

# How to Use the Condo and CyEnce Clusters

Glenn R. Luecke

Director of HPC Education & Professor of Mathematics

April 11, 2018

## Online Information and Help

If you experience problems and would like help, send an email to [hpc-help@iastate.edu](mailto:hpc-help@iastate.edu) with detailed information about the problem. Online documentation for Condo and CyEnce can be found at:

- <http://www.hpc.iastate.edu/guides/condo>
- <http://www.hpc.iastate.edu/guides/cyence>

## Configuration of the Condo Cluster (installed spring 2015)

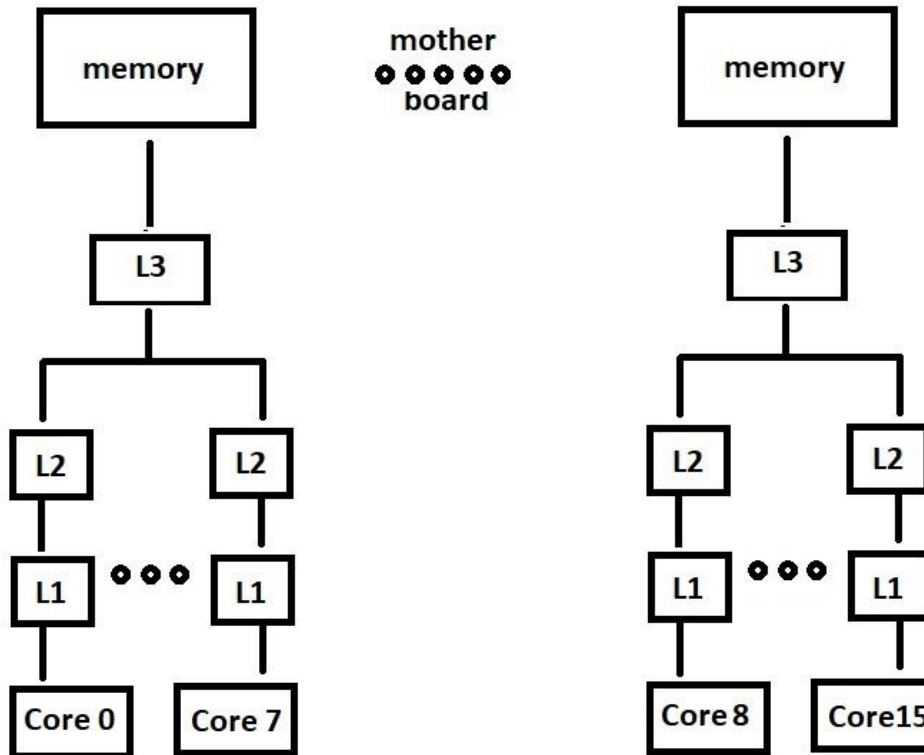
1. **188 compute nodes** each with two 8-core, 2.6 GHz Intel Haswell E5-2640 v3 processors (16 cores/node), **128 GB of memory** and 2.5 TB of available local disk (\$TMPDIR). The peak theoretical performance of this processor is  $144/8 = 18.0$  **Gflops/core** (with turbo boost). Turbo boost is ON for all nodes. Use the “cpuinfo” command to find additional information. The following gives the cache structure for each node in the cluster:
  - L1 32 KB one for each core (no sharing among cores)
  - L2 256 KB one for each core (no sharing among cores)
  - L3 20 MB one for each processor (shared among the 8 cores on the processor)
2. **One GPU node** with 2 Nvidia Tesla K20c GPUs.
3. **One “large” memory node** with four 8-core, 2.6 GHz Intel Ivy Bridge E5-4620 v2 processors (32 cores/node), 1 TB of memory and 1.8 TB of available local disk (\$TMPDIR). The peak theoretical performance of this processor is  $173/8 = 21.6$  **Gflops/core** (with turbo boost).
4. **One “huge” memory node** with four, 10-core, 2.2 GHz Intel Ivy Bridge E7-4830 v2 processors (40 cores/node), 2 TB of memory and 1.3 TB of available local disk (\$TMPDIR). The peak theoretical performance of this processor is  $77/10 = 7.7$  **Gflops/core** (with turbo boost).
5. All nodes and storage are connected via Intel/Qlogic’s QDR **40 Gb/s InfiniBand switch**.
6. 256 TB shared **temporary** disk space named /ptmp.
7. 756 TB of RAID-6 long-term **NFS** disk space under quota
8. **Operating System:** Red Hat Enterprise Linux version 7 with Intel’s Hyper-Threading turned OFF
9. **Slurm** for resource and job management.
10. **Head Node** (condo2017.its.iastate.edu).
11. **Data Transfer Node** (condodtn)

