MPI on Cirrus

Reusing this material

Access

- Cirrus: **`ssh -XY user@cirrus-msc.epcc.ed.ac.uk`**
  - you must use this dedicated MSc login node

- You can access systems using ssh from anywhere
  - Trivial for Linux
  - Mac: enable the X server (xquartz) to display any graphics
  - Windows: need to install an X server program, e.g. MobaXterm

Useful files and templates

- Take a copy of **`MPP-templates.tar`**
  - see the course web pages
- unpack: **`tar -xvf MPP-templates.tar`**
- Crib sheets for MPI programs available on course web pages

Setting up Cirrus environment

- Load the Message-Passing Toolkit
  - module load mpt
- Load the Intel Compilers
  - module load intel-compilers-17
- To automate, add these lines to your ".bash_profile" file

```
[user@cirrus] gedit ~/.bash_profile
```

Compiling MPI Programs on Cirrus

- C programmers use: **`mpicc -cc=icc`**
- C++ programmers use: **`mpicxx -cxx=icpc`**
- Fortran programmers use: **`mpif90`**

- There is nothing magic about these MPI compilers!
  - simply wrappers which automatically include various libraries etc
  - compilation done by standard (e.g. Intel) compilers
    - icc, icpc and ifort

- You can use the supplied Makefiles for convenience
  - **`make -f Makefile_c`**
  - **`make -f Makefile_cc`**
  - **`make -f Makefile_f90`**

- Easiest to make a copy of one of these called “Makefile”
  - also need to change the line “MF=“ in the Makefile itself

Running interactively on Cirrus

- Timings will not be reliable
  - shared with other users, many more processes than processors
  - but **very useful** during development and for debugging

- **`mpirun –n 4 ./mpiprog.exe`**
  - runs your code on 4 processes

- NOTE
  - output might be buffered
  - if your program crashes, you may see no output at all

- May need to explicitly flush prints to screen
  - **`FLUSH(6)`**
  - **`fflush(stdout);`**

Running batch jobs on Cirrus

- Run via a batch system
  - Cirrus uses Portable Batch System (PBS); submit script that launches your program

- In **`MPP-templates/`** is a standard batch script: **`cirrusmpi.pbs`**
  - make a copy of this file with a name that matches your executable, e.g.
  - **`user@cirrus$ cp cirrusmpi.pbs hello.pbs`**

- To run on 4 processors: **`qsub hello.pbs`**
  - use reserved queue during lab sessions, e.g. qsub –q R12345 hello.pbs
  - automatically runs executable called “hello”
  - output will appear in a file called **`hello.pbs.oXXXXX`**
  - can follow job progress using **`qstat`** or **`qstat –u $USER`**
  - script also times your program using the Unix “time” command
  - full instructions included as comments in the template
  - no need to alter the script - just rename it as appropriate
    - e.g. to run a program “pingpong” make another copy called “pingpong.pbs”

- Charging code for this course: `tc004`

Cirrus idiosyncrasies

- By default, MPI wrappers are not in your path

```
user@cirrus$ mpicc
-bash: mpicc: command not found
```

- To access correct version: **`module load mpt`**
  - defaults to GNU compilers: gcc, g++ and gfortran
  - in batch system, job launcher is called **`mpiexec_mpt`**
- Intel compilers: **`module load intel-compilers-17`**
  - add these to end of your **`.bash_profile`** file in home directory
  - to check you have the right version (similarly for mpif90)

```
user@cirrus$ which mpicc
/opt/hpe/hpc/mpt/mpt-2.16/bin/mpicc
```

  - **`mpif90`** automatically picks up the Intel Fortran compiler
  - to use Intel C [C++] compilers: **`mpicc -cc=icc [-cc=icpc]`**

C++ Interface

- MPI is not an OO interface
  - however, can be called from C++
- Originally had different function calls, e.g.
  - `MPI::Intracomm comm;`
  - `...`
  - `MPI::Init();`
  - `comm = MPI::COMM_WORLD;`
  - `rank = comm.Get_rank();`
  - `size = comm.Get_size();`
- Compiler is called **`mpicxx`**
  - see **`hello.cc`** and **`Makefile_cc`**

