001\_v01

Version	Date	Authors	Description of Change
01	August 3, 2012		Initial release
02	September 10, 2012	EK, SP	Minor edits, clarity regarding solver optimizations

# TABLE OF CONTENTS

- Maximus Technology for ANSYS Mechanical** ..... 4
  - Prerequisite Skills ..... 4
  - What This Document Contains..... 5
  - Benefits of Maximus Technology ..... 5
  - OEMs for Maximus Technology ..... 5
- NVIDIA Maximus Technology** ..... 6
  - Maximus Computing Advantage ..... 6
  - Elements of Maximus Technology ..... 6
  - Basic Maximus Configuration ..... 8
- Maximus for ANSYS Mechanical** ..... 9
  - Enabling ANSYS Mechanical for Maximus ..... 9
    - Monitoring GPU Activity with Maximus Configuration Utility..... 11
    - Supported Solver Types ..... 12
    - Model Considerations ..... 12
  - Troubleshooting..... 13
  - References..... 13

# MAXIMUS TECHNOLOGY FOR ANSYS MECHANICAL

This document describes the basic settings, configuration, and monitoring of an NVIDIA® Maximus-enabled workstation for ANSYS Mechanical solvers.



This document does not replace any documentation provided by ANSYS for software offerings specific to CAE software. Refer to the documentation provided by ANSYS for ANSYS software configurations.

This document does not explain the fundamentals of ANSYS usage or the discipline of CAE.

## PREREQUISITE SKILLS

This document is intended for persons responsible for optimizing a Maximus-enabled workstation for ANSYS Mechanical. It is assumed the audience is familiar with, or has skilled experience with the following:

- ▶ ANSYS Mechanical CAE Software
- ▶ Computer Aided Engineering (CAE)
- ▶ Graphics Processing Unit (GPU) functionality
- ▶ Modern workstation terminology
- ▶ Hardware connectivity
- ▶ Physical system building skills
- ▶ Thermal and electrical workstation system internals
- ▶ Microsoft Windows configuration

## WHAT THIS DOCUMENT CONTAINS

This document provides an introduction to NVIDIA Maximus technology and how to enable your workstation and ANSYS Mechanical to use Maximus.

- ▶ *NVIDIA Maximus Technology* starting on page 6 describes the benefits of NVIDIA Maximus technology, its key elements, and the basic system requirements for enabling Maximus on your workstation.
- ▶ *Maximus for ANSYS Mechanical* starting on page 9 focuses on using Maximus with ANSYS Mechanical. It explains how to enable Maximus for ANSYS Mechanical and monitor GPU activity. It identifies the solvers that are enabled for Maximus, describes model considerations, and troubleshooting tips when a solver workload does not perform to your expectations.
- ▶ *References* on page 13 provides useful references to related documentation.

## BENEFITS OF MAXIMUS TECHNOLOGY

For a comprehensive overview of Maximus technology, its benefits, and how it is being used, go to <http://www.nvidia.com/maximus>.

## OEMS FOR MAXIMUS TECHNOLOGY

A list of OEMs that carry Maximus platforms are listed at <http://www.nvidia.com/maximus>.

# NVIDIA MAXIMUS TECHNOLOGY

This section describes the benefits of Maximus-enabled workstations and applications, the key elements of Maximus technology, and the basic configuration for a Maximus-enabled workstation.

## MAXIMUS COMPUTING ADVANTAGE

In the past, workstation architectures forced professionals to do graphics-intensive and compute-intensive work serially; often offline. NVIDIA Maximus technology represents a revolution for these professionals by enabling both tasks to be performed concurrently without experiencing any drop in performance.

For example; a designer can work on design iteration **B** while running a simulation on design iteration **A**. Because these tasks are performed concurrently, it is possible to explore ideas faster and converge more quickly on the best possible answers.

## ELEMENTS OF MAXIMUS TECHNOLOGY

Maximus is an enabling technology that brings together the professional 3D graphics capability of NVIDIA Quadro® GPUs with the massive parallel computing capabilities of the NVIDIA Tesla™ C2075 companion processors. Figure 1 illustrates the advantages of the Tesla processor.

