

Tips and Tricks for Shared Memory Parallelization

Dirk Colbry
colbrydi@msu.edu
Research Specialist
Institute for Cyber-Enabled Research

MICHIGAN STATE UNIVERSITY © 2012 Michigan State University Board of Trustees. ICER

Pleasantly Parallel

MICHIGAN STATE UNIVERSITY ICER

Loosely Coupled

MICHIGAN STATE UNIVERSITY ICER

Tightly Coupled

MICHIGAN STATE UNIVERSITY ICER

Shared Memory

- Different threads communicate through pointers to the same memory access
- Problems can occur if different threads write the same memory at the same time
- Flags (also called locks and/or semaphores) are used to allow only one thread to access memory at the same time

MICHIGAN STATE UNIVERSITY ICER

Intel10

Machine (24GB)
 NUMANode #0 (12GB)
 Socket #0
 L3 #0 (12MB)
 L2 #0 (256KB) L2 #1 (256KB) L2 #2 (256KB) L2 #3 (256KB)
 L1 #0 (32KB) L1 #1 (32KB) L1 #2 (32KB) L1 #3 (32KB)
 Core #0 Core #1 Core #2 Core #3
 PU #0 PU #1 PU #2 PU #3 PU #4 PU #5
 PU #6 PU #7
 NUMANode #1 (12GB)
 Socket #1
 L3 #1 (12MB)
 L2 #4 (256KB) L2 #5 (256KB) L2 #6 (256KB) L2 #7 (256KB)
 L1 #4 (32KB) L1 #5 (32KB) L1 #6 (32KB) L1 #7 (32KB)
 Core #4 Core #5 Core #6 Core #7
 PU #8 PU #9 PU #10 PU #11 PU #12 PU #13
 PU #14 PU #15

MICHIGAN STATE UNIVERSITY ICER

amdgpu

The diagram illustrates a row of 10 AMDGPU nodes. Each node is represented by a small box containing a grid of 16 squares, representing the processor layout on the node.

MICHIGAN STATE UNIVERSITY
ICER

NUMA

The diagram illustrates a NUMA (Non-Uniform Memory Access) architecture. It shows a grid of 10 nodes, each with a grid of processors. The nodes are interconnected, representing a shared memory system where access time varies based on the physical location of the processor and memory.

MICHIGAN STATE UNIVERSITY
ICER

Shared memory submission scripts

- Typically one node with multiple processors per node (ppn)
 - #PBS -l nodes=1:ppn=8
- Different programs use different methods to tell them how many processors to use
 - Command line arguments
 - Environment variables

MICHIGAN STATE UNIVERSITY
ICER

Example: shared memory Script

- Bowtie uses shared memory parallelization
- Get the bowtie example


```
> module load powertools
> getexample bowtie
```
- Change to the bowtie directory


```
> cd ./bowtie
```
- Look at the submission script


```
> less ./bowtie.qsub
```
- Run the job


```
> qsub bowtie.qsub
```

MICHIGAN STATE UNIVERSITY
ICER

OpenMP

- Common Shared Memory parallelization
- Single program runs in many threads
- Really easy to pick loops that are parallel and split them into multi threads
- Minor modifications to code that can be written not to affect single

MICHIGAN STATE UNIVERSITY
ICER

simpleOMP.qsub example

```
#!/bin/bash -login
#PBS -l walltime=00:01:00
#PBS -l nodes=1:ppn=8,feature=gbe

cd ${PBS_O_WORKID}
OMP_NUM_THREADS=`cat ${PBS_NODES} | wc -l`
export NUM_OMP_THREADS
./simpleOMP

qstat -f ${PBS_JOBID}
```

MICHIGAN STATE UNIVERSITY
ICER

MKL

- Does shared memory Paralliation built in:
 - FFTW
 - BLAS

simpleOMP.qsub example

```
#!/bin/bash -login
#PBS -l walltime=00:01:00
#PBS -l nodes=1:ppn=8,feature=gbe

module load MKL
cd ${PBS_O_WORKID}
MKL_NUM_THREADS=${PBS_NUM_PPN}
export MKL_NUM_THREADS
./mklJob

qstat -f ${PBS_JOBID}
```