



CST STUDIO SUITE™ 2011

GPU Computing Guide



Contents

| | | |
|----------|---|-----------|
| 1 | Nomenclature | 3 |
| 2 | Technical Requirements | 3 |
| 2.1 | Supported Hardware | 3 |
| 2.1.1 | Tesla 20 (Fermi) Series | 4 |
| 2.1.2 | Tesla 8 and 10 Series | 8 |
| 3 | NVIDIA Drivers Download and Installation | 11 |
| 3.1 | Installation on Windows | 12 |
| 3.2 | Installation on Linux | 12 |
| 3.3 | Verifying correct installation of the Tesla Cards | 15 |
| 4 | Uninstalling NVIDIA Drivers | 16 |
| 4.1 | Uninstall Procedure on MS Windows | 16 |
| 4.2 | Uninstall Procedure on Linux | 16 |
| 5 | Switch on GPU Computing | 17 |
| 5.1 | Interactive Simulations | 17 |
| 5.2 | Simulations in Batch Mode | 18 |
| 6 | Usage Guidelines | 19 |
| 6.1 | Disabling the Error Correction Code (ECC) Feature | 19 |
| 6.2 | The Tesla Compute Cluster (TCC) Mode | 19 |
| 6.2.1 | Enabling the TCC Mode | 19 |
| 6.2.2 | Disabling the TCC mode | 20 |
| 6.2.3 | Disabling the Exclusive Mode | 20 |
| 6.3 | Combined MPI Computing and GPU Computing | 21 |
| 6.4 | Service User | 21 |
| 6.5 | Running Multiple Simulations at the Same Time | 21 |
| 6.6 | Video Card Drivers | 21 |
| 6.7 | Operating Conditions | 21 |
| 6.8 | Latest CST Service Pack | 21 |
| 7 | Troubleshooting Tips | 21 |
| 7.1 | General | 21 |
| 7.2 | Known Issues | 22 |
| 8 | Unsupported Features | 22 |
| 9 | History of Changes | 23 |

1 Nomenclature

The following section explains the nomenclature used in this document.

- command** Commands you have to enter either on a command prompt (cmd on MS Windows[®] or your favorite shell on Linux) are typeset using typewriter fonts.

- <...>** Within commands the sections you should replace according to your environment are enclosed in "<...>". For example "<CST_DIR>" should be replaced by the directory where you have installed CST STUDIO SUITE[™] (e.g. "c:\Program Files\CST STUDIO SUITE").



Those icons indicate that the following section of the text applies only to a certain operating system:

 = MS Windows[®]

 = Linux

2 Technical Requirements

This section provides you with information about the hardware as well as the software requirements necessary to successfully perform CST simulations on your GPU hardware. Please ensure that your system is configured according to the following points to avoid problems during the setup or operation of CST STUDIO SUITE[™] on your system.

2.1 Supported Hardware

Please note CST STUDIO SUITE[™] currently only supports the usage of 1, 2 or 4 GPU cards.

The following tables contain some basic information about the hardware currently supported by the GPU Computing feature of CST STUDIO SUITE[™] , as well as the requirements for the host system equipped with the hardware. Please be aware that this document is not a recommendation for which hardware to buy. The hardware currently recommended by CST can be found in the FAQ section of the CST support website (FAQ No.3). To ensure compatibility, please consult your hardware vendor to obtain a list of supported motherboards or systems. Please note that a 64 bit computer architecture is required for GPU Computing.

2.1.1 Tesla 20 (Fermi) Series

CST STUDIO SUITE™ 2011 supports the following GPU devices of the Tesla 20 series:



- C2050/M2050
- C2070/M2070
- C2075/M2075 (support added in Service Pack 6)
- M2090 (support added in Service Pack 6)

The C2050, C2070, C2075 cards are supplied with a display link. The desired performance of the GPU card is however deteriorated when using this link, since the performance mode, called the TCC (Tesla Compute Cluster) mode of the GPU card is disabled when the link is used. It is therefore recommended to have an additional NVIDIA graphics card if you need to connect a monitor to the machine, or to use an onboard graphics chip for graphics output if the motherboard of your machine has such a device.

The following table highlights the Tesla 20 series of GPU cards that is supported in CST STUDIO SUITE™ 2011.

