Based on original ideas by Mike Giles

OpenCL Practical 1 – getting started

This practical gives an introduction to OpenCL programming using a very simple example. The main objectives are to learn about:

- the way in which an application consists of a host code to be executed on the CPU, plus kernel code to be executed on the GPU
- how to copy data between the OpenCL device (GPU) and the host (CPU)
- how to check any error codes returned by OpenCL

1. Log in to the head node, e.g. ssh username@gpu.hector.ac.uk
2. Execute the command “cp -r ~crsadmin/opencl_course ~”
3. Change to the prac1 directory: “cd ~/opencl_course/prac1”
4. Look at the Makefile to see how it works then type “make”
5. Submit jobs to the GPUs via the queue manager using ‘qsub’, e.g. “qsub jobSub1”
6. Keep track of where your jobs are in the queue with “qstat”
7. Have a look at the output that’s produced in jobSub1.oxxx
8. Read through the vadd.c source file and err_code.c which decodes OpenCL error codes
9. Try introducing errors into vadd.c, such as trying to allocate too much memory or enqueueing 0 kernels, and improve the code to catch and interpret the error codes using the err_code() routine
10. If you have spare time, look at the NVIDIA SDK OpenCL examples in ~crsadmin/NVIDIA_GPU_Computing_SDK/OpenCL/. A good one to try is “oclDeviceQuery”