## Configuration of the CyEnce cluster (installed summer 2013 along with hpc-class)

1. **240 compute nodes** each with two 8-core, 2.0 GHz Intel Sandy Bridge E5-2650 processors (16 cores/node), **128 GB of memory** and 2.5 TB of available local disk (\$TMPDIR). The peak theoretical performance of this processor is  $179/8 = 22.4$  **Gflops/core** (with turbo boost). Turbo boost is on for all nodes. Use the “cpuinfo” command to find additional information. The cache structure for all nodes is the same as for the Condo cluster.
2. **One large memory (“fat”) node** with four, 8-core, 2.0GHz Intel Sandy Bridge processors (32 cores/node), 1 TB of memory and 2.5 TB of available local disk (\$TMPDIR).
3. **16 GPU nodes** configured as a compute node plus 2 NVIDIA K20c’s.
4. **16 Phi nodes** configured as a compute node plus 2 Intel Phi co-processors.
5. All nodes and storage are connected via Intel/Qlogic’s QDR **40 Gb/s InfiniBand switch**.

6. **288TB** shared **temporary** disk space named /ptmp.
7. **768TB** of RAID-6 long-term **NFS** disk space under quota
8. **Operating System:** Red Hat Enterprise Linux version 7 with Intel's Hyper-Threading turned OFF
9. **Slurm** for resource and job management.
10. **Head Node** (cyence), **Data Transfer Node** (cyencedtn)

### Compute Node



## Obtaining an account & logging onto the Condo cluster

Those who have purchased both nodes and storage can use the Condo cluster. Nodes and storage can be purchased by filling out the Condo Cluster Purchase Form at [https://hpc-ga.its.iastate.edu/condo\\_purchase/](https://hpc-ga.its.iastate.edu/condo_purchase/). To obtain an account send an email to [hpc-help@iastate.edu](mailto:hpc-help@iastate.edu). LAS, Mathematics and Statistics have purchased nodes and storage for their researchers. To obtain an account, contact

- For LAS send an email to [researchit@iastate.edu](mailto:researchit@iastate.edu)
- For mathematics send an email to Professor Hal Schenck, [hschenck@iastate.edu](mailto:hschenck@iastate.edu)
- For Statistics send an email to Michael Brekke [mbrekke@iastate.edu](mailto:mbrekke@iastate.edu)

For security **Google Authenticator** is required for logging on to Condo. Instructions on how to install and use Google Authenticator are on the Condo web page:

<http://www.hpc.iastate.edu/guides/condo/access-and-login>

The access to the Condo cluster is managed through ASW access lists, and PIs can add or delete users to/from the list for their group. Once users have been added to the access list, they should login

to the Condo cluster as described below. At the first attempt to login a Google Authenticator account will be created for them, and the email with the instructions will be sent to their email address at ISU. **To access Condo from off-campus machines, use the ISU VPN**, see <http://www.it.iastate.edu/howtos/vpn>.

To logon issue: `ssh username@condo2017.its.iastate.edu`, where “username@” can be omitted if your username on the machine from which you are issuing the ssh is the same as that on Condo. Enter ISU password when prompted for a password. If you need an X-session, change “ssh” to “ssh -X”. The vi (vim), iMacs and xemacs editors are available.

## Obtaining an account and logging onto the CyEnce cluster

PIs on the NSF MRI grant and their students and collaborators can use the CyEnce cluster. To obtain an account, PIs are to send an email to [hpc-help@iastate.edu](mailto:hpc-help@iastate.edu).

For security **Google Authenticator** is required for logging on to CyEnce. Instructions on how to install and use Google Authenticator are on the CyEnce web page: <http://www.hpc.iastate.edu/guides/cyence/login-and-file-transfer>

The access to the CyEnce cluster is managed through ASW access lists, and PIs can add or delete users to/from the list for their group. Once users have been added to the access list, they should login to the CyEnce cluster as described below. At the first attempt to login a Google Authenticator account will be created for them, and the email with the instructions will be sent to their email address at ISU. **To access CyEnce from off-campus machines, use the ISU VPN**, see <http://www.it.iastate.edu/howtos/vpn>.

To logon issue: `ssh username@cyence.its.iastate.edu`, where “username@” can be omitted if your username on the machine from which you are issuing the ssh is the same as that on CyEnce. Enter ISU password when prompted for a password. If you need an X-session, change “ssh” to “ssh -X” when logging onto discover and share. The vi (vim), emacs and xemacs editors are available.

## Creating the Work and /ptmp Directories

For both Condo and CyEnce, data is stored on two separate sets of storage servers: NFS storage for permanent data storage (/work) and parallel storage (/ptmp) for temporary data storage and for parallel I/O. Along with the home directories, the /work and /ptmp storage are accessible from all nodes.