GPU Hardware Properties

Host System Requirements

| | | |
|---|---|---|
| Hardware Type |  NVIDIA Tesla M2090 |  NVIDIA Tesla C2075/M2075 (do not use the provided display link) |
| CST version required | 2011 SP 6 | 2011 SP 6 |
| Number of GPUs | 1 | 1 |
| Max. Problem Size | approx. 60 million mesh cells | approx. 60 million mesh cells |
| Size | 4.36" x 10.5" 11.1cm x 26.67cm two slot width | 4.36" x 10.5" 11.1cm x 26.67cm two slot width |
| Memory | 6 GB GDDR5 | 6 GB GDDR5 |
| Power Consumption | 225 W (max.) | 215 W (max.) requires two auxiliary power connectors |
| PCI Express Requirements | 1x PCIe Gen 2 (x16 electrically) | 1x PCIe Gen 2 (x16 electrically) |
| Power Supply of Host System | min. 750 W ¹ | min. 750 W ¹ |
| Min. RAM of Host System | 24 GB | 24 GB |
| Recommended Host System for the Use with the Hardware | Please ask your hardware vendor! | Please ask your hardware vendor! |
| Recommended OS | RedHat EL 5 (x64) Windows 7 (x64) Win Server 2008 R2 (x64) | RedHat EL 5 (x64) Windows 7 (x64) Win Server 2008 R2 (x64) |
| Vendor Information | www.nvidia.com | www.nvidia.com |

¹**Important:** The specifications shown assume that only one adapter is plugged into the machine. If you would like to plug in two or more adapters you will need a better power supply (1000W or above) as well as more RAM. Additionally, you need to provide a sufficient cooling for the machine. Each Tesla card takes power from the PCI Express host bus as well as the 8-pin and the 6-pin PCI Express power connectors. This is an important consideration while selecting power supplies.

CST assumes no liability for any problems caused by this information.



| | | | | |
|---|--|---|---|----------------------------------|
| GPU Hardware Properties | Hardware Type |  NVIDIA Tesla C2050/M2050 (do not use the provided display link) |  NVIDIA Tesla C2070/M2070 (do not use the provided display link) | |
| | CST version required | 2011 Release | 2011 Release | |
| | Number of GPUs | 1 | 1 | |
| | Max. Problem Size | approx. 30 million mesh cells | approx. 60 million mesh cells | |
| | Size | 4.36" x 10.5" 11.1cm x 26.67cm two slot width | 4.36" x 10.5" 11.1cm x 26.67cm two slot width | |
| | Memory | 3 GB GDDR5 | 6 GB GDDR5 | |
| | Power Consumption | 247 W (max.) requires two auxiliary power connectors | 247 W (max.) requires two auxiliary power connectors | |
| | Host System Requirements | PCI Express Requirements | 1x PCIe Gen 2 (x16 electrically) | 1x PCIe Gen 2 (x16 electrically) |
| | | Power Supply of Host System | min. 750 W ¹ | min. 750 W ¹ |
| | | Min. RAM of Host System | 12 GB | 24 GB |
| Recommended Host System for the Use with the Hardware | | Please ask your hardware vendor! | Please ask your hardware vendor! | |
| Recommended OS | | RedHat EL 5 (x64) Windows 7 (x64) Win Server 2008 R2 (x64) | RedHat EL 5 (x64) Windows 7 (x64) Win Server 2008 R2 (x64) | |
| Vendor Information | www.nvidia.com | www.nvidia.com | | |

¹**Important:** The specifications shown assume that only one adapter is plugged into the machine. If you would like to plug in two or more adapters you will need a better power supply (1000W or above) as well as more RAM. Additionally, you need to provide a sufficient cooling for the machine. Each Tesla C2050 or C2070 takes power from the PCI Express host bus as well as the 8-pin and the 6-pin PCI Express power connectors. This is an important consideration while selecting power supplies.

CST assumes no liability for any problems caused by this information.

GPU Hardware Properties



Host System Requirements

| | | |
|-----------------------------|---|---|
| Hardware Type |  NVIDIA Tesla S2050 NextIO vCore Express S2050 |  NextIO vCore Express S2070 |
| CST version required | 2011 Release | 2011 Release |
| Number of GPUs | 4 | 4 |
| Max. Problem Size | approx. 120 million mesh cells | approx. 240 million mesh cells |
| Size | 1U rack-mount system | 1U rack-mount system |
| Memory | 12 GB GDDR5 | 24 GB GDDR5 |
| Power Consumption | 1200 W (max.) | 1200 W (max.) |
| PCI Express Requirements | 2x PCIe Gen 2 (x16 electrically) | 2x PCIe Gen 2 (x16 electrically) |
| Power Supply of Host System | min. 750 W | min. 750 W |
| Recommended Host System | Please ask your hardware vendor! | Please ask your hardware vendor! |
| Recommended OS | RedHat EL 5 (x64) Win Server 2008 R2 (x64) | RedHat EL 5 (x64) Win Server 2008 R2 (x64) |
| Vendor Information | www.nvidia.com www.nextio.com | www.nvidia.com www.nextio.com |

CST assumes no liability for any problems caused by this information.

