



LearnIT Express is on its way!

- The demonstration will start at 12:30 pm.
- To use the Live Chat:
 - Maximize your browser window so you can see the postings.
 - If you refresh the browser page, you'll need to log in again.
 - After the demonstration, staff will remain available in the chat for a few minutes to answer additional questions.

Today's topic:
Mills HPC Cluster

LearnIT Express: Increase Matlab Throughput

Matlab and Threads

- Many Matlab native functions use a math multi-threaded library. This is called *implicit multiprocessing*.
- You can define workers to do *explicit multiprocessing* with the Matlab parallel toolbox.

Grid Engine and Threads

- Increase throughput of your jobs by telling Grid Engine the number of threads you will be using.
- `qlogin` now allows the thread parallel environment for interactive jobs.

```
qlogin -pe threads NSLOTS
```



Threads and Cores

Goal: keep one thread per core

- Matlab does this, but it doesn't consider other jobs competing for the same cores.
- Use Grid Engine to reserve slots, so Matlab can utilize the cores effectively.

Reserving Slots

qstatgrp

CLUSTER	QUEUE	CQLOAD	USED	RES	AVAIL	TOTAL
it_css-dev.q		0.00	0	0	72	72
it_css-qrsh.q		0.00	0	0	72	72
it_css.q		0.00	0	0	72	72
it_css.q+		0.00	0	0	72	72

qlogin -pe threads 10

CLUSTER	QUEUE	CQLOAD	USED	RES	AVAIL	TOTAL
it_css-dev.q		0.00	0	0	72	72
it_css-qrsh.q		0.00	10*	0	62*	72
it_css.q		0.00	0	0	72	72
it_css.q+		0.00	0	0	72	72

* 10 slots were moved from AVAIL to USED columns

Workers and Slots

`matlabpool` function will start Matlab workers that are scheduled by a local scheduler (12 for the parallel toolbox scheduler).

To make this work well with the Grid Engine Scheduler use:

```
-pe threads 12
```

with `qlogin`, `qsub` or in the script file.

Implicit Threads and Slots

Matlab native functions use threads implicitly, but we can estimate the number of slots by timing a test run.

```
real    2m30.213s
user    25m51.183s
sys     0m1.544s
```

```
user/real = 1551.18/150.21 = 10.32
```

```
-pe threads 11
```

Matlab tic/toc commands

To measure the elapsed time inside a matlab function use the `tic` and `toc` commands.

```
Elapsed time is 3.276713 seconds.
```

Compare the times:

```
real    0m5.776s
user    0m33.303s
sys     0m0.220s
```

Fraction of time between tic/toc 3.28/5.78



Summary

- Know your task and the number of cores you will need.
- Tell Grid engine by specifying a threads parallel environment.
- You and your cluster group can do many tasks on your nodes.

Demonstration Ainvb

The Matlab function `Ainvb.m` is in my matlab directory on lustre file system.

```
cd $MYWORKDIR/matlab  
qlogin -pe threads 12
```

On the compute node:

```
vpkg_require matlab  
matlab & top
```

Timing results

`matlab`

1200% cpu for 16.02 seconds

`matlab -singleCompThread`

100% cpu for 145.68 seconds

More information

- **IT Research Computing**

<http://www.it.udel.edu/research-computing>

- Alerts
- Announcements
- Resources: Training

Contact the IT Support Center

- **Email:** consult@udel.edu

If you make the first line of the e-mail message

Type=Cluster-Mills

your question will be routed more quickly.

- **Phone:** (302) 831-6000
- **Text:** (302) 722-6820