Currently each user has 5 GB of home directory space, /home/<username>. In addition, each group has a directory shared by all members of the group. Files should be stored in the work directory and not in the home directory. The group working directories are located in /work on Condo and in /work1, /work2 on CyEnce. To find your group issue “groups” on the head node. To find the location of your work directory on CyEnce issue: **ls -d /work\*/<group>**

The cdw and cds utilities have been created to make the creation of the /work and Lustre/ptmp directories easy. Issue

**cdw** to create the /work directory, and issue  
**cds** to create your directory in /ptmp

Note that /ptmp is temporary storage, and files older than 90 days will be automatically deleted. Files you want to keep should be copied from your directory in /ptmp to your work directory. Please delete your /ptmp files when you have finished with them so that others can utilize this storage.

Normally, all jobs should be submitted either from the work directory or from /ptmp/<group>/<username>. Some applications may run faster from /ptmp, since files are striped across multiple volumes, though many factors may affect IO performance. To take full advantage of striping, files should have size of at least 1MB.

For the one-node IO intensive jobs we recommend using local scratch space on the compute nodes. The users will have to explicitly copy data from their group working directories to \$TMPDIR within the job script before executing commands, and then copy output data back to the group working directories at the end of the script. We have not tested whether or not MPI parallel IO can help IO performance.

## Using the Data Transfer Node (DTN)

On CyEnce the DTN is named “cyencedtn” and on Condo it is named “condodtn”. The DTNs have 10 Gbit/second connections to the campus network and to the internet allowing for fast transfer of data to on-campus and off-campus machines. The DTN should be used to transfer data to and from a cluster. The DTN can also be used to compile large programs when compilation on the head node fails with a segmentation fault.

To use the DTN, ssh from the head node to the DTN. All files and directories on the DTN are identical to those on the head node. There are several ways of transferring data to and from the DTN. If you have files on MyFiles (<https://www.it.iastate.edu/services/storage/myfiles>), you can copy files to/from MyFiles and the DTN. When logged onto a DTN, issue 'kinit' and enter your university password to get access to MyFiles. After entering your password, you will be able to cd to /myfiles/Users/\$USER and to /myfiles/<deptshare>, where <deptshare> is your departmental share name; this is usually your department or college’s shortname, such as engr or las. Since files are not backed up, users are encouraged to use MyFiles to back up critical files. To copy all files in your home directory on condo, issue:

```
rsync -pvr /home/$USER/ /myfiles/Users/$USER/condo
```

Backed-up files can be viewed by going to /myfiles/Users/\$USER/condo.

You can also use scp and rsync to transfer files between computers. If you want to copy many files and/or directories and/or large files, rsync is recommended over scp. An advantage of using rsync instead of scp is that rsync will only copy files if they have been changed. The following are typical examples using scp and rsync. One can replace “scp” or “scp -r” with “rsync -ae ssh” in the following examples.

**Example 1:** To copy ~username/mydir/prog.c from hpc-class to ./mydir/prog.c in your current directory on the DTN, issue the following on the DTN:

```
scp username@hpc-class.its.iastate.edu:mydir/prog.c ./mydir/prog.c OR  
rsync -ae ssh username@hpc-class.its.iastate.edu:mydir/prog.c ./mydir/prog.c
```

**Example 2:** To copy the entire “mydir” directory in example 1 to the current directory on the DTN, issue the following on the DTN: (The “/” is not needed, but without it the “.” is easy to miss.)

```
scp -r username@hpc-class.its.iastate.edu:mydir ./ OR
rsync -ae ssh username@hpc-class.its.iastate.edu:mydir ./
```

**Example 3:** To copy all the files in “mydir” that end in .c from hpc-class to DTN, issue the following on the DTN:

```
scp username@hpc-class.its.iastate.edu:mydir/*.c ./newdir OR
rsync -ae ssh username@hpc-class.its.iastate.edu:mydir/*.c ./newdir
```

**Example 4:** To copy prog1.f90 in your current directory on the DTN to prog2.f90 in your home directory on hpc-class issue the following on the DTN:

```
scp prog1.f90 username@hpc-class.its.iastate.edu: prog2.f90 OR
rsync -ae ssh prog1.f90 username@hpc-class.its.iastate.edu: prog2.f90
```

If you're using Windows computer, you probably don't have scp available. In this case you can download free SFTP/SCP client WinSCP.