C++ interface is now removed

Must therefore cross-call to C

Documentation

- MPI Standard available online
  - See: http://www.mpi-forum.org/docs/
  - currently version 3.1

- Available in printed form
  - http://www.hlrs.de/mpi/mpi31/

- Man pages available on Cirrus and ARCHER
  - must use the C style of naming: **`man MPI_Routine_name`**, e.g.:
  - **`user@computer$ man MPI_Init`**

MPI Books

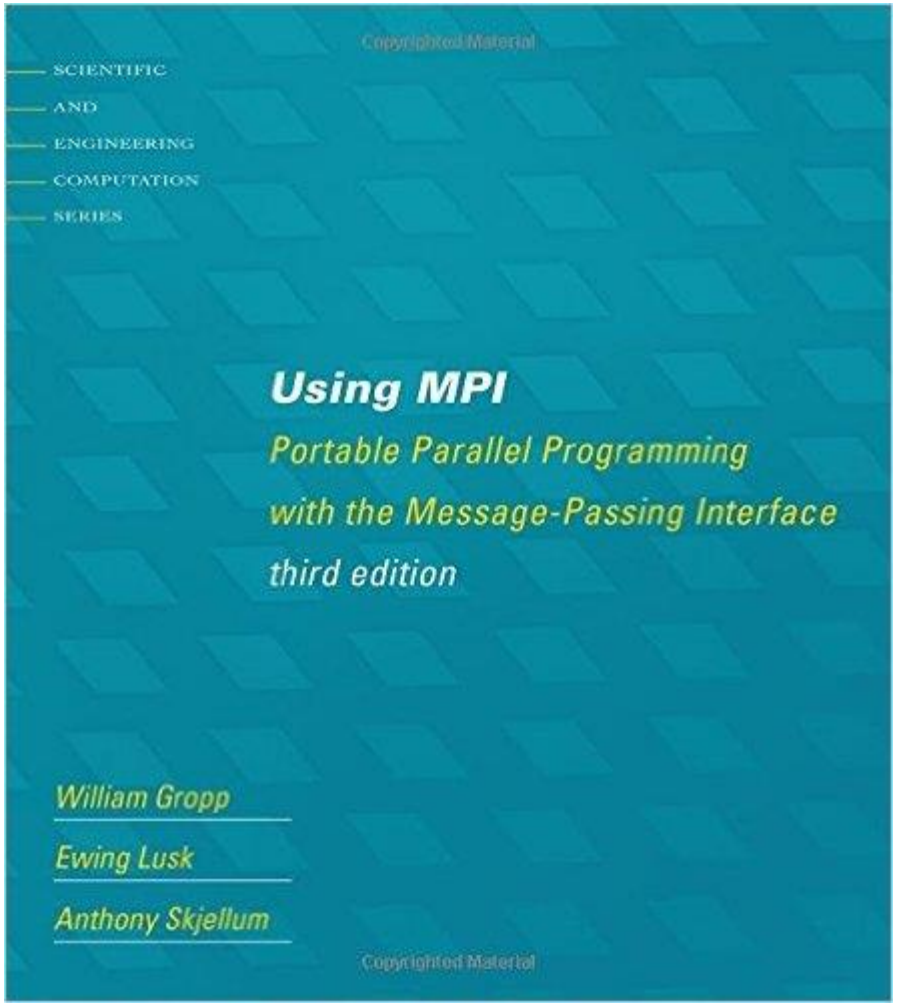
Copyrighted Material
SCIENTIFIC
AND
ENGINEERING
COMPUTATION
SERIES
Using MPI
Portable Parallel Programming
with the Message-Passing Interface
third edition
William Gropp
Ewing Lusk
Anthony Skjellum
Copyrighted Material

Exercise: Hello World

The minimal MPI program

- See Exercise 1 on the exercise sheet
- Write an MPI program that prints a message to the screen
- Main purpose is to get you compiling and running parallel programs on ness
  - also illustrates the SPMD model and use of basic MPI calls
- We supply some very basic template code
  - you need to add appropriate calls to compute rank and size