

LAB

NUMA Control for Hybrid Applications



Hang Liu

October, 23rd, 2012

What you will learn

- Using numactl in execution of serial, MPI and a 2x# (2 tasks each with # threads) hybrid code
- For on site users, untar the file numahybrid.tar
 - cd (Start in your home directory.)
 - tar xvf ~train00/numahybrid.tar (extract files)
 - cd numahybrid
- For remote users, download the lab_numahybrid_external.tar, and refer the README in it about how to experiment numactl on your local system.

numactl_serial on Lonestar

The memory intensive daxpy code is run on two different sockets using local, interleave and off-socket-memory policies. The commands below make the daxpy executable and run it with numa control commands. See the job script and the table on the next page for the numa options. Run the job and report the times and relative performance.

- Execute:

```
cd numactl_serial (change directory to numactl_serial)
```

```
make
```

```
qsub job (submits job)
```

numactl_serial on Lonestar

- From the job output fill in the table.

Command	Time (secs)
No numactl options	
numactl -l --physcpubind 0	
numactl -l --physcpubind 6	
numactl -l --physcpubind 3	
numactl -l --physcpubind 9	
numactl -i all --physcpubind 0	
numactl -i all --physcpubind 6	
numactl -i all --physcpubind 3	
numactl -i all --physcpubind 9	
Numactl -m 0 --physcpubind 11	

Rank the performance of local, interleave, and off-Socket-memory policies.

1.)

2.)

3.)

4.)

Why ?

numactl_2x1, numactl_2x2 on Lonestar

The daxpy code is run as 2 tasks in a node (2x1) and 2 tasks with 2 threads in a node(2x2). “cd” down to directories numactl_2x1 and numactl_2x2, respectively, and follow the instructions below.

- Execute:

```
cd numactl_2x1 or numactl_2x2
```

```
make
```

```
qsub job (submits job)
```

numactl_2x1, numactl_2x2 on Lonestar

- From the job output fill in the table.

Command (4x1)	Time (secs)
<no numactl> numactl -l	
numactl -i all	
numactl tacc_affinity	

Command (4x4)	Time (secs)
<no numactl> numactl -l	
numactl -i all	
numactl tacc_affinity	

Rank the 2x1 performance

1.)

2.)

3.)

4.)

Rank the 2x2 performance of.

1.)

2.)

3.)

4.)

Why ?