

Introduction to CUDA and Parallel Programming

Sunday, 2 December 2018 13:30 (1h 30m)

Description:

The goal of this workshop is to introduce students to CUDA and provide them with an understanding of parallel programming. CUDA is more than a few new keywords. One must understand SIMD and the pitfalls of serialization.

Students will leave with basic CUDA skills and some OpenACC knowledge plus useful machine learning and big data tools as well. My “from Hello World to exascale machine learning in one slide” will also be covered as data parallel training fits on GPUs nicely.

Target Audience:

Anyone with C/C++ programming skills in the Unix environment who wishes to learn about parallel programming and CUDA. The material will be 60% beginner, 30% intermediate, and 10% advanced.

Prerequisites:

C/C++ along with an ability to edit and compile programs in a Unix environment

Special requirements:

Users will have access to a CHPC system with GPUs.

Attendees should bring their own laptops. The ability to view pdf or PowerPoint files is required.

Outline of full syllabus:

08:00 Registration

09:00 Introduction and morning talk (30 minutes)

Login details and extracting the workshop material (15 minutes)

Section 01: Parallel intro and a first CUDA program)

Section 02: Profiling on a GPU

10:30 Morning Refreshment Break

11:00 Section 03: More CUDA and the Thrust Interface

Section 04: “From Hello World to TF/s machine learning”

12:30 Lunch

13:30 Afternoon talk

Section 05: Controlling parallel resources

Section 06: C++ objects and transparent host/GPU data movement

15:00 Afternoon Refreshment Break

15:30 Section 07: Task level parallelism on a GPU

Section 08: Managing big data, CUDA as a scripting language via dynamic load/link

17:00 End of Day

Additional Comments:

Students can work at their own pace.

Introductory students will learn the basics of CUDA and the profiler as well as how to think in parallel and understand the impact of parallel hardware on performance.

Intermediate/advanced students will hone their thinking about parallel programming and the limitations and advantages of GPU hardware. Extra credit exercises will challenge them.

All students will learn how to use machine learning and the ability to explore this hot field and leave with a tool that allows them to train and predict using their own data sets and neural network architectures. Further, they learn how to work with and collaborate using big data.

Presenter Biography

Rob Farber was a pioneer in the field of neural networks while on staff as a scientist in the Theoretical Division at Los Alamos National Laboratory.

He is active in the field and works with companies and national laboratories as a consultant plus teaches about HPC and AI technology worldwide.

Rob can be reached at info@techenablement.com

<http://www.techenablement.com/rob-farber/>

Primary author: FARBER, Rob (TechEnablement)

Presenter: FARBER, Rob (TechEnablement)

Session Classification: Introduction to CUDA