Cornell University
Center for Advanced Computing

Hybrid Programming with OpenMP and MPI

Steve Lantz
Senior Research Associate
Cornell CAC

Workshop: Introduction to Parallel Computing on Ranger, May 20, 2010

Based on materials developed by Kent Milfeld at TACC

RAM Arrangement on Ranger

- *Distributed memory*
  - each node has its own local memory
  - not directly addressable from other nodes
- *Multichip nodes*
  - each node has 4 chips (sockets)
- *Multicore chips*
  - each chip (socket) has 4 cores
- *Memory spans entire nodes*
  - node's full local memory is addressable from any core in any socket
- *Memory is attached to sockets*
  - 4 cores sharing the socket have fastest access to attached memory

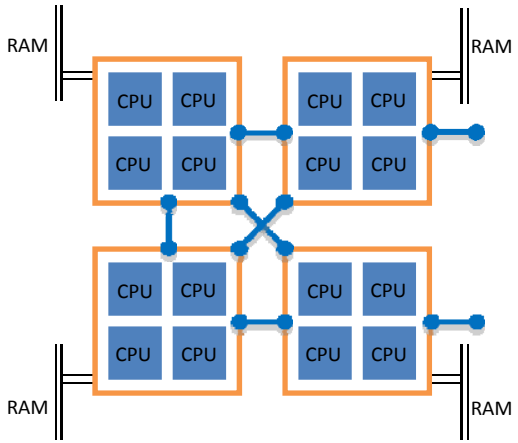

Dealing with NUMA

How do we deal with NUMA (Non-Uniform Memory Access)?

Two standard models for parallel programs -

- Threads for shared memory
  - parent process uses pthreads or OpenMP to fork multiple threads
  - threads share the same virtual address space
  - also known as SMP = Symmetric MultiProcessing
- Message passing for distributed memory
  - processes use MPI to pass messages (data) between each other
  - each process has its own virtual address space

If we attempt to combine both types of models –

- ***Hybrid programming***
  - try to exploit the whole shared/distributed memory hierarchy

Why Hybrid? Or Why Not?

Why hybrid?

- Eliminates domain decomposition at node
- Automatic memory coherency at node
- Lower (memory) latency and data movement within node
- Can synchronize on memory instead of barrier

Why not hybrid?

- An SMP algorithm created by aggregating MPI parallel components on a node (or on a socket) may actually run slower
- Possible waste of effort

Motivation for Hybrid

- Balance the computational load
- Reduce memory traffic, especially for memory-bound applications