**Globus Online** is the recommended method for transferring data since once it has been installed, it is easy to use and fast. Globus Online is available on CyEnce and Condo. For information on how to use Globus Online, see the video at <http://www.hpc.iastate.edu/guides/globus-online> . Notice that you will need to install the Globus Online client on your laptop/workstation. If it is an ISU machine, you will need to have the administrator of the machine authorize the installation of the Globus Online client.

## Modules for Software Management

The web site for each machine contains a listing of the software available. Groups are responsible for paying for their own specific software. Software installed by HPC Support is managed by using environment modules. The following are commonly-used module commands:

```
module spider - displays all available modules
module spider intel - lists all available versions of the intel software
module spider gcc - lists all available versions of the GNU software
module load <module_name> - loads specified module
module unload <module_name> - unloads specified module
module list - lists all loaded modules
module purge - unloads all loaded modules
```

It is a good programming practice to remove all modules and then add the modules you need:

```
module purge
module load intel - if the Intel compilers and MPI will be used
module load intel/17.4 - if version 17.4 is desired
module load gcc - if the GNU compilers will be used
module load gcc/6.3.0 - specifies specific version to be loaded
module load openmpi - if OpenMPI will be used
module load allinea - if the Allinea tool software will be used
```

## Compilers

The Intel (recommended), GNU and PGI compilers are available: Intel Fortran (ifort), Intel C/C++ (icc), GNU fortran (gfortran), GNU C/C++ (gcc), PGI (pgfortran) and PGI C/C++ (pgCC). Compiler options can be found by issuing

```
man <compiler_name>
```

Free format Fortran files must be named something.f90, C files must be named something.c and C++ files must be named something.cc. When one successfully compiles a program, the executable will be placed in the file named “a.out” in your current directory. Use the **-o** option if you want the executable to go to a different file, e.g ifort **-o prog** program.f90 will produce the executable “prog” when the compile is successful.

Compiles are usually done on the head node. To prevent users from running computations on the head node, memory limits on the head node are set lower than on the compute nodes. This may result in segmentation fault when compiling large programs. In this case you can compile your program on the DTN.

## Compiling for Debugging for Intel Compilers

The **-g** option produces symbolic debug information that allows one to use a debugger. The **-debug** option turns on many of the debugging capabilities of the compiler. The **-qopenmp** option (-fopenmp for GNU compilers) allows the recognition of OpenMP directives/pragmas and should not be used if it is not needed. The debuggers available are Allinea’s ddt, Intel’s Inspector XE (inspxe-cl and inspxe-gui), GNU’s gdb and PGI’s pghpf. The recommended debugger is ddt. The following shows how to perform compiling for debugging with automatic linking of Intel’s MPI libraries and producing an executable named “prog”. Note, if you do not need the MPI library, change “mpiifort” to “ifort”, “mpiicc” to “icc” and “mpiicpc” to “icpc”. Intel’s MPI Checker provides special checking for MPI programs. To use **MPI Checker** add the **-check-mpi** to the mpirun command.

```
mpiifort -g -debug -check all -traceback -o prog prog.f90
mpiicc -g -debug -check-pointers=rw -check-uninit -traceback -o prog prog.c
mpiicpc -g -debug -check-pointers=rw -check-uninit -traceback -o prog prog.cc
```

### Remarks:

- The **-check all** option turns on all run-time error checking for Fortran. (There appears to be no equivalent option for Intel’s C/C++ compiler.)
- The **-check-pointers=rw** option enables checking of all indirect accesses through pointers and all array accesses for Intel’s C/C++ compiler but not for Fortran.
- The **-check-uninit** option enables uninitialized variable checking for Intel’s C/C++ compiler. This functionality is included in Intel’s Fortran **-check all** option.
- The **-traceback** option causes the compiler to generate extra information in the object file to provide source code traceback information when a severe error occurs at run-time.

## Compiling for High Performance

The Intel compiler will generate code for the host processor being used with the **-xHost** or equivalently **-march=native** option. Intel recommends the **-O2** compiler option as the optimization option that will usually give the best performance and this option is on by default. Intel also provides the **-O3**, **-ipo**, **-fast** and **-Ofast** options that turn on many optimizations, but they sometimes slow

down program execution. The **-ipo** option turns on interprocedural analysis between procedures. For best performance, experiment with the **-O2**, **-O3**, **-ipo**, **-Ofast** and **-fast** compiler options and find out which gives the best performance for your application. Choosing good compiler options can significantly reduce execution time. If the application uses both MPI and OpenMP, then one must add the **-qopenmp** option to the following:

```
mpiifort -xHost -o prog prog.f90
mpiicc   -xHost -o prog prog.c
mpiicpc  -xHost -o prog prog.cc
```

On the Condo cluster the fat and huge nodes have older processors than the head node's processors. Thus compile with the **-xHost** compiler option should be done on the fat/huge node when compiling programs that will run on those nodes. For the fat/huge nodes compile commands should be added to the job script so that the compiles will be performed on the fat/huge nodes.