2.1.2 Tesla 8 and 10 Series

Notice: The older adapters based on the NVIDIA Tesla 8 and 10 series are still supported in CST STUDIO SUITE™ 2011. However, you need to upgrade the drivers to the version listed in section 3 of this document in order to use this hardware. Those drivers are compatible with CST STUDIO SUITE™ 2010 as well.

| | | | |
|--------------------------|--|---|--|
| GPU Hardware Properties | Hardware Type |  NVIDIA Quadro FX 5800 (provides display link) |  NVIDIA Tesla C1060/ M1060 (no display link) |
| | Number of GPUs | 1 | 1 |
| | Max. Problem Size | approx. 40 million mesh cells | approx. 40 million mesh cells |
| | Size | 4.36" x 10.5" 11.1cm x 26.67cm two slot width | 4.36" x 10.5" 11.1cm x 26.67cm two slot width |
| | Memory | 4 GB GDDR3 | 4 GB GDDR3 |
| | Power Consumption | 200 W (max.) requires two auxiliary power connectors | 200 W (max.) requires two auxiliary power connectors |
| Host System Requirements | PCI Express Requirements | 1x PCIe Gen 2 (x16 electrically) | 1x PCIe Gen 2 (x16 electrically) |
| | Power Supply of Host System | min. 750 W ¹ | min. 750 W ¹ |
| | Min. RAM of Host System | 12 GB | 12 GB |
| | Recommended Host System for the Use with the Hardware | Please ask your hardware vendor! | Please ask your hardware vendor! |
| | Recommended OS | RedHat EL 5 (x64) Windows 7 (x64) Win Server 2008 R2 (x64) | RedHat EL 5 (x64) Windows 7 (x64) Win Server 2008 R2 (x64) |
| | Vendor Information | www.nvidia.com | www.nvidia.com |

¹**Important:** The specifications shown assume that only one adapter is plugged into the machine. If you would like to plug in two or more adapters you will need a better power supply (1000W or above) as well as more RAM. Additionally, you need to provide a sufficient cooling for the machine. Each Tesla C1060 requires either two 6-pin power connectors or one 8-pin power connector. This is an important consideration while selecting power supplies.

CST assumes no liability for any problems caused by this information.

| | | | |
|--------------------------|---|--|---|
| GPU Hardware Properties | Hardware Type |  NVIDIA Quadro Plex 2200 D2 (provides display link) |  NVIDIA Tesla S1070 |
| | Number of GPUs | 2 | 4 |
| | Max. Problem Size | approx. 80 million mesh cells | approx. 160 million mesh cells |
| | Size | 20.55" x 5.94" x 9.49" 52cm x 15cm x 24cm | 1U rack-mount system |
| | Memory | 2x 4 GB GDDR3 | 4x 4 GB GDDR3 |
| Host System Requirements | Power Consumption | 640 W (max.) requires two auxiliary power connectors | 800 W (max.) requires two auxiliary power connectors |
| | PCI Express Requirements | 1x PCIe Gen 2 (x16 electrically) | 2x PCIe Gen 2 (x16 electrically) |
| | Power Supply of Host System | min. 750 W | min. 750 W |
| | Min. RAM of Host System | 24 GB | 48 GB |
| | Recommended Host System for the Use with the Hardware | Please ask your hardware vendor! | Please ask your hardware vendor! |
| | Recommended OS | RedHat EL 5 (x64) Win Server 2008 R2 (x64) Windows 7 (x64) | RedHat EL 5 (x64) Win Server 2008 R2 (x64) |
| | Vendor Information | www.nvidia.com | www.nvidia.com |

CST assumes no liability for any problems caused by this information.

3 NVIDIA Drivers Download and Installation

An appropriate driver is required in order to use the GPU hardware. Please download the driver appropriate to your GPU hardware and operating system using the links in the following table. Those drivers are verified for use with our software. Other driver versions provided by NVIDIA might also work but it is highly recommended to use the versions verified by CST.

| GPU Type | <u>Win XP</u> <u>Win Server 2003 R2</u> | Win Server 2008 R2 | <u>Win Vista</u> <u>Win 7</u> | RHEL 4.x/5.x |
|--------------------|--|--------------------------|----------------------------------|--------------------------|
| FX5600/C870 | download | not available | not available | download |
| QuadroPlex 1000 IV | download | not available | not available | download |
| C1060/M1060 | download | download | download | download |
| FX5800 | download | download | download | download |
| QuadroPlex 2200 D2 | download | download | download | download |
| C2050/M2050 | not supported | download | download | download |
| C2070/M2070 | not supported | download | download | download |
| S1070 | download | download | not supported | download |
| S2050/S2070 | not supported | download | not supported | download |
| C2075/M2075/M2090 | not supported | download | download | download |

Please note:

- The statement "not supported" in the table above means that the combination of operating system and hardware is not supported by CST for some reason (e.g. because NVIDIA does not provide drivers for the configuration).
- The statement "not available" in the table above means that there are currently no drivers available from NVIDIA which work well together with CST STUDIO SUITE™ for the corresponding configuration. Thus, the configuration is not available.