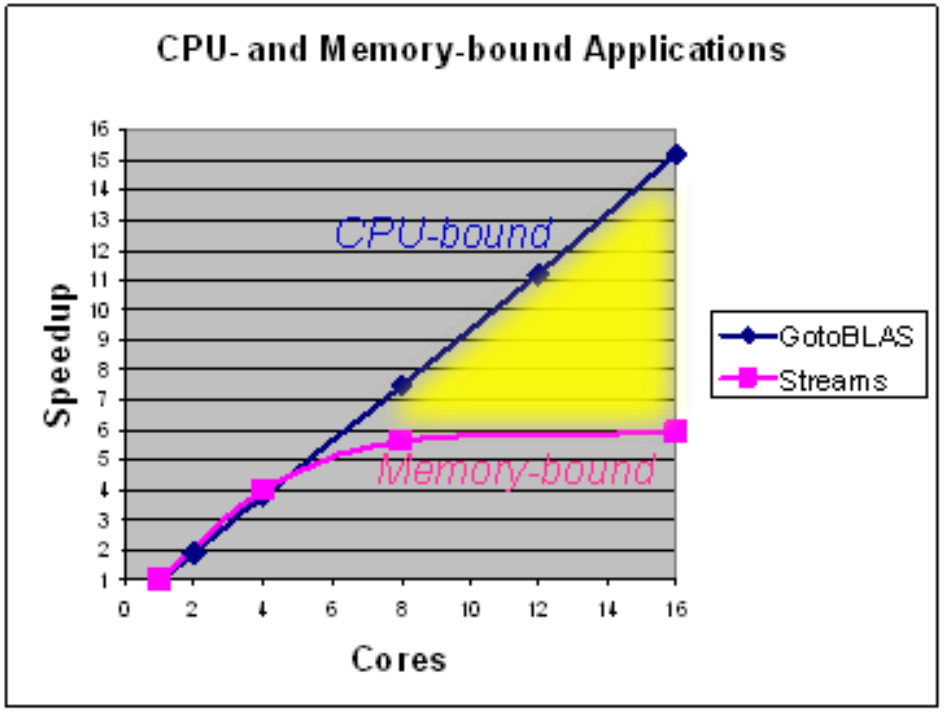

Two Views of a Node

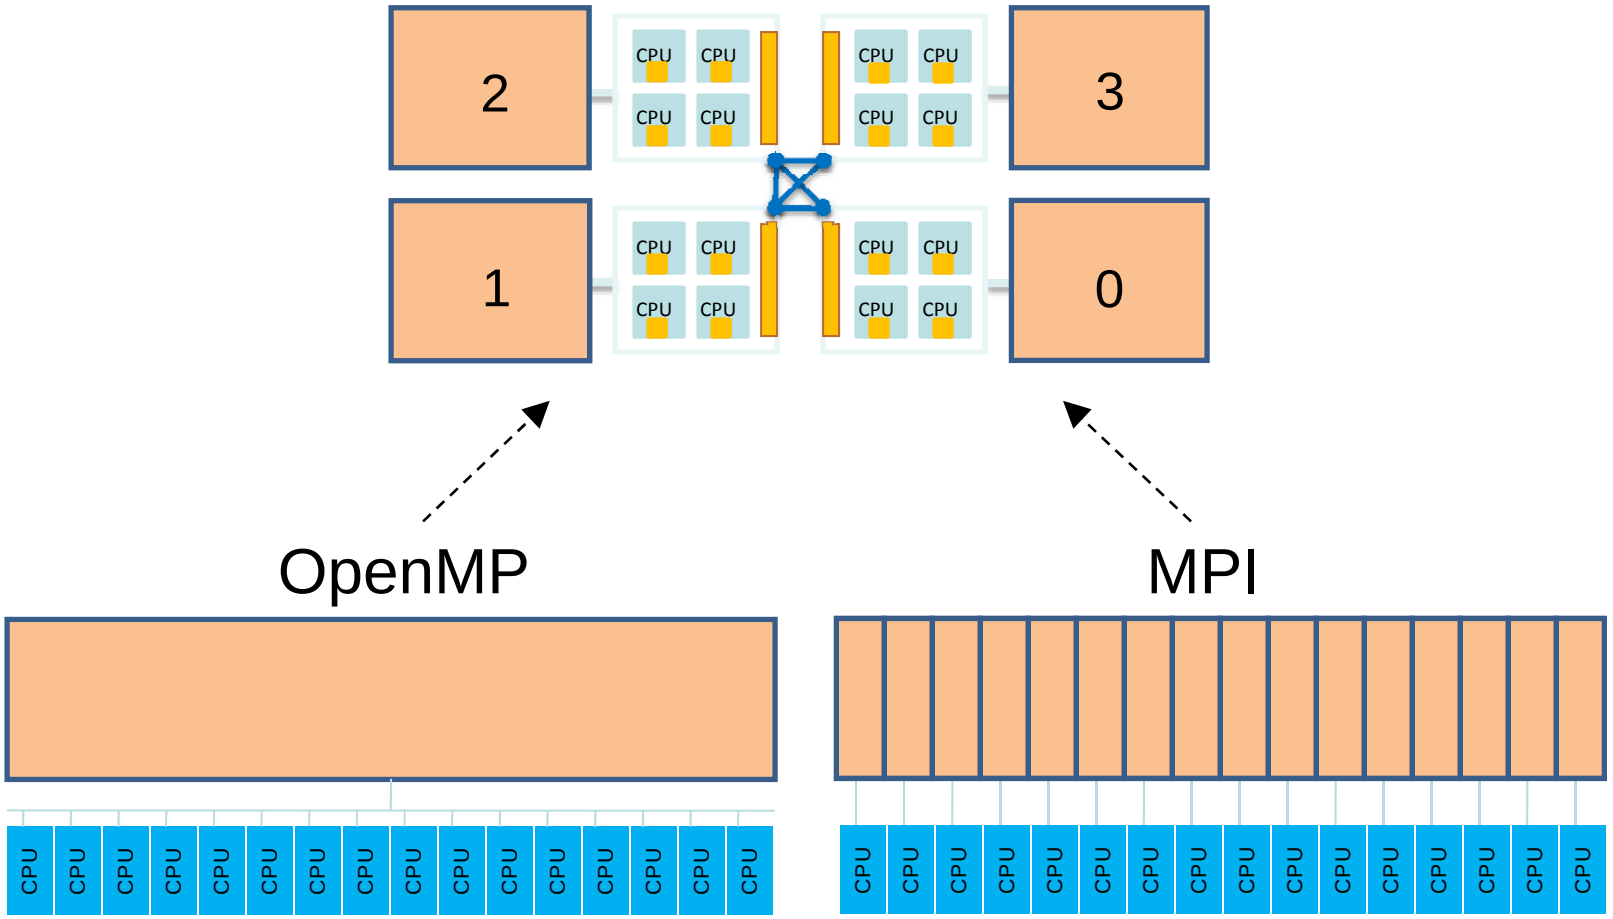

Two Views = Two Ways to Write Parallel Programs

- OpenMP (or pthreads) only
  - launch one process per node
  - have each process fork one thread (or maybe more) per core
  - share data using shared memory
  - can't share data with a different process (except maybe via file I/O)
- MPI only
  - launch one process per core, on one node or on many
  - pass messages among processes without concern for location
  - (maybe create different communicators intra-node vs. inter-node)
  - ignore the potential for any memory to be shared
- *With hybrid OpenMP/MPI programming, we want each MPI process to launch multiple OpenMP threads that can share local memory*

Some Possible MPI + Thread Configurations

- Treat each *node* as an SMP
  - launch a single MPI process per node
  - create parallel threads sharing full-node memory
  - typically want 16 threads/node on Ranger, e.g.
- Treat each *socket* as an SMP
  - launch one MPI process on each socket
  - create parallel threads sharing same-socket memory
  - typically want 4 threads/socket on Ranger, e.g.
- No SMP, ignore shared memory (all MPI)
  - assign an MPI process to each core
  - in a master/worker paradigm, one process per node may be master
  - not really hybrid, may at least make a distinction between nodes