Note: to compile MPI code with the GNU compilers use **-O2** or **-O3** optimization:

```
mpif90 -O2 -o prog prog.f90
mpicc   -O2 -o prog prog.c
mpicxx  -O2 -o prog prog.cc
```

For information about GNU compiler optimizations, see <https://gcc.gnu.org/onlinedocs/gcc/Optimize-Options.html>

## Intel's Optimization Report

Intel's compiler option **-qopt-report** will generate a report that gives the optimizations performed by the compiler. If the program compiled is prog.f90 then the optimization report will be in file prog.optrpt.

## Intel's Message Checker

Intel's Message Checker tool checks for MPI errors at run-time. To use Message Checker, compile with the **-g** option and then add the **-check-mpi** option on the **mpirun** command.

## How to Run MPI and OpenMP Jobs

On HPC clusters computations should be performed on the compute nodes and not on the head node. Special programs called resource managers, workload managers or job schedulers are used to allocate compute nodes to users' jobs. The Slurm Workload Manager is used for this purpose.

Condo, CyEnce and the student cluster use the Slurm job scheduler. Jobs are submitted using the sbatch command. For example, if the script file is named "prog.script", then the job is submitted for execution by issuing

```
sbatch          prog.script - for regular compute nodes
sbatch -p fat   prog.script - for the 1 TB node (-p can also be used within the script)
sbatch -p huge prog.script - for the 2 TB node (-p can also be used within the script)
```

The program will be compiled on the first compute node assigned to this job.

The compute nodes for Condo and CyEnce have 128 GB of memory and 2 processors with each processor having 8 cores. Thus, each node has 16 cores and 128 GB of memory (128/16 = 8 GB per core). For help creating scripts on CyEnce and Condo (respectively) see:

[http://hpcgroup.public.iastate.edu/HPC/CyEnce/CyEnce\\_script\\_writer.html](http://hpcgroup.public.iastate.edu/HPC/CyEnce/CyEnce_script_writer.html)  
<http://www.hpc.iastate.edu/guides/condo/slurm-job-script-generator-for-condo>

One can also translate existing PBS scripts to Slurm using the `pbs2sbatch` command.

It is sometimes useful to interactively debug a program by obtaining nodes that are dedicated to running your job. This can be done using the `salloc` command. For example, to obtain one dedicated node for 1 hour issue:

```
salloc -N 1 -n 16 -t 1:00:00
```

To obtain 2 dedicated nodes for an hour issue: **`salloc -N 2 -n 32 -t 1:00:00`**

What follows below provides background information to help you run your MPI and MPI/OpenMP programs using scripts on Condo and CyEnce.

## The “`mpirun -perhost n`” Option

The *-perhost n* option on the `mpirun` command means that *n* MPI processes are assigned per node, cycling through nodes in a round-robin fashion. If one wants 16 consecutive MPI processes assigned to each node, one should use *-perhost 16* (the default value is 16 for CyEnce, Condo and the student cluster). If one wants 1 MPI process for each of the two processors/sockets on a node, then use *-perhost 2*. In order for this option to work, first you need to execute the following command:

```
export setenv I_MPI_JOB_RESPECT_PROCESS_PLACEMENT=disable
```

## Pinning MPI Processes to Cores

During program execution, the operating system may move MPI processes to different cores on the same node. Remember that each processor/socket on a node has its own memory and that memory accesses from one socket to the memory of the other socket takes longer than memory accesses to the socket's memory. Memory is allocated using “first touch” and memory allocations remain in the same physical memory throughout the execution of the MPI job. By turning on MPI process pinning, MPI processes cannot be moved to other cores during program execution. MPI process pinning is on by default. If Intel changes this default, memory pinning can be enabled by using the *-genv* option with the `mpirun` command. The following are sometimes useful when using the *-genv* option with the `mpirun` command.

- **`-genv I_MPI_PIN on`** Turns on processor pinning. It is on by default.
- **`-genv I_MPI_PIN_PROCESSOR_LIST 0-15`** Pins MPI process 0 to core 0, process 1 to core 1, . . . , process 15 to core 15. This is the default assignment.
- **`-genv I_MPI_PIN_PROCESSOR_LIST 0,8`** Pins MPI processes to cores 0 and 8. This is useful when running programs using both MPI and OpenMP.



- **-genv I\_MPI\_PIN\_DOMAIN socket** Pins MPI processes to processors/sockets and is useful for MPI/OpenMP programs since this option insures that the execution of OpenMP threads cannot be moved to different sockets.
- **-genv I\_MPI\_DEBUG 4** - prints process pinning information.

