Exercise Introduction

Contents

Task 1: Hello World with Slurm and MPI (10 min) **1**

Task 1: Hello World with Slurm and MPI (10 min)

One of the most common ways of using a supercomputer is to allocate resources for an application via a HPC-scheduler and to parallelize the application using MPI. The GWDG utilizes Slurm as its HPC-scheduler.
The Message Passing Interface (MPI) is a standard for inter process communication on a single or distributed system. It provides efficient methods for running a single program over a number of systems. However, the MPI methods are relatively low level such that the programmer is responsible for implementing communication patterns and determining how and when processes will exchange data. An in-depth tutorial to MPI is not part of this course but it is still meaningful to understand the basics.
In this exercise you will compile a C program for usage with MPI and run it on multiple nodes using Slurm.

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

int main(int argc, char** argv) {
  MPI_Init(NULL, NULL);

  // Get number of processes
  int world_size;
  MPI_Comm_size(MPI_COMM_WORLD, &world_size);

  // Get rank of process
  int world_rank;
  MPI_Comm_rank(MPI_COMM_WORLD, &world_rank);

  // Get name of processor
  char processor_name[MPI_MAX_PROCESSOR_NAME];
  int name_len;
  MPI_Get_processor_name(processor_name, &name_len);

  // Print hello world message
  printf("Hello world from processor %s, rank %d out of %d processors\n",
  processor_name, world_rank, world_size);

  // Finalize the MPI environment
  MPI_Finalize();
}
```

Complete the following tasks:

1. Connect to the HPC cluster using SSH, see the previous exercise for instructions
2. Create a new file using `nano mpi-hello-world.c`
3. Insert the code from the listing into the file

4. Save and exit nano, CTRL + o to save and CTRL + x to exit
5. Load the *openmpi* module using `module load openmpi`
6. Compile your program `mpicc -o mpi-hello-world mpi-hello-world.c`
7. Test it `mpirun -np 2 mpi-hello-world`
8. Run it via Slurm `srun --nodes=2 --tasks-per-node=16 mpi-hello-world`
   It might take a moment for nodes to be allocated for you
9. Observe the output, what do you notice?