CS554 Project Ideas

GeMTC:Mon – Monitoring GeMTC on NVIDIA Hardware

Overview

GeMTC is a CUDA based GPGPU framework which allows Many-Task Computing (MTC) workloads to run efficiently on NVIDIA GPUs. Users call a C API, which allows their task to be scheduled and run on a GPU. Users push a TaskDescription to the GPU, which is a structure containing the application they would like to run, and any parameters for that application. Applications that run inside of GeMTC are referred to as Micro-Kernels.

The successful completion of this project will achieve the following steps:

- Log important GeMTC operations that occur on both the CPU and GPU.
- Calculate important dynamic metrics:
    - How long did an application (Micro-Kernel) run?
    - How long was a worker busy/idle?
- Visualize metrics and automatically push them to a web browser for quick viewing.

Relevant Systems and Reading Material

- GeMTC – http://datasys.cs.iit.edu/projects/GeMTC
- Xeon Phi - http://software.intel.com/en-us/mic-developer
- Swift – http://swift-lang.org
- Scott J. Krieder, Justin M. Wozniak, Timothy Armstrong, Michael Wilde, Daniel S. Katz, Benjamin Grimmer, Ian T. Foster, Ioan Raicu. “Design and Evaluation of the GeMTC Framework for GPU-enabled Many-Task Computing”, ACM HPDC 2014; http://datasys.cs.iit.edu/publications/2014_HPDC14_GeMTC.pdf

Preferred/Required Skills

- Required: C
- Preferred: CUDA, Threaded programming

Evaluation

You will be responsible for running several test applications that utilize this software stack as well as some high level tests similar to those found in the HPDC 2014 paper, and monitor the resource usage. Experiments can be done on a single node, on a single GPU on the Jarvis cluster.

Project Mentor

Ioan Raicu, iraicu@cs.iit.edu