3.1 Installation on Windows

After you have downloaded the installer executable please start the installation procedure by double clicking on the installer executable. After a quick series of pop-up windows, the NVIDIA InstallShield Wizard will appear. Press the "Next" button and driver installation will begin (The screen may turn black momentarily.). You may receive a message indicating that the hardware has not passed Windows logo testing (see fig. 1). In case you get this warning select "Continue Anyway". The "Wizard Complete" window will appear.



Figure 1: Warning regarding Windows Logo test

Select "Yes, I want to restart my computer now" and click the "Finish" button. It is recommended that you run the HWAccDiagnostics tool after the installation to confirm that the driver has been successfully installed. Please use HWAccDiagnostics_AMD64.exe which can be found in the AMD64 directory of the installation folder.

3.2 Installation on Linux

1. Login on the Linux machine as root.
2. Make sure that the adapter has been recognized by the system using the command `/sbin/lspci | grep -i nvidia`
If you do not see any settings try to update the PCI hardware database of your system using the command `/sbin/update-pciids`
3. Stop the X-Server by running in a terminal the command (You may skip this step if you are working on a system without X-server)
`telinit 3`

4. Install the NVIDIA graphics driver. Follow the instructions of the setup script. In most cases the installer needs to compile a specific kernel module. If this is the case the gcc compiler and Linux kernel headers need to be available on the machine.
5. Restart the X-server by running the command (You may skip this step if you are working on a system without X-server)
`telinit 5`
6. You may skip this step if a X-server is installed on your system. If no X-server is installed on your machine the NVIDIA kernel module will not be loaded automatically. Additionally, the device files for the GPUs will not be generated automatically. The following commands will perform the necessary steps to use the hardware for GPU Computing. It is recommended to append this code to your `rc.local` file such that it is executed automatically during system start.

```
# Load nvidia kernel module
modprobe nvidia

if [ "$?" -eq 0 ]; then

    # Count the number of NVIDIA controllers found.
    N3D=$(/sbin/lspci | grep -i NVIDIA | grep "3D controller" | wc -l)
    NVGA=$(/sbin/lspci | grep -i NVIDIA | grep "VGA compatible controller" | wc -l)

    N=$(expr $N3D + $NVGA - 1)
    for i in $(seq 0 $N); do
        mknod -m 666 /dev/nvidia$i c 195 $i;
    done

    mknod -m 666 /dev/nvidiactl c 195 255

fi
```

Please note:

- When asked whether you would like to install the "32-bit compatibility OpenGL libraries" please select "yes" as some components of the CST software need those 32-bit libraries.