Adding “**command time -v**” to mpirun when in the bash, dash and ksh shells provides an easy way to time your program and to estimate its parallel efficiency by looking at the “Elapsed (wall clock) time (h:mm:ss or m:ss): 0:01.73”. Replacing “mpirun” with “command time -v mpirun” in the following example scripts will cause the timing information to be placed in in the prog.errors file.

## Unlimiting the Stacksize

If one runs out of stack memory, then a segmentation fault error will be produced and the program will not run. It is difficult to debug such an error. It is therefore recommended to always remove this stacksize limit by issuing

```
ulimit -s unlimited
```

## Script Writing

As mentioned above, jobs are submitted using the Slurm sbatch command by issuing

```
sbatch prog.script
```

where prog.script is the script file for the program named prog.f90 (or C/C++) equivalent. The first 6 lines of a script should look like the following:

```
#!/bin/bash
#SBATCH -o ./prog.out    - where the program output goes
#SBATCH -e ./prog.errors - where the program errors go
#SBATCH -N 4            - number of nodes for the job
#SBATCH -n 64          - number of cores for the job
#SBATCH -t 3:20:00     - maximum job time is 3 hours, 20 min. 0 seconds
```

Suppose one wants to run a MPI program, prog.f90, with 64 MPI processes using all 64 cores on 4 nodes with a 3 hour and 20 minutes time limit with the output written to ./prog.out and error messages written to ./prog.errors. Setting N to 4 means 4 nodes are reserved for your job. Since pinning MPI processes to cores is the default, the **-genv I\_MPI\_PIN on** option need not be specified. (I suggest you save this script as prog.script64 so you know that the script is for prog.f90 with 64 MPI processes.) Script 1 contains the complete script.

## Script 1: 4 nodes with 64 MPI processes/node

```

#!/bin/bash
#SBATCH -o ./prog.out
#SBATCH -e ./prog.errors
#SBATCH -N 4
#SBATCH -n 64
#SBATCH -t 3:20:00
module purge
module load intel
module load allinea
ulimit -s unlimited
mpiifort -xHost -o prog prog.f90
export setenv I_MPI_JOB_RESPECT_PROCESS_PLACEMENT=disable
mpirun -perhost 16 -np 64 ./prog

```

## Script 2: 8 nodes with 8 MPI processes/node

If the above does not provide enough memory per MPI process, then the **perhost 8** option on the mpirun command can be used so that there will be 8 (instead of 16) MPI processes per node. To keep the same number of MPI processes, the script specifies 8 instead of 4 nodes and 128 cores instead of 64 cores. In the following script MPI processes are pinned to the first 4 cores on each of the two processors/sockets on each node.

```

#!/bin/bash
#SBATCH -o ./prog.out
#SBATCH -e ./prog.errors
#SBATCH -N 8
#SBATCH -n 128
#SBATCH -t 3:20:00
module purge
module load intel
module load allinea
ulimit -s unlimited
mpiifort -xHost -o prog prog.f90
export setenv I_MPI_JOB_RESPECT_PROCESS_PLACEMENT=disable
mpirun -perhost 8 -genv I_MPI_PIN_PROCESSOR_LIST 0-3,8-11 -np 64 ./prog

```

## Running MPI & OpenMP Programs

OpenMP is used for parallelization of shared memory computers and MPI for parallelization of distributed (and shared) memory computers. Since memory is shared on a node, one can parallelize with OpenMP within nodes and MPI between nodes. One could specify 1 MPI process per node and use 16 OpenMP threads to parallelize within a node. However, since there are two processors/sockets per node and since each processor has memory physically close to it, it is generally recommended to use 2 MPI processes per node and use 8 OpenMP threads for the 8 cores on each processor/socket. One can set the number of threads by either setting the OMP\_NUM\_THREADS environment variable or by calling omp\_set\_num\_threads in the program. OpenMP parallelization requires the insertion of directives/pragmas into a program and then compiled using the **-qopenmp** option (-fopenmp on GNU compilers). To configure script 3 to use the same number of cores as scripts 1 means that we want to use all 64 cores in 4 nodes and 2 MPI processes per node for a total of 8 MPI processes.

To keep OpenMP threads active, set `OMP_WAIT_POLICY=ACTIVE`. To keep threads executing on the same socket set `OMP_PLACES=sockets` and `OMP_PROC_BIND=close`. There is a default OpenMP stacksize and jobs that require more memory will produce a segmentation fault error message and will not run. The OpenMP stacksize can be increased using the `OMP_STACKSIZE` environment variable. For the student cluster, I recommend setting the OpenMP `stacksize=1G` and `=2G` for the Condo and CyEnce clusters.

