DEVELOPMENT AND OPTIMIZATION OF THE NVIDIA CUDA APPLICATION FOR MILKYWAY@HOME

By

Anthony J. Waters

An Abstract of a Thesis Submitted to the Graduate

Faculty of Rensselaer Polytechnic Institute

in Partial Fulfillment of the

Requirements for the Degree of

MASTER OF SCIENCE

Major Subject: COMPUTER SCIENCE

The original of the complete thesis is on file
in the Rensselaer Polytechnic Institute Library

Approved:

Boleslaw Szymanski, Thesis Adviser

Rensselaer Polytechnic Institute
Troy, New York

April 2010
(For Graduation May 2010)

ABSTRACT

Milkyway@Home takes advantage of maximum likelihood fitting in order to generate a model of fitting the Sagittarius Drawf Tidal Streams to the stars within the Milkyway Galaxy. The maximum likelihood calculation is simplified into a convolution integral and solved using the Gaussian quadrature method. Once a likelihood value is calculated it is used, along with the calculated parameters, in order to refine the model.

A disadvantage of this process is that it takes an exorbitant amount of CPU power and time. On a single core CPU, it was estimated that it would take 312 days to complete the analysis of a single stream. In order to get around the issue first grid computing was tried, through SALSA and MPI. Then, volunteer computing, through the BOINC framework, was experimented with. While analyzing a single stream through volunteer computing decreased the total amount of time spent, it was still time consuming in order to generate a single likelihood value.

To tackle this issue and decrease the time spent generating a single likelihood value GPU applications were explored. Initially there were concerns surrounding the fact that GPUs only supported single precision arithmetic, as opposed to double precision which is required. However, recent advances in the hardware have exposed double precision arithmetic in the next generation of GPUs. Therefore, this thesis explores the performance of a double precision Nvidia GPU application. In addition, a set of performance bottlenecks in the application are identified and a set of optimizations were fully explored and benchmarked. Lastly, the different applications that are used within Milkyway@Home are compared against each other in order to gain some insight on the computational power within consumer GPUs.