



Contribution ID: 110

Type: not specified

Supporting grid-enabled GPU workloads using rCUDA and StratusLab

Wednesday, 28 March 2012 14:20 (20 minutes)

Impact

The first outcome shall be to provide a standardised addition to the current GLUE-schema to enable the resource centres to advertise GPGPU resources.

The second outcome shall be to provide access to powerful grid-enabled services which can supporting massively parallel codes running on multiple commodity GPGPUs.

The expected positive effects for the user community are greater levels of resources, and shorter runtimes for certain application workloads.

Overview (For the conference guide)

The increasing capabilities of general purpose graphics processing units (GPGPUs) over the past few years has resulted in a huge increase in their exploitation by all the major scientific disciplines where massively parallel processing capabilities are desired. However, there are two major problems in supporting grid access to such resources: Firstly, there is currently no standardised way for resource centres to advertise/publish availability of these resources. Secondly, there are deficiencies in current batch scheduling systems that ensure exclusive access to those resources.

We present the results of an initial investigation into grid-enabling access to many general purpose graphics processing units (GPGPUs) distributed over a local cluster. We exploit two distinct virtualising technologies - rCUDA and StratusLab. The hybrid approach is used to achieve greater levels of parallelism, and to provide the necessary GPGPU resource isolation.

This is currently a work in progress.

Description of the Work

The first aspect of this work looks at using or extending the current grid GLUE-schema in order to advertise GPGPU compute resources. It will propose a baseline set of desired attributes.

The second aspect focuses on using rCUDA and StratusLab to provision the multi-GPGPU workernode environments. StratusLab provides the infrastructure to create/deploy the virtual machines with the desired rCUDA environment.

rCUDA is a GPGPU virtualisation layer, developed by research groups at Universidad Politécnic de Valencia and Universidad Jaime I. Castellón in Spain, that provides a framework which enables access to multiple remote GPGPUs distributed over a network of machines. It uses a socket based client/server architecture, and it provides a large subset of the CUDA 4.0 API capabilities. From the users or applications perspective, a set of remote GPGPU devices in a cluster can be accessed as if it they were set of local devices. This should allow

applications to achieve greater levels of parallelism and to allow application developers to avoid requiring the use of additional layers of parallelisation frameworks, such as MPI.

The StratusLab project is developing an open-source cloud distribution that allows grid and non-grid resource centres to offer and to exploit an “Infrastructure as a Service” (IaaS) cloud.

This hybrid approach shall also be used to investigate the feasibility of avoiding many deficiencies in many batch job scheduling systems. For instance, MAUI - a very popular scheduler - does not have any support for GPGPU resources. Moreover, on nodes where GPGPU resources are available, MAUI cannot ensure exclusive access to the GPGPU resource, therefore distinct user processes may concurrently try to access the same device. The rCUDA/StratusLab approach can avoid this contention by providing a single job slot on a virtual machine with dedicated virtual access to many GPGPU devices.

Conclusions

Recent advances in hardware and software virtualisation capabilities have made it possible to customise hardware and software environments for a huge variety of applications. Grid infrastructures have capitalized on many of these advances, for example, through the use of grid-enabled virtual machines which provide well-known and trusted user services and environments.

At the same time, there are still deficiencies in providing grid-enabled access to generic computing capabilities, such as GPGPUs. This work looks at exploring existing virtualisation services to provision safe and coherent grid-enabled access to commodity GPGPU resources.

Primary author: WALSH, John (Trinity College Dublin)

Co-authors: Prof. COGHLAN, Brian (Trinity College Dublin); Dr O’CALLAGHAN, David (Trinity College Dublin)

Presenter: WALSH, John (Trinity College Dublin)

Session Classification: Clouds: Infrastructure