

Mogon2: deeplearning GPU partition

-Overview

Srijit Paul

Johannes Gutenberg University, Mainz

Overview

- 1 Introduction
- 2 Compilation
- 3 Running
- 4 Lattice QCD specific codebases

section I

- 1 Introduction
 - Hardware & Softwares available
- 2 Compilation
- 3 Running
- 4 Lattice QCD specific codebases

Hardware & Softwares available

- Mainz: Mogon2, 2×8 V100 GPUs.
 - (1) Memory throughput: 900 GB/sec.
 - (2) 16, 32 GB HBM2 memory.
 - (3) 5120 CUDA cores.
- CUDA: software development.
 - (1) GNU compilers
`module load toolchain/gompic/2020b`
 - (2) Intel compilers
`module load toolchain/iccifortcuda/2020a`
- ML: Python code development.
 - (1) Pytorch
`module load devel/PyTorch/1.9.0-fosscuda-2020b`
 - (2) Tensorflow
`module load lib/TensorFlow/2.4.1-fosscuda-2020b`

section 2

- 1 Introduction
 - Hardware & Softwares available
- 2 Compilation**
- 3 Running
- 4 Lattice QCD specific codebases

Compilation

CUDA development.

- Compiler: `nvcc -arch=sm_70`
- Profiler: `nvprof --metrics achieved_occupancy,dram_read_throughput,dram_utilization,dram_write_throughput,l2_read_throughput,l2_write_throughput,issued_ipc,..`
- In-code profilers: CUPTI hooks.

Sample code here: <https://github.com/srijitpaul/SAXPY>¹

Python Machine Learning code development.

- Load required Pytorch/tensorflow modules and create virtual environment for specific packages.

¹More here <https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html>

section 3

- 1 Introduction
 - Hardware & Softwares available
- 2 Compilation
- 3 Running**
- 4 Lattice QCD specific codebases

Prerequisites

```
[spaul@login21]$ groups  
kph m2_dgxqcd m2_gitter m2_him_th him-thfl
```

Typical script

```
#!/bin/bash  
  
#SBATCH -J jobname  
#SBATCH -A m2_dgxqcd  
#SBATCH --partition="deeplearning"  
#SBATCH --gres=gpu:8  
#SBATCH --ntasks=8  
# Number of cpu cores per task (OpenMP threads)  
#SBATCH -c 4  
#SBATCH --time=02:30:00
```


Prerequisites

```
[spaul@login21]$ groups  
kph m2_dgxqcd m2_gitter m2_him_th him-thfl
```

Typical script

```
#!/bin/bash  
  
#SBATCH -J jobname  
#SBATCH -A m2_dgxqcd  
#SBATCH --partition="deeplearning"  
#SBATCH --gres=gpu:8  
#SBATCH --ntasks=8  
# Number of cpu cores per task (OpenMP threads)  
#SBATCH -c 4  
#SBATCH --time=02:30:00
```

section 4

- 1 Introduction
 - Hardware & Softwares available
- 2 Compilation
- 3 Running
- 4 Lattice QCD specific codebases

Lattice QCD specific codebases: wit

Please send your email ids to join,

```
https://github.com/MainzLattice/wit
```

Salient features:

- CPU support: openQCD solvers + QDP data structures.
Grid solvers + Grid data structures.
- GPU support: QUDA (CUDA based) solvers, Grid(sycl based) solvers.
- In production: 2-pt Vector current correlators. Distillation for H -dibaryon.

Looking forward to your contributions.