

Accelerating High-Throughput Computing through OpenCL

Andrei Dafinoiu, Joshua Higgins, Violeta Holmes
High-Performance Computing Research Group
University of Huddersfield
Huddersfield, United Kingdom



Overview

Introduction

Resources

Motivation

Experiment Design

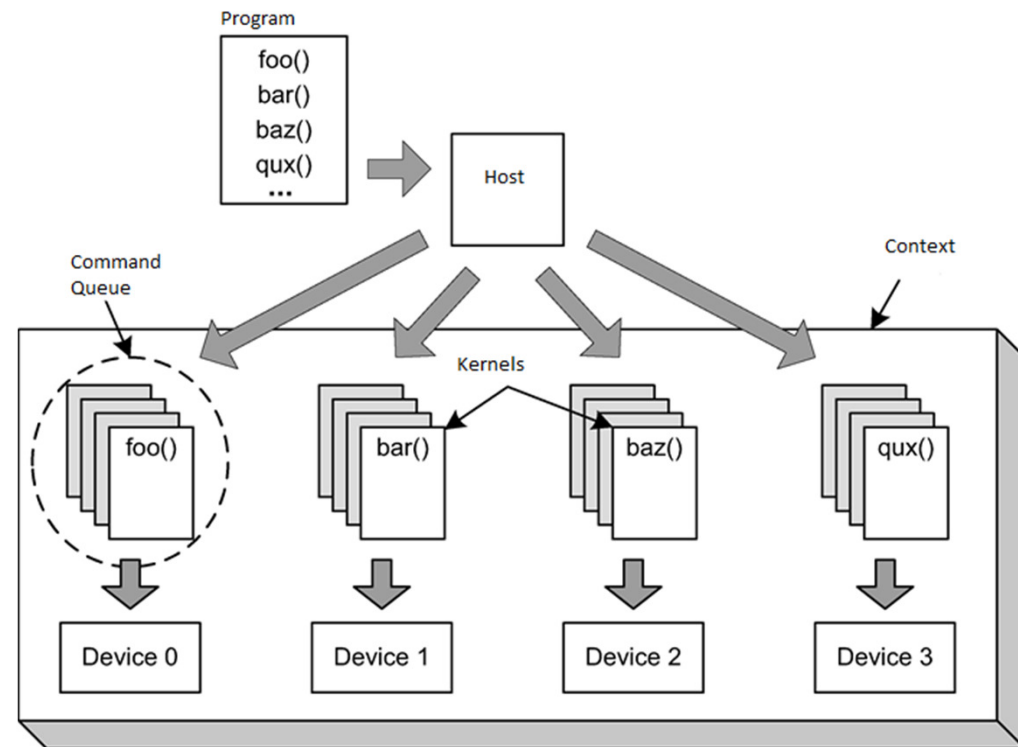
Results and Performance

Conclusions



OpenCL

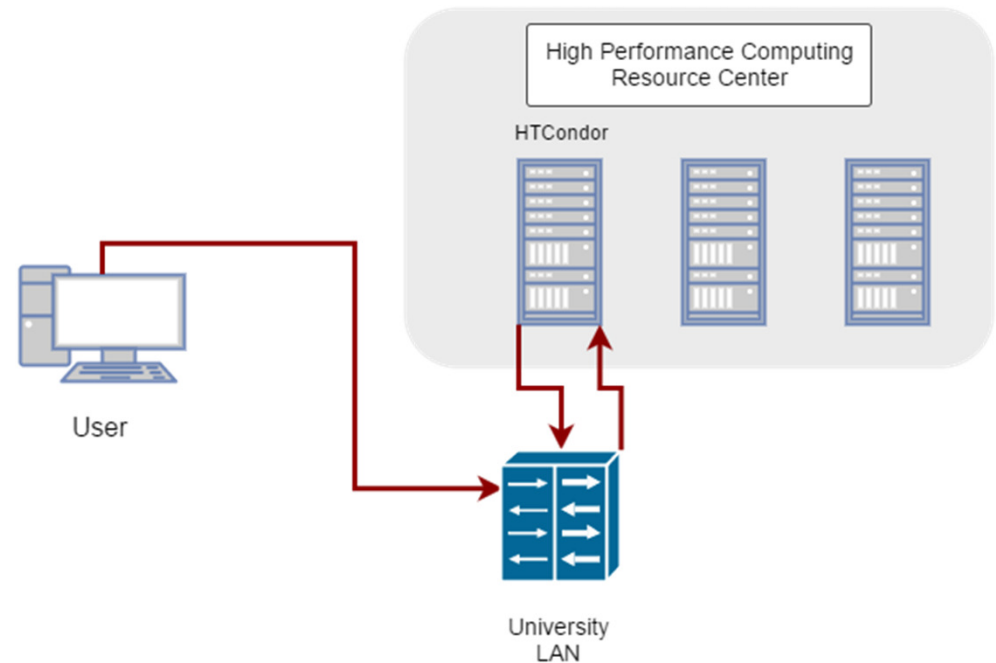
- OpenCL -> A programming framework for heterogeneous compute platforms
- Supports CPU, GPU, DSP, and other accelerators
- Programming API based on C/C++




Introduction



- High-Throughput Computing
- HTCCondor
- QGGCondor




Aim of project

- To expand the capabilities of the QGGCondor.
 - To evaluate the efficiency and flexibility of OpenCL for use within a heterogeneous HTC environment.
 - To increase the visibility of GPGPU computing
- 


HTCondor Setup

ClassAdds, what are they ?