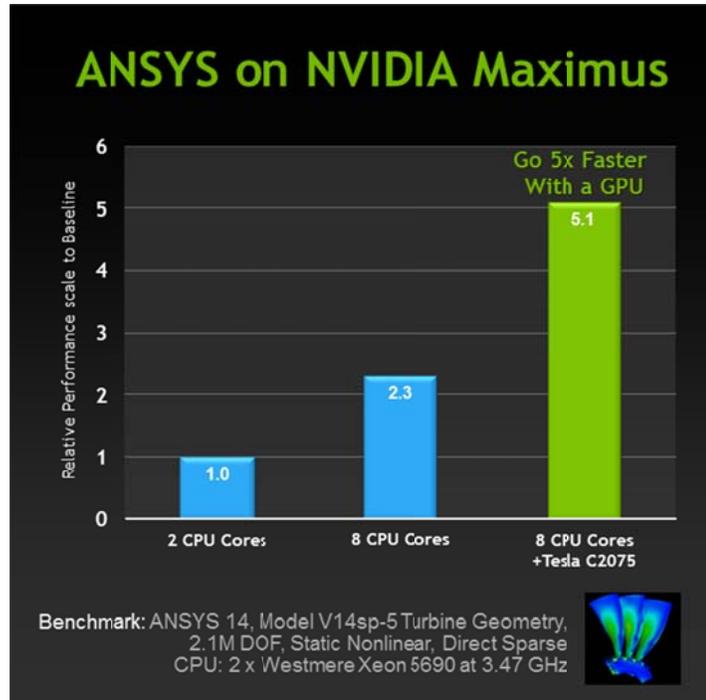


Figure 1. Tesla Performance

Figure 2 shows the performance improvements possible with ANSYS Mechanical on a Maximus-enabled workstation. The Maximus configuration can consist of any of the Quadro cards in the performance scaling chart plus one Tesla card.

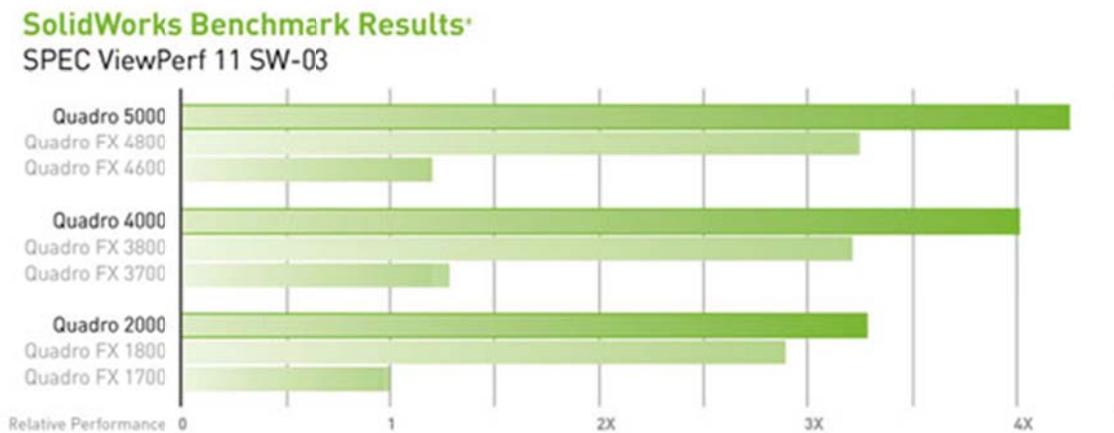


Figure 2. SolidWorks Scaling Chart

## BASIC MAXIMUS CONFIGURATION

This section describes hardware and software requirements for a Maximus-enabled workstation running ANSYS Mechanical. Only the basic requirements are covered in this document. For further details about upgrading an eligible workstation to a Maximus configuration, refer to the *NVIDIA Maximus System Builders' Guide for Microsoft Windows 7-64* document.

Check that your system satisfies the following software and hardware requirements:

- ▶ Microsoft Windows 7 – 64 bit operating system.
- ▶ ANSYS Mechanical 14 with HPC License Pack. At the release date of this document, one HPC pack enables eight CPU cores and one entire GPU.
- ▶ One NVIDIA Quadro card installed in the first x16 (x16 electrical) PCIe slot of the host computer.
- ▶ One NVIDIA Tesla card installed in the second x16 (x16 electrical) PCIe slot of the host computer



After ANSYS Mechanical becomes *multi-GPU aware* and if it is physically possible, you can have more than one Tesla processor in your system. At the release date of this document, ANSYS Mechanical is still *single-GPU aware*. Follow future announcements from ANSYS and NVIDIA regarding multi-GPU awareness.

