



Getting Started

NVIDIA CUDA 2.2 Installation and Verification on Mac OS X

Getting Started with CUDA

Table of Contents

Chapter 1. Introduction	1
CUDA—Supercomputing on Desktop Systems	1
System Requirements	2
About This Document.....	2
Chapter 2. Installing CUDA	3
Verify CUDA-Capable GPU	3
Download the CUDA Software.....	4
Verify the Correct Version of Mac OS X	4
Verify That gcc Is Installed	5
CUDA Software	5
Installing the Software.....	5
Verify the Installation.....	6
Compiling the Examples.....	6
Running the Binaries	6
Chapter 3. Additional Considerations	9
Compiling for Hardware Emulation.....	9
What’s Next?.....	9

Getting Started with CUDA



Chapter 1. Introduction

CUDA—Supercomputing on Desktop Systems

NVIDIA® CUDA™ is a general purpose parallel computing architecture introduced by NVIDIA. It includes the CUDA Instruction Set Architecture (ISA) and the parallel compute engine in the GPU. To program to the CUDA architecture, developers can, today, use C, one of the most widely used high-level programming languages, which can then be run at great performance on a CUDA enabled processor.

The CUDA architecture and its associated software were developed with several design goals in mind:

- ❑ Provide a small set of extensions to standard programming languages, like C, that enable a straightforward implementation of parallel algorithms. With CUDA and C for CUDA, programmers can focus on the task of parallelization of the algorithms rather than spending time on their implementation.
- ❑ Support heterogeneous computation where applications use both the CPU and GPU. Serial portions of applications are run on the CPU, and parallel portions are offloaded to the GPU. As such, CUDA can be incrementally applied to existing applications. The CPU and GPU are treated as separate devices that have their own memory spaces. This configuration also allows simultaneous computation on both the CPU and GPU without contention for memory resources.

CUDA-capable GPUs have hundreds of cores that can collectively run thousands of computing threads. Each core has shared resources, including registers and memory. The on-chip shared memory allows parallel tasks running on these cores to share data without sending it over the system memory bus.

This guide will show you how to install and check the correct operation of C for CUDA and CUDA runtimes.

System Requirements

To use CUDA on your system, you will need the following installed:

- ❑ CUDA-enabled GPU
- ❑ Mac OS X v. 10.5.2 or later
- ❑ The gcc compiler and toolchain installed via Xcode
- ❑ CUDA software (available at no cost from <http://www.nvidia.com/cuda>)

About This Document

This document is intended for readers familiar with the Mac OS X environment and the compilation of C programs from the command line. You do not need previous experience with CUDA or experience with parallel computation.

Chapter 2. Installing CUDA

The installation of CUDA on a system running the appropriate version of Mac OS X consists of four simple steps:

- ❑ Verify the system has a CUDA-capable GPU.
- ❑ Download the CUDA software.
- ❑ Verify the correct version of the operating system and that a version of gcc is installed via Xcode.
- ❑ Install the CUDA software.

Test your installation by compiling and running one of the sample programs in the CUDA software to validate that the hardware and software are running correctly and communicating with each other.

Verify CUDA-Capable GPU

Many NVIDIA products today contain CUDA-enabled GPUs. These include:

- ❑ NVIDIA GeForce® 8, 9, and 200 series GPUs
- ❑ NVIDIA Tesla™ computing solutions
- ❑ Many of the NVIDIA Quadro® products

An up-to-date list of CUDA-enabled GPUs can be found on the NVIDIA CUDA Web site at http://www.nvidia.com/object/cuda_learn_products.html. The Release Notes for the CUDA Toolkit also contain a list of supported products.

To verify which video adapter your Mac OS X system uses, under the **Apple** menu select **About This Mac**, click the **More Info ...** button, and then select **Graphics/Displays** under the **Hardware** list.

Note: It is possible to develop CUDA software in the absence of a CUDA-enabled GPU. You can test the software in an emulation mode described later in this document. Naturally, performance on this platform is far less than on the CUDA-enabled processor, so the emulated hardware should not be used for release versions and performance tuning.

Download the CUDA Software

The CUDA software is available from the main CUDA download site at http://www.nvidia.com/object/cuda_get.html.

Choose the platform you are using, click **Search**, and download the CUDA Toolkit and the CUDA SDK (software development kit). The following sections describe the installation and configuration of version 2.2 of these packages.

Verify the Correct Version of Mac OS X

CUDA v. 2.2 requires an Intel-based Mac running Mac OS X v. 10.5.2 or later. To check which version you have, go to the Apple menu on the desktop and select **About This Mac**. You should see a dialog box similar to Figure 1.