Creating Hybrid Configurations

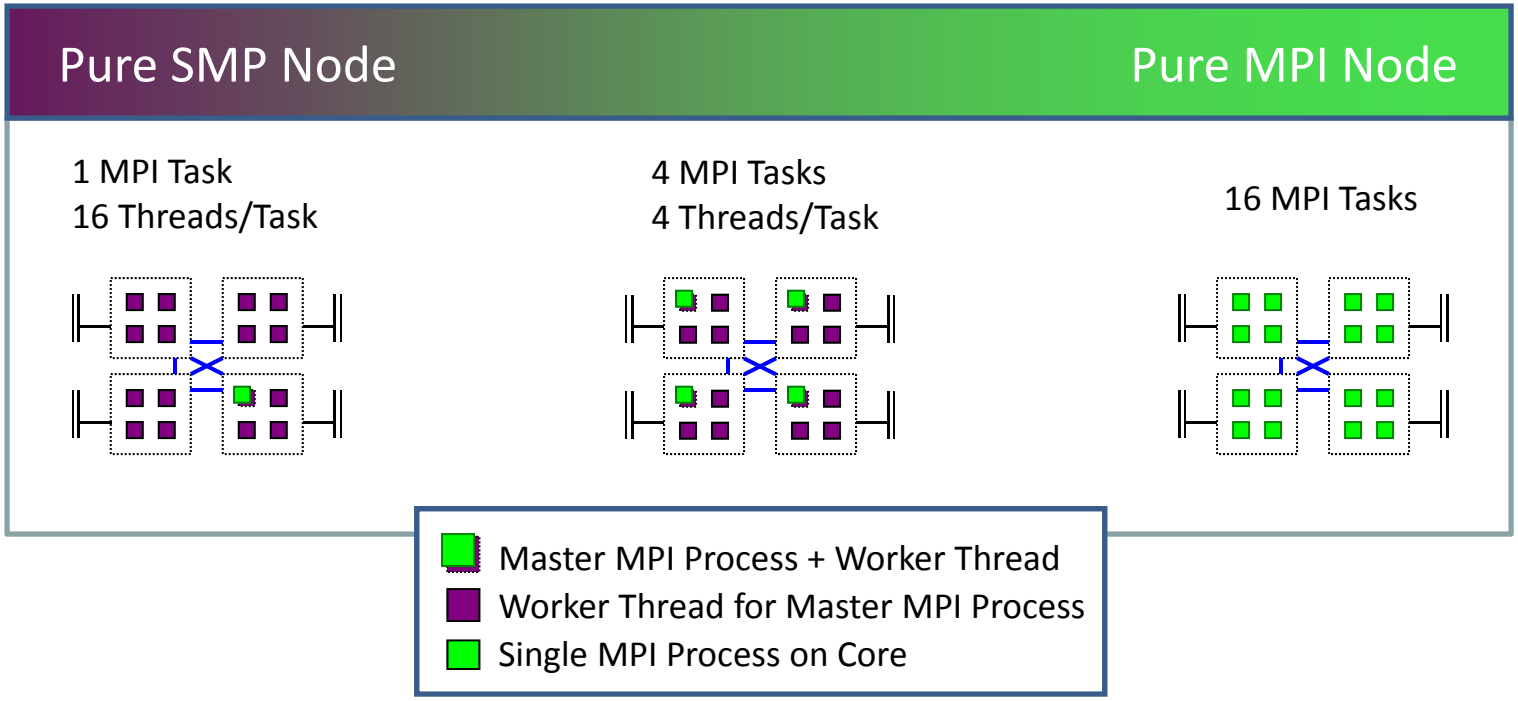

To achieve configurations like these, we must be able to:

- Assign to each process/thread an *affinity* for some set of cores
- Make sure the *allocation* of memory is appropriately matched

NUMA Operations

Where do processes, threads, and memory allocations get assigned?

- If memory were completely uniform, there would be no need to worry about questions like, “where do processes go?”
- Only for NUMA is the placement of processes/threads and allocated memory (NUMA control) of any importance

The default NUMA control is set through policy

- The policy is applied whenever a process is executed, or a thread is forked, or memory is allocated
- These are all events that are directed from within the kernel

NUMA control is managed by the kernel.

NUMA control can be changed with numactl.

Process Affinity and Memory Policy

- One would like to set the *affinity* of a process for a certain socket or core, and the *allocation* of data in memory relative to a socket or core.
- Individual users can alter kernel policies
  (setting Process Affinity and Memory Policy == PAMPer)
  - users can PAMPer their own processes
  - root can PAMPer any process
  - careful, libraries may PAMPer, too!
- Means by which Process Affinity and Memory Policy can be changed:
  - dynamically on a running process (knowing process id)
  - at start of process execution (with wrapper command)
  - within program through F90/C API