At the release date of this document, Tesla C2075 is the only supported compute card for ANSYS Mechanical.

NVIDIA recommends that no display device be connected to the Tesla C2075 DVI display output.

- ▶ NVIDIA Quadro/Tesla Driver 275.89, or newer *ANSYS-certified driver*, correctly installed. Refer to <http://www.nvidia.com> for a list of drivers for download
- ▶ Correct installation and cabling with power connectors (as needed) of all NVIDIA graphics cards. If you purchased your system from an OEM with the NVIDIA cards pre-installed, no action is needed.



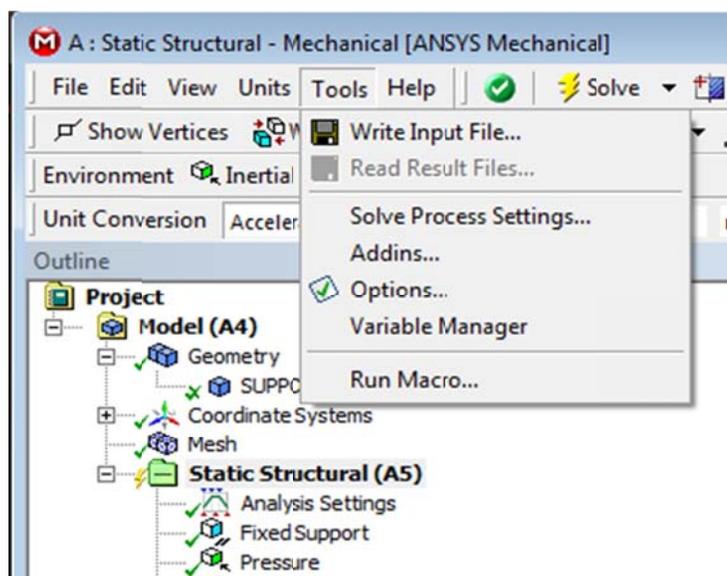
No NVIDIA SLI (Scalable Link Interface) ribbon cable is necessary or required for a Maximus configuration.

This section explains how to enable Maximus for ANSYS Mechanical and monitor GPU activity. The solvers that can use Maximus are listed and characteristics of models that need to be considered are identified. The section also contains trouble-shooting tips for a solver workload that does not perform to your expectations.

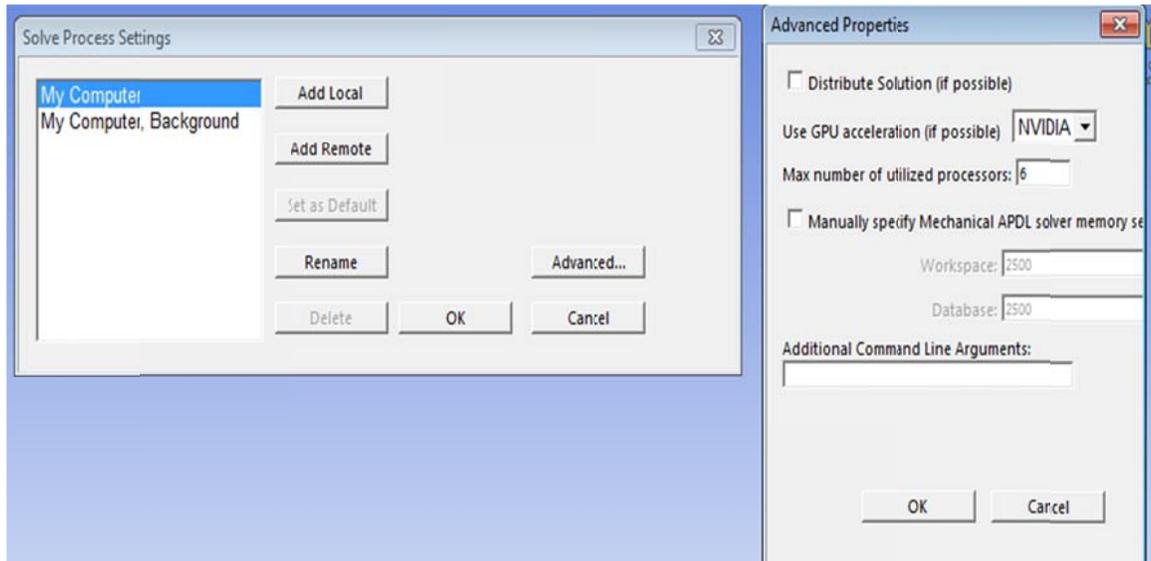
## ENABLING ANSYS MECHANICAL FOR MAXIMUS

Use the following procedure to enable Maximus for ANSYS Mechanical:

1. Select Tools from the main menu
2. Select Solve Process Settings... to display the Solve Process Settings menu.



3. Check that My Computer is selected on the *Solve Process Settings* menu.
4. Click Advanced... to display the *Advanced Properties* dialog menu.