To run CUDA, you need to check for version 10.5.2 or later.

Figure 1. About This Mac Dialog Box

Note: New versions of CUDA can require later versions of Mac OS X, so always verify that you are running the right release for the version of CUDA you are using.

Verify That gcc Is Installed

The gcc compiler and toolchain are installed via installation of Xcode. The Xcode development environment is found on the Xcode Developer Tools DVD that ships with new Mac systems and with Leopard, if you buy the operating-system upgrade. When installing Xcode, the package that contains gcc and the necessary tools is called Developer Tools Essentials.

CUDA Software

To run CUDA programs, you will need the following CUDA software:

- ❑ The CUDA Toolkit
- ❑ The CUDA SDK

The CUDA Toolkit contains the tools needed to compile and build a CUDA application in conjunction with the compilation driver. It includes tools, libraries, header files, and other resources.

The CUDA SDK includes sample projects that provide source code and other resources for constructing CUDA programs.

Before installing these packages, you should read the Release Notes bundled with each, as those notes provide details on installation and software functionality.

Installing the Software

Follow these few steps for a successful installation. For information not listed here, see the documentation in `/usr/local/cuda/doc`.

Download the CUDA software from http://www.nvidia.com/object/cuda_get.html and save the .pkg files.

Uninstall previous versions of the CUDA Toolkit and CUDA SDK if they have previously been installed. Do this by deleting the files from `/usr/local/cuda`, their default installation location. Adjust accordingly if you placed the files elsewhere. (If you wish to keep the files so you can compile for different versions of CUDA, then rename the existing directories and modify your makefile accordingly.)

Install the CUDA Toolkit by installing the downloaded .pkg file. The CUDA Toolkit installation defaults to `/usr/local/cuda`. Several environment variables need to be defined in the installation:

PATH needs to add `/usr/local/cuda/bin`. In addition, **DYLD_LIBRARY_PATH** needs to contain `/usr/local/cuda/lib`. The typical way to

place these values in your environment is with the following commands:

```
export PATH=/usr/local/cuda/bin:$PATH
export DYLD_LIBRARY_PATH=/usr/local/cuda/lib:
$DYLD_LIBRARY_PATH
```

To make these settings permanent, place them in `~/.bash_profile`.

The CUDA SDK is in the second .pkg file. It should be installed. The installation process will place the files in `/Developer/CUDA`.

Verify the Installation

Before proceeding, it's important to verify that the CUDA programs can find and communicate correctly with the CUDA-enabled hardware. To do this, you will need to compile and run some of the included sample programs.

Compiling the Examples

The version of the CUDA Toolkit can be checked by running `nvcc -v` in a terminal window. `nvcc` is the command to run the compiler driver that compiles CUDA programs. It calls the gcc compiler for C code and the NVIDIA PTX compiler for the CUDA code.

NVIDIA includes sample programs in source form in the CUDA SDK. You should compile them all by changing to `/Developer/CUDA` and typing `make`. The resulting binaries will be installed in the directory `/Developer/CUDA/bin/darwin/release`. If the compilation fails because it doesn't find gcc, the compiler should be installed via an Xcode installation.

Running the Binaries

The sample projects use libraries pointed to by `DYLD_LIBRARY_PATH`, as described earlier, so make sure it points to the right directory.

In addition, the executables need to find `libcutil.a` in `/Developer/CUDA/lib` and the corresponding include file in `/Developer/CUDA/common/inc`. They also need to access graphics libraries in `/Developer/CUDA/common/lib/darwin`. If you have installed the CUDA software as explained earlier, these locations are defaults and the programs should run without difficulty.