Using numactl

```
numactl <option socket(s)/core(s)> ./a.out
```

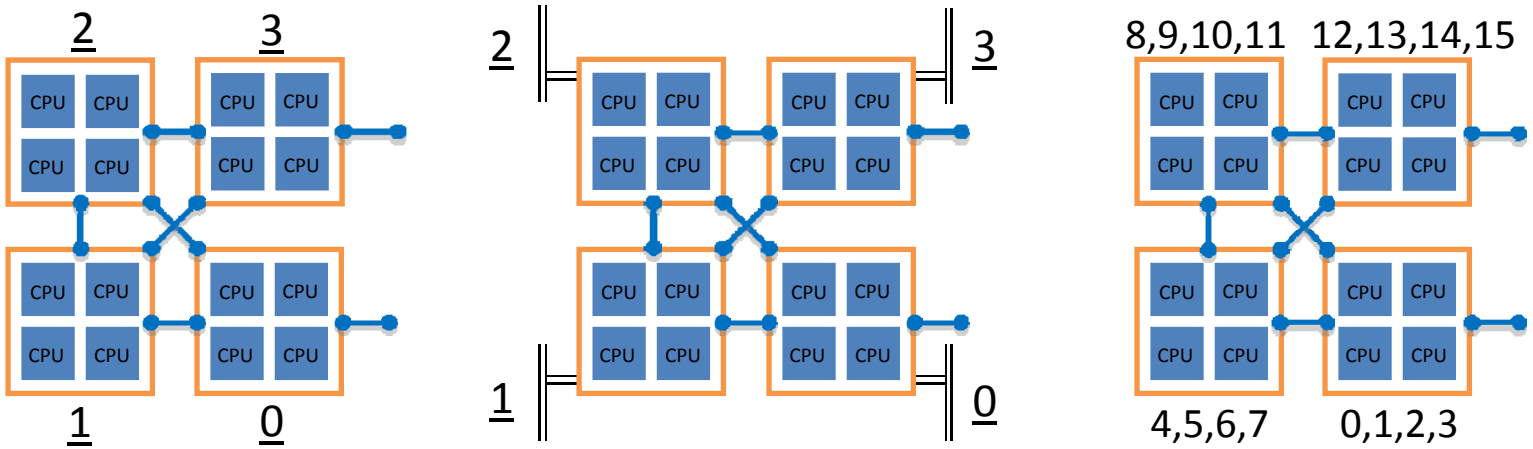

For a Process: Socket Control	For a Process's Memory: Socket Control	For a Process: Core Control
socket assignment -N	memory allocation -l -i --preferred –m (local, interleaved, pref., mandatory)	core assignment -C

Quick Guide to numactl

Socket Affinity	-N	{0,1,2,3}	Execute process on cores of this (these) socket(s) only.
Memory Policy	-l	no argument	Allocate on current socket; fallback to any other if full.
Memory Policy	-i	{0,1,2,3}	Allocate round robin (interleave) on these sockets. No fallback.
Memory Policy	--preferred=	{0,1,2,3} select one	Allocate on this socket; fallback to any other if full.
Memory Policy	-m	{0,1,2,3}	Allocate only on this (these) socket(s). No fallback.
Core Affinity	-C	{0,1,2,3,4,5,6,7, 8,9,10,11,12,13, 14,15}	Execute process on this (these) core(s) only.

SMP Nodes

Hybrid batch script for 16 threads/node

- Make sure **1 process per node** is created
- Specify **total cores allocated** by batch (nodes x 16)
- Set number of **threads for each process**
- PAMPering at **job level**
  - controls behavior for ALL tasks
  - no simple/standard way to control thread-core affinity

job script (Bourne shell)	job script (C shell)

`...`
`#! -pe 1way 192`
`...`
`export OMP_NUM_THREADS=16`
`ibrun numactl -i all ./a.out`
	`...`
	`#! -pe 1way 192`
	`...`
	`setenv OMP_NUM_THREADS 16`
	`ibrun numactl -i all ./a.out`

SMP Sockets

Hybrid batch script for 4 tasks/node, 4 threads/task

Example script setup for a square (6x6 = 36) processor topology...

- Make sure **4 processes per node** are created (one per socket)
- Specify **total cores allocated** by batch (nodes x 16)
- Specify **actual cores used** with MY_NSLOTS
- Set number of **threads for each process**
- PAMPering at **process level**, must create script to manage affinity

job script (Bourne shell)	job script (C shell)

`...`
`#! -pe 4way 48`
`export MY_SLOTS=36`
`export OMP_NUM_THREADS=4`
`ibrun numa.sh`
	`...`
	`#! -pe 4way 48`
	`setenv MY_NSLOTS 36`
	`setenv OMP_NUM_THREADS 4`
	`ibrun numa.csh`

Script for Socket Affinity

- Example script to extract rank for numactl options, a.out execution (TACC MPI systems always assign ranks sequentially on a node)
- No simple/standard way to control thread-core affinity

numa.sh

```
#!/bin/bash
export MV2_USE_AFFINITY=0
export MV2_ENABLE_AFFINITY=0

#TasksPerNode
TPN=`echo $PE|sed 's/way//'`
[ ! $TPN ] && echo TPN null!
[ ! $TPN ] && exit 1

sk=$(( $PMI_RANK % $TPN ))

numactl -N $sk -m $sk ./a.out
```

numa.csh

```
#!/bin/csh
setenv MV2_USE_AFFINITY 0
setenv MV2_ENABLE_AFFINITY 0

#TasksPerNode
set TPN=`echo $PE|sed 's/way//'`
if(! ${%TPN}) echo TPN null!
if(! ${%TPN}) exit 0

@ sk = $PMI_RANK % $TPN

numactl -N $sk -m $sk ./a.out
```