5. Select NVIDIA from the drop-down list of the Use GPU acceleration (if possible) field.
6. Click OK.
7. Open the ANSYS Mechanical **solve.out** file to check that Maximus is enabled. Search for the text **GPU ACCELERATOR OPTION ENABLED**. If this text string is not displayed, Maximus is not enabled.

```

* software license agreement and FAR 12.212 (for non-DOD      *
* licenses).                                                  *
*                                                              *
*****
          ***** ANSYS COMMAND LINE ARGUMENTS *****
BATCH MODE REQUESTED (-b)   = NOLIST
INPUT FILE COPY MODE (-c)  = COPY
      6 PARALLEL CPUS REQUESTED
START-UP FILE MODE         = NOREAD
STOP FILE MODE             = NOREAD
GPU ACCELERATOR OPTION ENABLED

00000000          VERSION=WINDOWS x64          RELEASE= 13.0          UP20101012
CURRENT JOBNAME=file 14:22:46  OCT 24, 2011 CP=          0.811
    
```

Working together with ANSYS Mechanical software, the Maximus driver automatically ensures that ANSYS Mechanical runs on the Tesla GPU. This setting provides the best performance.

## Monitoring GPU Activity with Maximus Configuration Utility

🗨️ NVIDIA Maximus Configuration Utility (MCU) is supported for Maximus-enabled workstations only. Download the MCU from <http://www.nvidia.com/maximus>. The MCU is accessible from in the NVIDIA Control Panel starting with the 304 release of the Quadro drivers.

The NVIDIA Maximus Configuration Utility (MCU) is a separate graphical software utility that provides convenient GPU processing controls. The MCU provides GPU memory and utilization monitors for all supported GPUs in a Maximus-enabled system. Typically, MCU is used to ensure that ANSYS Mechanical is using the system GPUs correctly. Figure 3 shows the MCU menu page.

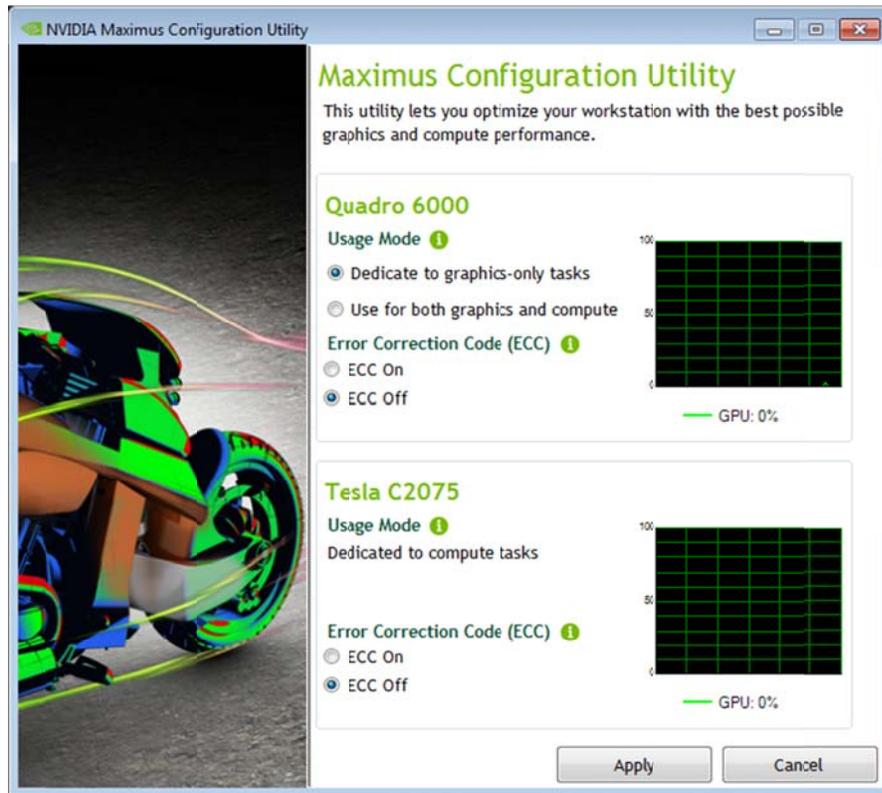


Figure 3. NVIDIA Maximus Configuration Utility

In typical ANSYS Mechanical workflows, you have the option to turn ECC (memory error correction) off. Note that when ECC is on, available memory on the Tesla board is reduced by 13 percent.