### Script 3: 4 nodes with 2 MPI processes/node + OpenMP

```
#!/bin/bash
#SBATCH -o ./prog.out
#SBATCH -e ./prog.errors
#SBATCH -N 4
#SBATCH -n 64
#SBATCH -t 3:20:00
module purge
module load intel
module load allinea
ulimit -s unlimited
export setenv I_MPI_JOB_RESPECT_PROCESS_PLACEMENT=disable
export OMP_WAIT_POLICY=ACTIVE
export OMP_PROC_BIND=close
export OMP_PLACES=sockets
export OMP_STACKSIZE=2G
mpifort -xHost -qopenmp -o prog prog.f90
mpirun -perhost 2 -genv I_MPI_PIN_DOMAIN socket -np 8 ./prog
```

### Script 3: Remarks

If one wants to pin threads to cores and not allow threads to move between cores, then replace

```
export OMP_PROC_BIND=close
export OMP_PLACES=sockets
```

with

```
export OMP_PROC_BIND=true
```

Calling MPI routines outside of OpenMP parallel regions does not require the replacement of `mpi_init` with `mpi_init_thread`. However, when placing MPI calls inside parallel regions, then one must replace the call to `mpi_init(ierror)` with `mpi_init_thread(required, provided, ierror)` where *required* is the desired level of thread support and *provided* is the thread support provided. The following are the values *required* can have:

- `MPI_THREAD_SINGLE`: only one thread will execute.
- `MPI_THREAD_FUNNELED`: The process may be multi-threaded, but only the main thread will make MPI calls (all MPI calls are funneled to the main thread).

- `MPI_THREAD_SERIALIZED`: The process may be multi-threaded, and multiple threads may make MPI calls, but only one at a time: MPI calls are not made concurrently from two distinct threads (all MPI calls are serialized).
- `MPI_THREAD_MULTIPLE`: Multiple threads may call MPI, with no restrictions.

For best performance, use the least general option possible. Notice that one can overlap communication and computation by having one thread execute MPI routines and the other threads are free to do computations. There is also a possibility of reducing memory requirements over a pure MPI implementation since processor cores share a common memory.

## Using the Large Memory Nodes

On CyEnce, there is one large memory node, called “fat” that has 1 TB of memory and 4 processor/sockets (32 total cores). Condo has two large memory nodes: one is called “fat” and has 1 TB of memory and 32 cores and the other is called “huge” and has 2 TB of memory and 40 cores. The recommended way to use the large memory nodes is to use 32 or 40 OpenMP threads by setting `omp_set_num_threads(32 or 40)` in the program. One must use ***sbatch -p fat*** to submit jobs for the 1 TB memory node and ***sbatch -p huge*** for the 2 TB node. Notice the compiler command inside the script so that the program will be compiled on the large memory node where it will execute. This is important since these nodes use different processors from the regular compute nodes.

## Script 4: OpenMP Using a Large Memory Node

```
#!/bin/bash
#SBATCH -o ./prog.out
#SBATCH -e ./prog.errors
#SBATCH -N 1
#SBATCH -n 32 (or 40)
#SBATCH -t 3:20:00
module purge
module load intel
module load allinea
ulimit -s unlimited
export OMP_WAIT_POLICY=ACTIVE
export OMP_PROC_BIND=close
export OMP_PLACES=sockets
export OMP_STACKSIZE=2G
ifort -xHost -qopenmp -o prog prog.f90
./prog
```

## Job Queue Commands

1. **sbatch prog.script** - submits the script in file prog.script
2. **squeue** - status of all jobs executing and queued
3. **squeue -u <your\_userid>** status of only your jobs – useful when there are many jobs
4. **sinfo** - status of all queues and the current queue structure
5. **scancel #** - deletes the job with number #. One can only delete one’s jobs.

## How many nodes should one use?

The answer depends on the problem size and how well your application has been parallelized. Often applications will require a minimum number of nodes due to large memory requirements. Once this minimum number of nodes has been established, one must decide how many nodes to use for running the application. For example, let's take an MPI parallelization of the Jacobi iteration with  $N = 4 \cdot 1024$  and  $N = 64 \cdot 1024$  using differing numbers of nodes. Both of these problem sizes can be run on a single node, so the question is how many nodes should one use. The following numbers were obtained running on ISU CyEnce cluster in the fall of 2013.

