

CUDA Teaching Center: High Performance Computational Science (CTC/HPCS) Workshop

CUDA Teaching Center WMU, HPCS Laboratory CEAS WMU¹

Department of Computer Science, Western Michigan University

Workshop: Room C-208 CEAS Parkview Campus

1. Morning session 9:00 - 12:00: HPCS Laboratory

1.1. Overview

Cluster architecture:	network fabric	multicore nodes	Accelerators
Interconnect:	Infiniband, GigE		GPUs, (Xeon Phi)
Parallel Environments:	MPI/Open MPI	OpenMP	CUDA, OpenACC, OpenCL
High level language: Python			
Simple examples:	π program(s)	series summation	dot product (OpenACC)

Local projects

1.2. Using the cluster

ssh, putty, login, compiling and excuting on Linux, basic Linux commands

1.3. Parallel Environments

MPI and OpenMP examples

1.4. Job scheduler

PBS/torque: sample scripts, incl. MPI, interactive session

2. Afternoon session 13:00 - 16:00: CUDA

2.1. GPU/CUDA architecture

2.2. Simple CUDA examples

Dot product, dot product with arbitrarily long vectors, Monte Carlo

2.3. Advanced features

Dynamic parallelism (K20)

2.4. More complex applications, packages

Large, sparse linear systems, Finite Elements, Molecular Dynamics, Quasi-MC

Packages: CURAND, CUBLAS, CUSPARSE, CUSP, Abaqus, Calculix, ANSYS, OpenFOAM, Fluent

2.5. Q & A

¹ **Acknowledgments.** We acknowledge the support from NVIDIA under the CUDA Teaching Center Award at Western Michigan University, and from the National Science Foundation under Award Number 1126438.