Welcome to Nvidia HPC SDK Training

Jan 12-13, 2022

NERSC: Helen He, Chris Daley
OLCF: Suzanne Parete-Koon, Tom Papatheodore, Sherry Ray
ALCF: Yasaman Ghadar, Ray Loy

Introduction

- NVIDIA HPC SDK is a Comprehensive Suite of Fortran, C, and C++ Development Tools and Libraries
  - **NVIDIA compiler is default and recommended compiler for Perlmutter GPU**
- Hands-on training provided by NVIDIA to NERSC/OLCF/ALCF users
  - Thanks **Jeff Larkin, Brent Leback, Max Katz, Matt Stack, Robert Searles**
- Topics include:
  - GPU architecture and HPC SW developer considerations
  - Standard Language Acceleration and Libraries
  - OpenACC
  - OpenMP offload
  - CUDA
  - Profiling tools

Some Logistics (1)

- Muted upon joining Zoom due to large number of attendees
- Please change your name in Zoom session as “first_name last_name”
  - Click “Participants”, then “More” next to your name to rename
- Live “Captions” and “View Full Transcripts” are enabled
- Please ask your questions in Slack (preferred) instead of Zoom chat
  - https://tinyurl.com/nvhpc-jan2022-slack

Some Logistics (2)

- Slides will be uploaded to the #presentations in Slack
- Videos will be available at a later date
- Hands-on exercises
  - git clone https://github.com/olcf/NVHPC2022
  - NERSC users: use Perlmutter
  - OLCF users: use Summit
  - Other users: use NERSC training account on Perlmutter
- Please help us with answering the survey afterwards
  - https://tinyurl.com/nvhpc-jan2022-survey

Agenda (1)

Day 1, 9:00 am - 12:30 pm (Pacific time), January 12, Wednesday

Time (PDT)	Topic	Presenters
9:00 am	Introduction (10 min)	NERSC/OLCF/ALCF
9:10 am	Introduction and Updates to NVHPC SDK (55 min)	Jeff Larkin
10:05 am	GPU Architecture and HPC SW Developer Considerations (25 min)	Brent Leback / Max Katz
10:30 am	Break (15 min)	
10:45 am	NVHPC Standard Language Acceleration, Libraries (30 min)	Brent Leback / Max Katz / Matt Stack
11:15 am	Introduction to Nsight Systems and Nsight Compute (30 min)	Max Katz
11:45 am	Break + Lab + Demo: stdpar, profiling (45 min)	Brent Leback / Max Katz / Matt Stack/ Robert Searles
12:30 pm	End of Day 1	

Agenda (2)

Day 2, 9:00 am - 12:30 pm (Pacific time), January 13, Thursday

Time (PDT)	Topic	Presenters
9:00 am	NVHPC OpenACC (45 min)	Brent Leback
9:45 am	NVHPC OpenMP (45 min)	Brent Leback
10:30 am	Break (15 min)	
10:45 am	Introduction to CUDA (30 min)	Max Katz
11:15 am	Putting it Together (15 min)	Brent Leback
11:30 am	Break + Lab + Demo: OpenACC, OpenMP, and more (60 min)	Brent Leback / Max Katz / Matt Stack / Robert Searles
12:30 pm	End of Day 2	

NERSC Perlmutter Usage Info

- NERSC users are added to the ntrain4 project
- Training account expires Jan 18
- Perlmutter Compute node reservations. To access, use:
  - Jan 12, 10:30-13:30: #SBATCH --reservation=nvhpc_day1 -A ntrain4_g
  - Jan 13, 10:30-13:30: #SBATCH --reservation=nvhpc_day2 -A ntrain4_g
- Outside of reservations, use:
  - #SBATCH -A ntrain4_g or #SBATCH -A <your_account>_g
- Perlmutter documentation is here
  - https://docs.nersc.gov/systems/perlmutter/

NVHPC on Perlmutter Cheat sheet

To use NVHPC: *module load nvidia/21.9* (loaded by default)

