

Using PBS to interface with the SLURM job scheduler

The nan.claruk.edu Microway cluster uses a newer job scheduler process called SLURM rather than the older TORQUE based Qsub as on the older Clark HPC cluster. However, Qsub is still installed on nan.claruk.edu and can be used to submit jobs.

Rather than simply submitting all jobs to the 'clark' queue as on the old cluster SLURM provides us with a range of queues useful for different situations.

```
[juapark@nan qsub_practice]$ qstat -q
```

Queue	Memory	CPU	Time	Walltime	Node	Run	Que	Ln	State
month-long-cpu	--	--		744:00:00	6	0	0	--	E R
week-long-cpu	--	--		168:00:00	8	0	0	--	E R
day-long-cpu	--	--		24:00:00	--	0	0	--	E R
short-cpu	--	--		00:30:00	--	0	1	--	E R
interactive-cpu	--	--		48:00:00	1	0	0	--	E R
month-long-gpu	--	--		744:00:00	1	0	0	--	E R
week-long-gpu	--	--		168:00:00	1	0	0	--	E R
day-long-gpu	--	--		24:00:00	--	0	0	--	E R
short-gpu	--	--		00:30:00	--	0	0	--	E R
interactive-gpu	--	--		48:00:00	1	0	0	--	E R

- Short/day/week/month queues represent how much walltime your job will need. In our old cluster, jobs that could finish in 10 minutes could stand in queue behind jobs that would take days and months to finish, slowing down the output rate in times of high congestion. To know which queue to use, look at the "walltime" column. This represents the maximum walltime a job in the queue can take, shorter jobs are more suitable to shorter queues.
- There are queues that use CPU and queues that use the GPU. The cluster has 16 CPU nodes and 2 GPU nodes. The advantage of using GPU nodes is that they're very powerful, which is suitable for high resource jobs. However, the downside is that there are only 2 nodes, so queue times might vary largely depending on the usage.

Rather than allowing Qsub to default to the old "clark" queue See the example PBS script below for a job to be submitted in the new cluster.

```
#!/bin/bash
#PBS -N sample_job
#PBS -q short-cpu
#PBS -j oe
#PBS -o clark.$PBS_JOBID
#PBS -l nodes=1:ppn=1
#PBS -l walltime=00:00:10
```

By specifying the SLURM queue using the "#PBS -q" value Qsub will be able to send the job to the appropriate queue. From here the job can be submitted using the standard qsub and qstat commands.