

Introductory Lab

This lab assignment is intended to help you get your feet wet with the basics of Python, numpy and GPU programming.

The problems get harder as you move along, and there is more than enough work here to keep participants at any level busy for the allotted time. So don't worry if you don't make it all the way to the end, and feel free to jump to the things that interest you most. Do come back and retry the problems that you didn't have time for.

Getting Started

Follow these steps to get started:

1. Along with this lab assignment, you should have been handed a snippet of paper with a user name and a password. Sit down at the workstation you've been assigned and turn it (and its monitor) on.
2. Wait for a menu to appear. The menu in question says something like "FOG" at the bottom of the screen.
3. Pick the bottom entry (should say "Linux") in the menu using the cursor keys and hit Enter.
4. Wait for a login prompt, then enter your user name and password.
5. Wait for the desktop to appear, then click 'Accessories', then 'Terminal'.
6. In the terminal window, type 'gedit &' (without the quotes) to start an editor.
7. Using the editor, create a file 'hello-world.py' that says 'print "Hello World!"'.
8. In the terminal, type 'python hello-world.py' (again, no quotes) and hit Enter. You've just run your first Python script—congratulations! Time to get started on the assignment.

Problem 1: Python, Numpy and Project Euler

If we list all the natural numbers below 10 that are multiples of 3 or 5, we get 3, 5, 6 and 9. The sum of these multiples is 23.

Find the sum of all the multiples of 3 or 5 below $n = 1000$.

- a) Write an iterative Python program to do this. (*Hints:* `range(5)` generates the list `[0, 1, 2, 3, 4]` '`%`' is modulo.)
- b) Write a numpy-vectorized program to do this.
- c) What happens if you increase n above? What is the relative speed of each solution? What is the scaling with n ? Look up Python's `time`¹ library to do your timing. (*Hint:* Use `xrange`, not `range` for large values of n .)

¹<http://docs.python.org/library/time.html>

- d) (Slightly advanced:) The solution to this problem can be expressed in a single line (of 50-odd characters) in Python. How? (*Hints:* Search the Python documentation for the words ‘generator expression’. Python has a function ‘sum’.)

Problem 2: Simple GPU Programming with PyOpenCL

Write a PyOpenCL code that solves Problem 1. You may start from the simple example in the [PyOpenCL documentation](#)².

- a) Think of a way to parallelize this. In your first attempt, you may use some CPU and some GPU computation.
- b) Gradually increase n . You likely want to make sure that you are using 64-bit integers (always ‘long int’ in OpenCL) beyond a certain value of n to avoid integer overflow.

What performance results do you expect? Benchmark your solution using synchronization (`queue.finish()`) and Python’s `time`³ library. Do your measurements match what you expect? Also try measuring using PyOpenCL Events. Are the results any different?

- c) Investigate the use of `pyopencl.array` and `pyopencl.elementwise`.
- d) Benchmark the array-based solution. What do you observe? In each of your codes, what machine feature determines your performance?
- e) Investigate the use of `pyopencl.reduction` to move the computation entirely onto the GPU.

Hint: Use your solution to Problem 1 to check your answers.

Problem 3: Better know a GPU

Let us try and measure a few properties of the GPU you are using.

- a) Try to measure the transfer bandwidth between host and device. Is it the same in both the reading and writing direction?
- b) Try to measure the device memory bandwidth. How does this number depend on alignment? How on the stride from one work item to the next? Does the number you measured match the published peak numbers? (see [Wikipedia](#)⁴)
- c) Can you measure the maximum flop rate? What obstacles do you encounter?

Hint: See [this article](#)⁵ for inspiration.

²<http://documen.tician.de/pyopencl>

³<http://docs.python.org/library/time.html>

⁴<http://is.gd/fldxGv>

⁵<http://dx.doi.org/10.1109/ISPASS.2010.5452013>