Basic Hybrid Program Template

Start with MPI initialization

(Serial regions are executed by the master thread of the MPI process)

Create OMP parallel regions within each MPI process

- MPI calls may be allowed here too
- MPI rank is known to all threads

Call MPI in single-threaded regions

Finalize MPI

```
MPI_Init
...
MPI_Call
...
    OMP parallel
    ...
    MPI_Call
    ...
    end parallel
...
MPI_Call
...
MPI_Finalize
```

Types of MPI Calls Among Threads

Single-threaded messaging

- Call MPI from a serial region
- Call MPI from a single thread within a parallel region

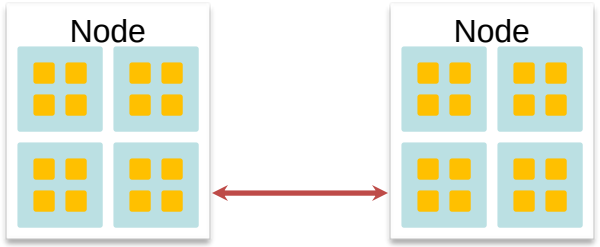

Rank to rank

Multi-threaded messaging

- Call MPI from multiple threads within a parallel region
- Requires an implementation of MPI that is thread-safe

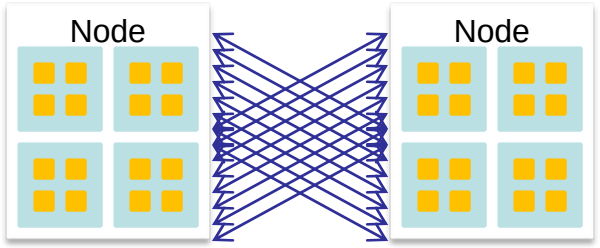

rank-thread ID to rank-thread ID

MPI-2 and Thread Safety

- **Consider thread safety when calling MPI from threads**
- Use MPI_Init_thread to select/determine the level of thread support
  - Supported in MPI-2, substitute for the usual MPI_Init
- Thread safety is identified/controlled by MPI's provided types
  - *Single* means no multi-threading
  - *Funneled* means only the master thread can call MPI
  - *Serialized* means multiple threads can call MPI, but only 1 call can be in progress at a time
  - *Multiple* means MPI is thread safe
- Monotonic values are assigned to parameters

  MPI_THREAD_SINGLE < MPI_THREAD_FUNNELED < MPI_THREAD_SERIALIZED < MPI_THREAD_MULTIPLE

MPI-2's MPI_Init_thread

Syntax:

```
call MPI_Init_thread(                                  irqd, ipvd, ierr)
int MPI_Init_thread (int *argc, char ***argv, int rqd, int *pvd)
int MPI::Init_thread(int& argc, char**& argv, int rqd)
```

- Input: **`rqd`**, or "required" (integer)
  - Indicates the desired level of thread support
- Output: **`pvd`**, or "provided" (integer)
  - Indicates the available level of thread support
- If thread level **`rqd`** is supported, the call returns **`pvd = rqd`**
- Otherwise, **`pvd`** returns the highest provided level of support

MPI-2 Thread Support Levels

Support Levels	Description
MPI_THREAD_SINGLE	Only one thread will execute.
MPI_THREAD_FUNNELED	Process may be multi-threaded, but only the main thread will make MPI calls (calls are “funneled” to main thread). *Default*
MPI_THREAD_SERIALIZE	Process may be multi-threaded, and any thread can make MPI calls, but threads cannot execute MPI calls concurrently (MPI calls are “serialized”).
MPI_THREAD_MULTIPLE	Multiple threads may call MPI, with no restriction.