- Using native Nvidia compilers: *nvfortran, nvc, nvc++, nvcc*

  or using Cray compiler wrappers*: ftn, cc, CC*
  - Stdpar: -stdpar -I<some_libs>
  - OpenMP: *-mp=gpu -gpu=cc80*
  - OpenACC: *-acc=gpu -gpu=cc80*
  - CUDA*: -cuda* (or use nvcc to build)
- Add -Minfo flag when compile to show detailed information
- Set runtime environment variable NVCOMPILER_ACC_NOTIFY to
  - 1: Show kernel launches information
  - 2: Show data transfers between CPU and Devices
  - 3: Show both above
  - works for OpenACC or OpenMP offload programs, not stdpar or CUDA
- Submit Slurm batch scripts with *sbatch my_batch_script*

Sample Compile and Run Batch Script

Stdpar code:

% nvc++ -stdpar=gpu -gpu=cc80 -Minfo -o code.exe code.cpp

or % CC -stdpar=gpu -gpu=cc80 -Minfo -o code.exe code.cpp

OpenMP offload code:

% nvfortran -mp=gpu -gpu=cc80 -Minfo -o code.exe code.f90

or % ftn -mp=gpu -gpu=cc80 -Minfo -o code.exe code.f90

- Compiling/Building software for Perlmutter
  - https://docs.nersc.gov/systems/perlmutter/#compilingbuilding-software
- Running Jobs on Perlmutter GPU
  - https://docs.nersc.gov/systems/perlmutter/running-jobs/

```
% cat my_batch_script
#!/bin/bash
#SBATCH -N 1
#SBATCH -C gpu
#SBATCH -G 1
#SBATCH -t 15
#SBATCH -c 128
#SBATCH -A ntrain4_g
#SBATCH --reservation=nvhpc_day1
#SBATCH --ntasks-per-node=1
#SBATCH --gpus-per-task=1

export NVCOMPILER_ACC_NOTIFY=3
srun ./code.exe

% sbatch my_batch_script
```

Hands-On NVHPC Perlmutter Tips

- Nsight-systems and nsight-compute profiling
  - https://docs.nersc.gov/performance/readiness/#profiling-tools
  - run with nsys profile --stats=true ./code.exe
- Recommend to install NX (NoMachine)
  - Improves X-forwarding when using GUI tools
  - Instructions at: https://docs.nersc.gov/connect/nx/
  - Good for the profiling tools hands-on with GUI

OLCF Summit Usage Info

- Summit reservations have been set up for this training. Users will need the follow in their batch scripts to access the reservation:

  On 1/12/22 #BSUB -U NVIDIA_SDK_1

  On 1/13/22 #BSUB -U NVIDIA_SDK_2

- Summit documentation is here https://docs.olcf.ornl.gov/systems/summit_user_guide.html

NVHPC on Summit Cheat sheet

To load NVHPC: *module load nvhpc/21.9*

- Compile with *nvfortran, nvc, nvc++*
  - Enable OpenMP: *-mp=gpu -gpu=cc70*
  - Enable OpenACC: *-acc=gpu -gpu=cc70*
  - Add -Minfo flag when compiling to show detailed information
- Compiler docs for Summit: https://docs.olcf.ornl.gov/systems/summit_user_guide.html#compilers

Running Jobs

- Example Summit batch script:
- https://docs.olcf.ornl.gov/systems/summit_user_guide.html#batch-script
  - Submit batch scripts with: *bsub your_batch_script.lsf*

Hands-On NVHPC Summit Tips

- Nsight-systems and nsight-compute profiling: https://docs.olcf.ornl.gov/systems/summit_user_guide.html?highlight=nsight#nvidia-nsight-compute
  - Make sure your write your profiles to GPFS
  - You cannot use the Nsight GUI on Summit, you will need to use scp or Globus to download your profile.
  - You will need to download the Nsight UI for your local machine to read the reports
    - Two ways to get it:
      - download the UI for your local system There are versions for Windows, Mac, and Linux (x86). You will need to register as an NVIDIA developer first.
      - Download the CUDAToolkit https://developer.nvidia.com/cuda-downloads.

Thank you!


NERSC
OAK RIDGE
National Laboratory
LEADERSH
COMPUTIN
FACILITY
Argonne
NATIONAL LABORATORY
Leadership
Computing
Facility