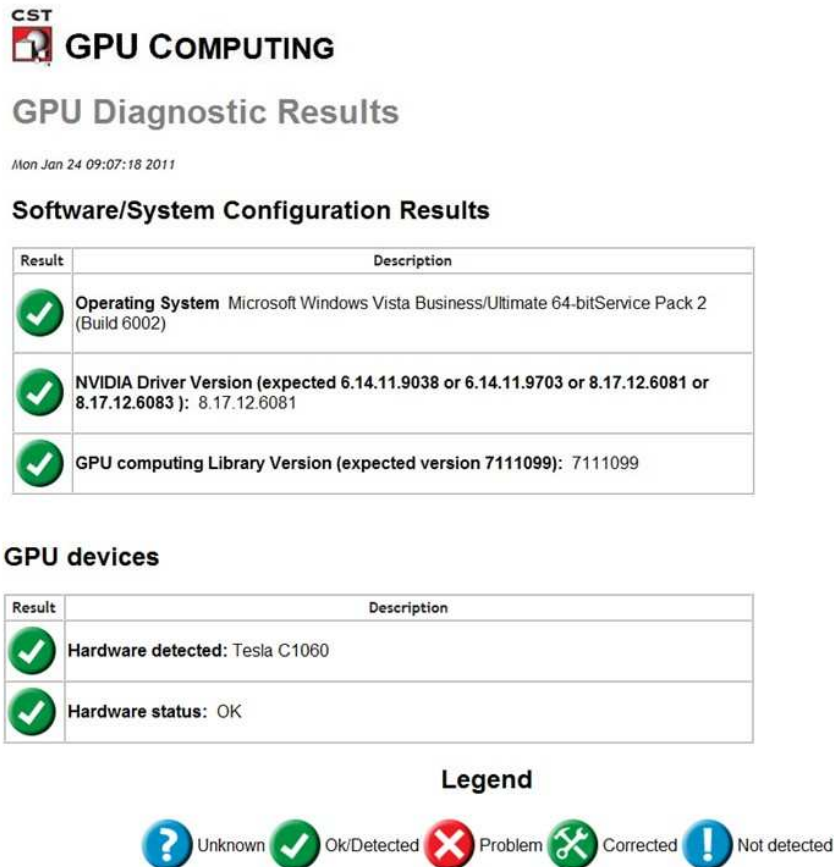


Figure 2: Install 32-bit compatible OpenGL libraries

- If you encounter problems during restart of the X-server please check chapter 8 "Common Problems" in the file README.txt located at `/usr/share/doc/NVIDIA_GLX-1.0`. Please also consider removing existing sound cards or deactivating onboard sound in the BIOS. Furthermore, make sure you are running the latest BIOS version.
- After installation, if the X system reports an error like `no screen found`, please check Xorg log files in `/var/log`. Open the log files in an editor and search for "PCI". According to the number of hardware cards in your system you will find entries of the following form: `PCI: (0@7:0:0)`. In `/etc/X11`, open the file `xorg.conf` in an editor and search for "nvidia". After the line `BoardName "Quadro FX 5800"` (or whatever card you are using) insert a new line that reads `BusID "PCI:7:0:0"` according to the entries found in the log files before. Save and close the `xorg.conf` file and type `startx`. If X still refuses to start, try the other entries found in the Xorg log files.
- You need the installation script to uninstall the driver. Thus, if you want to be able to uninstall the NVIDIA software you need to keep the installer script.
- Be aware of the fact that you need to reinstall the NVIDIA drivers if your kernel is updated as the installer needs to compile a new kernel module in this case.

3.3 Verifying correct installation of the Tesla Cards

As a final test to verify that all Tesla Cards have been correctly installed, the following test can be executed: Log in to the machine and execute the HWAccDiagnostics_AMD64 program found in the AMD64 subfolder of your CST installation (Windows) or in the folder LinuxAMD64 on a Linux system. The output of the tool should look similar to the following picture if the installation was successful.






CST GPU COMPUTING



GPU Diagnostic Results

Mon Jan 24 09:07:18 2011

Software/System Configuration Results

| Result | Description |
|---|---|
|  | Operating System Microsoft Windows Vista Business/Ultimate 64-bitService Pack 2 (Build 6002) |
|  | NVIDIA Driver Version (expected 6.14.11.9038 or 6.14.11.9703 or 8.17.12.6081 or 8.17.12.6083): 8.17.12.6081 |
|  | GPU computing Library Version (expected version 7111099): 7111099 |

GPU devices

| Result | Description |
|---|---------------------------------------|
|  | Hardware detected: Tesla C1060 |
|  | Hardware status: OK |

Legend






 Unknown  Ok/Detected  Problem  Corrected  Not detected

Figure 3: Output of HWAccDiagnostics_AMD64.exe tool.

4 Uninstalling NVIDIA Drivers

4.1 Uninstall Procedure on MS Windows

To uninstall NVIDIA drivers, select "NVIDIA Drivers" from the "Add or Remove Programs" list and press the "Change/Remove" button (see fig. 4). After the uninstall process has finished you will be prompted to reboot.

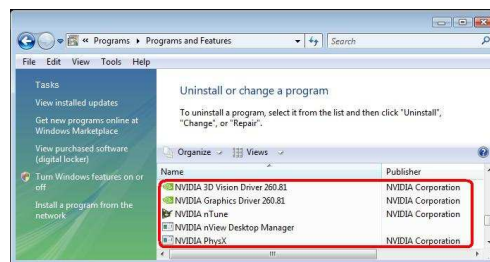


Figure 4: "Add or Remove Programs" dialog on Windows

4.2 Uninstall Procedure on Linux

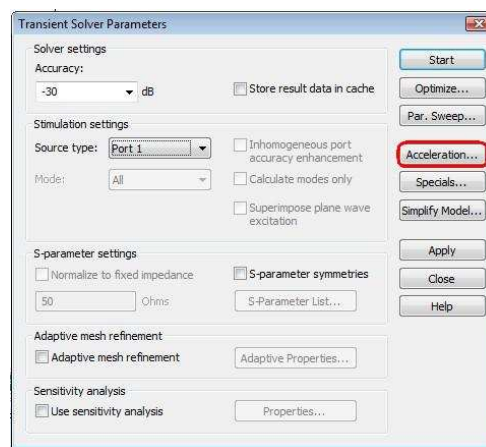
Start the installer with the "--uninstall" option. This requires root permissions.

5 Switch on GPU Computing

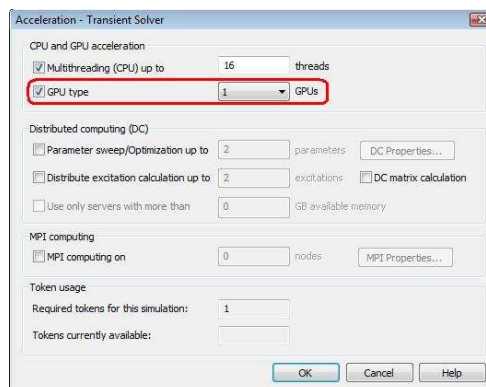
5.1 Interactive Simulations

GPU Computing needs to be enabled via the acceleration dialog box of CST MICROWAVE STUDIO[®] interface before running a simulation. To turn on GPU Computing:

1. Open the dialog of the solver. (Solve→Transient Solver).
2. Click on the Acceleration button.



3. Switch on "GPU Computing" and specify how many GPUs should be used for this simulation. Please note that the maximum number of GPUs available for a simulation depends upon the number of tokens in your license.



