

MPI: Practical session 1

Note: These practicals are based on the “Further Work” section at the end of Dr Rutter’s slides. Extra (or at least different) detail may be found there.

1 My First MPI Program

Open your favourite text-editor, and type in the following:

```
#include <stdio.h>
#include <mpi.h>

int main(int argc, char *argv[]){
    int rank,nproc;

    MPI_Init (&argc,&argv);

    MPI_Comm_size (MPI_COMM_WORLD,&nproc);
    MPI_Comm_rank (MPI_COMM_WORLD,&rank);

    std::cout << "Hello C++ parallel world, I am process "
               << rank << " out of " << nproc << std::endl;

    MPI_Finalize (); return 0;
}
```

Save the file as `HelloWorld.C` and then, at the command-line, type the following:

```
mpic++ HelloWorld.C -o HelloWorld
./HelloWorld
```

Now type the following:

```
mpirun ./HelloWorld
```

In order to run on different numbers of processors, use the `-np` option, e.g.:

```
mpirun -np 8 ./HelloWorld
```

Try different numbers of processors, up to and even beyond the number of physical cores on your chosen computer.

2 Using different MPI implementations

There are multiple implementations of the MPI standard. One or more may be installed on whatever computer you are using and they must be used consistently. For example, `apollo` has two implementations installed. The default version is `OpenMPI`, but `MPICH` is also available:

```
mpicxx.mpich HelloWorld.C -o HelloWorld
./HelloWorld
```

Now type the following:

```
mpirun.mpich ./HelloWorld
```

In order to run on different numbers of processors, use the `-np` option, e.g.:

```
mpirun.mpich -np 8 ./HelloWorld
```

Do you notice any differences between MPICH and OpenMPI? The differences you can see are purely cosmetic. There are more subtle differences in how the two implementations are written, and hence their respective performance for particular cases. There may also be (usually) non-obvious bugs that exist in one but not the other.

3 Mixing different MPI implementations

Since MPI implementations are apparently interchangeable, try the following:

```
mpicxx.openmpi HelloWorld.C -o HelloWorld
mpirun.mpich -np 8 ./HelloWorld
```

What does this tell you? Try the same thing with `openmpi` and `mpich` reversed.

When you come to use computers with many more implementations of MPI installed, you will need to be very careful that your compilation commands and execution commands are consistent. Even different versions of OpenMPI may not have compatible `mpiruns`.

4 MPI Wrappers

As you will have realised, MPI is essentially a header file and some pre-compiled library functions (rather than a compiler feature). If you type:

```
mpicxx -show HelloWorld.C -o HelloWorld
```

then you will see that `mpicxx` actually just calls a compiler with particular options which point it to the header and library files.

5 Parametric studies

OpenMPI sets environment variables when it runs, which can be used outside MPI programs. For example, try:

```
mpirun -np 4 /bin/bash -c 'echo I am rank $OMPI_COMM_WORLD_RANK of $OMPI_COMM_WORLD_SIZE'
```

Note the single-quotes, and the fact that we have to spawn a new shell.

Consider how this might be used to run an embarrassingly parallel parametric study, where the individual simulations are serial.

The environment variables are implementation specific; in order to find those for MPICH, try:

```
mpirun.mpich -np 1 env
```