Parallel Programming 2013/14- Lab exercise 1

This lab exercise introduces the basic steps of composing a CUDA program. The content has been adapted from that of Dan Ghica [1] and further material from the nVidia CUDA training website [2].

1. Find a machine in LG04 with a "CUDA" sticker and log in. Only these machines have been equipped with the necessary hardware (an nVidia GPU) and software (CUDA Toolkit)

2. Open a terminal and type in the command:

   ```
   module load CUDA
   ```

   This step configures the build environment.

3. Navigate to the directory:

   ```
   /bham/pd/packages/SL6/x86_64/cuda-4.0.17-sdk/C/bin/linux/release
   ```

   Which contains the binaries of several sample CUDA programs. Execute the sample program `deviceQuery` and observe the output; you should see information regarding your GPU device. You can experiment with the rest of the sample programs as well.

4. We're going to develop a simple GPU-accelerated array reversal program. Download the program skeleton:

   ```
   http://www.cs.bham.ac.uk/~hxt/2013/parallel-programming/lab1.cu
   ```

   Your task is to fill out the missing parts of this skeleton, you may refer CUDA documentation (on the web) to figure out the correct API calls. The crux of this exercise is to code up the kernel (just 3 lines).

5. Congratulations! You have developed your first GPU-accelerated (or not?) program.

References

[1] Dan Ghica. Programming Massively Parallel Architectures. Available at `http://www.cs.bham.ac.uk/~drg/cuda/`, 2013.

[2] nVidia. CUDA Zone - Education and Training. Available at `https://developer.nvidia.com/cuda-training`, 2013.