The MCU provides simple controls to enable or disable computational processing on an installed Quadro GPU. Use this feature to better tune the system for a particular workflow need.

## Supported Solver Types

ANSYS Mechanical uses a companion Tesla GPU with a single job per GPU. The following ANSYS Mechanical solvers are Maximus-enabled:

- ▶ Direct sparse solver for SMP and distributed ANSYS
- ▶ PCG and JCG iterative solvers for SMP and distributed ANSYS

## Model Considerations

Simulation models exist in a wide variety of sizes normally measured in degrees of freedom (DOF). Though all models may exhibit some level of GPU acceleration benefit, consideration should be given to the following factors to achieve maximum acceleration in ANSYS Mechanical:

- ▶ Models that deploy the direct sparse solver:
  - Models with approximately 500 thousand to 8 million DOF typically yield the most accelerated performance.
  - All model sizes are supported, but for very large models beyond 8 million DOF, some work might exceed the 6GB GPU memory and therefore stay on the CPU.
  - Models should always run in-core (system memory) to eliminate I/O.
- ▶ Models that deploy the iterative PCG or JCG solver:
  - Models with approximately 500 thousand to 5 million DOF typically yield the most accelerated performance.
  - Model size must not exceed the 6GB memory or they will not run, and will need to be restarted on CPUs-only (ANSYS Mechanical will provide a message in this case).
  - There is no out-of-core option (models always runs in-core for iterative solvers)
  - The MSAV option should be turned off, otherwise the GPU is deactivated. ANSYS Workbench will automatically set MSAVE for models over 100,000 nodes.



Solid structures always provide better performance than shell structures.

## TROUBLESHOOTING

There may be times when a solver workload does not perform to your expectations. Following is a list of common items that typically hinder optimal performance of a solver on a Maximus-enabled workstation:

- ▶ The Tesla C2075 is not set to handle compute tasks.
- ▶ ECC is turned ON for the Tesla C2075, or the Quadro 6000, or both. While enabling ECC improves accuracy, it slows down performance.
- ▶ The disk subsystem I/O rate is too slow or bandwidth is constrained.
- ▶ There is not enough scratch disk space in the system.
- ▶ The job unexpectedly runs out of core memory.
- ▶ There is not enough memory in the system.
- ▶ The ANSYS license does not provide support for Maximus.
- ▶ The simulation job is too small. See “Model Considerations” on page 10.
- ▶ MSAV option is set to ON.
- ▶ Shell structure models are being used.
- ▶ More CPUs are being used than necessary for the simulation job. More CPUs do not necessarily add linearly to overall wall clock time performance.

## REFERENCES

- ▶ ANSYS Support Documentation: Refer to <http://www.ansys.com/Support/Documentation>
- ▶ NVIDIA Maximus Configuration Guide: Refer to *NVIDIA Maximus System Builders' Guide for Microsoft Windows 7-64*

## Notice

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

## HDMI

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

## ROVI Compliance Statement

NVIDIA Products that support Rovi Corporation's Revision 7.1.L1 Anti-Copy Process (ACP) encoding technology can only be sold or distributed to buyers with a valid and existing authorization from ROVI to purchase and incorporate the device into buyer's products.

This device is protected by U.S. patent numbers 6,516,132; 5,583,936; 6,836,549; 7,050,698; and 7,492,896 and other intellectual property rights. The use of ROVI Corporation's copy protection technology in the device must be authorized by ROVI Corporation and is intended for home and other limited pay-per-view uses only, unless otherwise authorized in writing by ROVI Corporation. Reverse engineering or disassembly is prohibited.

## OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

## Trademarks

NVIDIA, the NVIDIA logo, Tesla, and Quadro are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

## Copyright

© 2012 NVIDIA Corporation. All rights reserved.