OpenCL Practical 2 – kernels

This practical will give you a chance to practice modifying simple kernels. You will also get to use OpenCL's vector data types. The main objectives are to learn about:

- how to modify kernels and make required changes in the host code
- how to use OpenCL's vector data types
- how to time your OpenCL code

1. Log in to the head node, e.g. ssh username@gpu.hector.ac.uk
2. Change to the prac2 directory: “cd ~/opencl_course/prac2”
3. Look at the Makefile to see how it works then type “make”
4. Submit jobs to the GPUs via the queue manager using ‘qsub’, e.g. “qsub jobSub2”
5. Keep track of where your jobs are in the queue with “qstat”
6. Have a look at the output that’s produced in jobSub2.oxxx

7. Try making changes to the OpenCL kernel at the top of vadd.c, such changing the size of the arrays. Don’t forget to make any necessary changes to the host code
8. Add a new OpenCL kernel, “vmul()” that takes the same parameters as vadd() but performs a vector multiply. Change the host code to run the vadd first kernel then your new vmul kernel
9. If you have spare time, look at the NVIDIA SDK OpenCL examples in ~crsadmin/NVIDIA_GPU_Computing_SDK/OpenCL/