Once these required files are in place, go to `/Developer/CUDA/bin/darwin/release` and run `deviceQuery`. If CUDA is installed and configured correctly, the output for `deviceQuery` should look similar to Figure 2.

```

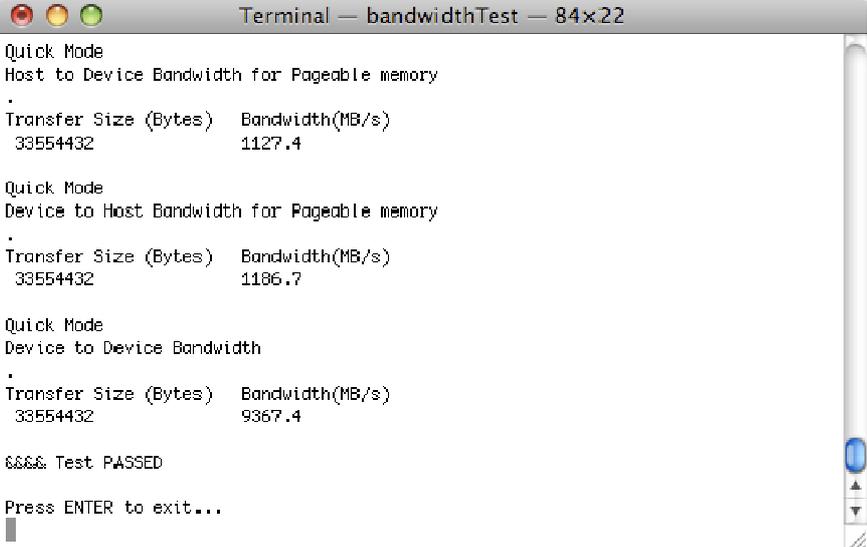
Terminal — deviceQuery — 80x21
There is 1 device supporting CUDA
Device 0: "GeForce 8600M GT"
Major revision number:      1
Minor revision number:      1
Total amount of global memory: 268238848 bytes
Total amount of constant memory: 65536 bytes
Total amount of shared memory per block: 16384 bytes
Total number of registers available per block: 8192
Warp size:                  32
Maximum number of threads per block: 512
Maximum sizes of each dimension of a block: 512 x 512 x 64
Maximum sizes of each dimension of a grid: 65535 x 65535 x 1
Maximum memory pitch:      262144 bytes
Texture alignment:         256 bytes
Clock rate:                 750000 kilohertz
Test PASSED
Press ENTER to exit...

```

Figure 2. Valid Results from Sample CUDA `deviceQuery` Program

Note that the parameters for your CUDA device will vary. The key lines are the first and second ones that confirm a device was found and what model it is. Also, the next-to-last line, as indicated, should show that the test passed.

Running the `bandwidthTest` program ensures that the system and the CUDA device are able to communicate correctly. Its output is shown in Figure 3.



```
Terminal — bandwidthTest — 84x22
Quick Mode
Host to Device Bandwidth for Pageable memory
.
Transfer Size (Bytes)  Bandwidth(MB/s)
33554432                1127.4

Quick Mode
Device to Host Bandwidth for Pageable memory
.
Transfer Size (Bytes)  Bandwidth(MB/s)
33554432                1186.7

Quick Mode
Device to Device Bandwidth
.
Transfer Size (Bytes)  Bandwidth(MB/s)
33554432                9367.4

$$$$ Test PASSED

Press ENTER to exit...
|
```

Figure 3. Valid Results from Sample CUDA `bandwidthTest` Program

Note that the measurements for your CUDA device description will vary from system to system. The important point is that you obtain measurements, and that the second-to-last line (highlighted) confirms that all necessary tests passed.

Should the tests not pass, make sure you have an NVIDIA GPU on your system that supports CUDA and make sure it is properly installed.

To see a graphical representation of what CUDA can do, run the sample `particles` executable.

To compile programs for emulation (on systems that have no CUDA-capable graphics hardware available), see the next section.



Chapter 3. Additional Considerations

Compiling for Hardware Emulation

The previous section explained how to compile and build the included files. To recap, simply go to the `/Developer/CUDA` directory, type `make`, and the resulting binaries will be installed in `/Developer/CUDA/bin/darwin/release`.

To see the individual steps in the build process, be sure to enable the verbose option on `make`.

On systems without a CUDA-capable GPU, it will be necessary to use an emulated GPU. CUDA enables the creation of binaries for an emulated hardware environment by using `make emu=1` from the command line in `/Developer/CUDA`. The resulting binaries will be placed in `/Developer/CUDA/bin/darwin/emurelease`.

What's Next?

Now that you have CUDA-capable hardware and the software installed, you can examine and enjoy the numerous included programs. To begin using CUDA to accelerate the performance of your own applications, consult the *CUDA Programming Guide*, located in `/usr/local/cuda/doc`.

For tech support on programming questions, consult and participate in the bulletin board and mailing list at <http://forums.nvidia.com/index.php?showforum=71>.

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 or otherwise under any patent or patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all information previously supplied. NVIDIA Corporation products are not authorized for use as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

Trademarks

NVIDIA, the NVIDIA logo, CUDA, GeForce, NVIDIA Quadro, and Tesla are trademarks or registered trademarks of NVIDIA Corporation. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2009 NVIDIA Corporation. All rights reserved.



NVIDIA.