Example: Single-Threaded MPI Calls

Fortran

```
include 'mpif.h'
program hybsimp


call MPI_Init(ie)
call MPI_Comm_rank(...irk,ie)
call MPI_Comm_size(...isz,ie)
!Setup shared mem, comp/comm

!$OMP parallel do
  do i=1,n
    <work>
  enddo

!Compute & communicate
call MPI_Finalize(ierr)
end
```

C

```
#include <mpi.h>
int main(int argc,
  char **argv) {
int rank, size, ie, i;
ie= MPI_Init(&argc,&argv[]);
ie= MPI_Comm_rank(...&rank);
ie= MPI_Comm_size(...&size);
//Setup shared mem, comp/comm

#pragma omp parallel for
  for(i=0; i<n; i++){
    <work>
  }

// compute & communicate
ie= MPI_Finalize();
}
```

Funneled MPI Calls via Master

- Must have support for **MPI_THREAD_FUNNELED or higher**
- Best to **use OMP_BARRIER**
  - there is no implicit barrier in the master workshare construct, OMP_MASTER
  - in the example, the master thread will execute a single MPI call within the OMP_MASTER construct
  - all other threads will be sleeping

Example: Funneled MPI Calls via Master

Fortran

```
include 'mpif.h'
program hybmas



!$OMP parallel

  !$OMP barrier
  !$OMP master

  call MPI_<Whatever>(...,ie)
  !$OMP end master
  !$OMP barrier

!$OMP end parallel
end
```

C

```
#include <mpi.h>
int main(int argc,
  char **argv) {
int rank, size, ie, i;

#pragma omp parallel
{
  #pragma omp barrier
  #pragma omp master
  {
    ie= MPI_<Whatever>(...);
  }
  #pragma omp barrier


}
}
```

Serialized MPI Calls and OpenMP

- Must have support for **MPI_THREAD_SERIALIZED or higher**
- Best to **use OMP_BARRIER only at beginning**, since there is an implicit barrier in the SINGLE workshare construct, OMP_SINGLE
  - Example is the simplest one: any thread (not necessarily master) will execute a single MPI call within the OMP_SINGLE construct
  - All other threads will be sleeping

Example: Serialized MPI Calls and OpenMP

Fortran

```
include 'mpif.h'
program hybsing


call MPI_Init_thread( &
MPI_THREAD_SERIALIZED,ipvd,ie)
!$OMP parallel

  !$OMP barrier
  !$OMP single

  call MPI_<Whatever>(...,ie)
  !$OMP end single
  !Don't need OMP barrier
!$OMP end parallel
end
```

C

```
#include <mpi.h>
int main(int argc,
  char **argv) {
int rank, size, ie, i;
ie= MPI_Init_thread(
MPI_THREAD_SERIALIZED,ipvd);
#pragma omp parallel
{
  #pragma omp barrier
  #pragma omp master
  {
    ie= MPI_<Whatever>(...);
  }
  //Don't need omp barrier
}
}
```

Overlapping Work & MPI Calls

- One core is capable of saturating the lanes of the PCIe network link...
  - Why use all cores to communicate?
  - Instead, **communicate using just one** or several cores
  - Can **do work with the rest** during communication
- Must have support for **MPI_THREAD_FUNNELED or higher** to do this
- Can be difficult to manage and load-balance!

Example: Overlapping Work & MPI Calls

Fortran

```
include 'mpif.h'
program hybsing


!$OMP parallel

  if (ithread .eq. 0) then
  call MPI_<Whatever>(...,ie)
  else
    <work>
  endif

!$OMP end parallel
end
```

C

```
#include <mpi.h>
int main(int argc,
  char **argv) {
int rank, size, ie, i;
#pragma omp parallel
{
  if (thread == 0){
    ie= MPI_<Whatever>(...);
  }
  if(thread != 0){
    <work>
  }
}
}
```

Multiple Threads Calling MPI

- Thread ID as well as rank can be used in communication
- Technique is illustrated in multi-thread “ping” (send/receive) example

