

Hands-on with Intel Xeon Phi

Lab 3: Symmetric Computing MPI on MIC and CPUs



Bill Barth
Kent Milfeld
Dan Stanzione





Lab 3

- What you will learn:
 - How to compile MPI code with Intel IMPI library.
 - Launch MPI code on MIC only (**native mode**), and on both CPUs and MIC (**symmetric mode**).
- What you will do:
 - Run in MPI symmetric mode using **mpiexec.hydra**
 - Run in MPI symmetric mode using **ibrun.symm**
 - Run an MPI code with offload.

Lab 3

MPI Computing on MIC and CPUs



Part 0 – Grab the Lab Files

- Login to Stampede.

```
$ ssh <username>@stampede.tacc.utexas.edu
```

- Untar the mic_mpi.tar file (in ~train00) into your directory:

```
$ tar xvf ~train00/mic_mpi.tar
```

- Use idev to launch a 1-node interactive session in the development queue

```
$ idev
```

Part 1 – Native MPI Execution (mpiexec.hydra)



- Move into the mpi/symmetric dir. Change to the Intel impi environment:

```
$ cd mpi/symmetric
$ module swap mvapich2 impi
```
- Create a MIC and CPU binary for pi_hybrid. (or use *.F90 files)

```
$ mpicc -openmp pi_hybrid.c -mmic -o pi_hybrid.mic
$ mpicc -openmp pi_hybrid.c -o pi_hybrid.cpu
```
- Source setup file to add IMPI environment variables:

```
$ source ./setup_mic_env.sh
```

Execute on MIC only (Native): 8 tasks and 30 threads/task:

```
$ mpiexec.hydra \
  -env OMP_NUM_THREADS 30 \
  -env LD_LIBRARY_PATH $MIC_LD_LIBRARY_PATH \
  -env I_MPI_PIN_MODE mpd \
  -n 8 -host mic0 \
  ./pi_hybrid.mic
```

#This is rather complicated!
or use ./do_native script

5

Part 1 –Symmetric MPI Execution (mpiexec.hydra)



Symmetric Execution (with hybrid on MIC)

Execute on MIC + 16 tasks on the host (CPU):

```
$ mpiexec.hydra \  
  -n 16 -host localhost ./pi_hybrid.cpu : \  
  -env OMP_NUM_THREADS 30 \  
  -env LD_LIBRARY_PATH $MIC_LD_LIBRARY_PATH \  
  -env I_MPI_PIN_MODE mpd \  
  -n 8 -host mic0 \  
  ./pi_hybrid.mic #Even more complicated!  
                  # or use ./do_sym script
```

- Hmm, why do they call this symmetric 😊 ?



Part 2 – Symmetric MPI (ibrun.symm)

- Let's try this with TACC's `ibrun.symm` launcher:

- 1st create environment variables for the MIC:

```
$ export MIC_PPN=8
$ export MIC_OMP_NUM_THREAD=30
$ export MIC_MY_NSLOTS=8
```

- Now use `ibrun.symm` to launch a MIC (`-m`) binary.

```
$ ibrun.symm -m pi_hybrid.mic
```

- Now launch MIC (`-m`) and CPU (`-c`) binaries:

```
$ ibrun.symm -m pi_hybrid.mic -c pi_hybrid.cpu
```

See `do_ibrun` scripts for setting env. & doing launch.

7



Part 3–MPI with Offload

- Change back to mvapich2 MPI. Swap and look 4 mvapich2 in module list.
`$ module swap impi mvapich2; module list`
- Move into the offload directory:
`$ cd $HOME/mpi/offload`
- Create an offload binary for pi_offload.
`$ mpicc -O3 -openmp pi_offload.c -o pi_offload`
- Set offload reporting on, # of threads, etc.
`$ export OFFLOAD_REPORT=2`
`$ export MIC_ENV_PREFIX=MIC MIC_OMP_NUM_THREADS=120`
`$ export MIC_KMP_AFFINITY=scatter`
- Launch MPI-offload code with ibrun, use 2 MPI tasks:
`$ ibrun -n 2 -o 0 ./pi_offload` #-o offset in host list
See do_off script for setting env. & doing launch.
- Compare pi_offload.c and pi_hybrid.c codes in the offload and symmetric directories.