5.2 Simulations in Batch Mode

If you start your simulations in batch mode (e.g. via an external job queuing system) there is a command line switch (`-withgpu`) which can be used to switch on the GPU Computing feature. The command line switch can be used as follows:

In Windows:

```
"<CST_INSTALL_DIR>/CST Design Environment.exe" -m -r -withgpu=<NUMBER_OF_GPUS> "<FULL_PATH_TO_CST_FILE>"
```

In Linux:

```
"<CST_INSTALL_DIR>/cst.design.environment" -m -r -withgpu=<NUMBER_OF_GPUS> "<FULL_PATH_TO_CST_FILE>"
```

6 Usage Guidelines

6.1 Disabling the Error Correction Code (ECC) Feature

The Tesla 20 hardware comes with an ECC feature. However, this feature deteriorates the performance of the GPU hardware. Therefore, we recommend disabling the feature. The ECC feature can be disabled using the command line. Please note, that on Windows Vista, Windows 7 and Windows Server 2008 R2, the following commands have to be run as administrator.

1. Locate the file `nvidia-smi`. This file is typically found in "`c:\Program Files\NVIDIA Corporation\NVSMI`" or in `/usr/bin` on Linux.
2. Open up a command prompt/terminal window and navigate to this folder.
3. Execute the following command:
`nvidia-smi -L`
4. Please note down how many GPUs are found.
5. For each of the GPUs, please execute the following command:
`nvidia-smi -g <number_of_the_GPU_card> -e 0`
6. Reboot.

6.2 The Tesla Compute Cluster (TCC) mode

6.2.1 Enabling the TCC Mode

When available, the GPUs have to operate in TCC mode¹. Please enable the mode, if not yet enabled. Please note, that on Windows Vista, Windows 7 and Windows Server 2008 R2, the following commands have to be run "as administrator".

1. Locate the file `nvidia-smi.exe`. This file is typically found in "`c:\Program Files\NVIDIA Corporation\NVSMI`".
2. Open up a command prompt and navigate to this folder.
3. Execute the following command:
`nvidia-smi -L`
4. Please note down how many GPUs are found.
5. For each of the GPUs, please execute the following command:
`nvidia-smi -g <number_of_the_GPU_card> -dm 1`
6. Reboot.

¹The TCC mode is available on all Tesla cards. Graphics cards, such as the NVIDIA Quadro cards (FX 5600 and FX 5800) and NVIDIA Quadro Plex cards (1000 IV and 2200 D2) do not have this mode.

6.2.2 Disabling the TCC mode

If available, this feature should always be enabled. However, under certain circumstances you may need to disable this mode.

1. Locate the file `nvidia-smi.exe`. This file is typically found in `"c:\Program Files\NVIDIA Corporation\NVSMI"`.
2. Open up a command prompt and navigate to this folder.
3. Execute the following command:
`nvidia-smi -L`
4. Please note down how many GPUs are found.
5. For each of the GPUs, please execute the following command:
`nvidia-smi -g <number_of_the_GPU_card> -dm 0`
6. Reboot.

6.2.3 Disabling the Exclusive Mode

This mode has to be disabled in order to use CST STUDIO SUITE™ . To test if this mode is switched on, please do the following:

1. Locate the file `nvidia-smi.exe`. This file is typically found in `"c:\Program Files\NVIDIA Corporation\NVSMI"` or in `/usr/bin` on Linux.
2. Open up a command prompt and navigate to this folder.
3. Execute the following command:
`nvidia-smi -L`

If the response to the above command is "0", then the card is not being used in an exclusive mode. If the response to the above command is "1", then the card is being operated in the exclusive mode. In this case, please execute the following commands in order to disable this mode:

1. Locate the file `nvidia-smi.exe`. This file is typically found in `"c:\Program Files\NVIDIA Corporation\NVSMI"` or `/usr/bin` on Linux.
2. Open up a command prompt and navigate to this folder.
3. Execute the following command:
`nvidia-smi -L`
4. Please note down how many GPUs are found.
5. For each of the GPUs, please execute the following command:
`nvidia-smi -g <number_of_the_GPU_card> -c 0`
6. There is no need to reboot.

6.3 Combined MPI Computing and GPU Computing

For combined MPI Computing and GPU Computing the TCC mode of the GPU hardware must be enabled (see [6.2](#)).

