

CUDA

The CUDA (Compute Unified Device Architecture) SDK can be used to write programs running on NVIDIA GPUs. All workstations in CIP2 have CUDA installed, the workstations have aliases "cuda1"- "cuda9".



Hardware

We have one **GeForce GTX 960** with CUDA capability 5.2, one **GeForce GTX 780** with CUDA capability 3.5, one **GeForce GT 640** with CUDA capability 3.0, one **GeForce GTX 560 Ti** with 2.1, one **GeForce GTX 470** with 2.0, and a few **GeForce 210** with 1.2.

```
-- cuda1.cip.ph.tum.de / cip2sandy1
Device 0: "GeForce GTX 960"
  CUDA Driver Version / Runtime Version      7.5 / 6.5
  CUDA Capability Major/Minor version number: 5.2
  ( 8) Multiprocessors, (128) CUDA Cores/MP: 1024 CUDA Cores
-- cuda2.cip.ph.tum.de / cip2sandy2
Device 0: "GeForce GTX 780"
  CUDA Driver Version / Runtime Version      7.5 / 6.5
  CUDA Capability Major/Minor version number: 3.5
  (12) Multiprocessors, (192) CUDA Cores/MP: 2304 CUDA Cores
-- cuda3.cip.ph.tum.de / ivy
Device 0: "GeForce GT 640"
  CUDA Driver Version / Runtime Version      7.5 / 6.5
  CUDA Capability Major/Minor version number: 3.0
  ( 2) Multiprocessors, (192) CUDA Cores/MP: 384 CUDA Cores
-- cuda4.cip.ph.tum.de / cip2iseven1
Device 0: "GeForce GTX 470"
  CUDA Driver Version / Runtime Version      7.5 / 6.5
  CUDA Capability Major/Minor version number: 2.0
  (14) Multiprocessors, ( 32) CUDA Cores/MP: 448 CUDA Cores
-- cuda5.cip.ph.tum.de / cip2iseven2
Device 0: "GeForce GT 630"
  CUDA Driver Version / Runtime Version      7.5 / 6.5
  CUDA Capability Major/Minor version number: 3.5
  ( 2) Multiprocessors, (192) CUDA Cores/MP: 384 CUDA Cores
-- cuda6.cip.ph.tum.de / cip2iseven3
Device 0: "GeForce GT 610"
  CUDA Driver Version / Runtime Version      6.5 / 6.5
  CUDA Capability Major/Minor version number: 2.1
  ( 1) Multiprocessors, ( 48) CUDA Cores/MP: 48 CUDA Cores
-- cuda7.cip.ph.tum.de / cip2iseven4
Device 0: "GeForce GTX 560 Ti"
  CUDA Driver Version / Runtime Version      6.5 / 6.5
  CUDA Capability Major/Minor version number: 2.1
  ( 8) Multiprocessors, ( 48) CUDA Cores/MP: 384 CUDA Cores
-- cuda8.cip.ph.tum.de / cip2iseven5
Device 0: "GeForce 210"
  CUDA Driver Version / Runtime Version      6.5 / 6.5
  CUDA Capability Major/Minor version number: 1.2
  ( 2) Multiprocessors, (  8) CUDA Cores/MP: 16 CUDA Cores
-- cuda9.cip.ph.tum.de / cip2smart
Device 0: "GeForce 210"
  CUDA Driver Version / Runtime Version      6.5 / 6.5
  CUDA Capability Major/Minor version number: 1.2
  ( 2) Multiprocessors, (  8) CUDA Cores/MP: 16 CUDA Cores
```

You can find out about the CUDA capabilities of the card in your workstation with the deviceQuery program from the samples directory:

```
gi32rog@cip2sandy2:~$ /mount/share/cuda-samples/1_Utillities/deviceQuery
/deviceQuery
/mount/share/cuda-samples/1_Utillities/deviceQuery/deviceQuery Starting...
```

CUDA Device Query (Runtime API) version (CUDA static linking)

Detected 1 CUDA Capable device(s)

```
Device 0: "GeForce GTX 560 Ti"
  CUDA Driver Version / Runtime Version      6.5 / 6.5
  CUDA Capability Major/Minor version number: 2.1
  Total amount of global memory:             2047 MBytes (2146631680
  bytes)
  ( 8 ) Multiprocessors, ( 48 ) CUDA Cores/MP: 384 CUDA Cores
  GPU Clock rate:                           1645 MHz (1.64 GHz)
  Memory Clock rate:                         2004 Mhz
  Memory Bus Width:                          256-bit
  L2 Cache Size:                             524288 bytes
  Maximum Texture Dimension Size (x,y,z)     1D=(65536), 2D=(65536,
  65535), 3D=(2048, 2048, 2048)
  Maximum Layered 1D Texture Size, (num) layers 1D=(16384), 2048 layers
  Maximum Layered 2D Texture Size, (num) layers 2D=(16384, 16384), 2048
  layers
  Total amount of constant memory:           65536 bytes
  Total amount of shared memory per block:   49152 bytes
  Total number of registers available per block: 32768
  Warp size:                                 32
  Maximum number of threads per multiprocessor: 1536
  Maximum number of threads per block:       1024
  Max dimension size of a thread block (x,y,z): (1024, 1024, 64)
  Max dimension size of a grid size (x,y,z): (65535, 65535, 65535)
  Maximum memory pitch:                     2147483647 bytes
  Texture alignment:                         512 bytes
  Concurrent copy and kernel execution:      Yes with 1 copy engine(s)
  Run time limit on kernels:                 Yes
  Integrated GPU sharing Host Memory:        No
  Support host page-locked memory mapping:   Yes
  Alignment requirement for Surfaces:        Yes
  Device has ECC support:                    Disabled
  Device supports Unified Addressing (UVA):  Yes
  Device PCI Bus ID / PCI location ID:       1 / 0
  Compute Mode:
    < Default (multiple host threads can use ::cudaSetDevice() with device
  simultaneously) >
```

```
deviceQuery, CUDA Driver = CUDART, CUDA Driver Version = 6.5, CUDA Runtime
Version = 6.5, NumDevs = 1, Device0 = GeForce GTX 560 Ti
Result = PASS
```

Local users can also try the **graphical examples** in /mount/share/cuda-samples, there are the binaries, the source code and quite a lot of documentation. Compiling your own programs has become a lot easier since CUDA5, in the past we needed to make lots of modifications to the Makefiles, but now there are just a few changes. After you have copied the source files to your home directory, change the line

```
INCLUDES      := -I$(CUDA_INC_PATH) -I. -I.. -I../common/inc
```

to

```
INCLUDES      := -I$(CUDA_INC_PATH) -I. -I.. -I$(CUDA_PATH)/samples/common/inc
```

(necessary because the includes of the samples are in the cuda tree, not in your HOME.) Also, you should comment out (or delete) the lines that make a copy of the finished binary,

```
#      mkdir -p ../../bin/$(OSLOWER)/$(TARGET)
#      cp $@ ../../bin/$(OSLOWER)/$(TARGET)
```

This way you should be able to compile everything by typing make. (If you are compiling some of the graphical examples, you also need to replace

```
-L../common/lib/$(OSLOWER)
```

with

```
-L/opt/cuda/samples/common/lib/$(OSLOWER)
```

etc.)

Documentation

Excellent Documentation is available at the NVIDIA homepage. You should check out [this exentsive site](#). But for starters it should be enough to read [Getting_Started](#). Last but not least, don't forget to checkout the [webinars](#).