For  $N = 4 \cdot 1024$ , it is best to use 1 node for this program if one wants to make best use of the allocation and use 8 nodes to minimize the execution time, see Table 1. For  $N = 64 \cdot 1024$ , using 64 nodes gives the shortest time and the cost for the allocation is not much different from the best value using only 1 node, see Table 2.

| Number of Nodes | Seconds for $N = 4 \cdot 1024$ | Node-Seconds for $N = 4 \cdot 1024$ |
|-----------------|--------------------------------|-------------------------------------|
| 1               | 3.3                            | 3.3                                 |
| 2               | 2.3                            | 4.6                                 |
| 4               | 1.3                            | 5.2                                 |
| 8               | 0.8                            | 6.4                                 |
| 16              | 1.8                            | 28.8                                |

Table 1: Jacobi iteration with  $N = 4 \cdot 1024$  using different numbers of nodes.

| Number of Nodes | Seconds for $N = 64 \cdot 1024$ | Node-Seconds for $N = 64 \cdot 1024$ |
|-----------------|---------------------------------|--------------------------------------|
| 1               | 875.6                           | 875.6                                |
| 2               | 442.4                           | 884.8                                |
| 4               | 224.7                           | 898.8                                |
| 8               | 113.8                           | 910.4                                |
| 16              | 59.4                            | 950.4                                |
| 32              | 32.6                            | 1,043.2                              |
| 64              | 17.2                            | 1,100.8                              |

Table 2: Jacobi iteration with  $N = 64 \cdot 1024$  using different numbers of nodes.

## Using Intel's Math Kernel Library

Intel provides highly optimized scientific library routines for Fortran and some C versions in their Math Kernel Library (MKL). You should call these routines from your program whenever possible to increase the performance of your application. For documentation see <http://software.intel.com/en-us/articles/intel-math-kernel-library-documentation>. Intel provides a tool to help one select the correct link `-L` and `-l` options to add to a Fortran or C/C++ compile/link command and to correctly set environmental variables. This tool is available at: <http://software.intel.com/en-us/articles/intel-mkl-link-line-advisor>. To link the MKL serial optimized libraries compile your program with `ifort -mkl=sequential`.

## Cluster Software

A listing of the software on each cluster can be found on the cluster web site:

- <http://www.hpc.iastate.edu/guides/condo>
- <http://www.hpc.iastate.edu/guides/cyence>

Groups are responsible for installing and paying for their own application software. There is a module facility available for software management of installed software.

## Accelerator Nodes

The student cluster, CyEnce and Condo clusters have accelerator nodes. Information on how to use the accelerator nodes can be found on the web sites of these clusters.

## The Student Cluster (hpc-class)

The purpose of the student cluster is to support HPC instruction at ISU. Compute nodes are identical to the compute nodes on Condo and CyEnce but with half the memory. There are 4 GPU nodes each with a single K20. There are 4 Phi nodes each with a single Phi co-processor. For information see: <http://www.hpc.iastate.edu/education/>

## ISU HPC Courses

Parallel computing techniques taught in the courses listed below will allow students to take discipline specific courses that teach parallel algorithms relevant to the discipline.

**Purpose:** The purpose of math 424 and 525 is to train students to be able to effectively use on-campus High Performance Computing (HPC) machines. Students without Fortran or C programming experience should take 424 and then 525. Students with Fortran or C programming experience should take only 525.

### 1. Introduction to High Performance Computing: Math/ComS/CprE 424, 3 credit. Offered fall semesters.

**Description:** This course is designed for students who have no Unix experience and no experience programming scientific applications in Fortran or C. Course Topics: Unix, serial programming of scientific applications for high performance (e.g. stride 1 array accesses, vectorization) and OpenMP for high performance shared memory parallelization. An oral and written semester project is required of each student. The semester project gives students the opportunity to apply concepts learned to their research interests. Professor Glenn Luecke will be the instructor fall 2018.

**Prerequisites:** Basic knowledge of linear algebra or permission of the instructor.

**Meeting times and room:** Tuesdays and Thursdays from 2:10 to 4:00 in 449 Carver. 2 hours of lecture and 2 hours of lab each week.

### 2. High Performance Computing: Math/ComS/CprE 525, 3 credit. Offered spring semesters.

**Description:** High Performance Computing (HPC) concepts and terminology, HPC machine architecture, how to use debugging and performance tools, advanced parallel programming with MPI, OpenMP and possibly OpenACC and/or CUDA. An oral and written semester project is required of

each student. The semester project gives students the opportunity to apply concepts learned to their research interests. Professor Luecke will be the instructor spring 2019.

**Prerequisites:** math 424, experience programming in Fortran or C, or permission of the instructor.

**Meeting times and room:** Tuesdays and Thursdays from 11:00 to 1:00 in 449 Carver. 2 hours of lecture and 2 hours of lab each week.