6.4 Service User

If you are using GPU Computing via the CST Distributed Computing system and your DC Solver Server runs Windows Vista/Server 2008 R2/7 then the DC Solver Server service must be started using the local administrator account (the CST STUDIO SUITE™ installer installs the service by default using the correct account).

6.5 Running Multiple Simulations at the Same Time

GPU Computing does not support running multiple simulations at the same time.

6.6 Video Card Drivers

Please use only the drivers recommended in this document or by the hardware diagnostics tool (See section [3.3](#)). They have been tested for compatibility with CST AG products.

6.7 Operating Conditions

CST AG recommends that GPU Computing is operated in a well ventilated temperature controlled area. For more information, please contact your hardware vendor.

6.8 Latest CST Service Pack

Download and install the latest CST Service Pack prior to running a simulation or HWAccDiagnostics.

7 Troubleshooting Tips

7.1 General

The following troubleshooting tips may help if you experience problems.

- If you experience problems during the installation of the NVIDIA driver on the Windows operating system please try to boot Windows in "safe mode" and retry the driver installation.
- Please note that CST STUDIO SUITE™ cannot run on GPU devices when they are in "exclusive mode". Please refer to section [6.2.3](#) on how to disable this mode.

- If you are using an external GPU device (e.g. S2050) ensure that the PCI connector cable is securely fastened to the host interface card.
- Uninstall video drivers for your existing graphics adapter prior to installing the new graphics adapter.
- Make sure the latest motherboard BIOS is installed on your machine. Please contact your hardware vendor support for more information about the correct BIOS version for your machine.
- Use the HWAccDiagnostics tool to find out whether your GPU hardware and your driver is correctly recognized.

7.2 Known Issues

- The CUDA 3.2 runtime may cause problems on some Windows 7, Windows Server 2008 R2, or Windows Vista systems. In many cases this has been observed when the host computer had 48 GB or more RAM. The GPU setup fails at an early stage. The drivers listed in section 3 have been verified to not show this problem. We recommend strongly to download the drivers listed in this section.

If you have tried the points above with no success please contact CST technical support (info@cst.com).

8 Unsupported Features

Currently the transient solver of CST MICROWAVE STUDIO® can benefit from the GPU Computing feature. However, the following features of the transient solver are currently not supported for GPU Computing:

- Subgridding
- Lossy Boundaries
- Gyrotropic Materials
- Waveguide Ports Touching Open Boundaries
- Tabulated Surface Impedance
- Ohmic Sheet
- Corrugated Wall

9 History of Changes

The following changes have been applied to the document in the past.

| Date | Description |
|--------------|---|
| Jan. 24 2011 | 2011 Version completed |
| Feb. 21 2011 | 6.2 applies only to Windows OS |
| Feb. 22 2011 | Removed Tesla 8 + Win Server 2008 R2 combination because of not functioning drivers. |
| Feb. 24 2011 | New sub section on Known Issues (7.2) was added. Problem with Cuda 3.2 described. |
| Mar. 07 2011 | Extended the ECC section (6.1) and TCC sections (6.2). Filled in which OS/Driver combinations are not supported by CST |
| Mar. 10 2011 | Updated the drivers list. The drivers now include the 267.17 driver (Windows) and the 19.36 driver for Linux |
| Mar. 31 2011 | Added instructions on how to disable the exclusive mode of the GPUs |
| Apr. 13 2011 | Included a link to FAQ 349, which is a workaround for the CUDA 3.2/Windows 7 problem |
| Apr. 13 2011 | Removed the asterisk, which indicates that the M2070 has now been fully verified to run on a Windows Server 2008 R2 OS |
| May. 02 2011 | <p>Added Tabulated surface impedance, Ohmic sheet and Corrugated Wall to the list of unsupported features</p> <p>Removed warnings about onboard graphics chipsets in several places.</p> <p>Added NextIO vCore Express S2050/S2070 devices to the table of supported Tesla 20 cards.</p> <p>Added equivalent Linux commands to several sections, where previously only Windows instructions were present.</p> |

| Date | Description (Cont.) |
|---------------|---|
| May. 09 2011 | <p>Removed references to the NVIDIA control panel. Corrected a few style mistakes. Added "safe mode" installation as an option in the Troubleshooting Tips section. In the "Known Issues" section, the CUDA 3.2 runtime error has now been expanded to other Windows operating systems. Removed the detailed information on overcoming the CUDA 3.2 runtime error, as this is dealt with in an FAQ. C2050 is now verified to work with Win Server 2008 R2.</p> |
| Juni. 10 2011 | <p>Added a note on linux installation and 32 bit OpenGL libraries</p> |
| July. 07 2011 | <p>Updated the recommended driver for the S2050/S2070 and Win Server 2008 R2. This should fix the problem seen with more than 32 GB of memory. Also updated the Known Issues section to explain that this problem can be resolved through upgrading of the drivers</p> |
| Aug. 05 2011 | <p>Made several smaller changes and updated the driver list.</p> |
| Oct. 13 2011 | <p>Added 2075/2090 GPUs to list of supported hardware.</p> |

