

HPC introduction Spring'19

Exercise block 1: Building an environment

Set up some example programs (see also Tutorial §3.1):

- Log on to the cluster
- Go to your data directory
`cd $VSC_DATA`
- Copy the examples for the tutorial to that directory: type
`cp -r /apps/antwerpen/tutorials/Intro-HPC/examples .`

Exercises:

1. Follow the text from the tutorial in §4.1 “Modules” and try out the module command. Note that `module spider` and `module keyword` are still missing from this text as those are not available on all VSC clusters.
2. Load the cluster module `leibniz/2018b` and check the difference between `module av` and `module spider`.
 - a. Try to find modules for the HDF5 library. Did you notice the difference between `module av` and `module spider`?
3. Which module offers support for VNC? Which command does that module offer?
4. Which module(s) offer the CMake tool?
5. What is the difference between the `HDF5/1.8.20-intel-2018b-MPI` and `HDF5/1.8.20-intel-2018b-noMPI` (well, you can guess it from the name but how can you find more information?)

Exercise block 2: Job basics

1. Tutorial §4.3: `examples/Running-batch-jobs` contains a short bash script (`fibonacci.pbs`) that will print Fibonacci numbers. Extend the script to request 1 core on 1 node with 2Gb of (virtual) memory and with a wall time of 1 minute in the job script and run the job. Note which files are generated
2. Extend the script to give the job a name of your choice. Run and note which files are generated.
3. While the job is running, try some of the `q`-commands. To do this, change your job script first to extend the walltime to 10 minutes and put a `sleep 300` at the end because otherwise the job will finish so quickly that you don't get the chance to see anything.

Exercise block 3: Starting various job types

Think for yourself what material might apply best to the research that you will be doing, and select some of the exercises below:

Starting an interactive job

With the help of the example scripts in `examples/Running-interactive-jobs`:

1. Start an interactive session with 1 core on 1 node and run the primes.py example (end of §5.2 in the tutorial)
2. If you have an X-server on your PC:
 - a. How would you run message.py on one of the login nodes?
 - b. Run message.py in an interactive job (so on one of the compute nodes)

Starting parallel applications

The module vsc-tutorial (2018b Intel toolchain) contains some compiled “Hello, world!”-style programs to demonstrate starting OpenMP, MPI or hybrid MPI/OpenMP-programs: omp_hello, mpi_hello and mpi_omp_hello. These commands are documented through man pages. These commands require very little memory and execution time, but are excellent to see where your regular OpenMP, MPI or hybrid application would run.

1. Write a job script that runs the omp_hello command on 7 cores.
2. Write a job script that runs the mpi_hello command on all cores of 2 nodes.
3. Write a job script that starts the mpi_omp_hello hybrid application using 8 MPI processes with 7 cores each. Use the torque-tools module in your job script.

Workflows

1. Try the example from the slides for yourself.
2. Extend the example:
 - a. A first job writes the number 10 to the file “sim1.txt”
 - b. Job two executes after the first. It reads the number from “sim1.txt”, multiplies it by 2 and writes it to “sim2.txt”
 - c. To result of the second simulation, two different perturbations are applied: add 1 and 2 to the number in sim2.txt, and the third “simulation” that starts from that perturbed value multiplies the result by 2 and writes it in a file with the name sim3_pert_<...>.txt with <...> the perturbation.Use dependent jobs to make sure that the second “simulation” runs after the first one and that both “simulations” for the perturbations can run concurrently and after the second “simulation”.

Multi-job submission

1. Parameter sweep: Have a look at the example in examples/Multi-job-submission/par_sweep (and tutorial §12.1) and run this job through the Worker framework
2. Job array: Have a look at the example in examples/Multi-job-submission/job_array (and tutorial §12.2) and run this job through the Worker framework
3. Advanced feature of worker: Map-Reduce: Have a look at the tutorial text §12.3 and the code in examples/Multi-job-submission/map_reduce and try to run this code through the Worker framework