Old ClassAdd



HTCondor Proposed



New ClassAdd




Environment Setup

Reasoning:

- Unreliable environment for benchmarking purposes.
- Desire to execute benchmarks over the live system.

Method:

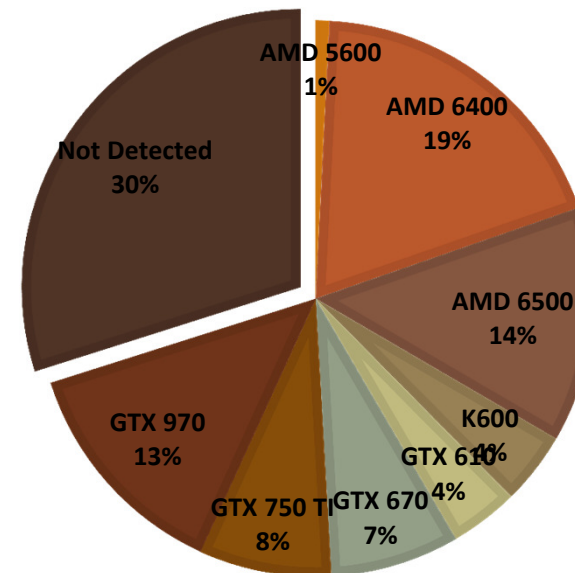
- Extracted hostnames of 1000 condor machines.
 - Script based generation of jobs.
- 

Resource Discovery

TABLE I
GPU LANDSCAPE

GPU	Nr
AMD 5600	8
NVIDIA Quadro K600	42
NVIDIA GTX 610	40
NVIDIA GTX 670	75
NVIDIA GTX 750 Ti	77
NVIDIA GTX 970	133
AMD 6500	137
AMD 6400	189
Not detected	299
Total	1000

GPU DISTRIBUTION



Experiment Design

Fast-Fourier Transform

1000 iterations -> to ensure precision

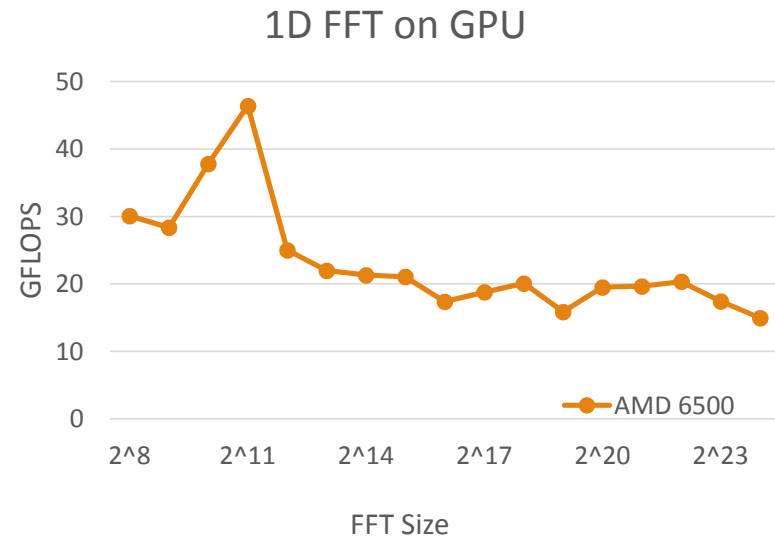

17 FFT sizes -> from 2^8 to 2^{24}

701 machines -> that reported GPUs


Resulting in:

- Aprox. 12 million FFT calculations
- Aprox. 28 thousand CPU hours

Results




QGGCondor GPU FFT Performance



Comparison with GPU cluster

- On average, a single compute GPU is negligibly faster.
- Newer-gen GPGPUs outperform older compute counterparts.



P.S: NVIDIA C2050 was released in 2011

Conclusion

- HTCondor GPU integration was successful with single entity ClassAdd definition.
- OpenCL based implementations over HTCondor are straightforward (without specific optimizations) however machines need dedicated graphics drivers installed.
- QGGCondor GPU performance varies greatly however the average is marginally slower than that of a dedicated compute GPU.

Thank you for your attention

Questions ?

