

ECE/CSC 506: Architecture of Parallel Computers

Program 1: GPU Programming

Due: Jan 21st, 2025 11:59pm

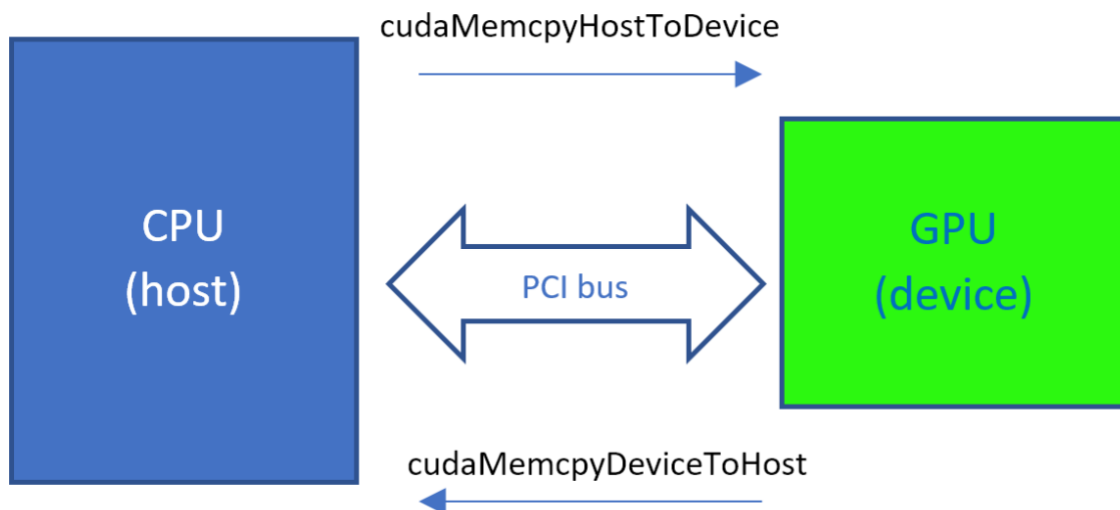
Introduction

GPUs outperform CPUs on applications like image processing, array processing, and plotting graphs. This advantage derives from GPUs' ability to process data in parallel rather than serially.

This assignment asks you to code vector calculations for a graph plot using CUDA code. This assignment has three objectives:

- To introduce you to CUDA.
- To introduce you to simulation environments
- To demonstrate how small modifications in the code can reduce the number of instructions processed by the processor, the number of cache accesses, and overall energy usage.

Background



GPUs generally require a host device to activate them. The CPU (host) will run most of the serial code, but for the parallelizable code, it invokes the GPU kernel.

Note these attributes of CUDA code:

- The kernel invocation is identified by `<<< blocksPerGrid, threadsPerBlock >>>`
- The GPU kernel function header is
`__global__ <ret type> <function name> (args....)`
- `cudaMemcpy` is used to copy the data in between device and host

Watch the following videos from the Heterogeneous Parallel Processing course to learn more about programming the GPU. You may take the embedded quizzes, but you are not required to. They will not be counted toward your grade.

- a. CUDA memory allocation and data movement API functions [19:37] [Watch](#)
- b. CUDA kernel-based SPMD parallel programming [19:09] [Watch](#)
- c. Kernel-based parallel programming, multidimensional kernel configuration [16:20] [Watch](#)
- d. Kernel-based parallel programming, basic matrix-matrix multiplication [17:17] [Watch](#)

Environment setup

To see the impact of the four parameters mentioned above on your code, we use `gpgpusim` to simulate the working of GPU. This simulator simulates the working of the **NVIDIA GeForce GTX 480**.

To set up the environment on your PC, follow these steps:

1. Download and Install Oracle VirtualBox from <https://www.virtualbox.org/wiki/Downloads>
2. Download the Ubuntu VM image with gpgpu simulator (`ece506_gpgpu-sim`) from here: https://drive.google.com/file/d/1N-zg9QpUh6_UpMbXutsl-57HZX0hp4Pa
3. Import the `ece506_gpgpu-sim.ova` using the Oracle VirtualBox environment. The `lspassword` is `cuda506`

NOTE: Do not update any installed packages in the Linux environment. The simulator requires a certain version of packages (e.g., `gcc`) to function smoothly. Kindly leave the packages untouched.

