

Bull user guide draft

1. Introduction to accelerators in HPC

1.1 GPU-computing

1.2 Many Integrated Core (MIC) -computing

2. Taito-GPU

2.1 Taito-GPU hardware and operating system

2.2 Getting access to Taito-GPU

2.3 Module and storage environment on Taito-GPU

2.4 Compiling & linking GPU-programs

2.4.1 Introduction

2.4.2 CUDA approach

2.4.3 OpenACC approach

2.4.3.1 C/C++ version of OpenACC DAXPY

2.4.3.2 Fortran version of OpenACC DAXPY

2.5 Running GPU-programs

2.5.1 Introduction

2.5.2 Running under GNU (plus CUDA 6.0) environment

2.5.3 Running under PGI (plus CUDA 6.0) environment

2.5.4 Running OpenACC version under PGI and CUDA 6.0 environment

2.5.4.1 Output from the C/C++ & OpenACC version of DAXPY

2.5.4.2 Output from the Fortran/OpenACC version of DAXPY

2.6 Deploying GPU + MPI-programs

2.7 Profiling GPU-programs

2.8 Debugging GPU-programs

3. Taito-MIC

3.1 Taito-MIC hardware and operating system

3.2 Getting access to Taito-MIC

3.3 Module and storage environment on Taito-MIC

3.4 Compiling & linking programs on Taito-MIC

3.4.1 Introduction

3.4.1.1 The recommended build directory

3.4.2 OpenMP builds (for host CPU and native MIC modes)

3.4.2.1 Host CPU-only build

3.4.2.2 Hello World in C : pure OpenMP, no offloads – for host CPU and native MIC runs

3.4.2.3 Native MIC build

3.4.3 Offload builds

3.4.3.1 Hello World in C : using OpenMP4.0 target -directive

3.4.3.2 OpenMP4.0 offload version of Fortran DAXPY

3.4.3.3 Important remark on building host CPU offload versions

3.5 Running programs on Taito-MIC

3.6 Deploying MPI-programs on Taito-MIC

3.7 Profiling programs on Taito-MIC

3.8 Debugging programs on Taito-MIC