Accelerating High-Throughput Computing through OpenCL

Andrei Dafinoiu Joshua Higgins Violeta Holmes

HPC Research Group,

University of Huddersfield, UK

High-Throughput Computing (HTC) focuses on computational tasks that require the use of resources over an extended period of time. One of the most attractive features of these systems is cycle-stealing, the ability to use workstations as compute nodes for as long as they remain idle.

OpenCL is a heterogeneous programming framework for the development of applications that span across multiple architectures, including CPUs, GPUs, DSPs, and other accelerators. While, from a computing standpoint, GPUs are being used as accelerators in the top supercomputers of the world, the general purpose GPUs (GPGPUs) available in typical workstations will not have a significant impact on the fast processing required in High Performance Computing. In HTC the fastest performance for individual units is not necessarily required, thus making it a better candidate for GPGPU based acceleration.

Previous work has shown that middleware vendors support the idea of GPU computing in HTC through built-in detection methods, as it is the case with HTCondor. Typically this relies on CUDA, as well established software already support CUDA. However, CUDA is limited to certain NVidia GPUs and also requires special packages to be installed on the system in order to work. In terms of general purpose PCs on campus, only a small number of workstations may be capable of supporting CUDA. OpenCL based detection is available in HTCondor, however, it is not yet mature technology.

Since OpenCL has wider device support, workstation specifications are not as relevant at the programming stage of development, making OpenCL a viable solution for environments that contain more than one architecture or operating system. Its ability to operate using base drivers promises a much faster integration with existing systems.

This paper presents an implementation of an OpenCL system within the 7000 core HTCondor pool at the University of Huddersfield. The goal of this work is to expand the capabilities of the system beyond CPU computing in a manner that exploits the GPGPUs but does not greatly affect the existing implementation. To evaluate the OpenCL based HTC system, GPGPUs were deployed in accelerating FFT computations.

The paper also discusses best practise implementations and considerations, limitations, and further development efforts for wider adoption of OpenCL.

References:

1. Stone, J. E., Gohara, D., & Shi, G. (2010). OpenCL: A parallel programming standard for heterogeneous computing systems. Computing in Science & Engineering, 12(3), 66-73. doi:10.1109/MCSE.2010.69

2. Guba, S., Ory, M., & Szeberényi, I. (2014). Harnessing wasted computing power for scientific computing.DOI: 10.1007/978-3-662-43880-0_56