Then click on the `ece506_gpgpu-sim` button at the left and then click the “Start” arrow and perform a “normal start”.

4. Launch a terminal and run the following commands:

- `cd ~/gpgpu-sim_distribution`
- `chmod +x setup_environment`
- `./setup_environment`
- `make clean`
- `make`
- `cd ~/vectorAdd`
- `make clean`
- `make`
- `chmod +x test_sim.sh`
- `./test_sim.sh`

Wait for the script to run in the background. This may take up to 3 minutes. If the output says “Test PASSED”, the environment works for CUDA programs, and you’re good to go! If not, repeat the steps.

5. Read the `vectorAdd.cu` code given in the `vectorAdd` directory. Understand the memory allocation, copying of data and memory freeing involved in the code. Also, understand how the CUDA kernel function works.

Experiment

The Ubuntu VM image contains a `submission` folder in the home directory, which has `Plot_1` and `Plot_2` as subfolders. Examine the `VectorAdd` folder in the home directory. It will help you understand the basics of CUDA programming.

Using the `Makefile` inside the `vectorAdd` folder as a reference, create your own `Makefile` for the folders `Plot1` and `Plot2`.

Now, program the GPU kernel to implement the following equations:

Plot 1:

$$C(x) = 3A(x)^4 + 2A(x)^3D(x)^2 - 8A(x)^2B(x) + 7A(x)^4B(x)^3 + 3A(x)B(x)^3 + 5A(x)^5 + 4B(x)^6 + 8$$

Plot 2:

$$C(x) = A(x)^5/D(x)^2 + 3A(x)^5 + 6A(x)^3B(x)^3/(C(x)^2D(x)^2) + 2A(x)^3B(x) + 9A(x)B(x)^5 - 7/D(x)^4$$

For the first plot, you only need to implement an optimized kernel. For the second plot, in addition to an optimized kernel, you also need to add appropriate code to accommodate vector D . To test your optimized kernel, modify the value of OPT, rebuild and run.

When running the simulator, [redirect its standard output](#) to a txt file and include it in the final submission. You need to collect data for both the base kernel and your optimized kernel.

From the simulator output, note the following for the analysis:

1. **Total instructions processed** (gpu_sim_insn)
2. **Total number of cache accesses** (L1I_total_cache_accesses+ L1D_total_cache_accesses) and
3. **Average energy usage** from the power report generated (note that power report log file will be created when you run the code. you'll need to check the power report every time you run the code to get the reading).

NOTE: You are free to consider other details of the output as well in your analysis.

Plot 1	Base code	Energy Optimized code	% improvement
Total instructions processed			
Total number of cache accesses			
Average Energy usage			

Two tables need to be created, one for **Plot 1** and one for **Plot 2**.

Analysis

The analysis should be done per plot. You need to compare the base code results with the energy-efficient code results and explain why you observe the changes.

1. Why is the total number of instructions less than in the base code?
2. Why are there fewer cache accesses?
3. What is the parameter that is compromised for saving energy? Why?
4. Which code would you prefer to use, and under which circumstances?
5. How did you find your first experience with a simulator?
6. Where do you think simulators will be helpful?

Submission format

Your submission should contain two folders and a report file in a zip archive. Inside each plot folder, rename the redirected output file and the generated `gpgpu_power_report_XXXX` file to reflect which kernel was used, exactly as shown below. The whole folder structure should look like the following:

- Plot_1
 - Plot_1_base.txt
 - Plot_1_base_power.log
 - Plot_1_opt.txt
 - Plot_1_opt_power.log
 - Plot_1.cu
 - config_fermi_islip.icnt
 - gpgpusim.config
 - gpuwattch_gtx480.xml
 - Makefile
- Plot_2
 - Plot_2_base.txt
 - Plot_2_base_power.log
 - Plot_2_opt.txt
 - Plot_2_opt_power.log
 - config_fermi_islip.icnt
 - gpgpusim.config
 - gpuwattch_gtx480.xml
 - Makefile
- report.pdf

Each folder should contain a `.cu` source file and a `Makefile`. Your program should be able to compile when “make” is issued in that folder.