

Batch Matlab

This example is based on using gplogin2, the login node for the new side of cluster.

We use lmod on gplogin2/3 as opposed to environment-modules on gplogin1.

All examples below are created in your home directory via copy-and-paste into your GP shell account:

Consider the following example, matlab_example.m

```
cat << 'EOF' > ~/matlab_example.m
[X,Y] = meshgrid(-2:.2:2);
Z = X .* exp(-X.^2 - Y.^2);
surf(X,Y,Z);
print('example-plot','-dpng');
exit;
EOF
```

Running the above matlab example ***WITHOUT*** Slurm: (this is how many people run on the login node which is **BAD!**)

```
cat << 'EOF' > ~/run.sh
#!/bin/bash

ml purge
ml matlab/R2017b

matlab -nodisplay -nodesktop -nosplash < matlab_example.m
EOF

chmod 755 ~/run.sh && ~/run.sh
```

running the above matlab example ***WITH*** Slurm:

```
cat << 'EOF' > ~/runv2.sh
#!/bin/bash
```

```
#SBATCH --job-name=my_matlab_job
#SBATCH --output=my_matlab_job.out
#SBATCH --error=my_matlab_job.err
#SBATCH --partition=brd2.4,has2.5,ilg2.3,m-c1.9,m-c2.2,nes2.8,sib2.9
#SBATCH --time=00:01:00
#SBATCH --nodes=1
#SBATCH --ntasks=16

ml purge
ml matlab/R2017b

matlab -nodisplay -nodesktop -nosplash < matlab_example.m
EOF

chmod 755 ~/runv2.sh
```

To submit the job

```
sbatch runv2.sh
```

This will create a file in your home directory

```
example-plot.png
```

The hardest part is determining how much resources your computation/simulation will need.

One has to pick an partition based on the computation. Usually people will want Intel CPUs, but we have AMD CPUs as an option.

- [Partitions](#)

The resources for your computation/simulation needs to be determined emperically.

Method 1: Make a Slurm submission script and guesstimate your resources required (CPU cores, number of nodes, walltime etc.).

Submit the job via sbatch and then analyze the efficiency of the job with seff and refine your scheduling parameters on the next run.

```
$ seff 10773
Job ID: 10773
Cluster: blueplanet
User/Group: santucci/staff
```

```
State: COMPLETED (exit code 0)
Nodes: 1
Cores per node: 32
CPU Utilized: 00:00:20
CPU Efficiency: 1.56% of 00:21:20 core-walltime
Job Wall-clock time: 00:00:40
Memory Utilized: 448.57 MB
Memory Efficiency: 2.74% of 16.00 GB
```

If you want to see which node was selected for the job look at the epilog output

```
$ cat slurm.epilog-10773
----- slurm.epilog -----
Job ID: 10773
User: santucci
Group: staff
Job Name: my_matlab_job
Partition: has2.5
QOS: normal
Account: staff
Reason: None,c-19-293
Nodelist: c-19-293
Command: /data11/home/santucci/runv2.sh
WorkDir: /data11/home/santucci
BatchHost: c-19-293
```

Method 2: Request an interactive shell and experiment to determine how much memory is required and how long it needs to run.

```
srun --pty --x11 -t 300 -n 1 -p <partition-list> bash -i
```

recommendations on profile are available @
https://www.nccs.nasa.gov/user_info/slurm/determine_memory_usage

If new to Slurm please see [this page](#).

Here are two quick reference guides that you will want to have handy:

<https://slurm.schedmd.com/pdfs/summary.pdf>

<https://www.chpc.utah.edu/presentations/SlurmCheatsheet.pdf>

Credit: inspiration for this example comes from <https://it.math.ncsu.edu/hpc/slurm/batch/matlab>

Revision #5

Created 12 May 2023 21:49:42 by Chris

Updated 12 May 2023 21:58:05 by Chris