Example: Multiple Threads Calling MPI

```
call mpi_init_thread( MPI_THREAD_MULTIPLE, iprovided,ierr)
call mpi_comm_rank(MPI_COMM_WORLD,  irank, ierr)
call mpi_comm_size( MPI_COMM_WORLD,nranks, ierr)
…
!$OMP parallel private(j, ithread, nthreads)
  nthreads=OMP_GET_NUM_THREADS()
  ithread   =OMP_GET_THREAD_NUM()
  call pwork(ithread, irank, nthreads, nranks…)
  if(irank == 0) then
    call mpi_send(ithread,1,MPI_INTEGER, 1,   ithread, MPI_COMM_WORLD, ierr)
  else
    call mpi_recv(          j,1,MPI_INTEGER, 0,   ithread, MPI_COMM_WORLD, istat, ierr)
    print*, "Yep, this is ",irank," thread ", ithread," I received from ", j
  endif
!$OMP END PARALLEL
end
```

Communicate between ranks.

Threads use tags to differentiate.

NUMA Control in Code

- Within a code, **Scheduling Affinity** and **Memory Policy** can be examined and changed through:
  - sched_getaffinity, sched_setaffinity
  - get_memorypolicy, set_memorypolicy
- To make scheduling assignments, set bits in a mask:

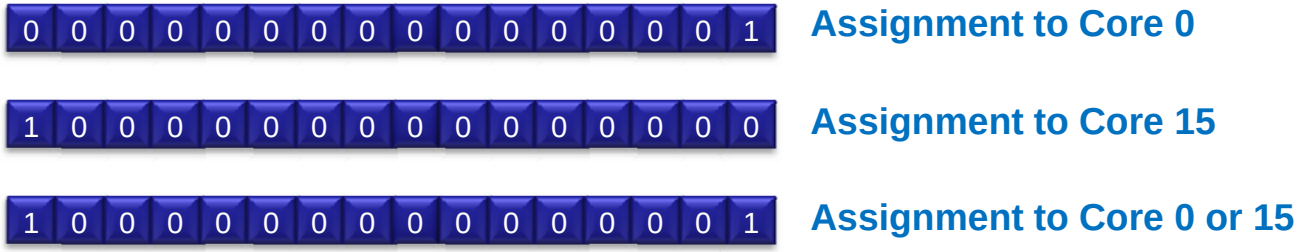

Code Example for Scheduling Affinity

```
...
#include <spawn.h>                  //C API parameters and prototypes
...
int icore=3;                        //Set core number
cpu_set_t cpu_mask;                 //Allocate mask
...
CPU_ZERO(     &cpu_mask);           //Set mask to zero
CPU_SET(icore,&cpu_mask);           //Set mask with core #

err = sched_setaffinity( (pid_t)0 ,          //Set the affinity
                         sizeof(cpu_mask),
                         &cpu_mask);
```

Conclusions and Future Prospects

- On NUMA systems like Ranger, placement and binding of processes and their associated memory are important performance considerations.
- Process Affinity and Memory Policy have a significant effect on pure MPI, pure OpenMP, and Hybrid codes.
- Simple numactl commands and APIs allow users to control process and memory assignments.
- Future prospects for hybrid programming:
  - 8-core and 16-core socket systems are on the way, so even more effort will be focused on process scheduling and data locality.
  - Expect to see more multi-threaded libraries; be alert for their potential interaction with your own multithreading strategy.

References

- Yun (Helen) He and Chris Ding, Lawrence Berkeley National Laboratory, June 24, 2004: Hybrid OpenMP and MPI Programming and Tuning (NUG2004).

  www.nersc.gov/nusers/services/training/classes/NUG/Jun04/NUG2004_yhe_hybrid.ppt

- Texas Advanced Computing Center: Ranger User Guide, see numa section. www.tacc.utexas.edu/services/userguides/ranger

- Message Passing Interface Forum: MPI-2: MPI and Threads (specific section of the MPI-2 report).

  http://www.mcs.anl.gov/research/projects/mpi/mpi-standard/mpi-report-2.0/node162.htm

- Intel Corp.: Thread Affinity Interface (Linux and Windows), from the Intel Fortran Compiler User and Reference Guides.

  http://www.intel.com/software/products/compilers/docs/fmac/doc_files/source/extfile/optaps_for/common/optaps_openmp_thread_affinity.htm