射频和天线设计培训课程推荐

易迪拓培训(www.edatop.com)由数名来自于研发第一线的资深工程师发起成立,致力并专注于微波、射频、天线设计研发人才的培养;我们于 2006 年整合合并微波 EDA 网(www.mweda.com),现已发展成为国内最大的微波射频和天线设计人才培养基地,成功推出多套微波射频以及天线设计经典培训课程和 ADS、HFSS 等专业软件使用培训课程,广受客户好评;并先后与人民邮电出版社、电子工业出版社合作出版了多本专业图书,帮助数万名工程师提升了专业技术能力。客户遍布中兴通讯、研通高频、埃威航电、国人通信等多家国内知名公司,以及台湾工业技术研究院、永业科技、全一电子等多家台湾地区企业。

易迪拓培训课程列表: <http://www.edatop.com/peixun/rfe/129.html>



射频工程师养成培训课程套装

该套装精选了射频专业基础培训课程、射频仿真设计培训课程和射频电路测量培训课程三个类别共 30 门视频培训课程和 3 本图书教材;旨在引领学员全面学习一个射频工程师需要熟悉、理解和掌握的专业知识和研发设计能力。通过套装的学习,能够让学员完全达到和胜任一个合格的射频工程师的要求...

课程网址: <http://www.edatop.com/peixun/rfe/110.html>

ADS 学习培训课程套装

该套装是迄今国内最全面、最权威的 ADS 培训教程,共包含 10 门 ADS 学习培训课程。课程是由具有多年 ADS 使用经验的微波射频与通信系统设计领域资深专家讲解,并多结合设计实例,由浅入深、详细而又全面地讲解了 ADS 在微波射频电路设计、通信系统设计和电磁仿真设计方面的内容。能让您在最短的时间内学会使用 ADS,迅速提升个人技术能力,把 ADS 真正应用到实际研发工作中去,成为 ADS 设计专家...



课程网址: <http://www.edatop.com/peixun/ads/13.html>



HFSS 学习培训课程套装

该套课程套装包含了本站全部 HFSS 培训课程,是迄今国内最全面、最专业的 HFSS 培训教程套装,可以帮助您从零开始,全面深入学习 HFSS 的各项功能和在多个方面的工程应用。购买套装,更可超值赠送 3 个月免费学习答疑,随时解答您学习过程中遇到的棘手问题,让您的 HFSS 学习更加轻松顺畅...

课程网址: <http://www.edatop.com/peixun/hfss/11.html>

CST 学习培训课程套装

该培训套装由易迪拓培训联合微波 EDA 网共同推出,是最全面、系统、专业的 CST 微波工作室培训课程套装,所有课程都由经验丰富的专家授课,视频教学,可以帮助您从零开始,全面系统地学习 CST 微波工作的各项功能及其在微波射频、天线设计等领域的设计应用。且购买该套装,还可超值赠送 3 个月免费学习答疑...

课程网址: <http://www.edatop.com/peixun/cst/24.html>



HFSS 天线设计培训课程套装

套装包含 6 门视频课程和 1 本图书,课程从基础讲起,内容由浅入深,理论介绍和实际操作讲解相结合,全面系统的讲解了 HFSS 天线设计的全过程。是国内最全面、最专业的 HFSS 天线设计课程,可以帮助您快速学习掌握如何使用 HFSS 设计天线,让天线设计不再难...

课程网址: <http://www.edatop.com/peixun/hfss/122.html>

13.56MHz NFC/RFID 线圈天线设计培训课程套装

套装包含 4 门视频培训课程,培训将 13.56MHz 线圈天线设计原理和仿真设计实践相结合,全面系统地讲解了 13.56MHz 线圈天线的工作原理、设计方法、设计考量以及使用 HFSS 和 CST 仿真分析线圈天线的具体操作,同时还介绍了 13.56MHz 线圈天线匹配电路的设计和调试。通过该套课程的学习,可以帮助您快速学习掌握 13.56MHz 线圈天线及其匹配电路的原理、设计和调试...

详情浏览: <http://www.edatop.com/peixun/antenna/116.html>



我们的课程优势:

- ※ 成立于 2004 年,10 多年丰富的行业经验,
- ※ 一直致力并专注于微波射频和天线设计工程师的培养,更了解该行业对人才的要求
- ※ 经验丰富的一线资深工程师讲授,结合实际工程案例,直观、实用、易学

联系我们:

- ※ 易迪拓培训官网: <http://www.edatop.com>
- ※ 微波 EDA 网: <http://www.mweda.com>
- ※ 官方淘宝店: <http://shop36920890.taobao.com>